

Official Cert Guide



Practice tests



Flash Cards



Review Exercises



Study Planner

CCNP

SPCOR 350-501

ciscopress.com

BRADLEY RIAPOLOV, CCIE® No. 18921
MOHAMMAD S. KHALIL, CCIE® No. 35484

FREE SAMPLE CHAPTER |



CCNP SPCOR 350-501 Official Cert Guide

Companion Website and Pearson Test Prep Access Code

Access interactive study tools on this book's companion website, including practice test software, review exercises, Key Term flash card application, a study planner, and more!

To access the companion website, simply follow these steps:

1. Go to www.ciscopress.com/register by December 31, 2027.
2. Enter the **print book ISBN: 9780135324806**.
3. Answer the security question to validate your purchase.
4. Go to your account page.
5. Click on the **Registered Products** tab.
6. Under the book listing, click on the **Access Bonus Content** link.

When you register your book, your Pearson Test Prep practice test access code will automatically be populated with the book listing under the Registered Products tab. You will need this code to access the practice test that comes with this book. You can redeem the code at **PearsonTestPrep.com**. Simply choose Pearson IT Certification as your product group and log into the site with the same credentials you used to register your book. Click the **Activate New Product** button and enter the access code. More detailed instructions on how to redeem your access code for both the online and desktop versions can be found on the companion website.

If you have any issues accessing the companion website or obtaining your Pearson Test Prep practice test access code, you can contact our support team by going to pearsonitp.ehelp.org.

This page intentionally left blank

CCNP SPCOR 350-501

Official Cert Guide

BRADLEY RIAPOLOV, CCIE No. 18921

MOHAMMAD KHALIL, CCDE No. 2023::57, CCIE No. 35484

CCNP SPCOR 350-501 Official Cert Guide

Bradley Riapolov

Mohammad Khalil

Copyright© 2025 Cisco Systems, Inc.

Published by:

Cisco Press

Hoboken, New Jersey

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.

For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions department, please visit www.pearson.com/global-permission-granting.html.

\$PrintCode

Library of Congress Control Number: 2024947468

ISBN-13: 978-0-13-532480-6

ISBN-10: 0-13-532480-7

Warning and Disclaimer

This book is designed to provide information about the Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam. 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.

Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

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.

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.

GM K12, Early Career and Professional Learning: Soo Kang	Copy Editor: CAH Editorial
Alliances Manager, Cisco Press: Caroline Antonio	Technical Editor: Brad Edgeworth
Director, ITP Product Management: Brett Bartow	Editorial Assistant: Cindy Teeters
Executive Editor: James Manly	Cover Designer: Chuti Prasertsith
Managing Editor: Sandra Schroeder	Composition: codeMantra
Development Editor: Christopher A. Cleveland	Indexer: Timothy Wright
Senior Project Editor: Tonya Simpson	Proofreader: Jennifer Hinchliffe



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

Bradley Riapolov is a seasoned technical solutions architect at Cisco Systems, currently serving in the MassScale Infrastructure Group. Since joining Cisco in 2008, he has been at the forefront of designing and deploying cutting-edge networking solutions, leveraging his extensive expertise to meet the complex demands of today's technology landscape.

Before his tenure at Cisco, Bradley gained invaluable experience working with various Fortune 500 companies, where he was instrumental in designing and operating small, medium, and large networks. His 20-year career is marked by a diverse background in successfully implementing technical campaigns across multiple industries, including transport, service provider, enterprise, industrial, and mobility networking.

In addition to his role at Cisco, Bradley is a recognized thought leader and educator. As a Cisco Press author and a distinguished speaker at Cisco Live, he has contributed to the body of knowledge in networking, sharing his insights and expertise through various publications. Moreover, his dedication to education is widely demonstrated, not only within Cisco but also through his role as a NetAcad instructor, where he has mentored and guided aspiring network professionals.

Bradley's dedication to excellence and his ability to simplify complex concepts have made him a respected figure in the networking community. His contributions continue to shape the future of networking, driving innovation and excellence in the industry. Bradley's refreshing no-nonsense approach to problem-solving has earned him a credible reputation among customers and peers alike.

Mohammad Khalil is an experienced service provider and enterprise expert, having worked in several service provider networks within the MENA region. Currently, he is a leader with the Cisco Competitive Win Center team covering enterprise architecture and working closely with Cisco sales/technical teams on designing their solutions.

Mohammad is a passionate networking expert, following technology trends and certifications development by writing several work labs and design scenarios. He was honored to be one of the SMEs for the CCIE Service Provider blueprint, SME for updated content of the CCIE Enterprise, and president of the Jordan IPv6 Council (part of the IPv6FORUM).

About the Technical Reviewer

Brad Edgeworth, CCIE No. 31574 (R&S and SP), is an SD-WAN technical solutions architect at Cisco Systems. Brad is a distinguished speaker at Cisco Live, where he has presented on various topics. Before joining Cisco, Brad worked as a network architect and consultant for various Fortune 500 companies. Brad's expertise is based on enterprise and service provider environments, with an emphasis on architectural and operational simplicity. Brad holds a bachelor of arts degree in computer systems management from St. Edward's University in Austin, Texas. Brad can be found on X (formerly Twitter) as @BradEdgeworth.

Dedications

To my beautiful wife, Evelina, who has supported me throughout the development of this book. You have challenged my assumptions, sharpened my ideas, and made some parts truly exceptional.

—*Bradley Riapolov*

I want to dedicate this book with affection to my wife, Qamar, and my precious children, Sireen and Saeed, who have supported me and motivated me to accomplish this effort. To my wise father, who believed in me and encouraged me all the way. Lastly, to my second home Estarta, which provided me with the support and environment to walk through.

—*Mohammad Kbalil*

Acknowledgments

I would like to thank the Cisco Press team, especially James Manly and Christopher Cleveland, for their patience, guidance, and consideration.

I would like to thank the technical editor, colleague, and friend, Brad Edgeworth, for his time, technical expertise, and readiness to generously share knowledge without expecting anything in return.

I would like to thank Bikram Gandhok, Cisco Certification Program Manager, for driving Cisco Certifications to stay relevant and impactful.

I would like to thank Kent Dailey, my steadfast accomplice, who has navigated both challenges and victories with equal zeal.

I would like to thank my mentors, Marty Fierbaugh and Emerson Moura, for helping me gain insights that extend beyond my own perceptions.

Finally, I would like to thank my co-author, Mohammad, for sharing the burden of this challenging endeavor.

—Bradley Riapolov

I would like to massively thank our development editor, Chris Cleveland, for his guidance and the contribution he provided. James Manly's support was remarkable; a big thanks for his kindness.

I want to thank our technical editor, Brad, for his great insights, patience, and kindness for sharing his amazing practical experience and guiding us throughout the content development.

And I don't want to forget to thank my friend, Adeeb Al-Husseini, for his time sharing his practical experience and insights.

Finally, I would like to thank my partner, Bradley, for his seamless cooperation, contribution, and kindness to share his ideas and thoughts with me.

—Mohammad Khalil

Contents at a Glance

Introduction xxx

Part I Architectures

Chapter 1 Service Provider Architectures 2

Chapter 2 Software Architectures 40

Chapter 3 Service Provider Virtualization 60

Part II Routing

Chapter 4 Routing Fundamentals 72

Chapter 5 IS-IS 108

Chapter 6 OSPF 158

Chapter 7 BGP Fundamentals 228

Chapter 8 BGP Optimization and Convergence 316

Chapter 9 Multicast 352

Part III Transport Protocols

Chapter 10 MPLS Fundamentals 452

Chapter 11 MPLS L2VPNs 494

Chapter 12 MPLS L3VPNs 554

Chapter 13 Advanced MPLS Services 622

Chapter 14 MPLS Traffic Engineering 666

Chapter 15 Segment Routing 724

Part IV Service Provider Security

Chapter 16 Securing Control Plane 764

Chapter 17 Securing Management Plane 836

Chapter 18 Securing Data Plane 854

Part V Critical Operational Elements

Chapter 19 IPv6 Transitions 882

Chapter 20 High Availability Designs 904

Chapter 21 Quality of Service 936

Part VI Automation and Orchestration

Chapter 22 Automation and Assurance 976

Part VII Final Preparation

Chapter 23 Final Preparation 1012

Chapter 24 CCNP SPCOR (350-501) Exam Updates 1014

Part VIII Appendix

Appendix A Answers to the “Do I Know This Already?” Quizzes and Review Questions 1018

Index 1032

Online Elements

Appendix B Memory Tables

Appendix C Memory Tables Answer Key

Appendix D Study Planner

Glossary of Key Terms

Contents

	Introduction	xxx
Part I	Architectures	
Chapter 1	Service Provider Architectures	2
	“Do I Know This Already?” Quiz	2
	Foundation Topics	4
	Network Core Architectures	4
	Metro Ethernet	5
	<i>E-Line (Ethernet Line Service)</i>	5
	<i>E-LAN (Ethernet LAN Service)</i>	7
	<i>E-Tree (Ethernet Tree Service)</i>	7
	<i>Metro Ethernet Forum (MEF)</i>	8
	MPLS	8
	Unified MPLS	10
	Segment Routing	13
	<i>Segment Routing Traffic Engineering (SRTE)</i>	14
	<i>SRv6 (Segment Routing IPv6)</i>	15
	Transport Technologies	16
	DOCSIS	16
	xDSL	17
	TDM	19
	Dense Wavelength-Division Multiplexing (DWDM)	21
	xPON	24
	Routed Optical Networking	27
	Mobility (Packet Core, RAN xhaul Transport for 5G vRAN and ORAN Transport)	31
	Packet Core	32
	xHaul	33
	Functional Splits Options	34
	O-RAN	36
	Exam Preparation Tasks	37
	Review All Key Topics	38
	Define Key Terms	38
	Review Questions	38
	Bibliography	39

Chapter 2	Software Architectures	40
	“Do I Know This Already?” Quiz	40
	Foundation Topics	41
	Software Architectures	41
	IOS	42
	IOS XE	42
	IOS XR	47
	Exam Preparation Tasks	58
	Review All Key Topics	58
	Define Key Terms	58
	Command Reference to Check Your Memory	58
	Review Questions	59
Chapter 3	Service Provider Virtualization	60
	“Do I Know This Already?” Quiz	60
	Foundation Topics	61
	Virtualization Technologies	61
	NFV Infrastructure	61
	VNF Workloads	63
	Containers	63
	Application Hosting	68
	Exam Preparation Tasks	70
	Review All Key Topics	70
	Define Key Terms	70
	Command Reference to Check Your Memory	70
	Review Questions	71
Part II	Routing	
Chapter 4	Routing Fundamentals	72
	“Do I Know This Already?” Quiz	72
	Foundation Topics	74
	IP Routing	74
	Changing Natural Routing Protocol Behaviors	76
	Route Maps	77
	Route Policy Language (RPL)	88
	Prefix Lists	100
	Prefix Sets	103
	Exam Preparation Tasks	105

Review All Key Topics	105
Define Key Terms	105
Command Reference to Check Your Memory	106
Review Questions	107

Chapter 5 IS-IS 108

“Do I Know This Already?” Quiz	109
Foundation Topics	110
Implement IS-IS (IPv4 and IPv6)	110
IS-IS Topology	111
Basic IS-IS Configuration	112
IS-IS Single Topology and Multitopology	121
IS-IS Adjacencies	129
IS-IS Network Types	131
IS-IS Metrics	137
Route Advertisement	141
Overload Bit	146
Authentication	148
Back to IS-IS Areas	152
Troubleshooting IS-IS	155
Exam Preparation Tasks	155
Review All Key Topics	155
Define Key Terms	156
Command Reference to Check Your Memory	156
Review Questions	157
Bibliography	157

Chapter 6 OSPF 158

“Do I Know This Already?” Quiz	158
Foundation Topics	160
Implement OSPFv2	160
OSPF Overview	162
Designated Router and Backup Designated Router	169
DR and BDR Priorities	176
OSPF Timers	179
OSPF Network Types	181
Route Advertisement	186
OSPF Stubby Areas	193
OSPF Totally Stubby Areas	196

OSPF Not-So-Stubby Areas (NSSAs)	196
OSPF Totally Not-So-Stubby Areas	199
OSPF Multi-Area Adjacencies	200
OSPF Virtual Links	206
Implement OSPFv3	209
Configuring OSPFv3	211
Implementing OSPFv3 Authentication	218
OSPFv3 Multiple Instances	221
Troubleshooting OSPF	225
Exam Preparation Tasks	226
Review All Key Topics	226
Define Key Terms	226
Command Reference to Check Your Memory	226
Review Questions	227
Bibliography	227

Chapter 7 BGP Fundamentals 228

“Do I Know This Already?” Quiz	228
Foundation Topics	231
Introduction to BGP	231
Autonomous System (AS)	231
BGP ASN Representation	236
BGP Messages	238
BGP Neighbor States	239
BGP Address Families	244
BGP Prefix Advertisement	246
BGP Path Attributes	248
BGP Path Selection	255
Weight	261
Locally Originated	263
AS Path	263
Origin	267
MED	268
BGP Multipath	273
eBGP vs. iBGP	274
IGP Next Hop	275
Oldest Route	275
Router ID	275
Accumulated IGP (AIGP)	275

Aggregate Address	277
Redistribution	280
Communities	282
Community Lists	290
BGP Peer Templates	291
Cost Community	292
Regular Expressions	294
Filter Lists	295
Mitigating the Split-Horizon Rule with Route Reflection and Confederations	297
Route Reflection	297
Confederations	304
BGP Loop Prevention	306
Allowas-in	307
AS Override	308
Troubleshooting	308
ACL/Firewalls	308
TTL	309
Authentication	310
Autonomous System	310
Different Address Families	310
MTU	311
Exam Preparation Tasks	311
Review All Key Topics	311
Define Key Terms	312
Command Reference to Check Your Memory	312
Review Questions	314
Bibliography	314
Chapter 8 BGP Optimization and Convergence	316
“Do I Know This Already?” Quiz	316
Foundation Topics	317
Minimum Route Advertisement Interval (MRAI)	317
Fast Peering Session Deactivation	319
Next-Hop Tracking	322
BGP and IGP Interaction	326
BGP Route Dampening	328
Prefix Independent Convergence	333

BGP Shadow Route Reflector	340
BGP Best External	341
High Availability: BGP Graceful Restart	344
High Availability: BGP Non-Stop Routing	347
Exam Preparation Tasks	348
Review All Key Topics	348
Define Key Terms	349
Command Reference to Check Your Memory	349
Review Questions	350
Bibliography	350

Chapter 9 Multicast 352

“Do I Know This Already?” Quiz	352
Foundation Topics	354
Multicast	354
IGMP	356
IGMPv1	357
IGMPv2	357
IGMPv3	358
Implementing IGMP	358
MLD	379
Multicast Routing Protocol Types	392
PIM-SM Multicast Distribution Trees	392
Rendezvous Points	410
<i>PIM Auto-RP</i>	410
<i>PIM Bootstrap Router (BSR)</i>	416
Bidirectional PIM	418
Source-Specific Multicast (SSM)	424
PIMv6	433
MLD	434
PIMv2	437
Static RP	439
BSR	442
Exam Preparation Tasks	447
Review All Key Topics	447
Define Key Terms	448
Command Reference to Check Your Memory	448
Review Questions	450
Bibliography	451

Part III Transport Protocols**Chapter 10 MPLS Fundamentals 452**

- “Do I Know This Already?” Quiz 452
- Foundation Topics 454
- MPLS Fundamentals 454
 - Reserved Labels 455
 - Ethertype 455
 - Additions of MPLS 456
 - MPLS Label Operations 456
 - MPLS Label Types 457
 - Label Assignment 458
 - MPLS LDP Autoconfig 465
 - MPLS Label Assignment 467
 - MPLS Advertise Labels 469
- MPLS LDP Session Protection 477
 - MPLS LDP Session Authentication 480
- MPLS LDP IGP Synchronization 483
- MPLS OAM 488
- Exam Preparation Tasks 491
- Review All Key Topics 491
- Define Key Terms 492
- Command Reference to Check Your Memory 492
- Review Questions 493
- Bibliography 493

Chapter 11 MPLS L2VPNs 494

- “Do I Know This Already?” Quiz 494
- Foundation Topics 496
- Metro Ethernet 496
 - E-Line (Ethernet Line Services) 496
 - E-LAN (Ethernet LAN Services) 497
 - E-Tree (Ethernet Tree Services) 497
 - E-Access 498
- MPLS AToM 498
- VPLS 502
 - Discovery and Signaling 505
 - VPLS Signaled with LDP (Manual) 505
 - VPLS Signaled with LDP (Autodiscovery) 509
 - VPLS Signaled with BGP 513

H-VPLS	519
ITU-T G.8032 Ethernet Ring Protection Switching	525
<i>CFM Protocols and Link Failures</i>	526
<i>Ethernet Connectivity Fault Management</i>	526
<i>Customer Service Instance</i>	527
<i>Maintenance Domain</i>	528
<i>Ethernet CFM Maintenance Domain</i>	528
<i>Maintenance Point</i>	529
<i>Maintenance Association</i>	529
<i>Maintenance Endpoints</i>	530
<i>CFM Messages</i>	531
<i>Cross-Check Function</i>	532
<i>Ethernet CFM and Ethernet OAM Interaction</i>	533
Provider Bridges (802.1ad)	535
802.1ad Ports	537
Service Provider Bridges	538
S-Bridge Component	538
C-Bridge Component	539
NNI Port	539
Provider Backbone Bridging (PBB)	539
PBB-EVPN Components	541
EVPN	543
Next-Generation Solutions for L2VPN	545
<i>CE Multihoming</i>	545
<i>Frame Duplication</i>	545
<i>MAC Flip-Flopping</i>	545
MPLS-Based Data Plane EVPN	547
Exam Preparation Tasks	550
Review All Key Topics	550
Define Key Terms	551
Command Reference to Check Your Memory	551
Review Questions	552
Bibliography	552
Chapter 12 MPLS L3VPNs	554
“Do I Know This Already?” Quiz	554
Foundation Topics	556
MPLS L3VPN	556

Virtual Routing and Forwarding (VRF)	557
MPLS LDP	562
Multiprotocol-BGP (MP-BGP)	564
MPLS L3VPN Static PE-CE Routing	567
MPLS L3VPN OSPF PE-CE Routing	577
MPLS L3VPN OSPF Routing, Backdoor Link	585
OSPF Down Bit	589
MPLS L3VPN EIGRP PE-CE Routing	591
MPLS L3VPN BGP PE-CE Routing	595
BGP Allowas-in	596
BGP AS-Override	597
Route Target Filtering Mechanisms	598
Import and Export Maps	601
Route Target Constraint	604
MPLS L3VPN Route Reflector Deployment	607
Multicast VPN	610
Multicast VPN Service Types	611
Multicast Distribution Trees	611
Draft-Rosen	611
PIM/GRE mVPN: Routing Information Distribution Using PIM C-instances	613
NG Multicast for L3VPN–BGP/MPLS mVPN (NG mVPN)	614
Requirements for Support of PIM-SM SSM in an mVPN	614
BGP/MPLS mVPN: Carrying Multicast mVPN Routing Information Using C-Multicast Routes	615
Exam Preparation Tasks	617
Review All Key Topics	617
Define Key Terms	618
Command Reference to Check Your Memory	618
Review Questions	619
Bibliography	620
Chapter 13 Advanced MPLS Services	622
“Do I Know This Already?” Quiz	622
Foundation Topics	624
Unified MPLS	624
MPLS L3VPN Shared Services (Internet)	629

MPLS L3VPN Internet Access Option 1: VRF-Specific Default Route	629
MPLS L3VPN Internet Access Option 2: Separate PE-CE Interface	634
MPLS L3VPN Internet Access Option 3: Extranet with Internet-VRF	635
MPLS L3VPN Internet Access Option 4: VRF-Aware NAT	639
MPLS Inter-AS L3VPN	640
Inter-AS Option A	641
Inter-AS Option B	643
Inter-AS Option C	649
Inter-AS Option AB	652
Carrier Supporting Carrier (CsC)	654
Quality of Service (QoS)	659
MPLS QoS	660
Uniform Mode	661
Pipe Mode	661
Short Pipe Mode	662
Exam Preparation Tasks	662
Review All Key Topics	662
Define Key Terms	663
Command Reference to Check Your Memory	663
Review Questions	664
Bibliography	664

Chapter 14 MPLS Traffic Engineering 666

“Do I Know This Already?” Quiz	666
Foundation Topics	667
MPLS Traffic Engineering Fundamentals	667
OSPF Extensions for MPLS TE	670
IS-IS Extensions for MPLS TE	671
Constrained Shortest Path First	672
Resource Reservation Protocol (RSVP) Operation in MPLS TE	673
<i>RSVP PATH Message</i>	674
<i>RSVP RESV Message</i>	675
<i>RSVP Error Messages</i>	676
<i>RSVP Tear Messages</i>	676
<i>How to Place Traffic into a TE Tunnel</i>	678
Forwarding Adjacency	687
MPLS Fast Reroute (FRR)	695
FRR Terminology	696

Link Protection	697
Node Protection	697
FRR on IOS XR	703
MPLS TE QoS	712
The DiffServ-TE Solution	712
Maximum Allocation Model (MAM)	713
Russian Dolls Model (RDM)	714
DiffServ for DS-TE	715
DS-TE Modes	715
Class-Based Tunnel Selection (CBTS)	716
Policy-Based Tunnel Selection	720
Exam Preparation Tasks	721
Review All Key Topics	721
Define Key Terms	722
Command Reference to Check Your Memory	722
Review Questions	723
Bibliography	723
Chapter 15 Segment Routing	724
“Do I Know This Already?” Quiz	724
Foundation Topics	726
Segment Types	728
Global Segments	728
Local Segments	730
IGP Segments	731
<i>IGP Prefix Segments</i>	731
<i>IGP Adjacency Segment</i>	733
<i>Combining IGP Segments</i>	734
Segment Routing Control Plane	735
IS-IS Control Plane	735
OSPFv2 Control Plane	737
BGP Control Plane	739
SRv6 Control Plane	742
<i>SRv6 (Segment Routing over IPv6) Header</i>	742
<i>SRv6 Node Roles</i>	744
<i>SRv6 Micro-Segment SID (uSID)</i>	744
<i>SRv6/MPLS L3 Service Interworking Gateway</i>	745
<i>Co-existence with LDP</i>	746

Segment Routing Traffic Engineering	748
Segment Routing Policies	749
SR Policies and Candidate Paths	750
Binding-SID (BSID)	750
Flex-Algo	751
TI-LFA	753
<i>Terms from Remote LFA Technology</i>	754
<i>Classic LFA Limitations</i>	756
PCE-PCC Architecture	757
Exam Preparation Tasks	760
Review All Key Topics	760
Define Key Terms	760
Command Reference to Check Your Memory	760
Review Questions	762
Bibliography	763

Part IV Service Provider Security

Chapter 16 Securing Control Plane 764

“Do I Know This Already?” Quiz	764
Foundation Topics	766
CoPP	768
LPTS	772
Keeping LDP Safe	778
LDP Authentication	778
Label Advertisement Control (Outbound Filtering)	781
Label Acceptance Control (Inbound Filtering)	781
Label Allocation Filtering	783
Keeping BGP Safe	789
BGP Authentication	790
BGP TTL-Security	793
BGP Route Filtering	803
BGP Maximum-prefix	811
BGP Prefix Suppression	812
BGPsec	815
BGP Flowspec	822
Exam Preparation Tasks	830
Review All Key Topics	830

Define Key Terms	831
Command Reference to Check Your Memory	831
Review Questions	833
Bibliography	833

Chapter 17 Securing Management Plane 836

“Do I Know This Already?” Quiz	836
Foundation Topics	838
Management Plane Protection Fundamentals	838
Tracebacks	843
AAA and TACACS	846
REST APIs	847
DDoS	849
Exam Preparation Tasks	851
Review All Key Topics	851
Define Key Terms	851
Command Reference to Check Your Memory	851
Review Questions	852
Bibliography	852

Chapter 18 Securing Data Plane 854

“Do I Know This Already?” Quiz	854
Foundation Topics	855
Unicast Reverse Path Forwarding (uRPF)	855
Access Control Lists (ACLs)	865
Remote Triggered Black Hole (RTBH) Filtering	866
Media Access Control Security (MACsec)	872
Exam Preparation Tasks	879
Review All Key Topics	879
Define Key Terms	879
Command Reference to Check Your Memory	879
Review Questions	880
Bibliography	880

Part V Critical Operational Elements

Chapter 19 IPv6 Transitions 882

“Do I Know This Already?” Quiz	882
Foundation Topics	884
Introduction	884
NAT44	884

- CGNAT 887
- NAT64 888
 - Stateless NAT64 889
 - Stateful NAT64 892
 - DNS64 895
- DS-Lite 897
- MAP-E 898
- MAP-T 899
- Exam Preparation Tasks 900
- Review All Key Topics 900
- Define Key Terms 900
- Command Reference to Check Your Memory 900
- Review Questions 901
- Bibliography 901

Chapter 20 High Availability Designs 904

- “Do I Know This Already?” Quiz 904
- Foundation Topics 905
- NSF 906
 - IS-IS NSF 907
 - OSPF NSF 910
- NSR 913
- BFD 917
- Link Aggregation 921
- Exam Preparation Tasks 933
- Review All Key Topics 933
- Define Key Terms 933
- Command Reference to Check Your Memory 933
- Review Questions 934
- Bibliography 935

Chapter 21 Quality of Service 936

- “Do I Know This Already?” Quiz 936
- Foundation Topics 938
- Traffic Classification 940
- Traffic Policing 946
- Traffic Shaping 958

Congestion Avoidance	961
Traffic Marking	966
IPv6 Flow Label	971
Exam Preparation Tasks	972
Review All Key Topics	972
Define Key Terms	973
Command Reference to Check Your Memory	973
Review Questions	973
Bibliography	974

Part VI Automation and Orchestration

Chapter 22 Automation and Assurance 976

“Do I Know This Already?” Quiz	976
Foundation Topics	978
REST APIs	978
What Is SOAP?	980
What Is RPC?	981
Model-Driven Programmability	982
Client/Server Model	984
Clients and Servers with NETCONF	985
Understanding Transport Protocols and Data Models	985
NETCONF	988
NETCONF Capabilities	988
Understanding NETCONF Operations	989
YANG	990
gRPC	991
Network Services Orchestrator (NSO)	992
How Does It Work?	993
Benefits	993
NSO Use Cases	994
Network Service Orchestration Tools	995
Secure ZTP	995
Secure ZTP Components	996
Initial Setup	996
How Secure ZTP Works	997
NetFlow/IPFIX	999

Streaming Telemetry	1001
SNMP	1004
SNMPv2c	1005
SNMPv3	1005
Ansible and Terraform	1008
Exam Preparation Tasks	1009
Review All Key Topics	1009
Define Key Terms	1009
Command Reference to Check Your Memory	1009
Review Questions	1010
Bibliography	1010

Part VII Final Preparation

Chapter 23 Final Preparation 1012

Hands-on Activities	1012
Suggested Plan for Final Review and Study	1012
Summary	1013

Chapter 24 CCNP SPCOR (350-501) Exam Updates 1014

The Purpose of This Chapter	1014
About Possible Exam Updates	1015
Impact on You and Your Study Plan	1015
News About the Next Exam Release	1016
Updated Technical Content	1016

Part VIII Appendix

Appendix A Answers to the “Do I Know This Already?” Quizzes and Review Questions 1018

Index	1032
-------	------

Online Elements

Appendix B Memory Tables

Appendix C Memory Tables Answer Key

Appendix D Study Planner

Glossary of Key Terms

Command Syntax Conventions

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

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

Preface

I'm not an academic or engineer with a PhD, just an ordinary person like many of you. Back in 2000, when I lost my job, my college roommate's brother-in-law handed me a book on networking (*The CCNA Certification Guide*). I started my career at the very bottom, schlepping printers. I worked with various customers for eight years, thinking I was a pretty smart engineer until I joined Cisco, where smart engineers were as common as raindrops in a storm. Throughout my career, I've been fortunate to work with many networks, from small to massive, for over a quarter of a century. The insights in this book come from someone who's built real networks and learned from clever mentors, mistakes, and experiences.

