



IP Routing on Cisco IOS®, IOS XE, and IOS XR

An Essential Guide to Understanding
and Implementing IP Routing Protocols

ciscopress.com

Brad Edgeworth, CCIE No. 31574
Aaron Foss, CCIE No. 18761
Ramiro Garza Rios, CCIE No. 15469

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

IP Routing on Cisco IOS, IOS XE, and IOS XR

An Essential Guide to Understanding and Implementing IP Routing Protocols

Brad Edgeworth, CCIE No. 31574
Aaron Foss, CCIE No.18761
Ramiro Garza Rios, CCIE No. 15469

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

IP Routing on Cisco IOS, IOS XE, and IOS XR

Brad Edgeworth, Aaron Foss, Ramiro Garza Rios

Copyright © 2015 Cisco Systems, Inc.

Cisco Press logo is a trademark of 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 December 2014

Library of Congress Control Number: 2014957562

ISBN-13: 978-1-58714-423-3

ISBN-10: 1-58714-423-9

Warning and Disclaimer

This book is designed to provide information about Cisco IOS, IOS XE, and IOS XR. 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.

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.

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.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Business Operation Manager,
Cisco Press: Jan Cornelssen

Acquisitions Editor: Denise Lincoln

Managing Editor: Sandra Schroeder

Senior Development Editor: Christopher Cleveland

Project Editor: Seth Kerney

Copy Editor: Keith Cline

Technical Editors: Richard Furr, Pete Lumbis

Editorial Assistant: Vanessa Evans

Book Designer: Gary Adair

Cover Designer: Mark Shirar

Composition: Trina Wurst

Indexer: Heather McNeill

Proofreader: Apostrophe Editing Services



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco Logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, 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

Brad Edgeworth, CCIE No. 31574 (R&S & SP), has been with Cisco since 2011, working as a Systems Engineer and a Technical Leader. Brad is a distinguished speaker at Cisco Live, where he has presented on IOS XR. Before joining Cisco, Brad worked as a network architect and consultant for various Fortune 500 companies. Brad's other certifications include Cisco Certified Design Professional (CCDP) and Microsoft Certified Systems Engineer (MCSE). Brad has been working in the IT field for the past 18 years, with an emphasis on enterprise and service provider environments from an architectural and operational perspective. Brad holds a bachelor of arts degree in computer systems management from St. Edward's University in Austin, Texas.

Aaron Foss, CCIE No. 18761 (R&S & SP), is a High Touch Engineer with Cisco's Focused Technical Support (FTS) organization. He works with large service providers to troubleshoot issues relating to Multiprotocol Label Switching (MPLS), quality of service (QoS), and IP routing protocols. Aaron has more than 15 years of experience designing, deploying, and troubleshooting IP networks. He holds a bachelor of science degree in management information systems from Rochester Institute of Technology.

Ramiro Garza Rios, CCIE No. 15469 (R&S, SP, and Security), is a Senior Network Consulting Engineer with Cisco Advanced Services. His current role consists of planning, designing, implementing, and optimizing next-generation (NGN) service provider networks in the United States. He has been with Cisco for more than 8 years and has 14 years of networking industry experience. Before joining Cisco, Ramiro was a Network Consulting and Presales Engineer for a Cisco Gold Partner in Mexico, where he was involved in the planning, design, and implementation of many enterprise and service provider networks. He holds a bachelor of science degree in electronic engineering from the Instituto Tecnológico de Reynosa and lives with his wife and four children in Cary, North Carolina.

About the Technical Reviewers

Richard Furr, CCIE No.9173 (R&S & SP), is a Technical Leader with Cisco's Technical Assistance Center (TAC). For the last 13 years, Richard has worked for Cisco TAC and High Touch Technical Support (HTTS) organizations, supporting service providers and large enterprise environments with a focus on troubleshooting routing protocols, MPLS, IP multicast and QoS.

Pete Lumbis, CCIE No. 28677 (R&S) and CCDE 20120003, is an expert in routing technologies including Border Gateway Protocol (BGP), MPLS, and multicast. He spent five years working in the Cisco TAC as the Routing Protocols Escalation Engineer supporting all of Cisco's customers. Most recently, Pete is focused on network design and architecture at Microsoft. Pete has been a distinguished speaker at Cisco Live on routing fast convergence and IOS routing internals.

Dedications

This book is dedicated to my loving wife Tanya, who has endured and supported me through all my endeavors.

-Brad

I would like to dedicate this book to my supportive wife, Anne, and to my children, Ashley, Benny, and Clara, for giving up some weekend time with Dad so that I could write this book.

-Aaron

I would like to dedicate this book to my wonderful and beautiful wife, Mariana, and to my children Ramiro, Frinee, Felix, and Lucia for their love, patience, sacrifice, and support while writing this book.

To my parents, Ramiro Garza and Blanca Dolores Rios, for their continued support, love, encouragement, guidance, and wisdom.

And most importantly, I would like to thank God for all His blessings in my life.

-Ramiro

Acknowledgments

Brad Edgeworth:

A special thank you goes to Norm Dunn, Jocelyn Lau, Brett Bartow, and Denise Lincoln for making this book possible.

A debt of gratitude goes to my co-authors, Aaron and Ramiro. You accepted the challenge of helping me write this book. Little did you know that this project would become your second job. Some of the book's best chapters were supposed to be small, but exploded in to mini-novels to cover the topic properly. Your knowledge and dedication to this project are appreciated more than you will ever know.

To our technical editors, Richard and Pete: Thank you for finding all of our mistakes. In addition to your technical accuracy, your insight into the technologies needed by Cisco customers versus crazy ninja router tricks has kept the size of the book manageable and the content relevant.

Aaron, Ramiro, and I want to thank the Cisco Press team for their assistance and insight throughout this project. Chris Cleveland, you have been a pleasure to work with, and your attention to detail is simply amazing. It has been an educational experience for the three of us.

A special thanks to the Cisco HTTS RP and IOS XR teams, who continuously educate those about routing protocols. A special recognition to Hunter, Yigal, and Jimmy—you guys are rock stars!

Many people within Cisco have provided feedback and suggestions to make this a great book. Thanks to all who have helped in the process, especially Umair Arshad, Heather Bunch, Luc de Ghein, David Roehsler, Faraz Shamim, Craig Smith, and Mobeen Tahir.

Aaron Foss:

I would like to thank my co-authors Brad and Ramiro for their amazing collaboration on this project. Brad, you have an extraordinary determination and drive that I admire greatly; and Ramiro, your technical knowledge and ability to make us laugh throughout the process of writing this book was much appreciated.

Finally, I want to acknowledge my manager, Zulfiqar Ahmed, for supporting me and encouraging me to undertake this book endeavor.

Ramiro Garza Rios:

I would like to thank God for giving me the opportunity to work on this book. I would like to acknowledge my co-author Brad for the inception of this book and for being persistent until it became a reality. I would also like to acknowledge both of my co-authors, Aaron and Brad, for the great teamwork, dedication, and valuable input provided throughout the project.

Contents at a Glance

Foreword xxviii

Introduction xxix

Part I Network Fundamentals

Chapter 1 Introduction to the Operating Systems 1

Chapter 2 IPv4 Addressing 29

Chapter 3 How a Router Works 67

Part II Routing Protocols

Chapter 4 Static Routing 91

Chapter 5 EIGRP 125

Chapter 6 OSPF 191

Chapter 7 Advanced OSPF 241

Chapter 8 IS-IS 315

Chapter 9 Advanced IS-IS 373

Chapter 10 Border Gateway Protocol (BGP) 407

Part III Advanced Routing Techniques

Chapter 11 Route Maps and Route Policy 467

Chapter 12 Advanced Route Manipulation 521

Chapter 13 Route Redistribution 551

Part IV Advanced BGP

Chapter 14 Advanced BGP 609

Chapter 15 BGP Best Path Selection 671

Part V Multicast

Chapter 16 IPv4 Multicast Routing 745

Chapter 17 Advanced IPv4 Multicast Routing 811

Part VI IPv6

Chapter 18 IPv6 Addressing 893

Chapter 19 IPv6 Routing 941

Chapter 20 IPv6 Multicast Routing 1007

Part VII High Availability

Chapter 21 High Availability Online

Appendixes

Appendix A Decimal to Hex to Binary Values Online

Appendix B BGP Attributes Online

Contents

Part I Network Fundamentals

Chapter 1 Introduction to the Operating Systems 1

IOS, IOS XE, and IOS XR Software Architecture 1

IOS 1

Kernel and OS Scheduling 2

Memory Management 2

Software Packaging 2

IOS XE 4

Kernel and OS Scheduling 4

Memory Management 4

IOS XR 5

Kernel and OS Scheduling 5

Memory Management 5

Software Packaging 6

Debugging 8

CLI and Configuration 8

IOS 9

User Mode 9

Privileged Mode 10

Global Configuration Mode 10

Configuration Archiving 11

Configuration Replace 13

IOS XR 14

Viewing Changes in the SysDB 17

Commit Label 18

Commit Replace 19

Failed Commits 20

Configuration Rollback 21

Commit Confirmed 22

Multiple Commit Options 23

Loading Files for Changes 24

Hierarchical Configuration 24

PWD 26

Root 26

Summary 27

References in This Chapter 27

Chapter 2 IPv4 Addressing 29

- IP Fundamentals 29
- Understanding Binary 31
- Address Classes 34
- Subnet Masks and Subnetting 35
 - Subnet Mask Purpose 36
 - Calculating Usable IP Addresses 37
 - Network Prefix Notation 38
 - Subnetting 38
 - Subnet Field* 39
 - Subnet Math* 41
 - Subnet Design 46
- Classless Interdomain Routing 49
 - Classful Versus Classless Routing 50
 - Classful Routing* 50
 - Classless Routing* 53
 - Variable-Length Subnet Masks 55
 - Summarization 56
- Private IP Addressing 58
- Special IP Addresses 59
- IPv4 Address Configuration 60
- Wildcard Subnet Masks 62
- Summary 64
- References in This Chapter 65

Chapter 3 How a Router Works 67

- IP Routing 67
 - Distance Vector Algorithms 69
 - Enhanced Distance Vector Algorithm 70
 - Link-State Algorithms 70
 - Path Vector Algorithm 71
 - Routing Table 72
 - Prefix Length* 73
 - Administrative Distance* 73
 - Metrics* 75
 - Virtual Routing and Forwarding 76
- IP Packet Switching 83
 - Process Switching 84

Cisco Express Forwarding	85
<i>Software CEF</i>	87
<i>Hardware CEF</i>	88
Planes of Operation	89
References in This Chapter	90

Part II Routing Protocols

Chapter 4 Static Routing 91

Connected Networks	91
Secondary Connected Networks	94
Static Routing Fundamentals	94
Point-to-Point Interfaces	96
Broadcast Interfaces	98
Default Route	99
Floating Static Routing	103
Recursive Lookup	105
Multihop Routing	108
<i>Single Recursive Lookup</i>	108
<i>Multiple Recursive Lookups</i>	109
Problems with Static Route Recursion	112
Null Interface	116
Static VRF Routes	121
References in This Chapter	124

Chapter 5 EIGRP 125

EIGRP Fundamentals	125
EIGRP Neighbors	126
Inter-Router Communication	126
Forming EIGRP Neighbors	128
Classic EIGRP Autonomous System Configuration	131
IOS network Statement	132
IOS XR	134
Passive Interfaces	134
Sample Topology and Configuration	134
Confirmation of Interfaces	136
Verification of EIGRP Neighbor Adjacencies	139
Display of Installed EIGRP Routes	140
Router ID	141
EIGRP Terminology	142

Topology Table	143
Path Metric Calculation	145
Custom K Values	148
Interface Delay Settings	149
Load Balancing	151
EIGRP Wide Metrics	153
Failure Detection and Timers	155
Convergence	156
Stuck in Active	159
Stub	160
Design Considerations with EIGRP Stubs	164
Summarization	166
Interface-Specific Summarization	166
Summarization Metrics	171
Advertising a Default Route	172
Automatic Summarization	172
Authentication	174
Enabling Authentication on the interface	174
Key Chain Configuration	174
WAN Considerations	177
<i>IP Bandwidth Percent</i>	177
Split Horizon	179
Next-Hop Self	182
EIGRP Named Configuration	184
Address Family Instance Configuration	185
Address Family Interface Configuration	186
Address Family Topology Configuration	188
Summary	189
References in This Chapter	189

Chapter 6 OSPF 191

OSPF Fundamentals	191
Inter-Router Communication	193
OSPF Hello Packets	194
Router ID	195
Neighbors	196
Forming OSPF Neighbor Adjacencies	197

Basic OSPF Configuration	202
<i>IOS network Statement</i>	202
<i>IOS Interface Specific</i>	204
IOS XR	205
Passive Interfaces	205
Sample Topology and Configuration	206
Confirmation of Interfaces	208
Verification of OSPF Neighbor Adjacencies	209
Verification of OSPF Routes	211
Designated Router and Backup Designated Router	212
Designated Router Elections	214
DR and BDR Placement	216
Failure Detection	219
Hello Timer	219
Dead Interval Timer	219
Verifying OSPF Timers	220
OSPF Fast Packet Hellos	220
OSPF Network Types	221
Broadcast	221
Non-Broadcast	222
Point-to-Point Networks	224
Point-to-Multipoint Networks	225
Loopback Networks	229
Review of OSPF Network Types	231
OSPF Adjacency with Different OSPF Network Types	231
Link Costs	235
Authentication	236
IOS Support for OSPF Authentication	236
IOS XR Support for OSPF Authentication	236
Summary	239
References in This Chapter	240
Chapter 7	Advanced OSPF 241
Areas	241
Area ID	245
OSPF Route Types	246
<i>External OSPF Routes</i>	247
Link-State Announcements	249
LSA Age and Flooding	251

LSA Types	251
<i>LSA Type 1: Router Link</i>	252
<i>LSA Type 2: Network Link</i>	257
<i>LSA Type 3: Summary Link</i>	259
<i>LSA Type 5: External Routes</i>	263
<i>LSA Type 4: ASBR Summary</i>	265
<i>LSA Type 7: NSSA External Summary</i>	268
<i>LSA Type Summary</i>	270
OSPF Path Selection	270
Intra-Area Routes	271
Interarea Routes	272
External Route Selection	272
E1 and N1 External Routes	273
E2 and N2 External Routes	273
Equal Cost Multi-Path	274
Summarization of Routes	274
Interarea Summarization	276
External Summarization	280
Default Route	283
OSPF Stubby Areas	286
Stub Areas	286
Totally Stubby Areas	289
Not-So-Stubby Areas	292
Totally NSSA Areas	295
Virtual Links	298
Discontiguous Network	301
Multi-Area Adjacency	304
Prefix Suppression	308
Summary	313
References in This Chapter	314

Chapter 8 IS-IS 315

IS-IS Fundamentals	315
Areas	318
OSI Addressing	319
<i>Inter-Domain Part</i>	320
<i>Domain Specific Part</i>	321
<i>NET Addressing</i>	322

Inter-Router Communication	323
IS Protocol Header	325
TLVs	326
IS PDU Addressing	326
Hello Packets	327
Link-State Packets	329
<i>LSP Lifetime</i>	329
<i>LSP ID</i>	330
<i>LSP Sequence</i>	331
<i>Attribute Fields</i>	331
<i>LSP Packet and TLVs</i>	332
IS-IS Neighbor	333
<i>Ethernet</i>	333
<i>Point-to-Point</i>	338
Basic IS-IS Configuration	340
IOS	340
IOS XR	340
Sample Topology and Configuration	341
Confirmation of IS-IS Interfaces	343
Verification of IS-IS Neighbor Adjacencies	346
Verification of IS-IS Routes	347
Designated Intermediate System	348
DIS Elections	351
DIS Placement	352
Point-to-Point Adjacency on Broadcast Media	353
Link State Packet Database	355
Viewing the LSPDB	356
<i>Non-Pseudonode LSPs</i>	357
<i>Pseudonode LSPs</i>	358
Building the Topology	359
Viewing the Topology	360
SPF Calculations	361
Passive Interfaces	362
Removal of Hello Padding	364
Failure Detection	366
Hello Timer	366
Hello Multiplier and Holding Timer	367

Authentication	367
IS-IS Hello Authentication	367
IS-IS LSP Authentication	368
Summary	371
References in This Chapter	372

Chapter 9 Advanced IS-IS 373

Advanced IS-IS Routing	373
Route Leaking	377
Backbone Continuity	380
Loop Prevention	382
Router-Specific IS-IS Levels	384
Interface Specific IS-IS Levels	385
Path Selection	386
Equal Cost Multi-Path	387
Interface Metrics	387
Overload Bit	394
Summarization	396
Default Routes	400
Prefix Suppression	401
Summary	405
References in This Chapter	406

Chapter 10 Border Gateway Protocol (BGP) 407

BGP Fundamentals	408
Autonomous System Numbers	408
Path Attributes	409
Loop Prevention	409
Address Families	410
Inter-Router Communication	410
Open Messages	412
<i>Hold Time</i>	412
<i>BGP Identifier</i>	413
Keepalive Messages	413
Update Messages	413
Notification Messages	414
BGP Sessions	415

BGP Neighbor States	415
Idle State	415
Connect State	415
Active State	416
OpenSent State	416
OpenConfirm State	417
Established State	417
Basic BGP Configuration	418
IOS	419
IOS XR	420
Verification of BGP Sessions	421
Prefix Advertisement	425
Receiving and Viewing Routes	427
iBGP	431
iBGP Full-Mesh Requirement	432
Peering via Loopback Addresses	433
eBGP	438
eBGP and iBGP Topologies	442
Next-Hop Manipulation	444
iBGP Scalability	446
Route Reflectors	446
<i>Loop Prevention in Route Reflectors</i>	451
<i>Out-of-Band Route Reflectors</i>	453
Confederations	453
Failure Detection	459
Security	459
eBGP Multihop	459
TTL Security	461
Summary	463
References in This Chapter	465

Part III Advanced Routing Techniques

Chapter 11 Route Maps and Route Policy 467

Access Control Lists	467
Standard ACLs	468
Extended ACLs	469
<i>IGP Network Selection</i>	469
<i>BGP Network Selection</i>	470

- Prefix Matching 471
 - Prefix Lists 473
 - Prefix Sets 474
 - Regular Expressions 475
 - _ (Underscore)* 477
 - ^ (Caret)* 478
 - \$ (Dollar Sign)* 478
 - [] (Brackets)* 479
 - (Hyphen)* 479
 - [^] (Caret in Brackets)* 480
 - () (Parentheses and | Pipe)* 480
 - .* (Period) 481
 - +* (Plus Sign) 481
 - ?* (Question Mark) 481
 - ** (Asterisk) 482
 - Looking Glass and Route Servers 483
- AS_Path Access List 484
- IOS XR AS_Path Selection Options 484
 - is-local 485
 - length 485
 - unique-length 486
 - passes-through 486
 - neighbor-is 487
 - originates-from 487
 - AS Path Set 488
- Route Maps 488
 - Conditional Matching 490
 - Multiple Conditional Match Conditions* 491
 - Complex Matching* 491
 - Optional Actions 492
 - Continue 493
 - Route Map Examples 494
- Routing Policy Language 496
 - Route Policy Structure 496
 - Match Statements 497
 - Attribute Modification 498
 - Common Route Policy Structure 499

- Boolean Operators 504
 - Negation* 504
 - Conjunction* 504
 - Disjunction* 505
 - Order of Processing* 505
- Comparing Prefix Sets to Prefix Lists 506
- Parameterization 507
- Route Policy Nesting 510
- Original Value 511
- Editors 512
- RPL Examples 513
- RPL Verification 515
 - Redistribution RPL Verification* 516
 - BGP RPL Verification* 517
- References in This Chapter 519

Chapter 12 Advanced Route Manipulation 521

- Conditional Routing of Packets 521
 - Policy-Based Routing Configuration 522
 - Access-List-Based Forwarding Configuration 523
 - Local PBR 525
- Administrative Distance 526
 - Modifying EIGRP AD 528
 - Modifying OSPF AD 529
 - Modifying IS-IS AD 531
 - Modifying BGP AD 532
- Route Filtering and Manipulation 534
 - EIGRP Filtering by Prefix 534
 - EIGRP Filtering by Hop Count 538
 - EIGRP Offset Lists 538
 - OSPF Filtering (Local) 541
 - OSPF Filtering (Area) 543
 - IS-IS Filtering (Local) 546
 - BGP Filtering 546
 - Clearing BGP Connections 549
- Summary 550
- References in This Chapter 550

Chapter 13 Route Redistribution 551

- Redistribution Basics 553
 - Redistribution Is Not Transitive 553
 - Sequential Protocol Redistribution 555
 - Routes Must Exist in the RIB 555
 - Metrics 558
- Protocol-Specific Configuration 558
 - Source-Specific Behaviors 560
 - Connected Networks* 561
 - IS-IS* 561
 - BGP* 562
 - Destination-Specific Behaviors 563
 - EIGRP* 563
 - OSPF* 568
 - IS-IS* 576
 - BGP* 580
- Challenges with Redistribution 582
 - Route Feedback 583
 - Suboptimal Routing 584
 - Invalid Routing Tables 589
 - Routing Loops 590
 - Methods to Avoid Routing Loops 593
 - Prefix Filtering* 593
 - Tagging* 595
 - Increase Seed Metrics* 598
 - Administrative Distance* 601
 - Summarization on Redistributing Router* 603
 - Solutions to Redistribution Challenges 606
- Summary 606
- References in This Chapter 607

Part IV Advanced BGP**Chapter 14 Advanced BGP 609**

- BGP Communities 609
 - Enabling BGP Community Support 610
 - Well-Known Communities 611
 - Internet* 611
 - No_Export* 611

<i>No_Advertise</i>	614
<i>No_Export_SubConfed</i>	617
Conditionally Matching BGP Communities	620
<i>Community Set</i>	621
<i>Inline</i>	622
<i>Setting Private BGP Communities</i>	625
Route Summarization	628
Aggregate Address	629
Flexible Route Suppression	632
<i>Selective Prefix Suppression</i>	632
<i>Leaking Suppressed Routes</i>	634
Atomic Aggregate	637
Route Aggregation with AS_SET	639
Route Aggregation with Selective Advertisement of AS_Set	641
Default Route Advertisement	643
Default Route Advertisement Per Neighbor	644
Conditional Route Advertisement	645
Outbound Route Filtering	647
Backdoor Networks	649
Maximum Autonomous System	652
Maximum Prefix	654
Remove Private Autonomous System	656
Allow Autonomous System	658
Local Autonomous System	660
Configuration Scalability	664
IOS Peer Groups	664
IOS Peer Templates	665
IOS XR Configuration Templates	667
Summary	668
References in This Chapter	669
Chapter 15 BGP Best Path Selection	671
BGP Best Path Overview	672
Weight	673
Local Preference	679
Locally Originated via Network or Aggregate Advertisement	684
Accumulated Interior Gateway Protocol	686
Shortest AS_Path	694

Origin Type	700
Multi-Exit Discriminator	704
<i>Missing MED behavior</i>	709
<i>Always Compare Med</i>	711
<i>BGP Deterministic MED</i>	713
eBGP over iBGP	714
Lowest IGP Metric	718
Prefer the Oldest EBGP Path	720
Router ID	720
Minimum Cluster List Length	721
Lowest Neighbor Address	722
BGP ECMP	723
eBGP and iBGP Multipath	723
eiBGP Multipath	726
R1	729
R2	730
XR3	730
XR4	730
XR5	731
AS_Path Relax	731
Suboptimal Routing with Route Reflectors	733
Additional Route Reflector	734
Shadow Route Reflector	735
Shadow Session Route Reflector	738
BGP Add-Path	739
Summary	742
Further Reading	743

Part V Multicast

Chapter 16 IPv4 Multicast Routing 745

Multicast Fundamentals	745
Multicast Addressing	749
Layer 2 Multicast Addresses	752
Internet Group Management Protocol	753
IGMP Snooping	753
IGMPv2	756
IGMPv3	759

Multicast Distribution Trees	759
Source Trees	759
Shared Trees	760
Protocol Independent Multicast	762
PIM Dense Mode	765
PIM Sparse Mode	768
<i>PIM Shared and Source Path Trees</i>	768
<i>Shared Tree Join</i>	769
<i>Source Registration</i>	769
<i>PIM SPT Switchover</i>	771
<i>Designated Routers</i>	772
Rendezvous Points	772
Static RP	773
Auto-RP	773
<i>Candidate RPs</i>	773
<i>RP Mapping Agents</i>	774
PIM Bootstrap Router	775
<i>Candidate RPs</i>	775
Reverse Path Forwarding	776
PIM Forwarder	778
Basic Multicast Configuration	780
Configure Rendezvous Points	783
<i>Static RP</i>	784
<i>Auto-RP</i>	785
<i>BSR</i>	786
Multicast Verification	787
Bidirectional PIM	802
Bidir-PIM Designated Forwarder	804
Summary	808
References in This Chapter	809
Chapter 17 Advanced IPv4 Multicast Routing	811
Interdomain Multicast Routing	811
Multiprotocol BGP	812
Multicast Source Discovery Protocol	817
<i>MSDP Source Active Message Types</i>	818
<i>SA Messages</i>	819

<i>Keepalive Messages</i>	819
<i>MSDP Peers</i>	822
<i>MSDP Verification</i>	828
<i>MSDP Stub Networks</i>	831
Rendezvous Point Redundancy	833
Auto-RP with Multiple RPs	835
<i>Auto-RP Group Filtering</i>	836
BSR with Multiple RPs	840
<i>BSR Group Filtering</i>	843
<i>BSR RP Hash Algorithm</i>	845
Static RP with Multiple RPs	846
Anycast RP	847
Source Specific Multicast	850
SSM Mapping	857
<i>DNS SSM Mapping</i>	857
<i>Static SSM Mapping</i>	860
Multicast Security	862
Auto-RP Scoping	862
Multicast Boundaries	863
<i>Administratively Scoped Boundaries</i>	863
<i>Auto-RP Multicast Boundaries</i>	865
<i>BSR Multicast Boundaries</i>	866
Auto-RP Cisco-RP-Announce Message Filtering	867
PIM-SM Source Registration Filtering	867
PIM-SM Accept RP	868
PIM Neighbor Control	869
PIM Register Rate Limit	870
Multicast Traffic Engineering	871
RPF Rules	871
Static Mroutes	872
MBGP	875
Static IGMP Joins	882
Multicast Troubleshooting	886
Mtrace	887
Summary	889
References in This Chapter	891

Part VI IPv6

Chapter 18 IPv6 Addressing 893

IPv6 Address Structure	893
Text Representation Address Abbreviation	895
IPv6 Hexadecimal to Binary Conversion	896
IPv6 Address Types	898
Unicast	898
Global Unicast	900
Unique Local Unicast	904
Link-Local Unicast	905
Anycast	906
Multicast	908
Special IPv6 Addresses	911
Neighbor Discovery Protocol	912
Router, Prefix, and Parameter Discovery.	913
Redirect	919
IPv6 Stateless Address Autoconfiguration	920
<i>Extended Unique Identifier</i>	920
<i>SLAAC Router Configuration</i>	921
<i>RA Options for DNS</i>	923
Stateless DHCPv6	924
<i>IOS Stateless DHCPv6 Configuration</i>	924
<i>IOS XR Stateless DHCPv6 Configuration</i>	925
<i>Stateless DHCPv6 Verification</i>	926
Stateful DHCPv6, Relay Agent, and Relay Proxy	926
<i>IOS Relay Agent Configuration</i>	927
<i>IOS Relay Agent Verification</i>	928
<i>IOS XR Proxy Agent Configuration</i>	928
<i>IOS XR Proxy Agent Verification</i>	929
<i>IOS Stateful DHCPv6 Server Configuration</i>	930
<i>IOS XR Stateful DHCPv6 Server Configuration</i>	931
<i>Stateful DHCPv6 Server Verification</i>	932
IPv6 Address Resolution and Neighbor Unreachability Detection	934
Duplicate Address Detection	937
Summary	938
References in This Chapter	939

Chapter 19 IPv6 Routing 941

- Static Routing 941
 - Static Route Configuration 942
 - Static Route Reference Chart for IPv6 943
- EIGRPv6 944
 - EIGRPv6 Inter-Router Communication 944
 - EIGRPv6 Configuration 945
 - IOS EIGRPv6 Autonomous System Configuration (Classic)* 945
 - IOS EIGRPv6 Hierarchical Configuration (Named Mode)* 946
 - IOS XR EIGRPv6 Configuration* 946
 - EIGRPv6 Verification* 947
 - Summarization 950
 - Default Route 952
 - Route Filtering 953
 - EIGRP Configuration Command Reference Chart for IPv6 954
- OSPFv3 956
 - OSPFv3 Inter-Router Communication 957
 - OSPFv3 Link-State Advertisement 958
 - OSPFv3 LSA Flooding Scope 959
 - OSPFv3 Configuration 960
 - IOS OSPFv3 Configuration* 960
 - IOS XR OSPFv3 Configuration* 961
 - OSPFv3 Verification* 962
 - OSPFv3 Authentication 970
 - OSPFv3 Multiple Instances 973
 - OSPFv3 Configuration Command Reference Chart for IPv6 975
- Integrated IS-IS for IPv6 977
 - IS-IS Inter-Router Communication 978
 - IS-IS Type-Length-Value 978
 - IS-IS Topology Modes 978
 - IS-IS Configuration 979
 - IOS Base Configuration* 979
 - IOS XR Base Configuration* 980
 - IOS Topology Mode Configuration* 981
 - IOS XR Topology Mode Configuration* 981
 - Verification* 985
 - IS-IS Configuration Reference Chart for IPv6 987

Multiprotocol BGP for IPv6	989
Inter-Router Communication	989
BGP Configuration	991
<i>IOS Base Configuration</i>	991
<i>IOS XR Base Configuration</i>	992
BGP Verification	993
IPv6 over IPv4 BGP Sessions	998
BGP Configuration Command Reference Chart for IPv6	1001
IPv6 Route Redistribution	1002
Summary	1006
References in This Chapter	1006

Chapter 20 IPv6 Multicast Routing 1007

IPv6 Multicast Routing Overview	1007
IPv6 Multicast Address Mapping into MAC Address	1009
Enabling Multicast Routing	1010
Multicast Listener Discovery	1010
Protocol Independent Multicast	1015
PIM Sparse Mode	1015
<i>Static RP</i>	1017
<i>Bootstrap Router</i>	1018
<i>Embedded RP</i>	1021
IPv6 Multicast Verification Commands	1024
Reverse Path Forwarding	1030
Multicast Boundary Scope	1032
PIM Source Specific Multicast	1033
Summary	1034
References in This Chapter	1035

Index 1037

Part VII High Availability

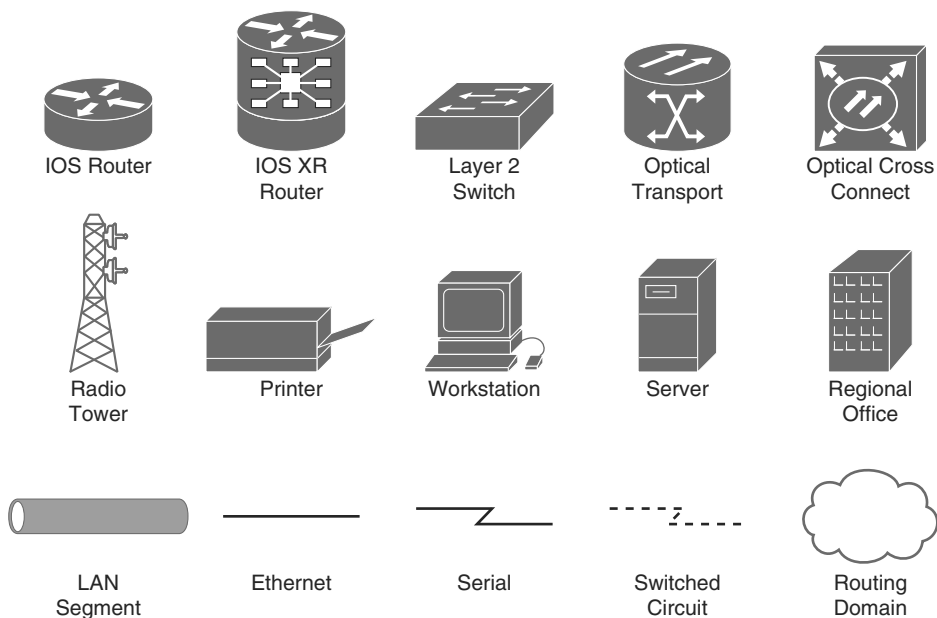
Chapter 21 High Availability Online

Appendixes

Appendix A Decimal to Hex to Binary Values Online

Appendix B BGP Attributes Online

Icons Used in This Book



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in Cisco's 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).
- Italics indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

Note This book covers multiple operating systems, and a differentiation of icons and router names indicate the appropriate operating system that is being referenced. IOS and IOS XE use router names like R1 and R2 and are referenced by the IOS router icon. IOS XR routers will use router names like XR1 and XR2 and are referenced by the IOS XR router icon.

