



Implementing Cisco IP Routing (ROUTE)

Foundation Learning Guide

CCNP ROUTE 300-101



ciscopress.com

Diane Teare
Bob Vachon
Rick Graziani

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide

Diane Teare
Bob Vachon
Rick Graziani

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide

Diane Teare, Bob Vachon, Rick Graziani

Copyright © 2015 Cisco Systems, 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

First Printing January 2015

Library of Congress Control Number: 2014957555

ISBN-13: 978-1-58720-456-2

ISBN-10: 1-58720-456-8

Warning and Disclaimer

This book is designed to provide information about Cisco CCNP routing. 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.

Special Sales

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

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

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

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

Business Operation Manager,
Cisco Press: Jan Cornelissen

Executive Editor: Mary Beth Ray

Managing Editor: Sandra Schroeder

Senior Development Editor: Christopher Cleveland

Project Editor: Mandie Frank

Copy Editor: Keith Cline

Technical Editor: Denise Donahue

Team Coordinator: Vanessa Evans

Designer: Mark Shirar

Composition: Trina Wurst

Indexer: Tim Wright

Proofreader: Paula Lowell



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, CCIIP, CCNA, CCNP, CCSP, CCRP, 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 Authors

Diane Teare, P.Eng, CCNP, CCDP, CCSI, PMP, is a professional in the networking, training, project management, and e-learning fields. She has more than 25 years of experience in designing, implementing, and troubleshooting network hardware and software, and has been involved in teaching, course design, and project management. She has extensive knowledge of network design and routing technologies. Diane is a Cisco Certified Systems Instructor (CCSI), and holds her Cisco Certified Network Professional (CCNP), Cisco Certified Design Professional (CCDP), and Project Management Professional (PMP) certifications. She is an instructor, and the Course Director for the CCNA and CCNP Routing and Switching curriculum, with one of the largest authorized Cisco Learning Partners. She was the director of e-learning for the same company, where she was responsible for planning and supporting all the company's e-learning offerings in Canada, including Cisco courses. Diane has a bachelor's degree in applied science in electrical engineering and a master's degree in applied science in management science. She authored or co-authored the following Cisco Press titles: the first edition of this book; the second edition of *Designing Cisco Network Service Architectures (ARCH)*; *Campus Network Design Fundamentals*; the three editions of *Authorized Self-Study Guide Building Scalable Cisco Internetworks (BSCI)*; and *Building Scalable Cisco Networks*. Diane edited the first two editions of the *Authorized Self-Study Guide Designing for Cisco Internetwork Solutions (DESGN)*, and *Designing Cisco Networks*.

Bob Vachon, is a professor at Cambrian College in Sudbury, Ontario, Canada, where he teaches Cisco networking infrastructure courses. He has more than 30 years of work and teaching experience in the computer networking and information technology field. Since 2001, Bob has collaborated as team lead, lead author, and subject matter expert on various CCNA, CCNA-S, and CCNP projects for Cisco and the Cisco Networking Academy. He also was a contributing author for the *Routing Protocols Companion Guide*, *Connecting Networks Companion Guide*, and authored the *CCNA Security (640-554) Portable Command Guide*. In his downtime, Bob enjoys playing the guitar, playing pool, and either working in his gardens or white-water canoe tripping.

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. Before teaching, Rick worked in IT for various companies, including Santa Cruz Operation, Tandem Computers, and Lockheed Missiles and Space Corporation. He holds a Master of Arts degree in computer science and systems theory from California State University Monterey Bay. Rick also works for the Cisco Networking Academy Curriculum Engineering team and has written other books for Cisco Press, including *IPv6 Fundamentals*. When Rick is not working, he is most likely surfing. Rick is an avid surfer who enjoys surfing at his favorite Santa Cruz breaks.

About the Technical Reviewer

Denise Donohue, CCIE No. 9566 (Routing and Switching), is a senior solutions architect with Chesapeake NetCraftsmen. Denise has worked with computer systems since the mid-1990s, focusing on network design since 2004. During that time, she has designed for a wide range of networks, private and public, of all sizes, across most industries. Denise has also authored or co-authored many Cisco Press books covering data and voice networking technologies and spoken at Cisco Live and other industry events.

Dedications

From Diane: This book is dedicated to my husband, Allan Mertin—thank you for your love, encouragement, and patience; to our extraordinary son, Nicholas—thank you for your love and for sharing as you discover the world; and to my parents, Syd and Beryl, for their inspiration.

From Rick: This book is dedicated to the Cabrillo College CIS/CS faculty, staff, administration, and especially students for giving me the privilege and honor to teach computer networking courses at such a wonderful institution. I would also like to thank all my family and friends for their love and support.

From Bob: This book is dedicated to my beautiful wife, Judy, and my girls, Lee-Anne, Joëlle, Brigitte, and Lilly. Thank you for your encouragement and for putting up with me while working on this project. I also dedicate this book to my students at Cambrian College and to my dean, Joan Campbell, for your continued support.

Acknowledgments

We want to thank many people for helping to put this book together:

The Cisco Press team: Mary Beth Ray, the executive editor, coordinated the whole project, steered the book through the necessary processes, and understood when the inevitable snags appeared. Sandra Schroeder, the managing editor, brought the book to production. Vanessa Evans was once again wonderful at organizing the logistics and administration. Chris Cleveland, the development editor, has been invaluable in coordinating and ensuring we all focused on producing the best manuscript.

We also want to thank Mandie Frank, the project editor, and Keith Cline, the copy editor, for their excellent work in getting this book through the editorial process.

The Cisco ROUTE course development team: Many thanks to the members of the team who developed the ROUTE course.

The technical reviewer: We want to thank the technical reviewer of this book, Denise Donahue, for her thorough review and valuable input.

Our families: Of course, this book would not have been possible without the endless understanding and patience of our families. They have always been there to motivate and inspire us and we are forever grateful.

From Diane: A few special thank yous are in order. First, to Brett Bartow (who invited me to first write with Cisco Press many years ago) and Mary Beth Ray, for the very warm welcome when I finally met you both in person and for continuing to involve me in your projects. Second, to Rick and Bob for including me in this book; it has been a great pleasure to work with you both!

From Rick: A special thank you to Mary Beth Ray for giving me the opportunity years ago to begin writing for Cisco Press, and for being such a wonderful friend. Also, thank you to my two good friends Diane and Bob for letting me work with you on this book.

From Bob: A special thank you to Mary Beth Ray and her team at Cisco Press for your continued support, your professionalism, and skills to make us look good. Also, a big thank you to my fellow co-authors, Diane and my good friend Rick, whom I've had the honor and pleasure to work with on numerous projects.

Contents at a Glance

Introduction	xxv
Chapter 1: Basic Network and Routing Concepts	1
Chapter 2: EIGRP Implementation	59
Chapter 3: OSPF Implementation	155
Chapter 4: Manipulating Routing Updates	267
Chapter 5: Path Control Implementation	327
Chapter 6: Enterprise Internet Connectivity	373
Chapter 7: BGP Implementation	423
Chapter 8: Routers and Routing Protocol Hardening	527
Appendix A: Answers to End of Chapter Review Questions	607
Appendix B: IPv4 Supplement	613
Appendix C: BGP Supplement	671
Appendix D: Acronyms and Abbreviations	697
Index	701

Contents

Introduction xxv

Chapter 1 Basic Network and Routing Concepts 1

Differentiating Routing Protocols	2
Enterprise Network Infrastructure	2
Role of Dynamic Routing Protocols	3
Choosing a of Dynamic Routing Protocols	5
IGP versus EGP	5
Types of Routing Protocols	7
Convergence	8
Route Summarization	9
Route Protocol Scalability	10
Understanding Network Technologies	10
Traffic Types	11
IPv6 Address Types	13
ICMPv6 Neighbor Discovery	14
Network Types	15
NBMA Networks	16
Routing Over the Internet	18
Connecting Remote Locations with Headquarters	18
Principles of Static Routing	19
<i>Configuring an IPv4 Static Route</i>	20
<i>Configuring a Static Default Route</i>	22
<i>Basic PPP Overview</i>	23
<i>PPP Authentication Overview</i>	23
<i>PPPoE</i>	26
<i>Basic Frame Relay Overview</i>	28
<i>VPN Connectivity Overview</i>	31
<i>MPLS-based VPNs</i>	31
<i>Tunneling VPNs</i>	32
<i>Hybrid VPNs</i>	32
Routing Across MPLS VPNs	32
Routing Over GRE Tunnel	34
Dynamic Multipoint Virtual Private Network	35
Multipoint GRE	36
NHRP	37
IPsec	39

Routing and TCP/IP Operations	40
MSS, Fragmentation, and PMTUD	40
IPv4 Fragmentation and PMTUD	41
Bandwidth Delay Product	41
TCP Starvation	42
Latency	42
ICMP Redirect	42
Implementing RIPng	43
RIP Overview	43
RIPv2 Overview	45
Configuring RIPng	47
<i>Basic RIPng Configuration</i>	47
<i>Propagating a Default Route</i>	50
Investigating the RIPng Database	53
Summary	55
Review Questions	56
Chapter 2 EIGRP Implementation	59
Establishing EIGRP Neighbor Relationships	60
EIGRP Features	60
EIGRP Features	62
EIGRP Operation Overview	63
Configuring and Verifying Basic EIGRP for IPv4	64
Manipulating EIGRP Timers	73
EIGRP Neighbor Relationship over Frame Relay	74
Establishing EIGRP over Layer 3 MPLS VPN	74
Establishing EIGRP over Layer 2 MPLS VPN	75
Building the EIGRP Topology Table	76
Building and Examining the EIGRP Topology Table	77
<i>Choosing the Best Path</i>	80
Exchange of Routing Knowledge in EIGRP	88
EIGRP Metric	88
EIGRP Metric Calculation	89
<i>EIGRP Wide Metrics</i>	90
EIGRP Metric Calculation Example	90
EIGRP Metric Calculation Example	91
EIGRP Path Calculation Example	92

Optimizing EIGRP Behavior	94
EIGRP Queries	95
EIGRP Stub Routers	96
Configuring EIGRP Stub Routing	97
<i>EIGRP Stub Options</i>	100
Stuck in Active	108
Reducing Query Scope by Using Summary Routes	109
Configuring EIGRP Summarization	110
<i>Determining the Summary Route</i>	116
<i>Obtaining Default Route</i>	120
Load Balancing with EIGRP	123
Configuring EIGRP Load Balancing	123
<i>EIGRP Load Balancing</i>	124
<i>EIGRP Load Balancing Across Unequal-Metric Paths</i>	126
Configuring EIGRP for IPv6	128
Overview of EIGRP for IPv6	128
Configuring and Verifying EIGRP for IPv6	129
<i>EIGRP for IPv6 Configuration</i>	130
<i>Determining the IPv6 Summary Route</i>	134
Named EIGRP Configuration	136
Introduction to Named EIGRP Configuration	136
Configuring Named EIGRP	137
<i>Address Families</i>	139
<i>EIGRP for IPv4 Address Family</i>	139
<i>EIGRP for IPv6 Address Family</i>	142
Named EIGRP Configuration Modes	148
Classic Versus Named EIGRP Configuration	150
Summary	151
Review Questions	152
Chapter 3 OSPF Implementation	155
Establishing OSPF Neighbor Relationships	155
OSPF Features	156
OSPF Operation Overview	157
Hierarchical Structure of OSPF	158
Design Restrictions of OSPF	160
OSPF Message Types	160
Basic OSPF Configuration	161
<i>Optimizing OSPF Adjacency Behavior</i>	170
<i>Using OSPF Priority in the DR/BDR Election</i>	174

<i>OSPF Behavior in NBMA Hub-and-Spoke Topology</i>	175
<i>The Importance of MTU</i>	177
<i>Manipulating OSPF Timers</i>	179
OSPF Neighbor Relationship over Point-to-Point Links	182
OSPF Neighbor Relationship over Layer 3 MPLS VPN	182
OSPF Neighbor Relationship over Layer 2 MPLS VPN	184
OSPF Neighbor States	184
OSPF Network Types	186
Configuring Passive Interfaces	187
Building the Link-State Database	187
OSPF LSA Types	188
Examining the OSPF Link-State Database	189
<i>OSPF Link-State Database</i>	190
<i>OSPF Type 2 Network LSA</i>	196
<i>OSPF Type 3 Summary LSA</i>	197
<i>OSPF Type 4 ASBR Summary LSA</i>	199
<i>OSPF Type 5 External LSA</i>	201
Periodic OSPF Database Changes	203
Exchanging and Synchronizing LSDBs	204
Synchronizing the LSDB on Multiaccess Networks	206
Running the SPF Algorithm	207
Configuring OSPF Path Selection	208
<i>OSPF Path Selection</i>	208
<i>OSPF Best Path Calculation</i>	210
<i>Default OSPF Costs</i>	211
Calculating the Cost of Intra-Area Routes	214
Calculating the Cost of Interarea Routes	214
Selecting Between Intra-Area and Interarea Routes	215
Optimizing OSPF Behavior	215
OSPF Route Summarization	216
Benefits of Route Summarization	217
Configuring OSPF Route Summarization	218
Summarization on ABRs	223
Summarization on ASBRs	224
OSPF Virtual Links	225
<i>Configuring OSPF Virtual Links</i>	227
Configuring OSPF Stub Areas	229
<i>OSPF Stub Areas</i>	230
<i>OSPF Totally Stubby Areas</i>	234

Cost of the Default Route in a Stub Area	236
The default-information originate Command	237
Other Stubby Area Types	238
OSPFv3	239
Configuring OSPFv3	240
<i>Implementing OSPFv3</i>	241
<i>OSPFv3 for IPv4 and IPv6</i>	246
Configuring Advanced OSPFv3	260
OSPFv3 Caveats	261
Summary	262
Review Questions	263
Chapter 4 Manipulating Routing Updates	267
Using Multiple IP Routing Protocols on a Network	267
Why Run Multiple Routing Protocols?	269
Running Multiple Routing Protocols	269
<i>Administrative Distance</i>	269
Multiple Routing Protocols Solutions	270
Implementing Route Redistribution	270
Defining Route Redistribution	270
Planning to Redistribute Routes	271
Redistributing Routes	271
Seed Metrics	272
<i>Default Seed Metrics</i>	273
Configuring and Verifying Basic Redistribution in IPv4 and IPv6	275
<i>Redistributing OSPFv2 Routes into the EIGRP Routing Domain</i>	276
<i>Redistributing OSPFv3 Routes into the EIGRP for IPv6 Routing Domain</i>	279
<i>Redistributing EIGRP Routes into the OSPFv2 Routing Domain</i>	281
<i>Redistributing EIGRP for IPv6 Routes into the OSPFv3 Routing Domain</i>	285
Types of Redistribution Techniques	287
<i>One-Point Redistribution</i>	287
<i>Multipoint Redistribution</i>	288
<i>Redistribution Problems</i>	289
<i>Preventing Routing Loops in a Redistribution Environment</i>	291
<i>Verifying Redistribution Operation</i>	292
Controlling Routing Update Traffic	292
Why Filter Routes?	292
Route Filtering Methods	293

Using Distribute Lists	294
<i>Configuring Distribute Lists</i>	294
<i>Distribute List and ACL Example</i>	295
Using Prefix Lists	297
<i>Prefix List Characteristics</i>	297
<i>Configuring Prefix Lists</i>	298
<i>Distribute List and Prefix List Example</i>	299
<i>Prefix List Examples</i>	300
<i>Verifying Prefix Lists</i>	301
<i>Manipulating Redistribution Using ACLs, Prefix Lists, and Distribute Lists</i>	302
Using Route Maps	305
<i>Understanding Route Maps</i>	305
<i>Route Map Applications</i>	305
<i>Configuring Route Maps</i>	306
<i>Route Map Match and Set Statements</i>	308
Configuring Route Redistribution Using Route Maps	310
<i>Using Route Maps with Redistribution</i>	310
<i>Manipulating Redistribution Using Route Maps</i>	311
<i>Mutual Redistribution without Route Filtering</i>	312
<i>Mutual Redistribution with Route Maps</i>	313
<i>Change Administrative Distance to Enable Optimal Routing</i>	315
Manipulating Redistribution Using Route Tagging	318
Caveats of Redistribution	319
Summary	320
References	323
Review Questions	323
Chapter 5 Path Control Implementation	327
Using Cisco Express Forwarding Switching	327
Control and Data Plane	328
Cisco Switching Mechanisms	328
Process and Fast Switching	332
Cisco Express Forwarding	333
Analyzing Cisco Express Forwarding	335
<i>Verify the Content of the CEF Tables</i>	335
<i>Enable and Disable CEF by Interface and Globally</i>	341
Understanding Path Control	343
The Need for Path Control	343

Implementing Path Control Using Policy-Based Routing	344
<i>PBR Features</i>	344
<i>Steps for Configuring PBR</i>	345
<i>Configuring PBR</i>	346
<i>Verifying PBR</i>	348
<i>Configuring PBR Example</i>	348
Implementing Path Control Using Cisco IOS IP SLAs	354
<i>PBR and IP SLA</i>	354
<i>IP SLA Features</i>	354
<i>Steps for Configuring IP SLAs</i>	356
<i>Verifying Path Control Using IOS IP SLAs</i>	360
<i>Configuring IP SLA Example</i>	361
<i>Configuring PBR and IP SLA Example</i>	364
Summary	369
References	370
Review Questions	370
Chapter 6	Enterprise Internet Connectivity 373
Planning Enterprise Internet Connectivity	374
Connecting Enterprise Networks to an ISP	374
<i>Enterprise Connectivity Requirements</i>	374
<i>ISP Redundancy</i>	375
Public IP Address Assignment	376
<i>The Internet Assigned Numbers Authority</i>	376
<i>Regional Internet Registries</i>	377
<i>Public IP Address Space</i>	377
Autonomous System Numbers	378
Establishing Single-Homed IPv4 Internet Connectivity	381
Configuring a Provider-Assigned IPv4 Address	381
DHCP Operation	382
Obtaining a Provider-Assigned IPv4 Address with DHCP	383
Configuring a Router as a DHCP Server and DHCP Relay Agent	384
NAT	385
<i>Configuring Static NAT</i>	388
<i>Configuring Dynamic NAT</i>	389
<i>Configuring PAT</i>	390
<i>Limitations of NAT</i>	392
NAT Virtual Interface	393
<i>Configuring NAT Virtual Interface</i>	393
<i>Verifying NAT Virtual Interface</i>	396

Establishing Single-Homed IPv6 Internet Connectivity	398
Obtaining a Provider-Assigned IPv6 Address	398
<i>Manual Assignment</i>	399
<i>Configuring Basic IPv6 Internet Connectivity</i>	399
<i>Stateless Address Autoconfiguration</i>	401
<i>DHCPv6 Operation</i>	402
<i>Stateless DHCPv6</i>	403
<i>Stateful DHCPv6</i>	404
<i>DHCPv6 Prefix Delegation</i>	405
NAT for IPv6	405
NAT64	405
NPTv6	405
IPv6 ACLs	405
<i>IPv6 ACL Characteristics</i>	406
<i>Configuring IPv6 ACLs</i>	406
Securing IPv6 Internet Connectivity	409
Improving Internet Connectivity Resilience	410
Drawbacks of a Single-Homed Internet Connectivity	410
Dual-Homed Internet Connectivity	410
<i>Dual-Homed Connectivity Options</i>	411
<i>Configuring Best Path for Dual-Homed Internet Connectivity</i>	411
Multihomed Internet Connectivity	413
Summary	415
References	417
Review Questions	418
Chapter 7 BGP Implementation	423
BGP Terminology, Concepts, and Operation	424
BGP Use Between Autonomous Systems	424
Comparison with Other Scalable Routing Protocols	425
BGP Path Vector Characteristics	426
BGP Characteristics	428
BGP Tables	430
BGP Message Types	431
<i>Open and Keepalive Messages</i>	431
<i>Update Messages</i>	433
<i>Notification Messages</i>	433
When to Use BGP	433
When Not to Use BGP	434

Implementing Basic BGP	435
BGP Neighbor Relationships	435
<i>External BGP Neighbors</i>	436
<i>Internal BGP Neighbors</i>	437
<i>iBGP on All Routers in a Transit Path</i>	438
Basic BGP Configuration Requirements	442
Entering BGP Configuration Mode	442
Defining BGP Neighbors and Activating BGP Sessions	443
Basic BGP Configuration and Verification	444
<i>Configuring and Verifying an eBGP Session</i>	445
<i>Configuring and Verifying an iBGP Session</i>	449
<i>Advertising Networks in BGP and Verifying That They Are Propagated</i>	450
<i>Using the Next-Hop-Self Feature</i>	457
<i>Understanding and Troubleshooting BGP Neighbor States</i>	458
<i>BGP Session Resilience</i>	460
<i>Sourcing BGP from Loopback Address</i>	461
<i>eBGP Multihop</i>	463
<i>Resetting BGP Sessions</i>	464
BGP Attributes and the Path-Selection Process	467
BGP Path Selection	467
<i>BGP Path-Selection Process</i>	468
<i>The Path-Selection Decision Process with a Multihomed Connection</i>	469
BGP Attributes	471
<i>Well-Known Attributes</i>	471
<i>Optional Attributes</i>	472
<i>Defined BGP Attributes</i>	472
<i>The AS-Path Attribute</i>	473
<i>The Next-Hop Attribute</i>	474
<i>The Origin Attribute</i>	475
<i>The Local-Preference Attribute</i>	475
<i>The Community Attribute</i>	475
<i>The MED Attribute</i>	476
<i>The Weight Attribute (Cisco Only)</i>	478
<i>Changing the Weight for All Updates from a Neighbor</i>	479
<i>Changing the Weight Using Route Maps</i>	479
Influencing BGP Path Selection	480
<i>Changing the Weight</i>	485

	<i>Changing Local Preference</i>	486
	<i>Setting the AS-Path</i>	488
Controlling BGP Routing Updates		491
	Filtering BGP Routing Updates	492
	<i>BGP Filtering Using Prefix Lists</i>	492
	<i>BGP Filtering Using AS-Path Access Lists</i>	494
	<i>BGP Filtering Using Route Maps</i>	496
	<i>Filtering Order</i>	498
	<i>Clearing the BGP Session</i>	498
	BGP Peer Groups	498
	<i>Peer Group Operation</i>	498
	<i>Peer Group Configuration</i>	500
	<i>Peer Group Configuration Example</i>	500
Implementing BGP for IPv6 Internet Connectivity		502
	MP-BGP Support for IPv6	502
	Exchanging IPv6 Routes over an IPv4 Session	504
	Exchanging IPv6 Routes over an IPv6 Session	506
	BGP for IPv6 Configuration and Verification	507
	<i>Initial State of Routers</i>	508
	<i>Enable eBGP IPv6 Route Exchange</i>	511
	<i>Enable iBGP IPv6 Route Exchange</i>	516
	Comparing IPv4 to Dual (IPv4/IPv6) BGP Transport	518
	BGP Filtering Mechanisms for IPv6	518
	<i>IPv6 Prefix List Filtering</i>	518
	<i>IPv6 Path Selection with BGP Local Preference</i>	519
Summary		520
References		522
Review Questions		523
Chapter 8	Routers and Routing Protocol Hardening	527
	Securing the Management Plane on Cisco Routers	528
	Securing the Management Plane	529
	Router Security Policy	530
	Encrypted Passwords	531
	<i>Use Strong Passwords</i>	532
	<i>Encrypting Passwords</i>	532
	Authentication, Authorization, Accounting	536
	<i>RADIUS and TACACS+ Overview</i>	536
	<i>Enabling AAA and Local Authentication</i>	538

<i>Enabling AAA RADIUS Authentication with Local User for Backup</i>	539
<i>Enabling AAA TACACS+ Authentication with Local User for Backup</i>	541
<i>Configuring Authorization and Accounting</i>	542
<i>Limitations of TACACS+ and RADIUS</i>	542
Use SSH Instead of Telnet	543
Securing Access to the Infrastructure Using Router ACLs	547
Implement Unicast Reverse Path Forwarding	549
<i>uRPF in an Enterprise Network</i>	550
<i>uRPF Examples</i>	550
<i>Enabling uRPF</i>	551
Implement Logging	551
Implementing Network Time Protocol	552
<i>NTP Modes</i>	552
<i>Enabling NTP</i>	554
<i>Securing NTP</i>	555
<i>NTP Versions</i>	556
<i>NTP in IPv6 Environment</i>	557
<i>Simple NTP</i>	557
Implementing SNMP	558
<i>SNMPv3</i>	561
<i>Enabling SNMPv3</i>	561
<i>Verifying SNMPv3</i>	562
Configuration Backups	563
<i>The archive Command</i>	563
Using SCP	565
<i>Enabling SCP on a Router</i>	565
Disabling Unused Services	567
Conditional Debugging	568
<i>Enabling Conditional Debugging</i>	569
Routing Protocol Authentication Options	570
The Purpose of Routing Protocol Authentication	570
<i>Plain-Text Authentication</i>	571
<i>Hashing Authentication</i>	572
Time-Based Key Chains	574
<i>Key Chain Specifics</i>	574
Authentication Options with Different Routing Protocols	575

Configuring EIGRP Authentication	576
EIGRP Authentication Configuration Checklist	577
Configuring EIGRP Authentication	577
<i>Configure EIGRP MD5 Authentication Mode</i>	578
<i>Configure EIGRP Key-Based Routing Authentication</i>	579
Configuring EIGRP for IPv6 Authentication	581
<i>Configure EIGRP for IPv6 MD5 Authentication Mode</i>	581
<i>Configuring Named EIGRP Authentication</i>	582
Configuring OSPF Authentication	583
OSPF Authentication	583
OSPF MD5 Authentication	584
<i>Configure OSPF MD5 Authentication</i>	584
<i>Configure OSPF MD5 Authentication on Interfaces</i>	585
<i>Configure OSPF MD5 Authentication in an Area</i>	586
OSPFv2 Cryptographic Authentication	587
<i>Configuring OSPFv2 Cryptographic Authentication</i>	587
<i>Configure OSPFv2 Cryptographic Authentication Example</i>	588
OSPFv3 Authentication	590
<i>Configuring OSPFv3 Authentication</i>	590
<i>Configuring OSPFv3 Authentication on an Interface Example</i>	591
<i>Configuring OSPFv3 Authentication in an Area Example</i>	592
Configuring BGP Authentication	593
BGP Authentication Configuration Checklist	594
BGP Authentication Configuration	594
BGP for IPv6 Authentication Configuration	596
Implementing VRF-Lite	597
VRF and VRF-Lite	597
Enabling VRF	597
Easy Virtual Network	601
Summary	603
References	604
Review Questions	604
Appendix A	Answers to End of Chapter Review Questions
Chapter 1	607
Chapter 2	608
Chapter 3	609
Chapter 4	610
Chapter 5	610

Chapter 6 611

Chapter 7 611

Chapter 8 612

Appendix B IPv4 Supplement 613

IPv4 Addresses and Subnetting Job Aid 614

Decimal-to-Binary Conversion Chart 614

IPv4 Addressing Review 618

 Converting IP Addresses Between Decimal and Binary 618

 Determining an IP Address Class 619

 Private Addresses 620

 Extending an IP Classful Address Using a Subnet Mask 620

 Calculating a Subnet Mask 621

 Calculating the Networks for a Subnet Mask 623

 Using Prefixes to Represent a Subnet Mask 624

IPv4 Access Lists 625

 IP Access List Overview 625

 IP Standard Access Lists 626

Wildcard Masks 628

Access List Configuration Tasks 629

IP Standard Access List Configuration 629

Implicit Wildcard Masks 630

Configuration Principles 631

Standard Access List Example 632

Location of Standard Access Lists 633

 IP Extended Access Lists 634

Extended Access List Processing 634

Extended IP Access List Configuration 635

Extended Access List Examples 642

Location of Extended Access Lists 643

Time-Based Access Lists 644

 Restricting Virtual Terminal Access 645

How to Control vty Access 645

Virtual Terminal Line Access Configuration 646

 Verifying Access List Configuration 647

IPv4 Address Planning 648

 Benefits of an Optimized IP Addressing Plan 648

 Scalable Network Addressing Example 650

 Non-scalable Network Addressing 651

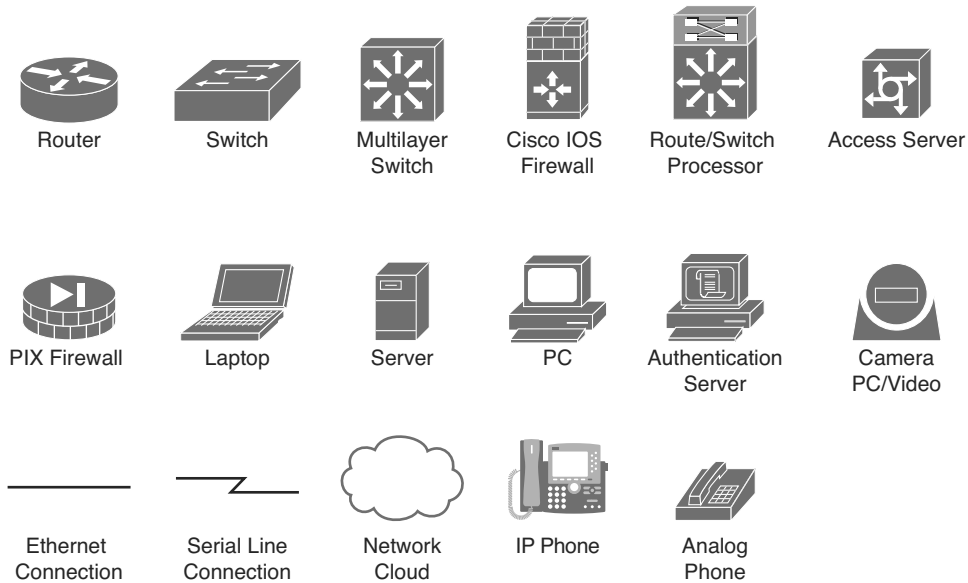
<i>Update Size</i>	651
<i>Unsummarized Internetwork Topology Changes</i>	652
<i>Summarized Network Topology Changes</i>	652
Hierarchical Addressing Using Variable-Length Subnet Masks	653
Subnet Mask	653
<i>Use of the Subnet Mask</i>	653
<i>Subnet Mask Example</i>	653
Implementing VLSM in a Scalable Network	654
VLSM Calculation Example	656
<i>LAN Addresses</i>	657
<i>Serial Line Addresses</i>	658
<i>Summary of Addresses Used in the VLSM Example</i>	661
Another VLSM Example	661
Route Summarization	662
Route Summarization Overview	662
Route Summarization Calculation Example	664
Summarizing Addresses in a VLSM-Designed Network	665
Route Summarization Implementation	666
Route Summarization Operation in Cisco Routers	666
Route Summarization in IP Routing Protocols	667
Classless Interdomain Routing	667
CIDR Example	668
Appendix C BGP Supplement	671
BGP Route Summarization	671
CIDR and Aggregate Addresses	671
Network Boundary Summarization	673
BGP Route Summarization Using the network Command	674
Creating a Summary Address in the BGP Table Using the aggregate-address Command	677
Redistribution with IGP	680
Advertising Networks into BGP	680
Advertising from BGP into an IGP	681
Communities	682
Community Attribute	682
Setting and Sending the Communities Configuration	682
Using the Communities Configuration	685

Route Reflectors	687
Route Reflector Benefits	689
Route Reflector Terminology	689
Route Reflector Design	690
Route Reflector Design Example	690
Route Reflector Operation	691
Route Reflector Migration Tips	692
Route Reflector Configuration	694
Route Reflector Example	694
Verifying Route Reflectors	695
Advertising a Default Route	695
Not Advertising Private Autonomous System Numbers	696

Appendix D Acronyms and Abbreviations 697

Index 701

Icons Used in This Book



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.

Configuration and Verification Examples

Most of the configuration and verification examples in this book were done using Cisco IOS over Linux (IOL) virtual environment (the same environment used in the ROUTE course). This environment runs the IOS software on Linux instead of on actual router and switch hardware. As a result, there are a few things to note for these configuration examples:

- All Ethernet-type interfaces on the devices are “Ethernet” (rather than “FastEthernet” or “GigabitEthernet”).
- All PCs used in the examples are actually running the IOL, so testing is done with IOS commands such as ping and traceroute.
- An interface always indicates that it is up/up unless it is shutdown. For example, if an interface on device 1 is shutdown, the interface on device 2, connected to that down interface on device 1, will indicate up/up (it does not reflect the true state).

Introduction

Networks continue to grow, becoming more complex as they support more protocols and more users. This book teaches you how to plan, implement, and monitor a scalable routing network. It focuses on using Cisco routers connected in LANs and WANs typically found at medium to large network sites.

In this book, you study a broad range of technical details on topics related to routing. First, basic network and routing protocol principles are examined in detail before the following IP Version 4 (IPv4) and IP Version 6 (IPv6) routing protocols are studied: Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP). Enterprise Internet connectivity is explored. Manipulating routing updates and controlling the path that traffic takes are examined. Best practices for securing Cisco routers are described.

Configuration examples and sample verification outputs demonstrate troubleshooting techniques and illustrate critical issues surrounding network operation. Chapter-ending review questions illustrate and help solidify the concepts presented in this book.

This book starts you down the path toward attaining your CCNP or CCDP certification, providing in-depth information to help you prepare for the ROUTE exam (300-101).

The commands and configuration examples presented in this book are based on Cisco IOS Release 15.1 and 15.2.

Who Should Read This Book?

This book is intended for network architects, network designers, systems engineers, network managers, and network administrators who are responsible for implementing and troubleshooting growing routed networks.

If you are planning to take the ROUTE exam toward your CCNP or CCDP certification, this book provides you with in-depth study material. To fully benefit from this book, you should have your CCNA Routing and Switching certification or possess the same level of knowledge, including an understanding of the following topics:

- A working knowledge of the OSI reference model and networking fundamentals.
- The ability to operate and configure a Cisco router, including:
 - Displaying and interpreting a router's routing table
 - Configuring static and default routes
 - Enabling a WAN serial connection using High-Level Data Link Control (HDLC) or Point-to-Point Protocol (PPP), and configuring Frame Relay permanent virtual circuits (PVCs) on interfaces and subinterfaces
 - Configuring IP standard and extended access lists
 - Managing network device security

- Configuring network management protocols and managing device configurations and IOS images and licenses
- Verifying router configurations with available tools, such as **show** and **debug** commands
- Working knowledge of the TCP/IP stack, for both IPv4 and IPv6, and the ability to establish and troubleshoot Internet and WAN connectivity with both protocols
- The ability to configure, verify, and troubleshoot basic EIGRP and OSPF routing protocols, for both IPv4 and IPv6

If you lack this knowledge and these skills, you can gain them by completing the Interconnecting Cisco Network Devices Part 1 (ICND1) and Interconnecting Cisco Network Devices Part 2 (ICND2) courses or by reading the related Cisco Press books.

ROUTE Exam Topic Coverage

Cisco.com has the following information on the exam topics page for the ROUTE exam, exam number 300-101 (available at <http://www.cisco.com/web/learning/exams/list/route2.html#~Topics>):

“The following topics are general guidelines for the content that is likely to be included on the practical exam. However, other related topics may also appear on any specific delivery of the exam. In order to better reflect the contents of the exam and for clarity purposes, the following guidelines may change at any time without notice.”

The referenced list of exam topics available at the time of writing of this book is provided in Table I-1.

The Cisco ROUTE course does not cover all the listed exam topics, and may not cover other topics to the extent needed by the exam because of classroom time constraints. The Cisco ROUTE course is not created by the same group that created the exam.

This book does provide information on each of these exam topics (except when the topic is covered by prerequisite material as noted), as identified in the “Where Topic Is Covered” column in Table I-1. This book’s authors provided information related to all the exam topics to a depth that they believe should be adequate for the exam. Do note, though, that because the wording of the topics is quite general in nature and the exam itself is Cisco proprietary and subject to change, the authors of this book cannot guarantee that all the details on the exam are covered.

As mentioned, some of the listed ROUTE exam topics are actually covered by the prerequisite material. The authors believe that readers would already be familiar with this material and so have provided pointers to the relevant chapters of the ICND1 and ICND2 Foundation Learning Guide (ISBN 978-1587143762 and 978-1587143779) Cisco Press books for these topics.