While my writing style might not adhere to academic conventions, it remains practical, mirroring my approach to problem-solving. It offers an opportunity to assess your problem-solving approach and perhaps discover new perspectives. I aim to help you pass a particularly challenging exam and impart valuable skills to those who build networks, not just academics, whom I deeply respect. So, expect my approach to differ from what you're used to, with a focus on hands-on, real-world scenarios.

What makes this exam so tough? Having worked across various networking domains, I find that service provider networks are especially complex compared to their enterprise, data center, or mobility counterparts. Cisco acknowledges this complexity in its exam structure. Expect tough questions, and be pleasantly surprised when you encounter the ones easier than that.

As you prepare, pay close attention to the exam blueprint and how its sections are worded. Questions deliberately fall into four categories: Describe, Compare, Configure, and Troubleshoot. "Describe" sections assume a solid grasp of general knowledge and concepts. "Compare" sections probe deeper, expecting candidates to differentiate between similar topics. "Configure" sections require advanced familiarity to configure or spot errors accurately. "Troubleshoot" sections demand the highest skill level, testing your ability to solve complex problems with twists.

What's my top tip for acing the exam on your first attempt? Practice. I have structured my portions of the content for you to follow in the book and the command line. In theory, there is no difference between theory and practice. I can tell you that in practice, there is. There's no substitute for hands-on keyboard time in mastering service provider networks.

Once, at a networking convention, I spotted someone wearing a humorous T-shirt defining "engineer" as "someone who does precision guesswork based on unreliable data provided by those of questionable knowledge." During the exam and throughout your career, you might feel like that. But remember, you're not alone. This book is dedicated to those who've tackled the seemingly impossible in the past, those fixing networks alongside me today, and those learning to do the same in the near future. May some of us meet and recognize each other.

—Bradley Riapolov

To add to what Bradley mentioned, we tried our best to add new terms within the book from brainstorming and use cases to make it more realistic from practical experience and to assist with some design guidelines for the service provider technologies.

—Mohammad Khalil

Introduction

The Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam is the required “core” exam for the CCNP Service Provider certifications. This exam tests your knowledge of implementing core service provider network technologies including architecture, services, networking, automation, quality of service, security, and network assurance.

Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) is a 120-minute exam.

TIP You can review the exam blueprint from Cisco’s website at <https://learningnetwork.cisco.com/s/spcor-exam-topics>.

This book gives you the foundation and covers the topics necessary to start your CCNP Service Provider or CCIE Service Provider journey.

The CCNP Service Provider Certification

The CCNP Service Provider certification is one of the industry’s most respected certifications. In order for you to earn the CCNP Service Provider certification, you must pass two exams: the SPCOR exam covered in this book (which covers core security technologies) and one of four available service provider concentration exams of your choice, so you can customize your certification to your technical area of focus.

TIP The SPCOR core exam is also the qualifying exam for the CCIE Service Provider certification. Passing this exam is the first step toward earning both of these certifications.

The following are the CCNP Service Provider concentration exams:

- Implementing Cisco Service Provider Advanced Routing Solutions (300-510 SPRI)
- Implementing Cisco Service Provider VPN Services (300-515 SPVI)
- Automating Cisco Service Provider Solutions (300-535 SPAUTO)
- Designing and Implementing Cisco Service Provider Cloud Network Infrastructure (300-540 SPCNI)

The CCIE Service Provider Certification

The CCIE Service Provider certification is one of the most admired, elite, and challenging certifications in the industry. The CCIE Service Provider program prepares you to be a recognized technical leader. In order to earn the CCIE Service Provider certification, you must pass the SPCOR 350-501 exam and an eight-hour, hands-on lab exam. The lab exam covers very complex network service provider network scenarios. These scenarios range from designing through implementing, operating, and optimizing dual-stack solutions (IPv4 and IPv6) of complex service provider networks.

Cisco considers ideal candidates to be those who possess the following:

- Extensive hands-on experience with Cisco’s Service Provider portfolio
- Experience deploying Cisco’s wide assortment of legacy and modern service provider technologies
- Deep understanding of multiple transport protocols and multitenant segmentation solutions
- Hands-on experience with MPLS networks and VPN solutions
- Configuring and troubleshooting QoS, mobility networking, device hardening, and general and access control
- Deep understanding of network automation and orchestration constructs

The Exam Objectives (Domains)

The Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam is broken down into five major domains. The contents of this book cover each of the domains and the subtopics included in them, as illustrated in the following descriptions.

The following table breaks down each of the domains represented in the exam.

Domain	Percentage of Representation in Exam
1: Architecture	15%
2: Networking	30%
3: MPLS and Segment Routing	20%
4: Services	20%
5: Automation and Assurance	15%
	Total 100%

Here are the details of each domain:

Domain 1: Architecture: This domain is covered in Chapters 1–3, 14, 16–18, and 21.

1.1 Describe service provider architectures

1.1.a Core architectures (Metro Ethernet, MPLS, unified MPLS, SR, SRTE, SRv6)

1.1.b Transport technologies (xDSL, DWDM, DOCSIS, TDM, and xPON)

1.1.c Mobility (packet core, RAN xhaul transport for 5G vRAN and ORAN transport)

1.1.d Routed optical network

1.2 Describe Cisco network software architecture

1.2.a IOS

1.2.b IOS XE

1.2.c IOS XR

- 1.3 Describe service provider virtualization
 - 1.3.a NFV infrastructure
 - 1.3.b VNF workloads
 - 1.3.c Containers
 - 1.3.d Application hosting
- 1.4 Describe QoS architecture
 - 1.4.a MPLS QoS models (Pipe, Short Pipe, and Uniform)
 - 1.4.b MPLS TE QoS (MAM, RDM, CBTS, PBTS, and DS-TE)
 - 1.4.c DiffServ and IntServ QoS models
 - 1.4.d Trust boundaries between enterprise and SP environments
 - 1.4.e IPv6 flow label
- 1.5 Configure and verify control plane security
 - 1.5.a Control plane protection techniques (LPTS and CoPP)
 - 1.5.b BGP-TTL security and protocol authentication
 - 1.5.c BGP prefix suppression
 - 1.5.d LDP security (authentication and label allocation filtering)
 - 1.5.e BGP sec
 - 1.5.f BGP flowspec
- 1.6 Describe management plane security
 - 1.6.a Traceback
 - 1.6.b AAA and TACACS
 - 1.6.c RestAPI security
 - 1.6.d DDoS
- 1.7 Implement data plane security
 - 1.7.a uRPF
 - 1.7.b ACLs
 - 1.7.c RTBH
 - 1.7.d MACsec

Domain 2: Networking: This domain is covered in Chapters 4–8, 19, and 20.

- 2.1 Implement IS-IS (IPv4 and IPv6)
 - 2.1.a Route advertisement
 - 2.1.b Area addressing
 - 2.1.c Single/Multitopology
 - 2.1.d Metrics
- 2.2 Implement OSPF (v2 and v3)
 - 2.2.a Neighbor adjacency
 - 2.2.b Route advertisement

- 2.2.c Multiarea (addressing and types)
- 2.2.d Metrics
- 2.3 Describe BGP path selection algorithm
- 2.4 Implement BGP (v4 and v6 for IBGP and EBGP)
 - 2.4.a Neighbors
 - 2.4.b Prefix advertisement
 - 2.4.c Address family
 - 2.4.d Path selection
 - 2.4.e Attributes
 - 2.4.f Redistribution
 - 2.4.g Additional Paths
 - 2.4.h PIC
- 2.5 Implement routing policy language and route maps (BGP, OSPF, IS-IS)
- 2.6 Troubleshoot routing protocols
 - 2.6.a Neighbor adjacency (BGP, OSPF, IS-IS)
 - 2.6.b Route advertisement (BGP, OSPF, IS-IS)
- 2.7 Describe IPv6 transition (NAT44, NAT64, CGNAT, MAP-T and DS Lite)
- 2.8 Implement high availability
 - 2.8.a NSF / graceful restart
 - 2.8.b NSR
 - 2.8.c BFD
 - 2.8.d Link aggregation

Domain 3: MPLS and Segment Routing: This domain is covered in Chapters 10, 14, and 15.

- 3.1 Implement MPLS
 - 3.1.a LDP sync
 - 3.1.b LDP session protection
 - 3.1.c LDP neighbors
 - 3.1.d Unified MPLS
 - 3.1.e MPLS OAM
- 3.2 Describe traffic engineering
 - 3.2.a IS-IS and OSPF extensions
 - 3.2.b RSVP functionality
 - 3.2.c FRR
- 3.3 Describe segment routing
 - 3.3.a Segment types
 - 3.3.b SR control plane (BGP, OSPF, IS-IS)
 - 3.3.c Segment routing traffic engineering

- 3.3.d TI-LFa
- 3.3.e PCE-PCC architectures
- 3.3.f Flexible algorithm
- 3.3.g SRv6 (locator, micro-segment, encapsulation, interworking gateway)

Domain 4: Services: This domain is covered in Chapters 9, 11–13, and 21.

- 4.1 Describe VPN services
 - 4.1.a EVPN
 - 4.1.b Inter-AS VPN
 - 4.1.c CSC
 - 4.1.d mVPN
- 4.2 Configure L2VPN and Carrier Ethernet
 - 4.2.a Ethernet services (E-Line, E-Tree, E-Access, E-LAN)
 - 4.2.b IEEE 802.1ad, IEEE 802.1ah, and ITU G.8032
 - 4.2.c Ethernet OAM
 - 4.2.d VLAN tag manipulation
- 4.3 Configure L3VPN
 - 4.3.a Intra-AS VPN
 - 4.3.b Shared services (extranet and Internet)
- 4.4 Implement multicast services
 - 4.4.a PIM (PIM-SM, PIM-SSM, and PIM-BIDIR, PIMv6)
 - 4.4.b IGMP v1/v2/v3 and MLD
- 4.5 Implement QoS services
 - 4.5.a Classification and marking
 - 4.5.b Congestion avoidance, traffic policing, and shaping

Domain 5: Automation and Assurance: This domain is covered in Chapter 22.

- 5.1 Describe the programmable APIs used to include Cisco devices in network automation
- 5.2 Interpret an external script to configure a Cisco device using a REST API
- 5.3 Describe the role of Network Services Orchestration (NSO)
- 5.4 Describe the high-level principles and benefits of a data modeling language, such as YANG
- 5.5 Describe configuration management tools, such as Ansible and Terraform
- 5.6 Describe Secure ZTP
- 5.7 Configure dial-in/out, TCP, TLS, and mTLS certificates using gRPC and gNMI
- 5.8 Configure and verify NetFlow/IPFIX
- 5.9 Configure and verify NETCONF and RESTCONF
- 5.10 Configure and verify SNMP (v2c/v3)

Steps to Pass the SPCOR Exam

There are no prerequisites for the SPCOR exam. However, students must have an understanding of networking and cybersecurity concepts.

Signing Up for the Exam

The steps required to sign up for the Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam:

1. Create a Certiport account at <https://www.certiport.com/portal/SSL/Login.aspx>.
2. Once you have logged in, make sure that “Test Candidate” from the drop-down menu is selected.
3. Click the **Shop Available Exams** button.
4. Select the **Schedule exam** button under the exam you wish to take.
5. Verify your information and continue throughout the next few screens.
6. On the **Enter payment and billing** page, click the **Add Voucher or Promo Code** button if applicable. Enter the voucher number or promo/discount code in the field below and click the **Apply** button.
7. Continue through the next two screens to finish scheduling your exam.

Facts About the Exam

The exam is a computer-based test. The exam consists of multiple-choice questions only. You must bring a government-issued identification card. No other forms of ID will be accepted. You can take the exam at a Pearson Vue center or online via the OnVUE platform. Visit the OnVUE page for your exam program: <https://home.pearsonvue.com/Test-takers/OnVUE-online-proctoring/View-all.aspx>

Once there, navigate to the FAQs section of the page, where you’ll find helpful information on everything from scheduling your exam to system requirements, testing policies, and more.

NOTE Refer to the Cisco Certification site at <https://cisco.com/go/certifications> for more information regarding this, and other, Cisco certifications.

About the *CCNP SPCOR 350-501 Official Cert Guide*

This book maps directly to the topic areas of the SPCOR exam and uses a number of features to help you understand the topics and prepare for the exam.

Objectives and Methods

This book uses several key methodologies to help you discover the exam topics that need more review, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. The book does not try to help you pass the exam only by memorization; it seeks to help you to truly

learn and understand the topics. This book is designed to help you pass the Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam by using the following methods:

- Helping you discover which exam topics you have not mastered
- Providing explanations and information to fill in your knowledge gaps
- Supplying exercises that enhance your ability to recall and deduce the answers to test questions
- Providing practice exercises on the topics and the testing process via test questions on the companion website

How to Use This Book

To help you customize your study time using this book, the core chapters have several features that help you make the best use of your time:

- **Foundation Topics:** These are the core sections of each chapter. They explain the concepts for the topics in that chapter.
- **Brainstorming:** These encourage you to actively apply their knowledge. You are invited to take a step beyond mere fact retrieval and engage deeply with the material you've covered. This is your opportunity to think critically about what you've learned and attempt to apply it on your own. Most of the time, we guide you through the process, but these exercises are designed to help you assess an exam question and provide an educated, well-reasoned answer. By participating in these sessions, you'll develop the skills to approach challenges with confidence and creativity, ensuring you're prepared to succeed independently.
- **Exam Preparation Tasks:** After the "Foundation Topics" section of each chapter, the "Exam Preparation Tasks" section lists a series of study activities that you should do at the end of the chapter:
 - **Review All Key Topics:** The Key Topic icon appears next to the most important items in the "Foundation Topics" section of the chapter. The Review All Key Topics activity lists the key topics from the chapter, along with their page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic, so you should review these.
 - **Define Key Terms:** Although the Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam may be unlikely to ask a question such as "Define this term," the exam does require that you learn and know a lot of cybersecurity terminology. This section lists the most important terms from the chapter, asking you to write a short definition and compare your answer to the glossary at the end of the book.



- **Review Questions:** Confirm that you understand the content you just covered by answering these questions and reading the answer explanations in Appendix A.
- **Web-based practice exam:** The companion website includes the Pearson Cert Practice Test engine, which allows you to take practice exam questions. Use it to prepare with a sample exam and to pinpoint topics where you need more study.

How This Book Is Organized

This book contains 22 core chapters—Chapters 1 through 22. Chapter 23 includes preparation tips and suggestions for how to approach the exam. Each core chapter covers a subset of the topics on the Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR 350-501) exam. The core chapters map to the SPCOR topic areas and cover the concepts and technologies you will encounter on the exam.

The Companion Website for Online Content Review

All the electronic review elements, as well as other electronic components of the book, exist on this book's companion website.

To access the companion website, which gives you access to the electronic content that accompanies this book, start by establishing a login at www.ciscopress.com and registering your book by December 31, 2027. To do so, simply go to www.ciscopress.com/register and enter the ISBN of the print book: 9780135324806. After you have registered your book, go to your account page and click the **Registered Products** tab. From there, click the **Access Bonus Content** link to get access to the book's companion website.

Note that if you buy the Premium Edition eBook and Practice Test version of this book from Cisco Press, your book will automatically be registered on your account page. Simply go to your account page, click the **Registered Products** tab, and select **Access Bonus Content** to access the book's companion website.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following these steps, please visit www.pearsonITcertification.com/contact and select the **Site Problems/Comments** option. Our customer service representatives will assist you.

How to Access the Pearson Test Prep (PTP) App

You have two options for installing and using the Pearson Test Prep application: a web app and a desktop app. To use the Pearson Test Prep application, start by accessing the registration code that comes with the book. You can access the code in these ways:

- You can get your access code by registering the print ISBN 9780135324806 on <https://www.ciscopress.com/register>. Make sure to use the print book ISBN, regardless of whether you purchased an eBook or the print book. After you register the book, your access code will be populated on your account page under the Registered Products tab. Instructions for how to redeem the code are available on the book's companion website by clicking the **Access Bonus Content** link.

- If you purchase the Premium Edition eBook and Practice Test directly from the Pearson IT Certification website, the code will be populated on your account page after purchase. Just log in at <https://www.ciscopress.com>, click **Account** to see details of your account, and click the **Digital Purchases** tab.

NOTE After you register your book, your code can always be found in your account under the Registered Products tab.

Once you have the access code, to find instructions about both the PTP web app and the desktop app, follow these steps:

- Step 1.** Open this book's companion website, as was shown earlier in this Introduction under the heading "The Companion Website for Online Content Review."
- Step 2.** Click the **Practice Exams** button.
- Step 3.** Follow the instructions listed there both for installing the desktop app and for using the web app.

Note that if you want to use the web app only at this point, just navigate to www.pearsonstestprep.com, establish a free login if you do not already have one, and register this book's practice tests using the registration code you just found. The process should take only a couple of minutes.

Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- **Study mode:** Allows you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps.
- **Practice Exam mode:** Locks certain customization options, as it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness.
- **Flash Card mode:** Strips out the answers and presents you with only the question stem. This mode is great for late-stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode does not provide the detailed score reports that the other two modes do, so you should not use it if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual

chapters, simply deselect all the chapters and then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus. Each exam bank comes complete with a full exam of questions that cover topics in every chapter. The two exams printed in the book are available to you as well as two additional exams of unique questions. You can have the test engine serve up exams from all four banks or just from one individual bank by selecting the desired banks in the exam bank area.

There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-answer questions, and whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

Updating Your Exams

If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software while connected to the Internet, it checks if there are any updates to your exam data and automatically downloads any changes that were made since the last time you used the software.

Sometimes, due to many factors, the exam data may not fully download when you activate your exam. If you find that figures or exhibits are missing, you may need to manually update your exams. To update a particular exam you have already activated and downloaded, simply click the **Tools** tab and click the **Update Products** button. Again, this is only an issue with the desktop Windows application.

If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply click the **Tools** tab and click the **Update Application** button. This ensures that you are running the latest version of the software engine.

Figure Credits

Figure 1-15 © ITU 2024

Figure 7-2, baranoski.ca

Figures 17-1 and 17-2 © 2024 Postman, Inc.

Cover, insta_photos/shutterstock

Segment Routing

This chapter covers the following exam topics:

3.3 Describe segment routing

- 3.3.a Segment types
- 3.3.b SR control plane (BGP, OSPF, IS-IS)
- 3.3.c Segment routing traffic engineering
- 3.3.d TI-LFA
- 3.3.e PCE-PCC architectures
- 3.3.f Flexible algorithm
- 3.3.g SRv6 (locator, micro-segment, encapsulation, interworking gateway)

Segment Routing represents a cutting-edge technology designed to enhance and optimize the use of MPLS-based and IPv6 networks. This innovative approach introduces a suite of tools and concepts that not only simplify network operations but also significantly enhance the flexibility available to network operators. By allowing for more granular control over packet paths, Segment Routing empowers network operators to efficiently navigate the intricacies of modern networks, adapt to dynamic requirements, and minimize the challenges associated with traditional MPLS-based architectures.

In essence, Segment Routing is a beacon of progress in the networking realm, offering a robust set of solutions to meet the most demanding network requirements. Its advent marks a transformative shift in how networks are managed, providing operators with a versatile toolkit to navigate the complexities of traffic engineering, and ensuring that networks can dynamically respond to changing conditions in real time. By simplifying the routing paradigm and enhancing the scalability of MPLS-based networks, Segment Routing emerges as a pivotal technology, poised to play a crucial role in the evolution of modern network architectures.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 15-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Review Questions.”

Table 15-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Segment Types	1–2
Segment Routing Control Plane	3–4
Segment Routing Traffic Engineering	5–8
PCE-PCC Architecture	9–10

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which of the following are viable options for assigning SRGB ranges? (Select all that apply.)
 - a. Globally
 - b. Per IGP
 - c. Dynamically
 - d. Statically
2. Which labels cannot be a part of SRGB? (Select all that apply.)
 - a. 15999
 - b. 16000
 - c. 24999
 - d. 1,048,576
3. Which TLVs play an important role in Segment Routing? (Select all that apply.)
 - a. 22
 - b. 41
 - c. 135
 - d. 163
4. Which of the following is not an option for the SR Control Plane?
 - a. RSVP-TE
 - b. IS-IS
 - c. OSPF
 - d. BGP
5. Which of the following is not a viable option to supply SR policy to the ingress router?
 - a. NETCONF
 - b. FIB

- c. CLI
 - d. PCEP
6. Which behavior properly describes an SR policy?
 - a. An SR policy uses Network Service Orchestrator (NSO) to program its data plane.
 - b. SR-PCE will use BGP-LS to program an SR policy into the ingress node.
 - c. An SR policy encodes the list of constraints into the MPLS headers.
 - d. An SR policy will use BSID entries to program its forwarding table.
 7. Which of the following describes TI-LFA functionality? (Select all that apply.)
 - a. TI-LFA must rely on LDP functionality to provide double-segment protection.
 - b. TI-LFA uses next-hop neighbor as the point of local repair.
 - c. TI-LFA preprograms the post-convergence path into a router's data plane.
 - d. TI-LFA can use P space to calculate the post-convergence path.
 8. Viable repair tunnel endpoints are found at which intersections?
 - a. At midpoints of extended spaces
 - b. At the intersection of point of local repair and double segments
 - c. At the intersection of PQ nodes
 - d. At the intersection of P and Q spaces
 9. Which component serves the needs of long-term network engineering and capacity planning?
 - a. Crosswork Network Controller
 - b. Crosswork Optimization Engine
 - c. WAN Automation Engine
 - d. Crosswork Cloud
 10. Which of the following protocols are used by PCEP to calculate network paths? (Select all that apply.)
 - a. IS-IS
 - b. BGP-LS
 - c. RSVP-TE
 - d. OSPF

Foundation Topics

Over time, a skilled engineer learns to make the most of the available resources. Over the past three decades, various strategies have been explored to maximize the potential of computer networks. The traditional approach to steering network traffic has been to manipulate the Interior Gateway Protocol (IGP) and rely solely on destination-based routing strategies. This only provided a restricted range of options for directing traffic within the network. The main limitation of using IGP metrics has been the “all or nothing” approach. Some of the links inevitably become congested while other links remain underutilized even though they are high bandwidth or low latency. IGP metrics simply lack optimization capabilities because they do not allow you to map different services to different paths. Despite years of featured improvements, such challenges persist and remain unsolvable with conventional IGP manipulations and destination-based routing strategies.

Thus, in the 1990s, the development of the **Label Distribution Protocol (LDP)** and Resource Reservation Protocol with Traffic Engineering extensions (RSVP-TE) marked a significant advancement in networking. These capable adjunct protocols effectively addressed specific challenges and provided network operators with crucial Traffic Engineering tools to overcome a wide variety of traffic-steering issues. Network operators could now manipulate how traffic flowed through the network to make the most of the available paths. Nevertheless, while providing network optimization, these protocols also introduced intricate challenges.

In the case of LDP, an additional process had to be created and maintained on the network, which led to a complex interaction with the Interior Gateway Protocol. LDP-IGP synchronization problems would cause traffic disruptions until these two protocols could settle on an agreement regarding the best way to forward traffic.

When dealing with RSVP-TE, reserving bandwidth accurately involves placing traffic within RSVP-TE tunnels. While this approach is feasible in smaller networks with minimal traffic engineering needs, it becomes exceedingly intricate on a larger scale. Managing hundreds of tunnels, their backup paths, and upkeep of a pertinent set of rules in the face of a dynamically evolving network posed formidable challenges, demanding considerable time and effort. Such scaling issues led many operators to limit their RSVP-TE deployments to the fast reroute (FRR) use cases. RSVP-TE is also not “ECMP-friendly,” so it can never use all IGP-derived paths, forcing the operator to create more tunnels. An additional aspect to consider is that RSVP-TE generates a persistent “always-on” network state where every router must account for available bandwidth and that state must be constantly monitored. This incurs a cost in terms of network compute resources and the hardware required to sustain this continuous state irrespective of whether the network experiences congestion or not.

Could network optimization be further improved? These were the experiences and thoughts of the designers behind Segment Routing. They proposed that a properly designed network should have enough capacity to effectively handle an expected volume of traffic without congestion, even in the presence of a probable set of independent failures. IGP coupled with ECMP can competently absorb the majority of the traffic volume. In less frequent instances of congestion, traffic engineering tools would address applications intolerant of such network bottlenecks. This represented a simpler and more resource-efficient approach, both in terms of hardware and human efforts.

They proposed that it is better to distribute labels associated with IGP-signaled prefixes within the IGP framework itself, rather than relying on a separate protocol such as LDP to perform this task. This would solve the LDP-IGP synchronization problem during network failures because there now would be a single source of truth (IGP) to find available network paths. The network would precalculate such paths even before the failure occurred. IGP can have such “preknowledge”—think of an EIGRP-feasible successor—which is aware of the best available network route even before the failure occurs.

Second, why not give the network operator the power to direct any packet to any path at the ingress router only (where the packet enters the network), without having to maintain the expensive and complex “always-on” state throughout the entire network domain? This would lower control plane pressure, conserve network resources, and give the operator the full flexibility to force any packet anywhere.

Cisco engineers formulated the Segment Routing concept, obtained approval from their management in 2012, presented the idea to IETF in March 2013, authored numerous IETF drafts, and significantly influenced the industry in 2015 with the introduction of Segment Routing on IOS XR platforms (IOS and NX-OS carry some of the SR features). The market positively responded with other vendors supporting this capable technology. It is also referred to as Source Routing outside of Cisco. As of the present moment, Cisco alone has documented more than 1200 operational Segment Routing production deployments.