Foreword

Service providers and even large, well-established enterprises, while they continue to sweat some legacy networking assets, they also realize the operational efficiencies gained by converging these disparate assets onto a common IP infrastructure. Furthermore, they generally understand the benefits of being able to offer new and innovative services with quicker time-to-market deployment with one unified converged IP backbone. Many service providers and enterprises have built out new IP backbones and are already realizing benefits of converged networking, but many have not, plus most have not realized the full potential of capability and revenue generation they can provide. This is where the need and demand for highly skilled IP network engineers becomes critical to the evolution of these IP network infrastructures, and where learning products such as Cisco career certifications and this Cisco Press resource shines in value.

This Cisco Press book is an excellent self-study resource to help aid candidates in preparing to pass exams associated with the CCNA Service Provider, CCNP Service Provider, and CCIE Service Provider career certifications. Exams associated with these Cisco certifications cover technology areas such as routing protocols (Enhanced Interior Gateway Routing Protocol [EIGRP], Open Shortest Path First [OSPF] Protocol, Intermediate System-to-Intermediate System [IS-IS] Protocol, and Border Gateway Protocol [BGP]), multicast, IPv6, and high availability. This book serves as a valuable aid in preparation in these areas. Furthermore, the book covers these topics across multiple Cisco operating system implementations, such as Cisco IOS and IOS XR, which are also covered within the noted Cisco career certifications. This resource can also aid in prepping candidates pursuing CCNA-CCNP-CCIE Routing and Switching certifications. Lastly, this book is valuable in general for learners looking to simply increase their technical understanding about how to configure routing protocols, multicast, IPv6, and high availability.

We hope and expect you'll find this book to be a valuable and frequently referenced technical aid, and a unique reference book for your personal library.

Norm Dunn

Senior Product Manager, Learning@Cisco

Global Product Management, Service Provider Portfolio

Cisco Systems, Inc.

Introduction

Within Cisco's Focused Technical Support (FTS) organization, a large number of questions about the IOS, IOS XE, and IOS XR operating systems are encountered on a daily basis. This book answers IP routing questions, in addition to covering the implementation and troubleshooting differences between the operating systems.

In alignment with the saying “a picture is worth a thousand words,” multiple illustrations are included in the chapters to explain the various concepts. All protocols are presented conceptually, with applicable illustrations, configurations, and appropriate output. The scope of this book evolved to include the IOS and IOS XE operating systems so that non-IOS XR users could benefit from the explanations on the routing protocols. The book's structure explains a concept, and then provides the configuration commands and verification of the feature in small, digestible nuggets of information.

This book's content was created in alignment with Learning@Cisco to address the demand for more efficient self-study content for the Cisco Career Service Provider Certifications.

This book encompasses content spread across multiple sources and presents them in a different perspective while covering updated standards and features that are found in enterprise and service provider networks.

Who Should Read This Book?

Network engineers, consultant, and students who want to understand the concepts and theory of EIGRP, OSPF, IS-IS, BGP, and multicast routing protocols on Cisco IOS, IOS XE, and IOS XR operating systems should read this book.

The book's content is relevant to network engineers in various stages of their career and knowledge. Every topic assumes minimal knowledge and explains the protocol from a ground-up perspective. For the advanced network engineers, relevant information on the routing protocol behavior is included. Differences in protocol behavior between IOS, IOS XE, and IOS XR are explicitly identified for each protocol.

How This Book Is Organized

Although this book could be read cover to cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover just the material that you need more work with. This book is organized into seven distinct sections.

Part I of the book provides a brief review of the operating systems, IP addressing, and networking fundamentals.

- **Chapter 1, “Introduction to the Operating Systems:”** This chapter provides a high-level comparison of the network operating system architectures. An overview of the CLI configuration is provided so that users are comfortable with logging in and configuring the routers.

- **Chapter 2, “IPv4 Addressing:”** This chapter explains the IPv4 addressing structure, the need for subnetting, and the techniques to differentiate a network address from a host address.
- **Chapter 3, “How a Router Works:”** This chapter explains the reasons for using a routing protocol, the types of routing protocols, and the logic a router uses for forwarding packets.

Part II of the book explains static routing, EIGRP, OSPF, IS-IS, and BGP routing protocols.

- **Chapter 4, “Static Routes:”** This chapter explains connected networks and static routes from the perspective of a router.
- **Chapter 5, “EIGRP:”** This chapter explains the EIGRP routing protocol and how distance vector routing protocols work.
- **Chapter 6, “OSPF:”** This chapter explains the basic fundamentals of the routing protocol, and its operational characteristics.
- **Chapter 7, “Advanced OSPF:”** This chapter explains the reason for breaking an OSPF routing domain into multiple areas, techniques for optimization, and how to determine the best path.
- **Chapter 8, “IS-IS:”** This chapter explains the history of the IS-IS routing protocol, along with the similarities and differences it has with OSPF.
- **Chapter 9, “Advanced IS-IS:”** This chapter explains multilevel routing in an IS-IS domain, optimization techniques, and the path selection process.
- **Chapter 10, “Border Gateway Protocol:”** This chapter explains the fundamental concepts of BGP sessions and route advertisement. The chapter covers the differences between external and internal peers.

Part III of the book explains the advanced routing concepts that involve routing policies and redistribution.

- **Chapter 11, “Route Maps and Route Policy Language:”** This chapter explains prerequisite concepts such as matching networks prefixes with an access control list (ACL), prefix list or BGP advertisements with regex queries. This chapter also explains how IOS and IOS XE route maps can manipulate traffic. The chapter then discusses how IOS XR’s route policy language was designed to provide clarity and scalability.
- **Chapter 12, “Advanced Route Manipulation:”** This chapter discusses policy-based routing, along with administrative distance manipulation, to modify route forwarding behavior. The chapter concludes by describing how to filter out specific routes from routing protocol participation.
- **Chapter 13, “Route Redistribution:”** This chapter explains the ability to inject network prefixes learned from one routing protocol into another routing protocol. The chapter provides a thorough coverage on the rules of redistribution, problems associated with mutual redistribution, and methods for remediation.

Part IV of the book revisits BGP and describes how prefix lists, route maps, route policies, and redistribution can be used for traffic engineering.

- **Chapter 14, “Advanced BGP:”** BGP communities, summarizations, and other router conservation techniques are explained in this chapter.
- **Chapter 15, “BGP Best Path Selection:”** This chapter provides a thorough explanation of the best path selection algorithm and the ramifications that the selection has for other routers in the autonomous system. BGP route reflectors are examined, along with suboptimal routing due to path information loss. The chapter concludes with an overview of the various techniques available to optimize traffic flows when using route reflectors.

Part V of the book explains multicast traffic, the benefits of multicast, and configuration.

- **Chapter 16, “IPv4 Multicast Routing:”** This chapter describes the benefits of multicast. Key multicast features such as Internet Group Management Protocol (IGMP), Protocol Independent Multicast (PIM), rendezvous points, multicast distribution trees are all discussed.
- **Chapter 17, “Advanced IPv4 Multicast Routing:”** Large multicast networks require additional features to provide scalability and reachability between routing domains and autonomous systems. This chapter explains the advanced features: Multicast Source Discovery Protocol (MSDP), Source Specific Multicast (SSM), multicast boundaries, and multicast BGP.

Part VI of the book explains the IPv6 address structure, the changes to the routing protocols, and IPv6 multicast routing.

- **Chapter 18, “IPv6 Addressing:”** This chapter describes the IPv6 address structure. The protocol stack’s neighbor discovery mechanisms are outlined, such as router advertisement messages, stateless address autoconfiguration, and duplicate address detection.
- **Chapter 19, “IPv6 Routing:”** This chapter outlines the subtle command structure and protocol mechanics changes between the IPv4 and IPv6 routing protocols.
- **Chapter 20, “IPv6 Multicast Routing:”** This chapter explains the fundamental differences between IPv4 and IPv6 multicast routing while emphasizing technologies like Multicast Listener Discovery (MLD), SSM, Embedded RP, and multicast boundaries.

Part VII, which can be found online at this book’s site, explains the concepts involved with improving the operational uptime of the network.

- **Chapter 21, “High Availability:”** This chapter describes the techniques available to improve network availability and provide fast routing convergence.

Final Words

This book is an excellent self-study resource to learn the routing protocols on Cisco IOS, IOS XE, and IOS XR operating systems. However, reading is not enough, and anyone who has obtained their CCIE will tell you that you must implement a technology to fully understand it. Our topologies are intentionally kept small to explain the routing concepts. We encourage the reader to re-create the topologies and follow along with the examples. A variety of resources are available that will allow you to practice the same concepts. Look online for the following:

- Online simulators at Learning@Cisco
- Online rack rentals
- Free demo versions of Cisco CSR 1000V (IOS XE)
- Free demo versions of Cisco IOS XRv (IOS XR)

Happy labbing!

This page intentionally left blank

How a Router Works

This chapter covers the following topics:

- IP routing
- IP packet switching
- Planes of operation

The previous chapters described that a router is necessary to transmit packets between network segments. This chapter explains the process a router uses to accomplish this task. By the end of this chapter, you should have a good understanding of how a router performs IP routing and IP packet forwarding between different network segments.

IP Routing

A router's primary function is to move an IP packet from one network to a different network. A router learns about nonattached networks through static configuration or through dynamic IP routing protocols.

Dynamic IP routing protocols distribute network topology information between routers and provide updates without intervention when a topology change in the network occurs. Design requirements or hardware limitations may restrict IP routing to static routes, which do not accommodate topology changes very well, and can burden network engineers depending on the size of the network. Routers try to select the best loop-free path in a network that forwards a packet to its destination IP address.

A network of interconnected routers and related systems managed under a common network administration is known as an *autonomous system*. The Internet is composed of thousands of autonomous systems spanning the globe.

The common dynamic routing protocols found in networks today are as follows:

- RIPv2 (Routing Information Protocol Version 2)
- EIGRP (Enhanced Interior Gateway Routing)
- OSPF (Open Shortest Path First) Protocol
- IS-IS (Intermediate System-to-Intermediate System) Protocol
- BGP (Border Gateway Protocol)

With the exception of BGP, the protocols in the preceding list are designed and optimized for routing within an autonomous system and are known as *internal gateway protocols* (IGPs). External gateway protocols (EGPs) route between autonomous systems. BGP is an EGP protocol but can also be used within an autonomous system. If BGP exchanges routes within an autonomous system, it is known as an *internal BGP* (iBGP) session. If it exchanges routes between different autonomous systems, it is known as an *external BGP* (eBGP) session.

Figure 3-1 shows an illustration of how one or many IGPs as well as iBGP can be running within an autonomous system and how eBGP sessions interconnect the various autonomous systems together.

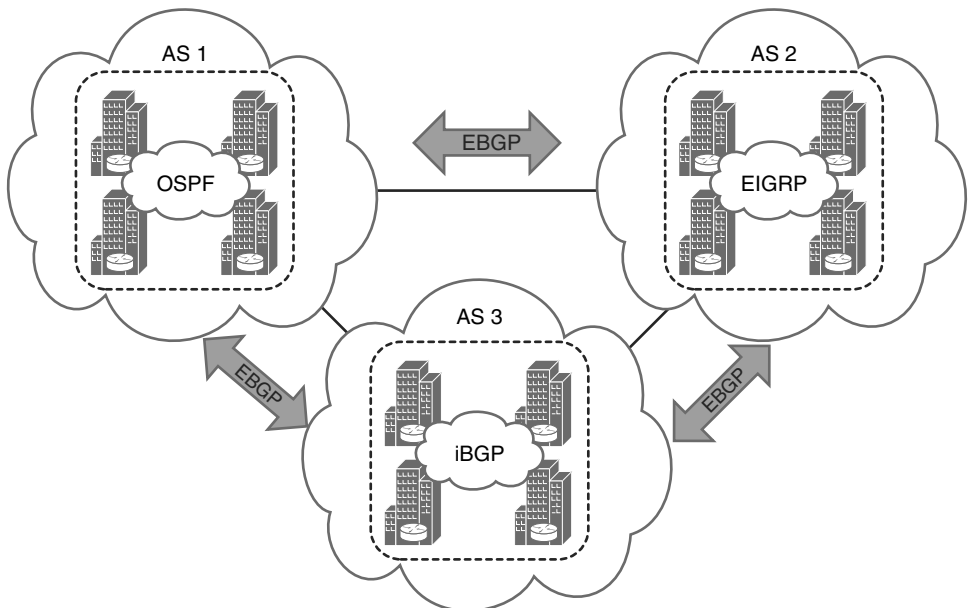


Figure 3-1 *Autonomous Systems and How They Interconnect*

EGPs and IGPs use different algorithms for path selection and are discussed in the following sections.

Distance Vector Algorithms

Distance vector routing protocols, such as RIP, advertise routes as vectors (distance, vector), where distance is a metric (or cost) such as hop count and vector is the next-hop router's IP used to reach the destination:

- **Distance:** The distance is the route metric to reach the network.
- **Vector:** The vector is the interface or direction to reach the network.

When a router receives routing information from a neighbor, it stores it in a local routing database as it is received and the distance vector algorithm (also known as *Bellman-Ford* and *Ford-Fulkerson* algorithms) is used to determine which paths are the best loop-free paths to each reachable destination. Once the best paths are determined, they are installed into the routing table and are advertised to each neighbor router.

Routers running distance vector protocols advertise the routing information to their neighbors from their own perspective, modified from the original route that it received. For this reason, distance vector protocols do not have a complete map of the whole network; instead, their database reflects that a neighbor router knows how to reach the destination network and how far the neighbor router is from the destination network. They do not know how many other routers are in the path toward any of those networks. The advantage of distance vector protocols is that they require less CPU and memory and can run on low-end routers.

An analogy commonly used to describe distance vector protocols is that of a road sign at an intersection that indicates the destination is 20 miles to the west; this information is trusted and blindly followed, without really knowing whether there is a shorter or better way to the destination or if the sign is even correct. Figure 3-2 illustrates how a router using a distance vector protocol views the network and the direction that R3 needs to go to reach the 192.168.1.0/24 subnet.




Figure 3-2 Distance Vector Protocol View of the Network

Enhanced Distance Vector Algorithm

The Diffused Update Algorithm (DUAL) is an enhanced distance vector algorithm that EIGRP uses to calculate the shortest path to a destination within a network. EIGRP advertises network information to its neighbors as other distance vector protocols do, but it has some enhancements as its name suggests. Some of the enhancements introduced into this algorithm compared to other distance vector algorithms are the following:

- Rapid convergence time for changes in the network topology.
- Only sends updates when there is a change in the network. It does not send full routing table updates in a periodic fashion like distance vector protocols.
- It uses hellos and forms neighbor relationships just like link-state protocols.
- It uses bandwidth, delay, reliability, load, and maximum transmission unit (MTU) size instead of hop count for path calculations.
- It has the option to load balance traffic across equal or unequal metric cost paths.

EIGRP is sometimes referred to as a *hybrid routing protocol* because it has characteristics of both distance vector and link-state protocols, as shown in the preceding list (for example, forming adjacencies with neighbor routers and relying on more advanced metrics such as bandwidth other than hop count for its best path calculations).

Link-State Algorithms

Link-state dynamic IP routing protocols advertise the link state and link metric for each of their connected links and directly connected routers to every router in the network. OSPF and IS-IS are two common link-state routing protocols found in enterprise and service provider networks. OSPF advertisements are called *link-state advertisements* (LSAs), and IS-IS uses link-state packets (LSPs) for its advertisements.

As a router receives an advertisement from a neighbor, it stores the information in a local database called the *link-state database* (LSDB), and advertises the link-state information on to each of its neighbor routers exactly as it was received. The link-state information is essentially flooded throughout the network from router to router unchanged, just as the originating router advertised it. This allows all the routers in the network to have a synchronized and identical map of the network.

Using the complete map of the network, every router in the network then runs the Dijkstra shortest path first (SPF) algorithm (developed by Edsger W. Dijkstra) to calculate the best shortest loop-free paths. The link-state algorithm then populates the routing table with this information.

Due to having the complete map of the network, link-state protocols usually require more CPU and memory than distance vector protocols, but they are less prone to routing loops and make better path decisions. In addition, link-state protocols are equipped with extended capabilities such as opaque LSAs for OSPF and TLVs (type/length/value)

for IS-IS that allows them to support features commonly used by service providers such as MPLS traffic engineering.

An analogy for link-state protocols is a GPS navigation system. The GPS navigation system has a complete map and can make the best decision as to which way is the shortest and best path to reach the destination. Figure 3-3 illustrates how R3 would view the network to reach the 192.168.1.0/24 subnet.




Figure 3-3 *Link-State Protocol View of the Network*

Path Vector Algorithm

A path vector protocol such as BGP is similar to a distance vector protocol; the difference is that instead of looking at the distance to determine the best loop-free path, it looks at various BGP path attributes. BGP path attributes include autonomous system path (AS_Path), Multi-Exit Discriminator (MED), origin, next hop, local preference, atomic aggregate, and aggregator. BGP path attributes are covered in Chapter 10, “BGP,” and Chapter 14, “Advanced BGP.”

A path vector protocol guarantees loop-free paths by keeping a record of each autonomous system that the routing advertisement traverses. Any time a router receives an advertisement in which it is already part of the autonomous system path, the advertisement is rejected because accepting the autonomous system path would effectively result in a routing loop.

Figure 3-4 illustrates this concept where autonomous system 1 advertises the 10.1.1.0/24 network to autonomous system 2. Autonomous system 2 receives this information and adds itself to the autonomous system path and advertises it to autonomous system 4. Autonomous system 4 adds itself to the path and advertises it to autonomous system 3. Autonomous system 3 receives the route advertisement and adds itself to the path as well. However, when autonomous system 3 advertises that it can reach 10.1.1.0/24 to autonomous system 1, autonomous system 1 discards the advertisement because the

autonomous system path (path vector) contained in the advertisement includes its autonomous system number (autonomous system 1). When autonomous system 3 attempts to advertise reachability for 10.1.1.0/24 to autonomous system 2, autonomous system 2 also discards it because the advertisement includes autonomous system 2 in the autonomous system path, too.

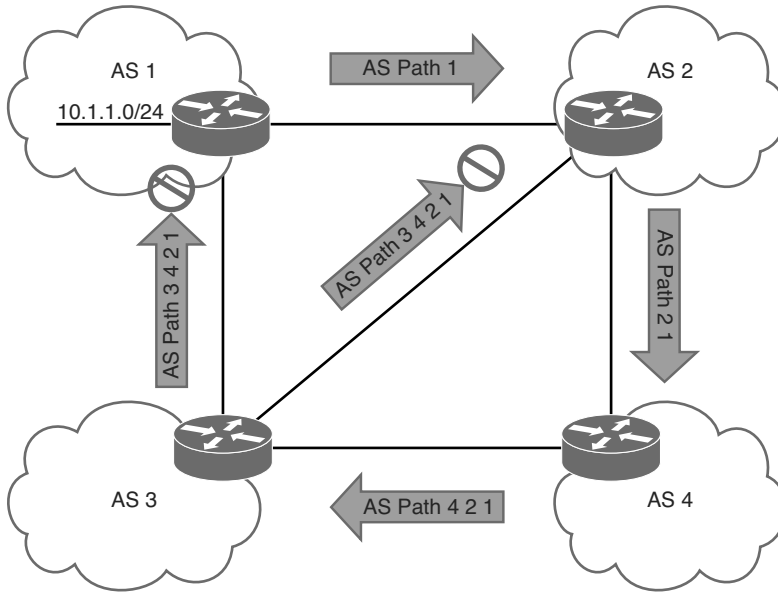


Figure 3-4 Path Vector Loop Avoidance

All BGP path attributes and how to manipulate them to influence the best path selection process are covered in Chapter 15, “BGP Best Path Selection.”

Routing Table

A router identifies the path a packet should take by evaluating the following components on a router:

- **Prefix length:** The prefix length represents the number of leading binary bits in the subnet mask that are in the on position.
- **Administrative distance:** Administrative distance (AD) is a rating of the trustworthiness of a routing information source. If a router learns about a route to a destination from more than one routing protocol and they all have the same prefix length, AD is compared. The preference is given to the route with the lower AD.
- **Metrics:** A unit of measure used by a routing protocol in the best path calculation.

Prefix Length

Let's look at a scenario of a router selecting a route when the packet destination is within the network range for multiple routes. Assume that a router has the following routes with various prefix lengths in the routing table:

- 10.0.3.0/28
- 10.0.3.0/26
- 10.0.3.0/24

Because each of these routes, also known as *prefix routes* or simply *prefixes*, has a different prefix length (subnet mask), they are considered to be different destinations, and they will all be installed into the routing table. This is represented in Table 3-1.

Table 3-1 *Representation of Routing Table*

Prefix	Subnet Range	Next Hop	Outgoing Interface
10.0.3.0/28	10.0.3.0 – 10.0.3.15	10.1.1.1	Gigabit Ethernet 1/1
10.0.3.0/26	10.0.3.0 – 10.0.3.63	10.2.2.2	Gigabit Ethernet 2/2
10.0.3.0/24	10.0.3.0 – 10.0.3.255	10.3.3.3	Gigabit Ethernet 3/3

If a packet needs to be forwarded, the route chosen depends on the prefix length, where the *longest prefix length* is always preferred. For example, /28 is preferred over /26, and /26 is preferred over /24. The following is an example using Table 3-1 as a reference:

- If a packet needs to be forwarded to 10.0.3.14, it would match all three routes, but it would be sent to next hop 10.1.1.1 and outgoing interface Gigabit Ethernet 1/1 because 10.0.3.0/28 has the longest prefix match.
- If a packet needs to be forwarded to 10.0.3.42, it would match 10.0.3.0/24 and 10.0.3.0/26, so the packet would be sent to 10.2.2.2 and outgoing interface Gigabit Ethernet 2/2 because 10.0.3.0/26 has the longest prefix match.
- If a packet needs to be forwarded to 10.0.3.100, it matches only 10.0.3.0/24, so the packet is sent to 10.3.3.3 and outgoing interface Gigabit Ethernet 3/3.

Administrative Distance

As each routing protocol receives updates and other routing information, it chooses the best path to any given destination and attempts to install this path into the routing table. Table 3-2 provides the default AD for the routing protocols covered in this book.

Table 3-2 *Routing Protocol Default Administrative Distances*

Routing Protocol	Default Administrative Distance
Connected	0
Static	1
eBGP	20
EIGRP summary route	5
EIGRP (internal)	90
OSPF	110
IS-IS	115
RIP	120
EIGRP (external)	170
iBGP	200

For example, if OSPF learns of a best path toward 10.0.1.0/24, it first checks to see whether an entry exists in the routing table. If it does not exist, the route is installed into the Routing Information Base (RIB). If the route already exists in the RIB, the router decides whether to install the route presented by OSPF based on the AD of the route in OSPF and the AD of the existing route in the RIB. If this route has the *lowest AD* to the destination (when compared to the other route in the table), it is installed in the routing table. If this route is not the route with the best AD, the route is rejected.

Consider another example on this topic. A router has OSPF, IS-IS, and EIGRP running, and all three protocols have learned of the destination 10.3.3.0/24 network with a different best path and metric.

Each of these three protocols will then attempt to install the route to 10.3.3.0/24 into the routing table. Because the prefix length is the same, the next decision point is the AD, where the routing protocol with the lowest AD installs the route into the routing table.

Because the EIGRP internal route has the best AD, it is the one installed into the routing table:

10.0.3.0/24	EIGRP	90 <<< Lowest AD Installed in Route Table
10.0.3.0/24	OSPF	110
10.0.3.0/24	IS-IS	115

The routing protocol or protocols that failed to install their route into the table (in this example, that would be OSPF and IS-IS) will hang on to this route to use it as a backup route and will tell the routing table process to report to them if the best path fails so that they can then try to reinstall this route.

For example, if the EIGRP route 10.0.3.0/24 installed in the routing table fails for some reason, the routing table process calls OSPF and IS-IS, and requests them to reinstall the route in the routing table. Out of these two protocols, the preferred route is chosen based on AD, which would be OSPF because of its lower AD.

The default AD might not always be suitable for a network; for instance, there might be a requirement to adjust it so that OSPF routes are preferred over EIGRP routes. However, changing the AD on routing protocols can have severe consequences, such as routing loops and other odd behavior in a network. It is recommended that the AD be changed only with extreme caution, and only after what needs to be accomplished has been thoroughly thought out. A good backup plan is recommended in case things do not turn out as planned.

Metrics

As discussed in the previous section, routes are chosen and installed into the routing table based on the routing protocol's AD. The routes learned from the routing protocol with the lowest AD are the ones installed into the routing table. If there are multiple paths to the same destination from a single routing protocol, these paths would have the same AD; for this case, the best path is selected within the routing protocol. Most protocols use the path with the best metric, but OSPF and IS-IS have additional logic that preempts the lowest metric.

If a routing protocol identifies multiple paths as a *best path*, and supports multiple path entries, the router installs the maximum number of paths allowed per destination. This is known as *equal-cost multipath* (ECMP) and provides load sharing across all links.

For example, Figure 3-5 illustrates a network running OSPF to reach the prefix 10.3.3.0/24. Router 1 (R1) has two equal-cost paths; therefore, it will install both in the routing table.