Table I-1 *ROUTE Exam Topic Coverage*

Topic #	Topic	Where Topic Is Covered
1.0	Network Principles	
1.1	Identify Cisco Express Forwarding concepts	
	FIB	Chapter 1
	Adjacency table	Chapter 1
1.2	Explain general network challenges	
	Unicast	ICND1 Chapter 5
	Out-of-order packets	ICND1 Chapter 9 (sequencing)
	Asymmetric routing	Chapter 1
1.3	Describe IP operations	
	ICMP unreachable and redirects	Chapter 1, and IPv6 in ICND1 Chapter 20
	IPv4 and IPv6 fragmentation	IPv4 in Chapter 1, IPv6 in Chapter 6 and ICND1 Chapter 20
	TTL	ICND1 Chapter 7 and Glossary
1.4	Explain TCP operations	
	IPv4 and IPv6 (P)MTU	IPv4 in Chapter 1, IPv6 in Chapter 6
	MSS	Chapter 1
	Latency	ICND1 Chapter 1
	Windowing	ICND1 Chapter 9
	Bandwidth-delay product	Chapter 1
	Global synchronization	ICND1 Chapter 9
1.5	Describe UDP operations	
	Starvation	Chapter 1
	Latency	Chapter 1
1.6	Recognize proposed changes to the network	
	Changes to routing protocol parameters	Chapter 4
	Migrate parts of a network to IPv6	Chapter 6
	Routing protocol migration	Chapter 4
2.0	Layer 2 Technologies	
2.1	Configure and verify PPP	
	Authentication (PAP, CHAP)	Chapter 1
	PPPoE (client side only)	Chapter 1

Topic #	Topic	Where Topic Is Covered
2.2	Explain Frame Relay	
	Operations	Chapter 1
	Point-to-point	Chapters 1, 2, and 3
	Multipoint	Chapters 1, 2, and 3
3.0	Layer 3 Technologies	
3.1	Identify, configure, and verify IPv4 addressing and subnetting	
	Address types (unicast, broadcast, multicast, and VLSM)	Appendix B
	ARP	Appendix B
	DHCP relay and server	Chapter 6
	DHCP protocol operations	Chapters 6 and ICND1 Chapter 16
3.2	Identify IPv6 addressing and subnetting	
	Unicast	Chapter 1
	EUI-64	Chapters 6 and ICND1 Chapter 20
	ND, RS/RA	Chapter 1
	Autoconfig (SLAAC)	Chapter 6
	DHCP relay and server	Chapter 6
	DHCP protocol operations	Chapter 6
3.3	Configure and verify static routing	Chapter 1
3.4	Configure and verify default routing	Chapter 1
3.5	Evaluate routing protocol types	
	Distance vector	Chapter 1
	Link state	Chapter 1
	Path vector	Chapter 1
3.6	Describe administrative distance	Chapter 4
3.7	Troubleshoot passive interfaces	Chapters 2 and 3
3.8	Configure and verify VRF-lite	Chapter 8
3.9	Configure and verify filtering with any protocol	Chapter 4
3.10	Configure and verify redistribution between any routing protocols or routing sources	Chapter 4
3.11	Configure and verify manual and autosummarization with any routing protocol	Chapters 1, 2, and 3
3.12	Configure and verify policy-based routing	Chapter 4
3.13	Identify suboptimal routing	Chapter 4

Topic #	Topic	Where Topic Is Covered
3.14	Explain route maps	Chapter 4
3.15	Configure and verify loop prevention mechanisms	
	Route tagging and filtering	Chapter 4
	Split horizon	Chapters 1 and 2
	Route poisoning	Chapter 1
3.16	Configure and verify RIPv2	Chapter 1
3.17	Describe RIPng	Chapter 1
3.18	Describe EIGRP packet types	Chapter 2
3.19	Configure and verify EIGRP neighbor relationship and authentication	Chapters 2 and 8
3.20	Configure and verify EIGRP stubs	Chapter 2
3.21	Configure and verify EIGRP load balancing	
	Equal cost	Chapter 2
	Unequal cost	Chapter 2
3.22	Describe and optimize EIGRP metrics	Chapter 2
3.23	Configure and verify EIGRP for IPv6	Chapter 2
3.24	Describe OSPF packet types	Chapter 3
3.25	Configure and verify OSPF neighbor relationship and authentication	Chapters 3 and 8
3.26	Configure and verify OSPF network types, area types, and router types	
	Point-to-point, multipoint, broadcast, nonbroadcast	Chapter 3
	LSA types, area type: backbone, normal, transit, stub, NSSA, totally stub	Chapter 3
	Internal router, backbone router, ABR, ASBR	Chapter 3
	Virtual link	Chapter 3
3.27	Configure and verify OSPF path preference	Chapter 3
3.28	Configure and verify OSPF operations	Chapter 3
3.29	Configure and verify OSPF for IPv6	Chapter 3
3.30	Describe, configure, and verify BGP peer relationships and authentication	
	Peer group	Chapter 7
	Active, passive	Chapter 7 (But there is no “passive” in BGP; it’s “established.”)
	States and timers	Chapter 7

Topic #	Topic	Where Topic Is Covered
3.31	Configure and verify eBGP (IPv4 and IPv6 address families)	
	eBGP	Chapter 7
	4-byte AS number	Chapter 6
	Private AS	Chapter 6
3.32	Explain BGP attributes and best-path selection	Chapter 7
4.0	VPN Technologies	
4.1	Configure and verify GRE	Chapter 1 for GRE tunnels; configuration and verification in ICND2 Chapter 5.
4.2	Describe DMVPN (single hub)	Chapter 1
4.3	Describe Easy Virtual Networking (EVN)	Chapter 8
5.0	Infrastructure Security	
5.1	Describe IOS AAA using local database	Chapter 8
5.2	Describe device security using IOS AAA with TACACS+ and RADIUS	
	AAA with TACACS+ and RADIUS	Chapter 8
	Local privilege authorization fallback	Chapter 8
5.3	Configure and verify device access control	
	Lines (VTY, AUX, console)	Chapter 8
	Management plane protection	Chapter 8
	Password encryption	Chapter 8
5.4	Configure and verify router security features	
	IPv4 access control lists (standard, extended, time-based)	Appendix B
	IPv6 traffic filter	Chapter 6
	Unicast reverse path forwarding	Chapter 8
6.0	Infrastructure Services	
6.1	Configure and verify device management	
	Console and vty	Chapter 8
	Telnet, HTTP, HTTPS, SSH, SCP	Chapter 8
	(T)FTP	Chapter 8

Topic #	Topic	Where Topic Is Covered
6.2	Configure and verify SNMP	
	v2	Chapter 8 and ICND2 Chapter 6
	v3	Chapter 8 and ICND2 Chapter 6
6.3	Configure and verify logging	
	Local logging, syslog, debugs, conditional debugs	Chapter 8 and ICND2 Chapter 6
	Timestamps	ICND2 Chapter 6
6.4	Configure and verify Network Time Protocol	
	NTP master, client, version 3, version 4	Chapter 8
	NTP authentication	Chapter 8
6.5	Configure and verify IPv4 and IPv6 DHCP	
	DHCP Client, IOS DHCP server, DHCP relay	Chapter 6
	DHCP options (describe)	Chapter 6
6.6	Configure and verify IPv4 Network Address Translation	
	Static NAT, dynamic NAT, PAT	Chapter 6
6.7	Describe IPv6 NAT	
	NAT64	Chapter 6
	NPTv6	Chapter 6
6.8	Describe SLA architecture	Chapter 5
6.9	Configure and verify IP SLA	
	ICMP	Chapter 5
6.10	Configure and verify tracking objects	
	Tracking object	Chapter 5
	Tracking different entities (for example, interfaces, IP SLA results)	Chapter 5
6.11	Configure and verify Cisco NetFlow	
	NetFlow v5, v9	ICND2 Chapter 6
	Local retrieval	ICND2 Chapter 6
	Export (configuration only)	ICND2 Chapter 6

How This Book Is Organized

The chapters and appendixes in this book are as follows:

- Chapter 1, “Basic Network and Routing Concepts,” begins with an overview of routing protocols that focuses on characteristics that describe their differences. It describes how limitations of different underlying technologies affect routing protocols, followed by a closer look at how Layer 2 and Layer 3 VPNs, including Dynamic Multipoint Virtual Private Network (DMVPN), affect routing protocols. RIPv2 and RIPv6 configuration are covered.
- Chapter 2, “EIGRP Implementation,” explains EIGRP neighbor relationships and how EIGRP chooses the best path through the network. Configuration of stub routing, route summarization, and load balancing with EIGRP are covered. Basic EIGRP for IPv6, including with route summarization is covered. The chapter concludes with a discussion of a new way of configuring EIGRP for both IPv4 and IPv6: named EIGRP.
- Chapter 3, “OSPF Implementation,” introduces basic OSPF and OSPF adjacencies, and explains how OSPF builds the routing table. OSPF summarization and stub areas are covered. The chapter concludes with the configuration of OSPFv3 using address families for IPv6 and IPv4.
- Chapter 4, “Manipulating Routing Updates,” discusses network performance issues related to routing and using multiple IP routing protocols on a network. Implementing route redistribution between different routing protocols is described, and methods of controlling the routing information sent between these routing protocols are explored, including using distribute lists, prefix lists, and route maps.
- Chapter 5, “Path Control Implementation,” starts by discussing the Cisco Express Forwarding (CEF) switching method. Path control fundamentals are explored, and two path control tools are detailed: policy-based routing (PBR) and Cisco IOS IP service-level agreements (SLAs).
- Chapter 6, “Enterprise Internet Connectivity,” describes how enterprises can connect to the Internet, which has become a vital resource for most organizations. Planning for a single connection to an Internet service provider (ISP), or redundant connections to multiple ISPs, is a very important task, and is covered first in the chapter. The details of single connections for IPv4 and IPv6 are then described. The chapter concludes with a discussion of using multiple ISP connections to improve Internet connectivity resilience.
- Chapter 7, “BGP Implementation,” describes how enterprises can use BGP when connecting to the Internet. This chapter introduces BGP terminology, concepts, and operation, and provides BGP configuration, verification, and troubleshooting techniques. The chapter describes BGP attributes and how they are used in the path selection process, and also introduces route maps for manipulating BGP path attributes and filters for BGP routing updates. The chapter concludes with a section on how BGP is used for IPv6 Internet connectivity.

- Chapter 8, “Routers and Routing Protocol Hardening,” discusses how to secure the management plane of Cisco routers using recommended practices. The benefits of routing protocol authentication are described and configuration of routing authentication for EIGRP, OSPF, and BGP is presented. The chapter concludes with Cisco VRF-lite and Easy Virtual Networking (EVN).
- Appendix A, “Answers to End of Chapter Review Questions,” contains the answers to the review questions that appear at the end of each chapter.
- Appendix B, “IPv4 Supplement,” provides job aids and supplementary information that are intended for your use when working with IPv4 addresses. Topics include a subnetting job aid, a decimal-to-binary conversion chart, an IPv4 addressing review, an IPv4 access lists review, IP address planning, hierarchical addressing using variable-length subnet masks (VLSMs), route summarization, and classless interdomain routing (CIDR).
- Appendix C, “BGP Supplement,” provides supplementary information on BGP covering the following topics: BGP route summarization, redistribution with interior gateway protocols (IGPs), communities, route reflectors, advertising a default route, and not advertising private autonomous system numbers.
- Appendix D, “Acronyms and Abbreviations” identifies abbreviations, acronyms, and initialisms used in this book.

This page intentionally left blank

OSPF Implementation

This chapter covers the following topics:

- Basic OSPF Configuration and OSPF Adjacencies
- How OSPF Builds the Routing Table
- Configuration of Summarization and Stub Areas in OSPF
- Configuration of OSPFv3 for IPv6 and IPv4

This chapter examines the Open Shortest Path First (OSPF) Protocol, one of the most commonly used interior gateway protocols in IP networking. OSPFv2 is an open-standard protocol that provides routing for IPv4. OSPFv3 offers some enhancements for IP Version 6 (IPv6). OSPF is a complex protocol that is made up of several protocol handshakes, database advertisements, and packet types.

OSPF is an interior gateway routing protocol that uses link-states rather than distance vectors for path selection. OSPF propagates link-state advertisements (LSAs) rather than routing table updates. Because only LSAs are exchanged instead of the entire routing tables, OSPF networks converge in a timely manner.

OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router-usable interfaces and reachable neighbors.

Establishing OSPF Neighbor Relationships

OSPF is a link-state protocol based on the open standard. At a high level, OSPF operation consists of three main elements: neighbor discovery, link-state information exchange, and best-path calculation.

To calculate the best path, OSPF uses the shortest path first (SPF) or Dijkstra's algorithm. The input information for SPF calculation is link-state information, which is exchanged

between routers using several different OSPF message types. These message types help improve convergence and scalability in multi-area OSPF deployments.

OSPF also supports several different network types, which enables you to configure OSPF over a variety of different underlying network technologies.

Upon completion of this section, you will be able to describe the main operational characteristics of the OSPF protocol and configure its basic features. You will also be able to meet following objectives:

- Explain why would you choose OSPF over other routing protocols
- Describe basic operation steps with link-state protocols
- Describe area and router types in OSPF
- Explain what the design limitations of OSPF are
- List and describe OSPF message types
- Describe OSPF neighbor relationship over point-to-point link
- Describe OSPF neighbor relationship behavior on MPLS VPN
- Describe OSPF neighbor relationship behavior over L2 MPLS VPN
- List and describe OSPF neighbor states
- List and describe OSPF network types
- Configure passive interfaces

OSPF Features

OSPF was developed by the Internet Engineering Task Force (IETF) to overcome the limitations of distance vector routing protocols. One of the main reasons why OSPF is largely deployed in today's enterprise networks is the fact that it is an open standard; OSPF is not proprietary. Version 1 of the protocol is described in the RFC 1131. The current version used for IPv4, Version 2, is specified in RFCs 1247 and 2328. OSPF Version 3, which is used in IPv6 networks, is specified in RFC 5340.

OSPF offers a large level of scalability and fast convergence. Despite its relatively simple configuration in small and medium-size networks, OSPF implementation and troubleshooting in large-scale networks can at times be challenging.

The key features of the OSPF protocol are as follows:

- **Independent transport:** OSPF works on top of IP and uses protocol number 89. It does not rely on the functions of the transport layer protocols TCP or UDP.
- **Efficient use of updates:** When an OSPF router first discovers a new neighbor, it sends a full update with all known link-state information. All routers within an OSPF area must have identical and synchronized link-state information in their OSPF

link-state databases. When an OSPF network is in a converged state and a new link comes up or a link becomes unavailable, an OSPF router sends only a partial update to all its neighbors. This update will then be flooded to all OSPF routers within an area.

- **Metric:** OSPF uses a metric that is based on the cumulative costs of all outgoing interfaces from source to destination. The interface cost is inversely proportional to the interface bandwidth and can be also set up explicitly.
- **Update destination address:** OSPF uses multicast and unicast, rather than broadcast, for sending messages. The IPv4 multicast addresses used for OSPF are 224.0.0.5 to send information to all OSPF routers and 224.0.0.6 to send information to DR/BDR routers. The IPv6 multicast addresses are FF02::5 for all OSPFv3 routers and FF02::6 for all DR/BDR routers. If the underlying network does not have broadcast capabilities, you must establish OSPF neighbor relationships using a unicast address. For IPv6, this address will be a link-local IPv6 address.
- **VLSM support:** OSPF is a classless routing protocol. It supports variable-length subnet masking (VLSM) and discontinuous networks. It carries subnet mask information in the routing updates.
- **Manual route summarization:** You can manually summarize OSPF interarea routes at the Area Border Router (ABR), and you have the possibility to summarize OSPF external routes at the Autonomous System Boundary Router (ASBR). OSPF does not know the concept of autosummarization.
- **Authentication:** OSPF supports clear-text, MD5, and SHA authentication.

Note The term *IP* is used for generic IP and applies to both IPv4 and IPv6. Otherwise, the terms *IPv4* and *IPv6* are used for the specific protocols.

Note Although there is some review, this chapter assumes that you have basic CCNA knowledge of OSPF. If you need a more thorough review of OSPF or other routing protocols, see the *Routing Protocols Companion Guide* (Cisco Press, 2014).

OSPF Operation Overview

To create and maintain routing information, OSPF routers complete the following generic link-state routing process, shown in Figure 3-1, to reach a state of convergence:

1. **Establish neighbor adjacencies:** OSPF-enabled routers must form adjacencies with their neighbor before they can share information with that neighbor. An OSPF-enabled router sends Hello packets out all OSPF-enabled interfaces to determine whether neighbors are present on those links. If a neighbor is present, the OSPF-enabled router attempts to establish a neighbor adjacency with that neighbor.

2. **Exchange link-state advertisements:** After adjacencies are established, routers then exchange link-state advertisements (LSAs). LSAs contain the state and cost of each directly connected link. Routers flood their LSAs to adjacent neighbors. Adjacent neighbors receiving the LSA immediately flood the LSA to other directly connected neighbors, until all routers in the area have all LSAs.
3. **Build the topology table:** After the LSAs are received, OSPF-enabled routers build the topology table (LSDB) based on the received LSAs. This database eventually holds all the information about the topology of the network. It is important that all routers in the area have the same information in their LSDBs.
4. **Execute the SPF algorithm:** Routers then execute the SPF algorithm. The SPF algorithm creates the SPF tree.
5. **Build the routing table:** From the SPF tree, the best paths are inserted into the routing table. Routing decisions are made based on the entries in the routing table.

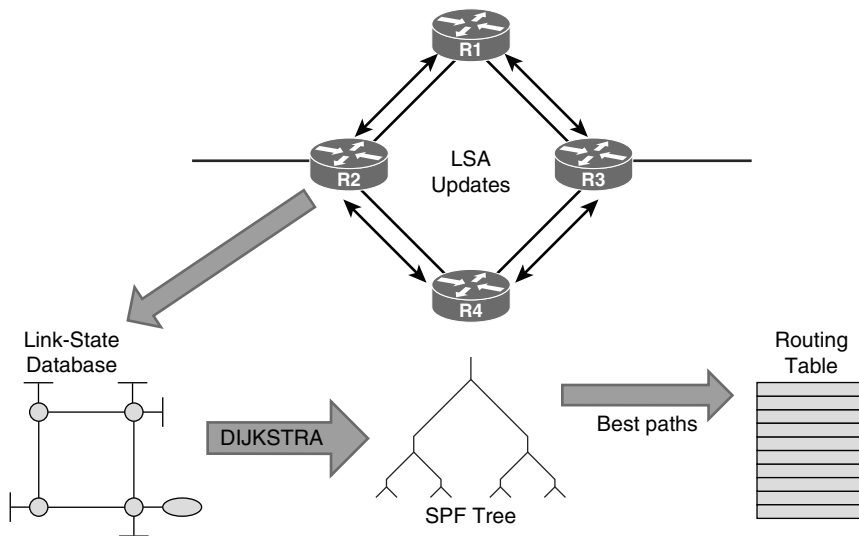


Figure 3-1 OSPF Operation

Hierarchical Structure of OSPF

If you run OSPF in a simple network, the number of routers and links are relatively small, and best paths to all destinations are easily deduced. However, the information necessary to describe larger networks with many routers and links can become quite complex. SPF calculations that compare all possible paths for routes can easily turn into a complex and time-consuming calculation for the router.

One of the main methods to reduce this complexity and the size of the link-state information database is to partition the OSPF routing domain into smaller units called *areas*, shown in Figure 3-2. This also reduces the time it takes for the SPF algorithm to

execute. All OSPF routers within an area must have identical entries within their respective LSDBs. Inside an area, routers exchange detailed link-state information. However, information transmitted from one area into another contains only summary details of the LSDB entries and not topology details about the originating area. These summary LSAs from another area are injected directly into the routing table and without making the router rerun its SPF algorithm.

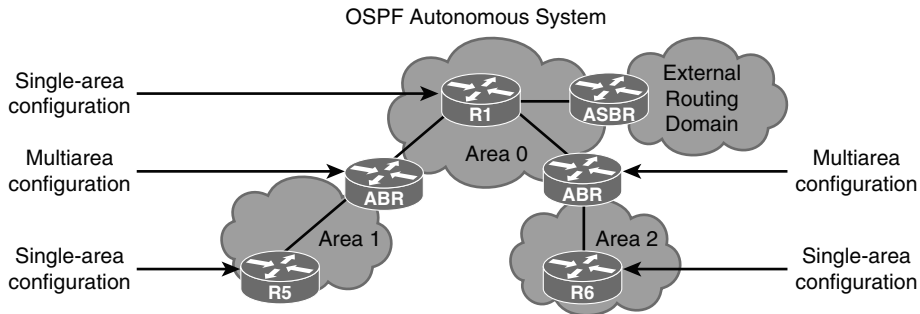


Figure 3-2 *OSPF Hierarchy*

OSPF uses a two-layer area hierarchy:

- **Backbone area, transit area or area 0:** Two principal requirements for the backbone area are that it must connect to all other nonbackbone areas and this area must be always contiguous; it is not allowed to have split up the backbone area. Generally, end users are not found within a backbone area.
- **Nonbackbone area:** The primary function of this area is to connect end users and resources. Nonbackbone areas are usually set up according to functional or geographic groupings. Traffic between different nonbackbone areas must always pass through the backbone area.

In the multi-area topology there are some special commonly used OSPF terms:

- **ABR:** A router that has interfaces connected to at least two different OSPF areas, including the backbone area. ABRs contain LSDB information for each area, make route calculation for each area and advertise routing information between areas.
- **ASBR:** ASBR is a router that has at least one of its interfaces connected to an OSPF area and at least one of its interfaces connected to an external non-OSPF domain.
- **Internal router:** A router that has all its interfaces connected to only one OSPF area. This router is completely internal to this area.
- **Backbone router:** A router that has at least one interface connected to the backbone area.

The optimal number of routers per area varies based on factors such as network stability, but in general it is recommended to have no more than 50 routers per single area.

Design Restrictions of OSPF

OSPF has special restrictions when multiple areas are configured in an OSPF routing domain or AS, as shown in Figure 3-3. If more than one area is configured, known as *multi-area OSPF*, one of these areas must be area 0. This is called the *backbone area*. When designing networks or starting with a single area, it is good practice to start with the core layer, which becomes area 0, and then expand into other areas later.

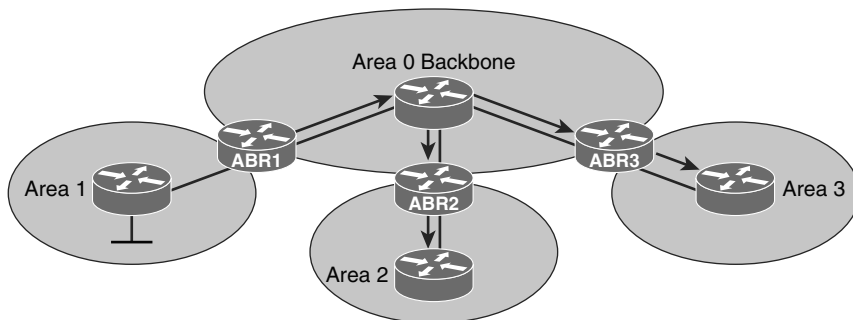


Figure 3-3 Multi-Area OSPF

The backbone has to be at the center of all other areas, and other areas have to be connected to the backbone. The main reason is that OSPF expects all areas to inject routing information into the backbone area, which distributes that information into other areas.

Another important requirement for the backbone area is that it must be contiguous. In other words, splitting up area 0 is not allowed.

However, in some cases, these two conditions cannot be met. Later in this chapter in the section, “OSPF Virtual Links,” you will learn about the use of virtual links as a solution.

OSPF Message Types

OSPF uses five types of routing protocol packets, which share a common protocol header. Every OSPF packet is directly encapsulated in the IP header. The IP protocol number for OSPF is 89.

- **Type 1: Hello packet:** Hello packets are used to discover, build, and maintain OSPF neighbor adjacencies. To establish adjacency, OSPF peers at both sides of the link must agree on some parameters contained in the Hello packet to become OSPF neighbors.
- **Type 2: Database Description (DBD) packet:** When the OSPF neighbor adjacency is already established, a DBD packet is used to describe LSDB so that routers can compare whether databases are in sync.
- **Type 3: Link-State Request (LSR) packet:** When the database synchronization process is over, the router might still have a list of LSAs that are missing in its database. The router will send an LSR packet to inform OSPF neighbors to send the most recent version of the missing LSAs.

- **Type 4: Link-State Update (LSU) packet:** There are several types of LSUs, known as LSAs. LSU packets are used for the flooding of LSAs and sending LSA responses to LSR packets. It is sent only to the directly connected neighbors who have previously requested LSAs in the form of LSR packet. In case of flooding, neighbor routers are responsible for re-encapsulation of received LSA information in new LSU packets.
- **Type 5: Link-State Acknowledgment (LSAck) packet:** LSAs are used to make flooding of LSAs reliable. Each LSA received must be explicitly acknowledged. Multiple LSAs can be acknowledged in a single LSAck packet.

Basic OSPF Configuration

This section explores how to configure and establish OSPF neighbor relationship. You will observe the impact of the interface MTU and OSPF hello/dead timer parameters on the OSPF neighbor relationship formation. In addition, you will learn what the roles are of the DR/BDR routers and how to control the DR/BDR election process.

The topology in Figure 3-4 shows five routers, R1 to R5. R1, R4, and R5 are already pre-configured, while R2 and R3 will be configured in this section.

R1, R4, and R5 are connected to common multiaccess Ethernet segment. R1 and R2 are connected over serial Frame Relay interface, and R1 and R3 are also connected over Ethernet link.

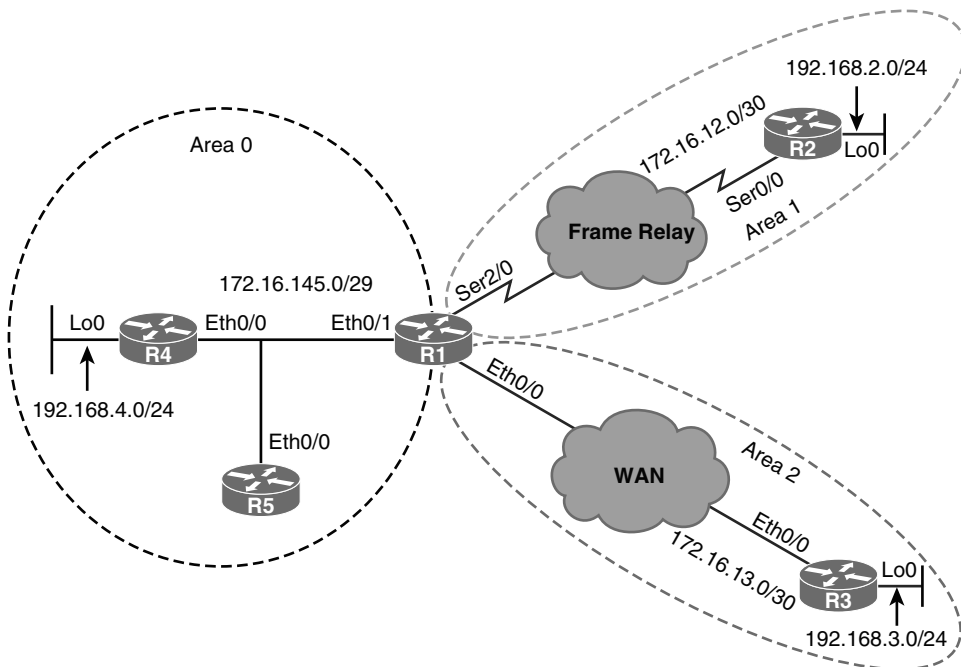


Figure 3-4 Topology for Basic OSPF Configuration

Example 3-1 begins the configuration of OSPF on WAN and LAN interfaces on R2 and R3. Use the process numbers 2 and 3 on R2 and R3, respectively.

Example 3-1 Configuration OSPF on R2 and R3

```
R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# router ospf 2
R2(config-router)# network 172.16.12.0 0.0.0.3 area 1
R2(config-router)# network 192.168.2.0 0.0.0.255 area 1

R3# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# router ospf 3
R3(config-router)# network 172.16.13.0 0.0.0.3 area 2
R3(config-router)# network 192.168.3.0 0.0.0.255 area 2
```

To enable the OSPF process on the router, use the **router ospf process-id** command. Process ID numbers between neighbors do not need to match for the routers to establish an OSPF adjacency. The OSPF process number ID is an internally used identification parameter for an OSPF routing process and only has local significance. However, it is good practice to make the process ID number the same on all routers. If necessary, you can specify multiple OSPF routing processes on a router, but you need to know the implications of doing so. Multiple OSPF processes on the same router is not common and beyond the scope of this book.

To define which interfaces will run the OSPF process and to define the area ID for those interfaces, use **network ip-address wildcard-mask area area-id** command. A combination of *ip-address* and *wildcard-mask* together allows you to define one or multiple interfaces to be associated with a specific OSPF area using a single command.

Cisco IOS Software sequentially evaluates the *ip-address wildcard-mask* pair specified in the **network** command for each interface as follows:

- It performs a logical OR operation between a *wildcard-mask* argument and the interface's primary IP address.
- It performs a logical OR operation between a *wildcard-mask* argument and the *ip-address* argument in the network command.
- The software compares the two resulting values. If they match, OSPF is enabled on the associated interface, and this interface is attached to the OSPF area specified.

This area ID is a 32-bit number that may be represented in integer or dotted-decimal format. When represented in dotted-decimal format, the area ID does not represent an IP address; it is only a way of writing an integer value in dotted-decimal format. For example, you may specify that an interface belongs to area 1 using **area 1** or **area 0.0.0.1** notation in the **network** command. To establish OSPF full adjacency, two neighbor routers must be in the same area. Any individual interface can only be attached to a single

area. If the address ranges specified for different areas overlap, IOS will adopt the first area in the **network** command list and ignore subsequent overlapping portions. To avoid conflicts, you must pay special attention to ensure that address ranges do not overlap.

In Example 3-2, the OSPF router IDs of R2 and R3 are configured using the **router-id** command.

Example 3-2 Configuration of OSPF Router IDs

```
R2(config-router)# router-id 2.2.2.2
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect

R3(config-router)# router-id 3.3.3.3
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect
```

The OSPF router ID is a fundamental parameter for the OSPF process. For the OSPF process to start, Cisco IOS must be able to identify a unique OSPF router ID. Similar to EIGRP, the OSPF router ID is a 32-bit value expressed as an IPv4 address. At least one primary IPv4 address on an interface in the up/up state must be configured for a router to be able to choose router ID; otherwise, an error message is logged, and the OSPF process does not start.

To choose the OSPF router ID at the time of OSPF process initialization, the router uses the following criteria:

1. Use the router ID specified in the **router-id ip-address** command. You can configure an arbitrary value in the IPv4 address format, but this value must be unique. If the IPv4 address specified with the **router-id** command overlaps with the router ID of another already-active OSPF process, the **router-id** command fails.
2. Use the highest IPv4 address of all active loopback interfaces on the router.
3. Use the highest IPv4 address among all active nonloopback interfaces.

After the three-step OSPF router ID selection process has finished, and if the router is still unable to select an OSPF router ID, an error message will be logged. An OSPF process that failed to select a router ID retries the selection process every time an IPv4 address becomes available. (An applicable interface changes its state to up/up or an IPv4 address is configured on an applicable interface.)

In Example 3-3, the OSPF routing process is cleared on R2 and R3 for the manually configured router ID to take effect.

Example 3-3 Clearing the OSPF Processes on R2 and R3

```
R2# clear ip ospf process
Reset ALL OSPF processes? [no]: yes
R2#
*Nov 24 08:37:24.679: %OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/0 from
```

```

FULL to DOWN, Neighbor Down: Interface down or detached
R2#
*Nov 24 08:39:24.734: %OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/0 from
LOADING to FULL, Loading Done

R3# clear ip ospf 3 process
Reset OSPF process 3? [no]: yes
R3#
*Nov 24 09:06:00.275: %OSPF-5-ADJCHG: Process 3, Nbr 1.1.1.1 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
R3#
*Nov 24 09:06:40.284: %OSPF-5-ADJCHG: Process 3, Nbr 1.1.1.1 on Ethernet0/0 from
LOADING to FULL, Loading Done

```

Once an OSPF router ID is selected, it is not changed even if the interface that is used to select it changed its operational state or its IP address. To change the OSPF router ID, you must reset the OSPF process with the **clear ip ospf process** command or reload the router.

In production networks, the OSPF router ID cannot be changed easily. Changing the OSPF router ID requires reset of all OSPF adjacencies, resulting in a temporary routing outage. The router also has to originate new copies of all originating LSAs with the new router ID.

You can either clear the specific OSPF process by specifying the process ID, or you can reset all OSPF processes by using the **clear ip ospf process** command.

The newly configured OSPF router ID is verified on R2 and R3 using **show ip protocols** commands in Example 3-4. Large output of this command can optionally be filtered using the pipe function, also shown in Example 3-4.

Example 3-4 Verifying the Router IDs on R2 and R3

```

R2# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 2"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 2.2.2.2
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.12.0 0.0.0.3 area 1
    192.168.2.0 0.0.0.255 area 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    1.1.1.1          110          00:02:55
  Distance: (default is 110)

```

```
R3# show ip protocols | include ID
Router ID 3.3.3.3
```

The OSPF neighborship on R2 and R3 is verified in Example 3-5 using the `show ip ospf neighbor` command.

Example 3-5 Verifying OSPF Neighborships on R2 and R3

```
R2# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/DR	00:01:57	172.16.12.1	Serial0/0

```
R3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/DR	00:00:39	172.16.13.1	Ethernet0/0

The command `show ip ospf neighbor` displays OSPF neighbor information on a per-interface basis. The significant fields of the outputs are as follows:

- **Neighbor ID:** Represents neighbor router ID.
- **Priority:** Priority on the neighbor interface used for the DR/BDR election.
- **State:** A Full state represents the final stage of OSPF neighbor establishment process and denotes that the local router has established full neighbor adjacency with the remote OSPF neighbor. DR means that DR/BDR election process has been completed and that the remote router with the router ID 1.1.1.1 has been elected as the designated router (DR).
- **Dead Time:** Represents value of the dead timer. When this timer expires, the router terminates the neighbor relationship. Each time a router receives an OSPF Hello packet from a specific neighbor, it resets the dead timer back to its full value.
- **Address:** Primary IPv4 address of the neighbor router.
- **Interface:** Local interface over which an OSPF neighbor relationship is established.

Example 3-6 verifies the OSPF-enabled interfaces on R2 and R3 using the `show ip ospf interface` command.

Example 3-6 Verifying the OSPF-Enabled Interfaces on R2 and R3

```
R2# show ip ospf interface
```

```
Loopback0 is up, line protocol is up
```

```
Internet Address 192.168.2.1/24, Area 1, Attached via Network Statement
```

```
Process ID 2, Router ID 2.2.2.2, Network Type LOOPBACK, Cost: 1
```

```
<Output omitted>
```

```
Serial0/0 is up, line protocol is up
```

```

Internet Address 172.16.12.2/30, Area 1, Attached via Network Statement
Process ID 2, Router ID 2.2.2.2, Network Type NON_BROADCAST, Cost: 64
<Output omitted>

```

```

R3# show ip ospf interface
Loopback0 is up, line protocol is up
Internet Address 192.168.3.1/24, Area 2, Attached via Network Statement
Process ID 3, Router ID 3.3.3.3, Network Type LOOPBACK, Cost: 1
<Output omitted>
Ethernet0/0 is up, line protocol is up
Internet Address 172.16.13.2/30, Area 2, Attached via Network Statement
Process ID 3, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 10
<Output omitted>

```

Output of the `show ip ospf interface` command shows you all interfaces enabled in the OSPF process. For each enabled interface, you can see detailed information such as OSPF area ID, OSPF process ID, and how the interface was included into the OSPF process. In the output, you can see that both interfaces on both routers were included via the `network` statement, configured with the `network` command.

In Example 3-7, the OSPF routes are verified in the routing table on R5 using the `show ip route ospf` command.

Example 3-7 Verifying the OSPF Routes on R5

```

R5# show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA 172.16.12.0/30 [110/74] via 172.16.145.1, 00:39:00, Ethernet0/0
O IA 172.16.13.0/30 [110/20] via 172.16.145.1, 00:19:29, Ethernet0/0
192.168.2.0/32 is subnetted, 1 subnets
O IA 192.168.2.1 [110/75] via 172.16.145.1, 00:07:27, Ethernet0/0
192.168.3.0/32 is subnetted, 1 subnets
O IA 192.168.3.1 [110/21] via 172.16.145.1, 00:08:30, Ethernet0/0
O 192.168.4.0/24 [110/11] via 172.16.145.4, 00:39:10, Ethernet0/0

```

Among the routes originated within the OSPF autonomous system, OSPF clearly distinguishes two types of routes: intra-area routes and interarea routes. Intra-area routes are routes that are originated and learned in the same local area. Code for the intra-area routes in the routing table is O. The second type is interarea routes, which originate in other areas and are inserted into the local area to which your router belongs. Code for the interarea routes in the routing table is O IA. Interarea routes are inserted into other areas on the ABR.

The prefix 192.168.4.0/24 is an example of intra-area route from the R5 perspective. It originated from router R4, which is part of the area 0, the same area as R5.

Prefixes from R2 and R3, which are part of area 1 and area 2, are shown in the routing table on R5 as interarea routes. Prefixes were inserted into area 0 as interarea routes by R1, which plays the role of ABR.

Prefixes 192.168.2.0/24 and 192.168.3.0/24 configured on the loopback interfaces of R2 and R3 are displayed in the R5 routing table as host routes 192.168.2.1/32 and 192.168.3.1/32. By default, OSPF will advertise any subnet configured on the loopback interface as /32 host route. To change this default behavior, you can optionally change OSPF network type on the loopback interface from the default loopback to point-to-point using the `ip ospf network point-to-point` interface command.

OSPF database routes on R5 are observed in Example 3-8 using the `show ip ospf route` command.

Example 3-8 OSPF Routes on R5

```
R5# show ip ospf route

      OSPF Router with ID (5.5.5.5) (Process ID 1)

      Base Topology (MTID 0)

      Area BACKBONE(0)

      Intra-area Route List
* 172.16.145.0/29, Intra, cost 10, area 0, Connected
   via 172.16.145.5, Ethernet0/0
*> 192.168.4.0/24, Intra, cost 11, area 0
   via 172.16.145.4, Ethernet0/0

      Intra-area Router Path List
i 1.1.1.1 [10] via 172.16.145.1, Ethernet0/0, ABR, Area 0, SPF 2

      Inter-area Route List
*> 192.168.2.1/32, Inter, cost 75, area 0
   via 172.16.145.1, Ethernet0/0
*> 192.168.3.1/32, Inter, cost 21, area 0
```



```

    via 172.16.145.1, Ethernet0/0
*> 172.16.12.0/30, Inter, cost 74, area 0
    via 172.16.145.1, Ethernet0/0
*> 172.16.13.0/30, Inter, cost 20, area 0
    via 172.16.145.1, Ethernet0/0

```

The **show ip ospf route** command clearly separates the lists of intra-area and interarea routes. In addition, output of the command displays essential information about ABRs, including the router ID, IPv4 address in the current area, interface that advertises routes into the area, and the area ID.

For interarea routes, the metric for the route (cost), the area into which the route is distributed, and the interface over which the route is inserted are displayed.

In Example 3-9, the OSPF neighbor adjacency and the associated OSPF packet types on R3 are observed using the **debug ip ospf adj** and **clear ip ospf process** commands. Disable **debug** when the OSPF session is reestablished.

Example 3-9 Observing Formation of OSPF Neighbor Adjacencies

```

R3# debug ip ospf adj
OSPF adjacency debugging is on
R3# clear ip ospf process
Reset ALL OSPF processes? [no]: yes
*Jan 17 13:02:37.394: OSPF-3 ADJ   Lo0: Interface going Down
*Jan 17 13:02:37.394: OSPF-3 ADJ   Lo0: 3.3.3.3 address 192.168.3.1 is dead, state
DOWN
*Jan 17 13:02:37.394: OSPF-3 ADJ   Et0/0: Interface going Down
*Jan 17 13:02:37.394: OSPF-3 ADJ   Et0/0: 1.1.1.1 address 172.16.13.1 is dead, state
DOWN
*Jan 17 13:02:37.394: %OSPF-5-ADJCHG: Process 3, Nbr 1.1.1.1 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
<Output omitted>
*Jan 17 13:02:37.394: OSPF-3 ADJ   Lo0: Interface going Up
*Jan 17 13:02:37.394: OSPF-3 ADJ   Et0/0: Interface going Up
*Jan 17 13:02:37.395: OSPF-3 ADJ   Et0/0: 2 Way Communication to 1.1.1.1, state 2WAY
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Backup seen event before WAIT timer
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: DR/BDR election
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Elect BDR 3.3.3.3
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Elect DR 1.1.1.1
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Elect BDR 3.3.3.3
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Elect DR 1.1.1.1
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: DR: 1.1.1.1 (Id) BDR: 3.3.3.3 (Id)
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Nbr 1.1.1.1: Prepare dbase exchange
*Jan 17 13:02:37.396: OSPF-3 ADJ   Et0/0: Send DBD to 1.1.1.1 seq 0x95D opt 0x52
flag 0x7 len 32
*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: Rcv DBD from 1.1.1.1 seq 0x691 opt 0x52
flag 0x7 len 32 mtu 1500 state EXSTART

```

```

*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: First DBD and we are not SLAVE
*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: Rcv DBD from 1.1.1.1 seq 0x95D opt 0x52
      flag 0x2 len 152 mtu 1500 state EXSTART
*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: NBR Negotiation Done. We are the MASTER
*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: Nbr 1.1.1.1: Summary list built, size 0
*Jan 17 13:02:37.397: OSPF-3 ADJ   Et0/0: Send DBD to 1.1.1.1 seq 0x95E opt 0x52
      flag 0x1 len 32
*Jan 17 13:02:37.398: OSPF-3 ADJ   Et0/0: Rcv DBD from 1.1.1.1 seq 0x95E opt 0x52
      flag 0x0 len 32 mtu 1500 state EXCHANGE
*Jan 17 13:02:37.398: OSPF-3 ADJ   Et0/0: Exchange Done with 1.1.1.1
*Jan 17 13:02:37.398: OSPF-3 ADJ   Et0/0: Send LS REQ to 1.1.1.1 length 96 LSA count
      6
*Jan 17 13:02:37.399: OSPF-3 ADJ   Et0/0: Rcv LS UPD from 1.1.1.1 length 208 LSA
      count 6
*Jan 17 13:02:37.399: OSPF-3 ADJ   Et0/0: Synchronized with 1.1.1.1, state FULL
*Jan 17 13:02:37.399: %OSPF-5-ADJCHG: Process 3, Nbr 1.1.1.1 on Ethernet0/0 from
      LOADING to FULL, Loading Done
R3# undebug all

```

An OSPF adjacency is established in several steps. In the first step, routers that intend to establish full OSPF neighbor adjacency exchange OSPF Hello packets. Both OSPF neighbors are in the Down state, the initial state of a neighbor conversation that indicates that no Hello's have been heard from the neighbor. When a router receives a Hello from the neighbor but has not yet seen its own router ID in the neighbor Hello packet, it will transit to the Init state. In this state, the router will record all neighbor router IDs and start including them in Hellos sent to the neighbors. When the router sees its own router ID in the Hello packet received from the neighbor, it will transit to the 2-Way state. This means that bidirectional communication with the neighbor has been established.

On broadcast links, OSPF neighbors first determine the designated router (DR) and backup designated router (BDR) roles, which optimize the exchange of information in broadcast segments.

In the next step, routers start to exchange content of OSPF databases. The first phase of this process is to determine master/slave relationship and choose the initial sequence number for adjacency formation. To accomplish this, routers exchange DBD packets. When the router receives the initial DBD packet it transitions the state of the neighbor from which this packet is received to ExStart state, populates its Database Summary list with the LSAs that describe content of the neighbor's database, and sends its own empty DBD packet. In the DBD exchange process, the router with the higher router ID will become master, and it will be the only router that can increment sequence numbers.