Segment Routing can be deployed on two data planes:

- MPLS data plane where segments will be encoded with MPLS labels.
- IPv6 data plane where segments will be encoded with IPv6 addresses.

Segment Types

Think of segments as a set of instructions. In Segment Routing, a source (an ingress router, as an example) chooses a certain path through the network and encodes the path in the packet header as an ordered list of instructions. These instructions are termed *segments* because they describe components of a divided whole.

In **SR-MPLS (Segment Routing based on MPLS data plane)**, such an identifier refers to an ordered list of segments represented by a stack of MPLS labels. When you instruct routers to follow these labels in a given sequence, the packets will take this path through the network. In **SRv6 (Segment Routing based on IPv6 data plane)**, it refers to an ordered list of segments encoded into a routing extension header. When you instruct routers to follow this list, the packets will flow via this path. In Figure 15-1, assigning such identifiers to routers can provide instructions for a specific path to send packets.




Figure 15-1 Assigning Segments to Network

Global Segments

Every router in a Segment Routing, or SR, domain understands such instruction and installs it in its forwarding table. This instruction is a domain-wide (watch the misleading name *global* because it is not known globally around the world) unique label value (a numerical number) that comes from the **Segment Routing Global Block (SRGB)** database. Table 15-2 shows how the **Label Switching Database (LSD)** carves the following default ranges (some can be changed) on Cisco routers running Segment Routing–capable software.

**Table 15-2** LSD Label Ranges

Label Range	Reserved for	Examples
0–15	Base special-purpose MPLS labels	0—IPv4 Explicit NULL 3—Implicit NULL
16–15,999	Static MPLS labels	LDP assigned
16,000–23,999	SRGB	Global Segments (i.e., SR)
24,000–1,048,575	Dynamic allocation	Adjacency Segments

In Figure 15-1, every node in the domain knows that label 16002 always and uniquely represents Router PE2, 16003 always represents Router PE3, and so on. These are referred to as Node SIDs (Segment Identifiers).

A Node SID is a type of Prefix SID, as it represents any prefix linked to a node. To send a set of segment routing instructions is to specify these labels (Node SIDs); 16002, 16003 literally means “send this traffic via shortest ECMP path to PE2, then same to PE3.” If the “uniqueness” rule is broken (two distinct routers are assigned the same label value), there will be an issue on your network because the nodes will not be able to accurately determine the appropriate path for routing traffic. (Technically, what happens is that IS-IS will prioritize the “first programmed” label and ignore the “second” one. This can get complex as to what “first programmed” means, as in when a router reboots or has failed, which router becomes “first programmed”? OSPF, on the other hand, will withdraw both SIDs. Don’t worry about this because this topic is highly unlikely to show up on the exam, but good to know.)

Global Segments are always distributed as a unique value via IGP (remember that SR no longer relies on LDP as the label distribution mechanism). This value must be unique and comes from a combination of a label range (SRGB) + index. The default SRGB range for Cisco routers is 16000–23999 (notice how it does not overlap with the LDP range) and the index is zero based (the first index = 0). In our scenario, we start adding index values to our routers (index 2 for PE2, index 3 for PE3), so we will come up with globally known unique values of 16002 and 16003 for these routers. Cisco gives an option to also define these as an absolute value. There is no difference in daily operations whether absolute or relative indexes are used.

There are two minimum requirements to enable Segment Routing:

- Configure the Segment Routing Global Block (SRGB)
- Enable Segment Routing and Node SID in the IGP (shown later in the chapter under respective protocols)

Example 15-1 demonstrates SRGB assignment on both IOS XE as well as IOS XR operating systems.

Example 15-1 SRGB Verification on IOS XE and IOS XR

```
IOS
PE2# show running-config segment routing
!
```

```

segment-routing
  global-block 16000 23999
!
PE2#

IOS XR
RP/0/0/CPU0:PE4# show running-config segment-routing
!
segment-routing
  global-block 16000 23999
!
RP/0/0/CPU0:PE4#

```

When you're configuring an SRGB block, the recommendation is to configure it globally (again misleading, in the global, i.e., router-wide configuration), not per individual IGP; this way, all IGP instances as well as BGP can use the global (i.e., router-wide) SRGB. This is important because, later on, if you choose to add another protocol and the SRGB range has to change, you will have to reload the router. To avoid this, it is better to assign the SRGB block globally in the router so that multiple protocols (including BGP) can use it, not just a specific IGP. You should have a homogenous SRGB block (same SRGB range) on all routers on the network. If you do not, you will have fun times building end-to-end LSPs.

Local Segments

This instruction is allocated and understood by the originating node. It is locally significant only (which aligns with the intended meaning regarding the local router). A locally allocated MPLS label would be a good example of a **local segment**. Sometimes a router has more than a single link for forwarding traffic. The operator can prefer one of these paths over the other. In Figure 15-2, a packet arrives at PE3 (via label 16003). PE3 intends to send this packet to PE6, and there are two shortest available paths—through P4 and P5. These links are identified with local labels 24100 and 24150. Selecting one of them will instruct the local router to pick the appropriate link.




Figure 15-2 *Assigning Local Segments to Network*

IGP Segments

Segments construct a path through a network. There are two building blocks distributed by IGP: Prefix Segments and Adjacency Segments.

IGP Prefix Segments

Think of an **IGP Prefix Segment** as a shortest path to the IGP prefix. Note that it is

- Known as Prefix-SID
- Associated with an IP prefix
- Represents an ECMP (Equal Cost Multi-Path)-aware shortest path to a prefix
- Likely a multihop path
- A Global Segment—known uniquely on the SR domain
- A Label (16000 + Index) that is advertised as an index
- Distributed via IGP (ISIS/OSPF)

Observe this behavior in Figure 15-3 (note that we removed two links to better illustrate our example), where different routers are used to forward traffic to P5 based on label 16005, destined to the loopback address of 10.1.100.5. For this Segment Routing domain, PE3, P4, PE6, and PE7 send traffic directly to P5 because when these routers look at the MPLS forwarding table, label 16005 will be associated with the one interface directly connecting to P5. The only exception is PE2 which will have two equal paths available due to ECMP. How did all the routers learn that prefix 10.1.100.5/32 is associated with label 16005? R5 has generated this value from the combination of the SRGB label base 16000, added its operator assigned index +5, and advertised this label via IGP.

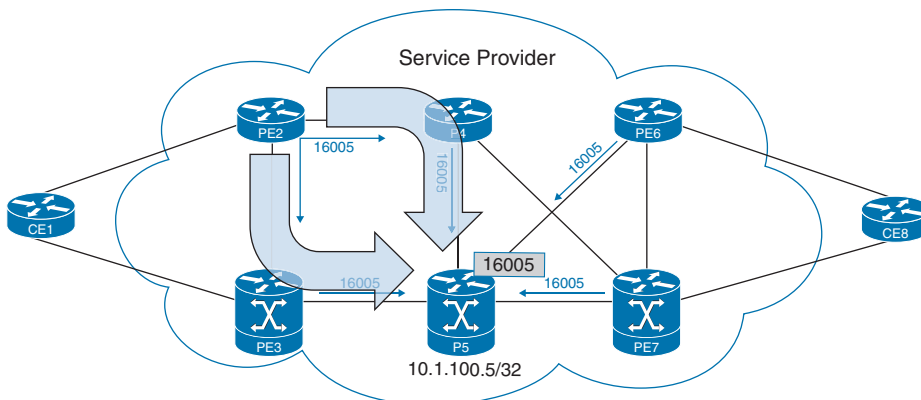


Figure 15-3 IGP Prefix Segment Behavior

I need to point something out here. Technically, there is a difference between a Node SID (which we just showed you) and a Prefix SID. A Node SID points to a router that can be a visiting point on the network. A Prefix SID is a label for a network prefix that is advertised by a router. This causes a point of confusion at times. Note that a special N flag is set to indicate that a SID represents a router (node) on the network. This is advanced and is unlikely

to show up on the exam, but we include Example 15-2 to clear up any confusion regarding the difference.

Example 15-2 Node SID and the N Flag

```

RP/0/0/CPU0:R2# show isis database R1.00-00 detail verbose
Thu May  9 13:03:44.757 UTC
IS-IS lab (Level-1) Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd  ATT/P/OL
The requested LSP R1.00-00 was not found in the IS-IS lab Level-1 LSP Database
IS-IS lab (Level-2) Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd  ATT/P/OL
R1.00-00             0x00000007   0xf6b7        1108 /1199         0/0/0
  Area Address:      49.0001
  NLPID:              0xcc
  NLPID:              0x8e
  IP Address:         1.1.1.1
  Metric: 10         IP-Extended 10.1.12.0/24
    Prefix Attribute Flags: X:0 R:0 N:0
  Metric: 10         IP-Extended 10.1.13.0/24
    Prefix Attribute Flags: X:0 R:0 N:0
  Metric: 0          IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
    Prefix Attribute Flags: X:0 R:0 N:1
  Hostname:          R1
  IPv6 Address:      2001:1:1:1::1
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:10:1:12::/64
    Prefix Attribute Flags: X:0 R:0 N:0
  Metric: 10         MT (IPv6 Unicast) IPv6 2001:10:1:13::/64
    Prefix Attribute Flags: X:0 R:0 N:0
  Metric: 0          MT (IPv6 Unicast) IPv6 2001:1:1:1:1/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
    Prefix Attribute Flags: X:0 R:0 N:1
  MT:                 Standard (IPv4 Unicast)
  MT:                 IPv6 Unicast                                0/0/0
  Metric: 10         IS-Extended R2.07
    Interface IP Address: 10.1.12.1
    Link Maximum SID Depth:
      Label Imposition: 10
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 P:0 weight:0 Adjacency-sid: 24001 System ID:R2
  Metric: 10         MT (IPv6 Unicast) IS-Extended R2.07
    Interface IPv6 Address: 2001:10:1:12::1
  Link Maximum SID Depth:
    Label Imposition: 10
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 P:0 weight:0 Adjacency-sid: 24003 System ID:R2
  Router Cap:        1.1.1.1 D:0 S:0

```

```

Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
SR Local Block: Base: 15000 Range: 1000
SR Algorithm:
  Algorithm: 0
  Algorithm: 1
Node Maximum SID Depth:
  Label Imposition: 10
RP/0/0/CPU0:R2#

```

This example from IS-IS shows the R2 is receiving SR information from R1, which has IPv4 (1.1.1.1/32) and IPv6 (2001:1:1:1::1/128) prefixes advertised with the N (Node) flag set to 1—a Node SID, not a Prefix SID.

IGP Adjacency Segment

An **IGP Adjacency Segment** is an identifier that describes a particular link between two routers. It is used to direct traffic over a specific link within the IGP routing domain. Routers can be given an instruction to forward based on the IGP adjacency. Note that adjacency segment is

- Known as Adj-SID
- Represents a hop over a specific link between two IGP-speaking routers
- Likely a one-hop path
- A Local Segment—significant only on a particular router
- Advertised as a label value
- Distributed via IGP (ISIS/OSPF)

Figure 15-4 illustrates this. Now, the packet has arrived at P5 and needs to travel further to PE6 (10.1.100.6). The operator has a choice to impose a local decision on P5 on which links to use—the direct link to PE6 or two-hop link via PE7. In this case, the link to PE7 is preferred (higher bandwidth, lower latency, encryption—take your pick), and label 24150 steers the traffic toward PE7.




Figure 15-4 IGP Prefix Segment Behavior with Adjacency Segments

Combining IGP Segments

By combining segment IDs, you can groom traffic on any path in the network:

1. Specify the sequential list of segment IDs in the packet header, known as a label stack with the top label being read first.
2. Path is *not* signaled, and per flow state is not created (as in RSVP-TE).
3. A single protocol (IS-IS, OSPF, BGP) distributes this instruction.

In Figure 15-5, a network operator instructs PE2 to steer traffic to PE6 by sending it to P5 first via two available ECMP paths. Once P5 gets the packet, the top label 16005 is removed, and P5 uses the direct link to PE6 by looking up this interface in the forwarding table where it is associated with the dynamic adjacency label 24100.




Figure 15-5 Combining IGP Segment IDs for Traffic Steering

By combining segment IDs (Prefix-SID and Adj-SID) in this way, you can put a packet on any path through the network, no matter how complex or unnatural this path may be. That is the power and essence of Segment Routing. At each hop, the top segment identifies the next hop. Segment IDs are stacked in sequential order at the top of the packet header. When the top segment ID contains the identity of another router, the receiving node uses equal cost multipaths (ECMP) to move the packet to the next hop. When the identity is the receiving router itself, the router will pop the top segment and perform the task required by the next segment.

NOTE Please note that I (Brad) am simplifying the mechanics of Segment Routing to its basic elements to facilitate a clear understanding of the fundamental concepts. In practice, multiple labels may be used, including transport labels and service labels that carry L3VPN traffic along with traffic engineering. Additionally, the number of labels a platform can handle depends on its hardware capabilities. Generally, Segment Routing can accommodate multiple labels in the label-switched path (LSP), with some platforms supporting up to 6, 9, or even 12 labels. However, most networks do not typically construct such elaborate paths, although it is possible and some customers have implemented them. Figure 15-6 shows a simple example of a label stack where PE2 assigns the inner service label 24192 for VPN traffic between CE1 and CE8. Labels are disposed along the way with PE6 associating this VPN label with the connection to CE8.

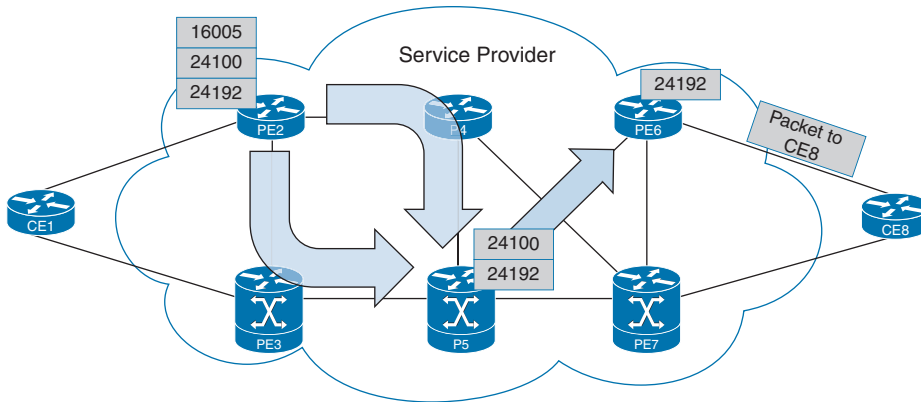


Figure 15-6 *Label Stack Example*

Segment Routing Control Plane

The control plane in Segment Routing (SR) plays a crucial role in managing how segment ID information is shared among network devices. Link-state Interior Gateway Protocol mechanisms distribute segment IDs on Segment Routing networks. Both OSPF and IS-IS include protocol extensions to support the distribution of segment IDs. These extensions enable routers to maintain a comprehensive database containing information about all nodes and adjacency segments. Because IGP is now responsible for distributing segment IDs, and labels in the case of the MPLS data plane, there's no need for a separate label distribution protocol, as mentioned earlier. Our control plane has become far simpler because it is working with only one source of truth—IGP—instead of having to reconcile both IGP and LDP information during failure events. It is important to note that the **Segment Routing control plane** can be applied to both MPLS and IPv6 data planes. In Cisco's documentation, this is referred to as SR-MPLS and SRv6, the former running on MPLS labels and the latter on IPv6 routing. Let's start by examining SR-MPLS and learn the details behind protocols that provide this unified well-organized improvement.

IS-IS Control Plane

The **IS-IS control plane** disseminates Segment Routing information within an autonomous system. Because LDP is not necessary, IS-IS will distribute both the prefixes and labels in the extensions built into IS-IS itself. This allows for seamless deployment of Segment Routing in existing MPLS networks. Rather than modifying the protocol itself, the designers "extended" its use by providing these additional protocol add-ons to carry information not originally intended by protocol designers. Think of a train of railway cars where the locomotive does not know the load being carried inside each car it is pulling. This way, new functionalities can be added to the protocol by adding new TLVs. IS-IS works exactly this way, because it understands how to transport such values for the use of Segment Routing. It uses Type-Length-Value (TLV) triplets along with sub-TLVs to encapsulate various information in its advertisements. It can support both IPv4 and IPv6 control planes and extends its reach to level-1, level-2, and multilevel routing. It is capable of providing MPLS penultimate hop popping (PHP) and explicit-null signaling as well. Several RFCs, including RFC 8667 and RFC 8402, describe the process of how Prefix-SID and Adj-SID are carried in sub-TLVs in great detail. Figure 15-7 shows the format of the Prefix-SID sub-TLV.


**Figure 15-7** IS-IS Prefix-SID Format

Table 15-3 shows the most significant TLVs you should be able to recognize on the exam.

**Key
Topic**
Table 15-3 IS-IS TLVs

TLV	Name	Description	Reference
2	IIS Neighbors	Shows all running interfaces to which IS-IS is connected, has a maximum metric of 6 with only 6 out of 8 bits used.	ISO 10589
10	Authentication	The information is used to authenticate IS-IS PDUs.	ISO 10589
22	Extended IS Reachability	Increases the maximum metric to 3 bytes (24 bits), addressing TLV 2 metric limitation.	RFC 5305
134	TE Router ID	MPLS Traffic Engineering router ID.	RFC 5305
135	Extended IP Reachability	Provides a 32-bit metric with an “up/down” bit for the route-leaking of L2 \neq L1.	RFC 5305
149	Segment Identifier/Label Binding	Advertises prefixes to SID/Label mappings. This functionality is called the Segment Routing Mapping Server (SRMS).	RFC 8867
222	MT-ISN	Allows for multiple-topology adjacencies.	RFC 5120
236	IPv6 Reachability	Describes network reachability through the specification of a routing prefix.	RFC 5308
242	IS-IS Router CAPABILITY	Allows a router to announce its capabilities within an IS-IS level or the entire routing domain.	RFC 7981

IS-IS allocates the SRGB along with the Adjacency-SIDs and advertises both in IS-IS for the enabled address-families. IS-IS enables MPLS forwarding for all non-passive interfaces. Example 15-3 shows commands necessary to turn on Segment Routing in IS-IS.

**Key
Topic**
Example 15-3 *Commands to Turn on Segment Routing in IS-IS*

```
IOS XR
RP/0/0/CPU0:PE4# show running-config router isis
router isis CCNP
  set-overload-bit on-startup 300
  is-type level-2-only
  net 49.0001.0000.0000.0004.00
  distribute link-state
  nsf ietf
```

```

log adjacency changes
lsp-gen-interval maximum-wait 10000 initial-wait 20 secondary-wait 200 level 2
lsp-refresh-interval 65000
max-lsp-lifetime 65535
address-family ipv4 unicast
  metric-style wide
  metric 100 level 2
  microloop avoidance
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
  spf-interval maximum-wait 2000 initial-wait 50 secondary-wait 200
  router-id Loopback0
  segment-routing mpls
!
address-family ipv6 unicast
  metric-style wide
  spf-interval maximum-wait 2000 initial-wait 50 secondary-wait 200
!
interface Loopback0
  passive
  address-family ipv4 unicast
    prefix-sid index 4
  !
!
mpls traffic-eng
RP/0/0/CPU0:PE4#

```

OSPFv2 Control Plane

Much like in IS-IS, OSPF does not rely on LDP to transmit prefix label information. It uses protocol extensions to distribute Segment Routing labels in the **OSPFv2 control plane**. OSPF relies on fixed-length link-state advertisements (LSAs) for its fundamental operations. Later, Opaque LSAs were introduced to expand to new protocol capabilities, accommodating features like Segment Routing and Traffic Engineering. These advertisements are flooded to OSPF neighbors opaquely, implying that even if a transit router lacks comprehension of this information (perhaps due to running older software), it will nonetheless indiscriminately transmit it to its neighboring routers. Multi-area functionality is supported, host loopback prefixes are advertised as IPv4 Prefix Segment IDs (Prefix-SIDs), and Adjacency Segment IDs (Adj-SIDs) are used for adjacencies. MPLS penultimate hop popping (PHP) and explicit-null signaling are also supported.

Note the format of Opaque LSAs in Figure 15-8.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
LS age										Options										9, 10, or 11																			
Opaque Type					Opaque ID																																		
Advertising Router																																							
LS Sequence Number																																							
LS Checksum																				Length																			
Opaque Information																																							

Figure 15-8 *Opaque LSA Format*

Opaque LSA types are identified by the topology flooding scope in Table 15-4. The most known of these (when it comes to Segment Routing) is type 10 LSAs, which distribute Traffic Engineering (TE) link attributes. It is often referred to as the *TE LSA*, yet it has other applications as well.

**Key
Topic**

Table 15-4 OSPF Opaque LSAs

LSA type	LSA Scope	Topology Flooding Scope	Reference
9	Link-local	Local network only	RFC 5250
10	Area-local	Only within an area	RFC 5250
11	Autonomous System	Domain-wide, same as AS-External type-5 LSAs	RFC 5250

Similar to IS-IS, OSPF will allocate and advertise the SRGB to its neighbors. It activates MPLS forwarding on all OSPF interfaces, excluding loopback interfaces, and assigns Adjacency-SIDs to these interfaces. Example 15-4 shows commands necessary to turn on OSPF for Segment Routing.

Example 15-4 *Commands to Turn on OSPF for Segment Routing*

```

IOS XR
RP/0/0/CPU0:PE4# show running-config router ospf
router ospf CCNP
  nsr
  distribute link-state
  log adjacency changes detail
  router-id 10.1.100.10
  segment-routing mpls
  segment-routing forwarding mpls
  fast-reroute per-prefix

```

```

fast-reroute per-prefix ti-lfa enable
affinity-map
  RED bit-position 0
!
nsf ietf
! Output omitted for brevity
address-family ipv4 unicast
area 0
  mpls traffic-eng
  segment-routing mpls
interface Loopback0
  passive enable
  prefix-sid index 10
!
interface HundredGigE0/0/0/0
  bfd minimum-interval 20
  bfd fast-detect
  bfd multiplier 3
  cost 200
  network point-to-point
!
interface HundredGigE0/0/0/1
  bfd minimum-interval 20
  bfd fast-detect
  bfd multiplier 3
  cost 200
  network point-to-point
!
!
mpls traffic-eng
RP/0/0/CPU0:PE4#

```

What about OSPFv3? While OSPFv3 has the potential to accommodate Segment Routing for IPv6 and utilize a native IPv6 data plane, specific extensions outlined in an IETF draft are required for implementation. It's worth noting that, at press time, these extensions have not been integrated into Cisco IOS XR and IOS XE.

BGP Control Plane

BGP also has the capability to function as the control plane for Segment Routing (SR), enabling prefix distribution throughout the network. In the context of Segment Routing, the **BGP control plane** distributes segment routing information between routers, enabling them to make forwarding decisions based on predefined segments. While seen less frequently than IS-IS and OSPF, BGP has been effectively used in practice in multiple large-scale web data centers. Such data centers can support over 100,000 services, profoundly influencing and challenging the scalability and operational efficiency of the underlying network

architectures. To meet the demands of high-intensity east-west traffic found in these compute clusters, operators frequently opt for variations of Clos or Fat-tree topologies. In these massive data center networks, symmetrical topologies with numerous parallel paths connecting two server-attachment points are common. It is in this context that BGP excels and arguably surpasses the IGP approach. The assertion that “BGP is a better IGP” challenges traditional viewpoints and has sparked conversations.

What would make BGP an attractive choice? Remember that these massive data centers seek maximum bandwidth to be transferred across the midpoint of the system. Such network structures are designed to be both highly scalable, cost-effective, and are constructed from affordable, low-end access-level switches. To maintain this level of scale, the designs call for a single protocol with simple behavior and wide vendor support. With the above in mind, when it comes to simplicity, BGP certainly has its advantages because it has less of a state machine and fewer data structures. This may not appear intuitive at first glance, but it does not take long to realize that the BGP RIB structure is simpler than those of Link-State Databases (LSDBs). There is a very clear picture of “which routing information is sent where.” There is a RIB-In and RIB-Out, a far easier construct for tracing exact routing paths than following link-state topology constraints with areas and levels. When it comes to operational troubleshooting, this is definitely a strength. Also, event propagation is more constrained in BGP because link failures have limited propagation scope. We can argue that BGP has more stability due to the reduced “event-flooding” domains. When it comes to traffic steering, BGP allows for per-hop Traffic Engineering that can be used for unequal cost Anycast load-balancing. In addition, BGP is widely supported by practically all vendors, so from the perspective of interoperability, BGP beats IGPs. We have been conditioned to perceive BGP as slow and suitable primarily for inter-domain routing, but it has no issues with demonstrating its adaptability and effectiveness in modern topologies. Therefore, it is advisable to approach BGP with an open mind, recognizing its potential to perform as well as, or even better than, traditional IGP alternatives in contemporary implementations.

BGP will advertise a BGP Prefix-SID associated with a prefix via BGP Labeled Unicast (BGP-LU) IPv4/IPv6 Labeled Unicast address-families. BGP Prefix-SID is a global SID, and the instruction forwards the packet over the ECMP-aware BGP best path to the associated prefix. RFC 8277 specifies that Label-Index TLV must be present in the BGP Prefix-SID attribute attached to IPv4/IPv6 Labeled Unicast prefixes. This 32-bit value represents the index value in the SRGB space and has the format illustrated in Figure 15-9.




Figure 15-9 BGP Prefix SID Advertised Format

The Prefix-SID for a locally originated BGP route can be set with a route-policy. Example 15-5 shows how to attach a label-index with **network** and **redistribute** commands.

Example 15-5 *Attaching a Label-Index via a Route-Policy*

IOS XR configuration with network command

```
route-policy SIDs($SID)
  set label-index $SID
end-policy
!
router bgp 100
  address-family ipv4 unicast
    network 10.1.100/4/32 route-policy SID(1)
    allocate-label all
```

IOS XR configuration with redistribute command

```
route-policy SIDs
  if destination in (10.1.100.4/32) then
    set label-index 1
  endif end-policy
!
router bgp 100
  address-family ipv4 unicast
    redistribute connected route-policy SIDs
    allocate-label all
```

One last thing regarding having BGP for a Segment Routing control plane. Remember the Anycast load-balancing I mentioned earlier in this section? Anycast allows different nodes to advertise the same BGP prefix. It is an application of Prefix SIDs to achieve anycast operations. Look at Figure 15-10, where I again moved some links around to represent a data center's spine-and-leaf architectures, with spines located at the top. PE2 and P4, while advertising their individual BGP Prefix-SIDs (16002 and 16004, respectively), have been made members of the same unicast set. Both of them advertise anycast prefix 10.1.100.24/32 with BGP-Anycast SID 20001. PE3 wants to send traffic to PE7 but would like to exclude spine PE6. BGP-Anycast SID 20001 will load-balance the traffic to any member of the Anycast set and then forward it to PE7.