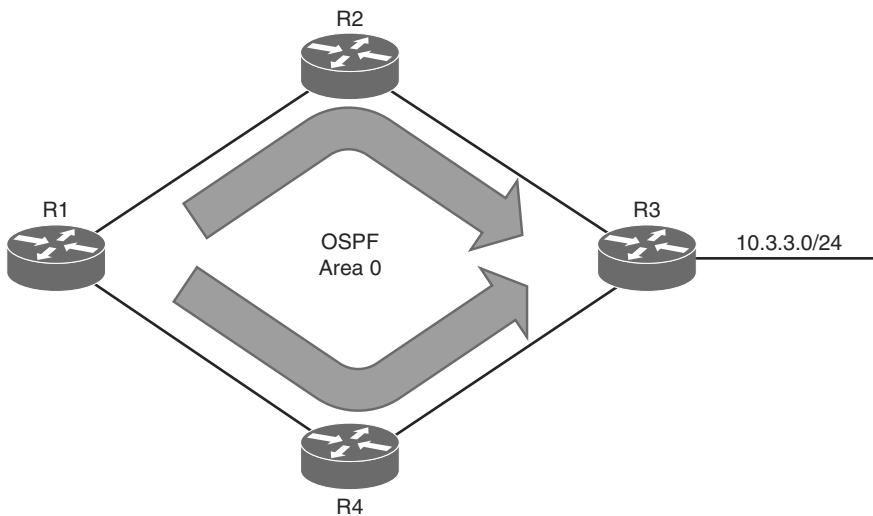


Figure 3-5 OSPF ECMP Technology

Example 3-1 confirms that both paths have been installed into the RIB, and because the metrics are identical, this confirms the router is using ECMP.

Example 3-1 *R1's Routing Table Showing the ECMP Paths to 10.3.3.0/24*

```
R1#show ip route
! Output omitted for brevity
O    10.3.3.0/24 [110/30] via 10.4.4.4, 00:49:12, GigabitEthernet0/0
                               [110/30] via 10.2.2.2, 00:49:51, GigabitEthernet0/2
```

Note Best path metric calculation and the default and maximum ECMP paths allowed for each routing protocol vary. This is covered in later routing protocol-related chapters.

Virtual Routing and Forwarding

Virtual Routing and Forwarding (VRF) is a technology that allows multiple independent virtual routing table and forwarding table instances to exist concurrently in a router. This can be leveraged to create segmentation between networks, which allows for overlapping IP addressing to be used even on a single interface (that is, using subinterfaces), and because the traffic paths are isolated, network security is increased and can eliminate the need for encryption and authentication for network traffic.

Service Providers with Multiprotocol Label Switching (MPLS) backbones typically use VRFs to create separate virtual private networks (VPNs) for their customers, and when used in this manner, VRFs are known as *VPN Routing and Forwarding*.

When VRF is not used in conjunction with MPLS, it is known as *VRF-Lite* (also termed *multi-VRF CE*, or *multi-VRF customer-edge device*). Because MPLS is beyond the scope of this book, only VRF-Lite is covered in this section and is referred to it simply as VRF.

The configurations in Example 3-2 should help clarify the VRF concept. Example 3-2 shows how configuring different interfaces with overlapping IP addresses and subnets is not allowed within a routing table, not even if they are both on different interfaces because they would end up in the same routing table and cause a conflict.

Example 3-2 *Overlapping IP Address Problems*

```
IOS
R1 (config)#interface GigabitEthernet0/1
R1 (config-if)#ip address 10.0.3.1 255.255.255.0
R1 (config-if)#interface GigabitEthernet0/3
R1 (config-if)#ip address 10.0.3.2 255.255.255.0
```

```
% 10.0.3.0 overlaps with GigabitEthernet0/1
```

```
IOS XR
```

```
RP/0/0/CPU0:XR2 (config)#interface gigabitEthernet 0/0/0/5
```

```
RP/0/0/CPU0:XR2 (config-if)#ipv4 address 10.0.3.1/24
```

```
RP/0/0/CPU0:XR2 (config-if)#commit
```

```
RP/0/0/CPU0:XR2 (config-if)#
```

```
RP/0/0/CPU0:XR2 (config)#interface gigabitEthernet 0/0/0/3
```

```
RP/0/0/CPU0:XR2 (config-if)#ipv4 address 10.0.3.2/24
```

```
RP/0/0/CPU0:XR2 (config-if)#commit
```

```
RP/0/0/CPU0:Jan 13 18:55:35.643 : ipv4_arm[189] : %IP-IP_ARM-3-CFLCT_FORCED_DOWN :
The IPv4 address 10.0.3.1/24 on GigabitEthernet0/0/0/5 conflicts with other IPv4
addresses and has been forced down
```

Note In IOS XR, the IP Address Repository Manager (IPARM) enforces the uniqueness of global IP addresses configured in the system. By default, when there is an IP address and subnet mask conflict, the lowest rack/slot/interface (that is, g0/0/0/3 is lower than g0/0/0/5) is the one that gets assigned the IP address. To change the default behavior, use the `ipv4 conflict-policy {static | highest-ip | longest-prefix}` command.

In older IOS releases, only single-protocol IPv4-only VRFs could be created. The command `ip vrf vrf-name` created a single-protocol VRF on the router and was activated on an interface with the command `ip vrf forwarding vrf-name` under the interface configuration mode.

In current IOS releases, a new configuration option allows the creation of multiprotocol VRFs that support both IPv4 and IPv6. Entering the command `vrf definition vrf-name` creates the multiprotocol VRF. Under VRF definition submode, the command `address-family {ipv4 | ipv6}` is required to specify the appropriate address family. The VRF is then associated to the interface with the command `vrf forwarding vrf-name` under the interface configuration submode.

Note The commands `ip vrf vrf-name` and `ip vrf forwarding vrf-name` will be available for a period of time before they are deprecated. To migrate any older IPv4-only VRFs to the new multiprotocol VRF configuration, you can use the `vrf upgrade-cli multi-af-mode {common-policies | non-common-policies} [vrf vrf-name]` command. When creating a new VRF, even if it is just an IPv4-only VRF, Cisco recommends using the multiprotocol VRF `vrf definition` and `vrf forwarding` commands.

In IOS, the following steps are required to create a VRF and assign it to an interface:

Step 1. Create a multiprotocol VRF.

The multiprotocol VRF routing table is created with the command `vrf definition` *vrf-name*.

Step 2. Identify the address family.

Initialize the appropriate address family with the command `address-family` (`ipv4` | `ipv6`). The address family can be IPv4, IPv6, or both.

Step 3. Specify the interface to be associated with the VRF.

Enter interface configuration submode and specify the interface to be associated with the VRF with the command `interface` *interface-type interface-number*.

Step 4. Associate the VRF to the interface.

The VRF is associated to the interface or subinterface by entering the command `vrf forwarding` *vrf-name* under interface configuration submode.

Step 5. Configure an IP address on the interface or subinterface.

The IP address can be IPv4, IPv6, or both. It is configured by entering the following commands:

```
IPv4
```

```
ip address ip-address subnet-mask [secondary]
```

```
IPv6
```

```
ipv6 address {ipv6-address/prefix-length | prefix-name sub-bits/
prefix-length}
```

Note On IOS nodes, the VRF needs to be associated to the interface first before configuring an IP address. If an IP address is already configured, and the VRF is associated to the interface, IOS will remove the IP address.

IOS XR supports only multiprotocol VRFs. The following steps are required to create a multiprotocol VRF and assign it to an interface on an IOS XR node:

Step 1. Create a multiprotocol VRF.

The multiprotocol VRF routing table is created with the command `vrf` *vrf-name*. The VRF name is arbitrary.

Step 2. Identify the address family.

Initialize the appropriate address family with the command `address-family` (`ipv4` | `ipv6`) `unicast`. The address family can be IPv4, IPv6, or both.

Step 3. Specify the interface to be associated with the VRF.

Enter interface configuration submode and specify the interface to be associated with the VRF with the command `interface` *interface-type interface-number*.

Step 4. Associate the VRF with an interface or subinterface.

The VRF is associated with the interface or subinterface by entering the command `vrf vrf-name` under interface configuration submode.

Step 5. Configure an IP address on the interface or subinterface.

The IP address can be IPv4, IPv6, or both. It is configured by entering the following commands:

IPv4

```
ipv4 address ipv4-address subnet-mask
```

IPv6

```
ipv6 address ipv6-address/prefix-length
```

Note For IOS XR, the VRF needs to be associated to the interface first before configuring an IP address; otherwise, the VRF configuration will not be accepted.

Figure 3-6 Illustrates two routers to help visualize the VRF routing table concept. One of the routers has no VRFs configured, and the other one has a management VRF named MGMT. This figure can be used as a reference for the following examples.




Figure 3-6 Comparison of a Router with no VRFs and a Router with a VRF

Table 3-3 provides a set of interfaces and IP addresses that overlap between the global routing table and the VRF. This information is used in the following examples.

Table 3-3 *Sample Interfaces and IP Addresses*

IOS Interface	IOS XR Interface	IP Address	VRF	Global
Gigabit Ethernet 0/1	Gigabit Ethernet 0/0/0/1	10.0.3.1/24		✓
Gigabit Ethernet 0/2	Gigabit Ethernet 0/0/0/2	10.0.4.1/24		✓
Gigabit Ethernet 0/3	Gigabit Ethernet 0/0/0/3	10.0.3.1/24	MGMT	
Gigabit Ethernet 0/4	Gigabit Ethernet 0/0/0/4	10.0.4.1/24	MGMT	

Example 3-3 shows how the IP addresses are assigned to the interfaces in the global routing table shown in Table 3-3.

Example 3-3 *IP Address Configuration in Global Routing Table*

```

IOS
R1(config)#interface GigabitEthernet0/1
R1(config-if)#ip address 10.0.3.1 255.255.255.0
R1(config)#interface GigabitEthernet0/2
R1(config-if)#ip address 10.0.4.1 255.255.255.0

IOS XR
RP/0/0/CPU0:XR1(config)#interface gigabitEthernet 0/0/0/1
RP/0/0/CPU0:XR1(config-if)#ipv4 address 10.0.3.1/24
RP/0/0/CPU0:XR1(config)#interface gigabitEthernet 0/0/0/2
RP/0/0/CPU0:XR1(config-if)#ipv4 address 10.0.4.1/24
RP/0/0/CPU0:XR1(config-if)#commit

```

Example 3-4 displays the global routing table with the command `show ip route` for IOS and `show route` for IOS XR to show the IP addresses configured in Example 3-3.

Example 3-4 *Output of Global Routing Table*

```

IOS
R1#show ip route
! Output omitted for brevity
    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C       10.0.3.0/24 is directly connected, GigabitEthernet0/1
L       10.0.3.1/32 is directly connected, GigabitEthernet0/1
C       10.0.4.0/24 is directly connected, GigabitEthernet0/2
L       10.0.4.1/32 is directly connected, GigabitEthernet0/2

IOS XR
RP/0/0/CPU0:XR1#show route

```



```
! Output omitted for brevity
```

```
C 10.0.3.0/24 is directly connected, 00:00:25, GigabitEthernet0/0/0/1
L 10.0.3.1/32 is directly connected, 00:00:25, GigabitEthernet0/0/0/1
C 10.0.4.0/24 is directly connected, 00:00:02, GigabitEthernet0/0/0/2
L 10.0.4.1/32 is directly connected, 00:00:02, GigabitEthernet0/0/0/2
```

Example 3-5 shows how the VRF named MGMT is created, two interfaces are associated with it, and the IP addresses in Table 3-3 are configured on the interfaces. These IP addresses overlap with the ones configured in Example 3-3, but there is no conflict because they are in a different routing table.

Example 3-5 VRF Configuration Example

IOS

```
R1 (config)#vrf definition MGMT
R1 (config-vrf)# address-family ipv4
R1 (config)#interface GigabitEthernet0/3
R1 (config-if)#vrf forwarding MGMT
R1 (config-if)#ip address 10.0.3.1 255.255.255.0
R1 (config)#interface GigabitEthernet0/4
R1 (config-if)#vrf forwarding MGMT
R1 (config-if)#ip address 10.0.4.1 255.255.255.0
```

IOS XR

```
RP/0/0/CPU0:XR1 (config)#vrf MGMT address-family ipv4 unicast
RP/0/0/CPU0:XR1 (config-vrf-af)#root
RP/0/0/CPU0:XR1 (config)#interface gigabitEthernet 0/0/0/3
RP/0/0/CPU0:XR1 (config-if)#vrf MGMT
RP/0/0/CPU0:XR1 (config-if)#ipv4 address 10.0.3.1/24
RP/0/0/CPU0:XR1 (config)#interface gigabitEthernet 0/0/0/4
RP/0/0/CPU0:XR1 (config-if)#vrf MGMT
RP/0/0/CPU0:XR1 (config-if)#ipv4 address 10.0.4.1/24
RP/0/0/CPU0:XR1 (config-if)#commit
```

Example 3-6 shows how the VRF IP addresses configured in Example 3-5 cannot be seen in the output of the `show ip route` command for IOS and the `show route` command for IOS XR; these commands display only the contents of the global routing table. To see a VRF routing table, the commands `show ip route vrf vrf-name` for IOS and `show route vrf {all | vrf-name}` for IOS XR should be used.

Example 3-6 *Output of Global Routing Table and VRF Routing Table*

```

R1#show ip route
! Output omitted for brevity
    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C       10.0.3.0/24 is directly connected, GigabitEthernet0/1
L       10.0.3.1/32 is directly connected, GigabitEthernet0/1
C       10.0.4.0/24 is directly connected, GigabitEthernet0/2
L       10.0.4.1/32 is directly connected, GigabitEthernet0/2

R1#show ip route vrf MGMT
! Output omitted for brevity
    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C       10.0.3.0/24 is directly connected, GigabitEthernet0/3
L       10.0.3.1/32 is directly connected, GigabitEthernet0/3
C       10.0.4.0/24 is directly connected, GigabitEthernet0/4
L       10.0.4.1/32 is directly connected, GigabitEthernet0/4

```

```

RP/0/0/CPU0:XR1#show route
! Output omitted for brevity
C       10.0.3.0/24 is directly connected, 00:12:44, GigabitEthernet0/0/0/1
L       10.0.3.1/32 is directly connected, 00:12:44, GigabitEthernet0/0/0/1
C       10.0.4.0/24 is directly connected, 00:12:21, GigabitEthernet0/0/0/2
L       10.0.4.1/32 is directly connected, 00:12:21, GigabitEthernet0/0/0/2

RP/0/0/CPU0:XR1#show route vrf MGMT
! Output omitted for brevity
C       10.0.3.0/24 is directly connected, 00:09:15, GigabitEthernet0/0/0/3
L       10.0.3.1/32 is directly connected, 00:09:15, GigabitEthernet0/0/0/3
C       10.0.4.0/24 is directly connected, 00:00:10, GigabitEthernet0/0/0/4
L       10.0.4.1/32 is directly connected, 00:00:10, GigabitEthernet0/0/0/4

```

In IOS, to display a quick summary of the usability status for each IP interface, in addition to all the IP addresses configured in the global routing table and all VRFs, the command **show ip interface brief** should be used. In IOS XR, the command **show ipv4 interface brief** only shows the IP addresses in the global routing table. To see the IP addresses in the global routing table and all VRFs, use the command **show ipv4 vrf all interface brief**. Example 3-7 provides sample output of these **show** commands.

Example 3-7 *Verification of Interfaces Status and IP Addresses*

```

R1#show ip interface brief

```

Interface	IP-Address	OK?	Method	Status	Protocol
GigabitEthernet0/1	10.0.3.1	YES	NVRAM	up	up

GigabitEthernet0/2	10.0.4.1	YES	NVRAM	up	up
GigabitEthernet0/3	10.0.3.1	YES	NVRAM	up	up
GigabitEthernet0/4	10.0.4.1	YES	NVRAM	up	up


```
RP/0/0/CPU0:XR2#show ipv4 interface brief
```

Interface	IP-Address	Status	Protocol
GigabitEthernet0/0/0/0	unassigned	Shutdown	Down
GigabitEthernet0/0/0/1	10.0.3.1	Up	Up
GigabitEthernet0/0/0/2	10.0.4.1	Up	Up


```
RP/0/0/CPU0:XR2#show ipv4 vrf all interface brief
```

Interface	IP-Address	Status	Protocol	Vrf-Name
GigabitEthernet0/0/0/0	unassigned	Shutdown	Down	default
GigabitEthernet0/0/0/1	10.0.3.1	Up	Up	default
GigabitEthernet0/0/0/2	10.0.4.1	Up	Up	default
GigabitEthernet0/0/0/3	10.0.3.1	Up	Up	MGMT
GigabitEthernet0/0/0/4	10.0.4.1	Up	Up	MGMT

VRF-Lite can provide similar functionality to that of virtual local-area networks (VLANs); however, instead of relying on Layer 2 technologies such as spanning tree, Layer 3 dynamic routing protocols can be used. Using routing protocols over Layer 2 technologies has some advantages such as improved network convergence times, dynamic traffic load sharing, and troubleshooting tools such as ping and traceroute.

IP Packet Switching

Chapter 2, “IP Addressing,” explained that devices on the same subnet could communicate directly with each other without the need of a router. The second layer of the OSI model, the data link layer, handles addressing beneath the IP protocol stack so that communication is directed between hosts. Network packets include the Layer 2 addressing with unique source and destination addresses for that segment. Ethernet commonly uses MAC addresses, and other data link layer protocols such as Frame Relay use an entirely different method of Layer 2 addressing.

The first routers would receive a packet, remove the Layer 2 information, and verify that the route exists for the destination IP address. If a matching route could not be found, the packet was dropped. If a matching route was found, the router would identify it and add new Layer 2 information to the packet. The Layer 2 source address would be the router’s outbound interface, and the destination information would be next hop’s Layer 2 address.

Figure 3-7 illustrates the concept where PC A is sending a packet to PC B via Ethernet connection to R1. PC A sends the packet to R1’s MAC address of 00:C1:5C: 00:00:02. R1 receives the packet, removes the Layer 2 information, and looks for a route to the

192.168.2.2 address. R1 identifies that connectivity to the 192.168.2.2 IP address is through Gigabit Ethernet 0/1. R1 adds the Layer 2 source address using its Gigabit Ethernet 0/1's MAC address 00:C1:5C:00:00:03 and a destination address for PC B of 00:00:00:00:00:04.




Figure 3-7 Layer 2 Addressing

Advancement in technologies has streamlined the process so that routers do not remove and add the Layer 2 addressing but simply rewrites them. IP packet switching or IP packet forwarding is the faster process of receiving an IP packet on an input interface and making a decision of whether to forward the packet to an output interface or drop it. This process is simple and streamlined for a router to be able to forward large amounts of packets.

When the first Cisco routers were developed, they used a mechanism called *process switching* to switch the packets through the routers. As network devices evolved, Cisco created Fast Switching and Cisco Express Forwarding (CEF) to optimize the switching process for the routers to be able to handle larger packet volumes. Fast Switching is deprecated in newer IOS releases and is not covered in this book.

Process Switching

Process switching, also referred to as *software switching* or *slow path*, is the switching mechanism in which the general-purpose CPU on a router is in charge of packet switching. In IOS, the `ip_input` process runs on the general-purpose CPU for processing incoming IP packets. Process switching is the fallback for CEF because it is dedicated for processing punted IP packets when they cannot be switched by CEF.

In IOS XR, the Network Input/Output (NetIO) process is the equivalent to the IOS `ip_input` process and is responsible for forwarding packets in software.

The type of packets that require software handling for both IOS and IOS XR include the following:

- Packets sourced or destined to the router (that is, control traffic, routing protocols)
- Packets that are too complex for the hardware to handle (that is, IP packets with IP options)
- Packets that require extra information that is not currently known (that is, Address Resolution Protocol [ARP] resolution, and so on)

Note Software switching is significantly slower than switching done in hardware. NetIO is designed to handle a very small percentage of traffic handled by the system. Packets are hardware switched whenever possible.

Figure 3-8 illustrates how a packet that cannot be CEF switched is punted to the CPU for processing. The `ip_input` process consults the routing table and ARP table to obtain the next-hop router's IP address, outgoing interface, and MAC address. It then overwrites the destination MAC address of the packet with the next-hop router's MAC address, overwrites the source MAC address with the MAC address of the outgoing Layer 3 interface, decrements the IP Time-To-Live (TTL) field, recomputes the IP header checksum, and finally delivers the packet to the next-hop router.

The routing table, also known as the *Routing Information Base (RIB)*, is built from information obtained from dynamic routing protocols, directly connected and static routes. The ARP table is built from information obtained from the ARP protocol. The ARP protocol is used by IP hosts to dynamically learn the MAC address of other IP hosts on the same subnet. For example, an IP host that needs to perform address resolution for another IP host connected by Ethernet can send an ARP request using a LAN broadcast address, and it then waits for an ARP reply from the IP host. The ARP reply includes the required Layer 2 physical MAC address information.

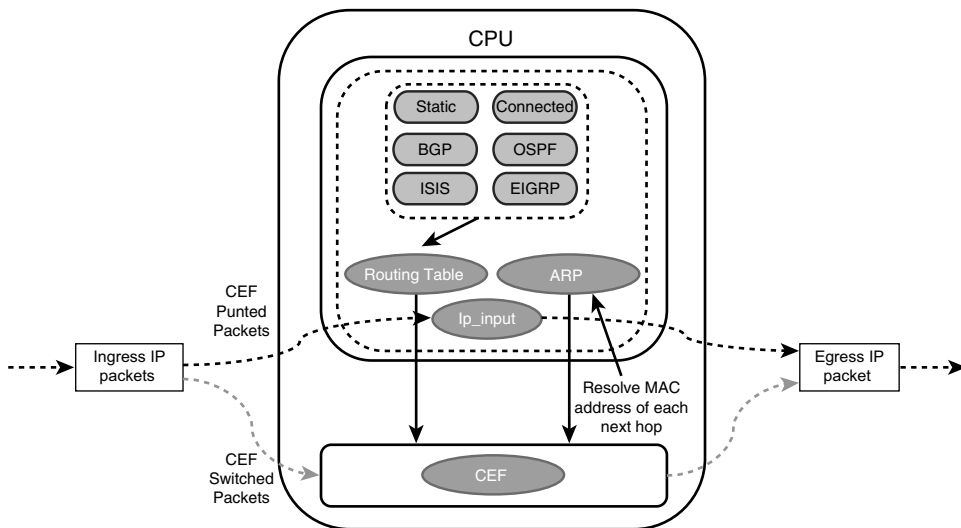


Figure 3-8 Process Switching

Cisco Express Forwarding

Cisco Express Forwarding (CEF) is a Cisco proprietary switching mechanism developed to keep up with the demands of evolving network infrastructures. It has been the default switching mechanism on most Cisco platforms that do all their packet switching using the general-purpose CPU (software based routers) since the 1990s, and it is the default

switching mechanism used by all Cisco platforms that use specialized application specific integrated circuits (ASICs) and network processing units (NPU) for high packet throughput (hardware-based routers).

The general-purpose CPU on the software-based and hardware-based routers is similar and perform all the same functions, the difference being that on software based routers the general-purpose CPU is in charge of all operations, including CEF switching (software CEF), and the hardware-based routers do CEF switching using *forwarding engines* that are implemented in specialized ASICs, TCAMs, and NPUs (hardware CEF). Forwarding engines provide the packet switching, forwarding, and route lookup capability to routers.

Given the low cost of the general-purpose CPUs, the price point of software-based routers will be much more affordable, but at the expense of total packet throughput.

When a route processor (RP) engine is equipped with a forwarding engine so that it can make all the packet switching decisions, this is known as a *centralized forwarding architecture*. If the line cards are equipped with forwarding engines so that they can make packet switching decision without intervention of the RP, this is known as a *distributed forwarding architecture*.

For a centralized forwarding architecture, when a packet is received on the ingress line card, it is transmitted to the forwarding engine on the RP. The forwarding engine examines the packet's headers and determines that the packet will be sent out a port on the egress line card, and forwards the packet to the egress line card to be forwarded.

For a distributed forwarding architecture, when a packet is received on the ingress line card, it is transmitted to the local forwarding engine. The forwarding engine performs a packet lookup, and if it determines that the outbound interface is local, it forwards the packet out a local interface. If the outbound interface is located on a different line card, the packet is sent across the switch fabric, also known as the *backplane*, directly to the egress line card, bypassing the RP.

Figure 3-9 illustrates a packet flowing across a centralized and a distributed forwarding architecture.

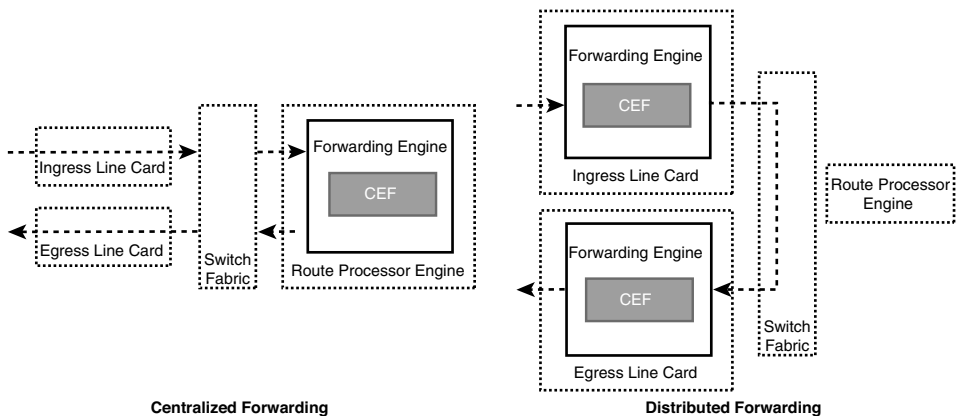


Figure 3-9 Centralized Versus Distributed Forwarding Architectures

Software CEF

Software CEF, also known as the software *Forwarding Information Base (FIB)*, consists of the following components:

- Forwarding Information Base:** The FIB is built directly from the routing table and contains the next-hop IP address for each destination IP in the network. It keeps a mirror image of the forwarding information contained in the IP routing table. When a routing or topology change occurs in the network, the IP routing table is updated, and these changes are reflected in the FIB. CEF uses the FIB to make IP destination prefix-based switching decisions
- Adjacency table:** The adjacency table is also known as the *Adjacency Information Base (AIB)*. It contains the MAC addresses and egress interfaces of all directly connected next hops, and it is populated with data from the ARP table and other Layer 2 protocol tables (that is, Frame Relay map tables).

Figure 3-10 illustrates how the CEF table is built from the routing table and the ARP table and how a packet is CEF switched through the router. When an IP packet is received, if there is a valid FIB and adjacency table entry for it, the router overwrites the destination MAC address of the packet with the next hop router's MAC address, overwrites the source MAC address with the MAC address of the outgoing Layer 3 interface, decrements IP TTL field, recomputes the IP header checksum, and finally delivers the packet to the next-hop router.

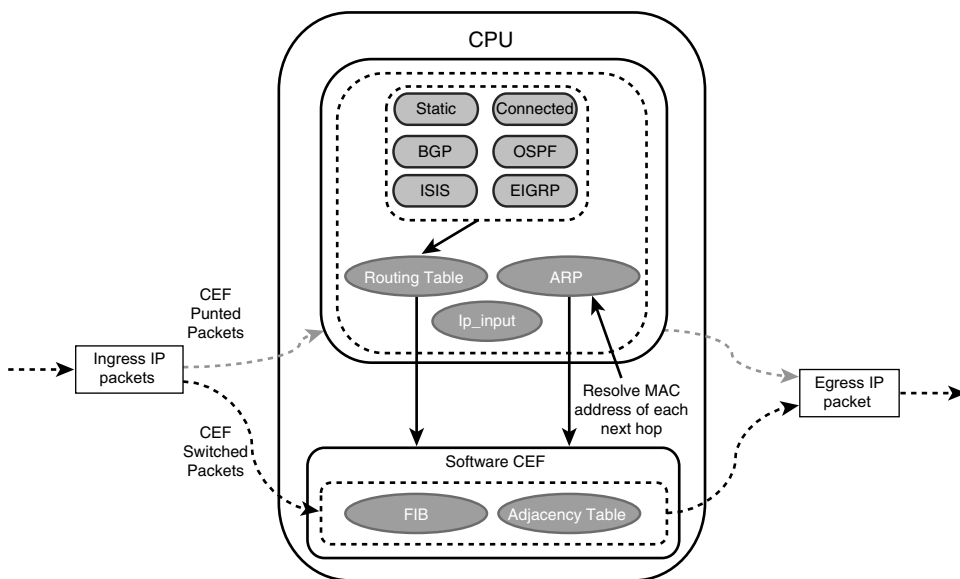


Figure 3-10 CEF Switching

Hardware CEF

The ASICs in hardware-based routers have a very high cost to design, produce, and troubleshoot. ASICs allow for very high packet rates, but the trade-off is that they are limited in their functionality because they are hardwired to perform specific tasks. There are routers equipped with NPUs that are designed to overcome the inflexibility of ASICs. Unlike ASICs, NPUs are programmable, and their firmware can be changed with relative ease.