With master/slave selection complete, database exchange can start. R3 will transit R1's neighbor state to Exchange. In this state, R3 describes its database to the R1 by sending DBD packets that contain the headers of all LSAs in the Database Summary list. The Database Summary list describes all LSAs in the router's database, but not the full content of the OSPF database. To describe the content of the database, one or multiple DBD packets may be exchanged. A router compares the content of its own Database

Summary list with the list received from the neighbor, and if there are differences, it adds missing LSAs to the Link State Request list. At this point, routers enter the Loading state. R3 sends an LSR packet to the neighbor requesting full content of the missing LSAs from the LS Request list. R1 replies with the LSU packets, which contain full versions of the missing LSAs.

Finally, when neighbors have a complete version of the LSDB, both neighbors transit to the Full state, which means that databases on the routers are synchronized and that neighbors are fully adjacent.

Optimizing OSPF Adjacency Behavior

Multiaccess networks, either broadcast (such as Ethernet) or nonbroadcast (such as Frame Relay), represent interesting issues for OSPF. All routers sharing the common segment will be part of the same IP subnet. When forming adjacency on multiaccess network, every router will try to establish full OSPF adjacency with all other routers on the segment. This may not represent an issue for the smaller multiaccess broadcast networks, but it may represent an issue for the nonbroadcast multiaccess (NBMA) networks, where in most cases you do not have full-mesh private virtual circuit (PVC) topology. This issue in NBMA networks manifests in an inability for neighbors to synchronize their OSPF databases directly among themselves. A logical solution in this case is to have a central point of OSPF adjacency responsible for the database synchronization and advertisement of the segment to the other routers, as shown in Figure 3-5.




Figure 3-5 *OSPF Adjacencies on Multiaccess Networks*

As the number of routers on the segment grows, the number of OSPF adjacencies increases exponentially. Every router must synchronize its OSPF database with every other router, and in the case of a large number of routers, this leads to inefficiency. Another issue arises when every router on the segment advertises all its adjacencies to other routers in the network. If you have full-mesh OSPF adjacencies, remaining OSPF routers will receive a large amount of redundant link-state information. Again, the solution for this problem is to establish a central point with which every other router forms adjacency and which advertises the segment as a whole to the rest of the network.

The routers on the multiaccess segment elect a designated router (DR) and backup designated router (BDR), which centralizes communications for all routers connected to the segment. The DR and BDR improve network functioning in the following ways:

- Reducing routing update traffic:** The DR and BDR act as a central point of contact for link-state information exchange on a multiaccess network; therefore, each router must establish a full adjacency with the DR and the BDR only. Each router, rather than exchanging link-state information with every other router on the segment, sends the link-state information to the DR and BDR only, by using a dedicated IPv4 multicast address 224.0.0.6 or FF00::6 for IPv6. The DR represents the multiaccess network in the sense that it sends link-state information from each router to all other routers in the network. This flooding process significantly reduces the router-related traffic on the segment.
- Managing link-state synchronization:** The DR and BDR ensure that the other routers on the network have the same link-state information about the common segment. In this way, the DR and BDR reduce the number of routing errors.

Only LSAs are sent to the DR/BDR. The normal routing of packets on the segment will go to the best next-hop router.

When the DR is operating, the BDR does not perform any DR functions. Instead, the BDR receives all the information, but the DR performs the LSA forwarding and LSDB synchronization tasks. The BDR performs the DR tasks only if the DR fails. When the DR fails, the BDR automatically becomes the new DR, and a new BDR election occurs.

In Example 3-10, the DR/BDR status on R1, R4, and R5 are observed using the **show ip ospf neighbor** command. Routers R1, R4, and R5 are all connected to the same shared network segment, where OSPF will automatically attempt to optimize adjacencies.

Example 3-10 Neighbor Status of R1, R4, and R5

R1# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	1	FULL/BDR	00:00:37	172.16.145.4	Ethernet0/1
5.5.5.5	1	FULL/DR	00:00:39	172.16.145.5	Ethernet0/1
2.2.2.2	1	FULL/DR	00:01:53	172.16.12.2	Serial2/0
3.3.3.3	1	FULL/DR	00:00:35	172.16.13.2	Ethernet0/0
R4# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/DROTHER	00:00:39	172.16.145.1	Ethernet0/0
5.5.5.5	1	FULL/DR	00:00:39	172.16.145.5	Ethernet0/0
R5# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/DROTHER	00:00:39	172.16.145.1	Ethernet0/0
4.4.4.4	1	FULL/BDR	00:00:35	172.16.145.4	Ethernet0/0

When R1, R4, and R5 start establishing OSPF neighbor adjacency, they first send OSPF Hello packets to discover which OSPF neighbors are active on the common Ethernet segment. After the bidirectional communication between routers is established and they are all in the OSPF neighbor 2-Way state, the DR/BDR election process begins. The OSPF Hello packet contains three specific fields used for the DR/BDR election: Designated Router, Backup Designated Router, and Router Priority.

The Designated Router and Backup Designate Router fields are populated with a list of routers claiming to be DR and BDR. From all routers listed, the router with the highest priority becomes the DR, and the one with the next highest priority becomes the BDR. If the priority values are equal, the router with the highest OSPF router ID becomes the DR, and the one with the next highest OSPF router ID becomes the BDR.

The DR/BDR election process takes place on broadcast and NBMA networks. The main difference between the two is the type of IP address used in the Hello packet. On the multiaccess broadcast networks, routers use multicast destination IPv4 address 224.0.0.6 to communicate with the DR (called AllDRRouters), and the DR uses multicast destination IPv4 address 224.0.0.5 to communicate with all other non-DR routers (called AllSPFRouters). On NBMA networks, the DR and adjacent routers communicate using unicast addresses.

The DR/BDR election process not only occurs when the network first becomes active but also when the DR becomes unavailable. In this case, the BDR will immediately become the DR, and the election of the new BDR starts.

In the topology, R5 has been elected as the DR and R4 as the BDR due to having the highest router ID values on the segment. R1 became a DROTHER. On the multiaccess segment, it is normal behavior that the router in DROTHER status is fully adjacent with DR/BDR and in 2-WAY state with all other DROTHER routers present on the segment.

In Example 3-11, the interface on R5 is shut down toward R1 and R4. Now, reexamine the DR/BDR status on R1 and R4. After the shutdown on the interface, wait until neighbor adjacencies expire before reexamining the DR/BDR state.

Example 3-11 R5's Ethernet 0/0 Interface Shutdown

```
R5(config)# interface ethernet 0/0
R5(config-if)# shutdown
*Dec  8 16:20:25.080: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Dec  8 16:20:25.080: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
```

```
R1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	1	FULL/DR	00:00:32	172.16.145.4	Ethernet0/1
2.2.2.2	1	FULL/DR	00:01:36	172.16.12.2	Serial2/0

3.3.3.3	1	FULL/DR	00:00:39	172.16.13.2	Ethernet0/0
---------	---	---------	----------	-------------	-------------

```
R4# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/BDR	00:00:33	172.16.145.1	Ethernet0/0

When R5's Ethernet 0/0 interface is shut down, the DR router on the segment becomes immediately unavailable. As a result, a new DR/BDR election takes place. The output of the `show ip ospf neighbor` command shows that R4 has become the DR and R1 the BDR.

Next, in Example 3-12, R5's interface toward R1 and R4 is enabled. Examine the DR/BDR status on R1, R4, and R5.

Example 3-12 R1's Ethernet 0/0 Interface Reenabled

```
R5(config)# interface ethernet 0/0
R5(config-if)# no shutdown
*Dec 10 08:49:26.491: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from
LOADING to FULL, Loading Done
*Dec 10 08:49:30.987: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Ethernet0/0 from
LOADING to FULL, Loading Done
```

```
R1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	1	FULL/DR	00:00:36	172.16.145.4	Ethernet0/1
5.5.5.5	1	FULL/DROTHER	00:00:38	172.16.145.5	Ethernet0/1
2.2.2.2	1	FULL/DR	00:01:52	172.16.12.2	Serial2/0
3.3.3.3	1	FULL/DR	00:00:33	172.16.13.2	Ethernet0/0

```
R4# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/BDR	00:00:30	172.16.145.1	Ethernet0/0
5.5.5.5	1	FULL/DROTHER	00:00:34	172.16.145.5	Ethernet0/0

```
R5# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/BDR	00:00:33	172.16.145.1	Ethernet0/0
4.4.4.4	1	FULL/DR	00:00:37	172.16.145.4	Ethernet0/0

When R5's Ethernet 0/0 interface is reenabled, a new DR/BDR election process will not take place even though R5 has the highest OSPF router ID on the segment. Once a DR

and BDR are elected, they are not preempted. This rule makes the multiaccess segment more stable by preventing the election process from occurring whenever a new router becomes active. It means that the first two DR-eligible routers on the link will be elected as DR and BDR. A new election will occur only when one of them fails.

Using OSPF Priority in the DR/BDR Election

One of the fields in the OSPF Hello packet used in the DR/BDR election process is the Router Priority field. Every broadcast and NBMA OSPF-enabled interface is assigned a priority value between 0 and 255. By default, in Cisco IOS, the OSPF interface priority value is 1 and can be manually changed by using the **ip ospf priority** interface command. When electing a DR and BDR, the routers view the OSPF priority value of other routers during the Hello packet exchange process, and then use the following conditions to determine which router to select:

- The router with the highest priority value is elected as the DR.
- The router with the second-highest priority value is the BDR.
- In case of a tie where two routers have the same priority value, router ID is used as the tiebreaker. The router with the highest router ID becomes the DR. The router with the second-highest router ID becomes the BDR.
- A router with a priority that is set to 0 cannot become the DR or BDR. A router that is not the DR or BDR is called a DROTHER.

The OSPF priority is configured on R1 using the **ip ospf priority** interface command, shown in Example 3-13. The OSPF process is cleared on R4 to reinitiate the DR/BDR election process. Setting the OSPF interface priority to a value higher than 1 will influence the DB/BDR election in favor of R1.

Example 3-13 *Configuring the OSPF Priority on an Interface*

```
R1(config)# interface ethernet 0/1
R1(config-if)# ip ospf priority 100

R4# clear ip ospf process
Reset ALL OSPF processes? [no]: yes
*Dec 10 13:08:48.610: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Dec 10 13:08:48.610: %OSPF-5-ADJCHG: Process 1, Nbr 5.5.5.5 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Dec 10 13:09:01.294: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from
LOADING to FULL, Loading Done
*Dec 10 13:09:04.159: %OSPF-5-ADJCHG: Process 1, Nbr 5.5.5.5 on Ethernet0/0 from
LOADING to FULL, Loading Done
```

In this example, the OSPF interface priority value is configured to 100. This influences the DR/BDR election, so that the R1 router will become DR after the OSPF process is cleared on the current DR, R4.

In Example 3-14, the `show ip ospf interface Ethernet 0/1` command on R1 verifies that it has been elected as a new DR.

Example 3-14 R1 Is the New DR

```
R1# show ip ospf interface ethernet 0/1
Ethernet0/1 is up, line protocol is up

  Internet Address 172.16.145.1/29, Area 0, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 10
  Topology-MTID      Cost      Disabled   Shutdown   Topology Name
        0           10         no         no         Base
  Transmit Delay is 1 sec, State DR, Priority 100
  Designated Router (ID) 1.1.1.1, Interface address 172.16.145.1
  Backup Designated router (ID) 5.5.5.5, Interface address 172.16.145.5
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 5
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 4.4.4.4
    Adjacent with neighbor 5.5.5.5 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

The Ethernet 0/1 interface on R1 has been assigned the OSPF priority value of 100, too, and when the new DR/BDR election process took place, the state of the R1 has become DR. The `show ip ospf interface` command on R1 shows that R1 is elected as the DR and that R5 is elected as the BDR. R1 is fully adjacent with two neighbors: R4 and R5.

OSPF Behavior in NBMA Hub-and-Spoke Topology

Special issues may arise when trying to interconnect multiple OSPF sites over an NBMA network. For example, if the NBMA topology is not fully meshed, a broadcast or multicast that is sent by one router will not reach all the other routers. Frame Relay and ATM are two examples of NBMA networks. OSPF treats NBMA environments like any other broadcast media environment, such as Ethernet; however, NBMA clouds are usually built as hub-and-spoke topologies using private virtual circuits (PVCs) or switched virtual

circuits (SVCs). The hub-and-spoke topology shown in Figure 3-6 means that the NBMA network is only a partial mesh. In these cases, the physical topology does not provide multiaccess capability, on which OSPF relies. In a hub-and-spoke NBMA environment, you will need to have the hub router acting as the DR and spoke routers acting as the DROTHER routers. On the spoke router interfaces, you want to configure an OSPF priority value of 0 so that the spoke routers never participate in the DR election.




Figure 3-6 Hub-and-Spoke Topology

In addition, OSPF is not able to automatically discover OSPF neighbors over an NBMA network like Frame Relay. Neighbors must be statically configured on at least one router by using the `neighbor ip_address` configuration command in the router configuration mode.

Example 3-15 shows setting the OSPF priority on R4's and R5's Ethernet 0/0 interfaces to 0 using the `ip ospf priority` interface command. Setting the OSPF interface priority to 0 prevents the router from being a candidate for the DR/BDR role.

Example 3-15 Setting the OSPF Priority to 0 on R4 and R5

```
R4(config)# interface ethernet 0/0
R4(config-if)# ip ospf priority 0

R5(config)# interface ethernet 0/0
R5(config-if)# ip ospf priority 0
```

Setting the OSPF priority value to 0 on the Ethernet 0/0 interfaces for R4 and R5 means that these two routers will not participate in the DR/BDR election and will not be eligible to become the DR/BDR. These routers will be DROTHER routers.

The state of the DR/BDR status on R1, R4, and R5 is shown in Example 3-16.

Example 3-16 DR/BDR States on R1, R4, and R5

```
R1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/DROTHER	00:00:36	172.16.145.4	Ethernet0/1
5.5.5.5	0	FULL/DROTHER	00:00:34	172.16.145.5	Ethernet0/1

2.2.2.2	1	FULL/DR	00:01:33	172.16.12.2	Serial2/0
3.3.3.3	1	FULL/DR	00:00:30	172.16.13.2	Ethernet0/0


```
R4# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	100	FULL/DR	00:00:37	172.16.145.1	Ethernet0/0
5.5.5.5	0	2WAY/DROTHER	00:00:37	172.16.145.5	Ethernet0/0


```
R5# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	100	FULL/DR	00:00:32	172.16.145.1	Ethernet0/0
4.4.4.4	0	2WAY/DROTHER	00:00:37	172.16.145.4	Ethernet0/0

The output of the `show ip ospf neighbor` commands on R1 shows that R1 is fully adjacent with R4 and R5 and that R4 and R5 have DROTHER functions. R4 is fully adjacent with the DR router R1, but it maintains a 2-Way state with its peer DROTHER router R5. Similarly, R5 is fully adjacent with DR R1 and maintains a 2-Way state with the DROTHER router R4. A 2-Way state between non-DR/BDR routers on the segment is normal behavior; they do not synchronize LSDBs directly, but over DR/BDR. By maintaining 2-Way state, DROTHER routers keep other DROTHER peers informed about their presence on the network.

The Importance of MTU

The IP MTU parameter determines the maximum size of an IPv4 packet that can be forwarded out the interface without fragmentation. If a packet with an IPv4 MTU larger than the maximum arrives at the router interface, it will be either discarded, if the DF bit in the packet header is set, or it will be fragmented. OSPF for IPv4 packets completely relies on IPv4 for the possible fragmentation. Although RFC 2328 does not recommend OSPF packet fragmentation, in some situations the size of the OSPF packet has greater value than the interface IPv4 MTU. If MTUs are mismatched between two neighbors, this could introduce issues with exchange of link-state packets, resulting in continuous retransmissions.

Note The interface command for setting the IPv6 MTU parameter is `ipv6 mtu`. An IPv6 router does not fragment an IPv6 packet unless it is the source of the packet.

To prevent such issues, OSPF requires that the same IPv4 MTU be configured on both sides of the link. If neighbors have a mismatched IPv4 MTU configured, they will not be able to form full OSPF adjacency. They will be stuck in the ExStart adjacency state.

In Example 3-17, the IPv4 MTU size on the R3 Ethernet 0/0 interface is changed to 1400.

Example 3-17 *Configuration of the IPv4 MTU on R3's Ethernet 0/0 Interface*

```
R3(config)# interface ethernet 0/0
R3(config-if)# ip mtu 1400
```

After the IPv4 MTU size is changed on R3's Ethernet 0/0 interface, this creates a mismatch between IPv4 MTU sizes on the link between R3 and R1. This mismatch will result in R3 and R1 not being able to synchronize their OSPF databases, and a new full adjacency between them will not be established. This is observed in Example 3-18 using the `debug ip ospf adj` command on R3. The OSPF process is cleared to reset adjacency, and debug is disabled when the OSPF session is reestablished.

Example 3-18 *Observing a Mismatched MTU*

```
R3# debug ip ospf adj
R3# clear ip ospf process
Reset ALL OSPF processes? [no]: yes
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Interface going Up
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: 2 Way Communication to 1.1.1.1, state 2WAY
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Backup seen event before WAIT timer
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: DR/BDR election
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Elect BDR 3.3.3.3
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Elect DR 1.1.1.1
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Elect BDR 3.3.3.3
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: Elect DR 1.1.1.1
*Jan 19 17:37:05.969: OSPF-3 ADJ   Et0/0: DR: 1.1.1.1 (Id)   BDR: 3.3.3.3 (Id)
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Nbr 1.1.1.1: Prepare dbase exchange
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Send DBD to 1.1.1.1 seq 0x21D6 opt 0x52
flag 0x7 len 32
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Rcv DBD from 1.1.1.1 seq 0x968 opt 0x52
flag 0x7 len 32 mtu 1500 state EXSTART
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Nbr 1.1.1.1 has larger interface MTU
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Rcv DBD from 1.1.1.1 seq 0x21D6 opt 0x52
flag 0x2 len 112 mtu 1500 state EXSTART
*Jan 19 17:37:05.970: OSPF-3 ADJ   Et0/0: Nbr 1.1.1.1 has larger interface MTU
R1# no debug ip ospf adj
```

The DBD packet carries information about largest nonfragmented packet that can be sent from the neighbor. In this situation, the IPv4 MTU values on different sides of the link are not equal. R3 will receive the DBD packet with an IPv4 MTU size of 1500, which is greater than its own MTU size of 1400. This will result in the inability of both R3 and R1 to establish full neighbor adjacency, and the output of the `debug` command will display that Nbr has a larger interface MTU message. Mismatched neighbors will stay in ExStart state. To form full OSPF adjacency, the IPv4 MTU needs to match on both sides of the link.

Note By default, the IPv6 MTU must also match between OSPFv3 neighbors. However, you can override this by using the `ospfv3 mtu-ignore` interface command.

In Example 3-19, the OSPF neighbor state is verified on R3 and R1.

Example 3-19 Verifying the OSPF Neighbor States

R3# show ip ospf neighbor						
Neighbor ID	Pri	State	Dead Time	Address	Interface	
1.1.1.1	1	EXSTART/BDR	00:00:38	172.16.13.1	Ethernet0/0	
R1# show ip ospf neighbor						
Neighbor ID	Pri	State	Dead Time	Address	Interface	
4.4.4.4	0	FULL/DROTHER	00:00:39	172.16.145.4	Ethernet0/1	
5.5.5.5	0	FULL/DROTHER	00:00:38	172.16.145.5	Ethernet0/1	
2.2.2.2	1	FULL/DR	00:01:55	172.16.12.2	Serial2/0	
3.3.3.3	1	EXCHANGE/DR	00:00:36	172.16.13.2	Ethernet0/0	
R1# show ip ospf neighbor						
Neighbor ID	Pri	State	Dead Time	Address	Interface	
4.4.4.4	0	FULL/DROTHER	00:00:38	172.16.145.4	Ethernet0/1	
5.5.5.5	0	FULL/DROTHER	00:00:31	172.16.145.5	Ethernet0/1	
2.2.2.2	1	FULL/DR	00:01:31	172.16.12.2	Serial2/0	
3.3.3.3	1	INIT/DROTHER	00:00:33	172.16.13.2	Ethernet0/0	

Mismatching interface IPv4 MTU sizes on opposite sides of the OSPF link results in the inability to form full adjacency. R3, which detected that R1 has higher MTU, keeps the neighbor adjacency in ExStart state. R1 continues to retransmit initial BDB packet to R3, but R3 cannot acknowledge them because of the unequal IPv4 MTU. On R1, you can observe how the OSPF neighbor relationship state with R3 is unstable. Adjacency gets to the Exchange state, but is then terminated, starting again from the Init state up to the Exchange state.

The recommended way to solve such issues is to make sure that the IPv4 MTU matches between OSPF neighbors.

Manipulating OSPF Timers

Similar to EIGRP, OSPF uses two timers to check neighbor reachability: the hello and dead intervals. The values of hello and dead intervals are carried in OSPF Hello packets and serve as a keepalive message, with the purpose of acknowledging the presence of the router on the segment. The hello interval specifies the frequency of sending OSPF Hello packets in seconds. The OSPF dead timer specifies how long a router waits to receive a Hello packet before it declares a neighbor router as down.

OSPF requires that both hello and dead timers be identical for all routers on the segment to become OSPF neighbors. The default value of the OSPF hello timer on multiaccess broadcast and point-to-point links is 10 seconds, and is 30 seconds on all other network types, including NBMA. When you configure the hello interval, the default value of the dead interval is automatically adjusted to four times the hello interval. For broadcast and point-to-point links, it is 40 seconds, and for all other OSPF network types, it is 120 seconds.

To detect faster topological changes, you can lower the value of OSPF hello interval, with the downside of having more routing traffic on the link. The **debug ip ospf hello** command enables you to investigate hello timer mismatch.

In Example 3-20, R1, the different hello/dead timer values on Ethernet 0/1, and Frame Relay Serial 2/0 interfaces are observed using the **show ip ospf interface** command.

Example 3-20 *Examining the Hello/Dead Timers on R1 Interfaces*

```
R1# show ip ospf interface ethernet 0/1
Ethernet0/1 is up, line protocol is up
  Internet Address 172.16.145.1/29, Area 0, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 10
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                  10        no            no            Base
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 5.5.5.5, Interface address 172.16.145.5
  Backup Designated router (ID) 4.4.4.4, Interface address 172.16.145.4
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
<Output omitted>

R1# show ip ospf interface serial 2/0
Serial2/0 is up, line protocol is up
  Internet Address 172.16.12.1/30, Area 1, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                  64        no            no            Base
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 2.2.2.2, Interface address 172.16.12.2
  Backup Designated router (ID) 1.1.1.1, Interface address 172.16.12.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
<Output omitted>
```

The default value of the OSPF hello interval on broadcast multiaccess (Ethernet) and point-to-point links is 10 seconds, and the default value of the dead interval is four times hello (40 seconds). Default values of the OSPF hello and dead timers on all other OSPF network types, including nonbroadcast (NBMA) like Frame Relay on the Serial 2/0 interface, are 30 seconds and 120 seconds, respectively.

On low-speed links, you might want to alter default OSPF timer values to achieve faster convergence. The negative aspect of lowering the OSPF hello interval is the overhead of more frequent routing updates causing higher router utilization and more traffic on the link.

In Example 3-21, the default OSPF hello and dead intervals on R1's Frame Relay Serial 2/0 interface are modified. You can change the OSPF by using the `ip ospf hello-interval` and `ip ospf dead-interval` interface commands.

Example 3-21 *Modifying the Hello and Dead Intervals on R1's Serial Interface*

```
R1(config)# interface serial 2/0
R1(config-if)# ip ospf hello-interval 8
R1(config-if)# ip ospf dead-interval 30
*Jan 20 13:17:34.441: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial2/0 from
FULL to DOWN, Neighbor Down: Dead timer expired
```

Once the default OSPF hello and dead interval values on the Frame Relay link are changed, both routers will detect hello timer mismatch. As a result, the dead timer will not be refreshed, so it will expire, declaring the OSPF neighbor relationship as down.

Note When you are changing only the OSPF hello interval, OSPF automatically changes the dead interval to four times the hello interval.

In Example 3-22, R2's default OSPF hello and dead timers on the Frame Relay Serial 0/0 interface are changed so that they match respective values configured on R1.

Example 3-22 *Modifying the Hello and Dead Intervals on R2's Serial Interface*

```
R2(config)# interface serial 0/0
R2(config-if)# ip ospf hello-interval 8
R2(config-if)# ip ospf dead-interval 30
*Jan 20 13:38:58.976: %OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/0 from
LOADING to FULL, Loading Done
```

When you are changing OSPF hello and dead timers on R2 so that they match the timers on R1, both routers on the link will be able to establish adjacency and elect the DR/BDR on the NBMA segment. Routers will then exchange and synchronize LSDBs and form full neighbor adjacency.

On R2, the OSPF neighbor state is verified by using the `show ip ospf neighbor detail` command, as demonstrated in Example 3-23.

Example 3-23 *Verifying the OSPF Neighbor States on R2*

```

R2# show ip ospf neighbor detail
Neighbor 1.1.1.1, interface address 172.16.12.1
  In the area 1 via interface Serial0/0
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 172.16.12.2 BDR is 172.16.12.1
  Poll interval 120
  Options is 0x12 in Hello (E-bit, L-bit)
  Options is 0x52 in DBD (E-bit, L-bit, O-bit)
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:26
  Neighbor is up for 00:14:57
  Index 1/1, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec

```

The output of the `show ip ospf neighbor detail` command confirms that full OSPF adjacency with R1 is established. The output also shows additional information about neighbor router ID, DR/BDR roles, and how long the neighbor session has been established.

OSPF Neighbor Relationship over Point-to-Point Links

Figure 3-7 shows a point-to-point network joining a single pair of routers. A T1 serial line that is configured with a data link layer protocol such as PPP or High-Level Data Link Control (HDLC) is an example of a point-to-point network.



Figure 3-7 *Point-to-Point link*

On these types of networks, the router dynamically detects its neighboring routers by multicasting its Hello packets to all OSPF routers, using the 224.0.0.5 address. On point-to-point networks, neighboring routers become adjacent whenever they can communicate directly. No DR or BDR election is performed; there can be only two routers on a point-to-point link, so there is no need for a DR or BDR.

The default OSPF hello and dead timers on point-to-point links are 10 seconds and 40 seconds, respectively.

OSPF Neighbor Relationship over Layer 3 MPLS VPN

Figure 3-8 shows a Layer 3 MPLS VPN architecture, where the ISP provides a peer-to-peer VPN architecture. In this architecture, provider edge (PE) routers participate in

customer routing, guaranteeing optimum routing between customer sites. Therefore, the PE routers carry a separate set of routes for each customer, resulting in perfect isolation between customers.




Figure 3-8 *Layer 3 MPLS VPN*

The following applies to Layer 3 MPLS VPN technology, even when running OSPF as a provider edge - customer edge (PE-CE) routing protocol:

- The customer routers should not be aware of MPLS VPN; they should run standard IP routing software.
- The core routers in the provider network between the two PE routers are known as the P routers (not shown in the diagram). The P routers do not carry customer VPN routes for the MPLS VPN solution to be scalable.
- The PE routers must support MPLS VPN services and traditional Internet services.

To OSPF, the Layer 3 MPLS VPN backbone looks like a standard corporate backbone that runs standard IP routing software. Routing updates are exchanged between the customer routers and the PE routers that appear as normal routers in the customer network. OSPF is enabled on proper interfaces by using the **network** command. The standard design rules that are used for enterprise Layer 3 MPLS VPN backbones can be applied to the design of the customer network. The service provider routers are hidden from the customer view, and CE routers are unaware of MPLS VPN. Therefore, the internal topology of the Layer 3 MPLS backbone is totally transparent to the customer. The PE routers receive IPv4 routing updates from the CE routers and install them in the appropriate virtual routing and forwarding (VRF) table. This part of the configuration, and operation, is the responsibility of a service provider.

The PE-CE can have any OSPF network type: point-to-point, broadcast, or even non-broadcast multiaccess.

The only difference between a PE-CE design and a regular OSPF design is that the customer has to agree with the service provider about the OSPF parameters (area ID, authentication password, and so on); usually, these parameters are governed by the service provider.

OSPF Neighbor Relationship over Layer 2 MPLS VPN

Figure 3-9 shows a Layer 2 MPLS VPN. The MPLS backbone of the service provider is used to enable Layer 2 Ethernet connectivity between the customer routers R1 and R2, whether an Ethernet over MPLS (EoMPLS) or Layer 2 MPLS VPN Ethernet service is used.




Figure 3-9 *Layer 2 MPLS VPN*

R1 and R2 thus exchange Ethernet frames. PE router PE1 takes the Ethernet frames that are received from R1 on the link to PE1, encapsulates them into MPLS packets, and forwards them across the backbone to router PE2. PE2 decapsulates the MPLS packets and reproduces the Ethernet frames on the link toward R2. EoMPLS and Layer 2 MPLS VPN typically do not participate in Shortest Tree Protocol (STP) and bridge protocol data unit (BPDU) exchanges, so EoMPLS and Layer 2 MPLS VPNs are transparent to the customer routers.

The Ethernet frames are transparently exchanged across the MPLS backbone. Keep in mind that customer routers can be connected either in a port-to-port fashion, in which PE routers take whatever Ethernet frame is received and forward the frames across the Layer 2 MPLS VPN backbone, or in a VLAN subinterface fashion in which frames for a particular VLAN—identified with subinterface in configuration—are encapsulated and sent across the Layer 2 MPLS VPN backbone.

When deploying OSPF over EoMPLS, there are no changes to the existing OSPF configuration from the customer perspective.

OSPF needs to be enabled, and network commands must include the interfaces that are required by the relevant OSPF area to start the OSPF properly.

R1 and R2 form a neighbor relationship with each other over the Layer 2 MPLS VPN backbone. From an OSPF perspective, the Layer 2 MPLS VPN backbone, PE1, and PE2 are all invisible.

A neighbor relationship is established between R1 and R2 directly, and it behaves in the same way as on a regular Ethernet broadcast network.

OSPF Neighbor States

OSPF neighbors go through multiple neighbor states before forming full OSPF adjacency, as illustrated in Figure 3-10.

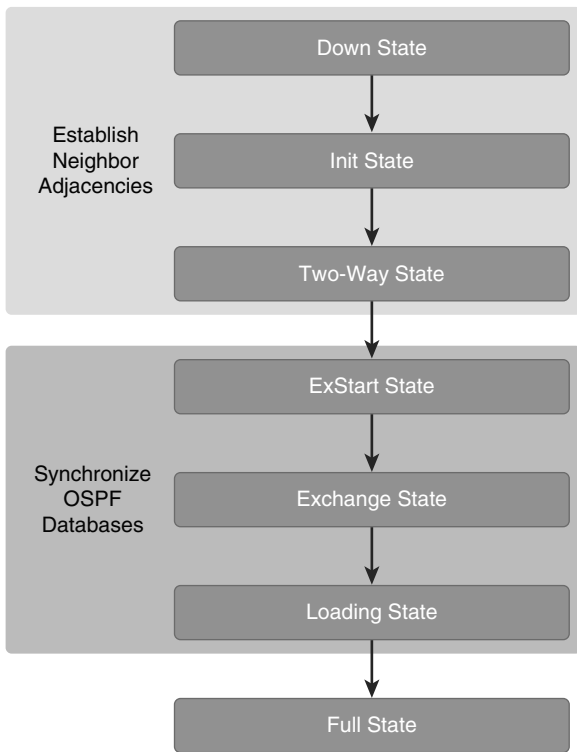


Figure 3-10 OSPF States

The following is a brief summary of the states that an interface passes through before becoming adjacent to another router:

- **Down:** No information has been received on the segment.
- **Init:** The interface has detected a Hello packet coming from a neighbor, but bidirectional communication has not yet been established.
- **2-Way:** There is bidirectional communication with a neighbor. The router has seen itself in the Hello packets coming from a neighbor. At the end of this stage, the DR and BDR election would have been done if necessary. When routers are in the 2-Way state, they must decide whether to proceed in building an adjacency. The decision is based on whether one of the routers is a DR or BDR or the link is a point-to-point or a virtual link.
- **ExStart:** Routers are trying to establish the initial sequence number that is going to be used in the information exchange packets. The sequence number ensures that routers always get the most recent information. One router will become the master and the other will become the slave. The primary router will poll the secondary for information.

- **Exchange:** Routers will describe their entire LSDB by sending database description (DBD) packets. A DBD includes information about the LSA entry header that appears in the router's LSDB. The entries can be about a link or about a network. Each LSA entry header includes information about the link-state type, the address of the advertising router, the link's cost, and the sequence number. The router uses the sequence number to determine the "newness" of the received link-state information.
- **Loading:** In this state, routers are finalizing the information exchange. Routers have built a link-state request list and a link-state retransmission list. Any information that looks incomplete or outdated will be put on the request list. Any update that is sent will be put on the retransmission list until it gets acknowledged.
- **Full:** In this state, adjacency is complete. The neighboring routers are fully adjacent. Adjacent routers will have similar LSDBs.

OSPF Network Types

OSPF defines distinct types of networks based on their physical link types, as shown in Table 3-1. OSPF operation on each type is different, including how adjacencies are established and which configuration is required.

Table 3-1 *OSPF Network Types*

OSPF Network Type	Uses DR/BDR	Default Hello Interval (sec)	Dynamic Neighbor Discovery	More than Two Routers Allowed in Subnet
Point-to-point	No	10	Yes	No
Broadcast	Yes	10	Yes	Yes
Nonbroadcast	Yes	30	No	Yes
Point-to-multipoint	No	30	Yes	Yes
Point-to-multipoint nonbroadcast	No	30	No	Yes
Loopback	No	—	—	No

These are the most common network types that are defined by OSPF:

- **Point-to-point:** Routers use multicast to dynamically discover neighbors. There is no DR/BDR election because only two routers can be connected on a single point-to-point segment. It is a default OSPF network type for serial links and point-to-point Frame Relay subinterfaces.
- **Broadcast:** Multicast is used to dynamically discover neighbors. The DR and BDR are elected to optimize the exchange of information. It is a default OSPF network type for Ethernet links.

- **Nonbroadcast:** Used on networks that interconnect more than two routers but without broadcast capability. Frame Relay and ATM are examples of NBMA networks. Neighbors must be statically configured, followed by DR/BDR election. This network type is the default for all physical interfaces and multipoint subinterfaces using Frame Relay encapsulation.
- **Point-to-multipoint:** OSPF treats this network type as a logical collection of point-to-point links even though all interfaces belong to the common IP subnet. Every interface IP address will appear in the routing table of the neighbors as a host /32 route. Neighbors are discovered dynamically using multicast. No DR/BDR election occurs.
- **Point-to-multipoint nonbroadcast:** Cisco extension that has the same characteristics as point-to-multipoint type except for the fact that neighbors are not discovered dynamically. Neighbors must be statically defined, and unicast is used for communication. Can be useful in point-to-multipoint scenarios where multicast and broadcast are not supported.
- **Loopback:** Default network type on loopback interfaces.

You can change OSPF network type by using the interface configuration mode command `ip ospf network network_type`.

Configuring Passive Interfaces

Passive interface configuration is a common method for hardening routing protocols and reducing the use of resources. The passive interface is supported by OSPF, and a sample configuration is shown in Example 3-24.

Example 3-24 *Passive Interface Configuration for OSPF*

```
Router(config)# router ospf 1
Router(config-if)# passive-interface default
Router(config-if)# no passive-interface serial 1/0
```

When you configure a passive interface under the OSPF process, the router stops sending and receiving OSPF Hello packets on the selected interface. The passive interface should be used only on interfaces where the router is not expected to form any OSPF neighbor adjacency. A specific interface can be configured as passive, or passive interface can be configured as the default. If the default option is used, any interfaces that need to form a neighbor adjacency must be exempted with the `no passive-interface` configuration command.

Building the Link-State Database

OSPF, as a link-state protocol, uses several different packets to exchange the information about network topology between routers. These packets are called *link-state*

advertisements (LSAs), and they describe the network topology in great detail. Each router stores the received LSA packets in the link-state database (LSDB). After LSDBs are synced between the routers, OSPF uses the shortest path first (SPF) algorithm to calculate the best routes. The best intra-area routes are calculated individually by each OSPF router. For the best interarea route calculation, the internal router must rely also on the best path information received from the ABRs.

Upon completing this section, you will be able to do the following:

- List and describe different LSA types
- Describe how OSPF LSAs are also reflooded at periodic intervals
- Describe the exchange of information in a network without a designated router
- Describe the exchange of information in a network with a designated router
- Explain when SPF algorithms occur
- Describe how the cost of intra-area routes is calculated
- Describe how the cost of interarea routes is calculated
- Describe rules selecting between intra-area and interarea routes

OSPF LSA Types

Knowing the detailed topology of the OSPF area is required for a router to calculate the best paths. Topology details are described by LSAs, which are the building blocks of the OSPF LSDB. Individually, LSAs act as database records. In combination, they describe the entire topology of an OSPF network area. Figure 3-11 shows a sample topology, highlighting the most common types of OSPF LSAs, which are described in further detail in the list that follows.

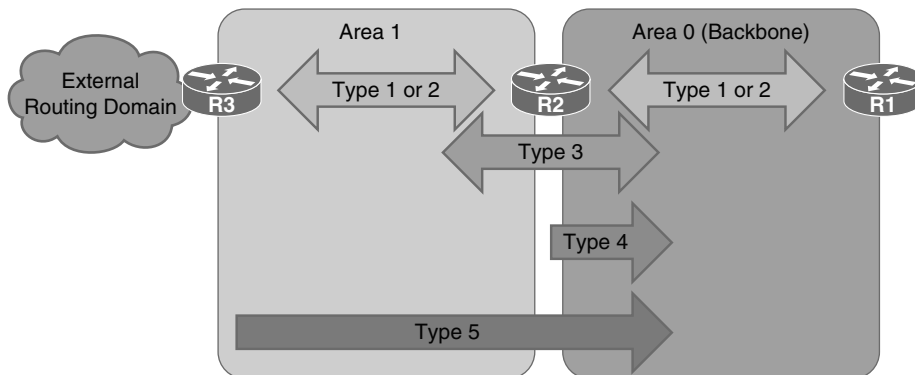


Figure 3-11 OSPF LSA Types

- **Type 1, Router LSA:** Every router generates router link advertisements for each area to which it belongs. Router link advertisements describe the state of the router links to the area and are flooded only within that particular area. For all types of LSAs, there are 20-byte LSA headers. One of the fields of the LSA header is the link-state ID. The link-state ID of the type 1 LSA is the originating router ID.
- **Type 2, Network LSA:** DRs generate network link advertisements for multiaccess networks. Network link advertisements describe the set of routers that are attached to a particular multiaccess network. Network link advertisements are flooded in the area that contains the network. The link-state ID of the type 2 LSA is the IP interface address of the DR.
- **Type 3, Summary LSA:** An ABR takes the information that it learned in one area and describes and summarizes it for another area in the summary link advertisement. This summarization is not on by default. The link-state ID of the type 3 LSA is the destination network number.
- **Type 4, ASBR Summary LSA:** The ASBR summary link advertisement informs the rest of the OSPF domain how to get to the ASBR. The link-state ID includes the router ID of the described ASBR.
- **Type 5, Autonomous System LSA:** Autonomous system external link advertisements, which are generated by ASBRs, describe routes to destinations that are external to the autonomous system. They get flooded everywhere, except into special areas. The link-state ID of the type 5 LSA is the external network number.

Other LSA types include the following:

- **Type 6:** Specialized LSAs that are used in multicast OSPF applications
- **Type 7:** Used in special area type NSSA for external routes
- **Type 8, 9:** Used in OSPFv3 for link-local addresses and intra-area prefix
- **Type 10, 11:** Generic LSAs, also called *opaque*, which allow future extensions of OSPF

Examining the OSPF Link-State Database

This section analyzes the OSPF LSDB and the different types of LSAs using the topology in Figure 3-12. All routers have already been preconfigured with OSPF. In the figure, R1 is an ABR between areas 0, 1, and 2. R3 acts as the ASBR between the OSPF routing domain and an external domain. LSA types 1 and 2 are flooded between routers within an area. Type 3 and type 5 LSAs are flooded when exchanging information about backbone and standard areas. Type 4 LSAs are injected into the backbone by the ABR because all routers in the OSPF domain need to reach the ASBR (R3).

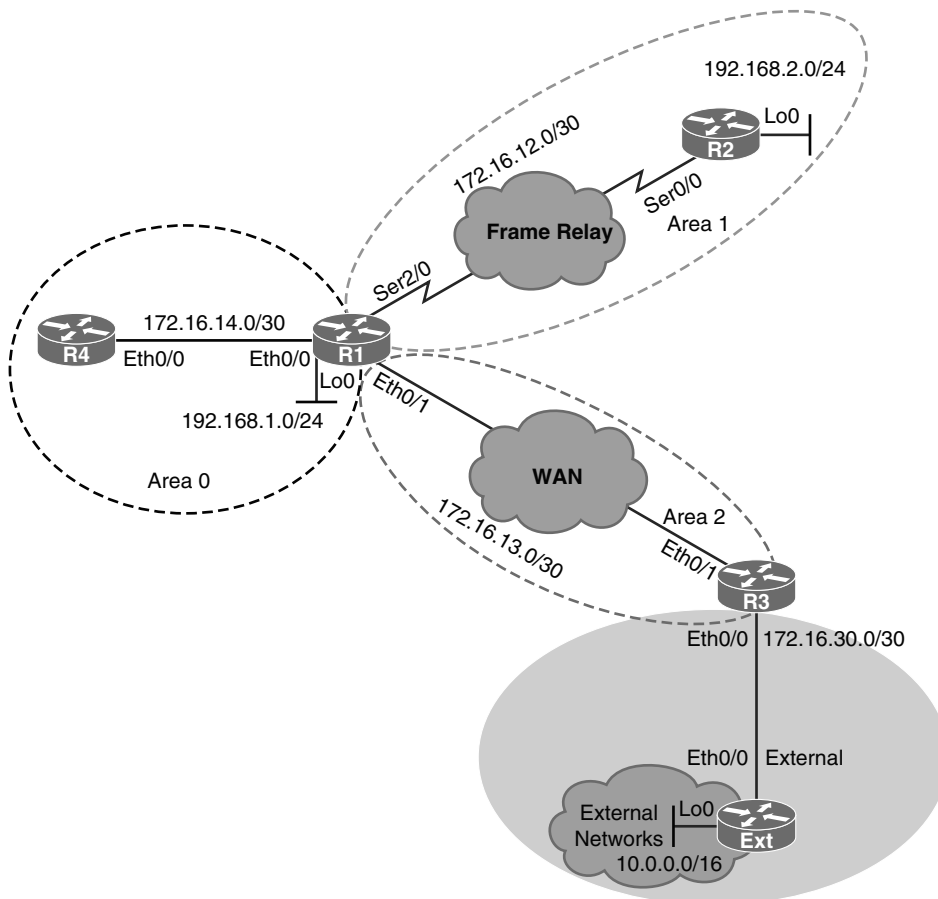


Figure 3-12 OSPF Topology

OSPF Link-State Database

Example 3-25 shows R4's routing table, which includes several OSPF routes because all the routers have already been configured.

Example 3-25 R4's Routing Table