Additionally, due to BGP Prefix-SID global label usage, BGP-LU local labels are going to be the same across all of the network's ASBRs. As a result, these Anycast loopbacks can be used as the next-hop for BGP-LU prefixes. That is pretty good resiliency! Nothing to scoff at, for sure.




Figure 15-10 BGP-SR Anycast Load-Balancing

SRv6 Control Plane

The **SRv6 control plane** manages the signaling, routing, and forwarding information for Segment Routing over IPv6 (SRv6) networks. It serves as the Segment Routing architecture tailored for the IPv6 data plane and extends to the value of IPv6, influencing future IP infrastructure deployments, whether in data centers, large-scale aggregation, or backbone networks. SRv6 functions as an extension of the Segment Routing architecture specifically designed for IPv6 networks. It introduces a source-routing mechanism by encoding instructions within the IPv6 packet header.

The use of IPv6 addresses to identify objects, content, or functions applied to objects opens up significant possibilities, particularly in the realm of chaining microservices within distributed architectures or optimizing content networking. Notably, stable networks, particularly in the Asia-Pacific region, have embraced SRv6, boasting tens of thousands of nodes on a single network as of the time of writing this book.

Fundamentally, SRv6 encodes topological and services paths into the packet header. The SRv6 domain does not hold any per-flow state for Traffic Engineering or network function virtualization (NFV). Sub-50ms path protection is delivered with TI-LFA. It natively delivers all services in the packet header, without any shims or overlays. IPv4's limitations have forced the industry to create extra tools to deal with its challenges. When IPv4 lacked sufficient address space, NAT was created to hide and conserve addresses. For engineered load-balancing, we have had to invent MPLS Entropy Label and VxLAN UDP. For separating discrete networks, MPLS VPNs along VxLAN were created. Since Traffic Engineering functions were missing in IPv4, RSVP-TE and SR-TE MPLS appeared. Network Service Header (NSH) overcame IPv4 service chaining limitations. All of the above is done natively in IPv6 and why so many service providers are turning to this technology.

SRv6 (Segment Routing over IPv6) Header

At the heart of SRv6 is the IPv6 Segment Routing Header (SRH). Figure 15-11 shows the IPv6 SRH replicated from RFC 8754. This header is added to IPv6 packets to implement Segment Routing on the IPv6 forwarding plane. SRH specifies an IPv6 explicit path, listing one or more intermediate nodes the packet should visit on the way to its final destination. The Segment Left field provides the number of transit nodes before traffic reaches its

destination. Then, the Segment List fields indicate the sequence of nodes in 128-bit IPv6 addresses to be visited from bottom to top. Segment List [n] shows the first node in the path; Segment List [0] shows the last node in the path.

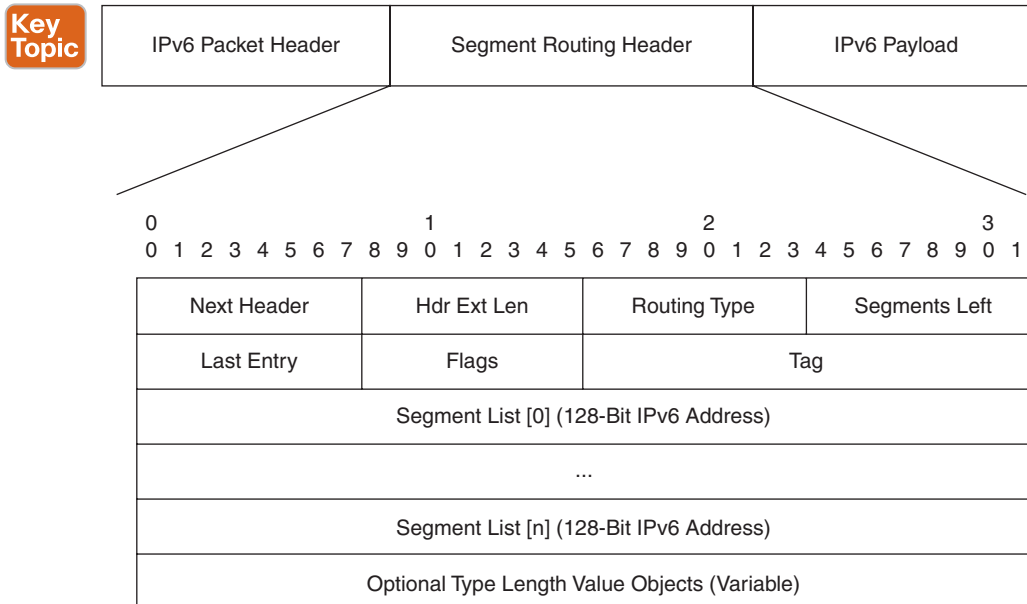


Figure 15-11 IPv6 Segment Routing Header Format

In SRv6, each segment is represented by an IPv6 address known as a segment identifier (SID). These SIDs play a crucial role in defining specific paths or instructions for forwarding packets throughout the network. It looks a lot like a 128-bit IPv6 address, but has different semantics because it consists of two parts, with Figure 15-12 providing the visualization:

- **Locator:** Represents an address of a specific SRv6 node performing the function.
- **Function:** Represents any possible network instruction bound to the node that generates the SRv6 SID (network instruction) and is executed locally on that particular node, specified by the locator bits.




Figure 15-12 IPv6 Segment Identifier

You now have the ability to send packets to a node (locator) and then instruct the node to execute an action (function). This is not a subtle difference! In SR-MPLS, IGP with extensions advertised the transport mechanism, and services (L2VPNs, L3VPNs) were signaled *independently* via LDP or MP-BGP. You could change your transport (from MPLS to SR) without affecting the upper protocols that ran on top of it. For the first time in the industry, transport and services instructions are coupled and signaled in the SID. You will see an example of this coming up shortly where an L3VPN is written into the SID.

SRv6 Node Roles

In the context of SRv6 (Segment Routing over IPv6) networks, different nodes play distinct roles in facilitating packet forwarding and processing. These roles include

- **Source Node:** This node has the capability to generate an IPv6 packet incorporating a Segment Routing Header (SRH), essentially forming an SRv6 packet. Alternatively, it serves as an ingress node that can apply an SRH to an existing IPv6 packet.
- **Transit Node:** Found along the SRv6 packet's path, the transit node functions without inspecting the SRH. The destination address of the IPv6 packet does not align with the transit node, and its role is primarily to forward the packet.
- **Endpoint Node:** Located within the SRv6 domain, this node acts as the termination point for the SRv6 segment. The destination address of the IPv6 packet containing an SRH corresponds to the endpoint node. The endpoint node executes the specific function associated with the SID bound to the segment.

SRv6 Micro-Segment SID (uSID)

Often referred to as *Micro-SID* or *Compressed SID*, the uSID feature is an extension of the SRv6 architecture. In SRv6, the micro segment identifier, or uSID, is a specialized form of Segment Routing where packets are marked with a compact identifier for precise forwarding. Unlike traditional SRv6, which might use longer segment identifiers for various purposes, uSID is specifically designed for efficient and granular traffic steering. It provides a more streamlined approach to segment routing, particularly useful for scenarios requiring fine-grained control and scalability enhancements.

Using the established SRv6 Network Programming framework, it can encode up to six SRv6 Micro-SID (uSID) instructions within a singular 128-bit SID address, termed a *uSID Carrier*. Moreover, this extension seamlessly integrates with the existing SRv6 data plane and control plane, requiring no modifications. Notably, it ensures minimal MTU overhead. For instance, when incorporating six uSIDs within a uSID carrier, it yields 18 source-routing waypoints with just 40 bytes of overhead in the Segment Routing Header. Look at Figure 15-13, which illustrates the usage of uSID. Pay attention to how the highlighted uSIDs correspond to router numbering/naming.

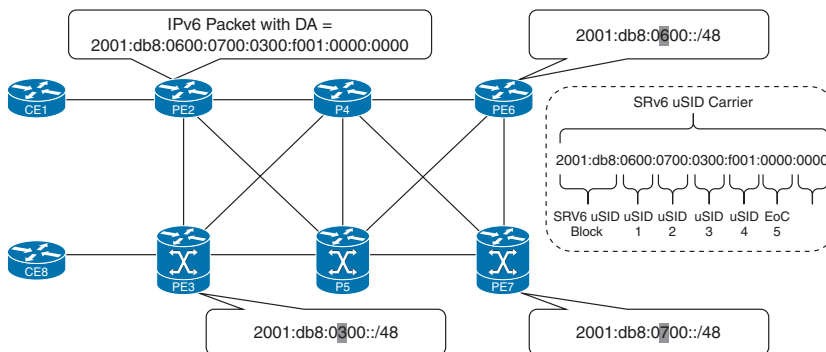


Figure 15-13 *uSID in Action*

The customer at CE1 is using the VPNv4 SP service to connect to a remote site CE8. Router PE2 sends traffic to VPNv4 CE8 to router PE3 via a traffic-engineered path visiting routers

PE6 and PE7 using a single (!) SRv6 SID (note that without uSID, a sequential Segment List would have to be specified). Let's unpack this:

1. PE2, PE6, PE7, and PE3 are SRv6 capable and are configured with 32-bit SRv6 block 2001:db8.
2. P4 and P5 run classic IPv6 forwarding and do not change the Destination Address.
3. PE6, PE7, and PE3 advertise their corresponding 2001:db8:0600::/48, 2001:db8:0700::/48, and 2001:db8:0300::/48 routes.
4. PE2 receives an IPv4 packet from CE1, encapsulates it, and sends an IPv6 packet with the destination address 2001:db8:0600:0700:0300:f001:0000:0000. This is an SRv6 uSID Carrier that contains a sequence of micro-SIDs (instructions 0600, 0700, 0300, f001, and 0000).
5. The 0600, 0700, and 0300 uSIDs are used to construct a traffic engineering path to PE3 with two stops along the way—PE6 and PE7. uSID f001 is a BGP-signaled instruction sent by PE3 indicating the VPNv4 service. uSID 0000 indicates the end of instructions.
6. What happens at P4? P4, running only classic IPv6, forwards the packet along the shortest path to PE6.
7. PE6 receives the packet, pops its own uSID 0600, and advances the micro-program by looking up the shortest path to the next Destination Address (DA) 2001:db8:0700::/48. Now the DA is 2001:db8:0700:0300:f001:0000:0000:0000. This behavior is called *shift and forward*.
8. PE7 receives the packet, pops its own uSID 0700, and advances the micro-program by looking up the shortest path to the next Destination Address (DA) 2001:db8:0300::/48. Now the DA is 2001:db8:0300:f001:0000:0000:0000:0000. Shift and forward again.
9. P5 forwards the packet to PE3, just like P4 did.
10. PE2 receives the packet and executes the VPNv4 function based on this own instruction f001. It decapsulates the IPv6 packet, performs IPv4 table lookup, and forwards the IPv4 packet to CE8.

SRv6/MPLS L3 Service Interworking Gateway

The SRv6/MPLS L3 Service Interworking Gateway facilitates the seamless extension of L3 services between MPLS and SRv6 domains, ensuring continuity in service delivery across both control and data planes. This feature enables interoperability between SRv6 L3VPN and existing MPLS L3VPN domains, offering a pathway for transitioning from MPLS to SRv6 L3VPN.

At the gateway node, the SRv6/MPLS L3 Service Interworking Gateway performs both transport and service termination tasks. It generates SRv6 VPN SIDs and MPLS VPN labels for all prefixes within the configured VRF for re-origination, as illustrated in Figure 15-14. The gateway supports traffic forwarding from the MPLS domain to the SRv6 domain by removing the MPLS VPN label, performing a destination prefix lookup, and applying the appropriate SRv6 encapsulation. Conversely, for traffic from the SRv6 domain to the MPLS domain, the gateway removes the outer IPv6 header, performs a destination prefix lookup, and applies the VPN and next-hop MPLS labels.

PE3 is the interworking gateway that has one leg in the SR-MPLS domain and the other in the SRv6 domain. It performs a translation service by popping the MPLS VPN label and looking up the destination prefix in the SRv6 domain. It encapsulates the payload in the

outer IPv6 header with P4's destination address. In the opposite direction, PE3 removes the outer IPv6 header, looks up the destination prefix, and pushes MPLS label 16002 for the BGP next-hop of PE2.

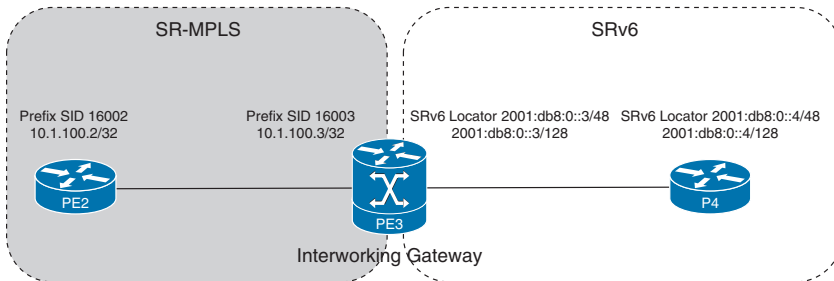


Figure 15-14 SR-MPLS SRv6 Interworking Gateway

Co-existence with LDP

It would be nice to never worry about LDP and RSVP, but the reality is that many of today's engineers will have to touch these older MPLS networks. A Segment Routing control plane can co-exist with the label-switched paths (LSPs) constructed with LDP or RSVP. The MPLS architecture allows for the simultaneous use of multiple label distribution protocols, including LDP, RSVP-TE, and others. The SR control plane can coexist alongside these protocols without any interaction. In Figure 15-15, we have removed some links in our network and have thus completely flattened it. This network runs a mix of both Segment Routing (SR) and Label Distribution Protocol (LDP). It is possible to establish an end-to-end seamless Multiprotocol Label Switching (MPLS) LSP, which will ensure interoperability between these two domains. To accomplish this, one or more nodes function as Segment Routing Mapping Servers (SRMS). These SRMS entities take on the responsibility of advertising SID mappings on behalf of nodes that are not SR-capable. This mechanism enables SR-capable nodes to learn about the SIDs assigned to non-SR-capable nodes without the need for explicit individual node configurations. Let's unpack this.




Figure 15-15 SR and LDP Domain Interoperability

Notice that this network runs both SR and LDP, which can be typical during network transitions and upgrades. PE2 and PE7 are exchanging BGP VPNv4 routes. PE2, PE3, and PE4 are SR-capable. PE4, PE5, PE6, and PE7 use LDP. How do these two domains talk to each other end to end? First, let's start from the LDP→SR direction, which is quite easy because SR-capable routers will automatically translate between LDP- and SR-based labels:

1. PE7 learns a service route (L3VPN route, for example) for customer prefix 172.16.1.0/24 with a service/VPN label of 30001.
2. PE7's BGP next-hop for this service label is associated with PE2's lo1 10.1.100.2/32.

3. PE7 finds LDP label binding 24016 from its neighbor PE6 for PE2's Forwarding Equivalence Class (FEC) of 10.10.100.2/32 and forwards the packet to PE6.
4. PE6 finds LDP label binding 24020 from its neighbor PE5 for PE2's FEC of 10.10.100.2/32, swaps 24016 for 24020, and forwards the packet to PE5.
5. PE5 finds LDP label binding 24036 from its neighbor PE4 for PE2's FEC of 10.10.100.2/32, swaps 24020 for 24046, and forwards the packet to P4.
6. P4 lacks an LDP binding originating from its next-hop PE3 for the FEC associated with PE1. What it does carry, though, is an SR node segment pointing to an IGP route leading to PE2. P4 engages in label merging, wherein it replaces its local LDP label (24036) for FEC PE2 with the corresponding SR node segment label, which is 16002.
7. PE3 pops label 16002 (assuming penultimate hop popping function is used) and forwards the packet to PE2.
8. PE2 receives the packet, looks up its service label of 30001, and drops the packet into the appropriate customer VRF.

We now have an end-to-end LDP⇄SR path. Simple. What about in the opposite direction? This is where we will encounter a problem going from SR⇄LDP. Can you take a moment to think what the problem would be by looking at Figure 15-14 before you examine Figure 15-15? PE2 needs to send traffic to 172.16.2.0/24 with service label 40001 that it received with the BGP next-hop of 10.1.100.7/32. Since PE2 only speaks SR, when it looks up the node segment for 10.1.100.7/32, what will it find in its label database? Nothing. Why? Because such label mapping does not exist on the network, since the operator has never configured it; therefore, no router advertises or receives this label mapping. There must be something to associate PE7's loopback with SR label mapping. The better answer here is Segment Routing Mapping Server (SRMS). All analogies finally break down, but it is possible to think of SRMS as a sort of route reflector for SR labels. Just as in BGP, we can centrally instruct all routers in our SR domain. Look at Figure 15-16.

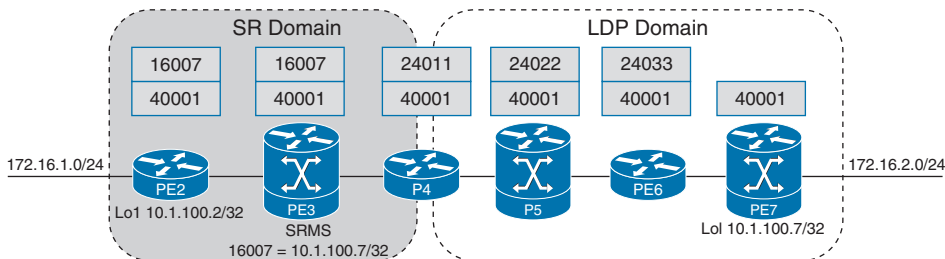


Figure 15-16 SR and LDP Domain Interoperability with SRMS

Walking back in the opposite direction looks like this:

1. PE3 is chosen as a Segment Routing Mapping Server (SRMS). In practice, it is recommended to have a redundant SRMS.
2. As PE7 lacks Segment Routing (SR) capability, you must create a mapping policy on the SRMS, which associates label 16007 with PE7's lo1 10.1.100.7/32.
3. Now, PE2 learns a service route (L3VPN route via BGP) for customer prefix 172.16.1.0/24 with a service/VPN label of 40001 with the BGP next-hop of 10.1.100.7/32.

4. PE2 finds an SR label binding 16007 it has received from the SRMS PE3 for PE7's FEC of 10.10.100.7/32 and forwards the packet to PE3 as the IGP next-hop.
5. PE3 finds an SR label binding 16007 pointing to its neighbor P4 as the IGP next-hop, swaps 16007 for 24011, and forwards to P4.
6. P4 does not have an SR label for PE7's IGP route, but it holds LDP label 24011 for this FEC. It swaps 16007 for 24011 (remember the process is called *label merge*) and forwards to P5.
7. P5 swaps 24011 for 24022 and forwards to PE6.
8. PE6 pops the label (due to PHP in this setup) and forwards to PE7.
9. PE7 receives the packet, looks up its service label of 40001, and drops the packet into the appropriate customer VRF.

What should you remember here? Segment Routing Mapping Server labels are only necessary in the SR→LDP direction. SR and LDP labels come from separate label database ranges (16000–23999 for SR and 24000+ for LDP), so unless the operator has deliberately violated this guidance, there is not a chance the network will be in a state of confusion, since the SR and LDP labels do not overlap each other. The network must maintain continuous SR connectivity in the SR domain. The network must also maintain continuous LDP connectivity in the LDP domain. If you understood the packet walkthrough, these points should be clear.

One last thing to know for completeness. By default, Cisco routers prefer LDP as the label imposition mechanism when the MPLS features are turned on. The way to enable SR for label imposition is shown in Example 15-6.

Example 15-6 *Segment Routing Label Imposition Preferred*

```
IS-IS

router isis 100
  address-family ipv4|6 unicast
    segment-routing mpls sr-prefer

OSPF

router ospf 100
  segment-routing mpls
  segment-routing sr-prefer
```

Segment Routing Traffic Engineering

RSVP-TE, despite its powerful Traffic Engineering capabilities, poses challenges in practical deployments due to its complexity. Managing backup tunnels, intricate configurations at scale, the absence of seamless inter-domain intelligence, and the complexities of steering traffic through methods like PBR or autoroute have resulted in various issues and limited widespread adoption. Simplifying these aspects is crucial for enhancing the usability and deployment scope of Traffic-Engineered networks. Enter Segment Routing policies. They are simple, automated, scalable, and carry support for a wide variety of functionalities including multidomain intelligence, which is provided by Path Computation Element (PCE) and Binding-SID (BSID)—more on this later.

Segment Routing Policies

In Segment Routing, there are no tunnels (the closest possible thing is Circuit-Style Segment Routing, which has policies to put traffic on the same A–Z path, akin to bidirectional co-routed LSPs—this is outside of the current exam’s scope). Instead, Segment Routing introduces the concept of Segment Routing Policies. These are typically deployed at ingress routers at the edge of the network and can force the packet to follow any desired path.

Key Topic

An SR Policy is fundamentally a sequence of segments. In its most basic structure, it is a sequence of IP waypoints presented in either SR-MPLS or SRv6 format (SID list), with the initial entry as the first destination to be visited. An SR Policy is uniquely identified by these attributes:

- **Headend:** An ingress router where the policy is implemented.
- **Tailend:** An egress router where the policy ends.
- **Color:** A numeric value that uniquely identifies multiple SR Traffic Engineering policies between the same pair of routers.

Figure 15-17 illustrates this best. PE2 needs to send traffic for prefixes 172.16.100.0/24 and 172.16.200.0/24 to the same PE7 router, since it is the egress point connecting these two networks. However, traffic destined for 172.16.100.0/24 must follow the top low-delay path due to latency requirements, and traffic for 172.16.200.0/24 must take the bottom low-cost path because the customer is not paying for the premium service. Try doing this with IGP alone! SR policies, on the other hand, easily differentiate traffic between the same pair of routers by steering them into differently colored policies (different numeric values) that properly groom traffic onto the desired paths.

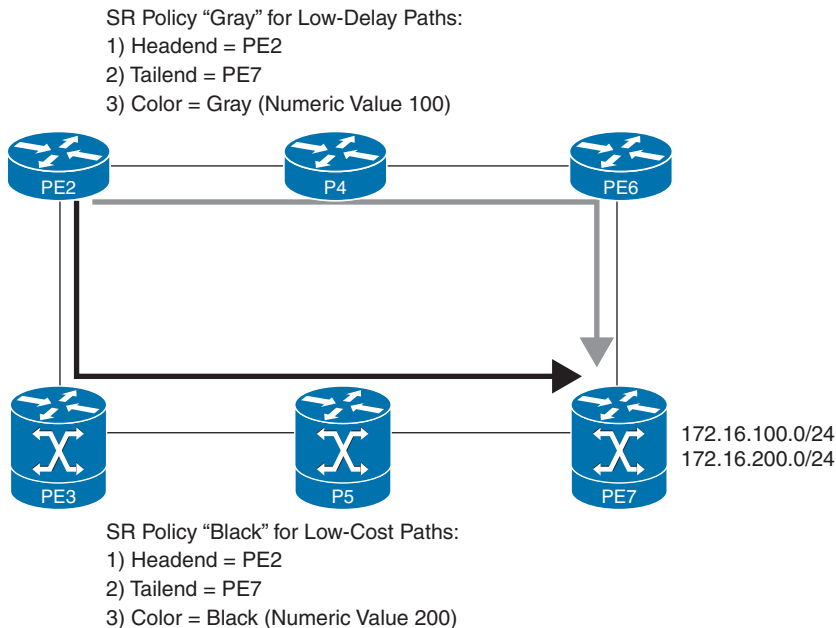


Figure 15-17 Segment Routing Policy Places Traffic on Diverse Paths

SR Policies and Candidate Paths

An SR Policy consists of one or more **Candidate Paths**. Each Candidate Path has a single SID-list or a set of weighted SID-list. Things to consider regarding a Candidate Path (the order is not important here):

1. Can be explicitly defined. The operator will provide the exact sequence of SIDs to be visited along the way to the destination.
2. Can be dynamically defined. The operator will provide optimization objectives (select only encrypted links) and constraints, a set of rules to follow (exclude links certain attributes, such as not meeting minimum delay).
3. Has a preference value (numeric, higher is preferred).
4. Is associated with a single Binding-SID (BSID, more on this later).
5. Can be supplied to the headend via
 - a. CLI
 - b. NETCONF
 - c. PCEP (Path Computation Element Protocol)
 - d. BGP
6. An SR Policy will select a single best Candidate Path and program it via BSID into the router's RIB/FIB forwarding table.

Key Topic

Binding-SID (BSID)

Binding Segment Identifier (BSID) is a SID value that is an opaque representation of a Segment Routing Policy. BSID shows a chosen path to upstream routers. It provides isolation and decoupling between distinct source-routed domains while increasing overall network scalability. Do not forget that SR Policies use BSID to program a router's forwarding table (just mentioned in point 6).

Note how in Figure 15-18 different routing/SR domains are involved. The list of SIDs to steer traffic onto the imagined low-delay path between DC1 and DC5 (DC2, PE2, P4, P4-ADJ-SID, PE7, DC5) can be long. A single BSID can represent the entire Segment Routing Policy sending it through the WAN Core domain, requiring only three SIDs (DC Primary, WAN Core SR Policy, DC Secondary). This reduces the number of segments imposed by the source.




Figure 15-18 Multidomain Use of Binding-SID

Additionally, this approach keeps one domain unaffected by routing changes in another domain, since BSID does not change during these events. Domain internal operations can be thus hidden (opaque) from each other, which can be beneficial to service providers who do not want to disclose the details of how they provide services to their customers.

Flex-Algo

Flex-Algo is the best way to do traffic engineering today. Flex-Algo, short for Flexible Algorithm, enhances Segment Routing Traffic Engineering (SRTE) by introducing additional segments with distinct properties compared to the Interior Gateway Protocol Prefix segments. It expands the SRTE capabilities by including customizable, user-defined segments in the toolbox. It can also use Segment Routing on-demand next hop (ODN) and Automated Steering to create traffic-engineered paths based on user intentions; these are outside of the scope of this book.

IETF has standardized algorithms 0 through 127. Routers run the default algorithm 0 as the IGP shortest path derived from the IGP metric. Additional algorithms 128 through 255 can be customized by network operators. They are known as *SR IGP Flexible Algorithms*, or Flex-Algo as the shorter version. It is called *flexible* because you can decide which metric you want to use in your intent.

In our earlier discussion of Prefix-SID in this chapter, we focused solely on explaining the default aspect of Prefix-SID behavior, specifically the one linked to algorithm 0. When you read (you really should) RFC 8867 and RFC 8665, you will notice that both IS-IS and OSPF include Prefix-SID sub-TLV algorithm field in the formats illustrated in Figure 15-19 and Figure 15-20.