The main advantage of the distributed forwarding architectures is that the packet throughput performance is greatly improved by offloading the packet switching responsibilities to the line cards. Packet switching in distributed architecture platforms is done via distributed CEF (dCEF), which is a mechanism in which the CEF data structures are downloaded to forwarding ASICs and the CPUs of all line cards so that they can participate in packet switching; this allows for the switching to be done at the distributed level, thus increasing the packet throughput of the router.

Software CEF in hardware-based platforms is not used to do packet switching as in software-based platforms; instead, it is used to program the hardware CEF, as shown in Figure 3-11.

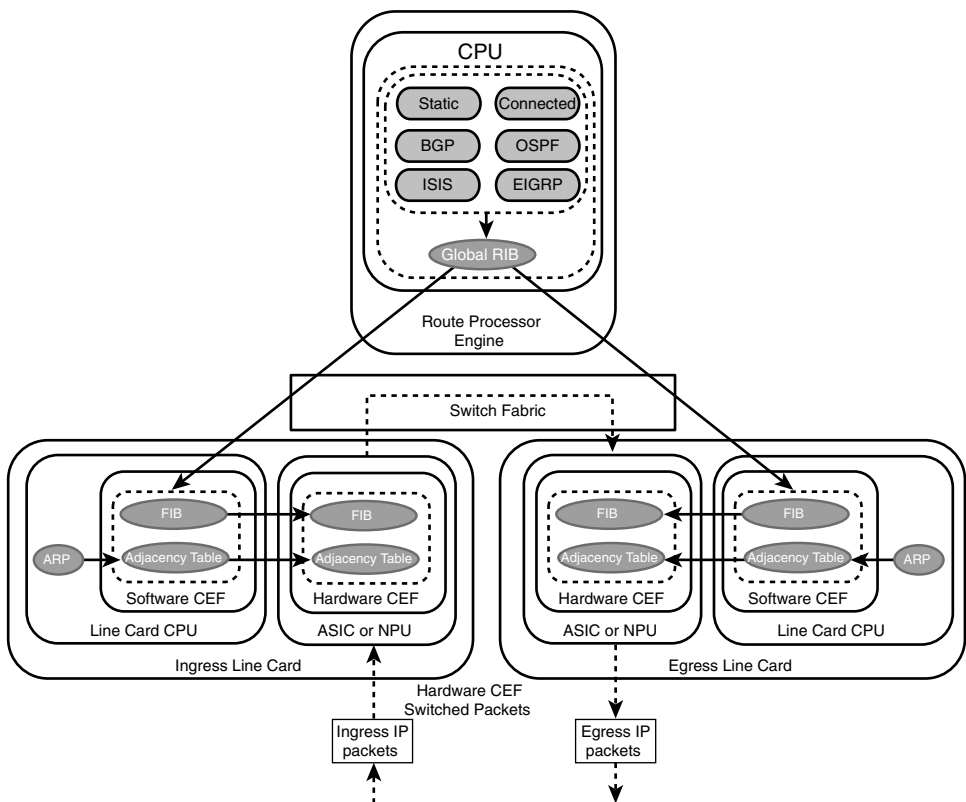


Figure 3-11 dCEF Hardware Switching

Figure 3-11 also illustrates how the RIB process interacts with the RIBs of the routing protocols. The RIB process is in charge of the calculation of best paths, alternative paths, and

the redistribution from different protocols and all these details merge into the global RIB (gRIB), where the best path for a destination network is installed. This is further distributed into the software CEF tables of different line cards, which is further mirrored into hardware CEF. The Switch Fabric is the backplane for all modules in the system. It creates a dedicated connection between all line cards and the route processors and provides fast data switching transmission between them.

In most distributed architecture platforms, if the incoming packet is control plane traffic or management traffic it is punted to the RP's CPU. The following list includes some examples of packets that are typically punted for processing by the RP's CPU or line card's CPU:

- Control traffic, such as BGP, OSPF, IS-IS, PIM, IGMP, and so on
- Management traffic, such as Telnet, SSH, SNMP, and so on
- Layer 2 mechanisms, such as CDP, ARP, LACP PDU, BFD, and so on
- Fragmentation, DF bit set, IP options set
- TTL expired
- ICMP echo request

Planes of Operation

A router is typically segmented into three planes of operation, each with a specific and clearly defined objective:

- **The control plane:** The control plane is the brain of the router. It consists of dynamic IP routing protocols (that is OSPF, IS-IS, BGP, and so on), the RIB, routing updates, in addition to other protocols such as PIM, IGMP, ICMP, ARP, BFD, LACP, and so on. In short, the control plane is responsible for maintaining sessions and exchanging protocol information with other router or network devices.

In centralized architecture platforms, the general-purpose CPU manages all control plane protocols. In distributed architecture platforms, routing protocols, and most other protocols, always run on the core CPU in the RPs or Supervisor engines, but there are other control plane protocols such as ARP, BFD, and ICMP that in some distributed architecture platforms have now been offloaded to the line card CPU.

- **The data plane:** The data plane is the forwarding plane, which is responsible for the switching of packets through the router (that is, process switching and CEF switching). In the data plane, there could be features that could affect packet forwarding such as quality of service (QoS) and access control lists (ACLs).
- **The management plane:** The management plane is used to manage a device through its connection to the network. Examples of protocols processed in the management plane include Simple Network Management Protocol (SNMP), Telnet, File Transfer Protocol (FTP), Secure FTP, and Secure Shell (SSH). These management protocols are used for monitoring and for command-line interface (CLI) access.

Figure 3-12 shows how the three planes of operation and how the processes are isolated from each other. In IOS XR, a process failure within one plane does not affect other processes or applications within that plane. This layered architecture creates a more reliable model than one with a monolithic architecture such as IOS, where failure of a single process may cause a failure of the whole system.

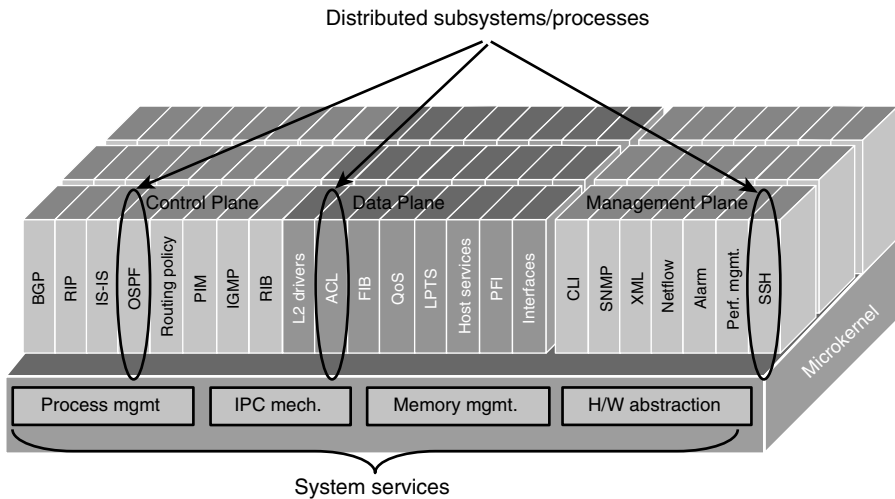


Figure 3-12 Separation of Control, Data, and Management Planes

Summary

This chapter provided an overview of the fundamentals of IP routing and IP switching and the control planes of operation. In summary, it showed how a router makes a forwarding decision, which consists of three basic components:

- The routing protocols, which are used to build the routing table (RIB)
- The routing table, which is used to program the switching mechanisms (that is, CEF)
- The switching mechanisms used to perform the actual packet forwarding

References in This Chapter

Bollapragada, Vijay, Russ White, and Curtis Murphy. *Inside Cisco IOS Software Architecture*. (ISBN-13: 978-1587058165).

Stringfield, Nakiya, Russ White, and Stacia McKee. *Cisco Express Forwarding*. (ISBN-13: 978-0-13-343334-0).

Tahir, Mobeen, Mark Ghattas, Dawit Birhanu, and Syed Natif Nawaz. *Cisco IOS XR Fundamentals*. (ISBN-13: 978-1-58705-271-2).

Doyle, Jeff and Jennifer Carroll. *Routing TCP/IP, Volume 1*, 2nd Ed. (ISBN-13: 978-1-58705-202-6).

This page intentionally left blank

Index

Symbols

* (asterisks) query modifier, 482
[] (brackets) query modifier, 479
^] (caret in brackets) query modifier, 480
^ (carets) query modifier, 478
\$ (dollar signs) query modifier, 478-479
- (hyphen) query modifier, 479-480
() (parentheses) query modifier, 480
. (period) query modifier, 481
| (pipe) query modifier, 480
+ (plus sign) query modifier, 481
? (question mark) query modifier, 481-482
_ (underscores) query modifier, 477-478
2-Way state (OSPF neighbors), 196
10.55.200.33/12 address range calculations, 45
172.16.2.3/23 address range calculations, 45
192.168.100.0/24 subnet challenge, 46

available hosts, 47
available subnets, 47
binary bit values, 47
final address allocation scheme, 49
LAN 1/LAN 2 subnet assignments, 47-48
LAN 3 assignments, 48
192.168.100.5/26 address range calculations, 45

A

abbreviations (IPv6 addresses), 895-896
ABF (access-list-based forwarding), 521
configuring, 523-525
overview, 521
ABRs (Area Border Routers), 242
totally stubby areas, 289
Type 3 LSAs, 259
accumulated interior gateway protocol. *See* AIGP
ACEs (access control entries), 467

ACLs (access control lists), 467

ABF, 521

*configuring, 523-525**overview, 521*

ACEs, 467

AS_Path, 484

categories, 468

extended

*BGP network selection, 470**defining, 469**IGP network selection, 469-470*

standard, 468-469

*ACL to network entries, 469-468**defining, 468***actions (route maps), 492****active neighbor field (OSPF Hello packet), 194****Active state (BGP neighbors), 416****AD (administrative distance), 73-76, 91**

default values, 526

floating static routes, verifying,
104-105

modifying, 527

*BGP, 532-534**EIGRP, 528-529**IS-IS, 531-532**OSPF, 529-531*redistribution routing loop preven-
tion, 601-603

suboptimal routing, 527, 584

Add-Path feature (BGP), 739-742**address families**

EIGRP configurations

*instances, 185-186**interfaces, 186-188**topology, 188*

identifiers (AFIs)

*BGP, 410**listing of, website, 812**MBGP, 812*

address-family command, 77

address-family ipv6 command, 954

address-family ipv6 unicast

autonomous-system as-number
command, 954address-family ipv6 unicast
command, 988, 1001**addresses**

anycast, 906-908

*duplicate address detection,
disabling, 908**topology, 906-908**updated topology, 907*

broadcast

*calculating, 43**defined, 37*

destination, 1012

host

*all-0s, 37**available, listing, 47**first usable, calculating, 43**last usable, calculating, 43*IP. *See* IP addresses**IPv6***abbreviations, 895-896**bit length values, 894**case, 894**components, 894**hexadecimal to binary conver-
sion, 896-898**hextets, 894**network boundaries, 894**unicast, 894*

- loopback, 911
- MAC
 - defined*, 752
 - multicast*, 752-753
- multicast
 - Administratively Scoped Block*, 751
 - GLOP Block*, 751
 - IANA assigned*, 749
 - Internetwork Control Block*, 750
 - IPv6*, 908-911
 - Layer 2*, 752-753
 - Local Network Control Block*, 750
 - reserved*, 750
 - Source-Specific Multicast Block*, 751
- NET, 322-323
 - common private*, 322
 - expanded structure*, 322
 - guidelines*, 323
 - minimal format*, 322
 - multiple*, 323
- next-hop
 - eBGP*, 444-445
 - MBGP for IPv6*, 997
- OSPF multicast, 193
- resolution, 934-936
- unicast
 - link-local (LLA)*, 905-906
 - unique local unicast (ULA)*, 904-905
- unspecified, 911
- Adj-RIB-Out table neighbor-specific view**, 429-430
- adjacencies (neighbors)**
 - EIGRP, 128
 - advertised networks update packet*, 131
 - hello packet capture*, 129
 - INIT flag and acknowledgment update packet*, 130
 - INIT flag set update packet*, 129
 - route prefixes update packet*, 130
 - sample topology*, 129
 - verification*, 139
 - EIGRPv6 verification, 949
 - IS-IS, 318
 - broadcast*, 333
 - compatibility chart*, 386
 - P2P*, 338-339
 - verification*, 346-347
 - OSPF, 197-202
 - DBD packet*, 199
 - debug perspective example*, 201-202
 - Hello packet with neighbors detected*, 198
 - Hello packet with no neighbors detected*, 197
 - hub-and-spoke topology*, 228
 - LSR*, 200
 - LSU*, 200
 - multi-area*, 304-308
 - network type compatibility*, 231-234
 - parameters, verifying*, 198-199
 - process*, 201
 - simple topology*, 197
 - verifying*, 209-210
 - OSPFv2 versus OSPFv3, 957

- Adjacency Information Base (AIB)**, 87
- admin mode (IOS XR)**, 14
- administrative distance**. *See* AD
- Administratively Scoped Block**, 751
- administratively scoped multicast boundaries**, 863-864
- advertisements**
 - BGP
 - conditional*, 645-647
 - default*, 643-644
 - default per neighbor*, 644-645
 - candidate RP (C-RP-Adv), 775
 - NDP RA messages, 913
- af-interface command**, 185
- af-interface interface-type interface-number command**, 954
- AFI (Authority and Format Identifier)**, 320
- AFIs (address family identifiers)**
 - BGP, 410
 - listing of, website, 812
 - MBGP, 812
- age (OSPF LSAs)**, 251
- aggregate-addressipv6-prefix/prefix-length command**, 1001
- aggregation**. *See also* summarization
 - addresses, 629-631
 - AS_Set, 639-641
 - atomic, 637-639
 - IPv6, 902-903
 - link aggregation (LAG), 722
 - selective AS_Set, 641-643
- AIB (Adjacency Information Base)**, 87
- AIGP (accumulated interior gateway protocol)**, 686-694
 - configuration, 689
 - guidelines, 688
 - MED conversion, 693-694
 - metric modifications, 688
 - neighbor sessions metrics verification, 690
 - PAs, exchanging, 687
 - processing logic, 692-719
- algorithms**
 - BGP best path, 672
 - AIGP*. *See* AIGP
 - AS_Path length*, 694-698
 - attribute classifications*, 673
 - attributes list*, 672
 - eBGP over iBGP*, 714-717
 - local preference*, 679-682
 - locally originated via network or aggregate advertisement*, 683-686
 - lowest IGP metric*, 718-720
 - lowest neighbor address*, 722
 - MED*. *See* MED, BGP best path selection
 - minimum cluster list length*, 721
 - oldest EBGP*, 720
 - origin types*, 700-703
 - RIDs*, 720-721
 - weight*, 673-679
 - BSR RP hash, 845-846
 - Dijkstra shortest path first, 70
 - distance vector, 69
 - enhanced distance vector, 70
 - link-state, 70-71
 - path vector, 71-72
 - RPF. *See* RPF
- allowing ASNs**, 658-660
- and keyword (conjunction operator)**, 41-42, 504
- any source multicast (ASM)**, 850

anycast IPv6 addresses, 906-908

- duplicate address detection,
 - disabling, 908
- RPs, 847-849
- topology, 906-908
- updated topology, 907

architecture

- centralized versus distributed forwarding, 86
- IOS, 1
 - kernel*, 2
 - memory management*, 2
 - run to completion scheduler*, 2
 - software packaging*, 2-4

IOS XE, 4

- kernel*, 4
- memory management*, 4
- scheduling*, 4

IOS XR, 5

- kernel*, 5
- memory management*, 6
- scheduling*, 5
- software packaging*, 6-7

IS-IS, 317

multicast, 745

OSPF, 192-193

prefix matching, 471

route policies, 499-504

- advanced conditional statement nesting*, 502
- bad RPL design*, 503
- conditional match and action statement*, 499
- conditional match with if-else logic*, 500
- conditional match with if-elseif logic*, 501

conditional match with if-elseif-else logic, 501

conditional statement nesting, 502

good RPL design, 503

if-elseif-else route-policy, 502

inline policy set expansion, 500

inline prefix filtering, 499

nested conditional match, 503

prefix set filtering, 500

RFC 1918 route map with if-else logic, 501

RPL, 496

archiving configurations, 11-13

area area-id authentication ipsec spi spi authentication-algorithm commands, 975

area area-id encryption ipsec spi spi esp encryption-algorithm command, 975

area area-id nssa command, 975

area area-id range ipv6-prefix/prefix-length command, 975

area area-id stub command, 975

area area-id stub no-summary command, 975

Area Border Routers. See ABRs

area ID field (OSPF Hello packet), 194

areas (IS-IS), 318-319

areas (OSPF)

ABRs, 242

backbone, 242

defined, 241

disadvantages, 242

failed route advertisement, 242

filtering, 543-546

IDs, 245-246

multi-area adjacencies, 304-308

- configuration*, 305-306
- inefficient topologies*, 304
- nonbroadcast interface verification*, 304
- OSPF neighborhood*, 306-307
- verification*, 306-308
- multi-area topology, 243-245
- NSSAs, 268-269
- OSPFv3 authentication/encryption, 972
- route types
 - external*, 247-248
 - intra-area/interarea*, 246-247
- stubby, 286
 - NSSAs, 292-295
 - stub areas*, 286-289
 - totally*, 289-292
 - totally NSSAs*, 295-298
 - types*, 286
- successful route advertisement, 243
- summarization, 276-280
- Type 1 LSA flooding, 252
- ASBRs (Autonomous System Boundary Routers)**, 247, 265-268
- AS-external LSAs**, 959
- ASM (any source multicast)**, 850
- ASNs (autonomous system numbers)**, 408-409
 - allowing, 658-660
 - local, 660-664
 - private, removing, 656-658
- AS_Paths**, 488
 - ACL (IOS), 484
 - BGP best path algorithm, 694-698
 - configuration*, 696
 - processing logic*, 698
 - IOS XR selection options, 484
 - is-local*, 485
 - length*, 485-486
 - neighbor-is*, 487
 - originates-from option*, 487
 - passes-through*, 486-487
 - unique-length*, 486
 - Assert PIM control message**, 764
 - AS_Set**, 641-643
 - as-set keyword**, 639-641
 - asterisk (*) query modifier**, 482
 - asynchronous mode (BFD)**
 - atomic aggregate attribute**, 637-639
 - attached bits (LSPs)**, 331
 - attachment points**, 516
 - Attempt state (OSPF neighbors)**, 196
 - attributes**
 - atomic aggregate, 637-639
 - BGP best path, 409
 - AIGP*. See *AIGP*
 - AS_Path length*, 694-698
 - classifications*, 673
 - eBGP over iBGP*, 714-717
 - listing of*, 672
 - locally originated via networks or aggregate advertisement*, 683-686
 - local preference*, 679-682
 - lowest IGP metric*, 718-720
 - lowest neighbor address*, 722
 - MED*. See *MED*, *BGP best path selection*
 - minimum cluster list length*, 721
 - oldest EBGP*, 720
 - origin*, 700-703
 - RIDs*, 720-721
 - weight*, 673-679

- eBGP/iBGP multipath, 725
 - eBGP prefix
 - listing of*, 441
 - local*, 441
 - remote*, 440
 - EIGRP, propagation, 146
 - LSPs, 331
 - route policies, modification, 498
 - authentication**
 - BGP, 462-463
 - EIGRP, 174
 - enabling*, 174
 - key chain configuration*, 174-177
 - IS-IS, 367
 - configuration*, 369-370
 - hello*, 368
 - LSP*, 368
 - types*, 367
 - OSPF, 236
 - IOS*, 236
 - IOS XR*, 237-239
 - types supported*, 236
 - OSPFv2 versus OSPFv3, 955
 - OSPFv3, 970-973
 - areas*, 972
 - interfaces*, 972
 - IPsec verification*, 973
 - authentication command, 975
 - authentication key-chain key-chain-name command, 186
 - authentication keychain key-chain-name command, 954
 - authentication mode command, 186
 - authentication options field (OSPF Hello packet), 194
 - Authority and Format Identifier (AFI), 320
 - Autonomous System Boundary Routers (ASBRs), 247, 265-268
 - autonomous systems, 67
 - EIGRP, 126
 - installed routes, displaying*, 140-141
 - interface verification*, 136-139
 - IOS*, 132-133
 - IOS XR*, 134
 - neighbor adjacencies verification*, 139
 - passive interfaces*, 134
 - sample topology*, 134-136
 - maximum BGP, 652-654
 - numbers (ASNs)
 - allowing*, 658-660
 - local*, 660-664
 - private, removing*, 656-658
 - Auto-RPs**, 773
 - candidate RPs, 773
 - Cisco-RP-announce message filtering, 867
 - configuring, 785-786
 - multicast boundaries, 865-866
 - RP MAs, 774
 - TTL scoping, 862
 - auto-summary command, 188
-
- B**
- backbone, 242, 380-382
 - backdoor networks, 649-652
 - backup connectivity, 101
 - bandwidth (EIGRP), 177-179

bandwidth-percent percent command, 186

BDRs (backup designated routers), 213

LSA distribution, 213

OSPF

elections, 214-216

placement, 216-219

Bellman-Ford algorithms, 69

BGP (Border Gateway Protocol), 68

AD, modifying, 532-534

address families, 410

Adj-RIB-Out table neighbor-specific view, 429-430

advertisements

conditional, 645-647

default, 643-644

default per neighbor, 644-645

aggregation

aggregate addresses, 629-631

AS_Set, 639-641

atomic, 637-639

selective AS_Set, 641-643

allow autonomous feature, 658-660

ASNs, 408-409

backdoor networks, 649-652

best path algorithm, 672

AIGP. See AIGP

AS_Path length, 694-698

attribute classifications, 673

attributes list, 672

eBGP versus iBGP, 714-717

locally originated via network or aggregate advertisement, 683-686

local preference, 679-682

lowest IGP metric, 718-720

lowest neighbor address, 722

MED. See MED, BGP best path selection

minimum cluster list length, 721

oldest EBGP, 720

origin types, 700-703

RIDs, 720-721

weight, 673-679

communities, 609

conditionally matching, 620-627

formats, 610

Internet, 611

No_Advertise, 614-617

No_Export, 611-614

No_Export_SubConfed, 617-620

support, 611

confederations, 453-458

AS100 BGP table, 457

configuring, 455-457

iBGP versus eBGP sessions, 454

NLRI, 458

topology, 453

configuring, 418-419

IOS, 419-420

IOS XR, 420-421

defined, 407

eBGP

BGP table, 440

configuration, 439

iBGP combinations, 442-444

iBGP sessions, compared, 438

local prefix attributes, 441

next-hop addresses, 444-445

prefix attributes, 441

remote prefix attributes, 440

topology, 438

- ECMP, 723
 - AS_Path relax feature*, 731-733
 - eBGP/iBGP multipath*, 723-726
 - eiBGP multipath*. See *eiBGP multipath*
- failure detection, 459
- fast convergence, 125, 189, 221, 315, 366, 459
- filtering, 546-548
- iBGP. See iBGP
- inter-router communication, 410-411
 - keepalive messages*, 413
 - message types*, 411
 - notification messages*, 414
 - open messages*, 412-413
 - update messages*, 413-414
- IPv4 neighbor output, 423-424
- local autonomous feature, 660-664
- loop prevention, 409-410
- maximum autonomous system, 652-654
- maximum prefix, 654-656
- Multiprotocol. See MBGP
- neighbors
 - AIGP metrics support, verifying*, 690
 - default advertisements*, 644-645
 - lowest address (path selection)*, 722
- neighbor states, 415
 - Active*, 416
 - Connect*, 415-416
 - Established*, 417
 - idle*, 415
 - OpenConfirm*, 417
 - OpenSent*, 416-417
- OSPF preferred over BGP for RPF calculation, 878
- outbound filtering, 647-649
- PAAs, 409
- peer configuration
 - IOS groups*, 664-665
 - IOS templates*, 665-666
 - IOS XR templates*, 667-668
- prefix advertisements, 425-427
- redistribution
 - destination-specific behaviors*, 580-582
 - source-specific behaviors*, 562-563
- regex reference topology, 476
- remove private autonomous system feature, 656-658
- routes
 - displaying*, 428-431
 - receiving*, 427
 - summary with prefixes*, 430
- RPL verification
 - BGP table*, 517-518
 - redistribution*, 516-517
- security, 459
 - authentication*, 462-463
 - eBGP multihop*, 459-461
 - TTL*, 461-462
- sessions
 - clearing*, 549
 - IPv6 over IPv4*, 998-1001
 - overview*, 415
 - verification*, 421-424
- suboptimal routing with RRs, 733-734
 - Add-Path feature*, 739-742
 - RRs, adding*, 734-735

- shadow RRs*, 735-737
- shadow session RRs*, 738-739
- summarization, 628, 632-637
- summary fields, 422
- table fields, 429
- bgp router-id router-id command**, 1001
- bidirectional connectivity verification**, 97-98
- Bidirectional Forwarding Detection**. *See* BFD
- Bidir-PIM (bidirectional PIM)**
 - designated forwarders, 804-808
 - overview, 802-808
- binary notation**
 - base 2 calculations, 31-32
 - bit values, 31-32, 47
 - to decimal conversion, 33
 - decimal to binary conversions, 32
 - defined, 31
 - IPv6 hexadecimal conversion, 896-898
 - subnet mask conversion, 36
- bits**
 - overload, 394-396
 - values, 31-32, 47
- bitwise AND operator**, 41-42
- boolean operators**, 504
 - conjunction, 504
 - disjunction, 505
 - negation, 504
 - order of processing, 505-506
- Bootstrap PIM control message**, 764
- bootstrap router**. *See* BSR
- Border Gateway Protocol**. *See* BGP
- boundaries**
 - IPv6, 894, 1032-1033
 - multicast, 863-866
 - administratively scoped*, 863-864
 - Auto-RP*, 865-866
 - BSR*, 866
- brackets ([]) query modifier**, 479
- broadcast addresses**
 - calculating, 43
 - defined, 37
- broadcast interfaces**
 - as P2P interfaces, configuring, 353-355
 - static routing, 98-99
- broadcast networks**
 - IS-IS, 326, 333-338
 - adjacency process*, 333
 - IIH hello with neighbors detected*, 337
 - IIH hello with no neighbors detected*, 335
 - parameter verification*, 336
 - simple topology*, 334
 - OSPF, 221
 - video feed, 747
- BSR (bootstrap router)**, 775-776
 - configuring, 786-787
 - IPv6 PIM-SM, 1018-1021
 - BSR election*, 1020
 - candidate RP*, 1021
 - embedded RP*, 1021-1024
 - IOS XR configuration*, 1020
 - RP cache*, 1020
 - multicast boundaries, 866
 - redundant RPs, 840-846
 - group filtering*, 843-845
 - hash algorithm*, 845-846

C

calculating

- broadcast addresses, 43
- CIDR best paths, 54-55
- EIGRP paths, 145
 - attribute propagation, 146*
 - custom K values, 148-149*
 - default K values, 145*
 - with definitions, 145*
 - formula, 145*
 - interface delay settings, 149-151*
 - interface metrics, 146*
 - load balancing, 151-153*
 - lowest link speed/cumulative delay, 146*
 - specific prefix, 147-148*
 - variance multiplier, 151-152*
 - wide metrics, 153-154*

host addresses

- first usable, 43*
- last usable, 43*

OSPF interface costs, 235-236

SPF

- IS-IS, 361-362*
- Type 1 LSAs, 256*

subnets

- bitwise AND operations, 41-42*
- magic number method, 42-45*

usable IP addresses, 37-38

wildcard subnet masks, 63

- /24 networks, 64*
- single IP hosts, 63*
- summary routes, 64*

candidate RP advertisement (C-RP-Adv), 764, 775

candidate RPs (C-RPs), 773, 1021-1024

caret in brackets [^] query modifier, 480

carets (^) query modifier, 478

CEF (Cisco Express Forwarding), 86

- centralized versus distributed forwarding architectures, 86

defined, 86

distributed, 88

hardware, 88-89

IPv6, enabling, 899

software, 87

tables

- nonrecursive multihop static routes, 107*
- static route recursion problem, 114*

centralized forwarding architecture, 86

CE router connectivity

- default route configuration, 100
- ISP to, 100

CIDR (classless interdomain routing), 35

best path calculation, 54-55

IP addresses interface assignments, 60

route summarization, 56-58

- no summarization example, 56*
- with summarization example, 57*
- summary prefixes, 57-58*

updates, 54

VLSMs

- classless routing update, 55*
- defined, 55*
- successful route convergence, 55-56*

Cisco

Express Forwarding. *See* CEF
 feature navigator tool website, 3

Cisco-RP-announce message filtering, 867

clarity (RPL), 496

classes (IP addresses)

A, 34
 B, 35
 C, 35
 D, 35
 history, 34
 listing of, 34

classful routing

best path confusion, 53
 discontinuous networks, 53
 local decision process, 50-51
 subnet masks
 uniform, 50-51
 variable-length, 52
 updates, 50

classic mode (EIGRPv6), 945

classless interdomain routing. *See* CIDR

clear bgp ipv6 unicast command, 1001

CLI (command-line interface)