```
R4# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```

+ - replicated route, % - next hop override

Gateway of last resort is not set

10.0.0.0/16 is subnetted, 1 subnets
O E2   10.0.0.0 [110/20] via 172.16.14.1, 00:46:48, Ethernet0/0
172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA   172.16.12.0/30 [110/74] via 172.16.14.1, 03:19:12, Ethernet0/0
O IA   172.16.13.0/30 [110/20] via 172.16.14.1, 03:19:12, Ethernet0/0
C      172.16.14.0/30 is directly connected, Ethernet0/0
L      172.16.14.2/32 is directly connected, Ethernet0/0
O      192.168.1.0/24 [110/11] via 172.16.14.1, 00:36:19, Ethernet0/0
O IA   192.168.2.0/24 [110/75] via 172.16.14.1, 00:47:59, Ethernet0/0

```

Notice the intra-area route 192.168.1.0/24 and interarea routes describing WAN links 172.16.12.0/30, 172.16.13.0/30, and the remote subnet 192.168.2.0/24 on R2. There is also routing information about an OSPF external route that is describing network 10.0.0.0/16. This route is injected into OSPF on R3, which has connectivity to external networks.

Example 3-26 displays the OSPF database on R4.

Example 3-26 R4's OSPF LSDB

```

R4# show ip ospf database

OSPF Router with ID (4.4.4.4) (Process ID 1)

Router Link States (Area 0)

Link ID          ADV Router      Age             Seq#            Checksum Link count
1.1.1.1          1.1.1.1        291            0x8000000B     0x00966C   2
4.4.4.4          4.4.4.4        1993           0x80000007     0x001C4E   1

Net Link States (Area 0)

Link ID          ADV Router      Age             Seq#            Checksum
172.16.14.2     4.4.4.4        1993           0x80000006     0x0091B5

Summary Net Link States (Area 0)

Link ID          ADV Router      Age             Seq#            Checksum
172.16.12.0     1.1.1.1        291            0x80000007     0x00C567
172.16.13.0     1.1.1.1        291            0x80000007     0x009CC5
192.168.2.0     1.1.1.1        1031           0x80000002     0x002E5D

```


Summary ASB Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
3.3.3.3	1.1.1.1	1031	0x80000002	0x0035EB	
Type-5 AS External Link States					
Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.0.0	3.3.3.3	977	0x80000002	0x000980	0

The OSPF database contains all LSAs that describe the network topology. The **show ip ospf database** command displays the content of the LSDB and verifies information about specific LSAs.

The output reveals the presence of different LSA types. For each LSA type, you can see which router advertised it, the age of the LSA, and the value of the link ID.

In Example 3-26, notice two different type 1 LSAs, or router link advertisements, generated by routers with router ID 1.1.1.1 and 4.4.4.4.

Example 3-27 displays the details of R4's type 1 LSAs

Example 3-27 R4 Type 1 LSA Details

```
R4# show ip ospf database router

          OSPF Router with ID (4.4.4.4) (Process ID 1)

          Router Link States (Area 0)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 321
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: 8000000B
Checksum: 0x966C
Length: 48
Area Border Router
Number of Links: 2

Link connected to: a Stub Network
(Link ID) Network/subnet number: 192.168.1.0
(Link Data) Network Mask: 255.255.255.0
Number of MTID metrics: 0
TOS 0 Metrics: 1
```

```

Link connected to: a Transit Network
(Link ID) Designated Router address: 172.16.14.2
(Link Data) Router Interface address: 172.16.14.1
Number of MTID metrics: 0
TOS 0 Metrics: 10

LS age: 2023
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: 80000007
Checksum: 0x1C4E
Length: 36
Number of Links: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 172.16.14.2
(Link Data) Router Interface address: 172.16.14.2
Number of MTID metrics: 0
TOS 0 Metrics: 10

```

Type 1 LSAs are generated by every router and flooded within the area. They describe the state of the router links in that area. R4 has two type 1 LSAs in the database: one received from R1 with router ID 1.1.1.1, and one that was generated by R4.

The content of the displayed LSA reveals that R1 is an ABR with two links. The output shows details for both links, to what kind of network the links are connected, and their settings, such as the IP configuration. Link can be connected to a stub, to another router (point-to-point), or to a transit network. The transit network describes Ethernet or NMBA segment, which can include two or more routers. If the link is connected to a transit network, the LSA also includes the info about the DR address.

The LSDB keeps copies of all LSAs, including those that were generated locally on the router. An example of a local LSA is the second advertisement that is displayed in the output. It includes the same topology parameters as the first LSA, but this time from a perspective of router R4.

OSPF identifies all LSAs using a 32-bit LSID. When generating a type 1 LSA, the router uses its own router ID as the value of LSID.

Using the **self-originate** command argument, Example 3-28 displays locally generated type 1 LSAs on R4.

Example 3-28 *Locally Generated Type 1 LSAs on R4*

```

R4# show ip ospf database router self-originate

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Router Link States (Area 0)

LS age: 23
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: 80000008
Checksum: 0x1A4F
Length: 36
Number of Links: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 172.16.14.2
(Link Data) Router Interface address: 172.16.14.2
Number of MTID metrics: 0
TOS 0 Metrics: 10

```

The output shows the type 1 LSA, which describes the interface that is enabled in OSPF area 0 on router R4.

R4 has an interface that is connected to a transit network; therefore, the DR information is also included. You can see that R4 is the DR on the segment.

Example 3-29 shows the OSPF database on router R2.

Example 3-29 *R2's OSPF LSDB*

```

R2# show ip ospf database

      OSPF Router with ID (2.2.2.2) (Process ID 1)

      Router Link States (Area 1)

Link ID          ADV Router      Age             Seq#            Checksum Link count
1.1.1.1          1.1.1.1        403            0x80000008     0x0097B7  1
2.2.2.2          2.2.2.2        1088           0x80000008     0x006E5C  2

      Net Link States (Area 1)

Link ID          ADV Router      Age             Seq#            Checksum

```

```

172.16.12.2    2.2.2.2      587          0x80000003 0x00A5B6

Summary Net Link States (Area 1)

Link ID      ADV Router   Age          Seq#         Checksum
172.16.13.0  1.1.1.1     403         0x80000007 0x009CC5
172.16.14.0  1.1.1.1     403         0x80000007 0x0091CF
192.168.1.0  1.1.1.1     403         0x80000002 0x00B616

Summary ASB Link States (Area 1)

Link ID      ADV Router   Age          Seq#         Checksum
3.3.3.3     1.1.1.1     1143        0x80000002 0x0035EB

Type-5 AS External Link States

Link ID      ADV Router   Age          Seq#         Checksum Tag
10.0.0.0    3.3.3.3     1089        0x80000002 0x000980 0

```

OSPF type 1 LSAs are exchanged only within OSPF areas. Router R2, which has interfaces that are configured in OSPF area 1, should not see any type 1 LSAs that were originated on R4. The output of the OSPF database from R2 confirms this. No type 1 LSA with the advertising router parameter set to 4.4.4.4 can be found in the LSDB.

Example 3-30 displays the LSAs on R1.

Example 3-30 R1's OSPF LSDB

```

R1# show ip ospf database

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

Router Link States (Area 0)
Link ID      ADV Router   Age          Seq#         Checksum Link count
1.1.1.1     1.1.1.1     445         0x8000000B 0x00966C 2
4.4.4.4     4.4.4.4     103         0x80000008 0x001A4F 1
<Output omitted>

Router Link States (Area 1)
Link ID      ADV Router   Age          Seq#         Checksum Link count
1.1.1.1     1.1.1.1     445         0x80000008 0x0097B7 1
2.2.2.2     2.2.2.2     1133        0x80000008 0x006E5C 2
<Output omitted>

Router Link States (Area 2)
Link ID      ADV Router   Age          Seq#         Checksum Link count
1.1.1.1     1.1.1.1     445         0x80000008 0x00DDA5 1
3.3.3.3     3.3.3.3     1131        0x8000000A 0x00521D 1
<Output omitted>

```

Notice that router R1 is the only router that is in multiple areas. As an ABR, its OSPF database includes type 1 LSAs from all three areas.

OSPF Type 2 Network LSA

Figure 3-13 shows a type 2 LSA, which is generated for every transit broadcast or NBMA network within an area.




Figure 3-13 OSPF Type 2 LSA

The DR of the network is responsible for advertising the network LSA. A type 2 network LSA lists each of the attached routers that make up the transit network, including the DR itself, and the subnet mask that is used on the link. The type 2 LSA then floods to all routers within the transit network area. Type 2 LSAs never cross an area boundary. The link-state ID for a network LSA is the IP interface address of the DR that advertises it.

Example 3-31 shows R4’s OSPF LSDB with a focus on the type 2 LSAs.

Example 3-31 R4’s Type 2 LSAs

```
R4# show ip ospf database

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Router Link States (Area 0)

Link ID        ADV Router    Age           Seq#           Checksum Link count
1.1.1.1        1.1.1.1      486          0x8000000B    0x00966C  2
4.4.4.4        4.4.4.4      142          0x80000008    0x001A4F  1

      Net Link States (Area 0)

Link ID        ADV Router    Age           Seq#           Checksum
172.16.14.2   4.4.4.4      142          0x80000007    0x008FB6

<Output omitted>
```

Notice that R4 has only one type 2 LSA in its LSDB. This is expected because there is only one multiaccess network in area 0.

Example 3-32 shows the details of a type 2 LSA on router R4.

Example 3-32 *R4's Type 2 LSA Details*

```

R4# show ip ospf database network

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Net Link States (Area 0)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 170
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 172.16.14.2 (address of Designated Router)
Advertising Router: 4.4.4.4
LS Seq Number: 80000007
Checksum: 0x8FB6
Length: 32
Network Mask: /30
      Attached Router: 4.4.4.4
      Attached Router: 1.1.1.1

```

The content of the displayed type 2 LSA describes the network segment listing the DR address, the attached routers, and the used subnet mask. This information is used by each router participating in OSPF to build the exact picture of the described multiaccess segment, which cannot be fully described with just type 1 LSAs.

OSPF Type 3 Summary LSA

ABRs do not forward type 1 and 2 LSAs between areas to improve OSPF scalability. However, other routers still need to learn how to reach interarea subnets in other areas. OSPF advertises these subnets on ABRs by using type 3 summary LSAs, as shown in Figure 3-14.




Figure 3-14 *OSPF Type 3 LSA*

The ABRs generate type 3 summary LSAs to describe any networks that are owned by an area to the rest of the areas in the OSPF autonomous system, as shown in the figure.

Summary LSAs are flooded throughout a single area only, but are regenerated by ABRs to flood into other areas.

Notice that the figure only illustrates how information is propagated from area 10 to the other areas. Type 3 LSAs are also advertised by ABRs in other direction, from area 20 to area 0, and from area 0 into area 10.

By default, OSPF does not automatically summarize groups of contiguous subnets. OSPF does not summarize a network to its classful boundary. A type 3 LSA is advertised into the backbone area for every subnet that is defined in the originating area, which can cause flooding problems in larger networks.

As a best practice, you can use manual route summarization on ABRs to limit the amount of information that is exchanged between the areas.

Example 3-33 displays R4's OSPF LSDB, with the focus on type 3 LSAs.

Example 3-33 R4's Type 3 LSAs

```
R4# show ip ospf database

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Router Link States (Area 0)
Link ID      ADV Router   Age         Seq#         Checksum Link count
1.1.1.1      1.1.1.1       583        0x8000000B  0x00966C  2
4.4.4.4      4.4.4.4       238        0x80000008  0x001A4F  1

      Net Link States (Area 0)
Link ID      ADV Router   Age         Seq#         Checksum
172.16.14.2  4.4.4.4       238        0x80000007  0x008FB6

      Summary Net Link States (Area 0)
Link ID      ADV Router   Age         Seq#         Checksum
172.16.12.0  1.1.1.1       583        0x80000007  0x00C567
172.16.13.0  1.1.1.1       583        0x80000007  0x009CC5
192.168.2.0  1.1.1.1      1322       0x80000002  0x002E5D

<Output omitted>
```

The LSDB on router R4 includes three different type 3 summary LSAs, all advertised into area 1 by the ABR R1.

Example 3-34 shows the details of R4's type 3 LSAs.

Example 3-34 R4's Type 3 LSA Details

```
R4# show ip ospf database summary

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Summary Net Link States (Area 0)

Routing Bit Set on this LSA in topology Base with MTID 0
```

```

LS age: 608
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 172.16.12.0 (summary Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000007
Checksum: 0xC567
Length: 28
Network Mask: /30
          MTID: 0          Metric: 64

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 608
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 172.16.13.0 (summary Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000007
Checksum: 0x9CC5
Length: 28
Network Mask: /30
          MTID: 0          Metric: 10

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 1348
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 192.168.2.0 (summary Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000002
Checksum: 0x2E5D
Length: 28
Network Mask: /24
          MTID: 0          Metric: 65

```

The output in the examples shows detailed information about three type 3 LSAs in the LSDB. Each type 3 LSA has a link-state ID field, which carries the network address, and together with the attached subnet mask describes the interarea network. Notice that all three LSAs were advertised by the router having router ID set to 1.1.1.1, which is the ABR router R1.

OSPF Type 4 ASBR Summary LSA

Figure 3-15 shows a type 4 summary LSA generated by an ABR only when an ASBR exists within an area. A type 4 LSA identifies the ASBR and provides a route to the

ASBR. The link-state ID is set to the ASBR router ID. All traffic that is destined to an external autonomous system requires routing table knowledge of the ASBR that originated the external routes.




Figure 3-15 OSPF Type 4 LSA

In the figure, the ASBR sends a type 1 router LSA with a bit (known as the *external bit*) that is set to identify itself as an ASBR. When the ABR (identified with the border bit in the router LSA) receives this type 1 LSA, it builds a type 4 LSA and floods it to the backbone, area 0. Subsequent ABRs regenerate a type 4 LSA to flood it into their areas.

Example 3-35 shows R4's OSPF LSDB with a focus on type 4 LSAs.

Example 3-35 R4's Type 4 LSAs

```
R4# show ip ospf database

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Router Link States (Area 0)

Link ID        ADV Router    Age         Seq#          Checksum Link count
1.1.1.1        1.1.1.1      666        0x8000000B   0x00966C  2
4.4.4.4        4.4.4.4      321        0x80000008   0x001A4F  1

      Net Link States (Area 0)

Link ID        ADV Router    Age         Seq#          Checksum
172.16.14.2    4.4.4.4      321        0x80000007   0x008FB6

      Summary Net Link States (Area 0)

Link ID        ADV Router    Age         Seq#          Checksum
172.16.12.0    1.1.1.1      666        0x80000007   0x00C567
172.16.13.0    1.1.1.1      666        0x80000007   0x009CC5
192.168.2.0    1.1.1.1      1405       0x80000002   0x002E5D

      Summary ASB Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
3.3.3.3	1.1.1.1	1405	0x80000002	0x0035EB
Type-5 AS External Link States				
Link ID	ADV Router	Age	Seq#	Checksum Tag
10.0.0.0	3.3.3.3	1351	0x80000002	0x000980 0

There is only one type 4 LSA present in the R4 OSPF database. The type 4 LSA was generated by ABR R1 and describing the ASBR with the router ID 3.3.3.3.

Example 3-36 shows the details of the type 4 LSA on R4.

Example 3-36 R4's Type 4 LSA Details

```
R4# show ip ospf database asbr-summary

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Summary ASB Link States (Area 0)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 1420
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 3.3.3.3 (AS Boundary Router address)
Advertising Router: 1.1.1.1
LS Seq Number: 80000002
Checksum: 0x35EB
Length: 28
Network Mask: /0
      MTID: 0          Metric: 10
```

A type 4 LSA contains information about the existence of the ASBR in the OSPF autonomous system. The information is advertised to R4 from R1, which recognizes the ASBR capability of R3 with a router ID of 3.3.3.3.

OSPF Type 5 External LSA

Figure 3-16 shows type 5 external LSAs used to describe routes to networks outside the OSPF autonomous system. Type 5 LSAs are originated by the ASBR and are flooded to the entire autonomous system.




Figure 3-16 OSPF Type 5 LSA

The link-state ID is the external network number. Because of the flooding scope and depending on the number of external networks, the default lack of route summarization can also be a major issue with external LSAs. Therefore, you should consider summarization of external network numbers at the ASBR to reduce flooding problems.

Example 3-37 shows R4's OSPF LSDB, with a focus on type 5 LSAs.

Example 3-37 R4's OSPF LSDB

```
R4# show ip ospf database
```

```
OSPF Router with ID (4.4.4.4) (Process ID 1)
```

```
Router Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	724	0x8000000B	0x00966C	2
4.4.4.4	4.4.4.4	380	0x80000008	0x001A4F	1

```
Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
172.16.14.2	4.4.4.4	380	0x80000007	0x008FB6

```
Summary Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
172.16.12.0	1.1.1.1	724	0x80000007	0x00C567
172.16.13.0	1.1.1.1	724	0x80000007	0x009CC5
192.168.2.0	1.1.1.1	1463	0x80000002	0x002E5D

```
Summary ASB Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
3.3.3.3	1.1.1.1	1463	0x80000002	0x0035EB

```
Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum Tag
10.0.0.0	3.3.3.3	1410	0x80000002	0x000980 0

The LSDB on R4 contains one external LSA describing external network 10.0.0.0, which was advertised into OSPF by router R3 with a router ID 3.3.3.3.

Example 3-38 shows the details of a type 5 LSA on R4.

Example 3-38 R4's Type 5 LSA Details

```
R4# show ip ospf database external

      OSPF Router with ID (4.4.4.4) (Process ID 1)

      Type-5 AS External Link States

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 1434
Options: (No TOS-capability, DC, Upward)
LS Type: AS External Link
Link State ID: 10.0.0.0 (External Network Number )
Advertising Router: 3.3.3.3
LS Seq Number: 80000002
Checksum: 0x980
Length: 36
Network Mask: /16
    Metric Type: 2 (Larger than any link state path)
    MTID: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

An external LSA on R4 describes the external network 10.0.0.0 with the subnet mask /16. The LSA is advertised by the R3 with a router ID 3.3.3.3. The zero forwarding address tells the rest of the routers in the OSPF domain that ASBR itself is the gateway to get to the external routes. Router R4 gathers information described in the type 5 LSA combined with the information received in the type 4 LSA, which describes the ASBR capability of router R3. This way, R4 learns how to reach the external networks.

Periodic OSPF Database Changes

Although OSPF does not refresh routing updates periodically, it does reflood LSAs every 30 minutes. Each LSA includes the link-state age variable, which counts the age of the LSA packet. When a network change occurs, the LSA's advertising router generates an updated LSA to reflect the change in the network topology. Each updated LSA

includes incremented sequence number so that other routers can distinguish an updated LSA from the old one.

If the LS age variable reaches 30 minutes, meaning that there was no updated LSA created in the last half an hour, it gets automatically regenerated with an increased sequence number and flooded through the OSPF autonomous system. Only the router that originally generated the LSA, the one with the directly connected link, will resend the LSA every 30 minutes.

The output of the OSPF LSDB reveals the value of the current link-state age timer for all LSAs. In a normally operating network, you will not see the age variable with values higher than 1800 seconds.

When an LSA reaches a max age of 60 minutes in the LSDB, it is removed from the LSDB, and the router will perform a new SPF calculation. The router floods the LSA to other routers, informing them to remove the LSA as well.

Because this update is only used to refresh the LSDB, it is sometimes called a *paranoid update*.

Exchanging and Synchronizing LSDBs

Once a bidirectional adjacency is formed, OSPF neighbors follow an exact procedure to synchronize the LSDBs between them.

When routers that are running OSPF are initialized, an exchange process using the hello protocol is the first procedure. The exchange process that happens when routers appear on the network is illustrated in the Figure 3-17 and detailed in the list that follows.

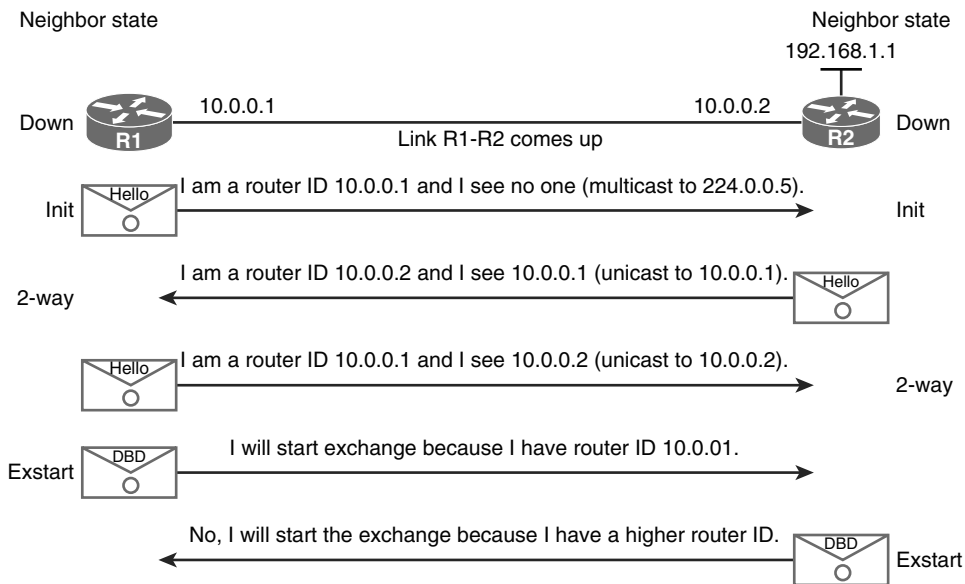


Figure 3-17 *Establishing Neighbor Adjacencies*

- Router R1 is enabled on the LAN and is in a down state because it has not exchanged information with any other router. It begins by sending a Hello packet through each of its interfaces that are participating in OSPF, even though it does not know the identity of the DR or of any other routers. The Hello packet is sent out using the multicast address 224.0.0.5.
- All directly connected routers that are running OSPF receive the Hello packet from router R1 and add R1 to their lists of neighbors. After adding R1 to the list, other routers are in the Init state.
- Each router that received the Hello packet sends a unicast reply Hello packet to R1 with its corresponding information. The neighbor field in the Hello packet includes all neighboring routers and R1.
- When R1 receives these Hello packets, it adds all the routers that had its router ID in their Hello packets to its own neighbor relationship database. After this process, R1 is in the 2-way state. At this point, all routers that have each other in their lists of neighbors have established bidirectional communication.

If the link type is a broadcast network, like Ethernet, a DR and BDR election occurs before the neighboring state proceeds to the next phase.

In the ExStart state, a master-slave relationship is determined between the adjacent neighbors. The router with the higher router ID acts as the master during the exchange process. In Figure 3-17, R2 becomes the master.

Routers R1 and R2 exchange one or more DBD packets while they are in the Exchange state. A DBD includes information about the LSA entry header that appears in the LSDB of the router. The entries can be about a link or a network. Each LSA entry header includes information about the link-state type, the address of the advertising router, the cost of the link, and the sequence number. The router uses the sequence number to determine the “newness” of the received link-state information.

When the router receives the DBD, it performs these actions, as shown in Figure 3-18:

- It acknowledges the receipt of the DBD using the LSAck packet.
- It compares the information that it received with the information that it has. If the DBD has a more up-to-date link-state entry, the router sends an LSR to the other router. When routers start sending LSRs, they are in the loading state.
- The other router responds with the complete information about the requested entry in an LSU packet. Again, when the router receives an LSU, it returns an LSAck.

The router adds the new link-state entries to its LSDB.

When all LSRs have been satisfied for a given router, the adjacent routers are considered synchronized. They are in a Full state, and their LSDBs should be identical. The routers must be in a Full state before they can route traffic.

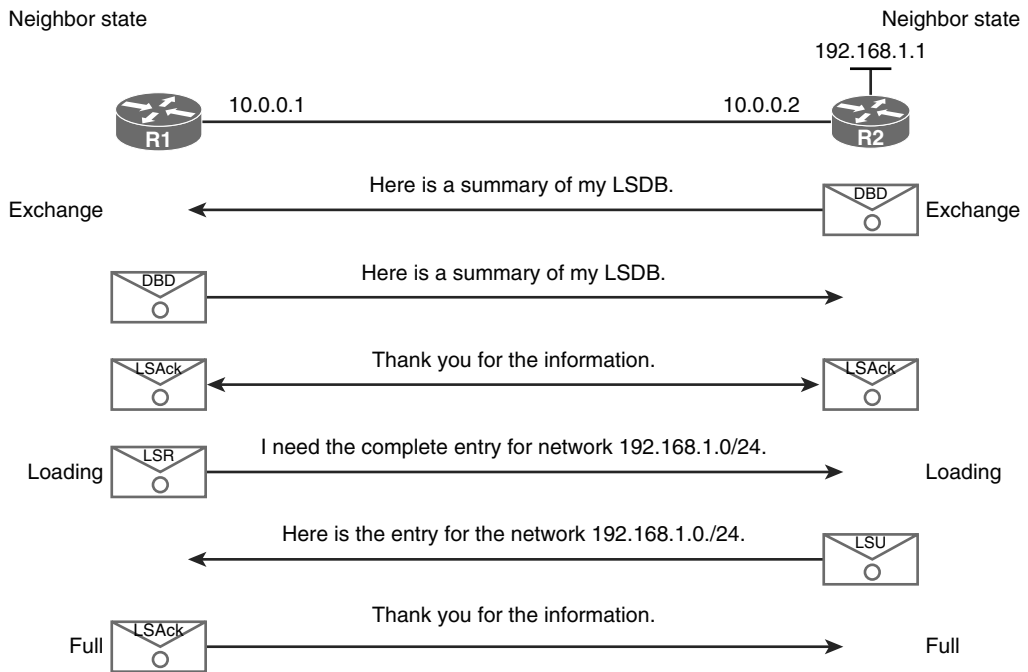


Figure 3-18 *Exchanging and Synchronizing a LSDB*

Synchronizing the LSDB on Multiaccess Networks

On multiaccess segments like Ethernet, OSPF optimizes the LSDB synchronization and the exchange of LSAs. When routers form a neighbor relationship on a multiaccess segment, the DR and BDR election takes place when routers are in the 2-Way state. The router with a highest OSPF priority, or highest router ID in case of a tie, is elected as a DR. Similarly, the router with the second highest priority or router ID becomes the BDR.

While the DR and BDR proceed in establishing the neighborhood with all routers on the segment, other routers establish full adjacency only with the DR and BDR. The neighbor state of other neighbors stays in the 2-Way state.

Non-DR routers exchange their databases only with the DR. The DR takes care to synchronize any new or changed LSAs with the rest of the routers on the segment.

In the flooding process that is illustrated in Figure 3-19, routers perform the following steps:

- Step 1.** A router notices a change in a link state and multicasts an LSU packet (which includes the updated LSA entry) to all OSPF DRs and BDRs at multicast address 224.0.0.6. An LSU packet may contain several distinct LSAs.
- Step 2.** The DR acknowledges receipt of the change and floods the LSU to others on the network using the OSPF multicast address 224.0.0.5.

- Step 3.** After receiving the LSU, each router responds to the DR with an LSAck. To make the flooding procedure reliable, each LSA must be acknowledged separately.
- Step 4.** The router updates its LSDB using the LSU that includes the changed LSA.




Figure 3-19 Synchronizing the LSDB on an Multiaccess Network

Running the SPF Algorithm

Every time there is a change in the network topology, OSPF needs to reevaluate its shortest path calculations. OSPF uses SPF to determine best paths toward destinations. The network topology that is described in the LSDB is used as an input for calculation. Network topology change can influence best path selection; therefore, routers must rerun SPF each time there is an intra-area topology change.

Interarea changes, which are described in type 3 LSAs, do not trigger the SPF recalculation because the input information for the best path calculation remains unchanged. The router determines the best paths for interarea routes based on the calculation of the best path toward the ABR. The changes that are described in type 3 LSAs do not influence how the router reaches the ABR; therefore, SPF recalculation is not needed.

You can verify how often the SPF algorithm was executed by using the `show ip ospf` command, as shown in Example 3-39. The output will also show you when the algorithm was last executed.

Example 3-39 Verifying OSPF Frequency of the SPF Algorithm

```
R1# show ip ospf | begin Area
Area BACKBONE(0) (Inactive)
  Number of interfaces in this area is 1
  Area has no authentication
  SPF algorithm last executed 00:35:04:959 ago
  SPF algorithm executed 5 times
  Area ranges are
  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
Area 1
    
```

Configuring OSPF Path Selection

In this section, we will analyze and influence how OSPF determines link costs to calculate the best path, continuing with the previous topology shown in Figure 3-20.

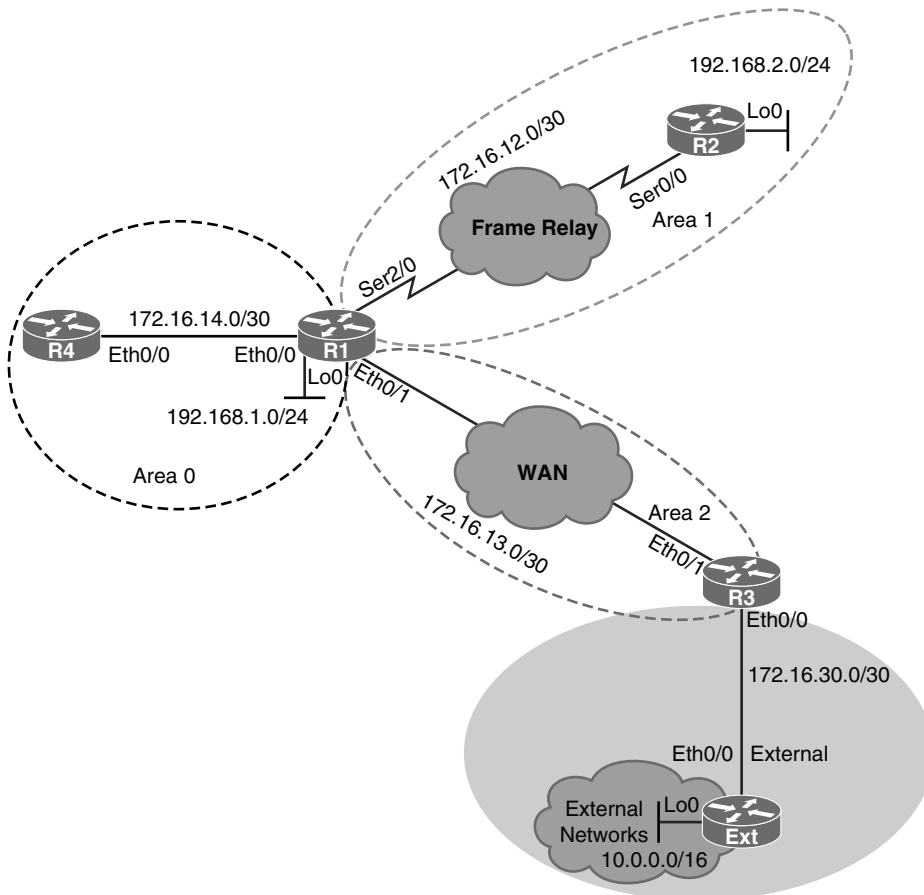


Figure 3-20 Topology for OSPF Path Selection

OSPF Path Selection

In Example 3-40, the output of the `show ip ospf` command verifies how many times the SFP algorithm was executed.

Example 3-40 *Verifying the SPF Calculations on R1*

```

R1# show ip ospf | begin Area
    Area BACKBONE(0)
        Number of interfaces in this area is 2
        Area has no authentication
        SPF algorithm last executed 00:02:17.777 ago
        SPF algorithm executed 3 times
        Area ranges are
        Number of LSA 7. Checksum Sum 0x0348C4
        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
<Output omitted>

```

The command output shows you how many times SPF has already run, together with the information about the last execution.

On R1, the link toward R4 is disabled and reenabled in Example 3-41. The number of SPF executions is verified afterward.

Example 3-41 *SPF Calculated on R1*

```

R1(config)# interface ethernet 0/0
R1(config-if)# shutdown
*Jan 31 12:33:20.617: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Ethernet0/0 from
FULL to DOWN, Neighbor Down: Interface down or detached
*Jan 31 12:33:22.613: %LINK-5-CHANGED: Interface Ethernet0/0, changed state to
administratively down
*Jan 31 12:33:23.617: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0,
changed state to down
R1(config-if)# no shutdown
*Jan 31 12:33:29.125: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to up
*Jan 31 12:33:30.129: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0,
changed state to up
*Jan 31 12:33:35.040: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Ethernet0/0 from
LOADING to FULL, Loading Done
R1(config-if)# do show ip ospf | begin Area
    Area BACKBONE(0)
        Number of interfaces in this area is 2
        Area has no authentication
        SPF algorithm last executed 00:00:07.752 ago
        SPF algorithm executed 5 times
        Area ranges are
        Number of LSA 7. Checksum Sum 0x033ACB
        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
<Output omitted>

```

Disabling the interface on R1 in area 0 triggers SPF calculation. Enabling the interface back into the OSPF triggers another SPF calculation. As a result, the counter displayed in the output has increased.

Link flap caused two recalculations of SPF algorithm. Frequent changes of link status can lead to frequent SPF calculation, which can utilize router resources.

OSPF Best Path Calculation

Once LSDBs are synchronized among OSPF neighbors, each router needs to determine on its own the best paths over the network topology.

When SPF is trying to determine the best path toward a known destination, it compares total costs of specific paths against each other. The paths with the lowest costs are selected as the best paths. The OSPF cost is an indication of the overhead to send packets over an interface. OSPF cost is computed automatically for each interface that is assigned into an OSPF process, using the following formula:

$$\text{Cost} = \text{Reference bandwidth} / \text{Interface bandwidth}$$

The cost value is a 16-bit positive number between 1 and 65,535, where a lower value is a more desirable metric. Reference bandwidth is set to 100 Mbps by default.

On high-bandwidth links (100 Mbps and more), automatic cost assignment no longer works. (It would result in all costs being equal to 1.) On these links, OSPF costs must be set manually on each interface.

For example, a 64-Kbps link gets a metric of 1562, and a T1 link gets a metric of 64. Cost is applied on all router link paths, and route decisions are made on the total cost of a path. The metric is only relevant on an outbound path; route decisions are not made for inbound traffic. The OSPF cost is recomputed after every bandwidth change, and the Dijkstra's algorithm determines the best path by adding all link costs along a path.

Example 3-42 reveals the interface bandwidth and the OSPF cost of the Frame Relay interface on R1.

Example 3-42 *Examining the Interface Bandwidth and OSPF Cost on R1*

```

R1# show interface serial 2/0
Serial2/0 is up, line protocol is up
  Hardware is M4T
  Internet address is 172.16.12.1/30
  MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255

```

```

Encapsulation FRAME-RELAY, crc 16, loopback not set
<Output omitted>

R1# show ip ospf interface serial 2/0
Serial2/0 is up, line protocol is up
  Internet Address 172.16.12.1/30, Area 1, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
         0          64          no            no            Base
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 2.2.2.2, Interface address 172.16.12.2
  Backup Designated router (ID) 1.1.1.1, Interface address 172.16.12.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
<Output omitted>

```

The first command in the output displays bandwidth of the serial interface, which connects R1 with R2. The second output shows that OSPF calculated the cost of 64 for this interface. The cost was calculated by dividing the reference bandwidth of 100 Mbps with the actual interface bandwidth.

Default OSPF Costs

OSPF calculates the default interface costs, based on the interface type and the default reference bandwidth, shown in Table 3-2.

Table 3-2 *Default OSPF Costs*

Link Type	Default Cost
T1 (1.544-Mbps serial link)	64
Ethernet	10
Fast Ethernet	1
Gigabit Ethernet	1
10-Gigabit Ethernet	1

The default reference bandwidth of 100 Mbps is not suitable to calculate OSPF costs for links faster than Fast Ethernet. All such links gets assigned cost of 1, and OSPF cannot optimally choose the shortest path as it treats all the high-speed links as equal.

To improve OSPF behavior, you can adjust reference bandwidth to a higher value by using the **auto-cost reference-bandwidth** OSPF configuration command.

In Example 3-43, the reference bandwidth on R1 is changed to 10 Gbps.

Example 3-43 *Modifying the Reference Bandwidth on R1*

```
R1(config)# router ospf 1
R1(config-router)# auto-cost reference-bandwidth 10000
% OSPF: Reference bandwidth is changed.
Please ensure reference bandwidth is consistent across all routers.
```

You can change the OSPF reference bandwidth under OSPF configuration mode by using the **auto-cost reference-bandwidth** command. The reference bandwidth value is inserted in megabits per second.

Notice also the warning that is displayed by the prompt. Only consistent reference bandwidth across OSPF domain ensures that all routers calculate the best paths correctly.

Example 3-44 highlights the OSPF link cost of R1's serial interface.

Example 3-44 *R1's OSPF Cost on Serial 2/0*

```
R1# show ip ospf interface serial 2/0
Serial2/0 is up, line protocol is up
Internet Address 172.16.12.1/30, Area 1, Attached via Network Statement
Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 6476
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                6476        no            no            Base
<Output omitted>
```

The changed OSPF reference bandwidth results in updated OSPF costs for all interfaces. The cost for Serial 2/0 interface has increased from 64 to 6476. The new cost was calculated based on reference bandwidth of 10 Gbps divided by the interface speed of 1.544 Mbps.