Figure 15-19 Algorithm Field in Prefix-SID Sub-TLV for IS-IS




Figure 15-20 Algorithm Field in Prefix-SID Sub-TLV for OSPF

This means that the operator can change the default algorithm 0 (IGP shortest-path) behavior on routers that are assigned to use a different algorithm (128–255) as different constraints (logical rules) will be imposed on the part of the network that participates in this algorithm. Let's look at Figure 15-21, where we break from the familiar-to-us topology.

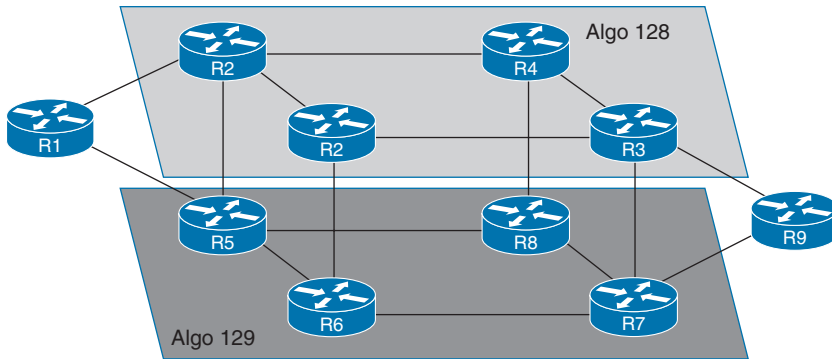


Figure 15-21 Flex- Algo Network Slicing

Suppose the operator wanted to impose different type behaviors on this network. Instead of using the default algo 0 (IGP shortest-path), the operator can define other algorithms that can minimize metrics other than shortest path, such as delay, for example. The operator can also combine this with rules to exclude links with certain properties (link-affinity, SRLG, encryption, etc.). Here is one example of what can be done:

1. The operator can define Flex- Algo 128 to prioritize IGP metric and avoid link-affinity “dark-gray” on the bottom.
2. The operator can define Flex- Algo 129 to delay metric and avoid link-affinity “light-gray” on the top.
3. Routers R1 and R9 would be added to participate in algo 0, 128, and 129.
4. Routers R1, R2, R3, and R4 would be added to participate in algo 0, 128.
5. Routers R5, R6, R7, and R8 would be added to participate in algo 0, 129.

Consider how powerful the network has become. The operator has bisected the network into two distinct profiles. The top part has routes based on the shortest path according to the IGP. The bottom half will route delay-sensitive traffic through the part of the network that uses dynamic link-delay measurement, which will be advertised by the IGP. Someone reroutes your optical underlay path? Does not matter. Someone moves a circuit without bothering to notify your department? Does not matter. The network will recalculate the best path according to the intent you had in mind. If you see the beauty of this approach, you will understand that the limitation of what can be done on a network at scale exists only in the minds of its architects. Low-cost delay-optimized paths (my personal SP operator nirvana, because this is where I make money) have now become a reality because we can now finally differentiate services on our infrastructure. It is my opinion that while SR policies are effective for carving dynamic paths on the network, the simplicity and flexibility of Flex- Algo allow the operator to easily slice the network into multiple planes that can be used to carry encrypted, application-dependent, dynamic delay-based, low-cost, and other intent-based traffic. You can even drain all network traffic to the bottom dark-gray plane of the network, run upgrades on the top light-gray plane of the network, and repeat the process again in the other direction—accomplishing zero downtime for your customers.

That is the power of Flex- Algo—operational simplicity and scale. You can finally manage massive networks with a simple picture in mind, rather than constructing hundreds of SR Policies for individual applications. Flex- Algo is applicable to SR-MPLS and SRv6. In the

case of SR-MPLS, you will get an extra label for the router's loopback. In SRv6, you get a different locator (recall our earlier discussion on this topic). SP operators are really beginning to heavily use this approach with SRv6 (IPv6).

Key Topic

Something you need to be aware of that is directly called out in the exam blueprint is *encapsulation*. SR-MPLS uses SRTE policies (on-demand or manual) to steer traffic into Flex-Algo. If you want traffic to follow a particular path, you specify a list of SIDs. When you specify the SID associated with Flex-Algo, the traffic takes that specific Flex-Algo plane. In contrast, SRv6 does not use policies. In SRv6, the ingress PE will directly encapsulate traffic based on the Service SID advertised in BGP. That Service SID (remember locator + function?) is a combination for the Algo locator and decapsulation function. Transport and Service become blended and are encoded into the transport intent (Algo locator). Transport intent and Service function are *encapsulated* into the same instruction. SRv6 is a much simpler approach to driving traffic intent.

TI-LFA

To date, TI-LFA is the number one reason why network operators deploy SR: they want an automated way to compute a backup path by IGP. No need to do MPLS-TE (traffic engineering tunnels) for fast reroute (FRR). **Topology-independent loop-free alternate (TI-LFA)** provides a simple, automatic, optimal, and topology-independent sub-50ms per-prefix protection to the network. It can protect Segment Routing, LDP, and IP traffic without relying on the construction of backup tunnels of any sort, as is the case in RSVP-TE. Whether IS-IS or OSPF is used, these protocols precompute a backup path for each active path per IP prefix destination. They run an SPF algorithm for the primary path and then automatically run the SPF again, *excluding* the primary path—deriving the backup path. IGP pre-installs this path in the data plane and immediately uses it once the active destination path is impacted. Be careful with analogies, because they all finally breakdown, but it can be helpful to think of how an EIGRP-feasible successor works. The router already knows what the post-convergence path will look like even before the failure occurs.

Figure 15-22 shows a fundamental TI-LFA operation from router PE2's perspective; once the protected PE4–P2 link fails, traffic is rerouted over the post-convergence path, which is known and preprogrammed before the link failure occurs. The recommendation is to enable this functionality on all routers in your Segment Routing domain. This approach creates automatic backup paths throughout the network without the burden of manually provisioning backup tunnel paths.

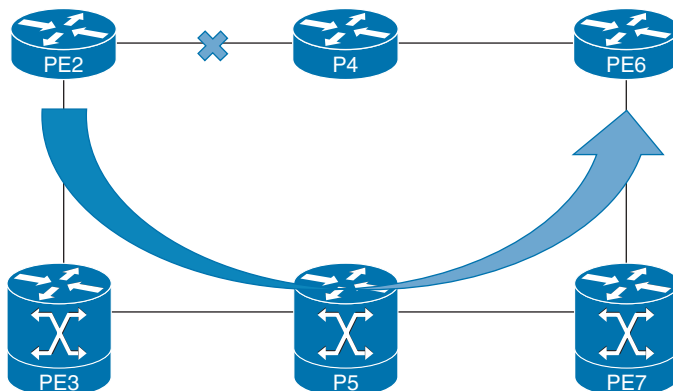


Figure 15-22 TI-LFA Operation

**Key
Topic**
Terms from Remote LFA Technology

RFC 7490 describes the following architectural reference areas to understand repair tunnel endpoints for link protection. While there is no concept of tunnels in Segment Routing (they have been replaced by policies), the same reference areas apply and are important for this exam.

P-Space

In Figure 15-23, we return to our topology and remove some of the internal links to create a ring topology to better understand these reference areas.

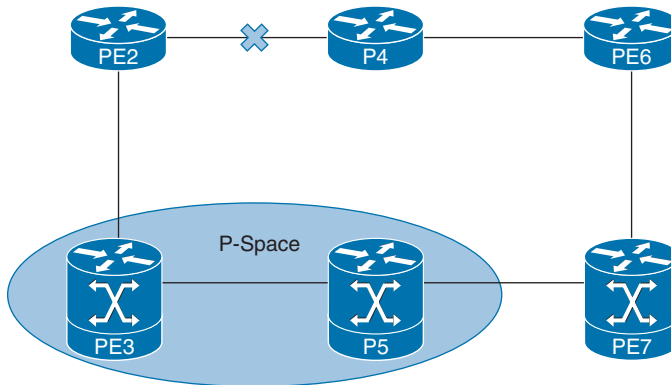


Figure 15-23 *P-Space Reference Area*

Reference areas are always seen from a perspective of a certain router with respect to a particular failed link. The way to look at reference areas depends on which router we're considering and the specific link it is protecting. In this example, router PE2 would like to protect the link between itself and router P4. The protected space (**P-Space**) of a router concerning a protected link refers to routers that PE2 can reach through the shortest paths without having to use the protected link. Which routers would those be? All link costs being equal here, only PE3 and P5 will be in P-Space. What about PE7? Not quite, as it is possible, due to ECMP, that PE2 can send a packet to PE7 through the top of the diagram through the protected link, thus disqualifying from being the shortest path. What would be the point of using a link that can potentially fail? Expressed in cost terms, P-Space contains a set of routers found on a shorter path than the path cost going through the protected link. In the case of PE7, it is equal and not shortest—thus, not a part of P-Space.

Q-Space

Q-space refers to a set of backup paths or alternate next hops that are precomputed for use during a failure. Figure 15-24 shows the other side of the protected PE2–P4 link from router P4's perspective, and the same rules apply again. When following the same rules, the set of routers reachable from P4 via the shortest path without possibly going through the protected link only include PE6 and PE7.

PQ Node

Viable repair tunnel endpoints are found at intersections of P- and Q-Spaces. In Figure 15-25, there is no common node that belongs to both reference areas and hence no viable repair tunnel endpoint is present.




Figure 15-24 Q-Space Reference Area




Figure 15-25 P-Space and Q-Space Reference Area

Extended P-Space

Because PE4 needs to repair the protected PE4–P2 link and reach any router in this ring topology without using the protected link, the concept of **Extended P-Space** was introduced. Extended P-Space is the union of each of PE4’s neighbors. In this case, this is router PE3 in Figure 15-26, whose P-Space contains routers P5 and PE7. By combining P-Spaces of PE2 and PE3, we extend PE2’s reach, and PE7 becomes a common point for P- and Q-Spaces. A **PQ node** of a node PE2, in relation to a protected link PE2–P4, is a node that belongs to both the P-space (or extended P-space) of PE2 for that link and the Q-space of P4 for the same link. PE7 is chosen as the repair tunnel endpoint. Why? Because repair tunnels are chosen from a set of PQ nodes.

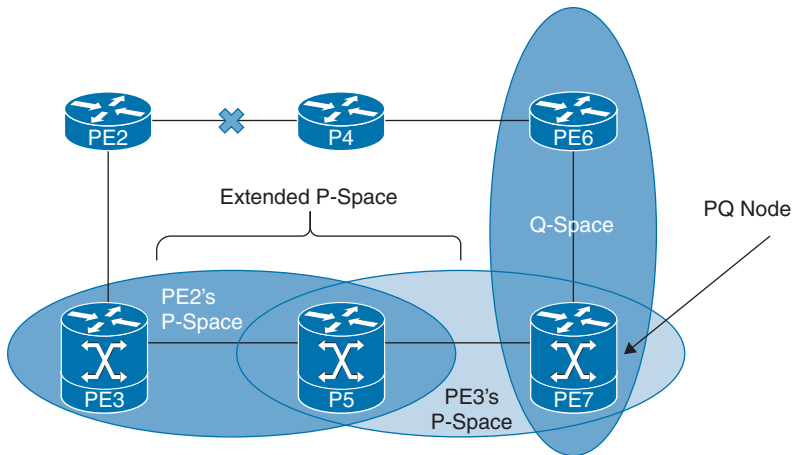


Figure 15-26 PQ Node and Extended P-Space

Classic LFA Limitations

Now, with the understanding of the reference points, Classic LFA's (loop-free alternate fast reroute, aka LFA-FRR) limitations become obvious. Note that I am not discussing LFA-FRR because it is not a part of the exam blueprint. I reference it here to highlight the advantages of TI-LFA. Figure 15-27 considers two such limitations.




Figure 15-27 Classic LFA Limitations Examples

First, LFA-FRR suffers from incomplete coverage, which makes it topology dependent (as opposed to TI-LFA, which is topology independent). Recall our discussion about the PQ node. PE2 protects the PE2–P4 link and sends traffic to PE6. When the PE2–P4 link fails, PE2 will send traffic to PE3. Before the network converges via IGP, PE3 has a problem, since the shortest path to PE8 is still through the failed PE2–P4 link and PE3 will send the traffic

back to PE2, looping the doomed packets. This is a real problem that, in the rLFA (Remote LFA) cases, can sometimes be solved by a Targeted LDP session, where PE2 would establish a remote LDP session with PE7, but this approach also has limitations that are outside of the scope of this exam. TI-LFA handles this topology though a “double-segment” coverage, where two labels are pushed (PE3, PE3-R5 ADJ-SID) to overcome this problem.

Second, notice the additional P9 router. Let’s suppose it is not a part of the network core or planned for capacity. Classic LFA will steer the traffic on this suboptimal backup path. Additional case-specific operator involvement would be necessary to avoid such undesired backup paths.

In contrast, a topology-independent loop free alternate (TI-LFA) provides 100 percent coverage and uses the post-convergence path as the fast reroute (FRR) backup path.

Key Topic

TI-LFA delivers significant improvements over the traditional loop-free alternate fast reroute (LFA-FRR) approach. TI-LFA uses a post-convergence path after a link failure occurs. This path is known before a failure occurs and is preprogrammed into the data plane. TI-LFA uses PQ nodes, or a combination of P and Q nodes located on the post-convergence path to compute backup paths. Traffic will be rerouted in sub-50ms on any topology.

While the blueprint does not focus on configuration of Segment Routing features, you need to know how to configure TI-LFA. So, here is your homework for this section: return to Example 15-3, which I took from a massive production lab we have within Cisco to show the latest technologies. Study it and locate the two highlighted commands that start with **fast-reroute**. I recommend you enable this on all provider facing links; they will provide “automagic” protection mechanisms for your entire network without having to build backup tunnels. Know that TI-LFA works seamlessly with Flex-Algo we discussed in the previous section.

NOTE One last thing that will not come up on the exam (as TI-LFA is positioned as a better alternative). Be aware that the TI-LFA concept is not new. RFC 4090 (*Fast Reroute Extensions to RSVP-TE for LSP Tunnels*) described this very technique in 2005, 10 years before Segment Routing hit the street) in Section 3.1 (One-to-one Backup method), albeit RSVP used signaling for tunnels and Segment Routing uses a label stack to guide the packet onto the new path. I’m here to assist you in preparing for your exam and provide valuable background information, without engaging in debates over differing viewpoints.

PCE-PCC Architecture

PCE-PCC architecture involves a Path Computation Element (PCE) that centrally computes optimal network paths and a Path Computation Client (PCC) that requests these paths, enabling efficient and scalable traffic engineering across the network. To take a step back, Segment Routing Traffic Engineering (SRTE) allows the network operator to force a packet anywhere on the network. The ingress router will contain the policy containing the operator’s intent. If the network is small like the basic topology we have been using for our examples, there are only a handful for routers that we have to individually configure with such policies. A great majority of service provider networks you are likely to encounter in your career will contain dozens, hundreds, or maybe thousands of nodes. The task of deploying a

uniform policy on such a distributed network domain becomes laborious and operationally costly. How do you scale this type of rollout? There are many other limitations operators have encountered on distributed SR (or RSVP-TE for that matter) networks. Among the more notable ones is stale policies, in which operators define a set of policies and in six months traffic patterns change, which leads to continuous “rinse-and-repeat” of policy redeployment. Another one would be applications requesting the best available path in real time—not something that can be automatically done with the SRTE approach we have described thus far. What about being able to offer paths that meet certain SLAs? There are many other ones.

From the beginning of Traffic Engineering, the need for a centralized optimization element that can dynamically adjust policies based on current network conditions was apparent. Enter Path Computation Element Protocol (PCEP). It was initially specified to support the classic RSVP-TE protocol. With the introduction of Segment Routing, PCEP has been extended to support SRTE. RFC 4655 defines multiple terms that support PCEP-based architecture. Of immediate interest to us are the following terms:

- **Path Computation Element (PCE)**, which is “an entity that can compute a network path or route based on a network graph, and of applying computational constraints during the computation. The PCE entity is an application that can be located within a network node or component, on an out-of-network server, etc...” (RFC 4655)
- **Path Computation Client (PCC)**, which is “a client application requesting a path computation to be performed by the Path Computation Element...” (RFC 4655)
- **Path Computation Element Protocol (PCEP)**, which is north-bound API capable, meaning that it can ingest information coming from the network routers (via BGP-LS updates, for example) and make real-time Traffic Engineering decisions based on current network conditions. This is extremely powerful and desired on modern networks.

Notice that you can run PCE on the router itself or can rely on another adjunct processor to perform this function. In the case of Cisco products, that would be the Crosswork Network Controller, which provides a wide assortment of functionalities that helps customers to simplify and automate intent-based network service provisioning, visualization, monitoring, and optimization in a multivendor network environment with a common GUI and API. In Cisco’s documentation, you will often encounter references to SR-PCE. When you see these, it will either be a router running PCE or the Crosswork Network Controller. It can also think of PCE as a BGP Route-Reflector for Segment Routing and associated services. The following is a partial list of its capabilities (get the overall picture, do not memorize these for the exam):

- Segment Routing (SR) policy provisioning with explicit intent (for example, bandwidth constraints, latency minimization, etc.).
- Services provisioning (for example, L2VPN, L3VPN services with associated segment routing policy).
- Collection of real-time performance information and network optimization to maintain the intent of the associated segment routing policy.
- Tactical optimization of the network during times of congestion.

- Assistance with migration to next-generation networks and technologies (for example, migration from RSVP-TE to SR-TE, implementing multicast with SR Tree-SID, embracing 5G network slicing, etc.).
- Monitoring and troubleshooting the health of L2VPN and L3VPN services through empirical data plane verification.
- Streamlining and automating network-focused Method of Procedure (MOP) for remediation and maintenance tasks.

Key Topic

What makes the SR-PCE controller so powerful is that it provided centralized SRTE visibility into multidomain topologies, something that SRTE routers are not able to deliver. North-bound APIs allow SR-PCE to compute paths in real time. Because of the above, the SR-PCE can construct SLA-aware path computations even across network domains while delivering end-to-end network topology awareness. Again, you should not view SR-PCE as a single all-overseeing device but rather think of a BGP Route-Reflector deployment model where intent is centrally disseminated.

Figure 15-28 shows a screenshot taken from the Crosswork Network Controller's GUI console.

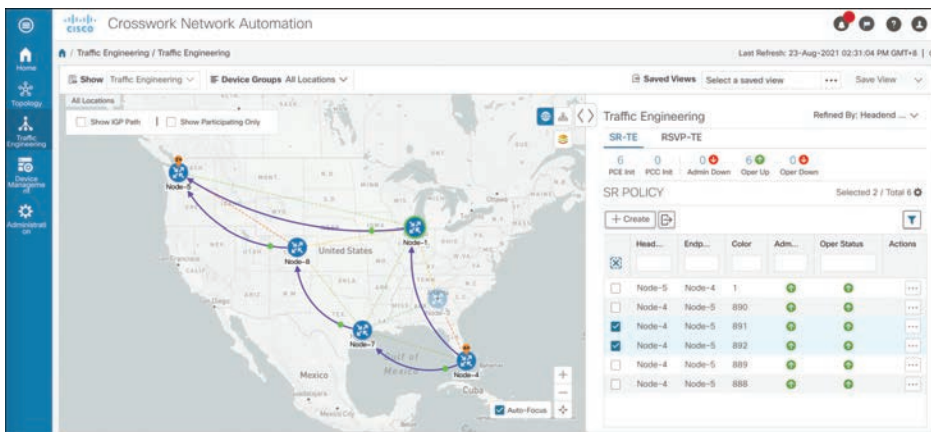


Figure 15-28 *Crosswork Network Controller*

The Cisco Crosswork Optimization Engine stands as a key element within the Crosswork Automation Suite, offering real-time network optimization capabilities. Network operators can enhance network utility and accelerate service deployment through dynamic Traffic Engineering and proactive optimization. Working seamlessly with the Crosswork Optimization Engine, the WAN Automation Engine (WAE) caters to diverse aspects of capacity management. It spans from long-term network engineering to capacity planning and Traffic Engineering, ensuring optimal network operation under various conditions. Furthermore, the WAN Automation Engine serves a valuable role in simulation analysis, aiding in the identification of potential network hotspots during failure scenarios.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have a few choices for exam preparation: the exercises here, Chapter 23, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep Software Online.

Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topic icon in the margin of the page. Table 15-5 lists a reference of these key topics and the page numbers on which each is found.



Table 15-5 Key Topics for Chapter 15

Key Topic Element	Description	Page Number
Table 15-2	LSD Label Ranges	729
Table 15-3	IS-IS TLVs	736
Example 15-3	Commands to Turn on Segment Routing in IS-IS	736
Table 15-4	OSPF Opaque LSAs	738
Figure 15-11	IPv6 Segment Routing Header Format	743
Figure 15-12	IPv6 Segment Identifier	743
Paragraph, Figure 15-17	Segment Routing Policy construction	749
Section	Binding-SID (BSID)	750
Paragraph	SRv6 encapsulation	753
Section	Terms from Remote LFA Technology	754
Paragraph	TI-LFA improvements	757
Paragraph	Crosswork considerations	759

Define Key Terms

Define the following key terms from this chapter and check the answers in the glossary:

BGP control plane, Binding Segment Identifier (BSID), Candidate Paths, Extended P-Space, Flex-Algo, Global Segments, IGP Adjacency Segment, IGP Prefix Segment, IS-IS Control Plane, Label Distribution Protocol (LDP), Label Switching Database (LSD), Local segment, OSPFv2 Control Plane, P-Space, Path Computation Client (PCC), Path Computation Element (PCE), Path Computation Element Protocol (PCEP), PCE-PCC architecture, PQ node, Q-space, Segment Routing, Segment Routing Control Plane, Segment Routing Global Block (SRGB), SR-MPLS (Segment Routing based on MPLS data plane), SRv6 (Segment Routing based on IPv6 data plane), SRv6 control plane, Topology-Independent Loop-free Alternate (TI-LFA)

Command Reference to Check Your Memory

This section includes the most important configuration and EXEC commands covered in this chapter. You might not need to memorize the complete syntax of every command, but you should be able to remember the basic keywords that are needed.

To test your memory of the commands, cover the right side of Table 15-6 with a piece of paper, read the description on the left side, and then see how much of the command you can remember.

The 350-501 exam focuses on practical, hands-on skills that are used by networking professionals. Therefore, you should be able to identify the commands needed to configure and test. Note that not all commands are fully covered in the chapter, but their presence in the table below should lead you to investigate them further to understand this technology.

Table 15-6 CLI Commands to Know

Task	Command Syntax
Define Segment Routing Global Block Range	RP/0/0/CPU0:P3(config)# segment-routing global-block
Configure IS-IS advertisements to BGP-LS	RP/0/0/CPU0:P3(config-isis)# distribute link-state
Configure IS-IS to generate and accept only new-style type-length-value (TLV) objects	RP/0/0/CPU0:P3(config-isis-af)# metric-style wide [transition] [level { 1 2 }]
Enable Segment Routing for IPv4 addresses with MPLS data plane	RP/0/0/CPU0:P3(config-isis-af)# segment-routing mpls
Enable topology-independent loop-free alternate (TI-LFA) path using the IP fast reroute (FRR) mechanism	RP/0/0/CPU0:P3(config-isis-if)# fast-reroute per-prefix RP/0/0/CPU0:P3 (config-isis-if)# fast-reroute per-prefix ti-lfa
Configure the Segment Routing Mapping Server (SRMS)	RP/0/0/CPU0:P3(config)# segment-routing mapping-server prefix-sid-map address-family ipv4 10.1.100.4/32 17000 range 100
Trace the routes to a destination in a Segment Routing network	RP/0/0/CPU0:P3# traceroute sr-mpls 10.1.100.2/32
Set the preference of Segment Routing (SR) labels over Label Distribution Protocol (LDP) labels	RP/0/0/CPU0:P3(config-isis-af)# segment-routing mpls [sr-prefer]
Specify or advertise the prefix (node) segment ID (SID) as an index value in IS-IS	RP/0/0/CPU0:P3(config)# router isis 100 RP/0/0/CPU0:P3(config-isis)# interface loopback0 RP/0/0/CPU0:P3(config-isis-if)# address-family ipv4 unicast RP/0/0/CPU0:P3(config-isis-if-af)# prefix-sid index 3
Specify or advertise the prefix (node) segment ID (SID) as an absolute value in OSPF	RP/0/0/CPU0:P3# configure RP/0/0/CPU0:P3(config)# router ospf 1 RP/0/0/CPU0:P3(config-ospf)# area 0 RP/0/0/CPU0:P3(config-ospf-ar)# interface loopback0 RP/0/0/CPU0:P3(config-ospf-ar-if)# prefix-sid absolute 16003

Task	Command Syntax
Specify the Binding SID (BSID) allocation behavior	<pre>RP/0/0/CPU0:P3# configure RP/0/0/CPU0:P3(config)# segment-routing RP/0/0/CPU0:P3(config-sr)# traffic-eng RP/0/0/CPU0:P3(config-sr-te)# binding-sid explicit fallback-dynamic RP/0/0/CPU0:P3(config-sr-te)# policy SAMPLE RP/0/0/CPU0:P3(config-sr-te-policy)# binding-sid mpls 1000</pre>
Configure SRv6-TE locator and Binding SID (BSID) behavior	<pre>RP/0/0/CPU0:P3# configure RP/0/0/CPU0:P3(config)# segment-routing traffic-eng RP/0/0/CPU0:P3(config-sr-te)# srv6 locator loc1 binding-sid dynamic behavior ub6-encaps-reduced</pre>
Globally enable SRv6	<pre>RP/0/0/CPU0:P3(config)# segment-routing srv6</pre>
Configure the SRv6 Locator	<pre>RP/0/0/CPU0:P3(config-srv6)# locators RP/0/0/CPU0:P3(config-srv6-locators)# locator myLoc1 RP/0/0/CPU0:P3(config-srv6-locator)# micro-segment behavior unode psp-usd RP/0/0/CPU0:P3(config-srv6-locator)# prefix 2001:0:8::/48</pre>

Review Questions

As a part of the review, we encourage you to provide a *single-sentence answer* (keep your answers as short as possible) to the following questions. If you struggle to complete this answer in a single sentence, this may indicate a lack of clarity or reveal gaps in your understanding. We have constructed these questions to help you consolidate this chapter's information and extract the essence of the covered content.

The answers to these questions appear in Appendix A. For more practice with exam format questions, use the Pearson Test Prep Software Online.