IOS modes, 9-11
 global configuration, 10-11
 privileged, 10
 user, 9-10
 IOS XR modes
 admin, 14
 EXEC, 14
 global configuration, 14

cluster lists (BGP), 721

commands

address-family, 77
 address-family ipv6, 954
 address-family ipv6 unicast, 988, 1001
 address-family ipv6 unicast autonomous-system as-number, 954
 af-interface, 185
 af-interface interface-type interface-number, 954
 aggregate-addressipv6-prefix/prefix-length, 1001
 area area-id authentication ipsec spi spi authentication-algorithm, 975
 area area-id encryption ipsec spi spi esp encryption-algorithm, 975
 area area-id nssa, 975
 area area-id range ipv6-prefix/prefix-length, 975
 area area-id stub, 975
 area area-id stub no summary, 975
 authentication, 975
 authentication key-chain key-chain-name, 186
 authentication keychain key-chain-name, 954
 authentication mode, 186
 auto-summary, 188
 bandwidth-percent percent, 186
 bgp router-id router-id, 1001
 clear bgp ipv6 unicast, 1001
 commit, 17
 commit label, 18
 default-information originate, 975, 988
 default-originate, 1001
 distribute-list prefix-list list-name interface-number, 954

distribute-list prefix-list list-name
 interface-type interface-number,
 954
 eigrp router-id router-id, 185
 eigrp stub, 185
 encryption, 975
 hello-interval seconds, 186
 hold-time seconds, 186
 if as-path, 497
 if as-path in, 497
 if as-path is-local, 497
 if destination in, 497
 if local-preference, 497
 if med, 497
 if next-hop in, 497
 if origin is, 497
 if tag, 497, 498
 instance instance-id, 975
 interface interface-type interface-
 number, 19.225, 19.579, 19.783
 ip address ip-address subnet-mask, 60
 ip ospf process-id area area-id, 204
 ipv4 address ip-address subnet
 mask, 60
 ipv4 address ipv4-address prefix-
 length, 60
 ipv6 authentication key-chain eigrp
 as-number key-chain-name, 954
 ipv6 authentication mode eigrp as-
 number md5, 954
 ipv6 router eigrp as-number, 954
 ipv6 router isis, 988
 ipv6 summary-address eigrp as-num-
 ber ipv6-prefix/prefix-length, 954
 ip vrf forwarding vrf-name, 77
 ip vrf vrf-name, 77
 isis ipv6 metric, 988
 match as-path acl-number, 490
 match ip address, 490
 match ip address prefix-list prefix-
 list-name, 490
 match local-preference, 490
 match metric, 490
 match tag tag-value, 490
 maximum-paths, 188
 metric, 988
 metric-style, 988
 metric weights TOS K1 K2 K3 K4
 K5, 186
 multitopology, 988
 neighbor, 1001
 neighbor remote-as as-number, 1001
 neighbor ip-address local-as-alter-
 nate-as-number,
 neighbor activate, 1001
 net network-entity-title, 988
 networkipv6-prefix/prefix-length,
 1001
 network network, 185
 next-hop-self, 186
 nssa, 975
 nssa [default-information-originate]
 no-summary, 975
 ospfv3 authentication, 975
 ospfv3 encryption, 975
 ospfv3 process-id area area-id, 975
 ospfv3 process-id area area-ID, 975
 passive-interface, 186
 prepend as-path, 498
 pwd, 26
 range ipv6-prefix/prefix-length, 975
 remote-as as-number, 1001
 replace as-path, 498

- root, 26
- route-policy route-policy-name, 954, 1001
- router bgp as-number, 1001
- router eigrp as-number, 132, 954
- router eigrp process-name, 954
- router isis, 988
- router ospfv3 process-id, 975
- set local-preference, 498
- set med, 498
- set next-hop, 498
- set origin, 498
- set weight, 498
- show bgp ipv6 unicast, 993
- show bgp ipv6 unicast neighbors, 993
- show bgp ipv6 unicast summary, 993
- show clns interface, 985
- show clns neighbor, 985
- show configuration, 15
- show configuration commit changes, 18
- show configuration commit changes incremental, 18
- show configuration commit list, 17
- show configuration failed, 20
- show configuration merge, 16
- show interface interface-type interface-number, 149
- show ip eigrp interface, 136
- show ip interface, 61
- show ip route vrf vrf-name, 81
- show ipv4 route vrf-name, 81
- show ipv6 mld groups, 1024
- show ipv6 mld interface, 1024
- show ipv6 mld traffic, 1024
- show ipv6 mrrib route, 1024
- show ipv6 mroute, 1024
- show ipv6 mroute active, 1024
- show ipv6 pim bsr candidate-rp, 1024
- show ipv6 pim bsr election, 1024
- show ipv6 pim bsr rp-cache, 1024
- show ipv6 pim group-map, 1024
- show ipv6 pim interface, 1024
- show ipv6 pim neighbor, 1024
- show ipv6 pim range-list, 1024
- show ipv6 pim topology, 1024
- show ipv6 pim traffic, 1024
- show ipv6 pim tunnel, 1024
- show ipv6 protocols, 985
- show ipv6 route bgp, 993
- show ipv6 route ospf, 962
- show ipv6 rpf ipv6-address, 1024
- show isis database, 985
- show isis interface, 985
- show isis neighbor, 985
- show isis neighbors, 985
- show mfib route, 798
- show mfib route rate, 1024
- show mfib route src-ip-address/group-address, 886
- show mld groups, 1024
- show mld interface, 1024
- show mld traffic, 1024
- show mrrib ipv6 route, 1024
- show mrrib route, 798
- show ospfv3 database, 962
- show ospfv3 interface, 962
- show ospfv3 neighbor, 962
- show ospfv3 virtual-links, 962
- show pim ipv6 bsr candidate rp, 1024
- show pim ipv6 bsr election, 1024
- show pim ipv6 bsr rp-cache, 1024
- show pim ipv6 group-map, 1024

show pim ipv6 interface, 1024
 show pim ipv6 neighbor, 1024
 show pim ipv6 range-list, 1024
 show pim ipv6 rpf ipv6-address, 1024
 show pim ipv6 topology, 1024
 show pim ipv6 traffic, 1024
 show pim ipv6 tunnel info, 1024
 show pim topology, 799
 show protocols ipv6, 985
 show route eigrp, 163-165
 show route ipv6 bgp, 993
 show route ipv6 ospf, 962
 show running-config, 15
 show the commit, 23
 single-topology, 988
 split-horizon, 186
 stub, 975
 stub no-summary, 975
 summary-address ipv6-prefix/prefix-length, 954, 975
 summary-address network, 186
 summary-address prefix/prefix-length, 988
 summary-metric network bandwidth delay reliability load MTU [AD], 188
 summary-prefix prefix/prefix-length, 988
 timers active-time, 188
 topology base, 185, 954
 variance variance-multiplier, 188
 vrf definition vrf-name, 77
 vrf upgrade-cli multi-af-mode, 77
commit command, 17
commit label command, 18

commits (configurations), 16-17
 changes, displaying, 18
 commit lists, displaying, 17
 confirming, 22-23
 failures, 20-21
 labels, 18-19
 lists, displaying, 17
 multiple commit options, 23-24
 replaces, 19-20
communities (BGP), 609
 conditionally matching, 620-627
 community sets, 621
 inline, 622-625
 private BGP communities, 625-627
 formats, 610
 support, 611
 well-known
 Internet, 611
 No_Advertise, 614-617
 No_Export, 611-614
 No_Export_SubConfed, 617-620
complete sequence number packets (CSNPs), 350
complex matching route maps, 491-492
conditional matching, 490
 BGP communities, 620-627
 community sets, 621
 inline, 622-625
 private BGP communities, 625-627
 complex matching, 491-492
 multiple match, 491

conditional routing

ABF

*configuring, 523-525**overview, 521*

BGP, 645-647

disadvantages, 522

PBR

*configuring, 522**local, 525-526**overview, 521***confederations (BGP), 453-458**

AS100 BGP table, 457

configuring, 455-457

iBGP versions eBGP sessions, 454

NLRI, 458

topology, 453

configurations

archiving, 11-13

ABF, 523-525

BGP, 418-419

*authentication, 462**confederations, 455-457**IOS, 419-420**IOS peer groups, 664-665**IOS peer templates, 665-666**IOS XR, 420-421**IOS XR templates, 667-668**maximum prefix, 655**ORF, 648*

BSR, 786-787

changes

*displaying in SysDB, 17-18**rolling back, 21-22*

commits, 16-17

*changes, displaying, 18**confirming, 22-23**failures, 20-21**labels, 18-19**list, displaying, 17**multiple commit options, 23-24**replaces, 19-20*

default routes

*CE router, 100**XR1/R3 example, 101-102*

DNS SSM mapping, 859

eBGP, 438, 439

*multihop sessions, 460**multipath, 724*

EIGRP

*bandwidth percent, 178**installed routes, displaying,
140-141**interface delays, 150**interface verification, 136-139**IOS, 132-133**IOS XR, 134**key chains, 174-177**named configurations. See
EIGRP, named configura-
tions**neighbor adjacencies verifica-
tion, 139**passive interfaces, 134**sample topology, 134-136**stubs, 162-163*

EIGRPv6

*classic mode, 945**IOS XR, 946-947**named mode, 946*

external route summarization, 282

files, loading, 24

hierarchical, 25-26

iBGP, 724

- interarea route summarization, 279
- IOS relay agents, 927-928
- IOS XR
 - relay proxy agents*, 928-929
 - two-stage commit*, 14-17
- IPv6
 - multicast*, 1025
 - static*, 942-943
 - unicast*, 899
- IS-IS, 340
 - authentication*, 369-370
 - broadcast interfaces as P2P*, 353-355
 - default routes*, 400
 - DISs*, 352
 - interfaces, verifying*, 343-346
 - IOS, 340
 - IOS XR, 340-341
 - metrics*, 391
 - neighbor adjacencies, verifying*, 346-347
 - overload bits*, 395
 - prefix suppression*, 403
 - route-leaking*, 377
 - route verification*, 347-348
 - summarization*, 399
 - topology example*, 341-343
- IS-IS for IPv6, 979
 - IOS base, 979-980
 - IOS topology mode, 981
 - IOS XR base, 980
 - IOS XR topology mode, 981-984
- MBGP, 815
 - IOS, 991-992
 - IOS XR, 992-993
- merging, 16
- multicast, 780-782, 786-787
- multiple recursive lookups, 110
- nested, exiting, 26
- nonrecursive multihop static routes, 107
- OSPF
 - IOS interface specific*, 204-205
 - IOS network statement*, 202-204
 - IOS XR, 205
 - multi-area adjacencies*, 305-306
 - NSSAs, 293-294
 - point-to-multipoint network*, 226
 - prefix suppression*, 310
 - stub areas*, 287-288
 - topology sample*, 206-208
 - totally NSSAs*, 297
 - totally stubby areas*, 290-291
 - virtual links*, 300
- OSPFv3
 - IOS, 960-961
 - IOS XR, 961
- PBR, 522
- prompts
 - IOS, 26
 - IOS XR, 26
- replacing, 13
- route reflectors, 449-450
- RPs, 783-787
 - Auto-RPs*, 785-786
 - static*, 784-785
- running, displaying, 15
- SSM, 852-853, 860
- stateful DHCPv6
 - IOS, 930
 - IOS XR, 931-932

- stateless DHCPv6
 - IOS*, 924-925
 - IOS XR*, 925-926
- target, displaying, 15
- VRFs
 - example*, 81
 - IOS*, 78
 - IOS XR*, 78-79
- confirmed parameter, 23
- confirming commits, 22-23
- conjunction operator, 504
- connected networks
 - defined, 91
 - local routes, 93
 - redistribution, 561
 - routing table sample, 92-93
 - sample topology, 92
 - secondary, 94
 - specific network routing detail
 - output sample, 93
- connections
 - BGP, clearing, 549
 - ES-IS/IS-IS, 316
 - OSPF prefix suppression, verifying, 312
- Connect state (BGP neighbors), 415-416
- continue keyword, 493-494
 - EIGRP, 158
- C-RP-Adv (candidate RP advertisement), 775
- C-RPs (candidate RPs), 773
- CSNPs (complete sequence number packets), 350

D

- DAD (duplicate address detection), 937-938
- databases
 - LSDB, 70, 192
 - OSPF Type 1 LSA fields*, 252
 - OSPFv3*, 965-969
 - LSPDB (IS-IS)
 - building topology*, 359-360
 - displaying*, 356-357
 - displaying topology*, 360-361
 - non-pseudonode LSPs*, 357-358
 - pseudonode LSPs*, 358-359
 - SPF calculations*, 361-362
 - without transit networks*, 404-405
- data plane, 89
- DBD packets
 - OSPF, 194
 - OSPFv3, 957
- dCEF (distributed CEF), 88
- dead interval field (OSPF Hello packet), 194
- dead interval timer, 219
- debugging IOS XR traces, 8
- default-information originate
 - command, 975, 988
- default-originate command, 1001
- default routes, 99-102
 - backup connectivity, 101
 - BGP advertisements
 - overview*, 643-644
 - per neighbor*, 644-645
 - CE router
 - configuration*, 100
 - routing tables*, 100-101

- EIGRPv6, 952-953
- IS-IS, 400-401
- OSPF summarization, 283-285
- XR1/R3
 - configuration example*, 101-102
 - routing tables*, 102
- Dense Mode (PIM), 765-767
- design
 - EIGRP stubs, 164-166
 - subnetting, 46
 - available host values, listing*, 47
 - available subnets, listing*, 47
 - binary bit values, listing*, 47
 - final address allocation scheme*, 49
 - LAN 1/LAN 2 subnet assignments*, 47-48
 - LAN 3 assignments*, 48
 - designated forwarders (DFs), 764, 804-808
 - designated intermediate systems. *See* DISs
 - designated router & backup designated router field (OSPF Hello packet), 194
 - designated routers. *See* DRs
 - destination addresses (MLD), 1012
 - destination-specific redistribution behaviors
 - BGP, 580-582
 - EIGRP, 563-566
 - EIGRP-to-EIGRP, 566-568
 - IS-IS, 576-578
 - IS-IS to IS-IS, 578-580
 - OSPF, 569-571
 - forwarding addresses*, 573-576
 - OSPF-to-OSPF*, 571-572
 - DFs (designated forwarders), 764, 804-808
 - DHCPv6
 - stateful, 926-927
 - IOS relay agent configuration*, 927-928
 - IOS relay agent verification*, 928
 - IOS server configuration*, 930
 - IOS XR relay proxy agent configuration*, 928-929
 - IOS XR relay proxy agent verification*, 929
 - IOS XR server configuration*, 931-932
 - server verification*, 932-934
 - topology*, 926
 - stateless, 924-926
 - IOS configuration*, 924-925
 - IOS XR configuration*, 925-926
 - verification*, 926
 - Diffused Update Algorithm (DUAL), 70
 - Dijkstra shortest path first algorithm, 70
 - discard routes, 399
 - discontiguous networks, 301-303
 - disjunction operator, 505
 - displaying
 - BGP
 - routes*, 428-431
 - configurations
 - changes in SysDB*, 17-18
 - commit changes*, 18
 - commit lists*, 17
 - running*, 15
 - target*, 15
 - EIGRP routes, 140-141
 - interfaces/IP addresses status verification, 82-83

- IOS configuration archive, 12
- IPv4 interfaces
 - specific*, 61
 - summaries*, 62
- IS-IS
 - LSPDB*, 356-359
 - LSP topology*, 360-361
- NDP parameters, 916-917
- OSPFv3 LSDB summary view, 969
- route policy states, 515-516
- VRFs, 81-82
- DISs (designated intermediate systems), 348**
 - configuration, 352
 - CSNPs, 350
 - elections, 351
 - logical drawing, 348
 - LSP advertisement, 349
 - placement, 352-353
 - pseudonode LSPs, 349-350
 - verification, 353
- distance vector algorithms, 69**
- distributed CEF (dCEF), 88**
- distributed forwarding architecture, 86**
- distribute-list prefix-list list-name interface-number command, 954**
- distribute-list prefix-list list-name interface-type interface-number command, 954**
- distribution trees (multicast), 759**
 - shared trees, 761
 - source trees, 759-760
- DNS**
 - SLAAC RA options, 923-924
 - SSM mapping, 857-860
 - configuration*, 859
 - topology*, 858
 - verification*, 859
 - stateless DHCPv6, 924-926
 - IOS configuration*, 924-925
 - IOS XR configuration*, 925-926
 - verification*, 926
- DNSSEC (DNS Search List), 923**
- dollar sign (\$) query modifier, 478-479**
- Domain Specific Part (DSP), 320-321**
- done messages (MLD), 1010**
- dot-decimal notation**
 - binary conversions to, 33
 - converting to binary, 32
 - defined, 30
 - octet subtraction, 63
 - /24 networks*, 64
 - single IP hosts*, 63
 - summary routes*, 64
 - wildcard subnet masks, 62
- Down state (OSPF neighbors), 196**
- downstream interface, 763**
- DRs (designated routers), 213**
 - election, verifying, 795
 - OSPF, 213
 - elections*, 214-216
 - LSA distribution*, 213
 - placement*, 216-219
 - PIM-SM, 772
 - Type 2 LSAs, 257-259
- DSP (Domain Specific Part), 320-321**
- DUAL (Diffused Update Algorithm), 70**
- dual stacks, 893**
- duplicate address detection (DAD), 937-938**
- dynamic routing protocols, 68**
 - administrative distance, 73-76
 - distance vector, 69
 - enhanced distance vector, 70

- link-state, 70-71
- listing of, 68
- path vector, 71-72
- routing table metrics, 75-76

E

eBGP

- BGP table, 440
- configuration, 439
- eiBGP multipath, 726-733
 - core deciding path*, 728
 - edge routers*, 727
 - topology*, 728-731

iBGP

- combinations*, 442-444
- sessions, compared*, 438
- multihop sessions, 459-461
- multipath, 723-726
 - attributes*, 725
 - configuration*, 724
- next-hop addresses, 444-445
- oldest path, 720
- over iBGP, 714-717
 - configuration*, 715
 - processing logic*, 717
 - routing table*, 717

prefix attributes

- listing of*, 441
- local*, 441
- remote*, 440

topology, 438

ECMP (equal-cost multipath), 75

BGP, 723

- AS_Path relax feature*, 731-733
- eBGP/iBGP multipath*, 723-726
- eiBGP multipath*, 726-733

IS-IS, 387

- OSPF path selection, 274

EGPs (exterior gateway protocols)

- defined, 68
- path selection algorithms
 - distance vector*, 69
 - enhanced distance vector*, 70
 - link-state*, 70-71
 - path vector*, 71-72

eiBGP multipath, 726-733

- core deciding path, 728
- edge routers, 727
- topology, 728-731

EIGRP (Enhanced Interior Gateway Routing Protocol), 68, 125

- AD, modifying, 528-529
- authentication, 174
 - enabling*, 174
 - key chain configuration*, 174-177
- autonomous system configuration, 126
 - installed routes, displaying*, 140-141
 - interface verification*, 136-139
- IOS, 132-133
- IOS XR, 134
- neighbor adjacencies verification*, 139
- passive interfaces*, 134
- sample topology*, 134-136
- convergence, 158
- feasibility condition, 142
- feasible distance (FD), 142
- feasible successor, 142
- filtering
 - hop counts*, 538
 - offset lists*, 538-541
 - prefix*, 534-537

- inter-router communication, 127
- link failures, 156-157
- metrics, 145
 - attribute propagation*, 146
 - custom K values*, 148-149
 - default K values*, 145
 - with definitions*, 145
 - formula*, 145
 - interface delay settings*, 149-151
 - interface metrics*, 146
 - load balancing*, 151-153
 - lowest link speed/cumulative delay*, 146
 - specific prefix*, 147-148
 - variance multiplier*, 151-152
 - wide*, 153-154
- named configurations, 184
 - address family instance configuration*, 185-186
 - address family interface configuration*, 186-188
 - address family topology configuration*, 188
 - benefits*, 184
- neighbors
 - advertised networks update packet*, 131
 - forming*, 128-131
 - hello packet capture*, 129
 - INIT flag and acknowledgment update packet*, 130
 - INIT flag set update packet*, 129
 - overview*, 126
 - requirements*, 128
 - route prefixes update packet*, 130
 - sample topology*, 129
 - verification*, 139
- packet types, 127
- PDMs, 126
- queries
 - boundaries, establishing*, 157
 - packets, handling*, 157
- redistribution, 563-566
- reference topology, 142
- reported distance (RD), 142
- RIDs, 141
- RTP, 127
- stubs, 160
 - default configuration*, 162-163
 - designs*, 164-166
 - inter-region traffic*, 161-162
 - regional/remote topology*, 161
 - route tables after link failures*, 163
- successors, 142
- summarization
 - automatic*, 172-174
 - hierarchical*, 166
 - interface-specific*, 166-171
 - metrics*, 171-172
- timers
 - hello/hold*, 155-156
 - SIA*, 159-160
- topology table, 143-144
- WANs, 177
 - bandwidth percent*, 177-179
 - next-hop self behavior*, 182-183
 - split horizons*, 179-182
- eigrp router-id router-id command**, 185
- eigrp stub command**, 185
- EIGRP-to-EIGRP redistribution**, 566-568

EIGRPv6, 944

configuration

classic mode, 945

IOS XR, 946-947

named mode, 946

configuration command reference
chart, 954

default routes, 952-953

inter-router communication, 944

route filtering, 953-954

summarization, 950-952

verification, 947-950

base configuration, 948-949

neighbor adjacency, 949-950

routing table, 950

show commands, 947

topology, 948

elections

DISs, 351

DRs/DBRs, 214-216

embedded RP, 1021-1024**encryption command, 975****end systems (ES), 316****enhanced distance vector
algorithms, 70****Enhanced Interior Gateway Routing
Protocol. *See* EIGRP****Equal Cost Multi-Path. *See* ECMP****ES (end systems), 316****Established state (BGP neighbors), 417****Ethernet link failures**

floating static routes, 104

static route recursion problems,
113-114

**EUI-IDs (extended unique identifiers),
920-921****Exchange state (OSPF neighbors), 196****EXEC mode (IOS XR), 14**

exiting nested configurations, 26

explicit IP addresses, 133

explicit subnets, 133

explicit tracking (MLD), 1014

ExStart state (OSPF neighbors), 196

extended ACLs

BGP network selection, 470

defining, 469

IGP network selection, 469-470

extended BGP communities, 610

extended unique identifiers (EUI-
IDs), 920-921

exterior gateway protocols. *See*
EGPs

external routes, 247-248

IS-IS, 386

OSPF path selection, 272

OSPF summarization, 280-283

concept, 280

configuration, 282

loop-prevention routers, 283

routing table before summarization, 281

topology, 281

Type 1, 273

Type 2, 273-274

F**failure detection**

BGP, 459

IS-IS, 366

hello multiplier, 367

hello timers, 366-367

holding timer, 367

- OSPF, 219
 - dead interval timer*, 219
 - fast-packet Hellos*, 220-221
 - Hello timer*, 219
 - timer verification*, 219
- failures
 - area route advertisement, 242
 - commit, 20-21
 - EIGRP
 - route tables after link failures with stubs*, 163
 - link*, 156-157
 - fast reroutes. *See* FRR
 - full-mesh iBGP link, 433
 - OSPF interarea summarization, 277
- fast-packet Hellos, 220-221
- fast reroute. *See* FRR
- FD (feasible distance), 142
- feasibility condition, 142
- feasible successors, 142
- feature navigator tool, 3
- feature sets, identifying, 2-3
- FHR (first-hop router), 763
- FIB (Forwarding Information Base), 87
- fields
 - BGP
 - summary*, 422
 - table*, 429
 - IGMPv2 messages, 757
 - IIH packets, 328
 - IS-IS neighbor states, 346
 - LSP lifetime, 329
 - NSAP
 - DSP*, 321
 - IDP*, 320
 - OSPF
 - Hello packet*, 194
 - LSDB*, 252
 - OSPFv3 Options bit field, 966
 - Type 3 LSAs, 262
 - Type 4 LSAs, 267
 - Type 5 LSAs, 265
 - Type 7 LSAs, 269
 - ULA, 904
- filter-autorp keyword, 864
- filtering routes, 534
 - Auto-RP Cisco-RP-announce messages, 867
 - Auto-RP group, 836-840
 - BGP, 546-548, 647-649
 - BSR multiple RPs, 843-845
 - EIGRP
 - hop counts*, 538
 - offset lists*, 538-541
 - prefix*, 534-537
 - EIGRPv6, 953-954
 - IS-IS, 546
 - OSPF
 - areas*, 543-546
 - local*, 541-543
 - PIM-SM source registration, 867-868
- Finite State Machine (FSM), 415
- first-hop router (FHR), 763
- first usable host addresses, calculating, 43
- flexible route suppression
 - leaking suppressed routes, 634-637
 - selective prefix, 632-634

floating static routes, 103-105

- AD verification, 104-105
- backups, configuring, 112
- defined, 103
- Ethernet link failures, 104
- XR1/R3
 - configuration, 103*
 - routing table, 103-104*

flooding

- OSPF LSAs, 251
 - Type 1, 252*
 - Type 5, 263*
- OSPFv3 scopes, 959-960

Ford-Fulkerson algorithms, 69**forwarding**

- access-list-based (ABF), 521
 - configuring, 523-525*
 - overview, 521*
- bidirectional detection (BFD)
- Cisco Express (CEF)
 - centralized versus distributed forwarding architectures, 86*
 - defined, 86*
 - distributed, 88*
 - hardware, 88-89*
 - IPv6, enabling, 899*
 - software, 87*
 - tables, 107-114*

continuous. *See* continuous forwarding

Information Base (FIB), 87

nonstop. *See* NSF

multicast forwarding information base (MFIB), 764

reverse path. *See* RPF

virtual routing and forwarding. *See* VRF

fragment IDs (LSP IDs), 330

frame relay interfaces, 222-223

FSM (Finite State Machine), 415

full-mesh topology (iBGP), 432-433

Full state (OSPF neighbors), 196

G

gateways, 30

general topology networks, 327

global configuration mode

IOS, 10-11

IOS XR, 14

global routing tables

IP address configuration, 80

output, 80-81

global unicast addresses. *See* GUAs

GLOP Block, 751

group membership LSAs, 959

groups

IPv6 multicast, 910-911

MLD, 1013

GUAs (global unicast addresses), 900

address allocation hierarchy,
900-901

prefixes

bit boundaries, 903

list website, 900

route aggregation, 902-903

subnetting chart, 903

H

hardware CEF, 88-89

hello authentication (IS-IS), 368

hello interval field (OSPF Hello packet), 194

hello-interval seconds command, 186

hello multiplier (IS-IS failure detection), 367

Hello packets

area IDs, 245

fast-packet, 220-221

IIH, 327-328

fields, 328

padding removing, 365

TLVs, 328

types, 327

IS-IS, 324

OSPF, 194-195, 245

OSPFv3, 957

Hello PIM control message, 764

hello timers

EIGRP, 155-156

IS-IS failure detection, 366-367

OSPF failure detection, 219

hexadecimal addresses

IPv6 binary conversion, 896-898

hextets, 894

hierarchy

configurations, 25-26

EIGRP summarization, 166

GUA address allocation, 900-901

IS-IS, 317

OSPF, 192-193

high availability

HO-DSP (High Order DSP), 321

hold timers, 367

BGP, 412

EIGRP, 155-156

hold-time seconds command, 186

hop counts (EIGRP filtering), 538

hosts

addresses

all-0s, 37

available, listing, 47

first usable, calculating, 43

last usable, calculating, 43

identifiers, 30

subnet mask reference chart, 40

hyphen (-) query modifier, 479-480

I

IANA (Internet Assigned Numbers Authority)

AFI listing website, 812

IPv6

address blocks, 911

multicast addresses, 910

multicast addresses, 749

private address ranges, 58

public ASNs, 408

special purpose reserved addresses, 59

iBGP

custom routing, 432

eBGP

combinations, 442-444

sessions, compared, 438

eiBGP multipath, 726-733

core deciding path, 728

edge routers, 727

topology, 728-731

full-mesh requirement, 432-433

multipath, 723-726

attributes, 725

configuration, 724

- over eBGP, 714-717
 - configuration*, 715
 - processing logic*, 717
 - routing table*, 717
- path attributes, 432
- peering via loopback addresses, 433-437
 - BGP table*, 437
 - configuration source*, 435-436
 - IPv4 session summary*, 437
- prefix advertisement behavior, 431
- scalability, 432, 446
 - confederations*, 453-458
 - loop prevention in route reflectors*, 451-452
 - out-of-band route reflectors*, 453
 - route reflectors*, 446-451
- ICMPv6, 919-920**
- identifying**
 - active licenses, 3-4
 - feature sets, 2-3
 - sender network IDs, 36-37
 - software versions, 2-3
 - universal image versions, 3-4
- IDI (Initial Domain Identifier), 320**
- Idle state (BGP neighbors), 415**
- IDP (Inter-Domain Part), 320**
- IDs**
 - extended unique (EUI-IDs), 920-921
 - LSPs, 330
 - OSPF areas, 245-246
 - router (RIDs)
 - BGP*, 413, 720-721
 - EIGRP*, 141
 - EIGRPv6*, 945
 - OSPF*, 196
 - OSPFv2 versus OSPFv3*, 955
 - sender network, 36-37
- IETF (Internet Engineering Task Force), 59**
- if as-path command, 497**
- if as-path in command, 497**
- if as-path is-local command, 497**
- if destination in command, 497**
- if local-preference command, 497**
- if med command, 497**
- if next-hop in command, 497**
- if origin is command, 497**
- if tag command, 497, 498**
- IGMP (Internet Group Management Protocol), 745, 753**
 - snooping, 753-756
 - static joins, 882-886
 - IGMPv2 configuration*, 884-885
 - IGMPv3 configuration*, 884-885
 - mroute table*, 885-886
 - v2, 757-758, 780-782
 - v3, 753, 759
- IGPs (interior gateway protocols)**
 - defined, 68
 - lowest IGP metric (BGP path selection), 718-720
 - path selection algorithms
 - distance vector*, 69
 - enhanced distance vector*, 70
 - link-state*, 70-71
 - path vector*, 71-72
- IIH (IS-IS hello) packets, 327-328**
 - fields, 328
 - IS-IS interface frequency, 393
 - padding, removing, 365
 - TLVs, 328
 - types, 327