In Example 3-45, the interface bandwidth is changed on R1's Serial 2/0 interface.

Example 3-45 *Changing the Interface Bandwidth on R1's Serial 2/0 Interface*

```
R1(config)# interface serial 2/0
R1(config-if)# bandwidth 10000
```

Changing the OSPF reference bandwidth influences the cost of all local interfaces included in the OSPF. Commonly, you will need to influence the cost just for a specific interface on the router. Using the **bandwidth** command, you can change how IOS treats a specific interface by default. Bandwidth setting changes the artificial value of the interface bandwidth that is derived by IOS based on the interface type. A manually set bandwidth value on the interface overrides the default value and is used by OSPF as input to the interface cost calculation.

Modifying the bandwidth not only influences OSPF but also other routing protocols like EIGRP, which takes the bandwidth into account when calculating the EIGRP metric.

The interface bandwidth and the OSPF cost of the serial interface on R1 are verified in Example 3-46.

Example 3-46 *Verifying the Interface Bandwidth and OSPF Cost on R1's Serial 2/0 Interface*

```
R1# show interfaces serial 2/0
Serial2/0 is up, line protocol is up
  Hardware is M4T
  Internet address is 172.16.12.1/30
  MTU 1500 bytes, BW 10000 Kbit/sec, DLY 20000 usec,
<Output omitted>
R1# show ip ospf interface serial 2/0
Serial2/0 is up, line protocol is up
  Internet Address 172.16.12.1/30, Area 1, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 1000
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
      0          1000        no            no            Base
<Output omitted>
```

The interface verification command displays the updated interface bandwidth, which was manually set to 10 Mbps. The change of the interface bandwidth is also reflected in the newly calculated OSPF cost, which is shown in the second output. The cost was calculated by dividing the reference bandwidth of 10000 Mbps with the configured bandwidth of 10 Mbps.

In Example 3-47, the OSPF cost of the serial interface link on R1 is changed using the `ip ospf cost` interface command.

Example 3-47 *Changing the OSPF Cost on an Interface*

```
R1(config)# interface serial 2/0
R1(config-if)# ip ospf cost 500
```

Using the OSPF interface configuration command, you can directly change the OSPF cost of specific interface. Cost of the interface can be set to a value between 1 and 65,535. This command overrides whatever value is calculated based on the reference bandwidth and the interface bandwidth.

The OSPF cost of the serial interface on R1 is verified in Example 3-48.

Example 3-48 *Verifying the OSPF Interface Costs on R1*

```
R1# show ip ospf interface brief
Interface      PID      Area          IP Address/Mask      Cost  State Nbrs F/C
Lo0            1        0             192.168.1.1/24       1     P2P   0/0
Et0/0          1        0             172.16.14.1/30       1000  DR    1/1
Se2/0          1        1             172.16.12.1/30       500   BDR   1/1
Et0/1          1        2             172.16.13.1/30       1000  BDR   1/1
```

To verify the OSPF cost, you can also use the **brief** keyword in the **show ip ospf interface** command. The verification command displays the summarized information on all OSPF-enabled interfaces, including the cost of the interface. You can notice the updated cost of the serial interface, which was manually configured in the previous step. In the output, you can observe the manually configured cost setting of the serial interface.

Calculating the Cost of Intra-Area Routes

To calculate the cost of intra-area routes, the router first analyzes OSPF database and identifies all subnets within its area. For each possible route, OSPF calculates the cost to reach the destination by summing up the individual interface costs. For each subnet, the route with the lowest total cost is selected as the best route.

Analyzing the topology in the Figure 3-21 from R1's perspective, notice that it can reach intra-area network A either via ABR1 or ABR2. The autonomous system path through ABR1 is associated with the lower cost, so it will be selected as the best path.

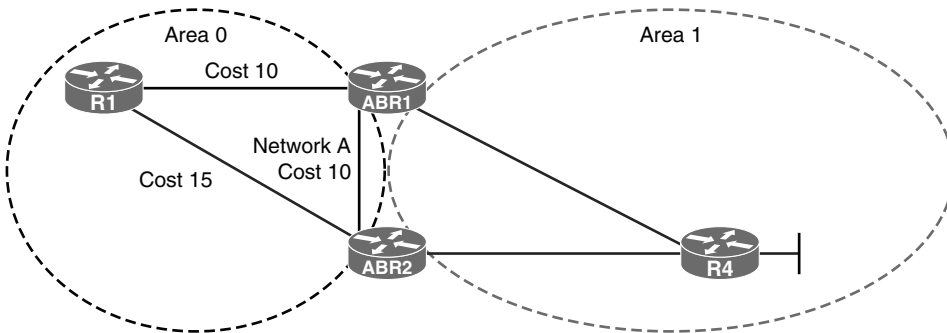


Figure 3-21 *Calculating the Cost of Intra-Area Routes*

In a scenario where two paths would have the same lowest total cost, both routes would be selected as the best paths and inserted in the routing table. As a result, a router would perform equal-cost load balancing.

Calculating the Cost of Interarea Routes

The internal OSPF router within an area receives only summarized info about interarea routes. As a result, the cost of an interarea route cannot be calculated the same way as for the intra-area routes.

When ABRs propagate information about the interarea routes with type 3 LSAs, they include their lowest cost to reach a specific subnet in the advertisement. The internal router adds its cost to reach a specific ABR to the cost announced in a type 3 LSA. Then it selects the route with the lowest total cost as the best route.

Router R1, in Figure 3-22, learns about network B from both ABRs. ABR2 in type 2 LSA reports the lowest cost to reach network B as 6, while ABR1 reports the cost of 21.

Router R1 determines the lowest cost to reach both ABRs and adds this cost to the one received in LSA. Router R1 selects the route via ABR2 as the total lowest cost route and tries to install it into the routing table.

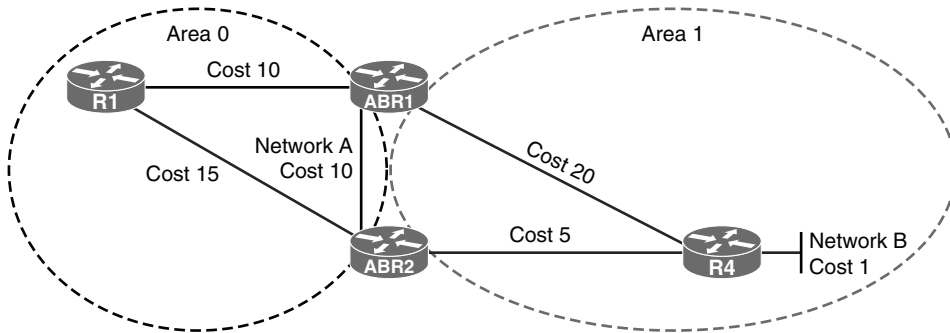


Figure 3-22 *Calculating the Cost of Interarea Routes*

Selecting Between Intra-Area and Interarea Routes

To eliminate the single point of failure on area borders, at least two ABRs are used in most networks. As a result, ABR can learn about a specific subnet from internal routers and also from the other ABR. ABR can learn an intra-area route and also an interarea route for the same destination. Even though the interarea route could have lower cost to the specific subnet, the intra-area path is always the preferred choice.

In the example topology in Figure 3-23, ABR1 learns about network B directly from a router R4 and also from the ABR2. Even though the interarea path has a cost of 16, the intra-area path with a total cost of 21 is selected as the best path.




Figure 3-23 *Selecting Between Intra-Area and Interarea Routes*

Optimizing OSPF Behavior

Scalability, improved CPU and memory utilization, and the ability to mix small routers with large routers are all the benefits of using proper route summarization techniques. A

key feature of the OSPF protocol is the ability to summarize routes at area and autonomous system boundaries.

Route summarization is important because it reduces the amount of the OSPF LSA flooding and the sizes of LSDBs and routing tables, which also reduces the memory and the CPU utilization on the routers. An OSPF network can scale to very large sizes, partially because of the route summarization.

The OSPF protocol defines several special-case area types, including stub areas, totally stubby areas, and NSSAs. The purpose of all three types of stub areas is to inject default routes into an area so that external and summary LSAs are not flooded. Stub areas are designed to reduce the amount of flooding, the LSDB size, and the routing table size in routers within the area. Network designers should always consider using stub area techniques when building networks. Stub area techniques improve performance in OSPF networks and allow the network to scale to significantly larger sizes.

Default routes reduce the routing table size, and also reduce the memory and the CPU utilization. OSPF injects a default route unconditionally or based on the presence of a default route inside the routing table.

This section defines different types of route summarization and describes the configuration commands for each type. It also describes the OSPF area types and the benefits of default routes.

Upon completing this section, you will be able to do the following:

- Describe the properties of OSPF route summarization
- Describe benefits of route summarization in OSPF
- Configure summarization on ABR
- Configure summarization on ASBR
- Configure the cost of OSPF default route
- Describe how you can use default routes and stub routing to direct traffic toward the Internet
- Describe the NSSA areas
- Configure the default route using the **default-information originate** command

OSPF Route Summarization

Route summarization is a key to scalability in OSPF. Route summarization helps solve two major problems:

- Large routing tables
- Frequent LSA flooding throughout the autonomous system

Every time that a route disappears in one area, routers in other areas also get involved in shortest-path calculation. To reduce the size of the area database, you can configure summarization on an area boundary or autonomous system boundary.

Normally, type 1 and type 2 LSAs are generated inside each area and translated into type 3 LSAs in other areas. With route summarization, the ABRs or ASBRs consolidate multiple routes into a single advertisement. ABRs summarize type 3 LSAs, and ASBRs summarize type 5 LSAs. Instead of advertising many specific prefixes, advertise only one summary prefix.

If the OSPF design includes many ABRs or ASBRs, suboptimal routing is possible. This is one of the drawbacks of summarization.

Route summarization requires a good addressing plan—an assignment of subnets and addresses that is based on the OSPF area structure and lends itself to aggregation at the OSPF area borders.

Benefits of Route Summarization

Route summarization directly affects the amount of bandwidth, CPU power, and memory resources that the OSPF routing process consumes. Without route summarization, every specific-link LSA is propagated into the OSPF backbone and beyond, causing unnecessary network traffic and router overhead.

With route summarization, only the summarized routes are propagated into the backbone (area 0), as illustrated in Figure 3-24. Summarization prevents every router from having to rerun the SPF algorithm, increases the stability of the network, and reduces unnecessary LSA flooding. Also, if a network link fails, the topology change is not propagated into the backbone (and other areas by way of the backbone). Specific-link LSA flooding outside the area does not occur.

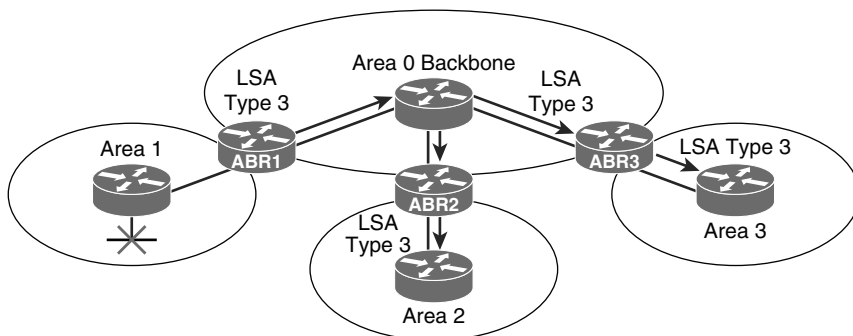


Figure 3-24 OSPF Route Summarization

Receiving a type 3 LSA into its area does not cause a router to run the SPF algorithm. The routes being advertised in the type 3 LSAs are appropriately added to or deleted from the router's routing table, but an SPF calculation is not done.

Configuring OSPF Route Summarization

In this section, we will implement route summarization on the area borders in an OSPF environment, shown in Figure 3-25. We will summarize the OSPF network using different subnet sizes and examine the impact of summarization on the OSPF database and routing.

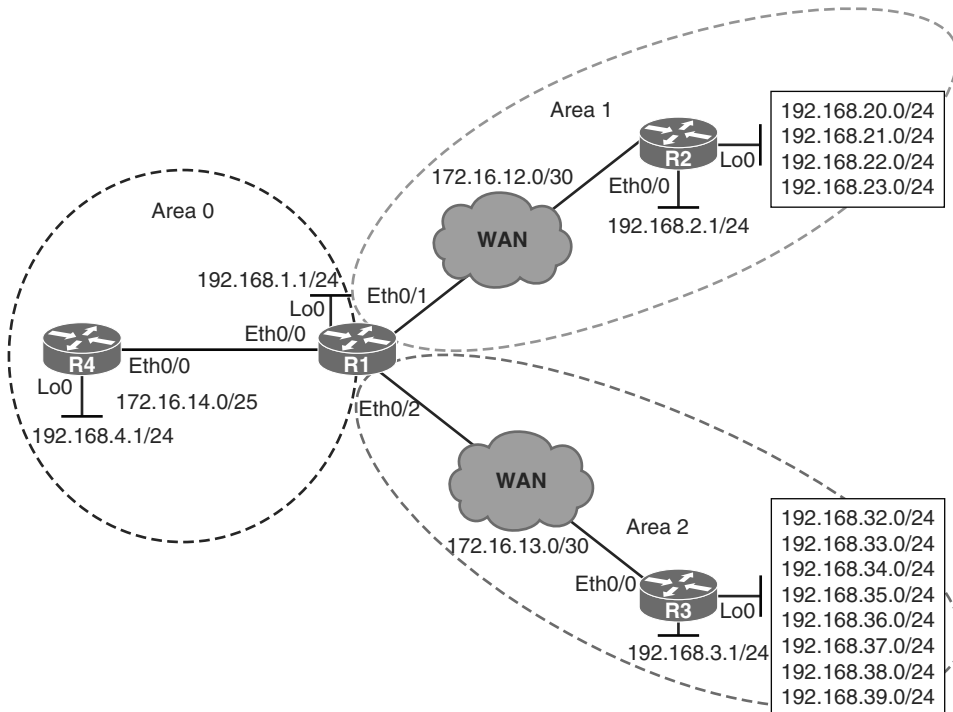


Figure 3-25 OSPF Route Summarization Topology

Example 3-49 displays OSPF routes in R1's routing table.

Example 3-49 OSPF Routes in R1's Routing Table

```
R1# show ip route ospf
<Output omitted>

O    192.168.2.0/24 [110/11] via 172.16.12.2, 00:41:47, Ethernet0/1
O    192.168.3.0/24 [110/11] via 172.16.13.2, 00:40:01, Ethernet0/2
O    192.168.4.0/24 [110/11] via 172.16.14.2, 00:38:09, Ethernet0/0
O    192.168.20.0/24 [110/11] via 172.16.12.2, 00:41:37, Ethernet0/1
O    192.168.21.0/24 [110/11] via 172.16.12.2, 01:03:46, Ethernet0/1
O    192.168.22.0/24 [110/11] via 172.16.12.2, 01:03:36, Ethernet0/1
O    192.168.23.0/24 [110/11] via 172.16.12.2, 01:03:26, Ethernet0/1
O    192.168.32.0/24 [110/11] via 172.16.13.2, 00:40:14, Ethernet0/2
O    192.168.33.0/24 [110/11] via 172.16.13.2, 00:57:01, Ethernet0/2
```

```
O 192.168.34.0/24 [110/11] via 172.16.13.2, 00:01:16, Ethernet0/2
O 192.168.35.0/24 [110/11] via 172.16.13.2, 00:01:06, Ethernet0/2
O 192.168.36.0/24 [110/11] via 172.16.13.2, 00:00:56, Ethernet0/2
O 192.168.37.0/24 [110/11] via 172.16.13.2, 00:00:46, Ethernet0/2
O 192.168.38.0/24 [110/11] via 172.16.13.2, 00:00:32, Ethernet0/2
O 192.168.39.0/24 [110/11] via 172.16.13.2, 00:00:18, Ethernet0/2
```

Apart from the loopback networks (192.168.x.0/24 where x is the router ID), notice the four Class C networks advertised by R2 (192.168.20.0/24 to 192.168.23.0/24) and eight Class C networks advertised by R3 (192.168.32.0/24 to 192.168.39.0/24).

Example 3-50 displays OSPF routes in R4's routing table.

Example 3-50 OSPF Routes in R4's Routing Table

```
R4# show ip route ospf
<Output omitted>

    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA 172.16.12.0/30 [110/20] via 172.16.14.1, 01:17:30, Ethernet0/0
O IA 172.16.13.0/30 [110/20] via 172.16.14.1, 01:17:30, Ethernet0/0
O    192.168.1.0/24 [110/11] via 172.16.14.1, 01:17:30, Ethernet0/0
O IA 192.168.2.0/24 [110/21] via 172.16.14.1, 00:49:23, Ethernet0/0
O IA 192.168.3.0/24 [110/21] via 172.16.14.1, 00:47:37, Ethernet0/0
O IA 192.168.20.0/24 [110/21] via 172.16.14.1, 00:49:08, Ethernet0/0
O IA 192.168.21.0/24 [110/21] via 172.16.14.1, 01:11:23, Ethernet0/0
O IA 192.168.22.0/24 [110/21] via 172.16.14.1, 01:11:13, Ethernet0/0
O IA 192.168.23.0/24 [110/21] via 172.16.14.1, 01:11:03, Ethernet0/0
O IA 192.168.32.0/24 [110/21] via 172.16.14.1, 00:47:50, Ethernet0/0
O IA 192.168.33.0/24 [110/21] via 172.16.14.1, 01:04:37, Ethernet0/0
O IA 192.168.34.0/24 [110/21] via 172.16.14.1, 00:02:26, Ethernet0/0
O IA 192.168.35.0/24 [110/21] via 172.16.14.1, 00:02:16, Ethernet0/0
O IA 192.168.36.0/24 [110/21] via 172.16.14.1, 00:02:06, Ethernet0/0
O IA 192.168.37.0/24 [110/21] via 172.16.14.1, 00:01:56, Ethernet0/0
O IA 192.168.38.0/24 [110/21] via 172.16.14.1, 00:01:43, Ethernet0/0
O IA 192.168.39.0/24 [110/21] via 172.16.14.1, 00:01:28, Ethernet0/0
```

Notice that the same networks are listed as interarea summary routes. They are being flooded into each area without any summarization on the area borders. You can see the respective routes that are received from the other areas on R2 and R3 as well.

Example 3-51 shows the OSPF database on R4.

Example 3-51 R4's OSPF LSDB

```
R4# show ip ospf database

    OSPF Router with ID (4.4.4.4) (Process ID 1)
```

Router Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	1110	0x80000006	0x008A7E	2
4.4.4.4	4.4.4.4	1406	0x80000005	0x00D915	2
Net Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
172.16.14.1	1.1.1.1	1373	0x80000003	0x004192	
Summary Net Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
172.16.12.0	1.1.1.1	553	0x80000008	0x00A5BC	
172.16.13.0	1.1.1.1	553	0x80000008	0x009AC6	
192.168.2.0	1.1.1.1	1541	0x80000006	0x0008B5	
192.168.3.0	1.1.1.1	3607	0x80000007	0x008C3A	
192.168.20.0	1.1.1.1	1541	0x8000000B	0x00376F	
192.168.21.0	1.1.1.1	1800	0x80000004	0x003A72	
192.168.22.0	1.1.1.1	1800	0x80000004	0x002F7C	
192.168.23.0	1.1.1.1	1800	0x80000004	0x002486	
192.168.32.0	1.1.1.1	3607	0x80000007	0x004C5D	
192.168.33.0	1.1.1.1	3607	0x80000008	0x003F68	
192.168.34.0	1.1.1.1	3607	0x80000002	0x00406C	
192.168.35.0	1.1.1.1	3607	0x80000002	0x003576	
192.168.36.0	1.1.1.1	3607	0x80000002	0x002A80	
192.168.37.0	1.1.1.1	3607	0x80000002	0x001F8A	
192.168.38.0	1.1.1.1	3607	0x80000002	0x001494	
192.168.39.0	1.1.1.1	3607	0x80000002	0x00099E	

Notice the corresponding LSA 3 updates for each interarea summary route received from R1.

In Example 3-52, R1 summarizes four networks (192.168.20.0/24 to 192.168.23.0/24) in area 1 and the eight networks (192.168.32.0/24 to 192.168.39.0/24) in area 2 using the appropriate address blocks.

Example 3-52 *Configuring Summarization on the ABR*

```
R1(config)# router ospf 1
R1(config-router)# area 1 range 192.168.20.0 255.255.252.0
R1(config-router)# area 2 range 192.168.32.0 255.255.248.0
```

OSPF is a classless routing protocol, which carries subnet mask information along with route information. Therefore, OSPF supports multiple subnet masks for the same major network, which is known as *variable-length subnet masking* (VLSM). OSPF supports

discontiguous subnets because the subnet masks are part of the LSDB. Network numbers in areas should be assigned contiguously to ensure that these addresses can be summarized into a minimal number of summary addresses.

In this scenario, the list of four networks advertised by R2 (192.168.20.0/24 to 192.168.23.0/24) in the routing table of the ABR can be summarized into one address block. The list of networks advertised by R3 (192.168.32.0/24 to 192.168.39.0/24) can also be aggregated by one summary address. All these networks will be summarized on the ABR R1. The block of addresses from 192.168.20.0 through 192.168.23.0/24 can be summarized using 192.168.20.0/22, and the block from 192.168.32.0 through 192.168.39.0/24 can be summarized using 192.168.32.0/21.

To consolidate and summarize routes at an area boundary, use the **area range** command in the router configuration mode. The ABR will summarize routes for a specific area before injecting them into a different area via the backbone as type 3 summary LSAs.

Example 3-53 examines the OSPF routing tables on R2, R3, and R4 with the route summarization on R1. Apart from the loopback networks, you will see the summary block of the other area, respectively.

Example 3-53 OSPF Summarized Routes in the Routing Table

```
R2# show ip route ospf
<Output omitted>

    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA    172.16.13.0/30 [110/20] via 172.16.12.1, 05:27:05, Ethernet0/0
O IA    172.16.14.0/25 [110/20] via 172.16.12.1, 05:07:35, Ethernet0/0
O IA    192.168.1.0/24 [110/11] via 172.16.12.1, 05:27:09, Ethernet0/0
O IA    192.168.3.0/24 [110/21] via 172.16.12.1, 01:24:16, Ethernet0/0
O IA    192.168.4.0/24 [110/21] via 172.16.12.1, 04:32:02, Ethernet0/0
O IA    192.168.32.0/21 [110/21] via 172.16.12.1, 00:57:42, Ethernet0/0
```

```
R3# show ip route ospf
<Output omitted>

    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA    172.16.12.0/30 [110/20] via 172.16.13.1, 05:25:50, Ethernet0/0
O IA    172.16.14.0/25 [110/20] via 172.16.13.1, 05:10:02, Ethernet0/0
O IA    192.168.1.0/24 [110/11] via 172.16.13.1, 05:25:50, Ethernet0/0
O IA    192.168.2.0/24 [110/21] via 172.16.13.1, 04:38:07, Ethernet0/0
O IA    192.168.4.0/24 [110/21] via 172.16.13.1, 04:34:29, Ethernet0/0
O IA    192.168.20.0/22 [110/21] via 172.16.13.1, 01:00:19, Ethernet0/0
```

```
R4# show ip route ospf
<Output omitted>
```

```

172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA 172.16.12.0/30 [110/20] via 172.16.14.1, 05:16:24, Ethernet0/0
O IA 172.16.13.0/30 [110/20] via 172.16.14.1, 05:16:24, Ethernet0/0
O 192.168.1.0/24 [110/11] via 172.16.14.1, 05:16:24, Ethernet0/0
O IA 192.168.2.0/24 [110/21] via 172.16.14.1, 04:48:17, Ethernet0/0
O IA 192.168.3.0/24 [110/21] via 172.16.14.1, 01:36:53, Ethernet0/0
O IA 192.168.20.0/22 [110/21] via 172.16.14.1, 01:10:29, Ethernet0/0
O IA 192.168.32.0/21 [110/21] via 172.16.14.1, 01:10:19, Ethernet0/0

```

In the routing table of R4, you will see the two summarized address blocks from areas 1 and 2.

Example 3-54 shows the OSPF database on the backbone router R4.

Example 3-54 R4's OSPF LSDB

```

R4# show ip ospf database
<Output omitted>

                Summary Net Link States (Area 0)

Link ID        ADV Router    Age         Seq#          Checksum
172.16.12.0    1.1.1.1      599        0x8000000B   0x009FBF
172.16.13.0    1.1.1.1      599        0x8000000B   0x0094C9
192.168.2.0    1.1.1.1      1610       0x80000009   0x0002B8
192.168.3.0    1.1.1.1      98         0x80000004   0x0001BD
192.168.20.0   1.1.1.1      599        0x8000000F   0x002085
192.168.32.0   1.1.1.1      98         0x80000005   0x009B0C

```

Notice the type 3 LSAs for the two summarized address blocks from areas 1 and 2. The type 3 LSAs for the specific networks are no longer in the database.

Example 3-55 displays the OSPF routing table on R1. Notice the two routes to the Null 0 interface. What is the purpose of these routes?

Example 3-55 OSPF Routes in R1's Routing Table

```

R1# show ip route ospf
<Output omitted>

O 192.168.2.0/24 [110/11] via 172.16.12.2, 01:18:25, Ethernet0/1
O 192.168.3.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O 192.168.4.0/24 [110/11] via 172.16.14.2, 01:18:25, Ethernet0/0
O 192.168.20.0/22 is a summary, 01:18:25, Null0
O 192.168.20.0/24 [110/11] via 172.16.12.2, 01:18:25, Ethernet0/1
O 192.168.21.0/24 [110/11] via 172.16.12.2, 01:18:25, Ethernet0/1
O 192.168.22.0/24 [110/11] via 172.16.12.2, 01:18:25, Ethernet0/1

```

```

O   192.168.23.0/24 [110/11] via 172.16.12.2, 01:18:25, Ethernet0/1
O   192.168.32.0/21 is a summary, 01:18:25, Null0
O   192.168.32.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.33.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.34.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.35.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.36.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.37.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.38.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2
O   192.168.39.0/24 [110/11] via 172.16.13.2, 01:18:25, Ethernet0/2

```

Cisco IOS Software creates a summary route to the Null0 interface when manual summarization is configured, to prevent routing loops. For example, if the summarizing router receives a packet to an unknown subnet that is part of the summarized range, the packet matches the summary route based on the longest match. The packet is forwarded to the Null0 interface (in other words, it is dropped), which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

Summarization on ABRs

OSPF offers two methods of route summarization:

- Summarization of internal routes performed on the ABRs
- Summarization of external routes performed on the ASBRs

Without summarization of internal routes, all the prefixes from an area are passed into the backbone as type 3 interarea routes. When summarization is enabled, the ABR intercepts this process and instead injects a single type 3 LSA, which describes the summary route into the backbone, shown in Figure 3-26. Multiple routes inside the area are summarized.




Figure 3-26 Type 3 Summary LSA

To consolidate and summarize routes at an area boundary, use the following command in router configuration mode:

```
area area-id range ip-address mask [advertise | not-advertise] [cost cost]
```


Table 3-3 shows the parameters used with this command. To remove the summarization, use the **no** form of this command.

Table 3-3 *area range Command Parameters*

Parameter	Description
<i>area-id</i>	Identifier of the area about which routes are to be summarized. It can be specified as either a decimal value or as an IP address.
<i>ip-address</i>	IP address.
<i>mask</i>	IP address mask.
advertise	(Optional) Sets the address range status to advertise and generates a type 3 summary LSA.
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.
cost <i>cost</i>	(Optional) Metric or cost for this summary route, which is used during OSPF SPF calculation to determine the shortest paths to the destination. The value can be 0 to 16,777,215.

An internal summary route is generated if at least one subnet within the area falls in the summary address range and the summarized route metric is equal to the lowest cost of all the subnets within the summary address range. Interarea summarization can only be done for the intra-area routes of connected areas, and the ABR creates a route to Null0 to avoid loops in the absence of more specific routes.

Summarization on ASBRs

Summarization can also be performed for external routes, as illustrated in Figure 3-27. Each route that is redistributed into OSPF from other protocols is advertised individually with an external LSA. To reduce the size of the OSPF LSDB, you can configure a summary for external routes. Summarization of external routes can be done on the ASBR for type 5 LSAs (redistributed routes) before injecting them into the OSPF domain. Without summarization, all the redistributed external prefixes from external autonomous systems are passed into the OSPF area. A summary route to Null0 is created automatically for each summary range.




Figure 3-27 Type 5 Summary LSA

To create aggregate addresses for OSPF at an autonomous system boundary, use the following command in router configuration mode:

```
summary-address {{ip-address mask} | {prefix mask}} [not-advertise] [tag tag]
```

The ASBR will summarize external routes before injecting them into the OSPF domain as type 5 external LSAs. Table 3-4 shows the parameters used with the **summary-address** command. To remove a the summarization, use the **no** form of this command.

Table 3-4 summary-address Command Parameters

Parameter	Description
<i>ip-address</i>	Summary address designated for a range of addresses.
<i>mask</i>	IP subnet mask used for the summary route.
<i>prefix</i>	IP route prefix for the destination.
<i>mask</i>	IP subnet mask used for the summary route.
not-advertise	(Optional) Suppress routes that match the specified prefix/mask pair. This keyword applies to OSPF only.
tag tag	(Optional) Tag value that can be used as a “match” value for controlling redistribution via route maps. This keyword applies to OSPF only.

It is recommended practice dictates implementing contiguous IP addressing to achieve optimal summarization results.

OSPF Virtual Links

OSPF’s two-tiered area hierarchy requires that if more than one area is configured, one of the areas must be area 0, the backbone area. All other areas must be directly connected to area 0, and area 0 must be contiguous. OSPF expects all nonbackbone areas to inject routes into the backbone, so that the routes can be distributed to other areas.

A virtual link is a link that allows discontinuous area 0s to be connected, or a disconnected area to be connected to area 0, via a transit area. The OSPF virtual link feature should be used only in very specific cases, for temporary connections or for backup after a failure. Virtual links should not be used as a primary backbone design feature.

The virtual link relies on the stability of the underlying intra-area routing. Virtual links cannot go through more than one area, nor through stub areas. Virtual links can only run through standard nonbackbone areas. If a virtual link needs to be attached to the backbone across two nonbackbone areas, two virtual links are required, one per area.

In Figure 3-28, two companies running OSPF have merged and a direct link does not yet exist between their backbone areas. The resulting area 0 is discontinuous. A logical link (virtual link) is built between the two ABRs, routers A and B, across area 1, a nonbackbone area. The routers at each end of the virtual link become part of the backbone and act as ABRs. This virtual link is similar to a standard OSPF adjacency, except that in a virtual link, neighboring routers do not have to be directly attached.




Figure 3-28 *Virtual Links Are Used to Connect a Discontinuous Area 0*

Figure 3-29 illustrates another example where a nonbackbone area is added to an OSPF network, and a direct physical connection to the existing OSPF area 0 does not yet exist. In this case, area 20 is added, and a virtual link across area 10 is created to provide a logical path between area 20 and the backbone area 0. The OSPF database treats the virtual link between ABR1 and ABR2 as a direct link. For greater stability, loopback interfaces are used as router IDs, and virtual links are created using these loopback addresses.




Figure 3-29 *Virtual Links Are Used to Connect an Area to the Backbone Area*

The hello protocol works over virtual links as it does over standard links, in 10-second intervals. However, LSA updates work differently on virtual links. An LSA usually refreshes every 30 minutes. However, LSAs learned through a virtual link have the DoNotAge (DNA) option set so that the LSA does not age out. This DNA technique is required to prevent excessive flooding over the virtual link.

Configuring OSPF Virtual Links

Use the following router configuration command to define an OSPF virtual link:

```
area area-id virtual-link router-id [authentication [message-digest| null]]
[hello-interval seconds] [retransmit-interval seconds] [transmit-
delay seconds] [dead-interval seconds] [[authentication-key
key] | [message-digest-key key-id md5 key]]
```

To remove a virtual link, use the **no** form of this command.

Table 3-5 describes the options available with the `area area-id virtual-link` command. Make sure that you understand the effect of these options before changing them. For instance, the smaller the hello interval, the faster the detection of topological changes, but the more routing traffic. You should be conservative with the setting of the retransmit interval, or the result is needless retransmissions. The value should be larger for serial lines and virtual links. The transmit delay value should take into account the interface's transmission and propagation delays.

Table 3-5 `area area-id virtual-link` Command Parameters

Parameter	Description
<i>area-id</i>	Specifies the area ID of the transit area for the virtual link. This ID can be either a decimal value or in dotted-decimal format, like a valid IP address. There is no default. The transit area cannot be a stub area.
<i>router-id</i>	Specifies the router ID of the virtual link neighbor. The router ID appears in the <code>show ip ospf</code> display. This value is in an IP address format. There is no default.
authentication	(Optional) Specifies an authentication type.
message-digest	(Optional) Specifies the use of MD5 authentication.
null	(Optional) Overrides simple password or MD5 authentication if configured for the area. No authentication is used.
hello-interval <i>seconds</i>	(Optional) Specifies the time (in seconds) between the hello packets that the Cisco IOS Software sends on an interface. The unsigned integer value is advertised in the Hello packets. The value must be the same for all routers and access servers attached to a common network. The default is 10 seconds.

Parameter	Description
<code>retransmit-interval seconds</code>	(Optional) Specifies the time (in seconds) between LSA retransmissions for adjacencies belonging to the interface. The value must be greater than the expected round-trip delay between any two routers on the attached network. The default is 5 seconds.
<code>transmit-delay seconds</code>	(Optional) Specifies the estimated time (in seconds) to send an LSU packet on the interface. This integer value must be greater than 0. LSAs in the update packet have their age incremented by this amount before transmission. The default value is 1 second.
<code>dead-interval seconds</code>	(Optional) Specifies the time (in seconds) that must pass without hello packets being seen before a neighboring router declares the router down. This is an unsigned integer value. The default is 4 times the default hello interval, or 40 seconds. As with the hello interval, this value must be the same for all routers and access servers attached to a common network.
<code>authentication-key key</code>	(Optional) Specifies the password used by neighboring routers for simple password authentication. It is any continuous string of up to eight characters. There is no default value.
<code>message-digest-key key-id md5 key</code>	(Optional) Identifies the key ID and key (password) used between this router and neighboring routers for MD5 authentication. There is no default value.

In the example in Figure 3-30, area 0 is discontinuous. A virtual link is used as a backup strategy to temporarily connect area 0. Area 1 is used as the transit area. Router A builds a virtual link to Router B, and Router B builds a virtual link to the Router A. Each router points at the other router's router ID.

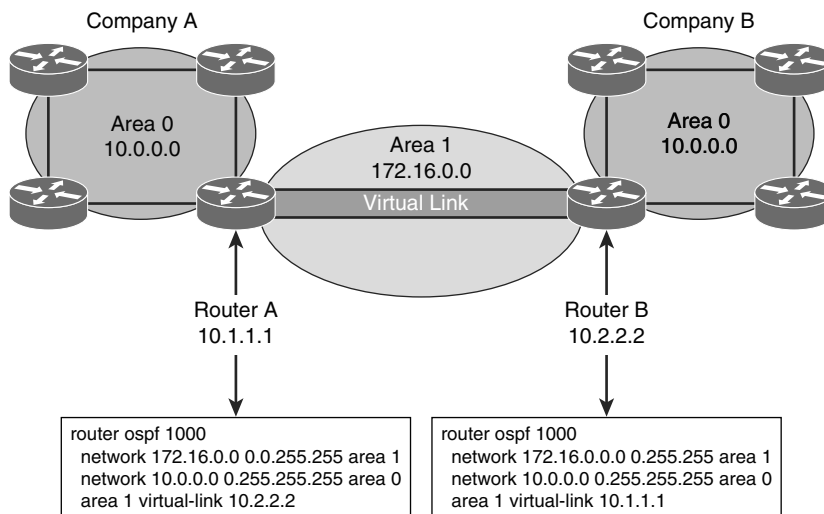


Figure 3-30 OSPF Virtual Link Configuration: Split Area 0

Figure 3-31 presents another example network. The configurations for routers R1 and R3 are provided in Example 3-56.

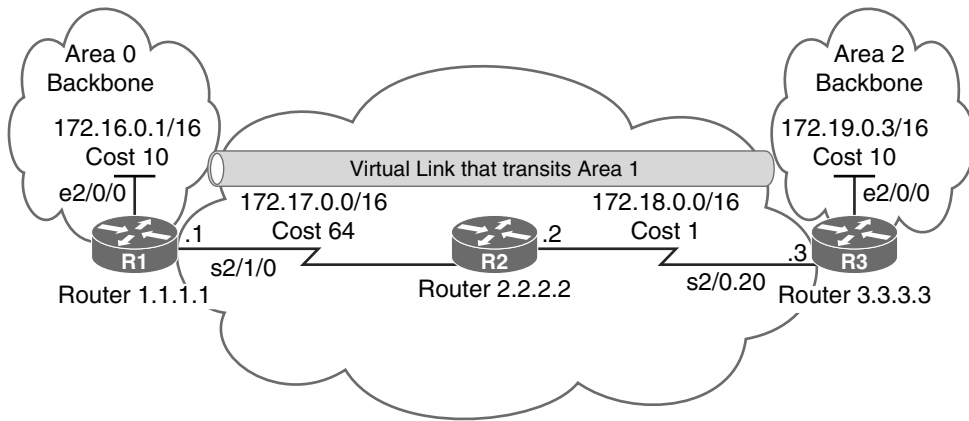


Figure 3-31 OSPF Virtual Link Across Area 1

Example 3-56 Configuring a Virtual Link Between R1 and R3

```
R1(config)# router ospf 2
R1(config-router)# area 1 virtual-link 3.3.3.3

R3(config)# router ospf 2
R3(config-router)# area 1 virtual-link 1.1.1.1
```

Configuring OSPF Stub Areas

In this section, you will learn how to implement special area types in an OSPF environment, using the topology in Figure 3-32. The stub and totally stubby areas are deployed to reduce the size of the OSPF database and routing table:

- **Stub area:** This area type does not accept information about routes external to the autonomous system, such as routes from non-OSPF sources. If routers need to route to networks outside the autonomous system, they use a default route, indicated as 0.0.0.0. Stub areas cannot contain ASBRs (except that the ABRs may also be ASBRs). The stub area does not accept external routes.
- **Totally stubby area:** This Cisco proprietary area type does not accept external autonomous system routes or summary routes from other areas internal to the autonomous system. If a router needs to send a packet to a network external to the area, it sends the packet using a default route. Totally stubby areas cannot contain ASBRs (except that the ABRs may also be ASBRs). A totally stubby area does not accept external or interarea routes.




Figure 3-32 *Topology for Stub and Totally Stubby Areas*

OSPF Stub Areas

Example 3-57 displays the OSPF routes in the routing tables of R2 and R3, including external OSPF routes.

Example 3-57 *OSPF Routes in R2's and R3's Routing Tables*

```
R2# show ip route ospf
<Output omitted>

    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA   172.16.13.0/30 [110/20] via 172.16.12.1, 00:56:16, Ethernet0/0
O IA   172.16.14.0/25 [110/20] via 172.16.12.1, 00:56:16, Ethernet0/0
O IA   192.168.1.0/24 [110/11] via 172.16.12.1, 00:56:16, Ethernet0/0
O IA   192.168.3.0/24 [110/21] via 172.16.12.1, 00:54:50, Ethernet0/0
O IA   192.168.4.0/24 [110/21] via 172.16.12.1, 00:46:00, Ethernet0/0
O E2   198.51.100.0/24 [110/20] via 172.16.12.1, 00:01:47, Ethernet0/0
O E2   203.0.113.0/24 [110/20] via 172.16.12.1, 00:01:47, Ethernet0/0

R3# show ip route ospf
<Output omitted>

    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA   172.16.12.0/30 [110/20] via 172.16.13.1, 00:53:58, Ethernet0/0
O IA   172.16.14.0/25 [110/20] via 172.16.13.1, 00:53:58, Ethernet0/0
```

```

O IA 192.168.1.0/24 [110/11] via 172.16.13.1, 00:53:58, Ethernet0/0
O IA 192.168.2.0/24 [110/21] via 172.16.13.1, 00:53:58, Ethernet0/0
O IA 192.168.4.0/24 [110/21] via 172.16.13.1, 00:45:10, Ethernet0/0
O E2 198.51.100.0/24 [110/20] via 172.16.13.1, 00:00:57, Ethernet0/0
O E2 203.0.113.0/24 [110/20] via 172.16.13.1, 00:00:57, Ethernet0/0

```

The two external routes, 198.51.100.0/24 and 203.0.113.0/24, are being redistributed into the OSPF domain by R4, which acts as the ASBR and provides Internet connectivity.