1. How does the implementation of Segment Routing enhance network scalability and simplify Traffic Engineering compared to traditional routing protocols?
2. In what ways can Segment Routing contribute to improved network resiliency and faster convergence times, especially in the face of dynamic changes or failures?
3. How can Segment Routing adapt to support emerging trends such as 5G networks, edge computing, and the increasing demand for network automation?
4. What specific scenarios or network topologies benefit the most from using the TI-LFA loop-free backup path mechanism?
5. Can you name one key benefit of integrating SRv6 to meet the evolving demands of modern applications, services, and emerging technologies?

Bibliography

- S. Bryant, C. Filsfils, S. Previdi, M. Shand, and N. So. RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*, <https://www.ietf.org/rfc/rfc7490.txt>, IETF, April 2015.
- P. Camarillo, Ed. RFC 8986, *Segment Routing over IPv6 (SRv6) Network Programming*, <https://www.ietf.org/rfc/rfc8986.txt>, IETF, February 2021.
- D. Dukas, Ed. RFC 8754, *IPv6 Segment Routing Header (SRH)*, <https://www.ietf.org/rfc/rfc8754.txt>, IETF, March 2020.
- A. Farrel, J.-P. Vasseur, and J. Ash. RFC 4655, *A Path Computation Element (PCE)-Based Architecture*, <https://www.ietf.org/rfc/rfc4655.txt>, IETF, August 2006.
- C. Filsfils. *Segment Routing, Part II: Traffic Engineering*, Self-published, 2019 (ISBN: 978-1095963135).
- C. Filsfils, K. Talaulikar, Ed., D. Voyer, A. Bogdanov, and P. Mattes. RFC 9256, *Segment Routing Policy Architecture*, <https://www.ietf.org/rfc/rfc9256.txt> IETF, July 2022.
- L. Ginsberg, Ed. RFC 8667, *IS-IS Extensions for Segment Routing*, <https://www.ietf.org/rfc/rfc8667.txt>, IETF, December 2019.
- LabN. RFC 5250, *The OSPF Opaque LSA Option*, <https://www.ietf.org/rfc/rfc5250.txt>, IETF, July 2008.
- J. Liste. *A Guide to a Successful Segment Routing Deployment*, Cisco Live Presentation, 2023.
- S. Litkowski, A. Bashandy, C. Filsfils, P. Francois, and B. Decraene. Internet Draft, *Topology Independent Fast Reroute Using Segment Routing*, <https://www.ietf.org/archive/id/draft-ietf-rtgwg-segment-routing-ti-lfa-11.html>, IETF, June 2023.
- P. Pan, G. Swallow, and A. Atlast, Eds. RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, <https://www.ietf.org/rfc/rfc4090.txt>, IETF, May 2005.
- S. Previdi, Ed. RFC 8665, *OSPF Extensions for Segment Routing*, <https://www.ietf.org/rfc/rfc8665.txt>, IETF, December 2019.
- S. Previdi. RFC 8670, *BGP Prefix Segment in Large-Scale Data Centers*, <https://www.ietf.org/rfc/rfc8670.txt>, IETF, December 2019.
- S. Previdi and L. Ginsberg, Eds. RFC 8402, *Segment Routing Architecture*, <https://www.ietf.org/rfc/rfc8402.txt>, IETF, July 2018.
- Redback Networks, Inc. RFC 5305, *IS-IS Extensions for Traffic Engineering*, <https://www.ietf.org/rfc/rfc5305.txt>, IETF, October 2008.
- E. Rosen. RFC 8277, *Using BGP to Mind MPLS Labels to Address Prefixes*, <https://www.ietf.org/rfc/rfc8277.txt>, IETF, October 2017.
- A. Roy. Internet Draft, *OSPFv3 LSA Extendibility*, <https://datatracker.ietf.org/doc/html/draft-ietf-ospf-ospfv3-lsa-extend-23>, IETF, January 2018.
- Segment Routing website: <https://www.segment-routing.net>

This page intentionally left blank



Index

Numerics

- 4G/5G, 31, 32–33
- 802.1ad, 535–537
 - C-bridge component, 539
 - NNI ports, 539
 - PBB (provider backbone bridging), 539–541
 - ports, 537–538
 - S-bridge component, 538–539
 - service provider bridges, 538

A

- AAA (authentication, authorization, and accounting), 846–847
- abort command, 56
- ABR (area border router), 11, 12
- AC (attachment circuit), 504
- ACL (access control list), 865–866, 946
 - classification, 844–845
 - troubleshooting, 308–309
- address families
 - BGP, 244–246
 - troubleshooting, 310–311
- address-family command, 115, 246
- adjacency/ies. *See also* IS-IS; OSPF (Open Shortest Path First)
 - authentication
 - on IOS*, 185
 - on IOS XR*, 184–185
 - debugging, 135
 - IS-IS, 129–131
 - multi-area
 - intra-area routing*, 201–202
 - on IOS*, 203–204
 - on IOS XR*, 204
 - OSPF auto-cost reference*, 202
 - OSPF costs*, 202–203
 - states, 163
- AFI (address family identifier), 111, 244
- AFRINIC (African Network Information Centre), 233
- aggregate-address command, 278–280
- AIGP (Accumulated IGP) path attribute, 275–276
- algorithm
 - constraint-based routing, 668
 - Dijkstra’s Shortest Path First, 162
 - SPF, 110
- Ansible, 1008
- API (application programming interface), 978–979
 - asynchronous, 979
 - northbound, 979, 985
 - REST (Representational State Transfer), 847–848, 979–980, 980
 - security, 979
 - SOAP (Simple Object Access Protocol), 980–981
 - southbound, 980, 985
 - synchronous, 979
- APNIC (Asia-Pacific Network Information Centre), 233

- application hosting, 68–70
- archiving configurations, 45–46
- areas
 - IS-IS, 129, 152–155
 - OSPF, 160–161
 - nonbackbone*, 162
 - not-so-stubby*, 196–199
 - stubby*, 193–196
 - totally not-so-stubby*, 199–200
- ARIN (American Registry for Internet Numbers), 232
- ASBR (autonomous system boundary router), 328, 624
- ASM (Any Source Multicast), 424
- ASN (autonomous system number), 236–237
 - Asdot format, 237–238
 - asplain to asdot+ conversion, 237
- asynchronous API, 979
- AToM (Any Transport over MPLS), 498–500
 - LFIB (Label Forwarding Information Base)
 - verifying output*, 501–502
 - xconnect endpoints*, 502
 - validating Layer 2 bindings, 502
 - VPWS validation, 500–501
 - xconnect configuration, 500
- attach bit, 143–145
- attack, DDoS (distributed denial-of-service), 849–851
- authentication
 - BGP, 310, 790–793
 - IS-IS, 148
 - configuring in IOS*, 148–149
 - configuring in IOS XR*, 149–150
 - LSP*, 150–152
 - LDP, 778–780
 - LDP session protection, 480–483
 - OSPFv3, 218–221

- token-based, 848–849
- automation. *See also* model-driven programmability
 - Ansible, 1008
 - gRPC, 991–992
 - NSO (Network Services Orchestrator), 992–995
 - RON (Routed Optical Networking), 31
 - Secure ZTP (Zero Touch Provisioning), 995–996
 - components*, 996
 - how it works*, 997–998
 - initial setup*, 996–997
- AS (autonomous system), 231–232
 - iBGP vs. eBGP, 233–234
 - key characteristics, 232
 - override, 308
 - troubleshooting, 310
- Auto-RP, 410–416

B

- backbone carrier, 654
- backhaul, 34
- BBU (baseband unit), 33
- BD (bridge domain), 541
- BDR (backup designated router), 162, 169–171
 - elections, 171–175
 - interface priority, 176–179
- Best Effort, 659
- Best External feature, BGP, 341–344
- BFD (Bidirectional Forwarding Detection), 327, 917–918
 - configurations, 919–921
 - Echo function, 918
- BGP (Border Gateway Protocol), 76, 231, 458
 - ABR (area border router), 12
 - ACL/firewalls, troubleshooting, 308–309

- Add-Path, 13
- address families, 244–246, 310–311
- AFI (address family identifier), 244
- ASN (autonomous system number), 236–237
 - Asdot format*, 237–238
 - Asplain to Asdot+ conversion*, 237
- as-path prepending, 263–266
 - inbound*, 267
 - outbound*, 266–267
- authentication, 790–793
- authentication, troubleshooting, 310
- AS (autonomous system), 231–232
 - key characteristics*, 232
 - troubleshooting*, 310
- Best External feature, 341–344
- community list, 290–291
- community/ies, 282–284
 - cost*, 292–293
 - GSHUT*, 290
 - Internet*, 289
 - no-advertise*, 285–286
 - no-export*, 285, 296–297
 - user-defined*, 286–289
- confederations, 304–306
- eBGP, neighbor relationship, 235–236
- EOR (End-of-RIB) message, 326
- fast peering session deactivation, 319–322
- GR (Graceful Restart), 344–346
- Graceful Restart extensions, 912–913
- hold and keepalive timers, 320
- iBGP, 233–234
 - and IGP interaction, 326–328
- label mapping, 11–12
- loop prevention, 306–307
 - Allow As*, 307
 - AS override*, 308
- maximum-prefix, 811–812
- messages, 238–239, 344
- MRAI (Minimum Route Advertisement Interval), 317–319
- multipath, 273–274
- neighbor states, 239–243
 - Established*, 243
 - Idle*, 243–244
- next-hop tracking, 322–326
- next-hop-self, 257–258
- NSR (non-stop routing), 347–348
- Origin Validation, 816
- path attributes, 248–249
 - AIGP (Accumulated IGP)*, 275–276
 - local preference*, 262–263
 - MED*, 268–273
 - origin*, 267–268
 - weight*, 261–263
- path selection, 255–261, 274–275
- peer template, 291–292
- PIC (Prefix-Independent Convergence), 13, 333–340
- prefix advertisement, 246–248
 - using network statement*, 250
 - using redistribution*, 253–255
- prefix exchange, 89
- prefix suppression, 812–815
- redistribution, 280
- route aggregation, 277–280
- route dampening, 328–333
- route filtering, 803
 - base BGP configuration*, 803–805
 - distribute list with IOX XE*, 805–807
 - route map using IOS XE*, 807–809
 - RPL (Route Policy Language)*, 809–811
- route map, 76, 84–86
- route policy, 89–91, 251–253

- route reflection, 297–304
 - routing table, 256, 258
 - scalability, 297
 - scanner, 322
 - security, 789–790
 - Segment Routing, control plane, 739–741
 - shadow route reflector, 340–341
 - speaker, 235
 - table validation, 250–251
 - TTL (time to live), troubleshooting, 309
 - TTL security, 235–236, 793–794, 799–802
 - one hop*, 796
 - three hops*, 798–799
 - two hops*, 795–798
 - verifying EBGp multihop*, 795
 - update message, 514
 - VPLS signaling, 513–515, 517
 - bridge domain definition*, 516
 - L2 instance verification*, 517–518
 - label assignment*, 518–521
 - learned routes*, 516–517
 - peering establishment*, 515–516
 - service instance interface association*, 516
 - bgp always-compare-med command**, 268–272
 - bgp asnotation dot command**, 238
 - bgp deterministic-med command**, 270–271
 - BGP Flowspec**, 822–824, 830
 - on the client, 828
 - configuration before and after on the client, 824–825
 - configuration before and after on the server, 825–827
 - IOS XE client, 828–829
 - NLRI types, 829
 - traffic filtering actions, 829
 - bgp graceful-restart command**, 345–346, 912–913
 - bgp rpki server command**, 816
 - bgp safe-ebgp-policy command**, 89
 - bgp suppress-inactive command**, 814–815
 - BGP-Free Core**, 476–477
 - BGPsec**, 821–822
 - BiDir-PIM**, 418–424
 - Boolean operators**, 100
 - botnet**, 849–850
 - bridge group**, 541
 - broadcast network, IS-IS**, 133–134
 - BSID (Binding-SID)**, 750–751
 - BSR (Bootstrap Router)**, 416, 442–447
-
- ## C
- Candidate Path**, 750
 - carrier-delay command**, 327–328
 - CBTS (Class-Based Tunnel Selection)**, 716–720
 - CCMs (continuity check messages)**, 532
 - CDN (content delivery networks)**, 5
 - CE (customer edge device)**, 456–457
 - certification, MEF (Metro Ethernet Forum)**, 8
 - CGNAT (Carrier-Grade NAT)**, 887–888
 - channel**, 21
 - Cisco Crosswork Optimization Engine**, 759
 - Cisco IOS**, 42
 - Cisco NOS**, 40
 - classic IOS XR**, 49
 - classification ACL**, 844–845
 - client/server model**, 984
 - CLNS (Connectionless Network Service)**, 110

- cloud, virtualization, 61**
- combining, IGP segments, 734–735**
- command/s**
 - abort, 56
 - address-family, 115, 246
 - aggregate-address, 278–280
 - BGP advertisement interval, 319
 - bgp always-compare-med, 268–272
 - bgp asnotation dot, 238
 - bgp deterministic-med, 270–271
 - bgp graceful-restart, 345–346, 912–913
 - bgp rpki server, 816
 - bgp safe-ebgp-policy, 89
 - bgp suppress-inactive, 814–815
 - Boolean operators, 98–100
 - carrier-delay, 327–328
 - commit, 58
 - commit replace, 54
 - copy startup-config running-config, 45
 - debug bgp ipv4 unicast updates, 258–259
 - debug bgp ipv4 updates, 301–303
 - debug ip bgp, 323
 - debug ip igmp, 369–372
 - debug ip igmp snooping, 361–362
 - debug ip ospf adjacency, 171–174
 - debug ip packet, 372–373
 - debug ip pim, 428–429
 - debug ip routing, 342–343
 - debug ip rsvp dump-messages path, 677–678
 - debug ipv6 mld, 383–386
 - debug isis adj-packets, 131, 135–137
 - debug mpls traffic-eng forwarding-adjacency, 691
 - do, 57
 - do debug ip igmp, 373–376
 - do show ip igmp groups, 377–379
 - do show ip mroute, 407
 - do show ip route, 153, 198–199
 - do show ip route isis, 80–82
 - do show ip route ospf, 195–197
 - do show isis database, 153–154
 - do show isis neighbor, 137
 - do show policy-map interface, 967–968
 - do show running-config, 79, 85
 - Docker, 65–68
 - domain-id type, 582
 - exit, 55
 - fast-external-fallover, 320–321
 - fast-reroute, 757
 - flush-delay-time, 911–912
 - interface loopback0, 138
 - ip access-list standard, 631
 - ip as-path access-list, 295
 - ip explicit-path, 680
 - ip igmp join-group, 362
 - ip nat inside source list, 631
 - ip nat pool, 631
 - ip ospf area, 167–168
 - ip ospf network point-to-point, 460
 - ip pim dr-priority, 409–410
 - ip pim sparse-mode, 362
 - ip prefix-list, 101–103
 - ip vrf vrf-name, 558–559
 - iperf, 69
 - ipv6 ospf authentication, 218–219
 - log-adjacency changes, 115
 - max-metric router-lsa, 609
 - mpls bgp forwarding, 648, 657–658
 - mpls ip, 465
 - mpls label protocol, 460
 - mpls ldp autoconfig, 465–466
 - mpls ldp discovery targeted-hello accept, 479
 - mpls traffic-eng reoptimize, 684
 - mpls traffic-eng tunnels, 679
 - neighbor ebgp-multihop, 235

- neighbor send-community, 291, 511
- no bgp default route-target filter, 647
- ping, 431, 633–634, 637–646
- ping mpls ipv4, 488–489
- platform qos marker-statistics, 968
- pwd, 55
- queue-limit, 961
- redistribute connected, 575–576
- root, 55
- router bgp 100, 233–236, 624–625, 637, 644–646, 650
- router bgp *autonomous system*, 233
- router isis, 83
- send-rp-announce, 411
- send-rp-discovery, 411
- service-policy input, 967–968
- show, 58
- show access-lists, 845
- show bfd interfaces location, 921
- show bfd session, 919–921
- show bgp ipv4 unicast, 342–343
- show bgp ipv4 unicast longer prefixes, 261
- show bgp l2vpn evpn summary, 549
- show bgp l2vpn vpls all, 513
- show bgp nexthops, 325–326
- show bgp summary, 236
- show bgp vpnv4 unicast all, 626, 627–628, 637
- show bgp vpnv4 unicast all command, 627–628
- show bgp vpnv4 unicast all summary, 566–567
- show bpg ipv4 unicast, 258, 334–340, 626
- show bpg ipv4 unicast dampening parameters, 330–333
- show bpg ipv4 unicast summary, 791–793
- show bridge-domain, 508, 523
- show bundle, 922
- show clns, 113–114, 118
- show configuration commit, 52
- show configuration failed, 55–56
- show install active summary, 41–51
- show install active summary command, 41–51
- show ip bgp, 812
- show ip bgp vpnv4 vrf, 638
- show ip igmp groups, 430–431
- show ip igmp snooping groups, 363–364
- show ip igmp snooping querier, 362
- show ip interface brief, 168
- show ip mfib, 432–433
- show ip mroute, 398, 403–405, 407–409, 432
- show ip nat translations, 640
- show ip ospf database, 199
- show ip ospf interface, 180
- show ip ospf interface brief, 168–169, 205, 464
- show ip ospf mpls ldp interface, 465, 466
- show ip ospf neighbors, 174–175, 179, 206
- show ip pim, 397
- show ip pim rp, 412, 419–421
- show ip route, 82, 138, 141, 142, 146–151, 148, 464
- show ip route isis, 83
- show ip route ospf, 86, 471–472, 582
- show ip route vrf, 463
- show ipv6 interface, 128
- show ipv6 interface brief, 211
- show ipv6 ospf neighbors, 219–220, 223–224
- show ipv6 route, 127, 128
- show ipv6 route ospf, 214
- show isis adjacency, 130
- show isis brief, 116
- show isis database, 133–134

- show isis database detail, 125–127, 143–145, 151–152
- show isis hostname, 130
- show isis instance, 117–118
- show isis interface, 118–120, 132
- show isis neighbor, 116
- show isis neighbor detail, 120, 128–129, 150
- show isis nsr, 913–917
- show isis protocol, 113, 138
- show isis protocols, 115–116
- show isis topology, 139
- show log, 116–117
- show mld groups, 386–387, 435–437
- show mpls forwarding, 571
- show mpls forwarding table, 463, 469, 471, 501, 570, 626–627
- show mpls l2transport binding, 508
- show mpls l2transport vc, 507, 512, 523
- show mpls label range, 467
- show mpls ldp bindings, 784–785
- show mpls ldp bindings local, 467–468
- show mpls ldp discovery, 461, 507
- show mpls ldp igp sync, 485–486
- show mpls ldp neighbor, 478–479, 501, 779–780
- show mpls traffic-eng forwarding-adjacency, 690–691
- show mpls traffic-eng tunnels, 681–682, 686
- show ospf database, 186, 188–189
- show ospf database router, 187–188
- show ospf virtual-links, 208–209
- show ospfv3 database router, 214–216
- show ospfv3 neighbors, 213
- show pim ipv6 group-map, 445–447
- show pim ipv6 neighbor, 438–439
- show pim ipv6 topology, 439–442
- show pim ipv6 tunnel info all, 439
- show policy-map control-plane, 769
- show route isis, 142
- show rpl route-policy, 98
- show running-config, 77, 124–125, 222–223, 395–397
- show running-config router bgp, 800–802
- show version, 50
- show vfi, 522
- show vrf, 560
- show xconnect all, 500
- suppress-signaling-protocol ldp, 515–516
- traceroute mpls, 490
- tunnel mpls traffic-eng forwarding-adjacency, 689
- tunnel mpls traffic-eng path-option command, 683
- vrf definition, 559
- xconnect encapsulation mpls, 499
- commit commands, 58**
- commit replace command, 54**
- community list, 290–291**
- community/ies, 282–284**
 - cost, 292–293
 - GSHUT, 290
 - Internet, 289
 - no-advertise, 285–286
 - no-export, 285, 296–297
 - user-defined, 286–289
- conditional statement, route policy, 91–94**
- confederations, BGP, 304–306**
- congestion avoidance, 961**
 - RED (Random Early Detection), 962–964
 - WRED (Weighted Random Early Detection), 964–966
- constraint-based routing, 668**
- container/s, 63–64**
 - application hosting, 68–70
 - Docker, 65–68

control plane

LDP, 778

Segment Routing, 735

BGP, 739–741*IS-IS*, 735–737*OSPFv2*, 737–739*SRv6*, 742

conversion, Asplain to Asdot+, 237

CoPP (Control Plane Policing), 768, 769

policy, 768–769

template, 770–772

verifying, 769

copy startup-config running-config command, 45

core architecture

Metro Ethernet, 5

E-LAN, 7*E-Line*, 5–6*E-Tree*, 7

MPLS, 8–9

intermediate LSR, 10*IP packet structure with MPLS encapsulation*, 9*labels*, 9–10*LSP (label switched path)*, 10*LSR (label switch router)*, 10*Unified*, 10–13

CoS (Class of Service), 944–945

cost community, 292–293

CPE (customer-premises equipment), 17

CPRI (Common Public Radio Interface), 33–34

CPU, protecting, 767–768

CsC (Carrier Supporting Carrier), 654–656

BGP configuration from PE and CE, 656

BGP VPNv4 output, 657

connectivity verification, 656

eBGP neighborhood building, 657

enabling MPLS on interface, IOS XR, 658

IOS XR host route, 657

traceroute output and label assignment, 658–659

CSPF (Constraint-based SPF), 672–673

customer carrier, 655

customer service instance, 527

C-VLAN (customer VLAN), 536

CWDM (coarse wavelength-division multiplexing), 22

D

data center, 5

data model, YANG, 986–987, 990–991

data plane protection, 854. *See also* uRPF (Unicast Reverse Path Forwarding)

ACLs (access control lists), 865–866

MACsec, 872–873

configuration and verification, 875–878*enabling on the first peer*, 874–875*header*, 873–874*tags*, 873

RTBH (Remote Triggered Black Hole) filtering, 866–867

bit bucket routes configured, 867–868*CE router verification*, 870*PE router verification*, 870–871*source-based*, 871–872*trigger router configuration on IOS XR*, 868–869*trigger router in action*, 869–870

uRPF (Unicast Reverse Path Forwarding), 855–856, 864

interface verification, 864–865

- Loose Mode*, 856–857
- operational modes*, 856
- setup*, 858–861
- Strict Mode*, 857–858, 863–864
- verification*, 861–863
- VRF Mode*, 865
- database
 - FRR (Fast Reroute), 700, 707
 - label switching, 728–729
 - LSA, Type 2, 188–189
- DDoS (distributed denial-of-service)
 - attack, 849–851
- debug bgp ipv4 updates command, 301–303
- debug ip bgp command, 323
- debug ip igmp command, 369–372
- debug ip igmp snooping command, 361–362
- debug ip ospf adjacency command, 171–174
- debug ip packet command, 372–373
- debug ip pim command, 428–429
- debug ip routing command, 342–343
- debug ip rsvp dump-messages path command, 677–678
- debug ipv6 mld command, 383–386
- debug isis adj-packets, 131
- debug isis adj-packets command, 135–137
- debug mpls traffic-eng forwarding-adjacency command, 691
- Dense mode, 392
- DevOps, 982, 1008
- DiffServ, 660, 712–713, 715–716
- Dijkstra’s Shortest Path First (SPF)
 - algorithm, 162
- DIS (Designated Intermediate System), 131–133
- discovery
 - LDP, 461–462
 - VPLS (Virtual Private LAN Service), 505
- distance-vector routing protocols, 75
- distribute list, IOS XE, 805–807
- DNS64, 895–897
- do command, 57
- do debug ip igmp command, 373–376
- do show ip igmp groups command, 377–379
- do show ip mroute command, 407
- do show ip route command, 153, 198–199
- do show ip route isis command, 80–82
- do show ip route ospf command, 195–197
- do show isis database command, 153, 154
- do show isis neighbor command, 137
- do show policy-map interface
 - command, 967–968
- do show running-config command, 79, 85
- Docker, 65–68
- DOCSIS (Data over Cable Service Interface Specification), 16
 - architecture, 16
 - CPE (customer-premises equipment), 17
 - HFC (hybrid fiber coaxial) network, 17
 - standards, 17
- domain-id type command, 582
- DR (designated router), 162, 169–171
 - elections, 171–175
 - interface priority, 176–179
- Draft-Rosen model, 611–613
- D-RAN (Distributed Radio Access Network), 33
- DS1 line, 21
- DS3 line, 21

DSCP (Differentiated Services Code Point), 943–945

DSL (Digital Subscriber Line), 17–18
 architecture, 18–19
 service provider offerings, 19

DS-Lite, 897–898

Dual-Rate Three-Color Marker, 954–956

dual-stack, single-topology IS-IS, 126–129

DWDM (dense wavelength-division multiplexing), 5, 21–24

dynamic NAT (Network Address Translation), 886

dynamic routing protocols, 75

E

E-Access, 498

eBGP (external Border Gateway Protocol), 233–234
 loop prevention, 306–307
 neighbor relationship, 235–236
 peering session deactivation, 319–322
 ttl-security feature, 235–236

egress LSR, 10

EIGRP (Enhanced Interior Gateway Routing Protocol), 75

E-LAN, 7

elections, DR/BDR, 171–175

E-Line, 5–6, 496
 multiple point-to-point service, 6
 point-to-point service, 6

encapsulation, 753

enhanced distance-vector routing protocols, 75

EOR (End-of-RIB) message, 326

ERP (Ethernet Ring Protection), 525

error messages, RSVP, 676

Established neighbor state, 243

Ethernet, 5. *See also* Metro Ethernet

Ethernet CFM (Connectivity Fault Management), 526–527
 CCMs (continuity check messages), 532
 cross-check function, 532–533
 customer service instance, 527
 Ethernet, 526
 loopback messages, 532
 maintenance association, 529–536
 maintenance domain, 528–529
 maintenance point, 529
 MEPs (maintenance endpoints), 530
 inward-facing, 530
 outward-facing, 530–531
 OAM interaction, 533–535
 traceroute messages, 532

Ethertype, 455–456, 539

E-Tree, 7, 497–498

E-UTRA (Evolved UMTS Terrestrial Radio Access), 34

EVC (Ethernet Virtual Circuit), 496–497, 533

Evolved IOS XR, 49

EVPN (Ethernet VPN), 541–544
 MPLS-based data plane, 547–550
 next-generation solutions for L2VPN, 545–547
 route types, 546–547

exit command, 55

explicit join, 393

Extended P-Space, 755–756

F

fast-external-fallover command, 320–321

fast-reroute commands, 757

FEC (Forwarding Equivalence Class), 463

FHR (First Hop Router), 356

FIB (Forwarding Information Base), 462, 464–465

filter list, 295–296

firewall, interaction with BGP, troubleshooting, 308–309

flapping route, 328–329

Flex-Algo, 751–753. *See also* algorithm

flows, 773–775, 999

flush-delay-time command, 911–912

fronthaul, 33–34

FRR (Fast Reroute), 14, 668, 695–696. *See also* MPLS TE (MPLS Traffic Engineering)

backup tunnel, 697, 699–700

database, 700

on IOS XR, 703, 710

- backup tunnel configuration*, 706
- CEs connectivity via traceroute*, 705
- database*, 707, 711
- label*, 707–708
- label stacking*, 706, 708–710
- LDP enable*, 704–705
- LFIB validation*, 705–706
- MPLS TE configuration*, 703–704
- next hop resolving via static route*, 704
- traceroute output*, 707
- traceroute outputs before and after failure*, 711

label tracking, 703

Link Protection, 697

MP (Merger Point), 696

NHOP (Next-Hop) backup tunnel, 696

NNHOP (Next-Next-Hop backup tunnel), 696

Node Protection, 697–698

PLR (Point of Local Repair), 696

traffic reroute, 700–702

tunnel configuration, 698–699

G

G.8032 ERP (Ethernet Ring Protection)