- include-connected key word, 1004
- incoming interfaces, 763
- incremental SPF. *See* iSPF
- INIT flag sets (EIGRP update packet), 129
- Initial Domain Identifier (IDI), 320
- Init state (OSPF neighbors), 196
- in keyword (ip multicast boundary access-list command), 864
- inline BGP community conditional matching, 622-625
- inline policy set expansion, 500
- inline prefix filtering route policy, 499
- instance instance-id command, 975
- instances
 - EIGRP address family configuration, 185-186
 - multiple
 - OSPFv2 versus OSPFv3*, 957
 - OSPFv3*, 973-975
- Integrated IS-IS, 323
- interarea routes, 246-247
 - IS-IS, 386
 - LSAs, 959
 - OSPF path selection, 272
 - OSPF summarization, 276-280
 - concept*, 276
 - configuring*, 279
 - example*, 278
 - failure*, 277
 - loop-prevention routes*, 280
 - metrics*, 277
 - routing table after summarization*, 279
 - routing table before summarization*, 278
- prefix LSAs, 959
- interdomain multicast routing
 - MBGP, 812
 - AFIs/SAFIs*, 812
 - BGP topology*, 813-814
 - configuration*, 815
 - multicast BGP session summary verification*, 817
 - unicast BGP session summary verification*, 816
 - update message packet capture*, 813
 - MSPD, 817-818
 - keepalive messages*, 819-821
 - peers*, 822-828
 - SA messages*, 818-819
 - stub networks*, 831-833
 - verification*, 828-831
- Inter-Domain Part (IDP), 320
- interface address mask field (OSPF Hello packet), 194
- interface interface-type interface-number command, 954, 975, 988
- interface priority field (OSPF Hello packet), 194
- interfaces
 - broadcast
 - as P2P, configuring*, 353-355
 - static*, 98-99
 - EIGRP
 - address family configuration*, 186-188
 - authentication, enabling*, 174
 - delay settings*, 149-151
 - path calculation metrics*, 146
 - summarization*, 166-171
 - verifying*, 136-139

- IP addresses, assigning
 - CIDR notation*, 60
 - full subnet masks*, 60
 - secondary addresses*, 60-61
 - SLAAC*, 895-924
- IPv4
 - specific, displaying*, 61
 - summaries, displaying*, 62
- IPv6 status, 899-900
- IS-IS
 - IIH frequency*, 393
 - metric*, 387-394
 - passive*, 361-362
 - specific levels*, 385-386
 - verifying*, 343-346
- loopback, 433-437
- null
 - defined*, 116
 - packet traces demonstrating routing loops*, 118-119
 - preventing routing loops*, 119-120
 - routing loops*, 117-118
- OSPF
 - areas. OSPF, areas*
 - costs*, 235-236
 - frame relay*, 222-223
 - non-broadcast*, 223, 301-303
 - P2P*, 225
 - point-to-multipoint network type*, 226-227
 - verifying*, 208-209
- OSPFv3
 - authentication/encryption*, 972
 - verification*, 965
- overlapping IP addresses, 76-77
- P2P
 - broadcast interfaces as, configuring*, 353-355
 - static routing*, 96-98
- passive
 - EIGRP*, 134
 - OSPF*, 205-206
- PIM
 - downstream*, 763
 - incoming*, 763
 - RPF*, 762
 - upstream*, 763
 - verification*, 791
- status verification, displaying, 82-83
- VRFs, assigning
 - IOS*, 78
 - IOS XR*, 78-79
- interior gateway protocols. *See* IGP
- intermediate systems (IS), 316
- Intermediate System-to-Intermediate System. *See* IS-IS
- Internet
 - Assigned Numbers Authority. *See* IANA
 - BGP community, 611
 - Engineering Task Force (IETF), 59
 - Group Management Protocol. *See* IGMP
- Internetwork Control Block, 750
- Internetworking Operating System. *See* IOS
- inter-region network connectivity, 161-162
- inter-router communication, 957
 - BGP, 410-411
 - keepalive messages*, 413
 - message types*, 411
 - notification messages*, 414

- open messages, 412-413*
 - sessions, 415*
 - update messages, 413-414*
- EIGRP, 127
- EIGRPv6, 944
- IPv6, 978
- IS-IS, 323-325
 - IIH packets, 327-328*
 - IS PDU addressing, 326-327*
 - IS protocol headers, 325*
 - LSPs, 329-332*
 - neighbors. See IS-IS, neighbors*
 - TLVs, 326*
- MBGP for IPv6, 989-990
- OSPF, 193-194
 - Hello packets, 194-195*
 - multicast addresses, 193*
 - packet types, 194*
- intra-area prefix, 959**
- intra-area routes, 246-247**
 - IS-IS, 386
 - OSPF path selection, 271-272
- IOS (Internetworking Operating System), 1**
 - administratively scoped multicast boundaries, 864
 - AS_Path ACL, 484
 - BGP
 - best path weight, 678*
 - configuration scalability, 664-666*
 - configuring, 419-420*
 - identifier, 413*
 - CLI modes, 9-11
 - global configuration, 10-11*
 - privileged, 10*
 - user, 9-10*
 - configurations
 - archiving, 11-13*
 - prompts, 26*
 - replacing, 13*
 - DNS SSM mapping, enabling, 858
 - eBGP configuration, 438
 - EIGRP
 - autonomous system configuration, 132-133*
 - RIDs, 141*
 - EIGRPv6
 - configuration, 945-946*
 - route filtering, 953-954*
 - IS-IS
 - configuring, 340*
 - route leaking, 377*
 - IS-IS for IPv6
 - base configuration, 979-980*
 - topology configuration, 981*
 - kernel, 2
 - memory management, 2
 - MBGP for IPv6 configuration, 991-992
 - MSDP
 - MSDP Compliance, 823*
 - peer configuration, 822*
 - OSPF
 - authentication, 236*
 - interface specific configuration, 204-205*
 - network statement, 202-204*
 - OSPFv3 configuration, 960-961
 - PBR, 521
 - configuring, 522*
 - local, 525-526*
 - overview, 521*
 - PIM-SM accept RP, 868-869

- prefix lists, 473
- private BGP communities, setting, 625-627
- process switching, 84
- run to completion scheduler, 2
- show ip mroute command flags, 794
- SLAAC configuration, 922
- software packaging, 2-4
- SSM, enabling, 851
- stateful DHCPv6
 - relay agent configuration*, 927-928
 - relay agent verification*, 928
 - server configuration*, 930
- stateless DHCPv6 configuration, 924-925
- static route, configuring, 96
- summary, 7
- VRFs, creating, 78
- IOSd, 4**
- IOS XE, 4**
 - kernel, 4
 - memory management, 4
 - scheduling, 4
 - summary, 7
- IOS XR, 5**
 - ABF, 521
 - configuring*, 523-525
 - overview*, 521
 - AS paths, 484
 - is-local*, 485
 - length*, 485-486
 - neighbor-is*, 487
 - originates-from*, 487
 - passes-through*, 486-487
 - sets*, 488
 - unique-length*, 486
 - BGP
 - best path weight*, 678
 - community conditional matching*, 621-625
 - configuring*, 420-421
 - identifier*, 413
 - peer configuration*, 667-668
 - commits
 - confirming*, 22-23
 - failures*, 20-21
 - labels*, 18-19
 - multiple commit options*, 23-24
 - replaces*, 19-20
 - configurations
 - change rollbacks*, 21-22
 - displaying changes in SysDB*, 17-18
 - files, loading*, 24
 - hierarchical*, 25-26
 - nested, exiting*, 26
 - processing changes*, 14-17
 - prompts*, 26
 - eBGP configuration, 438
 - EIGRP
 - autonomous system configuration*, 134
 - RIDs*, 141
 - EIGRPv6
 - configuration*, 946-947
 - route filtering*, 954
 - IPv6 BSR PIM-SM configuration, 1020
 - IS-IS
 - configuring*, 340-341
 - route leaking*, 377
 - IS-IS for IPv6
 - base configuration*, 980

- topology mode configuration, 981-984*
- kernel, 5
- MBGP for IPv6 configuration, 992-993
- memory management, 6
- modes
 - admin, 14*
 - EXEC, 14*
 - global configuration, 14*
- MSDP peer configuration, 822-823
- multicast boundary scope configuration, 1032
- OSPF
 - authentication, 237-239*
 - configuring, 205*
- OSPFv3 configuration, 961
- prefix
 - lists, 473*
 - sets, 474-475*
- process switching, 84
- scheduling, 5
- software packaging, 6-7
- SSM, enabling, 851
- stateful DHCPv6
 - relay proxy agent configuration, 928-929*
 - relay proxy agent verification, 929*
 - server configuration, 931-932*
- stateless DHCPv6 configuration, 925-926
- static routes, configuring, 96
- static VRF routes, 121-122
- summary, 7
- traces, 8
- VRFs, creating, 78-79

IP addresses

- assigning to interfaces
 - CIDR notation, 60*
 - full subnet masks, 60*
 - secondary addresses, 60-61*
 - SLAAC, 895-924*
- binary notation
 - base 2 calculations, 31-32*
 - bit values, 31-32*
 - to decimal conversion, 33*
 - decimal to binary conversions, 32*
 - defined, 31*
- classes
 - A, 34*
 - B, 35*
 - C, 35*
 - D, 35*
 - history, 34*
 - listing of, 34*
- defined, 30
- dot-decimal notation, 30
- explicit, 133
- gateway, 30
- global routing table configuration, 80
- history, 34
- interfaces, 60
- interesting octets, 43
- NAT, 58
- network/host identifiers, 30
- next-hop self, 182-183
- overlapping, 76-77
- private
 - defined, 58*
 - IPv4 address ranges, 58*
 - NAT, 58*

- secondary, assigning to interfaces, 60-61
- special-purpose reserved, 59
- status verification, displaying, 82-83
- subnet fields, 39
- subnet masks. *See* subnet masks
- usable, calculating, 37-38
- ip address ip-address subnet-mask command, 60**
- ip ospf process-id area area-id command, 204**
- IP packets, delivering, 30**
- IPv4**
 - BGP neighbor output, 421-424
 - interfaces
 - specific, displaying, 61*
 - summaries, displaying, 62*
 - IPv6
 - address resolution, compared, 934*
 - over IPv4 BGP sessions, 998-1001*
 - multicast routing. *See* multicast
 - multicast versus IPv6 multicast, 1008
 - session summary of loopback interfaces, 437
- ipv4 address ip-address subnet mask command, 60**
- ipv4 address ipv4-address prefix-length command, 60**
- IPv6**
 - address allocation hierarchy, 900-901
 - address structure
 - abbreviations, 895-896*
 - bit length values, 894*
 - case, 894*
 - components, 894*
 - hexadecimal to binary conversion, 896-898*
 - hextets, 894*
 - network boundaries, 894*
 - unicast, 894*
 - aggregation, 902-903
 - anycast, 906-908
 - duplicate address detection, disabling, 908*
 - topology, 906*
 - updated topology, 907*
 - CEF, enabling, 899
 - EIGRPv6, 944
 - classic mode configuration, 945*
 - configuration command reference chart, 954*
 - default routes, 952-953*
 - inter-router communication, 944*
 - IOS XR configuration, 946-947*
 - named mode configuration, 946*
 - route filtering, 953-954*
 - summarization, 950-952*
 - topology, 948*
 - verification, 947-950*
 - IS-IS, 977
 - configuration, 979*
 - configuration commands reference chart, 988*
 - inter-router communication, 978*
 - IOS base configuration, 979-980*
 - IOS topology mode configuration, 981*
 - IOS XR base configuration, 980*
 - IOS XR topology mode configuration, 981-984*
 - TLVs, 944-978*

- topology modes*, 978-979
- verification*, 985-987
- MBGP, 989
 - BGP configuration commands reference chart*, 1001
 - inter-router communication*, 989-990
 - IOS configuration*, 991-992
 - IOS XR configuration*, 992-993
 - IPv6 over IPv4 BGP sessions*, 998-1001
 - verification*. See *verification, MBGP for IPv6*
- multicast, 908-911
 - active traffic flows*, 1030
 - address format*, 909
 - common addresses*, 910
 - configuration*, 1025
 - enabling*, 1010
 - enhancements*, 1008
 - groups*, 910-911
 - IPv4 multicast comparison*, 1008
 - MAC address mapping*, 1009
 - MLD*, 1010-1015
 - PIM*, 1015
 - PIM-SM*. See *PIM-SM*
 - PIM-SSM*, 1033-1034
 - RPF*, 1030-1032
 - scope boundary*, 1032-1033
 - scope types*, 909
 - stream statistics, enabling*, 1029
 - topology*, 908
 - verification commands*, 1024
- NDP, 912
 - address resolution*, 934-936
 - default router preferences, modifying*, 916
 - disabling RA messages*, 918
 - duplicate address detection*, 937-938
 - ICMPv6 message types*, 913
 - NUD (neighbor unreachability detection)*, 936
 - parameters, displaying*, 916-917
 - prefix valid lifetime/preferred lifetime values*, 915-916
 - RA lifetime values*, 917-918
 - RA messages*, 913-914
 - reachability*, 936-937
 - redirects*, 919-920
 - stateful DHCPv6*. See *DHCPv6, stateful*
 - stateless address autoconfiguration*, 895-924
 - stateless DHCPv6*, 924-926
- OSPFv3, 955
 - authentication*, 970-973
 - configuration commands reference chart*, 975
 - flooding scopes*, 959-960
 - inter-router communication*, 957
 - IOS configuration*, 960-961
 - IOS XR configuration*, 961
 - LSAs*, 958
 - multiple instances*, 973-975
 - OSPFv2, compared*, 955-957
 - verification*, 962-970
- overview, 893
- redistribution, 1002-1005
 - IOS*, 1002
 - route tables*, 1003-1005
- special purpose addresses, 911

- static routing, 941-942
 - configuration commands reference chart, 943*
 - configuring, 942-943*
 - show commands reference chart, 943*
- subnetting chart, 903
- unicast, 898-900
 - configuring, 899*
 - global, 900-903*
 - interface status, 899-900*
 - link-local, 905-906*
 - scopes, 898-899*
 - unique local, 904-905*
- ipv6 authentication key-chain eigrp as-number key-chain-name command, 954
- ipv6 authentication mode eigrp as-number md5 command, 954
- ipv6 router eigrp as-number command, 954
- ipv6 router isis command, 988
- ipv6 summary-address eigrp as-number ipv6-prefix/prefix-length command, 954
- ip vrf forwarding vrf-name command, 77
- ip vrf vrf-name command, 77
- IS (intermediate systems), 316
- IS-IS (Intermediate System-to-Intermediate System), 68
 - AD, modifying, 531-532
 - areas, 318-319
 - authentication, 367
 - configuration, 369-370*
 - hello, 368*
 - LSP, 368*
 - types, 367*
 - backbone continuity, 380-382
 - broadcast interfaces as P2P, configuring, 353-355
 - configuring
 - interfaces, verifying, 343-346*
 - IOS, 340*
 - IOS XR, 340-341*
 - neighbor adjacencies, verifying, 346-347*
 - route verification, 347-348*
 - topology example, 341-343*
 - default routes, 400-401
 - DISs, 348-353
 - configuration, 352*
 - CSNPs, 350*
 - elections, 351*
 - logical drawing, 348*
 - LSP advertisement, 349*
 - placement, 352-353*
 - pseudonode LSPs, 349-350*
 - verification, 353*
 - failure detection, 366
 - hello multiplier, 367*
 - hello timers, 366-367*
 - holding timer, 367*
 - filtering, 546
 - hierarchy, 317
 - IIH padding, removing, 365
 - Integrated, 323
 - interarea topology, 374
 - interfaces
 - IIH frequency, 393*
 - metric, 387-394*
 - passive, 362-364*
 - specific levels, 385-386*
 - verifying, 343-346*

- inter-router communication, 323-325
 - IIH packets*, 327-328
 - IS PDU addressing*, 326-327
 - IS protocol headers*, 325
 - LSPs*, 329-332
 - TLVs*, 326
- IPv6, 977
 - configuration*, 979
 - configuration commands reference chart*, 988
 - inter-router communication*, 978
 - IOS base configuration*, 979-980
 - IOS topology mode configuration*, 981
 - IOS XR base configuration*, 980
 - IOS XR topology mode configuration*, 981-984
 - TLVs*, 944-978
 - topology modes*, 978-979
 - verification*, 985-987
- loop prevention, 382-384
- LSPDB, 317, 355
 - building topology*, 359-360
 - displaying*, 356-357
 - displaying topology*, 360-361
 - non-pseudonode LSPs*, 357-358
 - pseudonode LSPs*, 358-359
 - SPF calculations*, 361-362
- LSPs, 329-332
 - attribute fields*, 331
 - IDs*, 330
 - lifetime*, 329
 - sequence number*, 331
 - TLVs*, 332
 - types*, 329
- metrics
 - configuration*, 391
 - mismatch*, 389
 - narrow and wide*, 388
 - topology*, 391
 - verification*, 392
- neighbors, 333-339
 - adjacency capability chart*, 386
 - broadcast*, 333-338
 - P2P*, 338-339
 - verifying*, 346-347
- NET addresses, 322-323
 - common private*, 322
 - expanded structure*, 322
 - guidelines*, 323
 - minimal format*, 322
 - multiple*, 323
- NSAP, 320
 - DSP*, 321
 - IDP*, 320
 - structure*, 320
- OSPF commonalities, 316
- overload bits, 394-396
 - configuration*, 395
 - routing table*, 395-396
 - topology*, 394
- packets, 324
- path selection, 386
 - ECMP*, 387
 - interface metrics*, 387-394
 - processing order*, 387
- prefix
 - suppression*, 401-405
- redistribution, 576-578
- route-leaking, 377-380
 - configuring*, 377

- interarea TLVs*, 379-380
- routing table*, 378-379
- router specific levels, 384-385
- segmenting domains into multiple levels, 373-376
- SPT, 317
- suboptimal routing, 377
- summarization, 396
 - configuration*, 399
 - discard routes*, 399
 - metric*, 396-397
 - ranges*, 397
 - routing table*, 398
 - topology*, 398
- isis ipv6 metric command, 988
- IS-IS to IS-IS redistribution, 578-580
- is-local option (AS_Path), 485
- iSPF (incremental SPF)
- ISP to CE router connectivity, 100

J - K

- Join/prune PIM control message, 764
- keepalive messages, 413, 819-821
- kernels
 - IOS, 2
 - IOS XE, 4
 - IOS XR, 5
- key chains, 174-177
- keywords
 - and, 504
 - as-set, 639-641
 - continue, 493-494
 - filter-autorp, 864
 - in, 864
 - include-connected, 1004

- ip multicast boundary access-list command, 864
- not, 504
- or, 505
- out, 864
- rp-list, 867
- K values (EIGRP metrics)**
 - custom, 148-149
 - default, 145

L

- label parameter, 23
- LAG (link aggregation) bundles, 722
- last-hop router (LHR), 763
- last usable host addresses, calculating, 43
- Layer 2 addressing, 83-84
- Layer 2 multicast addresses, 752-753
- leaking
 - configuring, 377
 - interarea TLVs, 379-380
 - routing table, 378-379
 - suppressed routes, 634-637
- length option (AS_Path), 485-486
- LFA FRR (loop-free alternate FRR)
- LHR (last-hop router), 763
- licenses, 3-4
- lifetime field (LSPs), 329
- links
 - aggregation (LAG) bundles, 722
 - failures (EIGRP), 156-157, 163
 - link-state
 - acknowledgment packets*, 194, 957
 - advertisements*. See LSAs

- algorithms*, 70-71
- database*. See LSDB
- packet database*. See LSPDB
- packets*. See LSPs
- request (LSR) packets*, 194, 200
- update (LSU) packets*, 194, 200
- local addresses. See LLAs
- OSPF virtual, 298-301
 - configuring*, 300
 - verification*, 300-301
- Linux differentials**, 13
- LLAs (link-local addresses)**, 905-906
 - assignment, 906
 - format, 905
 - manually assigning, 906
- load balancing**
 - EIGRP, 151-153
 - RP groups, 836-840
 - configuration*, 837
 - group-to-RP mapping*, 837, 839-840
 - redundancy Auto-RP group-to-RP mappings*, 838
 - redundancy configuration*, 838
- loading configuration files**, 24
- Loading state (OSPF neighbors)**, 196
- local**
 - ASNs, 660-664
 - Network Control Block, 750
 - OSPF filtering, 541-543
 - PBR, 525-526
 - preferences (BGP), 679-682
 - configuration*, 680
 - processing logic*, 682
 - routing table*, 682
 - topology*, 679
 - routes, 93
- LocPrf field (BGP tables)**, 429
- looking glass servers**, 483
- loopback addresses (iBGP peering)**, 433-437
 - BGP table, 437
 - configuration source, 435-436
 - IPv4 session summary, 437
 - OSPF, 229-230
- loops**
 - packet traces demonstrating, 118-119
 - prevention, 119-120
 - BGP*, 409-410
 - IS-IS*, 382-384
 - route reflectors*, 451-452
 - redistribution
 - AD*, 601-603
 - overview*, 590-592
 - prefix filtering*, 593-595
 - seed metrics, increasing*, 598-600
 - summarization*, 603-605
 - tagging*, 595-598
 - topology, 117-118
- LSAs (link-state advertisements)**, 70, 191
 - DR/BDR distribution, 213
 - exponential LSA sessions per routers on same segment, 212
 - OSPF, 249-251
 - age*, 251
 - flooding*, 251
 - reduction through area segmentation*, 274
 - summary*, 270
 - Type 1*, 252-257
 - Type 2*, 257-259
 - Type 3*, 259-262

- Type 4*, 265-268
- Type 5*, 263-265
- Type 7*, 268-269
- OSPFv2 versus OSPFv3, 955
- OSPFv3, 958
 - flooding scopes*, 959-960
 - types*, 959
- LSDB (link-state database)**, 70, 192
 - OSPF Type 1 LSA fields, 252
 - OSPFv3, 965-969
- LSPDB (link-state packet database) (IS-IS)**, 317, 355
 - displaying, 356-357
 - non-pseudonode LSPs, 357-358
 - pseudonode LSPs, 358-359
 - SPF calculations, 361-362
 - topology
 - building*, 359-360
 - displaying*, 360-361
 - without transit networks, 404-405
- LSPs (link-state packets)**, 70
 - DISs
 - advertisement*, 349
 - pseudonode*, 349-350
 - IS-IS, 324, 329-332
 - attribute fields*, 331
 - authentication*, 368
 - building topology*, 359-360
 - displaying topology*, 360-361
 - IDs*, 330
 - lifetime*, 329
 - LSPDB, displaying*, 356-357
 - non-pseudonode*, 357-358
 - pseudonode*, 358-359
 - sequence number*, 331
 - TLVs*, 332
 - types*, 329

- LSRs (link-state request) packets**, 194, 200
- LSU (link-state update) packets**, 194, 200

M

MAC addresses

- defined, 752
- mapping, 1009
- multicast, 752-753

magic number method, 42-45

- address range calculation examples
 - 10.55.200.33/12*, 45
 - 172.16.2.3/23*, 45
 - 192.168.100.5/26*, 45

- broadcast addresses, calculating, 43
- first usable addresses, 43
- interesting IP address octets, 43
- interesting subnet octets, 42-43
- last usable host address, 43
- process overview, 44

management plane, 89

mapping (SSM), 857

- DNS, 857-860
 - configuration*, 859
 - topology*, 858
 - verification*, 859
- static, 860-862
 - configuring*, 860
 - verification*, 861

MAAs (RP mapping agents), 773

match as-path acl-number command, 490

matching length parameters (prefix matching), 471

match ip address command, 490

- match ip address prefix-list prefix-list-name command, 490
- match local-preference command, 490
- match metric command, 490
- match statements, 497
- match tag tag-value command, 490
- maximum-paths command, 188
- MBGP (Multiprotocol BGP), 812**
 - AFIs/SAFIs, 812
 - BGP
 - configuration commands*
 - reference chart, 1001*
 - topology, 813-814*
 - configuration, 815
 - IOS, 991-992
 - IOS XR, 992-993
 - inter-router communication, 989-990
 - IPv6 over IPv4 BGP sessions, 998-1001
 - multicast BGP session summary verification, 817
 - multicast traffic engineering, 875-882
 - as chosen RPF source of information, 880*
 - configuring, 876*
 - incongruent unicast/multicast data paths, 878-880*
 - MBGP table, 877*
 - multicast BGP as chosen RPF source of information, 880*
 - OSPF preferred over BGP for RPF calculation, 878*
 - RPF information, 877*
 - traceroute unicast traffic, 881*
 - unicast BGP as RPF source after multicast link failures, 882*
 - unicast BGP session summary verification, 816
 - update message packet capture, 813
 - verification
 - commands, 993*
 - IPv6 route table, 997*
 - IPv6 unicast BGP table, 996*
 - neighbor status, 996*
 - next-hop address selection, 997*
 - router configuration, 994-995*
- MDT, 851**
- MED (Multi-Exit Discriminator), 693**
 - AIGP metric conversion, 693-694
 - BGP best path selection, 704-714
 - always compare MED feature, 711-713*
 - configuration, 706*
 - deterministic, 713-714*
 - missing MED behaviors, 710-711*
 - outbound traffic, influencing, 705*
 - processing logic, 708*
 - routing table after modification, 707*
- memory, managing**
 - IOS, 2
 - IOS XE, 4
 - IOS XR, 6
- merging configurations, 16**
- messages**
 - Auto-RP Cisco-RP-announce filtering, 867
 - BGP, 411
 - keepalive, 413*
 - notification, 414*
 - open, 412-413*
 - update, 413-414*
 - IGMPv2 messages, 757

- keepalive, 819-821
- MLD, 1010
- PIM control, 764
- PIM-SM IPv6 multicast, 1015
- RA, 913-914
- SA, 818-819, 828
- metric command, 988**
- Metric field (BGP tables), 429**
- metrics**
 - AIGP**
 - configuration, 689*
 - guidelines, 688*
 - MED conversion, 693-694*
 - modifications, 688*
 - neighbor sessions metrics verification, 690*
 - processing logic, 692-719*
 - EIGRP**
 - attribute propagation, 146*
 - custom K values, 148-149*
 - default K values, 145*
 - with definitions, 145*
 - formula, 145*
 - interface delay settings, 149-151*
 - interface metrics, 146*
 - load balancing, 151-153*
 - lowest link speed/cumulative delay, 146*
 - specific prefix, 147-148*
 - summarization, 171-172*
 - variance multiplier, 151-152*
 - wide, 153-154*
- external routes
 - Type 1, 273*
 - Type 2, 273-274*
- interarea summarization, 277
- IS-IS**
 - configuration, 391*
 - interface, 387-394*
 - mismatch, 389*
 - narrow and wide, 388*
 - SPF, 361-362*
 - summarization, 396-397*
 - topology, 391*
 - verification, 392*
- lowest IGP (BGP path selection), 718-720
- OSPF interface costs, 235-236
- redistribution, 558
- routing tables, 75-76
- seed redistribution, 598-600, 606
- Type 3 LSAs, 261
- Type 4 LSAs, 265
- metric-style command, 988**
- metric weights TOS K1 K2 K3 K4 K5 command, 186**
- MFIB (multicast forwarding information base), 764**
- MLD (Multicast Listener Discovery), 1010-1015**
 - destination addresses, 1012
 - disabling, 1014
 - exclude filter mode, 1011
 - explicit tracking, 1014
 - groups, 1013
 - include filter mode, 1011
 - message types, 1010
 - querier election, 1012
 - query messages, 1012
 - static joins, 1015
 - versions, 1010

modes**IOS***global configuration, 10-11**privileged, 10**user, 9-10***IOS XR***admin, 14**EXEC, 14**global configuration, 14***MPLS (Multiprotocol Label Switching), 76****MRIB (multicast routing information base), 763****mroutes, 872-875****MSDP (Multicast Source Discovery Protocol), 817-818**

keepalive messages, 819-821

peers, 822-828

*IOS, 822**IOS XR, 822-823**MSDP and MBGP sessions
between PIM domains,
823-827**MSDP Compliance, 823**SA message packet captures, 828**SA RPF checks, 823**status, 829*