Area 0 is the backbone area. The backbone area is the central entity to which all other areas connect. All other areas connect to this area to exchange and route information. The OSPF backbone includes all the properties of a standard OSPF area.

Area 1 is a standard nonbackbone area, in which the type 5 LSAs are flooded from R1. This default area accepts link updates, route summaries, and external routes.

Area 2 is also a standard nonbackbone area. The type 5 LSAs are exchanged through the backbone area (R4 and R1) and the standard nonbackbone areas.

A critical design aspect arises in environments with thousands of external routes. The multitude of type 5 LSAs and the corresponding external routes consumes substantial resources. It also makes the network more difficult to monitor and manage.

Example 3-58 shows ABR R1's area 1 configured as a stub area. The stub area offers you a powerful method of reducing the size of the OSPF database and routing tables. This area does not accept information about routes that are external to the AS, such as routes from non-OSPF sources. Stub areas cannot contain ASBRs, except when ABRs are also ASBRs.

Example 3-58 *Configuring R1's Area 1 as a Stub Area*

```

R1(config)# router ospf 1
R1(config-router)# area 1 stub
%OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Ethernet0/1 from FULL to DOWN, Neighbor
Down: Adjacency forced to reset

```

Configuring a stub area reduces the size of the LSDB inside the area, resulting in reduced memory requirements for routers in that area. External network LSAs (type 5), such as those that are redistributed from other routing protocols into OSPF, are not permitted to flood into a stub area.

The **area stub** router configuration mode command is used to define an area as a stub area. Each router in the stub area must be configured with the **area stub** command. The Hello packets that are exchanged between OSPF routers contain a stub area flag that must match on neighboring routers. Until the **area 1 stub** command is enabled on R2 in this scenario, the adjacency between R1 and R2 will be down.

Example 3-59 shows R2's area 1 configured as a stub area. R2 is an internal router or leaf router in R2. Once you configure the area 1 as a stub on R2, the stub area flag in the OSPF Hello packets will start matching between R1 and R2. The routers establish an adjacency and exchange routing information.

Example 3-59 *Configuring R2's Area 1 as a Stub Area*

```
R2(config)# router ospf 1
R2(config-router)# area 1 stub
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from LOADING to FULL, Loading Done
```

Example 3-60 examines the OSPF routing table on R2 and verifies its connectivity to the Internet destinations 203.0.113.2 and 192.0.2.1. Why can you reach 203.0.113.2 and not 192.0.2.1, although both IP addresses exist on the upstream Internet router? _____

Example 3-60 *Verifying R2's Connectivity to the Internet*

```
R2# show ip route ospf
<Output omitted>

O*IA 0.0.0.0/0 [110/11] via 172.16.12.1, 00:19:27, Ethernet0/0
    172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
O IA   172.16.13.0/30 [110/20] via 172.16.12.1, 00:19:27, Ethernet0/0
O IA   172.16.14.0/25 [110/20] via 172.16.12.1, 00:19:27, Ethernet0/0
O IA   192.168.1.0/24 [110/11] via 172.16.12.1, 00:19:27, Ethernet0/0
O IA   192.168.3.0/24 [110/21] via 172.16.12.1, 00:19:27, Ethernet0/0
O IA   192.168.4.0/24 [110/21] via 172.16.12.1, 00:19:27, Ethernet0/0

R2# ping 192.0.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.0.2.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)

R2# ping 203.0.113.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Routing from a stub area to the outside is based on a default route (0.0.0.0). If a packet is addressed to a network that is not in the routing table of an internal router, the router automatically forwards the packet to the ABR (R1), which sends a 0.0.0.0 LSA. Forwarding the packet to the ABR allows routers within the stub to reduce the size of their routing tables, because a single default route replaces many external routes.

The routes that appear in the routing table of R2 include the default route and interarea routes, all designated with an IA in the routing table.

You can reach 203.0.113.2 because the 203.0.113.0/24 is being flooded as a type 5 LSA into the backbone area. The first leg of reachability is provided by the default route

injected into the stub area by the ABR. The second leg, through the backbone area, is ensured by the existing external route.

You cannot reach 192.0.2.1 because its network is not advertised into the OSPF domain as an external route. Despite the default route out of the stub area to the ABR, the ABR drops traffic to that destination because it does not have a path to the destination. This problem could be solved by advertising a default external route from the ASBR (R4) into the OSPF domain.

In Example 3-61, the ASBR (R4) is confirmed to have a default static route configured. The default route is then advertised into the OSPF domain.

Example 3-61 *Propagating a Default Route Using OSPF on R4*

```
R4# show ip route static
<Output omitted>
Gateway of last resort is 198.51.100.2 to network 0.0.0.0
S*    0.0.0.0/0 [1/0] via 198.51.100.2
R4(config)# router ospf 1
R4(config-router)# default-information originate
```

To be able to perform routing from an OSPF autonomous system toward external networks or toward the Internet, you must either know all the destination networks or create a default route. The most scalable and optimized way is through the use of a default route.

To generate a default external route into an OSPF routing domain, use the **default-information originate** router configuration command, as shown in Example 3-61. This command will generate a type 5 LSA for 0.0.0.0/0 when the advertising router already has a default route.

The ABR (R1), shown in Example 3-62, examines the injected default route in the OSPF routing table and database. Connectivity to the external destination 192.0.2.1 is verified with the **show ip ospf database** command.

Example 3-62 *Verifying R1's Default Route*

```
R1# show ip route ospf
<Output omitted>

Gateway of last resort is 172.16.14.2 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 172.16.14.2, 00:00:15, Ethernet0/0
O    192.168.2.0/24 [110/11] via 172.16.12.2, 19:08:02, Ethernet0/1
O    192.168.3.0/24 [110/11] via 172.16.13.2, 19:46:45, Ethernet0/2
O    192.168.4.0/24 [110/11] via 172.16.14.2, 19:46:45, Ethernet0/0
O E2 198.51.100.0/24 [110/20] via 172.16.14.2, 19:46:45, Ethernet0/0
O E2 203.0.113.0/24 [110/20] via 172.16.14.2, 19:46:45, Ethernet0/0
```

```

R1# show ip ospf database
<Output omitted>

                Type-5 AS External Link States

Link ID        ADV Router    Age          Seq#          Checksum Tag
-----
0.0.0.0        4.4.4.4      121         0x80000001  0x00C2DF 1
198.51.100.0  4.4.4.4      1131        0x80000027  0x0054B7 0
203.0.113.0   4.4.4.4      1131        0x80000027  0x00E943 0

R1# ping 192.0.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.0.2.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```

On the ABR, you can see the default route, injected into the backbone area as a type 5 LSA. It appears in the routing table with the symbols O (OSPF), * (default route), E2 (external type 2). You can also see the appropriate LSA 5 in the OSPF database.

Notice the external IP address 192.0.2.1 because the default route directs the traffic via the ASBR. The ASBR has a default static toward the upstream router.

In Example 3-63, connectivity from R2 in the stub area is verified to the external destination 192.0.2.1.

Example 3-63 Verifying R2's Connectivity to an External Destination

```

R2# ping 192.0.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.0.2.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```

Having flooded the default route as a type 5 LSA into the backbone area, you can now verify that R2 can reach the external IP address 192.0.2.1. The traffic to that destination first follows the default route injected into the stub area by the ABR, and then the default route injected into the backbone by the ASBR.

OSPF Totally Stubby Areas

Next, the ABR's (R1's) area 2 is configured as a totally stubby area, shown in Example 3-64.

Example 3-64 Configuring Area 2 as a Totally Stubby Area on the ABR

```

R1(config)# router ospf 1
R1(config-router)# area 2 stub no-summary
%OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Ethernet0/2 from FULL to
DOWN, Neighbor Down: Adjacency forced to reset

```

The totally stubby area is a Cisco proprietary enhancement that further reduces the number of routes in the routing table. A totally stubby area is a stub area that blocks external type 5 LSAs and summary type 3 and type 4 LSAs (interarea routes) from entering the area. Because it blocks these routes, a totally stubby area recognizes only intra-area routes and the default route of 0.0.0.0. ABRs inject the default summary link 0.0.0.0 into the totally stubby area. Each router picks the closest ABR as a gateway to everything outside the area.

Totally stubby areas minimize routing information further than stub areas and increase the stability and scalability of OSPF internetworks. Using totally stubby areas is typically a better solution than using stub areas, as long as the ABR is a Cisco router.

To configure an area as totally stubby, you must configure all the routers inside the area as stub routers. Use the **area stub** command with the **no-summary** keyword on the ABR to configure it as totally stubby. In this example, configuring the total stub on the ABR (R1) breaks the adjacency within area 2 until R3 is configured as a member of a stub area. The adjacency fails because the stub flag in the Hello packets does not match between R1 and R3.

Example 3-65 shows the configuration of an internal router or leaf router (R3) as a stub router in a totally stubby area.

Example 3-65 OSPF Routes in R1's Routing Table

```
R3 (config)# router ospf 1
R3 (config-router)# area 2 stub
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from LOADING to FULL, Loading Done
```

Once R3 in area 1 is configured as a stub, the stub area flag in the OSPF Hello packets will start matching between R1 and R3. The routers establish an adjacency and exchange routing information. R3 may or may not be configured with the **no-summary** keyword. The **no-summary** keyword has no effect when the router is not an ABR and thus does not advertise any interarea summaries.

Example 3-66 verifies R3's routing table and LSDB information in the totally stubby area.

Example 3-66 OSPF Routes in R1's Routing Table

```
R3# show ip route ospf
<Output omitted>
Gateway of last resort is 172.16.13.1 to network 0.0.0.0
O*IA 0.0.0.0/0 [110/11] via 172.16.13.1, 00:18:08, Ethernet0/0

R3# show ip ospf data
<Output omitted>
```

Summary Net Link States (Area 2)				
Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	1.1.1.1	1285	0x80000001	0x0093A6

```

R3# ping 192.0.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.0.2.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```

The leaf router (R3) in the totally stubby area has the smallest possible routing table. Only the intra-area routes are maintained. Interarea and external routes are not visible in the routing tables for each stub area, but are accessible via the intra-area default routes for that stub area. The ABR (R1) blocks interarea and external LSAs and inserts the default route instead.

Despite the minimal routing information about external reachability the leaf router can ping the outside address 192.0.2.1. The traffic to that destination first follows the default route injected into the totally stubby area by the ABR, and then the default route injected into the backbone by the ASBR (R4).

Cost of the Default Route in a Stub Area

By default, the ABR of a stub area will advertise a default route with a cost of 1. You can change the cost of the default route by using the **area default-cost** command. The *default-cost* option provides the metric for the summary default route that is generated by the ABR into the stub area.

To specify a cost for the default summary route sent into a stub or not so stubby area (NSSA), use the following command in router configuration mode:

```
area area-id default-cost cost
```

To remove the assigned default route cost, use the **no** form of this command. Table 3-6 shows the parameters available for this command.

Table 3-6 *Parameters for the area default-cost Command*

Parameter	Description
<i>area-id</i>	Identifier for the stub or NSSA. The identifier can be specified as either a decimal value or as an IP address.
<i>cost</i>	Cost for the default summary route used for a stub or NSSA. The acceptable value is a 24-bit number.

The `area default-cost` command is used only on an ABR attached to a stub or not-so-stubby area (NSSA). Use the `default-cost` option only on an ABR attached to the stub area. The `default-cost` option provides the metric for the summary default route generated by the ABR into the stub area.

The option of tuning the cost of the default route in the stub area is useful in stub areas with redundant exit points to the backbone area, as shown in Figure 3-33. The primary exit point can be configured using a lower cost. The secondary exit point would advertise a higher cost and thus attract external traffic only when the primary ABR fails. This distribution pattern applies only to external traffic. The traffic to interarea networks will follow the shortest path.

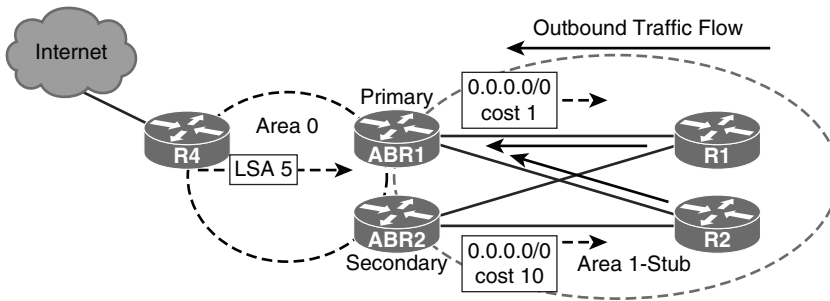


Figure 3-33 Cost of the Default Route in a Stub Area

The default-information originate Command

To generate a default external route into an OSPF routing domain, use the following command in router configuration mode:

```
default-information originate [always] [metric metric-value] [metric-type type-value] [route-map map-name]
```

To disable this feature, use the `no` form of this command. Table 3-7 shows the parameters available for this command.

Table 3-7 Parameters for the default-information originate Command

Parameter	Description
<code>always</code>	(Optional) Always advertises the default route regardless of whether the software has a default route.
<code>metric metric-value</code>	(Optional) Metric used for generating the default route. If you omit a value and do not specify a value using the <code>default-metric</code> router configuration command, the default metric value is 1. The value used is specific to the protocol.

Parameter	Description
metric-type <i>type-value</i>	(Optional) External link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: 1: Type 1 external route 2: Type 2 external route The default is type 2 external route.
route-map <i>map-name</i>	(Optional) Routing process will generate the default route if the route map is satisfied.

There are two ways to advertise a default route into a standard area. You can advertise 0.0.0.0/0 into the OSPF domain when the advertising router already has a default route. Use the **default-information originate** command to allow the ASBR to originate a type 5 default route inside the OSPF autonomous system. The default route must be in the routing table otherwise it will not be propagated by OSPF.

You can use different keywords in the configuration command to configure dependency on IP routing table entries. To advertise 0.0.0.0/0 regardless of whether the advertising router already has a default route, add the keyword **always** to the **default-information originate** command. The default route will be propagated by OSPF whether or not there is a default route.

Whenever you use the **redistribute** or the **default-information** command to redistribute routes into an OSPF routing domain, the router automatically becomes an ASBR. You can also use a route map to define dependency on any condition inside the route map. The **metric** and **metric-type** options allow you to specify the OSPF cost and metric type of the injected external route.

Other Stubby Area Types

The NSSA is a nonproprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

Redistribution into an NSSA creates a special type of LSA known as a type 7 LSA, which can exist only in an NSSA. An NSSA ASBR (router ASBR1 in the Figure 3-34) generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain. Type 7 LSAs have a propagate (P) bit in the LSA header to prevent propagation loops between the NSSA and the backbone area. The NSSA retains the majority of other stub area features. An important difference is the default behavior regarding the default route. ABR must be configured with additional commands before it starts announcing it into the NSSA area.




Figure 3-34 NSSA Area

The type 7 LSA is described in the routing table as an O N2 or O N1 (N means NSSA). N1 means that the metric is calculated like external type 1 (E1); N2 means that the metric is calculated like external type 2 (E2). The default is O N2.

Note The E1 metric adds external and internal costs together to reflect the whole cost to the destination. The E2 metric takes only the external cost, which is reflected in the OSPF cost.

The totally NSSA feature is an extension to the NSSA feature like the totally stubby feature is an extension to the stub area feature. It is a Cisco proprietary feature that blocks type 3, 4, and 5 LSAs. A single default route replaces both inbound-external (type 5) LSAs and summary (type 3 and 4) LSAs in the totally NSSA area. The ABRs for the totally NSSA area must be configured to prevent the flooding of summary routes for other areas into the NSSA area. Only ABRs control the propagation of type 3 LSAs from the backbone. If an ABR is configured on any other routers in the area, it will have no effect at all.

To configure an area as an NSSA, you must configure all routers inside the area for NSSA functionality. The `area nssa` router configuration mode command is used to define each router in the NSSA area as not-so-stubby. Totally NSSA functionality requires one more step; you must configure each ABR for totally NSSA functionality. The `area nssa` command with the `no-summary` keyword is used to define the ABR as totally not-so-stubby.

OSPFv3

OSPF is a widely used IGP in IPv4, IPv6, and dual-stack (IPv4/IPv6) environments. The OSPF upgrade to support IPv6 generated a number of significant changes to how the protocol behaves. Understanding the differences between OSPFv2 and OSPFv3 is required for the successful deployment and operation of an IPv6 network using OSPF for routing. This section describes OSPFv3, the IPv6-capable version of the OSPF routing protocol, including its operations, configuration, and commands.

Upon completing this section, you will be able to do the following:

- Implement OSPFv3 in a dual-stack (IPv4/IPv6) environment
- Configure external route summarization and load balancing in OSPFv3
- Explain the limitations and where you need to be careful when configuring OSPFv3

Configuring OSPFv3

In this section, you will learn how to implement OSPFv3 in a dual-stack (IPv4/IPv6) environment. Using the IPv6 topology in Figure 3-35 for IPv6 and Figure 3-36 for IPv4, routers R2, R3, and R4 have been completely preconfigured. R1 has been preconfigured with the necessary IPv4/IPv6 addresses, but does not have any routing protocol configuration. On R1, we will first configure OSPFv3 for IPv6 in the traditional way, in which a dedicated OSPF process serves the IPv6 protocol. Then we will migrate the configuration to the newest configuration approach, in which a single OSPFv3 process serves both address families, IPv4 and IPv6.

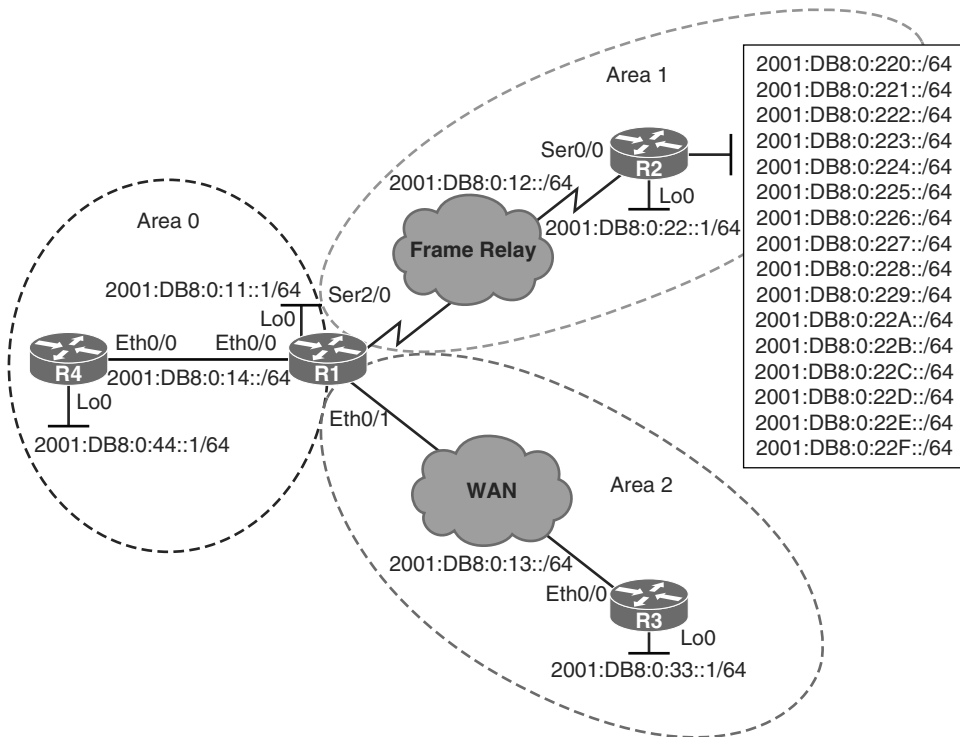


Figure 3-35 IPv6 Topology OSPFv3




Figure 3-36 IPv4 Topology OSPFv3

Implementing OSPFv3

Example 3-67 shows R1 enabled for IPv6 unicast routing and starting an IPv6 OSPF router process with ID 1. R1 is configured with a router ID 1.1.1.1 and loopback 0 is as a passive interface.

Example 3-67 Enabling OSPFv3 on R1

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 router ospf 1
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# passive-interface Loopback0
```

OSPFv3 is the IPv6-capable version of the OSPF routing protocol. It is a rewrite of the OSPF protocol to support IPv6, although the foundation remains the same as in IPv4 and OSPFv2. The OSPFv3 metric is still based on interface cost. The packet types and neighbor discovery mechanisms are the same in OSPFv3 as they are for OSPFv2. OSPFv3 also supports the same interface types, including broadcast, point-to-point, point-to-multipoint, NBMA, and virtual links. LSAs are still flooded throughout an OSPF domain, and many of the LSA types are the same, though a few have been renamed or newly created.

Cisco IOS routers offer two OSPF configuration methods for IPv6:

- Using the traditional **ipv6 router ospf** global configuration command
- Using the new-style **router ospfv3** global configuration command

We will first examine the traditional configuration approach, and then migrate the configuration to the new style.

To start any IPv6 routing protocols, you need to enable IPv6 unicast routing using the **ipv6 unicast-routing** command. In the traditional configuration approach, the OSPFv3 and OSPFv2 processes run independently on a router. In the traditional way, the OSPF process for IPv6 is started using the **ipv6 router ospf** command.

The OSPF process for IPv6 does not require an IPv4 address to be configured on the router, but it does require a 32-bit value for the router ID, which uses IPv4 address notation. The router ID is defined using the **router-id** command. If the router ID is not specifically configured, the system will try to dynamically choose an ID from the currently active IPv4 addresses, using the same process as OSPFv2 does for IPv4. If there is no active IPv4 address, the process will fail to start.

In the **ipv6 router ospf** configuration mode, you can specify the passive interfaces (using the **passive-interface** command), enable summarization, and fine-tune the operation, but you do not enable the process on specific interfaces. There is no **network** command. To activate the OSPF process on required interfaces, you will need the **ipv6 ospf** command in the interface configuration mode.

In Example 3-68, R1 is enabled for the OSPF-for-IPv6 process on its active interfaces. Interface Loopback 0 and Ethernet 0/0 are assigned to area 0, Serial 2/0 to area 1, and Ethernet 0/1 to area 2. The **exit** interface command does not need to be used between interfaces. It is only used in this example to better illustrate that OSPF-for-IPv6 is enabled on each specific interface.

Example 3-68 *Enabling OSPFv3 on the Interface*

```
R1(config)# interface Loopback0
R1(config-if)# ipv6 ospf 1 area 0
R1(config-if)# exit
R1(config)# interface Ethernet0/0
R1(config-if)# ipv6 ospf 1 area 0
R1(config-if)# exit
R1(config)# interface Serial2/0
R1(config-if)# ipv6 ospf 1 area 1
R1(config-if)# exit
R1(config)# interface Ethernet0/1
R1(config-if)# ipv6 ospf 1 area 2
%OSPFv3-5-ADJCHG: Process 1, Nbr 4.4.4.4 on Ethernet0/0 from LOADING to FULL,
Loading Done
%OSPFv3-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Ethernet0/1 from LOADING to FULL,
Loading Done
```

To enable the OSPF-for-IPv6 process on an interface and assign that interface to an area, use the `ipv6 ospf ospf-process area area-id` command in the interface configuration mode. To be able to enable OSPFv3 on an interface, the interface must be enabled for IPv6. This occurs when the interface is configured with a global unicast IPv6 address.

Example 3-69 examines R1's OSPF adjacencies and routing table.

Example 3-69 R1's Adjacencies and Routing Table

```
R1# show ipv6 ospf neighbor

                OSPFv3 Router with ID (1.1.1.1) (Process ID 1)

Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
4.4.4.4          1    FULL/DR         00:00:37   3             Ethernet0/0
3.3.3.3          1    FULL/DR         00:00:35   4             Ethernet0/1
R1# show ipv6 route ospf
<Output omitted>
O   2001:DB8:0:33::/64 [110/11]
    via FE80::A8BB:CCFF:FE00:AD10, Ethernet0/1
O   2001:DB8:0:44::/64 [110/11]
    via FE80::A8BB:CCFF:FE00:AE00, Ethernet0/0
```

After enabling the OSPF process on IPv6 interfaces you can verify the adjacencies and the IPv6 routing table. You can selectively display the OSPF-populated part of the routing table if you use the `show ipv6 route` command with the `ospf` keyword.

Why is the OSPF adjacency with R2, via Serial2/0, not working? On NBMA interfaces, the NBMA network type is by default used in OSPF routing. On such links, at least one side needs to define the OSPF neighbor, similarly as in OSPFv2. The `neighbor` command in the IPv6 environment requires that the IPv6 link-local address is specified for the peer, instead of using an IPv6 global unicast address. The IPv6 link-local addresses start with the FE80 prefix. In this scenario, the link-local address of R2 is FE80::2.

Example 3-70 specifies the IPv6 neighbor, FE80::2, for OSPFv3 on the NBMA interface Serial 2/0.

Example 3-70 Specifying the Neighbor on an NBMA Interface

```
R1(config)# interface serial 2/0
R1(config-if)# ipv6 ospf neighbor FE80::2
%OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial2/0 from LOADING to FULL, Loading Done
```

OSPF adjacencies over NBMA links require that IPv6 connectivity for both the link-local and the global addresses is established. Depending on the transport network, this may require mapping of IPv6 addresses to Layer 2 circuit identifiers. In this scenario, R1 and R2 have been pre-configured with the necessary mappings. The relevant configuration on R1, including the neighbor address, is shown in Example 3-71:

Example 3-71 *R1's Partial Running-Config*

```

R1# show running-config interface serial 2/0
Building configuration...

Current configuration : 404 bytes
!
interface Serial2/0
 ip address 172.16.12.1 255.255.255.252
 encapsulation frame-relay
 ipv6 address FE80::1 link-local
 ipv6 address 2001:DB8:0:12::1/64
 ipv6 ospf 1 area 1
 ipv6 ospf neighbor FE80::2
 serial restart-delay 0
 frame-relay map ip 172.16.12.2 102 broadcast
 frame-relay map ipv6 2001:DB8:0:12::2 102 broadcast
 frame-relay map ipv6 FE80::2 102 broadcast
 no frame-relay inverse-arp
end

```

Example 3-72 examines the IPv6 OSPF database on R1.

Example 3-72 *R1's OSPF LSDB*

```

R1# show ipv6 ospf database

      OSPFv3 Router with ID (1.1.1.1) (Process ID 1)

      Router Link States (Area 0)

ADV Router    Age         Seq#         Fragment ID  Link count  Bits
1.1.1.1       854        0x80000003  0            1           B
4.4.4.4       871        0x80000002  0            1           None

      Net Link States (Area 0)

ADV Router    Age         Seq#         Link ID      Rtr count
4.4.4.4       871        0x80000001  3            2

      Inter Area Prefix Link States (Area 0)

ADV Router    Age         Seq#         Prefix
1.1.1.1       845        0x80000001  2001:DB8:0:12::/64
1.1.1.1       845        0x80000001  2001:DB8:0:13::/64
1.1.1.1       845        0x80000001  2001:DB8:0:33::/64

```

<Output omitted>

Link (Type-8) Link States (Area 0)					
ADV Router	Age	Seq#	Link ID	Interface	
1.1.1.1	870	0x80000001	3	Et0/0	
4.4.4.4	1056	0x80000002	3	Et0/0	

Intra Area Prefix Link States (Area 0)					
ADV Router	Age	Seq#	Link ID	Ref-lstyp	Ref-LSID
1.1.1.1	865	0x80000003	0	0x2001	0
4.4.4.4	871	0x80000003	0	0x2001	0
4.4.4.4	871	0x80000001	3072	0x2002	3

OSPFv3 (for IPv6) renames two LSA types and defines two additional LSA types that do not exist in OSPFv2 (for IPv4).

The two renamed LSA types are as follows:

- **Interarea prefix LSAs for ABRs (Type 3):** Type 3 LSAs advertise internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks summarized into one advertisement. Only ABRs generate summary LSAs. In OSPF for IPv6, addresses for these LSAs are expressed as prefix/prefix length instead of address and mask. The default route is expressed as a prefix with length 0.
- **Interarea router LSAs for ASBRs (Type 4):** Type 4 LSAs advertise the location of an ASBR. Routers that are trying to reach an external network use these advertisements to determine the best path to the next hop. ASBRs generate Type 4 LSAs.

The two new LSA types are as follows:

- **Link LSAs (Type 8):** Type 8 LSAs have local-link flooding scope and are never flooded beyond the link with which they are associated. Link LSAs provide the link-local address of the router to all other routers attached to the link. They inform other routers attached to the link of a list of IPv6 prefixes to associate with the link. In addition, they allow the router to assert a collection of option bits to associate with the network LSA that will be originated for the link.
- **Intra-area prefix LSAs (Type 9):** A router can originate multiple intra-area prefix LSAs for each router or transit network, each with a unique link-state ID. The link-state ID for each intra-area prefix LSA describes its association to either the router LSA or the network LSA. The link-state ID also contains prefixes for stub and transit networks.

Example 3-73 reexamines the OSPFv3 adjacencies and routing table on R1.

Example 3-73 *R1's OSPFv3 Adjacencies and Routing Table*

```

R1# show ipv6 ospf neighbor

                OSPFv3 Router with ID (1.1.1.1) (Process ID 1)

Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
4.4.4.4          1    FULL/DR         00:00:39   3             Ethernet0/0
2.2.2.2          1    FULL/DR         00:01:43   3             Serial2/0
3.3.3.3          1    FULL/DR         00:00:39   4             Ethernet0/1
R1# show ipv6 route ospf
<Output omitted>
O   2001:DB8:0:22::/64 [110/65]
    via FE80::2, Serial2/0
O   2001:DB8:0:33::/64 [110/11]
    via FE80::A8BB:CCFF:FE00:AD10, Ethernet0/1
O   2001:DB8:0:44::/64 [110/11]
    via FE80::A8BB:CCFF:FE00:AE00, Ethernet0/0
O   2001:DB8:0:220::/64 [110/65]
    via FE80::2, Serial2/0
O   2001:DB8:0:221::/64 [110/65]
    via FE80::2, Serial2/0
<Output omitted>

```

After enabling the OSPF on the NBMA interface, notice an additional adjacency across Serial 2/0 and multiple OSPF intra-area routes received via this interface.

OSPFv3 for IPv4 and IPv6

OSPFv3 does not only support exchange of IPv6 routes, but it also supports exchange of IPv4 routes.

The newest OSPFv3 configuration approach utilizes a single OSPFv3 process. It is capable of supporting IPv4 and IPv6 within a single OSPFv3 process. OSPFv3 builds a single database with LSAs that carry IPv4 and IPv6 information. The OSPF adjacencies are established separately for each address family. Settings that are specific to an address family (IPv4/IPv6) are configured inside that address family router configuration mode.

Running single OSPFv3 for both IPv4 and IPv6 is supported since Cisco IOS Software Release 15.1(3)S.

Example 3-74 shows R1's configuration of an OSPFv3 process using the new configuration style (**router ospfv3**), using process number 1, OSPF router ID 1.1.1.1, and making the Loopback 0 interface passive.

Example 3-74 *Configuring OSPFv3 Using the router ospfv3 Command*

```
R1(config)# router ospfv3 1
R1(config-router)# router-id 1.1.1.1
R1(config-router)# passive-interface Loopback0
```

The new-style OSPFv3 process is enabled using the **router ospfv3 process-number** command. Within the OSPF process configuration mode, the OSPF process ID is defined (using the **router-id ospf-process-ID** command), the passive interfaces are set, and per-process OSPF behavior can be tuned.

Example 3-75 displays the OSPFv3 router configuration on R1 using the **show running-config | section router** command. The old-style OSPF router configuration (**ipv6 router ospf**) has disappeared and has been replaced by the new-style **router ospfv3** with an address family sub-mode.

Example 3-75 *R1's OSPFv3 Configuration*

```
R1# show running-config | section router
router ospfv3 1
  router-id 1.1.1.1
  !
  address-family ipv6 unicast
    passive-interface Loopback0
    router-id 1.1.1.1
  exit-address-family
```

The router ID is displayed in the router configuration mode that is valid globally for all address families.

The **address-family ipv6 unicast** has been automatically created on R1. Cisco IOS Software has parsed the previous old-style OSPFv3 configuration and found that the OSPF process was enabled only for IPv6. Consequently, when you chose the new-style configuration, the IPv6 address family has been instantiated and the IPv4 address family does not show in the configuration.

The **passive-interface** configuration is actually a setting that is valid per address family. You can have dissimilar settings for IPv4 and IPv4. Therefore this command has been placed in the address family submode.

Example 3-76 verifies R1's OSPFv3 operation by verifying its adjacencies, routing table, and database. The OSPFv3 operation can be verified using the old-style commands (**show ipv6 ospf neighbor**, **show ipv6 ospf database**) or the new-style commands, such as **show ospfv3 neighbor** and **show ospfv3 database**.

Example 3-76 *R1's OSPFv3 Adjacencies, Routing Table, and LSDB*

```

R1# show ospfv3 neighbor

          OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)

Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
4.4.4.4          1    FULL/DR         00:00:37   3             Ethernet0/0
2.2.2.2          1    FULL/DR         00:01:44   3             Serial2/0
3.3.3.3          1    FULL/DR         00:00:35   4             Ethernet0/1

R1# show ipv6 route ospf
IPv6 Routing Table - default - 28 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, I - LISP
O 2001:DB8:0:22::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:33::/64 [110/11]
   via FE80::A8BB:CCFF:FE00:AD10, Ethernet0/1
O 2001:DB8:0:44::/64 [110/11]
   via FE80::A8BB:CCFF:FE00:AE00, Ethernet0/0
O 2001:DB8:0:220::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:221::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:222::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:223::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:224::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:225::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:226::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:227::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:228::/64 [110/65]
   via FE80::2, Serial2/0
O 2001:DB8:0:229::/64 [110/65]
   via FE80::2, Serial2/0

```

```

O 2001:DB8:0:22A::/64 [110/65]
  via FE80::2, Serial2/0
O 2001:DB8:0:22B::/64 [110/65]
  via FE80::2, Serial2/0
O 2001:DB8:0:22C::/64 [110/65]
  via FE80::2, Serial2/0
O 2001:DB8:0:22D::/64 [110/65]
  via FE80::2, Serial2/0
O 2001:DB8:0:22E::/64 [110/65]
  via FE80::2, Serial2/0
O 2001:DB8:0:22F::/64 [110/65]
  via FE80::2, Serial2/0

```

R1# show ospfv3 database

OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)

Router Link States (Area 0)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
1.1.1.1	793	0x80000006	0	1	B
4.4.4.4	135	0x8000000D	0	1	None

Net Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Rtr count
4.4.4.4	379	0x80000006	3	2

Inter Area Prefix Link States (Area 0)

ADV Router	Age	Seq#	Prefix
1.1.1.1	301	0x80000006	2001:DB8:0:12::/64
1.1.1.1	301	0x80000006	2001:DB8:0:33::/64
1.1.1.1	301	0x80000006	2001:DB8:0:13::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:220::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:221::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:222::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:223::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:224::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:225::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:226::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:227::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:228::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:229::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22A::/64

1.1.1.1	1301	0x80000004	2001:DB8:0:22B::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22C::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22D::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22E::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22F::/64

Link (Type-8) Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Interface
1.1.1.1	793	0x80000006	3	Et0/0
4.4.4.4	135	0x8000000B	3	Et0/0

Intra Area Prefix Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Ref-lstyp	Ref-LSID
1.1.1.1	793	0x80000006	0	0x2001	0
4.4.4.4	379	0x8000000F	0	0x2001	0
4.4.4.4	379	0x80000006	3072	0x2002	3

Router Link States (Area 1)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
1.1.1.1	793	0x80000007	0	1	B
2.2.2.2	1464	0x80000029	0	1	None

Net Link States (Area 1)

ADV Router	Age	Seq#	Link ID	Rtr count
2.2.2.2	1464	0x80000004	3	2

Inter Area Prefix Link States (Area 1)

ADV Router	Age	Seq#	Prefix
1.1.1.1	301	0x80000006	2001:DB8:0:33::/64
1.1.1.1	301	0x80000006	2001:DB8:0:13::/64
1.1.1.1	301	0x80000006	2001:DB8:0:11::1/128
1.1.1.1	301	0x80000006	2001:DB8:0:44::/64
1.1.1.1	301	0x80000006	2001:DB8:0:14::/64

Link (Type-8) Link States (Area 1)

ADV Router	Age	Seq#	Link ID	Interface
1.1.1.1	793	0x80000006	11	Se2/0
2.2.2.2	1962	0x80000029	3	Se2/0

Intra Area Prefix Link States (Area 1)

ADV Router	Age	Seq#	Link ID	Ref-lstyp	Ref-LSID
2.2.2.2	1464	0x80000040	0	0x2001	0
2.2.2.2	1464	0x80000004	3072	0x2002	3

Router Link States (Area 2)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
1.1.1.1	793	0x80000006	0	1	B
3.3.3.3	1901	0x8000002B	0	1	None

Net Link States (Area 2)

ADV Router	Age	Seq#	Link ID	Rtr count
3.3.3.3	376	0x80000006	4	2

Inter Area Prefix Link States (Area 2)

ADV Router	Age	Seq#	Prefix
1.1.1.1	301	0x80000006	2001:DB8:0:12::/64
1.1.1.1	301	0x80000006	2001:DB8:0:11::1/128
1.1.1.1	301	0x80000006	2001:DB8:0:44::/64
1.1.1.1	301	0x80000006	2001:DB8:0:14::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:220::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:221::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:222::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:223::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:224::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:225::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:226::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:227::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:228::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:229::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22A::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22B::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22C::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22D::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22E::/64
1.1.1.1	1301	0x80000004	2001:DB8:0:22F::/64

Link (Type-8) Link States (Area 2)

ADV Router	Age	Seq#	Link ID	Interface
------------	-----	------	---------	-----------

1.1.1.1	793	0x80000006	4	Et0/1	
3.3.3.3	1901	0x80000028	4	Et0/1	
Intra Area Prefix Link States (Area 2)					
ADV Router	Age	Seq#	Link ID	Ref-lstyp	Ref-LSID
3.3.3.3	376	0x8000002B	0	0x2001	0
3.3.3.3	376	0x80000006	4096	0x2002	4

Despite the change of the OSPFv3 configuration to the new-style approach, the OSPF connectivity has been retained. In fact, R1 now uses a mixed configuration: new-style process configuration, and old-style interface commands. Example 3-77 shows the old-style interface commands.

Example 3-77 *OSPFv3 Old-Style OSPF Configuration Commands*

```
interface Loopback0
  ipv6 ospf 1 area 0
!
interface Ethernet0/0
  ipv6 ospf 1 area 0
!
interface Ethernet0/1
  ipv6 ospf 1 area 2
!
interface Serial2/0
  ipv6 ospf 1 area 1
  ipv6 ospf neighbor FE80::2
```

In Example 3-78, R1 is enabled using the OSPFv3 IPv6 address family on the active interfaces using the new-style configuration approach. Once again, the `exit` interface command is not needed but used to make the configuration clearer.

Example 3-78 *OSPFv3 New-Style OSPF Configuration Commands*