CFM (Connectivity Fault Management). *See also* CFM (Connectivity Fault Management)

CCMs (continuity check messages), 532

inward-facing MEPs, 530

loopback messages, 532

maintenance association, 529–530

maintenance domain, 528

maintenance point, 529

MEPs, 530

MIPs, 531

outward-facing MEPs, 530–531

customer service instance, 527

Ethernet CFM, 526–529

instance, 526

loop avoidance, 525

R-APS (Ring Automatic Protection Switching), 525

RPL types, 525

global configuration mode, 44

global segments, 728–730

group filtering, IGMPv3, 376–379

gRPC, 984, 991–992

GSHUT community, 290

H

HCL (Hashicorp Config Language), 1008

headend, 679

header

MACsec, 873–874

MPLS, 455

Hello messages, 459

HFC (hybrid fiber coaxial) network, 17

- hierarchical command-line structure, IOS XR, 54–55
 - hierarchical MPLS. *See* Unified MPLS
 - hierarchical policing, 956–958
 - high availability, 905–906
 - BFD (Bidirectional Forwarding Detection), 917–918
 - configurations*, 919–921
 - Echo function*, 918
 - BGP GR (Graceful Restart), 344–346
 - BGP NSR (non-stop routing), 347–348
 - NSF (Non-Stop Forwarding), 906–907
 - Graceful Restart extensions for BGP*, 912–913
 - IS-IS*, 907–910
 - OSPF*, 910–913
 - NSR (non-stop routing), 913–917
 - H-VPLS (Hierarchical VPLS), 519–521
 - architecture, 522
 - bridge domain
 - definition*, 522
 - verifying*, 523
 - CE-to-CE connectivity, 524–525
 - checking VC l2transport, 523
 - core layer, 519
 - edge layer, 520
 - L2 VFI instance configuration, 523
 - label assignment, 524
 - split-horizon, 521–522
 - VFI definition, 521
 - hypervisor, 61
-
- IANA (Internet Assigned Numbers Authority), 355
 - iBGP (internal Border Gateway Protocol), 233–234
 - versus eBGP, 306–307
 - multipath, 274
 - ICMP (Internet Control Message Protocol), traffic policing, 947–953
 - Idle neighbor state, 243–244
 - IEEE 802.1ad. *See* 802.1ad
 - IGMP (Internet Group Management Protocol), 354
 - basic routing setup, 358–359
 - Leave message, 367–369
 - Membership Query message, 362
 - membership report, 364–366
 - reference diagram, 358
 - snooping, 359–366
 - IGMPv1, 357
 - IGMPv2, 357
 - IGMPv3, 358
 - group filtering, 376–379
 - turning on, 372–376
 - IGP (interior gateway protocol), 11
 - and BGP interaction, 326–328
 - metric, 325
 - segments, 731
 - Adjacency*, 733
 - combining*, 734–735
 - Prefix*, 731–733
 - implicit null label, 474
 - inbound as-path prepending, 267
 - infrastructure, 2, 17
 - inside global address, 885
 - inside local address, 885
 - Inter-AS MPLS L3VPN
 - option A, 641–643
 - option AB, 652–654
 - option B, 643–648
 - option C, 649–652
 - interface loopback0 command, 138
 - intermediate LSR, 10
 - Internet access, MPLS L3VPN
 - extranet with Internet-VRF, 635–638
 - separate PE-CE interface, 634–635

- VRF-aware NAT, 639–640
- VRF-specific default route, 629–634
- Internet community**, 289
- intra-area routing**, 201–202
- intradomain LSP (label switched path)**, 12–13
- IntServ**, 659
- inward-facing MEPS**, 530
- IOS**
 - IS-IS
 - authentication*, 148–149
 - multitopology, enabling*, 124
 - LDP configuration, 562–563
 - OSPF
 - adjacency authentication*, 185
 - multi-area adjacencies*, 203–204
 - starting*, 164–168
- IOS XE**, 45
 - archiving and replacing configurations, 45–46
 - BGP Flowspec client, 828–829
 - BGP route filtering
 - distribute list*, 805–807
 - route map*, 807–809
 - copying and pasting, 45
 - global configuration mode, 44
 - IS-IS configuration, 113–115, 121
 - MPP (Management Plane Protection), 838–840
 - as-path prepending, 264
 - PIEs (package installation envelopes), 48
 - privileged EXEC mode, 44
 - reverting configurations with a revert timer, 47
 - user EXEC mode, 43
- IOS XR**, 47
 - BGP
 - prefix exchanges*, 89
 - route filtering*, 809–811
 - weight attribute configuration*, 262
 - command options, 53–54
 - commit labels, 53
 - Evolved, 49
 - exiting the configuration mode, 56–57
 - export-map configuration, 604
 - FRR (Fast Reroute), 703, 710
 - backup tunnel configuration*, 706
 - CEs connectivity via traceroute*, 705
 - database*, 707, 711
 - label*, 707–708
 - label stacking*, 706, 708–710
 - LDP enable*, 704–705
 - LFIB validation*, 705–706
 - MPLS TE configuration*, 703–704
 - next hop resolving via static route*, 704
 - traceroute output*, 707
 - traceroute outputs before and after failure*, 711
 - hierarchical command-line structure, 54–55
 - IS-IS
 - authentication*, 149–150
 - configuration*, 115–120, 121
 - single-topology, enabling*, 124
 - l2transport interface configuration, 509
 - LNT, 50–51
 - microkernel architecture, 47
 - MPP (Management Plane Protection), 840–843
 - OSPF
 - adjacency authentication*, 184–185
 - multi-area adjacencies*, 204
 - starting*, 164
 - as-path prepending, 264–265

- SMU (software maintenance upgrade), 49
- software packages, 48
- two-stage commit, 51–54
- virtualized, 49
- VPLS L2VPN configuration, 509
- VPLS validation, 509
- VRF configuration, 560–561
- IoT (Internet of Things), 4**
- ip access-list standard command, 631**
- ip as-path access-list command, 295**
- ip explicit-path command, 680**
- ip igmp join-group command, 362**
- ip nat inside source list command, 631**
- ip nat pool command, 631**
- ip ospf area command, 167–168**
- ip ospf network point-to-point command, 460**
- ip pim dr-priority command, 409–410**
- ip pim sparse-mode command, 362**
- IP Precedence, 943–944**
- ip prefix-list command, 101–103**
- ip vrf vrf-name command, 558–559**
- iperf command, 69**
- IPFIX (IP Flow Information Export), 999–1000**
 - Exporter Map configuration, 1000
 - interface attachment, 1001
 - Monitor Map configuration, 999–1000
 - Sampler Map configuration, 1000
 - validation, 1001
- IPv6, 884**
 - Flow Label, 971–972
 - segment routing over, 15–16
 - transitions
 - CGNAT (Carrier-Grade NAT)*, 887–888
 - DNS64*, 895–897
 - DS-Lite*, 897–898
 - MAP-E*, 898–899
 - MAP-T*, 899–900
 - NAT*, 884–886
 - NAT64*, 888
 - outside NAT*, 887–888
 - PAT*, 886–887
 - stateful NAT64*, 892–895
 - stateless NAT64*, 8892
- ipv6 ospf authentication command, 218–219**
- i-SID (service instance identifier), 541**
- IS-IS, 108–109**
 - adjacencies, 129–131, 135
 - areas, 129, 152–155
 - attach bit, 143–145
 - authentication, 148
 - configuring in IOS*, 148–149
 - configuring in IOS XR*, 149–150
 - broadcast network, 133–134
 - CLNS (Connectionless Network Service), 110
 - configuration, 112–113
 - desired results*, 121
 - IOS XE*, 113–115
 - IOS XR*, 115–120
 - control plane, 735–737
 - DIS (Designated Intermediate System), 131–133
 - extensions for MPLS TE, 671–672
 - hierarchical design, 110
 - interface MTU mismatch, troubleshooting, 136–137
 - levels, 141
 - metrics, 137–141
 - multitopology, enabling on IOS, 124
 - NET (Network Entity Title) address, 111
 - network types, 131
 - NSF (Non-Stop Forwarding), 907–910
 - overload bit, 146–148
 - point-to-point network, 134

- pseudonode, 132
- route advertisement, 141–146
- route map, 77, 80–83
- Segment Routing, turning on, 736–737
- single-topology
 - dual-stack*, 126–129
 - enabling on IOS*, 124
 - enabling on IOS XR*, 124–126
- single-topology transition mode, 122–123
- SPF algorithm, 110
- TLV (Type/Length/Value) extensions, 110
- troubleshooting, 155
- ISP (Internet service providers), services, 4–5
- ITY-T G.8032 Ethernet Ring Protection Switching. *See* G.8032 ERP (Ethernet Ring Protection)

J-K

- join-group message, 356
- JSON (JavaScript Object Notation), 986, 987
- JWT (JSON Web Token), 848–849
- Keepalive message, BGP, 238
- kernel architecture, 42

L

- label mapping, BGP, 11–12
- label/s, 9–10
 - advertisement, 469–476
 - assignment, 458–464, 467–468
 - FRR (Fast Reroute), 707–708
 - header fields, 455
 - implicit null, 474
 - reserved, 455
 - stacking, 457–458

- switch router, 455
- transport, 564
- types, 457
- VPN, 564
- LACNIC (Latin America and Caribbean Network Information Centre), 233
- LACP (Link Aggregation Control Protocol), 925–933
- LAG (Link Aggregation Group), 921–924
- lambda, 21
- LDP (Label Distribution Protocol), 12, 458, 505, 562, 727
 - authentication, 778–780
 - autoconfig, 465–467
 - co-existence with SRv6, 746–748
 - configuring on IOS, 562–563
 - configuring on IOS XR, 563–564
 - control plane, 778
 - discovery parameters, 461
 - Hello messages, 459
 - IGP synchronization, 483–488
 - inbound filtering, 781–783
 - inbound label control, 471
 - label allocation filtering, 783–789
 - manipulating the router ID, 461
 - outbound filtering, 781
 - outbound label control, 469
 - session establishment, 459–461
 - session protection, 477–483
 - verifying discovery, 462
 - VPLS signaling
 - autodiscovery*, 509–513
 - manual*, 502–509
- Leave message, IGMP, 367–369
- LER (label edge router), 457
- levels, IS-IS, 141
- LFIB (Label Forwarding Information Base), 462
 - output, 469, 470
 - verifying output, 501–502

LHR (Last Hop Router), 356
LIB (Label Information Base), 462
Link Protection, 694–697
link-state routing protocols, 75–76, 108. *See also* IS-IS; OSPF (Open Shortest Path First)
Linux kernel, IOS XE, 42
local paths, BGP, 263
local preference attribute, 262–263
local segments, 730
log-adjacency changes command, 115
loop avoidance, ERP (Ethernet Ring Protection), 525
loop prevention, BGP, 306–307
 Allow As, 307
 AS override, 308
loopback messages, 532
loopback network, OSPF, 183–185
Loose uRPF, 856–857
LPTS (Local Packet Transport Services), 772–773
 configuring policer, 7778
 default policing, 7777
 flows, 773–775
 identification, 775–777
 keeping LDP safe, 778
LSA (link-state advertisement), 160–161, 162
 Opaque, 737–738
 OSPFv3, 210
 Type 1 router links, 186–188
 Type 3, 189
 Type 4, 191–192
 Type 5, 190–191
LSD (Label Switching Database), 728–729
LSP (label switched path), 10, 12, 624
 intradomain, 12–13
 MD5 authentication, 150–152
LSR (label switch router), 10, 457

egress, 10
 intermediate, 10

M

MACsec, 872–873
 configuration and verification, 875–878
 enabling on the first peer, 874–875
 header, 873–874
 tags, 873
maintenance association, 529–530
maintenance point, 529
MAM (Maximum Allocation Model), 713–714
MAP-E (Mapping of Address and Port Encapsulated), 898–899
MAP-T (Mapping of Address and Port Translated), 899–900
max-metric router-lsa command, 609
MDT (multicast distribution tree), 611
MED (Multi-Exit Discriminator) attribute, 268–273
MEF (Metro Ethernet Forum), 8
Membership Query message, 362
membership report, IGMP, 364–366
MEPs (maintenance endpoints)
 inward-facing, 530
 outward-facing, 530–531
message/s
 BGP, 238–239
 EOR (End-of-RIB), 326
 Open, 344
 BGP update, 514
 CFM (Connectivity Fault Management)
 CCMs (continuity check messages), 532
 loopback, 532
 traceroute, 532
 Hello, 459

- IGMP
 - Leave*, 367–369
 - Membership Query*, 362
- join-group, 355
- MLDv1, 379–380
- traceback, 843–844
- metrics**
 - IGP, 325
 - IS-IS, 137–141
- Metro Ethernet, 496**
 - E-Access, 498
 - E-LAN, 7, 497
 - E-Line, 5–6, 496
 - multiple point-to-point service*, 6
 - point-to-point service*, 6
 - E-Tree, 7, 497–498
- midhaul, 34**
- MIPS (maintenance intermediate points), 531**
- MLD (Multicast Listener Discovery), 379**
 - messages, 379–380
 - reference diagram, 380–381
- MLDv2**
 - enabling on IOS routers, 381–382, 395–397
 - joining multicast groups, 382–392
- mobility radio technology**
 - 4G/5G, 31
 - differences between 4G and 5G, 32–33
 - packet core, 32
 - xHaul, 33–35
- model-driven programmability, 982**
 - client/server model, 984
 - gRPC, 984
 - NETCONF, 984, 988
 - agent*, 985
 - capabilities*, 988–989
 - client/server communication*, 985
 - XML elements*, 989
 - RESTCONF, 984
 - YANG, 984
- MP (Merger Point), 696**
- MP-BGP, 564–565**
 - BGP VPNv4 configuration
 - on IOS*, 565
 - on IOS XR*, 566
 - ICMP connectivity, 565
 - VRF definition and interface interaction on IOS XR, 566
- MPLS (Multiprotocol Label Switching), 8–9, 454. See also AToM (Any Transport over MPLS)**
 - CE (customer edge device), 456–457
 - enabling on a physical interface, 460
 - Ethertype, 455–456
 - FEC (Forwarding Equivalence Class), 463
 - FIB (Forwarding Information Base), 462, 464–465
 - header, 455
 - IP packet structure with MPLS encapsulation, 9
 - label/s, 9–10, 14–15, 455
 - advertisement*, 469–476
 - assignment*, 458–464, 467–468
 - header fields*, 455
 - implicit null*, 474
 - reserved*, 455
 - stacking*, 457–458
 - types*, 457
 - LDP, 562
 - autoconfig*, 465–467
 - configuring on IOS*, 562–563
 - configuring on IOS XR*, 563–564
 - discovery parameters*, 461
 - IGP synchronization*, 483–488
 - inbound label control*, 471
 - manipulating the router ID*, 461
 - session establishment*, 459–461

- session protection*, 477–480
- verifying discovery*, 462
- LER (label edge router), 457
- LFIB (Label Forwarding Information Base), 462, 469, 470
- LIB (Label Information Base), 462
- LSP (label switched path), 10
- LSR (label switch router), 457
- OAM (Operation, Administration, and Maintenance), 488–491
- PHP (penultimate hop popping), 473–476
- ping, 488–489
- QoS, 660–661
 - Pipe mode*, 661–662
 - Short Pipe mode*, 662
 - Uniform mode*, 661
- RIB (Routing Information Base), 462, 464–465
- seamless, 277
- traceroute, 490
- Unified, 10–13. *See also* Unified MPLS
- VPLS (Virtual Private LAN Service), 505
- mpls bgp forwarding command**, 648, 657–658
- MPLS EXP field**, 945–946
- mpls ip command**, 465
- MPLS L3VPN**, 556, 576–577
 - BGP PE-CE routing, 595–596
 - BGP allowas-in*, 596–597
 - BGP AS-override*, 597
 - BGP AS-PATH loop prevention*, 596
 - CsC (Carrier Supporting Carrier), 654–656
 - BGP configuration from PE and CE*, 656
 - BGP VPNv4 output*, 657
 - connectivity verification*, 656
 - eBGP neighborhood building*, 657
 - enabling MPLS on interface*, IOS XR, 658
 - IOS XR host route*, 657
 - traceroute output and label assignment*, 658–659
- EIGRP PE-CE routing
 - EIGRP configuration*, 591
 - EIGRP routing table output*, 591–592
 - EIGRP topology output*, 592–593
 - SoO (Site of Origin)*, 593–594
- import/export maps, 601
 - export-map configuration on IOS XR*, 604
 - route-target import/export validation*, 601–603
- Inter-AS
 - option A*, 641–643
 - option AB*, 652–654
 - option B*, 643–648
 - option C*, 649–652
- Internet access, 629
 - extranet with Internet-VRF*, 635–638
 - separate PE-CE interface*, 634–635
 - VRF-aware NAT*, 639–640
 - VRF-specific default route*, 629–634
- NG mVPN, 614
- OSPF PE-CE routing, 577–578
 - BGP routing table for VRF SPCOR*, 584
 - BGP VPN v4 output*, 580–581
 - connectivity verification*, 579–580
 - down bit*, 589–591
 - mutual redistribution*, 579
 - OSPF basic configuration*, 577–578
 - OSPF database output*, 581

- OSPF domain ID manipulation, 582–583
- OSPF domain-ID configuration on IOS XR, 584–585
- OSPF instance per VRF definition, 578
- OSPF process ID and domain ID verification, 580
- OSPF process ID verification, 583
- OSPF router-ID configuration, 584
- route table for OSPF output, 581–582
- route table OSPF output verification, 583
- VRF related verification commands, 578–579
- OSPF PE-CE routing with backdoor link, 585–586
 - OSPF metric manipulation, 588
 - OSPF neighbors output, 587
 - OSPF network advertisement, 585–587
 - sham-link establishment, 586–587
- router reflector deployment, 607–610
- RT (route target) filtering, 598
 - BGP VPN, route reflector, 600–601
 - debug BGP VPNv4 updates output, 598
 - IOS XR BGP retain route target, 599
 - turning off BGP default route-target filter, 598–600
- RTC (Route Target Constraint), 604
 - BGP rtfiler neighborship establishment, 606–607
 - debug BGP VPNv4 unicast update output, 605–606
- static PE-CE routing, 567
 - BGP VPNv4 output verification, 571
 - connectivity verification, 570
 - default route configuration, 569
 - FIB output, 571–572, 573
 - ICMP failure, 573–575
 - label assignment for VPNv4, 572
 - LFIB output, 570–571
 - LFIB verification for a certain prefix, 572–573
 - redistribution under BGP for a specific VRF, 569
 - RIB output, 569–570
 - static route under VRF configuration, 568
 - traceroute output, 555–573
 - verifying connectivity through VRF, 568
- mpls label protocol command, 460
- MPLS label stacking, 612
- mpls ldp autoconfig command, 465–466
- mpls ldp discovery targeted-hello accept command, 479
- MPLS TE (MPLS Traffic Engineering), 667–670
 - constraint-based routing, 668
 - CSPF (Constraint-based SPF), 672–673
 - Forwarding Adjacency, 687–688
 - configuration, 689
 - debug output, 691–692
 - detailed output, 690–691
 - holdtime, 691
 - traceroute output, 688
 - FRR (Fast Reroute), 668, 695–696
 - backup tunnel, 697
 - backup tunnel configuration, 699–700
 - backup tunnel configuration, IOS XR, 706
 - CEs connectivity via traceroute, IOS XR, 705
 - database, 700, 707, 711

- on IOS XR, 703
- IOS XR, 710
- label, 707–708
- label stacking, 706
- label stacking, IOS XR, 708–710
- label tracking, 703
- LDP enable, IOS XR, 704–705
- LFIB validation, IOS XR, 705–706
- Link Protection, 697
- MP (Merger Point), 696
- MPLS TE configuration, IOS XR, 703–704
- next hop resolving via static route, IOS XR, 704
- NHOP (Next-Hop) backup tunnel, 696
- NNHOP (Next-Next-Hop backup tunnel), 696
- Node Protection, 697–698
- PLR (Point of Local Repair), 696
- traceroute output, IOS XR, 707
- traceroute outputs before and after failure, IOS XR, 711
- traffic reroute, 700–702
- tunnel configuration, 698–699
- IS-IS extensions, 671–672
- OSPF extensions, 670–671
- PCE (Path Computation Element), 668
- Per-VRF, 692–693, 694–695
- placing traffic in a tunnel, 678, 686
 - Autoroute Announce, 678
 - connectivity validation between CEs, 681
 - end-to-end connectivity, 678–679
 - forcing soft optimization, 684
 - Forwarding Adjacency, 678
 - headend, 679
 - interface tunnel with explicit path, IOS and IOS XE, 679–680
 - interface tunnel with explicit path, IOS XR, 680
 - IPG TE-related commands, 679
 - MPLS forwarding table output, 686–687
 - OSPF learned routes, 684
 - path option manipulation, 683–684
 - path options, 683
 - RSVP configurations, 679m
 - static routes for placing traffic on the tunnel, 685–686
 - tailend, 679
 - traceroute outputs between CEs, 685
 - traffic engineering tunnel details, 681–683
 - tunnel configuration, 684–685
 - tunnel details, 683
- QoS (Quality of Service), 712
 - CBTS (Class-Based Tunnel Selection), 716–720
 - DiffServ, 715–716
 - DS-TE, 712–713
 - MAM (Maximum Allocation Model), 713–714
 - PBTS (Policy-Based Tunnel Selection), 717–720
 - RDM (Russian Dolls Model), 714–715
- resource reservation, 668
- RSVP (Resource Reservation Protocol), 673–674
 - error messages, 676
 - PATH message, 674–675
 - RESV message, 675
 - tear messages, 676–678
- tunnels, 668
- mpls traffic-eng reoptimize command, 684**
- mpls traffic-eng tunnels command, 679**

MPP (Management Plane Protection), 838

in IOS XE, 838–840

in IOS XR, 840–843

MQC (Modular Quality of Service Command-Line Interface), 939–940. See also QoS (Quality of Service)

hierarchical policing, 956–958

tail drop, 961–962

MRAI (Minimum Route Advertisement Interval), 317–319**MTU (maximum transmission unit), troubleshooting, 311****multi-access network, OSPF, 169–170****multi-area adjacencies, 200–201**

costs, 202–203

intra-area routing, 201–202

on IOS, 203–204

OSPF auto-cost reference, 202

multicast, 352, 354, 354–355. See also MVPN (multicast VPN)

Any Source, 424

Dense mode, 392

FHR (First Hop Router), 356

IGMP (Internet Group Management Protocol), 356–357*basic routing setup, 358–359**Leave message, 367–369**membership report, 364–366**reference diagram, 358**snooping, 359–366*

IGMPv1, 357

IGMPv2, 357

IGMPv3, 358

*group filtering, 376–379**turning on, 372–376*

important address ranges, 355

join-group message, 356

LHR (Last Hop Router), 356

MAC-to-IP address mapping, 355

MLD (Multicast Listener Discovery), 379*messages, 379–380**reference diagram, 380–381***MLDv2***enabling on IOS routers, 381–382**joining multicast groups, 382–392**turning on, 395–397***PIM (Protocol-Independent Multicast), 354***Auto-RP, 410–416**Bidirectional, 418–424**BSR (Bootstrap Router), 416, 442–447**mapping a group to a certain RP, 418**SPT switchover, 405–410***PIM-SM, 392***connecting the multicast receiver, 400–402**connecting the multicast source, 398–400**multicast trees, 392–395**PIM SPT switchover, 405–410**RP trees, 395**RPs, 410***PIMv6, 433–437****Source-Specific, 424–425***initializing, 427–428**launching traffic from source, 431–433**turning on, 425–427*

static RP, 439–442

mVPN (multicast VPN), 610–611

BGP/MPLS, 615–617

Draft-Rosen model, 611–613

MDTs (multicast distribution trees), 611

NG multicast, 614

PIM/GRE, 613–614
 service types, 611

N

NAT (Network Address Translation)

Carrier-Grade, 887–888
 dynamic, 886
 inside global address, 885
 inside local address, 885
 outside, 887–888
 outside global address, 885
 outside local address, 885
 static, 885–886
 VRF-aware, 639–640

NAT64, 888

stateful, 892–895
 stateless, 8892

NCS 540 Fronthaul router, 37

neighbor ebgp-multihop command, 235

neighbor send-community command,
 291, 511

neighbor states, BGP, 239–243

Established, 243
 Idle, 243–244

NET (Network Entity Title) address, 111

NETCONF, 984, 988

agent, 985
 capabilities, 988–989
 client/server communication, 985
 XML elements, 989

NetFlow, 845–846, 999. *See also* IPFIX (IP Flow Information Export)

network

data model, YANG, 986
 flows, 999
 model-driven programmability,
 982. *See also* model-driven
 programmability
client/server model, 984

gRPC, 984

NETCONF, 984, 988–989

RESTCONF, 984

YANG, 984

orchestration and monitoring systems,
 984–985

security, 766

streaming telemetry, 1001–1004

**network statement, 164–168,
 249–250, 253, 267, 576**

network types, OSPF, 181–185

next-hop tracking, BGP, 322–326

next-hop-self statement, 257–258

**NFV (network function
 virtualization), 4**

**NFVI (Network Functions
 Virtualization Infrastructure),
 61–63**

NHOP (Next-Hop) backup tunnel, 696

**NNHOP (Next-Next-Hop backup
 tunnel), 696**

**no bgp default route-target filter
 command, 647**

no-advertise community, 285–286

Node Protection, 697–698

Node SID (Segment Identifier), 729

no-export community, 285, 296–297

nonbackbone areas, OSPF, 162

northbound API, 979, 985

NOS (network operating systems), 40.