SA messages, 818-819

stub networks, 831-833

verification, 828-831

*Mroute table, 829**peer status, 829**RPF check for source subnet, 831**RPF check for source subnet
without MBGP, 831**SA cache, 830***mtrace, multicast troubleshooting,
887-889**

with active source, 888

no active source, 888

PIM disabled, 889

reference topology, 887

multi-area adjacencies, 304-308

configuration, 305-306

inefficient topologies, 304

nonbroadcast interface verification,
304

OSPF neighborship, 306-307

verification, 306, 307-308

multicast, 745

addressing

*Administratively Scoped Block,
751**GLOP Block, 751**IANA assigned, 749**Internetwork Control Block, 750**Layer 2, 752-753**Local Network Control Block,
750**OSPF, 193**reserved addresses, 750**Source-Specific Multicast
Block, 751*

any source (ASM), 850

architecture, 745

configuration, 780-782, 786-787

distribution trees, 759

*shared trees, 761**source trees, 759-760*forwarding information base
(MFIB), 764**IGMP, 753***snooping, 753-756**v2, 757-758*

- v3*, 759
- versions*, 753
- IPv6, 908-911
 - active traffic flows*, 1030
 - address format*, 909
 - common addresses*, 910
 - configuration*, 1025
 - enabling*, 1010
 - enhancements*, 1008
 - groups*, 910-911
 - IPv4 multicast comparison*, 1008
 - MAC address mapping*, 1009
 - MLD*, 1010-1015
 - PIM*, 1015
 - PIM-SM*. See *PIM-SM*
 - PIM-SSM*, 1033-1034
 - RPF*, 1030-1032
 - scope boundary*, 1032-1033
 - scope types*, 909
 - show commands*, 1024
 - stream statistics, enabling*, 1029
 - topology*, 908
- MBGP, 812
 - AFIs/SAFIs*, 812
 - BGP topology*, 813-814
 - configuration*, 815
 - multicast BGP session summary verification*, 817
 - unicast BPG session summary verification*, 816
 - update message packet capture*, 813
- MSDP, 817-818
 - keepalive messages*, 819-821
 - peers*, 822-828
 - SA messages*, 818-819
 - stub networks*, 831-833
 - verification*, 828-831
- PIM, 762
 - bidirectional*, 802-808
 - control messages*, 764
 - Dense Mode (PIM-DM)*, 765-767
 - downstream*, 763
 - first-hop router (FHR)*, 763
 - forwarder*, 778-779
 - incoming interface*, 763
 - last-hop router (LHR)*, 763
 - multicast forwarding information base (MFIB)*, 764
 - multicast routing information base (MRIB)*, 763
 - multicast state*, 764
 - operating modes*, 764
 - outgoing interface (OIF)*, 763
 - outgoing interface list (OIL)*, 763
 - reference topology*, 762
 - RPF*. See *RPF*
 - Sparse Mode*, 768-772
 - upstream*, 763
- redundant RPs, 833-834
 - anycast*, 847-849
 - Auto-RP load balancing*, 836-840
 - Auto-RP with multiple RPs*, 835-840
 - BSR*, 840-846
 - static RP*, 846
- RPF, 776-777
- RPs, 772
 - Auto-RP*, 773-774
 - configuring*. See *RPs, configuring*

- PIM bootstrap router*, 775-776
 - static*, 773
- security, 862
 - Auto-RP Cisco-RP-announce message filtering*, 867
 - Auto-RP TTL scoping*, 862
 - boundaries*, 863-866
 - PIM neighbor control*, 869-870
 - PIM register rate limit*, 870
 - PIM-SM*, 867-869
- SSM, 850
 - configuration*, 852-853
 - group-to-PIM mode mapping*, 853
 - IOS, enabling*, 851
 - IOS XR, enabling*, 851
 - mapping*, 857-862
 - MDT creation*, 851
 - mroute table*, 855-857
 - RPF neighbors, identifying*, 854
 - topology*, 854
- stream statistics, 801
- traffic engineering
 - MBGP*, 875-882
 - RPF rules*, 871-872
 - static IGMP joins*, 882-886
 - static mroutes*, 872-875
 - unicast incongruent*, 871
- troubleshooting, -889
 - mtrace*, 887-889
 - show commands*, 886
- verification
 - active traffic flows*, 801
 - configuration*, 788-789
 - DR election*, 792, 795
 - IGMP*, 790-791
 - IOS show ip mroute flags*, 794
 - PIM group mapping*, 792
 - PIM interfaces*, 791
 - PIM topology*, 793
 - show mrib route/show mfib route flags*, 798
 - show pim topology flags*, 799
 - source registration between RP and FHR tree creation*, 796-798
 - SPT between RP and FHR tree creation*, 796-798
 - SPT path*, 800
 - SPT switchover*, 799-800
 - stream statistics*, 801
 - video feed, 748
- Multicast Listener Discovery. See MLD**
- multicast routing information base (MRIB)**, 763
- Multicast Source Discovery Protocol. See MSDP**
- Multi-Exit Discriminator. See MED**
- multihop**
 - static routes
 - multiple recursive lookups*, 109-111
 - overview*, 108
 - single recursive lookups*, 108-109
- multiple instances (OSPFv3)**, 973-975
- multiple match route maps**, 491-492
- multiple recursive static routes**, 109-111
 - beneficial, 111
 - configuring, 110
 - recursive static routes routing table, 110-111
 - routing table, 110

multiprotocolBGP. *See* MBGP

Label Switching (MPLS), 76

VRFs

IOS, 78*IOS XR*, 78-79*defined*, 77*IPv4-only migrations*, 77**multitopology command**, 988**N****named configurations (EIGRP)**, 184

address family configurations

instances, 185-186*interfaces*, 186-188*topology*, 188

benefits, 184

named mode (EIGRPv6), 946**narrow IS-IS metrics**, 388**NAT (Network Address Translation)**,
58**NBMA (non-broadcast multi-access)**,
222**NDP (Neighbor Discovery Protocol)**,
912

address resolution, 934-936

default router preferences, modify-
ing, 916

duplicate address detection, 937-938

ICMPv6 message types, 913

NUD, 936

parameters, displaying, 916-917

prefix valid lifetime/preferred life-
time values, 915-916

RA messages

advertisements, 913-914*disabling*, 918*lifetime values*, 917-918*packet capture*, 915*responses*, 914

reachability, 936-937

redirects, 919-920

SLAAC, 895-924

DNS RA options, 923-924*extended unique identifiers*,
920-921*router configuration*, 921-923

stateful DHCPv6, 926-927

IOS relay agent configuration,
927-928*IOS relay agent verification*, 928*IOS server configuration*, 930*IOS XR relay proxy agent con-
figuration*, 928-929*IOS XR relay proxy agent veri-
fication*, 929*IOS XR server configuration*,
931-932*server verification*, 932-934*topology*, 926

stateless DHCPv6, 924-926

IOS configuration, 924-925*IOS XR configuration*, 925-926
verification, 926**negation boolean operator**, 504**neighbor activate command**, 1001**neighbor command**, 1001**Neighbor Discovery Protocol**. *See*
NDP**neighbor ip-address local-as-alternate-
as-number command**,**neighbor-is option (AS_Path)**, 487**neighbor remote-as as-number
command**, 1001

neighbors

adjacencies. *See* adjacencies

BGP

AIGP metrics support, verifying, 690

default advertisements, 644-645

lowest address (path selection), 722

BGP states, 415

Active, 416

Connect, 415-416

Established, 417

Idle, 415

OpenConfirm, 417

OpenSent, 416-417

EIGRP

adjacencies. See adjacencies, EIGRP

convergence, 158

forming, 128-131

hello/hold timers, 155-156

link failures, 156-157

overview, 126

packets, 129-131

query boundaries, establishing, 157

query packets, handling, 157

requirements, 128

sample topology, 129

SIA timers, 159-160

verification, 139

EIGRPv6 adjacency verification, 949

IS-IS, 333-339

adjacency compatibility chart, 386

broadcast, 333-338

P2P, 338-339

verifying, 346-347

OSPF, 196

adjacencies. See adjacencies, OSPF

multi-area interface, 306-307

non-broadcast, 223

P2P interfaces verification, 225

states, 196, 210

Type 1 LSAs state, 256

PIM, 869-870

RPF, 763, 854

unreachability detection (NUD), 936

nested configurations, exiting, 26

nesting route policies, 510-511

NET addresses, 322-323

common private, 322

expanded structure, 322

guidelines, 323

minimal format, 322

multiple, 323

net network-entity-title command, 988

Network Address Translation (NAT), 58

Network field (BGP tables), 429

networkipv6-prefix/prefix-length command, 1001

Network Layer Reachability Information (NLRI), 458

network network command, 185

networks

backdoor, 649-652

broadcast, 326, 333-338

connected

defined, 91

local routes, 93

redistribution, 561

routing table sample, 92-93

sample topology, 92

- secondary*, 94
- specific network routing detail output sample*, 93
- convergence. *See* convergence
- discontiguous, 301-303
- general topology, 327
- identifiers (IP addresses), 30
- IDs, 36-37
- LSAs, 959
- operating system (NOS), 1
- OSPF, 221
 - areas*, 246-248
 - broadcast*, 221
 - listing of*, 231
 - loopback*, 229-230
 - neighbor adjacency compatibility*, 231-234
 - non-broadcast*, 222-223
 - P2P*, 224-225
 - point-to-multipoint*, 225-229
- P2P, 338-339
- prefix notation. *See* CIDR
- Service Access Point. *See* NSAP
- stub, 831-833
- network statement**
 - IOS EIGRP autonomous system configuration, 132-133
 - all interfaces, enabling*, 133
 - explicit IP addresses*, 133
 - explicit subnets*, 133
 - large subnet ranges*, 133
 - IOS OSPF configuration, 202-204
- next-hop addresses**
 - BGP, 429
 - eBGP, 444-445
 - EIGRP WANs, 182-183
 - MBGP for IPv6, 997
 - next-hop-self command, 186
 - nibbles, 894
 - NLRI (Network Layer Reachability Information), 458
 - No_Advertise BGP community, 614-617
 - node protection
 - No_Export BGP community, 611-614
 - No_Export_SubConfed BGP community, 617-620
 - non-broadcast multi-access (NBMA), 222
 - non-broadcast networks (OSPF), 222-223, 301-303
 - non-pseudonode LSPs, 357-358
 - nonrecursive multihop static routes, 107
 - nonstop routing. *See* NSR
 - NOS (network operating system), 1
 - notification messages, 414
 - not keyword (negation operator), 504
 - not-so-stubby areas. *See* NSSAs
 - NSAP (Network Service Access Point), 320
 - DSP, 321
 - IDP, 320
 - structure, 320
 - nssa command, 975
 - nssa [default-information-originate] no-summary command, 975
 - NSSAs (not-so-stubby areas), 268, 292-295, 959
 - concept, 292
 - OSPF
 - autonomous system*, 293
 - configuration*, 293-294

- routing table
 - after*, 294
 - before*, 293
 - totally, 295-298
 - concept*, 295
 - configuration*, 297
 - routing table after*, 298
 - routing table before*, 296
 - topology*, 296
 - Type 7 LSAs, 268-269
 - NUD (neighbor unreachability detection)**, 936
 - null interfaces**
 - defined, 116
 - routing loops, 117-118
 - packet traces demonstrating*, 118-119
 - preventing*, 119-120
- ## O
-
- octets**
 - dot-decimal octet subtraction, 63
 - /24 networks*, 64
 - single IP hosts*, 63
 - summary routes*, 64
 - interesting IP addresses, 43
 - subnet, 42-43
 - offset lists**, 538-541
 - OIF (outgoing interface)**, 763
 - OIL (outgoing interface list)**, 763
 - OpenConfirm state (BGP neighbors)**, 417
 - open messages**, 412-413
 - BGP identifier, 413
 - hold time, 412
 - OpenSent state (BGP neighbors)**, 416-417
 - Open Shortest Path First**. *See* OSPF
 - Open Systems Interconnection**. *See* OSI
 - operating systems summary**, 7
 - operators**
 - bitwise AND, 41-42
 - boolean, 504
 - conjunction*, 504
 - disjunction*, 505
 - negation*, 504
 - order of processing*, 505-506
 - Options bit field (OSPFv3)**, 966
 - ORF (outbound route filtering)**, 647-649
 - originates-from option (AS_Path)**, 487
 - origin attribute (BGP)**, 700-703
 - configuration, 701
 - processing logic, 703
 - or keyword (disjunction operator)**, 505
 - OSI (Open Systems Interconnection)**
 - addressing
 - NET*, 322-323
 - NSAP*, 320-321
 - end systems (ES), 316
 - ES-IS, 316
 - intermediate systems (IS), 316
 - IS-IS, 316
 - hierarchy*, 317
 - LSPDB*, 317
 - OSPF commonalities*, 316
 - SPT*, 317
 - IS-IS. *See* IS-IS
 - model, 315
 - protocol suite, 316

OSPF (Open Shortest Path First), 68

AD, modifying, 529-531

areas

ABRs, 242

backbone, 242

defined, 241

disadvantages, 242

external routes, 247-248

failed route advertisement, 242

filtering, 543-546

IDs, 245-246

*interarea summarization,
276-280*

*intra-area/interarea routes,
246-247*

multi-area topology, 243-245

NSSAs. See NSSAs

*successful route advertisement,
243*

Type 1 LSA flooding, 252

authentication, 236

IOS, 236

IOS XR, 237-239

types supported, 236

BFD

configurations

IOS interface specific, 204-205

*IOS network statement,
202-204*

IOS XR, 205

topology sample, 206-208

discontiguous networks, 301-303

DRs/BDRs, 213

elections, 214-216

LSA distribution, 213

placement, 216-219

exponential LSA sessions per routers
on same segment, 212

failure detection, 219

dead interval timer, 219

fast-packet Hellos, 220-221

Hello timer, 219

timers, verifying, 220

filtering

areas, 543-546

local, 541-543

hierarchical architecture, 192-193

interfaces

costs, 235-236

nonbroadcast, verifying, 302

verifying, 208-209

inter-router communication,
193-194

Hello packets, 194-195

multicast addresses, 193

packet types, 194

IS-IS commonalities, 316

LSAs, 191, 249-251

age, 251

flooding, 251

*reduction through area segmen-
tation, 274*

summary, 270

Type 1, 252-257

Type 2, 257-259

Type 3, 259-262

Type 4, 265-268

Type 5, 263-265

Type 7, 268-269

LSDB (link-state database), 192

multi-access networks topology, 212

multi-area adjacencies, 304-308

configuration, 305-306

- inefficient topologies*, 304
- nonbroadcast interface verification*, 304
- OSPF *neighborship*, 306-307
- verification*, 306-308
- neighbors, 196
 - adjacencies*. See *adjacencies*, OSPF
 - multi-area interface*, 306-307
 - non-broadcast*, 223
 - P2P interfaces verification*, 225
 - states*, 196, 210
 - Type 1 LSAs state*, 256
- network types, 221
 - broadcast*, 221
 - listing of*, 231
 - loopback*, 229-230
 - neighbor adjacency compatibility*, 231-234
 - non-broadcast*, 222-223
 - P2P*, 224-225
 - point-to-multipoint*, 225-229
- passive interfaces, 205-206
- path selection, 270-271
 - ECMP, 274
 - external routes*, 272
 - interarea routes*, 272
 - intra-area routes*, 271-272
 - Type 1 external routes*, 273
 - Type 2 external routes*, 273-274
- preferred over BGP for RPF calculation, 878
- prefix suppression, 308-312
 - configuration*, 310
 - connectivity verification*, 312
 - routing table*, 309, 311
 - Type 1 LSAs after*, 309, 311
 - Type 2 LSAs*, 310, 312
- processes, running, 193
- redistribution, 569-571
 - forwarding address*, 573-576
 - OSPF-to-OSPF*, 571-572
- RIB installed routes, verifying, 211-212
- RIDs, 196
- SPT, 192
- stubby areas, 286
 - NSSAs, 292-295
 - stub areas*, 286-289
 - totally*, 289-292
 - totally NSSAs*, 295-298
 - types*, 286
- summarization, 274-276
 - default*, 283-285
 - external*, 280-283
 - interarea*, 276-280
- version 2 versus version 3, 955-957
- virtual links, 298-301
 - configuring*, 300
 - verification*, 300-301
- OSPFv3, 955
 - authentication, 970-973
 - areas*, 972
 - interfaces*, 972
 - IPsec verification*, 973
 - configuration
 - commands reference chart*, 975
 - IOS, 960-961
 - IOS XR, 961
 - inter-router communication, 957
 - LSAs, 958
 - flooding scopes*, 959-960
 - types*, 959

- multiple instances, 973-975
- OSPFv2, compared, 955-957
- verification, 962-970
 - configuration*, 963-964
 - interfaces*, 965
 - LSDB*, 965-968
 - routes table*, 969-970
 - show commands*, 962
- ospfv3 authentication command, 975
- ospfv3 encryption command, 975
- ospfv3 process-id area area-id command, 975
- ospfv3 process-id area area-ID command, 975
- outbound route filtering (ORF), 647-649
- outgoing interface (OIF), 763
- outgoing interface list (OIL), 763
- out keyword (ip multicast boundary access-list command), 864
- out-of-band route reflectors, 453
- overload bits
 - IS-IS, 394-396
 - configuration*, 395
 - routing table*, 395-396
 - topology*, 394
 - LSPs, 331

P

- P2P (point-to-point) interfaces**
 - broadcast interfaces as, configuring, 353-355
 - static routing, 96-98
 - bidirectional connectivity verification*, 97-98
 - direct attachments*, 96-97

- routing table example*, 97
- serial connections*, 96

P2P (point-to-point) networks

- IS-IS neighbors, 338-339
- OSPF, 224-225
 - configuring*, 224
 - interface verification*, 225
 - neighbors on P2P interfaces*, *verifying*, 225

packages

- installation envelopes (PIEs)
- IOS, 2-4
- IOS XR, 6-7

packets

- CSNPs, 350
- EIGRP
 - advertised networks update*, 131
 - hello packet*, 129
 - INIT flag and acknowledgment*, 130
 - INIT flag set*, 129
 - queries, handling*, 156-157
 - route prefixes update*, 130
 - types*, 127

- EIGRPv6, 944

- Hello. *See* Hello packets

- IIH padding, removing, 365

- IP, delivering, 30

- IS-IS

- hello (IIH)*, 327-328

- IS PDU addressing*, 326-327

- IS protocol headers*, 325

- TLVs*, 326

- types*, 324

- link-state. *See* LSPs

- OSPF, 194

- OSPFv2 versus OSPFv3, 955
- OSPFv3, 957
- sequence number (SNPs), 324
- packet switching**
 - CEF, 86
 - centralized versus distributed forwarding architectures, 86*
 - defined, 86*
 - distributed, 88*
 - hardware, 88-89*
 - software, 87*
 - Layer 2 addressing, 83-84
 - process switching, 84-85
- parameterization**
 - route policies, 507-510
 - RPL, 496
- parameters**
 - NDP, displaying, 916-917
 - OSPF
 - show the commit command, 23
- parentheses () query modifier, 480**
- partial route calculation. *See* PRC**
- partition bits (LSPs), 331**
- PAs (path attributes)**
 - AIGP, exchanging, 687
 - BGP, 409
- passes-through option (AS_Path), 486-487**
- passive-interface command, 186**
- passive interfaces**
 - EIGRP, 134
 - IS-IS, 361-362
 - OSPF, 205-206
- Path and Origin field (BGP tables), 429**
- paths**
 - AS sets, 488
 - BGP algorithm, 672
 - AIGP. *See* AIGP
 - AS_Path length, 694-698*
 - AS_Path relax feature, 731-733*
 - attributes, 409, 672-673*
 - eBGP versus iBGP, 714-717*
 - ECMP. *See* ECMP, BGP
 - locally originated via network or aggregate advertisement, 683-686*
 - local preference, 679-682*
 - lowest IGP metric, 718-720*
 - lowest neighbor address, 722*
 - MED. *See* MED, BGP *best path selection*
 - minimum cluster list length, 721*
 - oldest EBGP, 720*
 - origin types, 700-703*
 - RIDs, 720-721*
 - weight, 673-679*
- EIGRP metrics, 145
 - attribute propagation, 146*
 - custom K values, 148-149*
 - default K values, 145*
 - with definitions, 145*
 - formula, 145*
 - interface delay settings, 149-151*
 - interface metrics, 146*
 - load balancing, 151-153*
 - lowest link speed/cumulative delay, 146*
 - specific prefix, 147-148*
 - variance multiplier, 151-152*
 - wide, 153-154*
- IS-IS selection, 386
 - ECMP, 387
 - interface metrics, 387-394*
 - processing order, 387*

- OSPF selection, 270-271
 - ECMP*, 274
 - external routes*, 272
 - interarea routes*, 272
 - intra-area routes*, 271-272
 - Type 1 external routes*, 273
 - Type 2 external routes*, 273-274
- selection algorithms
 - distance vector*, 69
 - enhanced distance vector*, 70
 - link-state*, 70-71
 - path vector*, 71-72
- vector algorithms, 71-72
- PBR (policy-based routing)**, 521
 - configuring, 522
 - local, 525-526
 - overview, 521
- PDMs (protocol-dependent modules)**, 126
- PDU (protocol data units)**, 324, 326
- peers**
 - BGP, 415, 664-665
 - MSDP, 822-828
 - IOS, 822
 - IOS XR, 822-823
 - MSDP and MBGP sessions between PIM domains*, 823-827
 - MSDP Compliance*, 823
 - SA message packet captures*, 828
 - SA RPF checks*, 823
 - status*, 829
- period (.) query modifier**, 481
- PIEs (package installation envelopes)**, 6
- PIM (Protocol Independent Multicast)**, 745, 762
 - bidirectional, 802-808
 - BSR, 775-776,
 - configuration*, 786-787
 - boundaries*, 866
 - configuring, 780-782
 - control messages, 764
 - Dense Mode (PIM-DM), 765-767
 - downstream, 763
 - first-hop router (FHR), 763
 - forwarder, 778-779
 - group mapping verification, 792
 - interfaces
 - downstream*, 763
 - incoming*, 763
 - outgoing interface (OIF)*, 763
 - RPF*, 762
 - upstream*, 763
 - verification*, 791
 - IPv6, 1015
 - last-hop router (LHR), 763
 - multicast
 - forwarding information base (MFIB)*, 764
 - routing information base (MRIB)*, 763
 - state*, 764
 - neighbor control, 869-870
 - operating modes, 764
 - outgoing interface list (OIL), 763
 - reference topology, 762
 - register rate limit, 870
 - RPF, 763, 777
 - RPs. *See* RPs
 - Sparse Mode. *See* PIM-SM
 - SSM. *See* PIM-SSM

- topology, verification, 793
- upstream, 763
- PIM-SM (PIM Sparse Mode), 768-772**
 - accept RP, 868-869
 - BSR, 1018-1021
 - DRs, 772
 - message types, 1015
 - multicast forwarding overview, 768
 - PIM tunnels, 1017
 - shared tree joins, 769
 - source registration, 769-770, 867-868
 - SPT switchover, 771
 - static RP configuration, 1017-1018
- PIM-SSM (PIM-Source Specific Multicast), 850, 1033-1034**
 - configuration, 852-853
 - group-to-PIM mode mapping, 853
 - IOS, enabling, 851
 - IOS XR, enabling, 851
 - mapping, 857-862
 - MDT creation, 851
 - mroute table, 855-857
 - RPF neighbors, identifying, 854
 - topology, 854
- pipe (|) query modifier, 480**
- placement**
 - DISs, 352-353
 - DRs/BDRs, 216-219
- planes of operation, 89-90**
- plus sign (+) query modifier, 481**
- point-to-multipoint networks, 225-229**
- point-to-point interfaces. *See* P2P interfaces**
- point-to-point networks. *See* P2P networks**
- policies (route)**
 - actions, 496
 - attachment points, 516
 - attribute modifications, 498
 - boolean operators, 504
 - conjunction, 504*
 - disjunction, 505*
 - negation, 504*
 - order of processing, 505-506*
 - conditional matching commands, 559
 - editors, 512-513
 - match statements, 497
 - nesting, 510-511
 - original value, 511-512
 - parameterization, 507-510
 - prefix sets versus prefix lists, 506
 - redistribution set actions, 559
 - states, displaying, 515-516
 - structure, 496-497, 499-504
 - advanced conditional statement nesting, 502*
 - bad RPL design, 503*
 - conditional match and action statement, 499*
 - conditional match with if-elseif-else logic, 501*
 - conditional match with if-elseif logic, 501*
 - conditional match with if-else logic, 500*
 - conditional statement nesting, 502*
 - good RPL design, 503*
 - if-elseif-else route-policy, 502*
 - inline policy set expansion, 500*
 - inline prefix filtering, 499*
 - nested conditional match, 503*

- prefix set filtering*, 500
- RFC 1918 route map with if-else logic*, 501
- verification
 - BGP table*, 517-518
 - redistribution*, 516-517
- policy-based routing**. *See* PBR, 521
- PRC (partial route calculation)**
- prefixes**
 - BGP advertisements, 425-427
 - eBGP attributes
 - listing of*, 441
 - local*, 441
 - remote*, 440
 - EIGRP
 - filtering*, 534-537
 - update packet*, 130
 - filtering, 593-595
 - GUAs, 900
 - independent convergence. *See* PIC
 - lengths, 73
 - lists, 473-474, 506
 - matching
 - basic pattern*, 471
 - ineligible matched prefixes*, 472
 - looking glass/route servers*, 483
 - matching length parameters*, 471
 - prefix lists*, 473-474
 - prefix sets*, 474-475
 - regex*. *See* *regex*
 - structure*, 471
 - maximum, 654-656
 - NDP valid lifetime/preferred lifetime values, 915-916
 - sets, 474-475
 - filtering route policy*, 500
 - prefix lists, compared*, 506
 - suppression, 632-634
 - IS-IS*, 401-405
 - OSPF*, 308-312
- prepend as-path command**, 498
- preventing loops**, 119-120
 - redistribution
 - prefix filtering*, 593-595
 - seed metrics, increasing*, 598-600
 - tagging*, 595-598
 - RRs, 451-452
- private ASNs**, 408, 656-658
- private BGP communities**, 610, 625-627
- private IP addresses**
 - defined, 58
 - IPv4 address ranges, 58
 - NAT, 58
- privileged mode (IOS)**, 10
- problems**. *See* **troubleshooting**
- processes**
 - IOSd, 4
 - OSPF, 193
 - scheduling
 - IOS*, 2
 - IOS XE*, 4
 - IOS XR*, 5
 - switching, 84-85
- prompts (configuration)**
 - IOS, 26
 - IOS XR, 26
- protocol data units (PDUs)**, 324, 326
- protocol-dependent modules (PDMs)**, 126

Protocol Independent Multicast. *See* PIM

protocols. *See also* names of specific protocols

administrative distance, 73-76

distance vector, 69

dynamic, 68

distance vector, 69

enhanced distance vector, 70

link-state, 70-71

listing of, 68

path vector, 71-72

enhanced distance vector, 70

link-state, 70-71

path vector, 71-72

RIPv1, 50

RIPv2, 54

routing table metrics, 75-76

pseudonode IDs (LSP IDs), 330

DISs, 349-350

IS-IS, displaying, 358-359

public ASNs, 408

pwd command, 26

Q

queries

boundaries, 157

messages (MLD), 1010, 1012

packets, 156-157

query modifiers (regex)

* (asterisks), 482

[] (brackets), 479

[^] (caret in brackets), 480

^ (carets), 478

\$ (dollar signs), 478-479

- (hyphen), 479-480

() (parentheses), 480

. (period), 481

| (pipe), 480

+ (plus sign), 481

? (question mark), 481-482

_ (underscores), 477-478

question mark (?) query modifier, 481-482

R

range ipv6-prefix/prefix-length command, 975

RAs (router advertisement)

NDP messages, 913-914

advertisements, 913

disabling, 918

lifetime values, 917-918

packet capture, 915

responses, 914

SLAAC DNS options, 923-924

RD (reported distance), 142

RDNSS (Recursive DNS Server), 923

reachability (NDP), 936-937

real-time operating system (RTOS), 5

recursion (static routes), 98-99, 105-107

multiple, 109-111

nonrecursive multihop configurations, 107

problems, 112-116

single, 108-109

topology, 105

XR1/R3 configuration, 105-106

Recursive DNS Server (RDNSS), 923

redirects (IPv6), 919-920

redistribution

challenges

*invalid routing tables, 589-590**route feedback, 583**solutions, 606**suboptimal routing, 584-588*

default seed metrics, 606

defined, 551

destination-specific behaviors

*BGP, 580-582**EIGRP, 563-566**EIGRP-to-EIGRP, 566-568**IS-IS, 576-578**IS-IS to IS-IS, 578-580**OSPF, 569-571**OSPF forwarding address,
573-576**OSPF-to-OSPF, 571-572*