```
R1(config)# interface Loopback 0
R1(config-if)# ospfv3 1 ipv6 area 0
R1(config-if)# exit
R1(config)# interface Ethernet 0/0
R1(config-if)# ospfv3 1 ipv6 area 0
R1(config-if)# exit
R1(config)# interface Serial 2/0
R1(config-if)# ospfv3 1 ipv6 area 1
R1(config-if)# exit
R1(config)# interface Ethernet 0/1
R1(config-if)# ospfv3 1 ipv6 area 2
```

The preferred interface mode command for the new style OSPFv3 configuration is the **ospfv3 process-id (ipv4|ipv6) area area-id** command. It allows you to selectively activate the OSPFv3 process for an address family (IPv4 or IPv6) on a given interface.

With the OSPFv3 address families feature, you may have two device processes per interface, but only one process per AF. If the IPv4 AF is used, an IPv4 address must first be configured on the interface. For IPv6 AF it is enough, if only IPv6 is enabled on the interface, as OSPFv3 uses link-local addresses. A single IPv4 or IPv6 OSPFv3 process running multiple instances on the same interface is not supported.

Example 3-79 verifies the resulting configuration and operation on R1. The interface configuration can be viewed using the **show running-config interface** command. The **include** keyword can be used to display only the interface commands that include a certain information.

Example 3-79 Verifying OSPFv3 Configuration and Operation on R1

```
R1# show running-config interface Loopback 0 | include ospf
ospfv3 1 ipv6 area 0
R1# show running-config interface Ethernet 0/0 | include ospf
ospfv3 1 ipv6 area 0
R1# show running-config interface Serial 2/0 | include ospf
ospfv3 1 ipv6 area 1
ospfv3 1 ipv6 neighbor FE80::2
R1# show running-config interface Ethernet 0/1 | include ospf
ospfv3 1 ipv6 area 2
R1# show ospfv3 neighbor
      OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)
Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
4.4.4.4          1    FULL/DR         00:00:32   3             Ethernet0/0
2.2.2.2          1    FULL/DR         00:01:48   3             Serial2/0
3.3.3.3          1    FULL/DR         00:00:31   4             Ethernet0/1
```

The configuration on R1 shows that the new-style interface-mode commands have replaced the old-style (**ipv6 ospf**) commands. The configuration of the NBMA interface (Serial 2/0) shows that the neighbor commands has been automatically updated with the **ospfv3 process-id ipv6 neighbor** command.

The OSPF operation has not been affected. The OSPFv3 adjacencies, database, and routing table are functional.

Example 3-80 shows R1 enabled the OSPFv3 process enabled for IPv4 and the Loopback 0 interface is configured as passive. To activate OSPFv3 for IPv4, you need to configure the **ospfv3 process-number ipv4 area area-id** command in the configuration mode of the desired interface.

Example 3-80 *Enabling OSPFv3 for IPv4*

```

R1(config)# interface Loopback0
R1(config-if)# ospfv3 1 ipv4 area 0
R1(config-if)# exit
R1(config)# interface Ethernet0/0
R1(config-if)# ospfv3 1 ipv4 area 0
R1(config-if)# exit
R1(config)# interface Ethernet0/1
R1(config-if)# ospfv3 1 ipv4 area 2
R1(config)# exit
R1(config-if)# interface Serial2/0
R1(config-if)# ospfv3 1 ipv4 area 1
R1(config-if)# ospfv3 1 ipv4 neighbor FE80::2
R1(config-if)# exit
R1(config)# router ospfv3 1
R1(config-router)# address-family ipv4 unicast
R1(config-router-af)# passive-interface Loopback0
%OSPFv3-5-ADJCHG: Process 1, IPv4, Nbr 0.0.0.0 on Serial2/0 from ATTEMPT to DOWN,
Neighbor Down: Interface down or detached
%OSPFv3-5-ADJCHG: Process 1, IPv4, Nbr 3.3.3.3 on Ethernet0/1 from LOADING to FULL,
Loading Done
%OSPFv3-5-ADJCHG: Process 1, IPv4, Nbr 4.4.4.4 on Ethernet0/0 from LOADING to FULL,
Loading Done

```

This way some (or all) of the links can be enabled for IPv4 forwarding and be configured with IPv4 addresses. For example, pockets of IPv4-only devices may exist around the edges running an IPv4 static or dynamic routing protocol. In that scenario, you could forward IPv4 or IPv6 traffic between these pockets. The transit device needs both IPv4 and IPv6 forwarding stacks (that is, a dual stack).

This feature allows a separate (possibly incongruent) topology to be constructed for the IPv4 address family. It installs IPv4 routes in the IPv4 RIB, and then the forwarding occurs natively. The OSPFv3 process fully supports an IPv4 AF topology and can redistribute routes from and into any other IPv4 routing protocol.

An OSPFv3 process can be configured to be IPv4 or IPv6. The **address-family** command is used to determine which AF will run in the OSPFv3 process. Once the address family is selected, you can enable multiple instances on a link and enable address-family-specific commands.

On the NBMA links, such as the interface Serial 2/0 in this scenario, you need to define the OSPF neighbor. In the new-style OSPFv3 you must configure the IPv6 link-local address of the peer as the OSPF neighbor. Both address families use IPv6 as the underlying transport.

Example 3-81 examines R1's OSPFv3 adjacencies. The OSPF adjacencies can be displayed using the **show ospfv3 neighbor** command for both address families.

Example 3-81 *R1's OSPFv3 Adjacencies for Both IPv4 and IPv6 Address Families*

```
R1# show ospfv3 neighbor

      OSPFv3 1 address-family ipv4 (router-id 1.1.1.1)

Neighbor ID    Pri   State           Dead Time   Interface ID  Interface
4.4.4.4        1    FULL/DR         00:00:34   3             Ethernet0/0
2.2.2.2        1    FULL/DR         00:01:38   3             Serial2/0
3.3.3.3        1    FULL/DR         00:00:36   4             Ethernet0/1

      OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)

Neighbor ID    Pri   State           Dead Time   Interface ID  Interface
4.4.4.4        1    FULL/DR         00:00:35   3             Ethernet0/0
2.2.2.2        1    FULL/DR         00:01:58   3             Serial2/0
3.3.3.3        1    FULL/DR         00:00:34   4             Ethernet0/1
```

In Example 3-82, R1's IPv4 routing table is displayed computed from the OSPFv3 database. The IPv4 routing table, computed from the OSPFv3 database, can be displayed using the `show ip route ospfv3` command. The `ospfv3` keyword filters the content of the routing table and displays only the OSPFv3 routes.

Note that command `show ip route ospf` will not show any routes.

Example 3-82 *R1's IPv4 Routing Table with OSPFv3 Routes*

```
R1# show ip route ospfv3
<Output omitted>

      192.168.2.0/32 is subnetted, 1 subnets
O       192.168.2.2 [110/64] via 172.16.12.2, 00:27:49, Serial2/0
      192.168.3.0/32 is subnetted, 1 subnets
O       192.168.3.3 [110/10] via 172.16.13.2, 00:30:08, Ethernet0/1
      192.168.4.0/32 is subnetted, 1 subnets
O       192.168.4.4 [110/10] via 172.16.14.4, 00:30:08, Ethernet0/0
```

The OSPFv3 database for R1 is examined in Example 3-83. A router maintains a single OSPFv3 database, which contains various LSAs. Some LSAs carry IPv4-related information, others carry IPv6-related information, and others carry mixed information. You have to examine specific LSA types to see which address family is described by a given LSA.

Note that old-style verification commands like `show ip ospf database` will not show any information.

Example 3-83 *R1's OSPFv3 LSDB*

```

R1# show ospfv3 database inter-area prefix
      OSPFv3 1 address-family ipv4 (router-id 1.1.1.1)
          Inter Area Prefix Link States (Area 0)
LS Type: Inter Area Prefix Links
Advertising Router: 1.1.1.1
<Output omitted>
Prefix Address: 172.16.12.0
Prefix Length: 30, Options: None
<Output omitted>

      OSPFv3 1 address-family ipv6 (router-id 1.1.1.1)
          Inter Area Prefix Link States (Area 0)
LS Type: Inter Area Prefix Links
Advertising Router: 1.1.1.1
<Output omitted>
Prefix Address: 2001:DB8:0:12::
Prefix Length: 64, Options: None
<Output omitted>

```

Example 3-84 shows the IPv6 routing tables for R3 and R4. Configuring the areas using one of the stub options can help reduce the size of the routing tables.

Example 3-84 *OSPFv3 Routes in the Routing Tables of R3 and R4*

```

R3# show ipv6 route ospf
<Output omitted>
OI 2001:DB8:0:11::1/128 [110/10]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:12::/64 [110/74]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:14::/64 [110/20]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:44::/64 [110/21]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:220::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:221::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:222::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:223::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1

```

```
OI 2001:DB8:0:224::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:225::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:226::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:227::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:228::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:229::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22A::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22B::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22C::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22D::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22E::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
OI 2001:DB8:0:22F::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
```

R4# show ipv6 route ospf

<Output omitted>

```
O 2001:DB8:0:11::1/128 [110/10]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:12::/64 [110/74]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:13::/64 [110/20]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:33::/64 [110/21]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:220::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:221::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:222::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:223::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
```

```

OI 2001:DB8:0:224::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:225::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:226::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:227::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:228::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:229::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22A::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22B::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22C::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22D::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22E::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI 2001:DB8:0:22F::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0

```

When you display the IPv6 routing tables on R3 and R4, you will see numerous OSPF interarea routes. Making the nonbackbone area 2 a stub area can reduce the size of the R3 routing table. Summarizing the interarea routes on the area border router can reduce R4's routing table in area 0.

In Example 3-85, ABR R1 and area 2 router R3 are configured to act as a totally stubby area for IPv6.

Example 3-85 *Area 2 Routers Configured as a Totally Stubby Area*

```

R1(config)# router ospfv3 1
R1(config-router)# address-family ipv6 unicast
R1(config-router-af)# area 2 stub no-summary
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from FULL to DOWN,
Neighbor Down: Adjacency forced to reset

R3(config)# router ospfv3 1
R3(config-router)# address-family ipv6 unicast
R3(config-router-af)# area 2 stub
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 1.1.1.1 on Ethernet0/1 from LOADING to FULL,
Loading Done

```

Features specific to an address family are configured for the given address family. The stubbiness or total stubbiness of an area, for example, could be enabled individually for IPv4, IPv6, or both. In this scenario, area 2 is configured as a stub area for the IPv6 address family.

OSPF uses a stub feature flag in the Hello packets. This flag must match between the neighbors for the adjacency to be established. The flag is exchanged individually for each address family. This example illustrates how the adjacency fails if only one side has the area configured as stub, and then succeeds when both R1 and R3 have matching configuration.

The IPv6 and IPv4 routing in area 2 is verified by examining the routing table of R3 in Example 3-86.

Example 3-86 Examining the Differences Between R3's IPv4 and IPv6 Routing Tables

```
R3# show ipv6 route ospf
<Output omitted>
OI  ::/0 [110/11]
    via FE80::A8BB:CCFF:FE00:AB10, Ethernet0/1
R3# show ip route ospfv3
<Output omitted>
O IA   172.16.12.0/30 [110/74] via 172.16.13.1, 00:09:55, Ethernet0/1
O IA   172.16.14.0/25 [110/20] via 172.16.13.1, 00:09:55, Ethernet0/1
      192.168.1.0/32 is subnetted, 1 subnets
O IA   192.168.1.1 [110/10] via 172.16.13.1, 00:09:55, Ethernet0/1
      192.168.2.0/32 is subnetted, 1 subnets
O IA   192.168.2.2 [110/74] via 172.16.13.1, 00:09:55, Ethernet0/1
      192.168.4.0/32 is subnetted, 1 subnets
O IA   192.168.4.4 [110/20] via 172.16.13.1, 00:09:55, Ethernet0/1
```

When viewing the OSPF routing table for IPv4 and IPv6, notice the difference in the area 2 operations between the two address families. Area 2 acts as a standard area for IPv4 and therefore you see all external and interarea routes received via the backbone area. Area acts as a totally stubby area for IPv6. Therefore you see a default route toward the ABR.

Example 3-87 summarizes the IPv6 networks advertised by R2 (2001:DB8:0:220::/64 to 2001:DB8:0:22F::/64) using the smallest possible address block.

Example 3-87 Summarizing an IPv6 Address Block on R1

```
R1(config)# router ospfv3 1
R1(config-router)# address-family ipv6 unicast
R1(config-router-af)# area 1 range 2001:DB8:0:220::/60
```

Like in IPv4, OSPFv3 supports IPv6 address summarization. Interarea routes can be summarized on area border routers using the `area area-id range` command in the desired

address family mode. In this scenario, a set of IPv6 network addresses are summarized using the address block 2001:DB8:0:220::/60.

Although not demonstrated in these examples, you can summarize external routes on the ASBRs. To perform such summarization for IPv6, you would use the **summary-prefix** command in the IPv4 address family router configuration mode.

Example 3-88 verifies the IPv6 summarization effects in the backbone area by viewing the IPv6 routing table on the backbone router R4. R4 contains the summary address 2001:DB8:0:220::/60 instead of the individual smaller networks.

Example 3-88 OSPF Routes in R1's Routing Table

```
R4# show ipv6 route ospf
<Output omitted>
O   2001:DB8:0:11::1/128 [110/10]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI  2001:DB8:0:12::/64 [110/74]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI  2001:DB8:0:13::/64 [110/20]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI  2001:DB8:0:22::/64 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI  2001:DB8:0:33::/64 [110/21]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
OI  2001:DB8:0:220::/60 [110/75]
    via FE80::A8BB:CCFF:FE00:AB00, Ethernet0/0
```

Configuring Advanced OSPFv3

OSPFv3 offers you a set of tools that is very similar to that of OSPFv2 to fine-tune the OSPFv3 functionality.

Networks on the ASBR can be summarized during redistribution into OSPFv3. To configure an IPv6 summary prefix in Open Shortest Path First Version 3 (OSPFv3), use the following command in OSPFv3 router configuration mode, IPv6 address family configuration mode, or IPv4 address family configuration mode:

```
summary-prefix prefix [ not-advertise | tag tag-value ] [ nssa-only]
```

To restore the default, use the **no** form of this command. Table 3-8 describes the command parameters.

Table 3-8 *Parameters for summary-prefix Command*

Parameter	Description
<i>prefix</i>	IPv6 route prefix for the destination.
not-advertise	(Optional) Suppresses routes that match the specified prefix and mask pair. This keyword applies to OSPFv3 only.
tag tag-value	(Optional) Specifies the tag value that can be used as a match value for controlling redistribution via route maps. This keyword applies to OSPFv3 only.
nssa-only	(Optional) Limits the scope of the prefix to the area. Sets the NSSA-only attribute for the summary route (if any) generated for the specified prefix.

Example 3-89 shows a sample configuration. Redistribution is discussed in Chapter 4.

Example 3-89 *Configuring the summary-prefix Command on an ASBR*

```
Router(config)# router ospfv3 1
Router(config-router)# address-family ipv6 unicast
Router(config-router-af)# summary-prefix 2001:db8:1::/56
```

Load-balancing behavior can also be controlled on OSPFv3 routers. To control the maximum number of equal-cost routes that a process for OSPFv3 routing can support, use the **maximum-paths** command in IPv6 or IPv4 address family configuration mode, shown in Example 3-90. The range in OSPFv3 is from 1 through 64.

Example 3-90 *maximum-paths Command Configured in Address Family Mode*