See also IOS; IOS XE; IOS XR

Cisco IOS, 42

IOS XE, 45

*archiving and replacing
 configurations, 45–46*

copying and pasting, 45

global configuration mode, 44

Linux kernel, 42

privileged EXEC mode, 44

user EXEC mode, 43

IOS XR, 47

- classic*, 49
 - command options*, 53–54
 - commit labels*, 53
 - Evolved IOS XR*, 49
 - exiting the configuration mode*, 56–57
 - hierarchical command-line structure*, 54–55
 - LNT*, 50–51
 - microkernel architecture*, 47
 - PIEs (package installation envelopes)*, 48
 - SMU (software maintenance upgrade)*, 49
 - software packages*, 48
 - two-stage commit*, 51–54
 - virtualized*, 49
 - kernel architecture, 42
 - Notification message, BGP, 238
 - not-so-stubby areas, 196–199
 - NSEL (NSAP Selector), 111
 - NSF (Non-Stop Forwarding), 906–907
 - Graceful Restart extensions for BGP, 912–913
 - IS-IS, 907–910
 - OSPF, 910–913
 - NSO (Network Services Orchestrator), 992–995
 - NSR (non-stop routing), 347–348, 913–917
- O**
-
- OADM (optical add/drop multiplexer), 22
 - OAM (Operation, Administration, and Maintenance), 526–527. *See also* CFM (Connectivity Fault Management)
 - OAM manager, 533–535
 - Opaque LSAs, 737–738
 - Open message, 238, 344
 - O-RAN (Open Radio Access Network), 36–37
 - origin attribute, 267–268
 - OSPF (Open Shortest Path First), 75–76, 108–109, 193
 - adjacency/ies
 - authentication*, 184–185
 - multi-area*, 200–206
 - states*, 163
 - areas, 160–161
 - not-so-stubby*, 196–199
 - stubby*, 193–196
 - totally not-so-stubby*, 199–200
 - BDR (backup designated router), 169–171
 - elections*, 171–175
 - interface priority*, 176–179
 - DR (designated router), 169–171
 - elections*, 171–175
 - interface priority*, 176–179
 - extensions for MPLS TE, 670–671
 - LSA
 - Opaque*, 737–738
 - Type 1*, 186–188
 - Type 2*, 188–189
 - Type 3*, 189
 - Type 4*, 191–192
 - Type 5*, 190–191
 - LSAs (link-state advertisements), 162
 - multi-access network, 169–170
 - network types, 181–185
 - nonbackbone areas, 162
 - NSF (Non-Stop Forwarding), 910–913
 - packet types, 162–163
 - route advertisement, 186–193
 - route map, 76, 86
 - SPF tree, 162
 - starting
 - with interface-specific commands*, 168–169

- on IOS*, 164–168
 - on IOS XR*, 164
 - tiebreaker for equal-cost external routes, 192–193
 - timers, 179–181
 - topology, 161
 - troubleshooting, 225
 - virtual links, 206–209
 - OSPFv2, control plane**, 737–739
 - OSPFv3**, 209–210
 - authentication, 218–221
 - configuring, 211–214
 - LSA types, 210
 - LSDB view, 216–217
 - multiple instances, 221–224
 - Options field bits, 214–216
 - Router LSA bits, 217–218
 - outbound as-path prepending, 266–267
 - outside global address, 885
 - outside local address, 885
 - outside NAT (Network Address Translation), 887–888
 - outward-facing MEPs, 530–531
 - overload bit, IS-IS, 146–148
- P**
-
- PAT (Port Address Translation)**, 886–887
 - path attributes**, 248–249
 - AIGP (Accumulated IGP), 275–276
 - local, 263
 - local preference, 262–263
 - MED (Multi-Exit Discriminator), 268–273
 - origin, 267–268
 - as-path prepending, 263–267
 - weight, 261–263
 - PATH message, RSVP**, 674–675
 - as-path prepending**, 263–266
 - inbound, 267
 - outbound, 266–267
 - path selection, BGP**, 255, 274, 275
 - path vector routing algorithm**, 76. *See also* BGP
 - PBB (provider backbone bridging)**, 539–541
 - configuration, 542–543
 - EVPN (Ethernet VPN) components, 541–543
 - verifying, 543
 - PBTS (Policy-Based Tunnel Selection)**, 717–720
 - PCE (Path Computation Element)**, 668
 - PCEP (Path Computation Element Protocol) controller**, 15
 - PCE-PCC architecture**, 757–759
 - peer template**, 291–292
 - Per-VRF MPLS TE**, 692–693, 694–695
 - PHP (penultimate hop popping)**, 455, 473–476
 - PIC (Prefix-Independent Convergence)**, 333–340
 - PIEs (package installation envelopes)**, 48
 - PIM (Protocol-Independent Multicast)**, 354
 - Auto-RP, 410–416
 - Bidirectional, 418–424
 - BSR (Bootsrap Router), 416, 442–447
 - mapping a group to a certain RP, 418
 - SPT switchover, 405–410
 - PIM-SM**, 412
 - connecting the multicast receiver, 400–402
 - connecting the multicast source, 398–400
 - explicit join, 393
 - multicast trees, 392–393
 - shared*, 394–395

- source-based*, 393–394
- requirements for support in an mVPN, 614
- RP trees, 395
- RP (rendezvous points), 410
- PIMv2, 437–439
- PIMv6, 433–437
- ping, 431, 633–634, 637–646
- ping mpls ipv4 command, 488–489
- platform qos marker-statistics command, 968
- PLE (Prive Line Emulation), 21
- PLR (Point of Local Repair), 696
- point-to-point network, IS-IS, 134
- policy. *See also* QoS (Quality of Service)
 - based QoS, 660
 - CoPP, 768–769
 - route, 251–253
 - Segment Routing, 749, 750
 - trigger, 869
- ports, 802.1ad, 537–538
- POTS (plain old telephone service), 18
- PQ Node, 754–755
- prefix advertisement, BGP, 246–248
 - using network statement, 250
 - using redistribution, 253–255
- prefix list, 100–103
- prefix set, 103–105
- priority queuing, 958
- privileged EXEC mode, 44
- provider bridge, 535–537
- pseudonode, 132
- pseudowire, 6
- P-Space, 754
- punt a packet, 491
- pwd command, 55

Q

- Q-in-Q encapsulation, 535
- QoS (Quality of Service), 659, 938–939
 - Best Effort, 659
 - CBTS (Class-Based Tunnel Selection), 716–720
 - congestion avoidance, 961
 - RED (Random Early Detection)*, 962–964
 - WRED (Weighted Random Early Detection)*, 964–966
 - DiffServ, 660
 - IntServ, 659
 - IPv6 Flow Label, 971–972
 - MPLS, 660–661
 - Pipe mode*, 661–662
 - Short Pipe mode*, 662
 - Uniform mode*, 661
 - MPLS TE
 - DiffServ-TE solution*, 712–713
 - MAM (Maximum Allocation Model)*, 713–714
 - RDM (Russian Dolls Model)*, 714s-716
 - PBTS (Policy-Based Tunnel Selection), 717–720
 - policy-based, 660
 - priority queuing, 958
 - traffic classification, 940–942
 - 802.1Q VLAN tag*, 942
 - CoS (Class of Service)*, 944–945
 - DSCP (Differentiated Services Code Point)*, 943–945
 - IP Precedence*, 943–944
 - matching on access-lists*, 946
 - MPLS EXP field*, 945–946
 - ToS (Type of Service) byte*, 943–945

- values for class of service,*
942–943
 - traffic marking, 966–967
 - imposing multiple markings,*
968–969
 - internal marking,* 969–971
 - marking packets directly on the policy map,* 967–968
 - traffic policing, 946–947
 - Dual-Rate Three-Color Marker,*
954–956
 - hierarchical,* 956–958
 - ICMP traffic,* 947–953
 - Single-Rate Three-Color Marker,*
953–954
 - Single-Rate Two-Color Marker,*
953
 - traffic shaping, 958–960
 - Q-Space, 754–755
 - queue-limit command, 961
- ## R
-
- RAN (Radio Access Network)
 - O-RAN (Open Radio Access Network), 36–37
 - split points, 35
 - ranging, xPON, 25–26
 - R-APS (Ring Automatic Protection Switching), 525
 - RD (route distinguisher), 557–558
 - RDM (Russian Dolls Model), 714–715
 - RED (Random Early Detection), 962–964
 - redistribute connected command, 575–576
 - redistribution, 280
 - regular expressions, 294–295
 - reserved labels, 455
 - REST API, 847–848, 979, 980
 - RESTCONF, 984
 - RESV message, RSVP, 675
 - RFC 1195, 138
 - RFC 1918, 866
 - RFC 2827, 866
 - RFC 3032, 476
 - RFC 3107, 11, 12, 624
 - RFC 3330, 866
 - RFC 3535, 983
 - RFC 3787, 143
 - RFC 3810, 380
 - RFC 4655, 758
 - RFC 5575, 829
 - RFC 7490, 754
 - RFC 8955, 829
 - RIB (Routing Information Base), 462, 464–465
 - ring ports, 525
 - RIPE NCC (Réseaux IP Européens Network Coordination Centre), 232
 - RIR (Regional Internet Registry), 232–233
 - ROA (Route Origin Authorization) records, 815–816
 - ROADM (reconfigurable optical add-drop multiplexer), 22
 - RON (Routed Optical Networking), 27–31
 - root command, 55
 - route advertisement
 - IS-IS, 141–146
 - OSPF, 186–193
 - route aggregation, BGP, 277–280
 - route dampening, BGP, 328–333
 - route map, 77
 - applying, 81–82
 - BGP, 84–86, 261–262
 - conditional matching, 78–79
 - IOS XE, 807–809
 - IS-IS, 80–83
 - modifying, 82–83

- OSPF, 86
- RPKI, 820–821
- sequence numbers, 77–78
- taking action on matched components, 79–80
- template, 77
- route policy, 251–253**
 - BGP, 89–91
 - conditional statement, 91–94
 - in configuration mode, 100
 - with an inline set, 95
 - prefix lists, 100–103
 - with a prefix set, 95–97
 - prefix sets, 103–105
- route reflection, 297–304**
 - MPLS L3VPN, 607–610
 - shadow, 340–341
- router. *See also* NOS (network operating systems)**
 - area border, 12
 - autonomous system boundary, 328, 624
 - Bootstrap, 416
 - CPU, protecting, 767–768
 - First Hop, 356
 - label edge, 457
 - label switch, 10, 455, 457
 - egress*, 10
 - intermediate*, 10
 - Last Hop, 356
 - NCS 540 Fronthaul, 37
 - OSPF, 162
 - streaming telemetry, 1001–1004
 - tracebacks, 843–844
- router bgp 100 command, 233, 234–235, 236, 624–625, 637, 644–646, 650**
- router bgp autonomous system command, 233**
- router isis command, 83**
- routing, 74–75**
- routing protocol/s, 155**
 - distance-vector, 75
 - dynamic, 75
 - enhanced distance-vector, 75
 - link-state, 75–76, 108. *See also* IS-IS; OSPF (Open Shortest Path First)
 - path vector, 76. *See also* BGP
- routing table, BGP, 256**
- RP (rendezvous point), 410, 418**
- RP trees, PIM-SM, 395**
- RPC (remote procedure call), 981–982**
- RPKI (Resource Public Key Infrastructure), 815**
 - architecture, 817
 - prefix download, 819
 - prefix marking and usage, 819–820
 - ROA (Route Origin Authorization) records, 815–816
 - route map, 820–821
 - router configuration, 818–819
 - servers, 818
- RPL (ring protection link), 525**
- RPL (Route Policy Language), 88, 809–811. *See also* route policy**
- RPs (rendezvous points), 410, 439–442**
- RSVP (Resource Reservation Protocol), 458, 673–674, 727**
 - error messages, 676
 - PATH message, 674–675
 - RESV message, 675
 - tear messages, 676–678
- RT (route target), 558, 598–601. *See also* MPLS L3VPN, RT (route target) filtering**
- RTBH (Remote Triggered Black Hole) filtering, 866–867**
 - bit bucket routes configured, 867–868
 - CE router verification, 870
 - PE router verification, 870–871

source-based, 871–872
 trigger router configuration on IOS XR, 868–869
 trigger router in action, 869–870
RTC (Route Target Constraint), 604.
See also MPLS L3VPN, RTC (Route Target Constraint)
run-to-completion scheduler, 42

S

scalability, BGP, 297
 scanner, BGP, 322
SDH (Synchronous Digital Hierarchy), 20–21
SDN (software-defined networking), 4, 14–15, 660
 seamless MPLS, 277
Secure ZTP (Zero Touch Provisioning), 995–996
 components, 996
 how it works, 997–998
 initial setup, 996–997
security. *See also* authentication; data plane protection
 API (application programming interface), 979
 BGP, 789–790. *See also* BGP (Border Gateway Protocol)
 maximum-prefix, 811–812
 prefix suppression, 812–815
 route filtering, 803–811
 TTL, 235–236, 793–802
 BGPsec, 821–822
 classification ACL, 844–845
 CoPP, 768, 769
 policy, 768–769
 template, 770–772
 verifying, 769
 LPTS (Local Packet Transport Services), 772–773
 configuring policer, 7778
 default policing, 7777
 flows, 773–775
 identification, 775–777
 keeping LDP safe, 778
 management plane, 838. *See also* MPP (Management Plane Protection)
 network, 766
 REST API, token-based authentication, 848–849

Segment Routing, 13–14, 727–728. *See also* SRTE (Segment Routing Traffic Engineering); SRv6 (segment routing over IPv6)
 BSID (Binding-SID), 750–751
 classic LFA limitations, 756–757
 control plane, 735
 BGP, 739–741
 IS-IS, 735–737
 OSPFv2, 737–739
 SRv6, 742
 Flex-Algo, 751–753
 Global Block, 728
 global segments, 728–730
 IGP Adjacency segment, 733
 IGP Prefix segment, 731–733
 IGP segments, 731, 734–735
 IS-IS, turning on, 736–737
 local segments, 730
 minimum requirements, 729
 Node SID (Segment Identifier), 729
 OSPF, turning on, 738–739
 over IPv6, 15–16
 policy, 749, 750
 terms from remote LFA technology
 Extended P-Space, 755–756
 PQ Node, 754–755
 P-Space, 754
 Q-Space, 754–755
 TI-LFA (topology-independent loop-free alternate), 753

- Traffic Engineering, 748
- send-rp-announce command, 411
- send-rp-discovery command, 411
- sequence numbers, route map, 77–78
- service provider networks, 2
- Service Provider Tag, 536
- service-policy input command, 967–968
- shadow route reflector, 340–341
- sham-link, 586–587
- shared infrastructure, 17
- shared tree, 394–395
- show access-lists command, 845
- show app-hosting detail command, 68–69
- show bfd interfaces location command, 921
- show bfd session command, 919–921
- show bgp ipv4 unicast command, 342–343
- show bgp ipv4 unicast longer prefixes command, 261
- show bgp l2vpn evpn summary command, 549
- show bgp l2vpn vpls all command, 513
- show bgp nexthops command, 325–326
- show bgp summary command, 236
- show bgp vpnv4 unicast all command, 626, 637
- show bgp vpnv4 unicast all summary command, 566–567
- show bpg ipv4 unicast command, 258, 334–340, 626
- show bpg ipv4 unicast dampening parameters command, 330–333
- show bpg ipv4 unicast summary command, 791–793
- show bridge-domain command, 508, 523
- show bundle command, 922
- show clns command, 113–114, 118
- show commands, 58
- show configuration commit command, 52
- show configuration failed command, 55–56
- show install active summary command, 48
- show ip bgp command, 812
- show ip bgp vpnv4 vrf command, 638
- show ip igmp groups command, 430–431
- show ip igmp snooping groups command, 363–364
- show ip igmp snooping querier command, 362
- show ip interface brief command, 168
- show ip mfib command, 432–433
- show ip mroute command, 398, 403–405, 407–409, 432
- show ip nat translations command, 640
- show ip ospf database command, 199
- show ip ospf interface brief command, 168–169, 205, 464
- show ip ospf interface command, 180
- show ip ospf mpls ldp interface command, 465, 466
- show ip ospf neighbors command, 174–175, 179, 206
- show ip pim command, 397
- show ip pim rp command, 412, 419–421
- show ip route bgp command, 813
- show ip route command, 82, 138, 141–142, 146–151, 148, 464
- show ip route isis command, 83
- show ip route ospf command, 86, 186, 471–472, 582
- show ip route vrf command, 463
- show ipv6 interface brief command, 211

- show ipv6 interface command, 128
- show ipv6 ospf neighbors command, 219–220, 223–224
- show ipv6 route command, 127, 128
- show ipv6 route ospf command, 214
- show isis adjacency command, 130
- show isis brief command, 116
- show isis database command, 133, 134
- show isis database detail command, 125–127, 143–145, 151–152
- show isis hostname command, 130
- show isis instance command, 117–118
- show isis interface command, 118–120, 132
- show isis neighbor command, 116
- show isis neighbor detail command, 120, 128–129, 150
- show isis nsr command, 913–917
- show isis protocol command, 113, 138
- show isis protocols command, 115–116
- show isis topology command, 139
- show mld groups command, 386–387, 435–437
- show mpls forwarding command, 571
- show mpls forwarding table command, 463, 469, 471, 501, 570, 626–627
- show mpls l2transport binding command, 508
- show mpls l2transport vc command, 507, 512, 523
- show mpls label range command, 467
- show mpls ldp bindings command, 784–785
- show mpls ldp bindings local command, 467–468
- show mpls ldp discovery command, 461, 507
- show mpls ldp igp sync command, 485–486
- show mpls ldp neighbor command, 478, 501, 779–780
- show mpls traffic-eng forwarding-adjacency command, 690–691
- show mpls traffic-eng tunnels command, 681–682, 686
- show ospf database command, 186, 188–189
- show ospf database router command, 187–188
- show ospf virtual-links command, 208–209
- show ospfv3 database router command, 214–216
- show ospfv3 neighbors command, 213
- show pim ipv6 bsr election command, 4445
- show pim ipv6 group-map command, 445–447
- show pim ipv6 neighbor command, 438–439
- show pim ipv6 topology command, 439–442
- show pim ipv6 tunnel info all command, 439
- show policy-map control-plane command, 769
- show route isis command, 142
- show rpl route-policy command, 98
- show running-config command, 77, 124–125, 222–223, 395–397
- show running-config router bgp command, 800–802
- show version command, 50
- show vfi command, 522
- show vrf command, 560
- show xconnect all command, 500
- SID (segment identifier), 15
- signaling, VPLS (Virtual Private LAN Service), 505
- single topology transition mode, 122–123
- Single-Rate Three-Color Marker, 953–954

Single-Rate Two-Color Marker, 953

single-topology IS-IS

dual-stack, 126–129

enabling on IOS, 123

enabling on IOS XR, 124–126

Slammer worm, 766

SMU (software maintenance upgrade), 49

SNMP (Simple Network Management Protocol), 983, 1004–1005

SNMPv2c, 1005

SNMPv3, 1005–1008

SOAP (Simple Object Access Protocol), 980–981

SONET (Synchronous Optical Networking), 20–21

SoO (Site of Origin), 593–594

source-based RTBH, 871–872

source-based tree, 393–394

southbound API, 980, 985

Sparse mode, 392

SPCOR 3501 exam

suggested plan for final review and study, 1012

updates, 1014–1016

speaker, BGP, 235

SPF algorithm, 110

SPF tree, 162

split-horizon rule, 297, 503

SRGB (Segment Routing Global Block), 728

SR-MPLS (Segment Routing based on MPLS data plane), 728

SRTE (Segment Routing Traffic Engineering), 14–15

SRv6 (segment routing over IPv6), 15–16, 728

co-existence with LDP, 746–748

control plane, 742

–MPLS L3 Service Interworking Gateway, 745–746

node roles, 744

SRH (Segment Routing Header), 742–743

uSID (micro segment identifier), 744–745

SSM (Source-Specific Multicast), 424–425, 428–431

initializing, 427–428

launching traffic from source, 431–433

turning on, 425–427

standards

DOCSIS, 17

DS-TE, 716

xPON (Passive Optical Network), 24

stateful NAT64, 892–895

stateless NAT64, 8892

statement. *See also* command/s

network, 164–168, 249–250, 253, 267, 576

next-hop-self, 257–258

static NAT (Network Address Translation), 885–886

static routing, 75

static RP, 439–442

STP (Spanning Tree Protocol), 7, 503

streaming telemetry, 1001–1004

Strict uRPF, 857–858, 863–864

stubby areas, 193–196

suppress-signaling-protocol ldp command, 515–516

switch, IGMP snooping, 359–366

synchronous API, 979

T

TACACS (Terminal Access Controller Access Control System), 846–847

tail drop, 961–962

tailend, 679

- TDM (time-division multiplexing), 19–21
- tear messages, RSVP, 676–678
- telcos, 4
- template
 - BGP peer, 291–292
 - CoPP, 770–772
 - route map, 77
- Terraform, 1008
- TI-LFA (topology-independent loop-free alternate), 753, 756–757
- timer/s
 - BGP hold and keepalive, 320
 - MPLS LDP session protection, 480
 - MRAI (Minimum Route Advertisement Interval), 317–319
 - OSPF (Open Shortest Path First), 179–181
- TLV (Type/Length/Value) extensions, 110, 736
- token-based authentication, 848–849
- tools
 - Ansible, 1008
 - NetFlow, 999
 - Terraform, 1008
- topology
 - IS-IS, 111–112. *See also* IS-IS
 - enabling single and multitopology TLV exchange*, 122–123
 - multi-*, 124
 - single*, 123
 - OSPF, 161
- ToS (Type of Service) byte, 943–945
- totally not-so-stubby areas, 199–200
- tracebacks, 843–844
- traceroute messages, 532
- traceroute mpls ipv4 command, 490
- traffic classification, 940–942
 - 802.1Q VLAN tag, 942
 - CoS (Class of Service), 944–945
 - DSCP (Differentiated Services Code Point), 943–945
 - IP Precedence, 943–944
 - matching on access-lists, 946
 - MPLS EXP field, 945–946
 - ToS (Type of Service) byte, 943–945
 - values for class of service, 942–943
- traffic engineering. *See* SRTE (Segment Routing Traffic Engineering)
- traffic marking, 966–967
 - imposing multiple markings, 968–969
 - internal marking, 969–971
 - marking packets directly on the policy map, 967–968
- traffic policing, 946–947
 - Dual-Rate Three-Color Marker, 954–956
 - hierarchical, 956–958
 - ICMP traffic, 947–953
 - Single-Rate Three-Color Marker, 953–954
 - Single-Rate Two-Color Marker, 953
- traffic shaping, 958–960
- transport labels, 564
- transport technology, 16
 - DOCSIS, 16
 - architecture*, 16
 - CPE*, 17
 - HFC (hybrid fiber coaxial) network*, 17
 - standards*, 17
 - DSL (Digital Subscriber Line), 17–18
 - architecture*, 18–19
 - service provider offerings*, 19
 - DWDM (dense wavelength-division multiplexing), 21–24
 - TDM (time-division multiplexing), 19–21

xPON (Passive Optical Network), 24, 25, 27
network protection modes, 26–27
ranging, 25–26
standards, 24

troubleshooting

BGP
ACL/firewalls, 308–309
address families, 310–311
authentication, 310
AS (autonomous system), 310
TTL (time to live), 309

IS-IS, 155
adjacencies, 135
interface MTU mismatch, 136–137

MTU (maximum transmission unit), 311

TSP (Tag Switching Protocol), 460

TTL (time to live)
 BGP security, 235–236, 793–794, 799–802
one hop, 796
three hops, 798–799
two hops, 795–798
verifying EBGP multihop, 795

troubleshooting, 309

tunnel mpls traffic-eng forwarding-adjacency command, 689

tunnel mpls traffic-eng path-option command, 683

tunnels, MPLS TE, 668

two-stage commit, 51–54

Type 1 LSA, 186–188

Type 2 LSA, 188–189

Type 3 LSA, 189

Type 4 LSA, 191–192

Type 5 LSA, 190–191

U

UMMT (Unified MPLS Mobile Transport), 628–629

Unified MPLS, 10–13, 624, 624
 BGP configuration, 624–625
BGP table output, 626
BGP VPNv4 output, 625–626, 627–628
LFIB output, 626–627
traceroute output, 625

LSP (label switched path), 624

Mobile Transport, 628–629

Update message, BGP, 238

updates, SPCOR 3501 exam, 1014–1016

uRPF (Unicast Reverse Path Forwarding), 855–856, 864
 interface verification, 864–865
 Loose Mode, 856–857
 operational modes, 856
 setup, 858–861
 Strict Mode, 857–858, 863–864
 verification, 861–863
 VRF Mode, 865

user EXEC mode, 43

user-defined community, 286–289

uSID (micro segment identifier), 744–745

V

VC (virtual circuit), 504

VFI (Virtual Forwarding Instance), 503–504

virtual links, 206–209

virtualization, 61
 container/s, 63–64
application hosting, 68–70

- Docker*, 65–68
- hypervisor, 61
- NFVI (Network Functions Virtualization Infrastructure), 61–63
- VNF (virtual network function) workload, 63
- virtualized XR, 49
- VLAN stacking, 535
- VNF (virtual network function) workload, 63
- VPLS (Virtual Private LAN Service), 7, 497. *See also* H-VPLS (Hierarchical VPLS)
- AC (attachment circuit), 504
- BGP signaling, 513–515, 517
 - bridge domain definition*, 516
 - L2 instance verification*, 517–518
 - label assignment*, 518–521
 - learned routes*, 516–517
 - peering establishment*, 515–516
 - service instance interface association*, 516
- discovery and signaling, 505
- LDP signaling
 - autodiscovery*, 509–513
 - manual*, 505–509
- limitations, 545–546
- MPLS, 505
- scale-limiting factors, 520
- STP (Spanning Tree Protocol), 503
- VC (virtual circuit), 504
- VFI (Virtual Forwarding Instance), 503–504
- VSI (Virtual Switching Instance), 504
- VPN labels, 564
- VRF (Virtual Routing and Forwarding), 557
 - attaching an interface, 559

- aware NAT, 639–640
- configuration, 558–561
- membership, 561–562
- RD (route distinguisher), 557–558
- RT (route target), 558
- verifying interface connectivity, 559
- vrf definition command, 559
- VRF Mode, uRPF, 865
- VSI (Virtual Switching Instance), 504

W

- wavelength, 21
- WDM (wavelength-division multiplexing), 21–22
- weight attribute, 261–263
- Wi-Fi, 5
- workload, VNF, 63
- WRED (Weighted Random Early Detection), 964–966

X-Y-Z

- x-auth-token, 848
- xconnect (cross-connect), 6
- xconnect encapsulation mpls command, 499
- xDSL. *See* DSL (Digital Subscriber Line)
- xHaul, 33–35
- XML (Extensible Markup Language), 986, 987
- xPON (Passive Optical Network), 24, 25, 27
 - network protection modes, 26–27
 - ranging, 25–26
 - standards, 24
- YANG, 984, 986–987, 990–991