IPv6, 1002-1005

*IOS, 1002**route tables, 1003-1005*

metrics, 558

RIB existence, 555-557

route selection conditional matching

command, 559

routing loops

*AD, 601-603**overview, 590-592**prefix filtering, 593-595**seed metrics, increasing,
598-600**summarization, 603-605**tagging, 595-598*

RPL, 516-517

sequential protocol, 555

set actions, 559

source

*protocols, 553**selection criteria, 606**specific behaviors, 561-563*

transitivity, 553-554

redundancy (RPs), 833-834

anycast, 847-849

Auto-RP load balancing, 836-840

*configuration, 837**group-to-RP mapping, 837,
839-840**redundancy Auto-RP group-to-
RP mapping, 838**redundancy configuration, 838*

Auto-RP with multiple RPs, 835-840

*configuration, 835**RP mapping, 836*

BSR, 840-846

*group filtering, 843-845**hash algorithm, 845-846*

static RPs, 846

regex (regular expressions), 475

BGP reference topology, 476

query modifiers

** (asterisks), 482**[] (brackets), 479**[^] (caret in brackets), 480**^ (carets), 478**\$ (dollar signs), 478-479**- (hyphen), 479-480**() (parentheses), 480**. (period), 481**| (pipe), 480**+ (plus sign), 481**? (question mark), 481-482**_ (underscores), 477-478*

- Register PIM control message, 764
- register rate limits (PIM), 870
- Register stop PIM control message, 764
- relay agents
 - IOS
 - configuration*, 927-928
 - verification*, 928
 - IOS XR proxy
 - configuration*, 928-929
 - verification*, 929
- remote-as as-number command, 1001
- removing
 - IIH padding, 365
 - private ASNs, 656-658
- rendezvous points. *See* RPs
- replace as-path command, 498
- replacing
 - commits, 19-20
 - configurations, 13
- reported distance (RD), 142
- report messages (MLD), 1010
- reserved multicast addresses, 750
- reverse path forwarding. *See* RPF
- RIB (Routing Information Base), 74
 - BGP routes, displaying, 431
 - OSPF installed routes, verifying, 211-212
- RID field (OSPF Hello packet), 194
- RIDs (router IDs)
 - BGP, 413, 720-721
 - EIGRP, 141
 - EIGRPv6, 945
 - OSPF, 196
 - OSPFv2 versus OSPFv3, 955
- RIPv1 (Routing Information Protocol Version 1), 50
- RIPv2 (Routing Information Protocol Version 2), 54, 68
- rollbacks (configuration changes), 21-22
- root command, 26,
- route flap dampening. *See* RFD
- route maps, 488
 - command syntax, 489
 - components, 489
 - conditional matching, 490
 - commands*, 559
 - complex matching*, 491-492
 - multiple match*, 491
 - continue keyword, 493-494
 - examples, 494-495
 - optional actions, 492
 - redistribution set actions, 559
 - sample, 489
- route-policy route-policy-name command, 954, 1001
- router advertisements. *See* RAs
- router bgp as-number command, 1001
- route reflectors. *See* RRs
- router eigrp as-number command, 132, 954
- router eigrp process-name command, 954
- router IDs. *See* RIDs
- router isis command, 988
- router isis instance-id command, router LSA bits, 968
- router ospfv3 process-id command, 975
- routers
 - autonomous systems, 67
 - dynamic routing protocols, 68
 - IS-IS specific levels, 384-385
 - LSAs, 959

- overview, 67
- path selection algorithms
 - distance vector*, 69
 - enhanced distance vector*, 70
 - link-state*, 70-71
 - path vector*, 71-72
- planes of operation, 89-90
- routing tables
 - administrative distance*, 73-76
 - metrics*, 75-76
 - overview*, 72
 - prefix lengths*, 73
- types (LSPs), 331
- VRFs. *See* VRFs
- Routing Information Base. *See* RIB**
- Routing Information Protocol Version 1 (RIPv1), 50**
- Routing Information Protocol Version 2 (RIPv2), 54, 68**
- routing policy language. *See* RPL**
- routing tables**
 - administrative distance, 73-76
 - after Ethernet link failure, 114
 - after OSPF interarea route summarization, 279
 - default routes
 - CE router*, 100-101
 - XR1/R3 example*, 102
 - EIGRPv6, 950
 - global
 - IP address configuration*, 80
 - output*, 80-81
 - invalid, 589-590
 - IS-IS
 - default routes*, 401
 - overload bits*, 395
 - prefix suppression XR1*, 402
 - route-leaking*, 378-379
 - summarization*, 398
 - metrics, 75-76
 - OSPF
 - area external routes*, 248
 - loopback*, 230
 - multi-area topology*, 246-247
 - point-to-multipoint*, 228
 - prefix suppression*, 308-312
 - stubby areas*, 287
 - totally NSSAs*, 298
 - totally stubby areas*, 291
 - overview, 72
 - prefix lengths, 73
 - recursive static routes, 99
 - static routes with P2P interfaces, 97
 - virtual. *See* VRFs
- RPF (reverse path forwarding), 762, 776-777**
- IPv6 multicast routing, 1030-1032
- MSDP**
 - check for source subnet*, 831
 - check for source subnet without MBGP*, 831
 - SA RPF checks*, 823
- multicast traffic engineering rules, 871-872
- neighbor, 763, 854
- OSPF preferred over BGP for RPF calculations, 878
- RPL (routing policy language), 496**
 - actions, 496
 - architecture, 496
 - attachment points, 516
 - attribute modification actions, 498
 - boolean operators, 504
 - conjunction*, 504

- disjunction*, 505
- negation*, 504
- order of processing*, 505-506
- editors, 512-513
- examples, 513-515
- match statements, 497
- original value, 511-512
- parameterization, 496, 507-510
- prefix sets versus prefix lists, 506
- route policy nesting, 510-511
- route policy structure, 499-504
 - advanced conditional statement nesting*, 502
 - bad RPL design*, 503
 - conditional match and action statement*, 499
 - conditional match with if-elseif-else logic*, 501
 - conditional match with if-elseif logic*, 501
 - conditional match with if-else logic*, 500
 - conditional statement nesting*, 502
 - good RPL design*, 503
 - if-elseif-else route-policy*, 502
 - inline policy set expansion*, 500
 - inline prefix filtering*, 499
 - nested conditional match*, 503
 - prefix set filtering*, 500
 - RFC 1918 route map with if-else logic*, 501
- states, displaying, 515-516
- structure, 496-497
- verification
 - BGP table*, 517-518
 - redistribution*, 516-517
- rp-list keyword**, 867
- RP mapping agents (MAs)**, 774
- RPs (rendezvous points)**, 772
 - anycast, 847-849
 - Auto-RPs, 773
 - candidate RPs*, 773
 - Cisco-RP-announce message filtering*, 867
 - configuring*, 785-786
 - RP MAs*, 774
 - candidate, 1021-1024
 - configuring, 783-787
 - Auto-RPs*, 785-786
 - static*, 784-785
 - embedded, 1021-1024
 - PIM
 - bootstrap router*, 775-776
 - PIM-SM accept*, 868-869
 - redundant, 833-834
 - anycast*, 847-849
 - Auto-RP load balancing*, 836-840
 - Auto-RP with multiple RPs*, 835-840
 - BSR*, 840-846
 - static*, 773, 846
 - configuring*, 784-785
 - PIM-SM*, 1017-1018
 - switchover. *See* continuous forwarding
- RRs (route reflectors)**
 - BGP suboptimal routing, 733-734
 - Add-Path feature*, 739-742
 - RRs, adding*, 734-735
 - shadow RRs*, 735-737
 - shadow session RRs*, 738-739
 - table*, 450
 - configurations, 449-450

- loop prevention, 451-452
- out-of-band, 453
- rules, 446-447, 450
- topology, 448
- RTOS (real-time operating system), 5**
- RTP (Reliable Transport Protocol), 127**
- running configurations**
 - displaying, 15
 - target configurations, merging, 16
- run to completion schedulers, 2**

S

- SAFIs (subsequent address family identifiers), 812**
- SA (source active) messages, 818-819**
 - cache, 830
 - packet capture, 828
 - SA RPF checks, 823
- scalability**
 - BGP configuration**
 - IOS peers, 664-666*
 - IOS XR peer templates, 667-668*
 - iBGP, 432, 446
 - confederations, 453-458*
 - loop prevention in route reflectors, 451-452*
 - out-of-band route reflectors, 453*
 - route reflectors, 446-451*
 - RPL, 496
- scheduling processes**
 - IOS, 2
 - IOS XE, 4
 - IOS XR, 5

- scopes**
 - Auto-RP TTL, 862
 - IPv6
 - multicast, 909*
 - multicast boundary, 1032-1033*
 - unicast, 898-899*
 - OSPFv3 flooding, 959-960
- secondary connected networks, 94**
- secondary IP addresses, 60-61**
- security**
 - BGP, 459
 - authentication, 462-463*
 - eBGP multihop, 459-461*
 - TTL, 461-462*
 - multicast, 862
 - Auto-RP Cisco-RP-announce message filtering, 867*
 - Auto-RP TTL scoping, 862*
 - boundaries, 863-866*
 - PIM, 869-870*
 - PIM-SM, 867-869*
- seed metrics, redistribution, 598-600, 606**
- selective AS_Set (aggregation), 641-643**
- SEL (Selector) fields, 321**
- sender network IDs, identifying, 36-37**
- sequence number packets (SNPs), 324, 331**
- sequential protocol redistribution, 555**
- serial connections, 96**
- servers**
 - prefix matching, 483
 - stateful DHCPv6
 - IOS configuration, 930*
 - IOS XR configuration, 931-932*
 - verification, 932-934*

sessions

- BGP
 - clearing*, 549
 - IPv6 over IPv4*, 998-1001
 - overview*, 415
 - verification*, 421-424
- shadow RR, 738-739
- set actions, 559
- set local-preference command, 498
- set med command, 498
- set next-hp command, 498
- set origin command, 498
- set weight command, 498
- shadow RRs, 735-737
- shared trees, 761
- Shortest Path First. *See* SPF
- show bgp ipv6 unicast command, 993
- show bgp ipv6 unicast neighbors command, 993
- show bgp ipv6 unicast summary command, 993
- show clns interface command, 985
- show clns neighbor command, 985
- show configuration command, 15
- show configuration commit changes command, 18
- show configuration commit changes incremental command, 18
- show configuration commit list command, 17
- show configuration failed command, 20
- show configuration merge command, 16
- show interface interface-type interface-number command, 149
- show ip eigrp interface command, 136
- show ip interface command, 61
- show ip mroute command, 794
- show ip route vrf vrf-name command, 81
- show ipv4 route vrf-name command, 81
- show ipv6 mld groups command, 1024
- show ipv6 mld interface command, 1024
- show ipv6 mld traffic command, 1024
- show ipv6 mrrib route command, 1024
- show ipv6 mroute active command, 1024
- show ipv6 mroute command, 1024
- show ipv6 pim group-map command, 1024
- show ipv6 pim interface command, 1024
- show ipv6 pim neighbor command, 1024
- show ipv6 pim range-list command, 1024
- show ipv6 pim topology command, 1024
- show ipv6 pim traffic command, 1024
- show ipv6 pim tunnel command, 1024
- show ipv6 protocols command, 985
- show ipv6 route bgp command, 993
- show ipv6 route ospf command, 962
- show ipv6 rpf ipv6-address command, 1024
- show isis database command, 985
- show isis interface command, 985
- show isis neighbor command, 985
- show isis neighbors command, 985
- show mfib route command, 798
- show mfib route rate command, 1024
- show mfib route src-ip-address/group-address command, 886
- show mld groups command, 1024
- show mld interface command, 1024
- show mld traffic command, 1024
- show mrrib ipv6 route command, 1024
- show mrrib route command, 798

- show ospfv3 database command, 962
- show ospfv3 interface command, 962
- show ospfv3 neighbor command, 962
- show ospfv3 virtual-links command, 962
- show pim ipv6 bsr candidate rp command, 1024
- show pim ipv6 bsr election command, 1024
- show pim ipv6 bsr rp-cache command, 1024
- show pim ipv6 group-map command, 1024
- show pim ipv6 interface command, 1024
- show pim ipv6 neighbor command, 1024
- show pim ipv6 range-list command, 1024
- show pim ipv6 rpf ipv6-address command, 1024
- show pim ipv6 topology command, 1024
- show pim ipv6 traffic command, 1024
- show pim ipv6 tunnel info command, 1024
- show pim topology command, 799
- show protocols ipv6 command, 985
- show route ipv6 bgp command, 993
- show route ipv6 ospf command, 962
- show running-config command, 15
- show the commit command, 23
- SIA (stuck-in-active) timers, 159-160
- single recursive static routes, 108-109
- single-stage commits, 11
- single-topology command, 988
- SixXS ULA generator tool website, 905
- SLAAC (stateless address autoconfiguration), 895-924
 - DNS RA options, 923-924
 - extended unique identifiers, 920-921
 - router configuration, 921-923
- slow path, 84-85
- SMUs (software maintenance upgrades), 6
- SNPs (sequence number packets), 324, 331
- software
 - CEF, 87
 - maintenance upgrades (SMUs), 6
 - packaging
 - IOS, 2-4
 - IOS XR, 6-7
 - switching, 84-85
 - versions, 2-3
- source active. *See* SA messages
- sources
 - protocol-specific behaviors
 - BGP, 562-563
 - connected networks, 561
 - IS-IS, 561
 - redistribution protocols, 553, 606
 - registration filtering, 867-868
 - trees, 759-760
- Source Specific Multicast. *See* SSM
- Sparse Mode (PIM). *See* PIM-SM
- special IPv6 addresses, 911
- special-purpose reserved IP addresses, 59
- SPF (Shortest Path First)
 - calculation throttling
 - incremental. *See* iSPF
 - IS-IS, 361-362
 - tree. *See* SPT

SPF tree. *See* SPT

split-horizon command, 186

split horizons (EIGRP WANs),
179-182

SPT (SPF tree)

IS-IS, 317, 361-362

path verification, 800

PIM-SM switchover, 771

switchover, 799-800

Type 1 LSA calculation, 256

OSPF, 192

SSM (Source Specific Multicast), 850

configuration, 852-853

enabling

IOS, 851

IOS XR, 851

group-to-PIM mode mapping, 853

mapping, 857

DNS, 857-860

static, 860-862

MDT creation, 851

mroute table, 855-857

RPF neighbors, identifying, 854

topology, 854

standard ACLs, 468-469

stateful DHCPv6, 926-927

IOS relay agents

configuration, 927-928

verification, 928

IOS XR relay proxy agents

configuration, 928-929

verification, 929

servers

IOS configuration, 930

IOS XR configuration, 931-932

verification, 932-934

topology, 926

stateful switchover. *See* SSO

stateless address autoconfiguration.

See SLAAC

stateless DHCPv6, 924-926

configuration

IOS, 924-925

IOS XR, 925-926

verification, 926

statements

match, 497

network

*IOS EIGRP autonomous system
configuration, 132-133*

*IOS OSPF configuration,
202-204*

State refresh PIM control

message, 764

states

BGP neighbors, 415

Active, 416

Connect, 415-416

Established, 417

Idle, 415

OpenConfirm, 417

OpenSent, 416-417

IS-IS neighbors, 346

multicast, 764

OSPF neighbors, 196, 210, 256

route policies, displaying, 515-516

static joins

IGMP, 882-886

MLD, 1015

static routes

advantages/disadvantages, 94-95

broadcast interfaces, 98-99

classifications, 95

- default routes, 99-102
 - backup connectivity*, 101
 - CE routers*, 100-101
 - XR1/R3 example*, 101-102
- example, 95
- floating, 103-105
 - AD verification*, 104-105
 - backups, configuring*, 112
 - defined*, 103
 - Ethernet link failures*, 104
 - XR1/R3 configuration*, 103
 - XR1/R3 routing table*, 103-104
- IPv6, 941-942
 - configuration commands reference chart*, 943
 - configuring*, 942-943
 - show commands reference chart*, 943
- mroutes, 872-875
- multihop
 - multiple recursive lookups*, 109-111
 - overview*, 108
 - single recursive lookups*, 108-109
- nonrecursive multihop, 107
- null interfaces
 - defined*, 116
 - packet traces*, 118-119
 - routing loops*, 117-120
- P2P interfaces, 96-98
 - bidirectional connectivity verification*, 97-98
 - direct attachments*, 96-97
 - routing table example*, 97
 - serial connections*, 96
- recursion problems, 112-116
 - CEF table*, 114
 - Ethernet link failure*, 113-114
 - floating static routes for backup*, 112
 - floating static route working as intended routing table*, 115-116
 - next-hop IP addressing with outbound interface correction*, 115
 - next-hop IP address with outbound interface after Ethernet link recovery routing table*, 116
 - stable links routing table*, 112-113
- recursive lookup, 105-107
 - nonrecursive multihop configurations*, 107
 - topology*, 105
 - XR1/R3 configuration*, 105-106
- RPs, 773, 846
 - configuring*, 784-785
 - PIM-SM*, 1017-1018
- SSM mapping, 860-862
 - configuring*, 860
 - verification*, 861
- VRF, 121-123
 - IOS*, 122-123
 - IOS XR*, 121-122
- stream statistics**, 1029
- stub areas**, 286-289
 - configuration*, 287-288
 - routing table*, 287
- stubby areas (OSPF)**, 286
 - NSSAs*, 292-295
 - autonomous systems example*, 293

- concept*, 292
- configuration*, 293-294
- routing table after*, 294
- routing table before*, 293
- stub areas, 286-289
 - configuration*, 287-288
 - routing table*, 287
- totally, 289-292
 - configuration*, 290-291
 - routing table after*, 291
 - routing table before*, 290
 - topology*, 289
- totally NSSAs, 295-298
 - concept*, 295
 - configuration*, 297
 - routing table after*, 298
 - routing table before*, 296
 - topology*, 296
- types, 286
- stub command**, 975
- stub no-summary command**, 975
- stubs**
 - EIGRP, 160
 - default configuration*, 162-163
 - designs*, 164-166
 - inter-region traffic*, 161-162
 - regional/remote topology*, 161
 - route tables after link failures*, 163
 - MSDP, 831-833
- stuck-in-active (SIA) timers**, 159-160
- subnet masks**
 - binary conversion, 36
 - classful routing
 - best path confusion*, 53
 - discontiguous networks*, 53
 - uniform*, 50-51
 - variable-length*, 52
 - defined, 30
 - interesting octets, 42-43
 - IP address interface assignments, 60
 - network prefix notation, 38
 - prefix lengths, 73
 - reference chart, 40
 - sender network IDs, identifying, 36-37
 - subnets, calculating
 - bitwise AND operations*, 41-42
 - magic number method*, 42-45
 - usable IP address, calculating, 37-38
 - variable-length. *See* VLSMs
 - wildcard, 62-63
 - calculating*, 63
 - dot-decimal notation*, 62
 - example*, 63
 - rules*, 62
- subnetting**
 - address range calculations
 - 10.55.200.33/12, 45
 - 172.16.2.3/23, 45
 - 192.168.100.5/26, 45
 - design, 46
 - available host values, listing*, 47
 - available subnets, listing*, 47
 - binary bit values, listing*, 47
 - final address allocation scheme*, 49
 - LAN 1/LAN 2 subnet assignments*, 47-48
 - LAN 3 assignments*, 48
 - explicit, 133
 - IP address fields, 39
 - IPv6 subnetting chart, 903

- overview, 38-39
 - subnet mask reference chart, 40
 - subnets, calculating
 - bitwise AND operations*, 41-42
 - magic number method*, 42-45
 - suboptimal routing**
 - AD, 527
 - example, 377
 - fixing, 377-380
 - redistribution challenges, 584-588
 - AD, 584
 - multipoint redistribution topology*, 584
 - resolving*, 588
 - route feedback*, 585
 - with RRs (BGP), 733-734
 - Add-Path feature*, 739-742
 - RRs, adding*, 734-735
 - shadow RRs*, 735-737
 - shadow session RRs*, 738-739
 - subsequent address family identifiers (SAFIs)**, 812
 - successors**
 - defined, 142
 - feasible, 142
 - summarization**
 - BGP, 628
 - aggregate addresses*, 629-631
 - AS_Set aggregation*, 639-641
 - atomic aggregate*, 637-639
 - flexible route suppression*, 632-637
 - selective AS_Set*, 641-643
 - EIGRP
 - automatic*, 172-174
 - hierarchical*, 166
 - interface-specific*, 166-171
 - metrics*, 171-172
 - EIGRPv6, 950-952
 - IS-IS, 396
 - configuration*, 399
 - discard routes*, 399
 - metric*, 396-397
 - ranges*, 397
 - routing table*, 398
 - topology*, 398
 - OSPF routes, 274-276
 - default*, 283-285
 - external*, 276-283
 - redistribution routing loop prevention, 603-605
 - summary-address ipv6-prefix/prefix-length command**, 954, 975
 - summary-address network command**, 186
 - summary-address prefix/prefix-length command**, 988
 - summary-metric network bandwidth delay reliability load MTU [AD] command**, 188
 - summary prefixes**, 57-58
 - summary-prefix prefix/prefix-length command**, 988
 - supernetting. *See* summarization
 - SysDB, 17-18
 - system ID (LSP IDs), 330
-
- T**
- tags (route)**, 595-598
 - target configurations**
 - displaying, 15
 - running configurations, merging, 16

TCP (BGP)

communication, 410

templates (peer)

IOS, 665-666

IOS XR, 667-668

timers**EIGRP**

hello/hold, 155-156

SIA, 159-160

IS-IS

hello, 366-367

holding, 367

OSPF

dead interval, 219

Hello, 219

verification, 219

timers active-time command, 188

Time-To-Live. *See* TTL

TLVs (Type, Length, Value)

IIH packets, 328

IPv6 IS-IS, 944-978

IS-IS route-leaking interarea, 379-380

LSPs, 332

overview, 326

topology base command, 185, 954

totally NSSAs, 295-298

concept, 295

configuration, 297

routing table after, 298

routing table before, 296

topology, 296

totally stubby areas, 289-292

configuration, 290-291

routing table after, 291

routing table before, 290

topology, 289

traces (IOS XR), 8

troubleshooting

IOS XR traces, 8

multicast, 886-889

mtrace, 887-889

show commands, 886

static route recursion, 112-116

CEF table, 114

Ethernet link failure, 113-114

floating static routes for backup, 112

floating static route working as intended routing table, 115-116

next-hop IP address with outbound interface after Ethernet link recovery routing table, 116

next-hop IP address with outbound interface correction, 115

stable links routing table, 112-113

TTL (Time-To-Live), 862

Auto-RP scoping, 862

BGP security, 461-462

eBGP multihop sessions, 461

two-stage commit (IOS XR), 14-17

Type, Length, Value. *See* TLVs

Type 1 external routes, 273

Type 1 LSAs, 252-257

areas

1234 example, 253-255

flooding, 252

LSDB fields, 252

OSPF

neighbor states, 256

prefix suppression, 309, 311

- SPF tree calculation, 256
- stub network link types, 256
- summary output, 252
- visualization, 257
 - before*, 310
 - after*, 312
- Type 2 external routes, 273-274**
- Type 2 LSAs, 257-259**
- Type 3 LSAs, 259-262**
 - conceptual, 259
 - detailed output, displaying, 261-262
 - fields, 262
 - summary output, 260
 - visualization, 262
- Type 4 LSAs, 265-268**
 - detailed output, 267
 - fields, 267
 - generic output, 266
 - metrics, 265
- Type 5 LSAs, 263-265**
 - detailed output, 264
 - fields, 265
 - flooding, 263
 - generic output, 263
- Type 7 LSAs, 268-269**
 - detailed output, 269
 - fields, 269
 - generic output, 268

U

- ULA (unique local unicast) addresses, 904-905
- underscores (`_`) query modifier, 477-478
- unequal load balancing, 151-153

unicast

- BGP session summary verification, 816
- IPv6, 898-900
 - addresses*, 894
 - configuring*, 899
 - global*, 900-903
 - interface status*, 899-900
 - link-local*, 905-906
 - scopes*, 898-899
 - unique local*, 904-905
- video feed, 746
- unique-length option (AS_Path), 486**
- unique local unicast addresses (ULA), 904-905**
- universal images, identifying, 3-4**
- unspecified addresses (IPv6), 911**
- updates**
 - BGP, 413-414
 - CIDR, 54
 - classful routing, 50
 - EIGRP packets
 - advertised networks*, 131
 - INIT flag and acknowledgment*, 130
 - INIT flag set*, 129
 - route prefixes*, 130
 - VLSM classless routing, 55
- upstream, 763**
- user mode (IOS), 9-10**

V

- variable-length subnet masks. *See* VLSMs
- variance variance-multiplier command, 188

verification**BGP**

neighbors supporting AIGP metrics, 690

ORF, 649

sessions, 421-424

bidirectional connectivity between routers, 97-98

DHCPv6

stateful, 932-934

stateless, 926

DIS, 353

DNS SSM mapping, 859

DR manipulation, 217

EIGRP

authentication, 176

hello/hold timers, 155

interfaces, 136-139, 149

key chain settings, 176

K values, 148-149

neighbor adjacencies, 139

EIGRPv6, 947-950

base configuration, 948-949

neighbor adjacency, 949-950

routing table, 950

show commands, 947

topology, 948

floating static routes AD, 104-105

IOS relay agents, 928

IOS XR relay proxy agents, 929

IPv6 multicast routing, 1024

IS-IS

broadcast parameters, 336

interfaces, 343-346, 392

IPv6, 985-987

neighbor adjacencies, 346-347

passive interfaces, 364

routes, 347-348

MBGP for IPv6

commands, 993

IPv6 route table, 997

IPv6 unicast BGP table, 996

neighbor status, 996

next-hop address selection, 997

router configuration, 994-995

MSDP, 828-831

Mroute table, 829

peer status, 829

RPF check for source subnet, 831

RPF check for source subnet without MBGP, 831

SA cache, 830

multi-area

adjacencies, 306

OSPF routes, 307-308

multicast

active traffic flows, 801

configuration, 786-787

DR election, 792, 795

IGMP interface, 790

IGMP joins, 791

IOS show ip mroute flags, 794

PIM group mapping, 792

PIM interfaces, 791

PIM topology, 793

show mrib route/show mfib route flags, 798

show pim topology flags, 799

source registration between RP and FHR tree creation, 796-798

SPT between RP and FHR creation, 796-798

- SPT path*, 800
- SPT switchover*, 799-800
- stream statistics*, 801
- OSPF
 - authentication*, 238, 239
 - interfaces*, 208-209
 - neighbor adjacencies*, 209-210
 - non-broadcast interfaces*, 223, 302
 - non-broadcast neighbors*, 223
 - point-to-multipoint network*, 227
 - prefix suppression connectivity*, 312
 - RIB installed routes*, 211-212
 - timers*, 219
 - virtual links*, 300-301
- OSPFv3, 962-970
 - configuration*, 963-964
 - interfaces*, 965
 - IPsec*, 973
 - LSDB*, 965-969
 - routes table*, 969-970
 - show commands*, 962
- RPL
 - BGP table*, 517-518
 - redistribution*, 516-517
- static SSM, 861
- unequal load balancing, 152-153
- unicast BGP session summary, 816
- versions**
 - active software, identifying, 2-3
 - MLD, 1010
 - universal image, identifying, 3-4
- virtual links (OSPF)**, 298-301
 - configuring, 300
 - verification, 300-301
- VLSMs (variable-length subnet masks)**
 - classless routing update, 55
 - defined, 55
 - successful route convergence, 55-56
- vrf definition vrf-name command**, 77
- VRFs (Virtual Routing and Forwarding)**, 76
 - configuring
 - example*, 81
 - IOS*, 78
 - IOS XR*, 78-79
 - defined, 76
 - displaying, 81-82
 - interfaces/IP addresses status verification, 82-83
 - IPv4-only, 77
 - Lite, 76
 - migrating IPv4-only to multiprotocol, 77
 - MPLS support, 76
 - multiprotocol, 77
 - overlapping IP addresses/interfaces between global routing table and VRFs
 - global routing tables*, 80-81
 - VRF configuration example*, 81
 - overlapping IP address problems, 76-77
 - routers with/without comparison, 79
 - static VRF routes, 121-123
 - IOS*, 122-123
 - IOS XR*, 121-122
- vrf upgrade-cli multi-af-mode command**, 77

W

WANS (EIGRP), 177

- bandwidth percent, 177-179
- next-hop self behavior, 182-183
- split horizons, 179-182

websites

- Cisco, 3
- GUA prefixes, 900
- IANA AFI listing, 812
- looking glass servers, 483
- private BGP community patterns, 610
- route servers, 483
- SixXS ULA generator tool, 905

weight (BGP), 673-679

- configuration, 674
- IOS, 678
- IOS XR, 678
- processing logic, 676
- tables, 429
- topology, 673

well-known BGP communities

- Internet, 611
- No_Advertise, 614-617
- No_Export, 611-614
- No_Export_SubConfed, 617-620

wide metrics

- EIGRP, 153-154
- IS-IS, 388

wildcard subnet masks, 62-63

- calculating, 63
 - /24 networks, 64*
 - single IP hosts, 63*
 - summary routes, 64*
- dot-decimal notation, 62
- example, 63
- rules, 62