```
Router(config)# router ospfv3 1
Router(config-router)# address-family ipv6 unicast
Router(config-router-af)# maximum-paths 8
```

OSPFv3 Caveats

The OSPF processes: traditional OSPFv2, traditional OSPFv3, and new OSPFv3 that uses the address families to supports both IP stacks, differ in the transport protocols.

The traditional OSPFv2 method, configured with the **router ospf** command, uses IPv4 as the transport mechanism. The traditional OSPFv3 method, configured with the **ipv6 router ospf** command, uses IPv6 as the transport protocol. The new OSPFv3 framework, configured with the **router ospfv3** command, uses IPv6 as the transport mechanism for both address families. Therefore, it will not peer with routers running the traditional OSPFv2 protocol.

The OSPFv3 address families feature is supported as of Cisco IOS Release 15.1(3)S and Cisco IOS Release 15.2(1)T. Cisco devices that run software older than these releases and

third-party devices will not form neighbor relationships with devices running the address family feature for the IPv4 address family because they do not set the address family bit. Therefore, those devices will not participate in the IPv4 address family SPF calculations and will not install the IPv4 OSPFv3 routes in the IPv6 Routing Information Base (RIB).

Summary

In this chapter, you learned about establishing OSPF neighbor relationships, building the OSPF link-state database, optimizing OSPF behavior, configuring OSPFv2 and OSPFv3. Some key points in this chapter are:

- OSPF uses a two-layer hierarchical approach dividing networks into a backbone area (area 0) and nonbackbone areas.
- For its operation, OSPF uses five packet types: Hello, DBD, LSR, LSU, and LSAck.
- OSPF neighbors go through several different neighbor states before adjacency results in Full state.
- OSPF elects DR/BDR routers on a multiaccess segment to optimize exchange of information.
- The most common OSPF network types are point-to-point, broadcast, nonbroadcast, and loopback.
- OSPF uses several different LSA types to describe the network topology.
- LSAs are stored in an LSDB, which is synchronized with every network change.
- OSPF calculates interface costs based on default reference bandwidth and interface bandwidth.
- Using SPF, OSPF determines the total lowest cost paths and selects them as the best routes.
- Intra-area routes are always preferred over interarea routes.
- Route summarization improves CPU utilization, reduces LSA flooding, and reduces routing table sizes.
- The **area range** command is used to summarize at the ABR. The **summary-address** command is used to summarize at the ASBR.
- Default routes can be used in OSPF to prevent the need for specific routes to each destination network.
- OSPF uses the **default-information originate** command to inject a default route.
- There are several OSPF area types: normal, backbone, stub, totally stubby, NSSA, and totally stubby NSSA.
- Use the **area area-id** command to define an area as stubby.

- Use the `area area-id stub` command with the `no-summary` keyword only on the ABR to define an area as totally stubby.
- For stub areas, external routes are not visible in the routing table, but are accessible via the intra-area default route.
- For totally stubby areas, interarea and external routes are not visible in the routing table, but are accessible via the intra-area default route.
- OSPFv3 for IPv6 supports the same basic mechanisms that OSPFv2 for IPv4, including the use of areas to provide network segmentation and LSAs to exchange routing updates.
- OSPFv3 features two new LSA types and has renamed two traditional LSA types.
- OSPFv3 uses link-local addresses to source LSAs.
- OSPFv3 is enabled per-interface on Cisco routers.
- New-style OSPFv3 and traditional OSPFv3 for IPv6, configured with `ipv6 router ospf`, can coexist in the network to provide IPv6 routing.

Review Questions

Answer the following questions, and then see Appendix A, “Answers to Review Questions,” for the answers.

1. What is the OSPF transport?
 - a. IP/88
 - b. TCP/179
 - c. IP/89
 - d. IP/86
 - e. UDP/520
2. An Area Border Router maintains _____.
 - a. A single database for all areas
 - b. A separate database for each area with which it is connected
 - c. Two databases: one for the backbone and one for all other areas
 - d. A separate routing table for each area
3. Which two methods does OSPF employ to conserve the computing resources?
 - a. Area-based segregation including stub areas
 - b. LSDB
 - c. Summarization
 - d. Redistribution
 - e. Network types

4. What is the difference between an LSA 3 and an LSA 4?
 - a. LSA 3 is a summary LSA, and LSA 4 is E1.
 - b. LSA 3 is E1, and LSA 4 is a summary.
 - c. LSA 3 is a summary for networks, and LSA 4 is a summary for ASBRs.
 - d. LSA 3 is a summary for ASBRs, and LSA 4 is a summary for networks.
5. Which two LSAs describe intra-area routing information?
 - a. Summary
 - b. External 1
 - c. External 2
 - d. Router
 - e. Network
6. An OSPF router receives an LSA and checks the sequence number of the LSA. This number matches the sequence number of an LSA that the receiving router already has. What does the receiving router do with the received LSA?
 - a. Ignores the LSA
 - b. Adds the LSA to the database
 - c. Sends the newer LSU to the source router
 - d. Floods the LSA to the other routers
7. What are the two reasons why route summarization is important?
 - a. Reduces LSA type 1 flooding
 - b. Reduces LSA type 3 flooding
 - c. Reduces the size of the routing table
 - d. Reduces the size of the neighbor table
8. Route summarization reduces the flooding of which two of the following LSA types?
 - a. Router
 - b. Network
 - c. Summary
 - d. External
 - e. NSSA
9. Stub area design can improve _____.
 - a. CPU utilization on routers in the stub
 - b. The number of adjacencies in the stub
 - c. Ability to reach outside networks
 - d. LSDB size on routers in the backbone
10. Which feature characterizes both OSPFv2 and OSPFv3?
 - a. Router ID in IPv4 format
 - b. Router ID in IPv6 format
 - c. Process activation using the `network` command
 - d. The same LSA types

11. Which address would you configure in the neighbor command to set up an OSPFv3 adjacency over an NBMA link?
- a. Local IPv4 address
 - b. Neighbor's IPv4 address
 - c. Interface link local IPv6 address
 - d. Local global IPv6 address
 - e. Neighbor's link-local IPv6 address
 - f. Neighbor's global IPv6 address
12. You can run a single OSPFv3 process using the **ipv6 router ospf** command to support a dual-stack environment. (True or false?)
- a. True
 - b. False
13. Which of the following is not a characteristic of link-state routing protocols?
- a. They respond quickly to network changes.
 - b. They broadcast every 30 minutes.
 - c. They send triggered updates when a network change occurs.
 - d. They may send periodic updates, known as link-state refresh, at long time intervals, such as every 30 minutes.
14. Link-state routing protocols use a two-layer area hierarchy composed of which two areas?
- a. Backbone area
 - b. Transit area
 - c. Regular area
 - d. Linking area
15. Which IPv4 address is used to send an updated LSA entry to OSPF DRs and BDRs?
- a. Unicast 224.0.0.5
 - b. Unicast 224.0.0.6
 - c. Multicast 224.0.0.5
 - d. Multicast 224.0.0.6
16. To ensure an accurate database, how often does OSPF flood (refresh) each LSA record?
- a. Every 60 minutes.
 - b. Every 30 minutes.
 - c. Every 60 seconds.
 - d. Every 30 seconds.
 - e. Flooding each LSA record would defeat the purpose of a link-state routing protocol, which strives to reduce the amount of routing traffic it generates.

- 17.** What kind of router generates LSA type 5 in a standard area?
 - a.** DR
 - b.** ABR
 - c.** ASBR
 - d.** ADR

- 18.** Where does a type 1 LSA flood to?
 - a.** To immediate peers
 - b.** To all other routers in the area where it originated
 - c.** To routers located in other areas
 - d.** To all areas

- 19.** How does a routing table reflect the link-state information of an intra-area route?
 - a.** The route is marked with O.
 - b.** The route is marked with I.
 - c.** The route is marked with IO.
 - d.** The route is marked with EA.
 - e.** The route is marked with O IA.

- 20.** Which type of external route is the default?
 - a.** E1.
 - b.** E2.
 - c.** E5.
 - d.** There is no default external route. OSPF adapts and chooses the most accurate one.

- 21.** How is the cost of an E1 external route calculated?
 - a.** The sum of the internal cost of each link the packet crosses
 - b.** The sum of the external cost and the internal cost of each link the packet crosses
 - c.** The external cost only
 - d.** The sum of all area costs, even those that are not used

Index

Numerics

2-Way neighbor state, 185

A

AAA (authentication, authorization, and accounting), 536-543

accounting, 542

authorization, 542

local authentication, 538-539

RADIUS, 536-541

TACACS+, 536-538, 541-542

ABRs (area border routers), 159

route summarization, 223-224

access-list tcp command, 638

access-list udp command, 640

accounting, 542

ACLs (access control lists), 625-648. *See also* route maps

AS-Path access lists, 494-496

distribute lists, 295-297

examples, 299-300

extended IP ACLs, 634-644

configuring, 635-643

examples, 642-643

location of, 643-644

implicit deny any statements, 296

infrastructure ACLs, 547-549

IPv6 ACLs, 405-409

OSPFv2 routes, redistributing into EIGRP, 303-304

route filtering, 294-297

standard IP ACLs, 626-634

configuring, 629-633

location of, 633-634

wildcard masks, 628-629

time-based ACLs, 644

verifying configuration, 647-648

vtv access, restricting, 645-647

active state (BGP), troubleshooting, 459-460

active timer (EIGRP), 108

address classes, 619-620

address families, 139-148

EIGRP for IPv4 address families, 139-142

EIGRP for IPv6 address families, 142-148

address family command, 139-140

address family topology configuration mode (named EIGRP), 148

address-family interface configuration mode (named EIGRP), 148

address-family ipv6 command, 142

adjacencies

- establishing, 169-170
- ExStart state, 177
- mismatched MTUs, 177-179
- on multiaccess networks, 170-171
- neighbor states, 184-186
- observing, 168-169

adjacency table, 335**administrative distance, 269-270**

- changing to enable optimal routing, 315-318

advanced distance vector protocols, 60.

See also distance vector protocols

advanced OSPFv3 configuration, 260-261**advertising**

- from BGP into an IGP, 681
- BGP networks, 450-457
- connected routes, 103-105
- networks into BGP, 680-681

af-interface command, 143**aggregate-address command, 677-679****aggregate addresses, 671-673****aggregator attribute, 672****antireplay protection, 39****applications for route maps, 305-306****archive command, 563-565****area 0, 231****area 1, 231****area 2, 231****area area-id virtual link command, 227****area ID, 163****areas**

- multi-area OSPF, 160
- stub areas, configuring, 229-234
 - cost of default route, 236-238*
 - NSSAs, 238-239*
 - totally stubby areas, 234-236*

AS (autonomous systems), 6

- EIGRP neighbor relationships, establishing, 65-72

route redistribution, 271

- default seed metrics, 273-275*
- one-point redistribution, 287-288*
- redistributing EIGRP for IPv6 routes into OSPFv3, 285-287*
- redistributing EIGRP routes into OSPF, 281-284*
- redistributing OSPFv2 routes into EIGRP, 276-279*
- redistributing OSPFv3 routes into EIGRP for IPv6, 279-281*
- seed metrics, 272-273*

ASBR summary LSAs, 189, 199-201**ASBRs (autonomous system boundary routers), 159**

- route summarization, 224-225

ASNs (autonomous system numbers), 378-380

- private ASNs, 696

AS-Path attribute, 473-474

- AS-Path access lists, 494-496

assigned TCP port numbers, 639-640**assigning IPv6 addresses, 398-405**

- manual assignment, 399
- SLAAC, 401

asymmetric routing, 5**atomic aggregate attribute, 672****attributes (BGP), 471-480**

- aggregator attribute, 672
- AS-Path attribute, 473-474
- atomic aggregate attribute, 672
- Community attribute, 475-476
- community attribute, 682
- format, 471
- influencing BGP path selection, 480-491
- Local Preference attribute, 475
- MED attribute, 476-478
- Next-Hop attribute, 474
- optional attributes, 472
- Origin attribute, 475
- type codes, 473

- Weight attribute, 478-480
- well-known attributes, 471-472
- authentication**
 - BGP, configuring, 593-596
 - EIGRP, configuring, 576-583
 - hashing authentication, 572-573
 - local authentication, 538-539
 - OSPF, configuring, 583-593
 - PPP, 23-26
 - RADIUS, 539-541
 - routing protocol authentication, 570-576
 - purpose of*, 570-573
 - TACACS+, 541-542
 - time-based key chains, 574-575
- authorization**, 542
- automatic summarization, EIGRP configuration**, 111-116

B

- backbone area, OSPF**, 159
- backbone routers**, 159
- backing up configuration files**, 563-565
- bandwidth**
 - EIGRP composite metric, calculating, 89
 - slow start, 42
- BDP (Bandwidth Delay Product)**, 41
- BDR (backup designated router), election process**, 170-175
- believability of routing protocols**, 269
- benefits**
 - of DMVPN, 35-36
 - of optimized IP addressing plan, 648-650
 - of OSPF route summarization, 217
- best path selection**
 - BGP, 467-470
 - influencing with attributes*, 480-491
 - convergence, 8-9

- EIGRP, 80-87
 - RD*, 80
- OSPF, 208-210
 - calculating best path*, 210-211
 - default interface costs*, 211-214
- BGP (Border Gateway Protocol)**
 - advertising networks, 450-457
 - attributes, 471-480
 - AS-Path attribute*, 473-474
 - Community attribute*, 475-476
 - format*, 471
 - influencing path selection*, 480-491
 - Local Preference attribute*, 475
 - MED attribute*, 476-478
 - Next-Hop attribute*, 474
 - optional attributes*, 472
 - Origin attribute*, 475
 - type codes*, 473
 - Weight attribute*, 478-480
 - well-known attributes*, 471-472
 - authentication, configuring, 593-596
 - best path selection, 467-470
 - BGP tables, 430-431
 - characteristics, 428-430
 - communities, 682-687
 - comparing with other routing protocols, 425-426
 - configuring
 - entering BGP configuration mode*, 442
 - requirements*, 442
 - default seed metric, 274
 - eBGP multihop, 463-464
 - eBGP sessions, configuring, 445-449
 - full-mesh environments, 440
 - hop-by-hop routing paradigm, 427
 - iBGP sessions, configuring, 449-450
 - interdomain routing, 424-425

- messages, 431-433
 - keepalive messages*, 431-432
 - notification messages*, 433
 - open messages*, 431-432
 - update messages*, 433
- neighbor relationships, 435-441
 - defining*, 443-444
 - external BGP neighbors*, 436-437
 - iBGP*, 438-439
 - internal BGP neighbors*, 437-438
 - speakers*, 435
 - troubleshooting*, 458-460
- neighbor states, 433
- next-hop-self feature, 457-458
- partial-mesh environments, 440-441
- path vector characteristics, 426-428
- peer groups, 498-502
- private ASNs, 696
- resetting BGP sessions
 - hard resets*, 464-465
 - soft resets*, 465-466
- route redistribution
 - advertising from BGP into an IGP*, 681
 - advertising networks into BGP*, 680-681
- route reflectors, 687-695
 - clusters*, 689
 - migration tips*, 692-694
 - verifying*, 695
- route summarization, 671-679
 - aggregate addresses*, 671-673
 - aggregate-address command*, 677-679
 - CIDR*, 671-673
 - network boundary summarization*, 673-674
 - network command*, 674-676
- routing updates, filtering, 492-498
- session resilience, 460-461
- sourcing from loopback address, 461-463

- TCP, 440
 - when to use, 433-434
- BGP tables, 430-431
- BGP4+, 379
- binary-to-decimal conversion, 618-619
- boundary routers, 268
- broadcast addresses, 11
- broadcast networks, 15, 186
- building EIGRP topology table, 76-94

C

- C networks, 75
- calculating
 - EIGRP composite metric, 88-92
 - inter-area route cost, 214-215
 - intra-area route cost, 214
 - IPv4 summary routes, 116-120
 - OSPF best path, 210-211
 - RD, 92-94
 - subnet masks, 621-624
 - VLSMs, 656-662
- CE (customer-edge) routers, 75
- CEF (Cisco Express Forwarding), 21, 330, 333-343
 - adjacency table, 335
 - disabling, 341-343
 - FIB table, 333-334
 - verifying, 335-341
- changing
 - administrative distance to enable optimal routing, 315-318
 - router IDs, 164-165
- CHAP (Challenge Handshake Authentication Protocol), 23-26
- characteristics
 - of BGP, 428-430
 - of prefix lists, 297-298
- CIDR (classless interdomain routing), 671-673, 667-669

Cisco IOS IP SLAs, 354-368

example configuration, 361

features, 354

operations

*configuring, 356-358**scheduling, 359*

path control, verifying, 360-361

responders, 355-356

sources, 354-355

targets, 354-355

tracking objects, configuring, 360

Cisco routers

configuration backups, 563-565

disabling unused services, 567

infrastructure ACLs, 547-549

management plane, securing, 529-530

router security policy, 530-531

classes of IPv4 addresses, 619-620**classic EIGRP configuration, comparing to named EIGRP, 150-151****classless routing protocols, 45****clusters, 689****commands**

access-list tcp command, 638

access-list udp command, 640

address family command, 139-140

address-family ipv6 command, 142

af-interface command, 143

archive command, 563-565

area area-id virtual link command, 227

debug eigrp packets hello command, 70-72

distribute-list in command, 295

distribute-list out command, 294-295

eigrp stub command, 100-102

encapsulation ppp interface command, 23

ip community-list command, 685

ip ospf neighbor detail command, 181

ip prefix list command, 298-299

ip route command, 20-21

match commands, 308

match community command, 685

neighbor default-originate command, 695

neighbor peer group command, 500

neighbor remote-as command, 443

neighbor send-community command, 683-684

network command, 674-676

no auto-summary command, 46

OSPFv3 new-style configuration commands, 252

ppp authentication command, 24

redistribute command, 276-277

redistribute static command, 121

route-map command, 306

router eigrp autonomous-system-number command, 65-66

router ospf process-id command, 162

router rip command, 46

set commands, 309

set community command, 682-683

show ip bgp neighbors command, 447-448

show ip cef command, 342

show ip eigrp interfaces command, 69

show ip eigrp neighbors command, 66-68

show ip eigrp traffic command, 79

show ip interface command, 343

show ip ospf database command, 192

show ip ospf neighbor command, 165

show ip ospf route command, 168

show ip protocols command, 46, 147

show ipv6 eigrp topology command, 133

show ipv6 interface brief command, 48

show ipv6 protocols command, 49

syntax conventions

undebg all command, 102

communities, 682-687**Community attribute, 475-476**

comparing

- IPv4 and dual BGP transport, 507-518
- named and classic EIGRP configuration, 150-151

composite metrics, 80

- EIGRP, 88-89
 - calculating, 89-92*

conditional debugging, 568-569**configuration backups, 563-565****configuration modes, named EIGRP, 148-150****configuring****ACLs**

- extended IP ACLs, 635-643*
- standard IP ACLs, 629-633*

BGP

- authentication, 593-596*
- eBGP sessions, 445-449*
- iBGP sessions, 449-450*
- peer groups, 498-502*
- requirements, 442*
- route reflectors, 694-695*

Cisco IOS IP SLAs

- example configuration, 361*
- operations, 356-358*
- tracking objects, 360*

DHCP, 384-385**Dynamic NAT, 389-390****EIGRP**

- authentication, 576-583*
- load balancing, 123-128*
- neighbor relationships, 65-72*
- route summarization, 110-120*

EIGRP for IPv6, 130-134**IPv4 static routes, 20-21****IPv6 ACLs, 406-409****IPv6 Internet connectivity, 399-400****MP-BGP, 507-518****named EIGRP, 136-151****NVI, 393-395****OSPF, 161-182**

- authentication, 583-593*
- passive interfaces, 187*
- route summarization, 218-223*
- router ID, 163*
- stub areas, 229-234*
- virtual links, 227-229*

OSPFv3, 240-246**PAT, 390-392****PBR, 345-348**

- example configuration, 348-353*

PPP, 23**PPPoE, 27-28****prefix lists, 298-299****RIPng, 47-50****route maps, 306-308****route redistribution**

- with route maps, 310-318*
- seed metrics, 273*

single-homed connections**DHCP, 382-383**

- provider-assigned IPv4 address, configuring, 381-382*

static default routes, 22**Static NAT, 388****connect neighbor state (BGP), 458****connected routes, advertising, 103-105****connecting enterprise networks to Internet, 374-375****redundancy, 374-375****console passwords, encrypting, 533-536****control plane, 328, 527****convergence, 8-9****EIGRP, 60****OSPF, 157****speed of, influencing, 8-9****converting between decimal and binary, 614, 618-619****cost**

- of inter-area routes, calculating, 214-215
- of intra-area routes, calculating, 214

D

data plane, 328, 527

databases

LSDB, 189-206

contents, displaying, 192

synchronizing, 204-205

OSPF exchange process, 169-170

RIPng, verifying, 53-54

DBD (Database Description) packets, 160

dead timer, manipulating, 179-182

debug eigrp packets hello command,
70-72

debugging

conditional debugging, 568-569

decimal-to-binary conversion, 618-619

default interface costs, OSPF, 211-214

default metrics, 272-273

default routes

cost of in stub areas, 236-238

obtaining, 120-123

propagating, 50-53

default seed metrics, 273-275

defining

BGP neighbor relationships, 443-444

route redistribution, 270-271

delay, EIGRP composite metric, 89

DHCP (Dynamic Host Configuration
Protocol), 382-383

configuring, 384-385

provider-assigned IPv4 address,
obtaining, 383-384

DHCPv6, 402-405

DHCPv6-PD, 405

stateful DHCPv6, 404

stateless DHCPv6, 403-404

Dijkstra's algorithm, 156

disabling

automatic summarization, 115-116

CEF, 341-343

unused services, 567

disadvantages

of single-homed connections, 410

of static routing, 20

displaying

EIGRP for IPv6 routing table, 133

LSDB contents, 192

distance vector protocols, 7

EIGRP

active timer, 108

authentication, configuring, 581-583

best path selection, 80-87

composite metric, 88-89

default routes, obtaining, 120-123

default seed metric, 274

DUAL, 76

feasibility condition, 91

features, 60-62

hello packets, 70-72

hello timer, 71

hold timers, 71

for IPv4, 64-72

load balancing, 123-128

multiple network layer support, 61

*neighbor relationships, 63-64,
66-69*

neighbor table, 63

over Frame Relay, 74

over Layer 2 MPLS VPN, 75-76

over Layer 3 MPLS VPN, 74-75

partial updates, 61

passive interfaces, 69-70

query packets, 79, 95-96

reply packets, 79

route summarization, 109

RTP, 62

SIA state, 108-109

topology table, 63

unequal metric load balancing, 62

VLSM, 61

wide metric, 90

- EIGRP for IPv6
 - configuring*, 133
 - IPv6 summary route, determining*, 134-136
 - routing table, displaying*, 133
- RIP, 43-45
- RIPng, 43-54
 - configuring*, 47-50
 - default routes, propagating*, 50-53
 - manual summarization*, 50
- RIPv2, 45-47
- distribute lists, 294-297**
 - EIGRP routes, redistributing into OSPF, 304-305
 - examples, 299-300
 - OSPFv2 routes, redistributing into EIGRP, 303-304
- distribute-list in command, 295**
- distribute-list out command, 294-295**
- DLCIs (data-link connection identifiers), 29**
- DMVPN (Dynamic Multipoint Virtual Private Network), 35-36**
 - IPsec, 39-40
 - NHRP, 37-38
- Down neighbor state, 185**
- DR (designated router), election process, 170-175**
- DUAL, 76**
 - FD, 80
 - feasible successors, 82-87
 - RD, 80
 - successor routes, 82-87
- dual BGP transport, comparing to IPv4, 507-518**
- dual multihomed connections, 375**
- dual-homed connections, 375, 410-413**
 - best path, *configuring*, 411-413
 - connectivity options, 411
- Dynamic NAT, 387**
 - configuring*, 389-390

E

- eBGP multihop, 463-464**
- eBGP sessions, configuring, 445-449**
- EGPs (exterior gateway protocols), 6**
- EIGRP (Enhanced Interior Gateway Routing Protocol)**
 - active timer, 108
 - authentication
 - configuring*, 576-583
 - automatic summarization
 - configuring*, 111-116
 - composite metric, 88-89
 - calculating*, 89-92
 - wide metric*, 90
 - convergence, 60
 - default routes
 - obtaining*, 120-123
 - default seed metric, 274
 - DUAL, 76
 - features, 60-62
 - hello packets
 - observing*, 70-72
 - hello timer, 71
 - hold timers, 71
 - for IPv4, 64-72
 - load balancing
 - across unequal-metric paths*, 126-128
 - configuring*, 123-128
 - multiple network layer support, 61
 - named EIGRP
 - address families*, 139-148
 - configuring*, 136-151
 - neighbor relationships, 63-64
 - establishing*, 65-72
 - over Frame Relay*, 74
 - verifying*, 66-69
 - neighbor table, 63
 - over Layer 2 MPLS VPN, 75-76

- over Layer 3 MPLS VPN, 74-75
- partial updates, 61
- passive interfaces, 69-70
- query packets, 79, 95-96
- redistribution
 - redistributed static routes, verifying, 122-123*
- reply packets, 79
- route redistribution, 287-288
 - multipoint redistribution, 288-289*
 - one-point redistribution, 287-288*
 - redistributing EIGRP routes into OSPF, 281-284, 304-305*
 - redistributing OSPFv2 routes into EIGRP, 276-279*
 - routing loops, preventing, 291*
 - troubleshooting, 289-291*
 - verifying, 292*
- route summarization, 109
 - configuring, 110-120*
 - configuring manual summarization, 116-120*
 - summary routes, determining, 116-120*
- RTP, 62
- SIA state, 108-109
- stub routing,
 - connected routes, advertising, 103-105*
 - receive-only stub routers, 105-107*
- topology table, 63
 - best path selection, 80-87*
 - building, 76-94*
 - FD, 80*
 - feasible successors, 81-87*
 - feasibility condition, 91*
 - RD, 80, 92-94*
 - redistribution, 88*
 - routing knowledge exchange, 88*
 - successor routes, 82-87*
- timers, manipulating, 73
 - unequal metric load balancing, 62
 - VLSM, 61
- EIGRP for IPv6, 128-136**
 - address families, 142-148
 - authentication, configuring, 581-583
 - configuring, 130-134
 - IPv6 summary route, determining, 134-136
 - redistributing EIGRP for IPv6 routes into OSPFv3, 285-287
 - redistributing OSPFv3 routes into EIGRP for IPv6, 279-281
 - routing table, displaying, 133
- eigrp stub command, 100-102**
- electing OSPF DR/BDR, 170-175**
- encapsulation ppp interface command, 23**
- encrypting passwords, 531-536**
- entering BGP configuration mode, 442**
- enterprise networks**
 - connecting to Internet, 374-375
 - connection types, 375*
 - redundancy requirements, 374-375*
 - inbound connectivity, 374
 - infrastructure, 2-3
 - outbound connectivity, 374
 - uRPF, 549-551
- established keyword, 639**
- established state (BGP), 458**
- establishing**
 - EIGRP neighbor relationships, 65-72
 - OSPF adjacencies, 169-170
- Ethernet, PPPoE, 26-28**
- EVN (Cisco Easy Virtual Network), 601-602**
- examples**
 - of CIDR, 668-669
 - of Cisco IOS IP SLA configuration, 361
 - of distribute lists, 299-300
 - of extended IP ACLs, 642-643
 - of PBR configuration, 348-353
 - of prefix lists, 300-301
 - of VLSM calculation, 656-662

Exchange neighbor state, 186
 exchange process, OSPF databases,
 169-170
 ExStart adjacency state, 177, 185
 extended IP ACLs, 634-644
 configuring, 635-643
 established keyword, 639
 examples, 642-643
 location of, 643-644
 external BGP neighbors, 436-437

F

falsification of routing information, 570
 fast switching, 329, 332-333
 FD (feasible distance), 80
 feasibility condition, 91
 feasible successors, 81-87
 features
 of Cisco IOS IP SLAs, 354
 of EIGRP, 60-62
 of OSPF, 156-157
 of PBR, 344-345
 FHRPs (first-hop redundancy
 protocols), 412
 FIB table, 333-334
 filtering routes
 ACLs, AS-Path access lists, 494-496
 distribute lists, 294-297
 examples, 299-300
 prefix lists, 297-305
 BGP filtering, 492-493
 characteristics of, 297-298
 configuring, 298-299
 example prefix list, 300-301
 verifying, 301-302
 reasons for, 292-293
 route maps, 305-318
 applications for, 305-306
 BGP filtering, 496-498
 configuring, 306-308

match statements, 308-309
 *mutual redistribution with route
 maps, 313-314*
 *mutual redistribution without
 route filtering, 312-313*
 route redistribution, 310-318
 sequence numbers, 307
 set statements, 308-309
 statements, 307

format of BGP attributes, 471
 fragmentation, 40-41
 Frame Relay NBMA networks, 16-17,
 28-31
 EIGRP neighbor relationships, 74
 Full neighbor state, 186
 full-mesh environments, 440

G-H

global unicast addresses, 13
 GRE (Generic Routing Encapsulation),
 routing over, 34-35
 hard resets, 464-465
 hashing authentication, 572-573
 HDLC (High-Level Data Link Control), 23
 Hello packets
 EIGRP, observing, 70-72
 OSPF, 160
 Router Priority field, 174-175
 hello timer
 EIGRP, 71
 manipulating, 73
 OSPF, manipulating, 179-182
 hierarchical structure of OSPF, 158-159
 hold timers (EIGRP), 71
 manipulating, 73
 hop-by-hop routing paradigm, 427
 HSRP (Hot Standby Router Protocol), 412
 hub-and-spoke topology, 20
 OSPF behavior in, 175-177
 hybrid VPNs, 32

-
- IANA (Internet Assigned Numbers Authority), 376-377**
 - ICANN (Internet Corporation for Assigned Names and Numbers), 376**
 - ICMP (Internet Control Message Protocol), 84**
 - messages, 637
 - Redirect messages, 42-43
 - type names, 637
 - ICMPv6 (Internet Control Message Protocol for IPv6), 14-15**
 - idle state (BGP), troubleshooting, 459**
 - IGPs (interior gateway protocols), 6**
 - IGRP (Interior Gateway Routing Protocol), 60**
 - implementing**
 - route summarization, 666
 - VLSM in scalable networks, 654-656
 - VRF-Lite, 597-601
 - implicit deny any statements, 296**
 - INARP (Inverse Address Resolution Protocol), 29**
 - inbound connectivity, 374**
 - IND (Inverse Neighbor Discovery), 29**
 - influencing**
 - BGP path selection, 480-491
 - speed of convergence, 8-9
 - infrastructure ACLs, 547-549**
 - infrastructure of enterprise networks, 2-3**
 - Init neighbor state, 185**
 - input requirements for routing protocol selection, 5**
 - inside global addresses, 386**
 - inside local addresses, 386**
 - interarea prefix LSAs for ABRs, 245**
 - interarea router LSAs for ASBRs, 245**
 - inter-area routes, 167-168**
 - calculating cost of, 214-215
 - interfaces**
 - CEF, enabling and disabling, 341-343
 - flapping, 663
 - seed metrics, 272-273
 - configuring, 273*
 - internal BGP neighbors, 437-438**
 - internal routers, 159**
 - Internet**
 - ISPs, connecting to enterprise networks, 374-375
 - routing over, 18
 - securing IPv6 Internet connectivity, 409
 - intra-area routes, 167-168**
 - calculating cost of, 214
 - ip community-list command, 685**
 - ip ospf neighbor detail command, 181**
 - ip prefix list command, 298-299**
 - ip route command, 20-21**
 - IPsec, 39-40**
 - IPv4**
 - ACLs, 625-648
 - extended IP ACLs, 634-644*
 - IP standard ACLs, 626-634*
 - standard IP ACLs, 626-634*
 - time-based ACLs, 644*
 - verifying configuration, 647-648*
 - vty access, restricting, 645-647*
 - address classes, 619-620
 - address planning, 648-653
 - CIDR, 667-669
 - converting between decimal and binary, 618-619
 - decimal-to-binary conversion, 614
 - EIGRP
 - configuring, 64-72*
 - summary routes, calculating, 116-120*
 - multicast addresses, 12
 - named EIGRP, address families, 139-142
 - nonscalable network addressing, 651-653
 - OSPFv3 for IPv4 and IPv6, configuring, 246-260
 - PMTUD, 41

- private addresses, 620
- scalable network addressing, 650-651
- static routing, 20-21
 - CEF, 21
- subnet masks, 620-625
 - calculating, 621-624
 - representing with prefixes, 624-625
- VLSM
 - calculation examples, 656-662
 - implementing in scalable networks, 654-656
 - route summarization, 665
 - subnet masks, 653-654
- IPv6**
 - ACLs, 405-409
 - address assignment, 398-405
 - manual assignment, 399
 - SLAAC, 401
 - DHCPv6, 402-405
 - stateful DHCPv6, 404
 - stateless DHCPv6, 403-404
 - EIGRP for IPv6, 128-136
 - authentication, configuring, 581-583
 - configuring, 130-134
 - IPv6 summary route, determining, 134-136
 - one-point redistribution, 287-288
 - redistributing OSPFv3 routes into EIGRP for IPv6, 279-281
 - routing table, displaying, 133
 - global unicast addresses, 13
 - ICMPv6, 14-15
 - IND, 29
 - Internet connectivity, configuring, 399-400
 - link-local addresses, 13
 - loopback addresses, 14
 - MP-BGP, 502-507
 - BGP filtering mechanisms, 518-520
 - configuring, 507-518
 - dual BGP transport, comparing to IPv4, 507-518
 - exchanging IPv6 routes over IPv4 sessions, 504-506
 - exchanging IPv6 routes over IPv6 sessions, 506-507
 - named EIGRP
 - address families, 139-148
 - comparing with classic EIGRP configuration, 150-151
 - configuration modes, 148-150
 - configuring, 136-151
 - NAT64, 405
 - NPTv6, 405
 - NTP, 557
 - OSPFv3 for IPv4 and IPv6
 - configuring, 246-260
 - securing Internet connectivity, 409
 - unique local addresses, 14
 - unspecified address, 14
- IS-IS (Intermediate System-to-Intermediate System), seed metrics, 274**
- ISPs (Internet service providers), 4**
 - connection types, 375
 - dual-homed connections
 - best path, configuring, 411-413
 - connectivity options, 411
 - multihomed connections, 413-415
 - provider aggregatable address space, 378
 - provider independent address space, 378
 - requirements for enterprise network to ISP connectivity, 374-375
 - single-homed connections
 - DHCP, 382-383
 - disadvantages of, 410
 - NAT, 385-393
 - provider-assigned IPv4 address, configuring, 381-382

J-K

K values, 89

keepalive messages (BGP), 431-432

key chains, 574-575

L

latency, 42

Layer 2 MPLS VPN

EIGRP, establishing, 75-76

OSPF neighbor relationships, 184

Layer 3 MPLS VPNs

EIGRP, establishing, 74-75

OSPF neighbor relationships, 182-183

leased lines, 28

LFN (long fat network), 41

limitations

of NAT, 392-393

of RADIUS and TACACS+, 542-543

link LSAs, 245

link-local addresses, 13

link-state protocols, 7

OSPF, 155-156

ABRs, 159

adjacencies, observing, 168-169

area ID, 163

ASBR summary LSAs, 189,
199-201

ASBRs, 159

authentication, configuring,
583-593

autonomous systems LSAs, 189

backbone routers, 159

best path selection, 208-210

configuring, 161-182

convergence, 157

database exchange process,
169-170

default seed metric, 274

design restrictions, 160

features, 156-157

hierarchical structure, 158-159

in hub-and-spoke topology,
175-177

inter-area routes, 167-168

internal routers, 159

intra-area routes, 167-168

LSDB, 189-206

manual summarization, 157

messages, 160-161

metrics, 157

multi-area OSPF, 160

network LSAs, 189

network types, 186-187

optimizing, 215-239

passive interfaces, configuring, 187

process ID number, 162

route summarization, 216-225

router IDs, 163-165

router LSAs, 189

SPF algorithm, 207-208

summary LSAs, 189, 197-199

virtual links, 225-229

OSPFv3, 239-262

advanced configuration, 260-261

configuring, 240-246

load balancing

EIGRP

across unequal-metric paths,
126-128

configuring, 123-128

unequal metric load balancing, 62

Loading neighbor state, 186

local authentication, 538-539

Local Preference attribute, 475

location of standard IP ACLs, 633-634

logging, 551-552

logical AND operation, 307

loopback addresses, 14

loopback networks, 187

LSAck (Link-State Acknowledgment)

packets, 161

LSAs (link-state advertisements)

- ASBR summary LSAs, 199-201
- autonomous systems LSAs, 189
- interarea prefix LSAs for ABRs, 245
- interarea router LSAs for ASBRs, 245
- link LSAs, 245
- network LSAs, 189, 196-197
- router LSAs, 189, 192-196
- summary LSAs, 189, 197-199

LSDB (link-state database), 189-206

- contents, displaying, 192
- paranoid updates, 204
- periodic database changes, 203-204
- synchronizing, 204-205
 - on multiaccess networks, 206*
- Type 1 LSAs, 192-196
- Type 2 LSAs, 196-197
- Type 3 LSAs, 197-199
- Type 4 LSAs, 199-201
- Type 5 LSAs, 201-203

LSR (Link-State Request) packets, 160**LSU (Link-State Update) packets, 161**

M

management plane, 527

- securing, 529-530

manipulating

- EIGRP timers, 73
- OSPF timers, 179-182

manual IPv6 address assignment, 399**manual summarization, 50**

- EIGRP, configuring, 116-120
- OSPF, 157

master/slave relationship, determining, 169-170**match commands (PBR), 346-347****match community command, 685****match statements (route maps), 308-309****MED attribute, 476-478****messages**

- BGP, 431-433
 - keepalive messages, 431-432*
 - notification messages, 433*
 - open messages, 431-432*
 - update messages, 433*

ICMP, 637

ICMP Redirect, 42-43

ICMPv6, 14-15

OSPF, 160-161

metrics

- administrative distance, 269-270
 - changing to enable optimal routing, 315-318*
- composite metrics, 80
- OSPF, 157
- seed metrics, 272-273
 - configuring, 273*
 - default seed metrics, 273-275*
- wide metric (EIGRP), 90

mGRE (Multipoint GRE), 36**MIB (Management Information Base), 559****migration tips for route reflectors, 692-694****mismatched MTUs, 177-179****mismatched OSPF hello timers, troubleshooting, 180****MP-BGP, 502-507**

- BGP filtering mechanisms, 518-520
- configuring, 507-518
- dual BGP transport, comparing to IPv4, 507-518
- IPv6 routes
 - exchanging over IPv4 sessions, 504-506*
 - exchanging over IPv6 sessions, 506-507*

MPLS (Multiprotocol Label Switching) VPN, 31-32

Layer 2 MPLS VPN

- EIGRP, establishing, 75-76*
- OSPF neighbor relationships, 184*

- Layer 3 MPLS VPNs
 - EIGRP, establishing*, 74-75
 - OSPF neighbor relationships*, 182-183
 - routing over, 32-34
 - MSS (Maximum Segment Size)**, 40-41
 - MTU (Maximum Transmission Unit)**, 40, 88
 - mismatched MTUs, 177-179
 - multiaccess networks**
 - LSDB synchronization, 206
 - OSPF adjacencies, 170-171
 - multi-area OSPF**, 160
 - multicast addresses**, 11
 - IPv4, 12
 - multihomed connections**, 375, 413-415
 - multiple network layer support**, EIGRP, 61
 - multipoint one-way redistribution**, 288
 - multipoint redistribution**, 288-289
 - multipoint two-way redistribution**, 289
 - multiprotocol routing**, 267-270
 - administrative distance, 269-270
 - reasons for, 269
 - solutions for, 270
 - mutual redistribution**
 - with route maps, 313-314
 - without route filtering, 312-313
- ## N
-
- NA (Network Advertisement) messages**, 15
 - named EIGRP**
 - address families, 139-148
 - versus classic EIGRP configuration, 150-151
 - configuration modes, 148-150
 - configuring, 136-151
 - NAT (Network Address Translation)**, 385-393
 - Dynamic NAT, configuring, 389-390
 - limitations of, 392-393
 - NVI, 393-397
 - configuring*, 393-395
 - verifying*, 396-397
 - PAT, configuring, 390-392
 - Static NAT, configuring, 388
 - NAT64**, 405
 - NBMA (Nonbroadcast Multiaccess) networks**, 15-17
 - Frame Relay, 16-17
 - OSPF database synchronization, 170
 - point-to-multipoint subinterfaces, 17
 - point-to-point subinterfaces, 17
 - ND (Neighbor Discovery) address resolution**, 14
 - need for path control**, 343-344
 - neighbor default-originate command**, 695
 - neighbor peer group command**, 500
 - neighbor relationships**
 - BGP, 435-441
 - defining*, 443-444
 - external BGP neighbors*, 436-437
 - iBGP*, 438-439
 - internal BGP neighbors*, 437-438
 - troubleshooting*, 458-460
 - EIGRP, 63-64
 - establishing*, 65-72
 - hello packet transmission, observing*, 70-72
 - over Frame Relay*, 74
 - query packets*, 95-96
 - verifying*, 66-69
 - OSPF
 - adjacencies, establishing*, 169-170
 - adjacencies, observing*, 168-169
 - neighbor states*, 184-186
 - over Layer 2 MPLS VPN*, 184
 - over Layer 3 MPLS VPN*, 182-183
 - over point-to-point links*, 182
 - verifying*, 165
 - neighbor remote-as command**, 443
 - neighbor send-community command**, 683-684

neighbor table, EIGRP, 63

network boundary summarization, 673-674

network command, 674-676

network LSAs, 189, 196-197

networks

- BGP, advertising, 450-457
- broadcast networks, 15
- C networks, 75
- latency, 42
- LFN, 41
- multiprotocol routing, 267-270
 - administrative distance*, 269-270
 - reasons for*, 269
 - solutions for*, 270
- NBMA, 15-17
 - Frame Relay NBMA networks*, 16-17
 - point-to-multipoint subinterfaces*, 17
 - point-to-point subinterfaces*, 17
- OSPF network types, 186-187
- point-to-point, 15

Next-Hop attribute, 474

next-hop-self feature (BGP), 457-458

NHRP (Next-Hop Routing Protocol), 37-38

no auto-summary command, 46

nonbackbone area, OSPF, 159

nonbroadcast networks, 187

nonscalable network addressing, 651-653

notification messages (BGP), 433

NPTv6, 405

NS (Neighbor Solicitation) messages, 14

NSSAs (not-so-stubby areas), 238-239

NTP (Network Time Protocol), 552-558

- in IPv6 environments, 557
- modes, 552
- SNTP, 557-558
- versions, 556-557

NVI (NAT Virtual Interface), 393-397

- configuring, 393-395
- verifying, 396-397

O

observing

- adjacencies, 168-169
- mismatched MTUs, 177-179

observing EIGRP hello packet transmission, 70-72

obtaining

- default routes, 120-123
- provider-assigned IPv4 address with DHCP, 383-384

one-point redistribution, 287-288

one-way redistribution, 288

open confirm state (BGP), 458

open messages (BGP), 431-432

open sent neighbor state (BGP), 458

operations (IP SLAs)

- configuring, 356-358
- scheduling, 359

optimal routing, enabling, 315-318

optimizing

- IP addressing, 648-650
- OSPF, 215-239

optional attributes (BGP), 472

Origin attribute, 475

OSPF (Open Shortest Path First), 155-156

- ABRs, 159
- adjacencies
 - establishing*, 169-170
 - ExStart state*, 177
 - mismatched MTUs*, 177-179
 - on multiaccess networks*, 170-171
 - neighbor states*, 184-186
 - observing*, 168-169
- area ID, 163
- ASBRs, 159
- authentication, configuring, 583-593

- backbone routers, 159
- BDR
 - election process, 170-175*
- best path selection, 208-210
 - calculating best path, 210-211*
 - default interface costs, 211-214*
- configuring, 161-182
- convergence, 157
- default seed metric, 274
- design restrictions, 160
- DRs, election process, 170-175
- features, 156-157
- Hello packets, Router Priority field, 174-175
- hierarchical structure, 158-159
- in hub-and-spoke topology, 175-177
- inter-area routes, 167-168
 - calculating cost of, 214-215*
- internal routers, 159
- intra-area routes, 167-168
 - calculating cost of, 214*
- LSAs
 - ASBR summary LSAs, 189, 199-201*
 - autonomous systems LSAs, 189*
 - network LSAs, 189, 196-197*
 - router LSAs, 189, 192-196*
 - summary LSAs, 189, 197-199*
- LSDB, 189-206
 - contents, displaying, 192*
 - periodic database changes, 203-204*
 - synchronizing, 204-205*
- manual summarization, 157
- messages, 160-161
- metrics, 157
- multi-area OSPF, 160
- neighbor relationships
 - over Layer 2 MPLS VPN, 184*
 - over Layer 3 MPLS VPN, 182-183*
 - over point-to-point links, 182*
 - verifying, 165*
- network types, 186-187
- optimizing, 215-239
- passive interfaces, configuring, 187
- process ID number, 162
- route redistribution
 - into EIGRP, 276-279, 303-304*
 - multipoint redistribution, 288-289*
 - one-point redistribution, 287-288*
 - redistributing EIGRP routes into OSPF, 281-284*
 - routing loops, preventing, 291*
 - troubleshooting, 289-291*
 - verifying, 292*
- route summarization
 - on ABRs, 223-224*
 - on ASBRs, 224-225*
 - benefits of, 217*
 - configuring, 218-223*
- router IDs
 - changing, 164-165*
 - configuring, 163*
 - selecting, 163*
- SPF algorithm, 207-208
- stub areas
 - configuring, 229-234*
 - cost of default route, 236-238*
 - totally stubby areas, 234-236*
- timers, manipulating, 179-182
- virtual links, 225-229
 - configuring, 227-229*
 - removing, 227*
- OSPFv3, 239-262**
 - advanced configuration, 260-261
 - caveats, 261-262
 - configuring, 240-246
 - interarea prefix LSAs for ABRs, 245
 - interarea router LSAs for ASBRs, 245
 - link LSAs, 245
 - new-style configuration commands, 252
 - one-point redistribution, 287-288

- redistributing EIGRP for IPv6 routes into OSPFv3, 285-287
- redistributing OSPFv3 routes into EIGRP for IPv6, 279-281
- outbound connectivity, 374**
- outbound redistribution, 271**
- outside global addresses, 386**
- outside local addresses, 386**

P

packet forwarding, 327-343

- CEF, 330, 333-343
 - adjacency table, 335*
 - disabling, 341-343*
 - FIB table, 333-334*
 - verifying, 335-341*

- control plane, 328
- data plane, 328
- fast switching, 329, 332-333
- process switching, 328, 332-333

packets. *See also* LSAs

- EIGRP
 - hello packet transmission, observing, 70-72*
 - queries, 95-96*
- fragmentation, 40-41
- OSPF, 160-161
 - Hello packets, 174-175*

PAP (Password Authentication Protocol), 23-26

parameters

- access-list tcp command, 638
- access-list udp command, 640
- address family command, 139-140
- address-family ipv6 command, 142
- area area-id virtual link command, 227
- distribute-list in command, 295
- distribute-list out command, 294-295
- eigrp stub command, 100-102
- ip prefix list command, 298-299

- ip route command, 20-21
- neighbor default-originate command, 695
- neighbor peer group command, 500
- neighbor remote-as command, 443
- redistribute command, 276-277
- route-map command, 306

paranoid updates, 204

partial updates, EIGRP, 61

partial-mesh environments, 440-441

passive interfaces

- EIGRP, 69-70
- OSPF, configuring, 187

passwords, encrypting, 531-536

PAT (Port Address Translation), 387

- configuring, 390-392

path control

- Cisco IOS IP SLAs, 354-368
 - features, 354*
 - operations, configuring, 356-358*
 - operations, scheduling, 359*
 - responders, 355-356*
 - sources, 354-355*
 - targets, 354-355*
 - tracking objects, configuring, 360*
 - verifying, 360-361*

- need for, 343-344

PBR, 344-353

- configuring, 345-348*
- example configuration, 348-353*
- features, 344-345*
- match commands, 346-347*
- set commands, 347*
- verifying, 348*

path vector protocols, BGP, 7, 423

- advertising networks, 450-457
- attributes, 471-480
 - AS-Path attribute, 473-474*
 - Community attribute, 475-476*
 - format, 471*
 - influencing path selection, 480-491*
 - Local Preference attribute, 475*

- MED attribute*, 476-478
- Next-Hop attribute*, 474
- optional attributes*, 472
- Origin attribute*, 475
- type codes*, 473
- Weight attribute*, 478-480
- well-known attributes*, 471-472
- authentication, configuring, 593-596
- best path selection, 467-470
- BGP tables, 430-431
- characteristics, 428-430
- communities, 682-687
- comparing with other routing protocols, 425-426
- configuring
 - entering BGP configuration mode*, 442
 - requirements*, 442
- eBGP multihop, 463-464
- eBGP sessions, configuring, 445-449
- full-mesh environments, 440
- hop-by-hop routing paradigm, 427
- iBGP sessions, configuring, 449-450
- interdomain routing, 424-425
- messages, 431-433
 - keepalive messages*, 431-432
 - notification messages*, 433
 - open messages*, 431-432
 - update messages*, 433
- neighbor relationships, 435-441
 - defining*, 443-444
 - external BGP neighbors*, 436-437
 - iBGP*, 438-439
 - internal BGP neighbors*, 437-438
 - speakers*, 435
 - troubleshooting*, 458-460
- neighbor states, 433
- next-hop-self feature, 457-458
- partial-mesh environments, 440-441
- path vector characteristics, 426-428
- peer groups, 498-502
- private ASNs, 696
- resetting BGP sessions
 - hard resets*, 464-465
 - soft resets*, 465-466
- route redistribution
 - advertising from BGP into an IGP*, 681
 - advertising networks into BGP*, 680-681
- route reflectors, 687-695
 - clusters*, 689
 - configuring*, 694-695
 - migration tips*, 692-694
 - verifying*, 695
- route summarization, 671-679
 - aggregate addresses*, 671-673
 - aggregate-address command*, 677-679
 - CIDR*, 671-673
 - network boundary summarization*, 673-674
 - network command*, 674-676
- routing updates, filtering, 492-498
- session resilience, 460-461
- sourcing from loopback address, 461-463
- TCP, 440
- when to use, 433-434
- PBR (policy-based routing)**, 306, 344-353
 - configuring, 345-348
 - example configuration*, 348-353
 - features, 344-345
 - match commands, 346-347
 - set commands, 347
 - verifying, 348
- PE (provider-edge) routers**, 75
- peer groups, 498-502
- periodic LSDB changes, 203-204
- planning for route redistribution, 271
- PMTUD (Path MTU Discovery), 41
- point-to-multipoint subinterfaces, 17

- point-to-point networks, 15, 186
 - OSPF neighbor relationships, 182
- PPP (Point-to-Point Protocol)
 - authentication, 23-26
 - configuring, 23
- ppp authentication command, 24
- PPPoE (Point-to-Point Protocol over Ethernet), 26-28
 - configuring, 27-28
- prefix lists, 297-305
 - BGP filtering, 492-493
 - characteristics of, 297-298
 - configuring, 298-299
 - EIGRP routes, redistributing into OSPF, 304-305
 - example prefix list, 300-301
 - verifying, 301-302
- preventing routing loops, 291
- private addresses, 620
- private ASNs, 696
- privileged EXEC password, encrypting, 533
- process ID number, 162
- process switching, 328, 332-333
- propagating default routes with RIPng, 50-53
- provider aggregatable address space, 378
- provider independent address space, 378
- public IP address assignment
 - IANA, 376-377
 - provider aggregatable address space, 378
 - provider independent address space, 378
 - public address space, 377-378
 - RIRs, 377

Q-R

- QoS (quality of service), 345
- query packets (EIGRP), 79, 95-96
- RA (Router Advertisement) messages, 14
- RADIUS, 536-541
 - limitations of, 542-543

- RD (reported distance), 80
 - calculating, 92-94
- reasons for filtering routes, 292-293
- reasons for multiprotocol routing, 269
- receive-only stub routers, 105-107
- Redirect messages, 15
- redistribute command, 276-277
- redistribute static command, 121
- redistribution, 88
 - redistributed static routes, verifying, 122-123
- redundancy
 - FHRPs, 412
 - requirements for enterprise network to ISP connectivity, 374-375
- reliability, calculating EIGRP composite metric, 89
- remote access, SSH, 543-547
- removing virtual links, 227
- reply packets (EIGRP), 79
- representing subnet masks with prefixes, 624-625
- requirements for BGP configuration, 442
- reserved UDP port numbers, 642
- resetting BGP sessions
 - hard resets, 464-465
 - soft resets, 465-466
- responders, 355-356
- restricting vty access, 645-647
- RFC 1918, 385
- RFC 3587, 13
- RFC 4760, 502
- RIP (Routing Information Protocol), 43-45
 - default seed metric, 274
- RIPng (Routing Information Protocol Next Generation), 43-54
 - configuring, 47-50
 - database, verifying, 53-54
 - default routes, propagating, 50-53
 - manual summarization, 50
- RIPv2, 45-47

RIRs (regional Internet registries), 13, 377

role of routing protocols, 3-5

route filtering

ACLs

AS-Path access lists, 494-496

IPv6 ACLs, 405-409

distribute lists, 294-297

examples, 299-300

prefix lists, 297-305

BGP filtering, 492-493

characteristics of, 297-298

configuring, 298-299

examples, 300-301

verifying, 301-302

reasons for, 292-293

route maps, 305-318

applications for, 305-306

BGP filtering, 496-498

configuring, 306-308

match statements, 308-309

mutual redistribution with route maps, 313-314

mutual redistribution without route filtering, 312-313

route redistribution, 310-318

sequence numbers, 307

set statements, 308-309

statements, 307

route maps, 305-318

applications for, 305-306

BGP filtering, 496-498

configuring, 306-308

match statements, 308-309

route redistribution, 310-318

mutual redistribution with route maps, 313-314

mutual redistribution without route filtering, 312-313

sequence numbers, 307

set statements, 308-309

statements, 307

route redistribution, 270-292

advertising from BGP into an IGP, 681

advertising networks into BGP, 680-681

caveats, 319-320

configuring using route maps, 310-318

defining, 270-271

EIGRP for IPv6 routes, redistributing into OSPFv3, 285-287

EIGRP routes, redistributing into OSPF, 281-284, 304-305

multipoint redistribution, 288-289

one-point redistribution, 287-288

OSPFv2 routes, redistributing into EIGRP, 276-279, 303-304

OSPFv3 routes, redistributing into EIGRP for IPv6, 279-281

outbound redistribution, 271

planning for, 271

route tagging, 318-320

routing loops, preventing, 291

seed metrics, 272-273

troubleshooting, 289-291

verifying, 292

route reflectors, 687-695

clusters, 689

configuring, 694-695

migration tips, 692-694

verifying, 695

route summarization, 9, 662-667

automatic summarization, configuring, 111-116

BGP, 671-679

aggregate addresses, 671-673

aggregate-address command, 677-679

CIDR, 671-673

network boundary summarization, 673-674

network command, 674-676

in Cisco routers, 666

- EIGRP, 109
 - configuring, 110-120*
 - summary routes, determining, 116-120*
- implementing, 666
- manual summarization, 50
 - configuring, 116-120*
 - OSPF, 157*
- OSPF, 216-225
 - on ABRs, 223-224*
 - on ASBRs, 224-225*
 - benefits of, 217*
 - configuring, 218-223*
- in routing protocols, 666-667
- summary routes, verifying, 118-120
- in VLSM-designed networks, 665
- route tagging, 318-320
- route-map command, 306
- router eigrp autonomous-system-number command, 65-66
- router IDs
 - changing, 164-165
 - configuring, 163
 - selecting, 163
- router LSAs, 189, 192-196
- router ospf process-id command, 162
- Router Priority field (OSPF Hello packets), 174-175
- router rip command, 46
- routers. *See also* Cisco routers
 - ABRs, 159
 - route summarization, 223-224*
 - ASBRs, 159
 - route summarization, 224-225*
 - backbone routers, 159
 - BDR election process, 170-175
 - boundary routers, 268
 - CE routers, 75
 - control plane, 527
 - data plane, 527
 - DHCP configuration, 384-385
 - DR election process, 170-175
 - management plane, 527
 - PE routers, 75
 - router security policy, 530-531
- routing
 - over GRE, 34-35
 - over MPLS VPNs, 32-34
 - over the Internet, 18
- routing loops, preventing, 291
- routing protocols
 - asymmetric routing, 5
 - authentication, 570-576
 - purpose of, 570-573*
 - time-based key chains, 574-575*
 - believability, 269
 - classless routing protocols, 45
 - convergence, 8-9
 - EIGRP, 60*
 - OSPF, 157*
 - distance vector protocols, 7
 - RIP, 43-45*
 - RIPng, 43-54*
 - RIPv2, 45-47*
 - EGPs, 6
 - IGPs, 6
 - IPv4 multicast addresses, 12
 - link-state protocols, 7
 - multiprotocol routing, 267-270
 - administrative distance, 269-270*
 - reasons for, 269*
 - solutions for, 270*
 - path vector protocols, 7
 - role of, 3-5
 - route summarization, 9, 666-667
 - scalability, 10
 - selecting, 5
- routing table (EIGRP for IPv6), displaying, 133
- RS (Router Solicitation) messages, 14
- RTP (Reliable Transport Protocol), 62

S

scalability

- IPv4 network addressing, 650-651
- OSPF, 215-239
 - route summarization, 216-225*
- of routing protocols, 10

scheduling Cisco IOS IP SLA operations, 359

SCP (Secure Copy), 565-567

security

- AAA, 536-543
 - accounting, 542*
 - authorization, 542*
 - local authentication, 538-539*
 - RADIUS, 536-541*
 - TACACS+, 536-538, 541-542*
- authentication
 - BGP, 593-596*
 - EIGRP, 576-583*
 - bashing authentication, 572-573*
 - OSPF, 583-593*
 - routing protocol authentication, 570-576*
 - time-based key chains, 574-575*

- conditional debugging, 568-569
- configuration backups, 563-565
- disabling unused services, 567
- encrypting passwords, 531-536
- EVN, 601-602
- infrastructure ACLs, 547-549
- IPv6 connectivity, 409
- logging, 551-552
- management plane, 529-530
- NTP, 552-558
 - in IPv6 environments, 557*
 - modes, 552*
 - SNTP, 557-558*
 - versions, 556-557*
- router security policy, 530-531

- SCP, 565-567
- SNMP, 558-563
 - MIB, 559*
 - versions, 559*

- SSH, 543-547
- uRPF, 549-551

seed metrics, 272-273

- configuring, 273
- default seed metrics, 273-275

selecting

- between inter-area and intra-area routes, 215
- router ID, 163
- routing protocols, 5

sequence numbers, 307

set commands (PBR), 347

set community command, 682-683

set statements (route maps), 308-309

show ip bgp neighbors command, 447-448

show ip cef command, 342

show ip eigrp interfaces command, 69

show ip eigrp neighbors command, 66-68

show ip eigrp traffic command, 79

show ip interface command, 343

show ip ospf database command, 192

show ip ospf neighbor command, 165

show ip ospf route command, 168

show ip protocols command, 46, 147

show ipv6 eigrp topology command, 133

show ipv6 interface brief command, 48

show ipv6 protocols command, 49

show ipv6 rip command, 53-54

SIA (stuck-in-active) state, 108-109

single-homed connections, 375

- DHCP, 382-383
 - configuring, 384-385*

disadvantages of, 410

IPv6 address assignment, 398-405

manual assignment, 399

SLAAC, 401

- stateful DHCPv6*, 404
- stateless DHCPv6*, 403-404
- NAT, 385-393
 - provider-assigned IPv4 address, configuring, 381-382
- SLAAC (stateless address autoconfiguration), 401
- slow start, 42
- SNMP (Simple Network Management Protocol), 558-563
 - MIB, 559
 - versions, 559
- SNTP (Simple NTP), 557-558
- soft resets, 465-466
- solutions for multiprotocol routing, 270
- sources (IP SLAs), 354-355
- sourcing BGP from loopback address, 461-463
- speakers (BGP), 435
- speed of convergence
 - EIGRP, 60
 - influencing, 8-9
- SPF (Shortest Path First) algorithm, 7, 156, 207-208
- split horizon, 16
- SSH (Secure Shell), 543-547
- standard IP ACLs, 626-634
 - configuring, 629-633
 - location of, 633-634
 - wildcard masks, 628-629
- stateful DHCPv6, 404
- stateless DHCPv6, 403-404
- statements (route map), 307
- static default routes, configuring, 22
- Static NAT, 387
 - configuring, 388
- static routing, 19-22
 - disadvantages of, 20
 - hub-and-spoke topology, 20
 - IPv4 static routes
 - CEF, 21
 - ip route command*, 20-21

- IPv4 static routes, configuring, 20-21
- static default routes, configuring, 22
- stub areas**
 - configuring, 229-234
 - cost of default route, 236-238
 - NSSAs, 238-239
 - totally stubby areas, 234-236
- stub routing**,
 - connected routes, advertising, 103-105
 - receive-only stub routers, 105-107
- subnet masks**, 620-625, 653-654
 - calculating, 621-624
 - representing with prefixes, 624-625
- subnets, determining summary routes**, 116-120
- suboptimal routing**, 289-291
- successor routes**, 81, 82-87
- summary LSAs**, 189, 197-199
- summary routes**
 - determining, 116-120
 - verifying, 118-120
- synchronizing LSDB**, 204-205
 - on multiaccess networks, 206
- system logging**, 530

T

- TACACS+**, 536-538, 541-542
 - limitations of, 542-543
- targets (IP SLAs)**, 354-355
- TCP (Transmission Control Protocol)**
 - assigned port numbers, 639-640
 - latency, 42
 - MSS, 40-41
 - port names, 639
 - slow start, 42
- time-based ACLs**, 644
- time-based key chains**, 574-575

timers

EIGRP

*active timer, 108**manipulating, 73*

OSPF, manipulating, 179-182

topology table, EIGRP, 63

building, 76-94

feasibility condition, 91

feasible successors, 81-87

RD, 92-94

redistribution, 88

routing knowledge exchange, 88

totally stubby areas, 234-236**tracking objects (IP SLAs), configuring, 360****traffic**

broadcast, 11

multicast, 11

unicast, 11

troubleshooting

BGP neighbor states, 458-460

mismatched MTUs, 177-179

mismatched OSPF hello timers, 180

route redistribution, 289-291

trustworthiness of routing protocols, 269**tunneling VPNs, 32**

DMVPN, 35-36

NHRP, 37-38

GRE, routing over, 34-35

IPsec, 39-40

two-way redistribution, 288**Type 1 LSAs, 189, 192-196****Type 1 packets, 160****Type 2 LSAs, 189, 196-197****Type 2 packets, 160****Type 3 LSAs, 189, 197-199****Type 3 packets, 160****Type 4 LSAs, 189, 199-201****Type 4 packets, 161****Type 5 LSAs, 189, 201-203****Type 5 packets, 161****Type 6 LSAs, 189****Type 7 LSAs, 189****Type 8 LSAs, 189****Type 9 LSAs, 189****Type 10 LSAs, 189****Type 11 LSAs, 189**

type codes for BGP attributes, 473

type names (ICMP), 637

U

UDP (User Datagram Protocol)

port names, 641

reserved port numbers, 642

undebug all command, 102**unequal metric load balancing, 62**

EIGRP, 126-128

unicast addresses, 11

global unicast addresses, 13

unique local addresses, 14**unspecified IPv6 address, 14****unused services, disabling, 567****update messages (BGP), 433****uRPF (Unicast Reverse Path Forwarding),
549-551**

V

VCs (virtual circuits), 29**verifying**

ACL configuration, 647-648

CEF, 335-341

EIGRP neighbor relationships, 66-69

iBGP sessions, 449-450

MP-BGP, 507-518

NVI, 396-397

OSPF neighbor relationships, 165

path control with Cisco IOS IP SLAs,
360-361

PBR, 348

- prefix lists, 301-302
- redistributed static routes, 122-123
- RIPng database, 53-54
- route redistribution, 292
 - redistributed OSPFv3 routes, 281*
- route reflectors, 695
- summary routes, 118-120
- versions
 - of NTP, 556-557
 - of SNMP, 559
- virtual links, 225-229
 - configuring, 227-229
 - removing, 227
- VLSM**
 - calculation examples, 656-662
 - EIGRP support for, 61
 - implementing in scalable networks, 654-656
 - route summarization, 665
 - subnet masks, 653-654
- VPLS (Virtual Private LAN Service), 32**
- VPNs (virtual private networks), 31-40**
 - hybrid VPNs, 32
 - mGRE, 36
 - MPLS VPNs, 31-32
 - tunneling VPNs, 32
 - DMVPN, 35-36*
 - GRE, 34-35*
 - IPsec, 39-40*
- VPWS (Virtual Private Wire Service), 32**
- VRF (Virtual Routing and Forwarding), 597**
- VRF-Lite, 597-601**
- VRRP (Virtual Router Redundancy Protocol), 412**
- vtv access, restricting, 645-647
 - passwords, encrypting, 533-536

W-X-Y-Z

- WANs, Frame Relay, 28
- Weight attribute, 478-480
- well-known attributes (BGP), 471-472
- when to use BGP, 433-434
- wide metric (EIGRP), 90
- wildcard masks, 628-629