

# Segment Routing for Service Provider and Enterprise Networks



[ciscopress.com](http://ciscopress.com)

FLORIAN DERAGISCH  
LEONIR HOXHA  
RENE MINDER  
MATTHYS RABE  
KATEEL VIJAYANANDA

FREE SAMPLE CHAPTER |



# Segment Routing for Service Provider and Enterprise Networks

---

Florian Deragisch (CCIE #47970)

Leonir Hoxha (CCIE #49534)

Rene Minder (CCIE #8003)

Matthys “Thys” Rabe (CCIE #4237)

Kateel Vijayananda

**Cisco Press**

Hoboken, New Jersey

# Segment Routing for Service Provider and Enterprise Networks

Florian Deragisch  
Leonir Hoxha  
Rene Minder  
Matthys Rabe  
Kateel Vijayananda

Copyright© 2025 Cisco Systems, Inc.

Published by:  
Cisco Press

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. 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](http://www.pearson.com/global-permission-granting.html).

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

## \$PrintCode

Library of Congress Control Number: 2024945941

ISBN-13: 978-0-13-823093-7

ISBN-10: 0-13-823093-5

## Warning and Disclaimer

This book is designed to provide information about segment routing. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors 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](mailto:feedback@ciscopress.com). Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

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

<b>GM K12, Early Career and Professional Learning:</b> Soo Kang	<b>Technical Editors:</b> Jakub Horn, Christian Schmutzer, Luc Andrew Burdet, Johan Gustawsson, Bram Van der Zwet
<b>Alliances Manager, Cisco Press:</b> Caroline Antonio	<b>Editorial Assistant:</b> Cindy Teeters
<b>Director, ITP Product Management:</b> Brett Bartow	<b>Designer:</b> Chuti Prasertsith
<b>Managing Editor:</b> Sandra Schroeder	<b>Composition:</b> codeMantra
<b>Development Editor:</b> Ellie C. Bru	<b>Indexer:</b> Timothy Wright
<b>Senior Project Editor:</b> Mandie Frank	<b>Proofreader:</b> Barbara Mack
<b>Copy Editor:</b> Kitty Wilson	



**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](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned 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. (1110R)

## About the Authors

All the authors of this book are integral members of Cisco's professional services and product sales organizations. They have built extensive experience and expertise throughout their careers, working closely with service providers and enterprise customers to design, implement, transform, and optimize cutting-edge network solutions.

**Florian Deragisch, CCIE #47970**, is a Technical Leader, working with large service provider and carrier-grade enterprise customers. He joined Cisco in 2012 as part of a graduate program, where he discovered his passion for service provider designs and technologies. After gaining extensive exposure to MPLS-based networks and services, he embraced the evolution toward segment routing with his first SR-MPLS deployment in 2018. More recently, he has focused on the migration and deployment of L2VPN/L3VPN SRv6 services to build simple and highly scalable network architectures. He holds a master's degree in electrical engineering and information technology from the Swiss Federal Institute of Technology in Zurich and a Cisco Internetwork Expert certification (CCIE #47970). When not busy with work, he enjoys traveling to explore new places, cultures, and food.

**Leonir Hoxha, CCIE #49534**, has been with Cisco Systems since 2013, taking on various roles on the Professional Services team and later on the Pre-sales team—from troubleshooting to designing and implementing large-scale networks with a focus on service provider technologies, specifically MPLS services. In his current role as a Solutions Architect, he supports service providers and enterprise customers by understanding their requirements and providing cutting-edge solutions. An active speaker at Cisco Live conferences, he has delivered numerous sessions on segment routing across Europe, the United States, and Australia. He holds a bachelor's degree in computer science and a Cisco Internetwork Expert certification (CCIE #49534). In his free time, he enjoys electronic music, a nod to his first job as a DJ during his teenage years.

**Rene Minder, CCIE #8003**, is a Senior Program Advisor and Solution Architect with over 25 years of experience in the IT industry. He has been responsible for architecture and delivery in more than 70 customer engagements, evolving their networks and management infrastructures as well as the processes for developing, testing, and deploying them. He has led end-to-end IT architecture projects encompassing everything from portals offering service self-administration capabilities, to IT applications that automatically configure and test changes, to invoicing. His efforts have led to significant improvements in customer satisfaction, operational efficiency, and overall agility. He holds Lifetime Emeritus status for his Cisco Internetwork Expert certification (CCIE #8003).

**Matthys “Thys” Rabe, CCIE #4237**, a Lifetime Emeritus Cisco Certified Internetwork Expert (CCIE #4237), is a Technical Leader at Cisco Systems and holds a diploma in electrical engineering (telecommunications). With more than 25 years of experience in IP and MPLS operations with various service providers in South Africa and Switzerland, he has spent the past 10 years as a Technical Support Engineer focused on Swisscom. Prior to working at Cisco Systems, he was part of the Core IP Engineering team with a Swiss-based mobile provider. When he's not working, he enjoys fishing with his brothers in various southern African countries.

**Kateel Vijayananda**, is a solutions architect at Cisco Systems and has more than 30 years of experience in the networking industry. His expertise includes IPv6, IP services, design of large-scale networks for enterprise and service provider customers, and QoS assurance in IP networks. He has been with Cisco Systems since 2001, involved in several projects for service providers to deploy IP-based services using MPLS and segment routing. Before joining Cisco, he worked at Swisscom, a service provider in Switzerland, where he was responsible for developing MPLS VPN services. He is the co-author of the book *Developing IP-Based Services: Solutions for Service Providers and Vendors*. He holds a master's degree in computer science from the University of Maryland at College Park and a PhD in computer science from the Swiss Federal Institute of Technology at Lausanne (EPFL). In his spare time, he enjoys traveling and cooking.

## About the Technical Reviewers

**Jakub Horn** has worked for more than 20 years at Cisco Systems and currently serves as a Principal Technical Marketing Engineer, specializing in cutting-edge technologies for service providers. Prior to this role, Jakub was a Network Consulting Engineer, delivering strategic solutions to global clients. His journey in tech began at IBM, where he honed his skills in networking and computer systems. Today, Jakub's expertise is centered on SRv6 technology, driving innovation in network architecture. As a passionate technologist, he continuously explores new advancements to shape the future of connectivity.

**Christian Schmutzer** is a Distinguished Engineer at Cisco Systems and has been with the company since 1998. Early in his career, he primarily worked on the design and deployment of large service provider backbones. He has been part of a business unit since 2005, serving as a technical subject matter expert for market-leading routing platforms such as the Cisco 7600 and ASR 9000. Since 2013, he has focused on packet/optical network architectures, future product definition, technology innovation, and leading customer deployments. He is the holder of several patents and the author of a series of IETF standards documents.

**Bram van der Zwet** is the Lead Architect for Network & Infrastructure at Swisscom, where he has been shaping the network architecture and technical strategy for Swisscom's IP and optical networks. His responsibilities extend to overseeing the physical infrastructure from Swisscom's IT and data centers down to the central offices in regional networks. He holds a degree from Delft University of Technology and is based in Bern, Switzerland. With more than 25 years of experience at Swisscom and a history of strategic roles driving innovation and excellence, he has become a key figure in the telecommunications industry.

**Johan Gustawsson** is a Senior Director within Cisco Data Center and Service Provider, focusing on driving the direction and strategy for routing and architectures. He has spent his entire career operating and building mass-scale networks, pioneering and driving market disruptions across routing and optical domains. Prior to joining Cisco, Johan was the Head of Network Architecture, Strategy, and Engineering at Arelion (formerly Telia Carrier), leading a globally distributed organization at the world's number-one-ranked Internet backbone. Johan holds a degree in Engineering from the KTH Royal Institute of Technology in Stockholm.

**Luc André Burdet** is a Senior Technical Leader in Engineering at Cisco, where he has been instrumental in driving innovation and strategic initiatives since May 2012. With more than 12 years of experience at Cisco, he focuses on advancing the company's engineering capabilities and leading key technical projects. He holds a master's degree from ETH Zürich and is based in Ottawa, Ontario. Luc André's technical expertise and leadership have established him as a pivotal figure in the networking industry, significantly contributing to Cisco's engineering excellence.

## Acknowledgments

First and foremost, we would like to thank our main reviewers, Jakub Horn and Christian Schmutzer, for their meticulous reviews and invaluable feedback. Their dedication and attention to detail have significantly enhanced the quality of this book.

We also extend our thanks to Luc Andre Burdet for his expertise in the chapter focused on Layer 2 VPN technologies, and to Bram Van der Zwet and Johan Gustawsson for reviewing the chapters on business opportunities and organizational considerations. Your feedback has been instrumental in ensuring the accuracy and relevance of the information presented.

Special thanks to Marcel Witmer for all the support around PLE and integrated visibility and to Christian Schmutzer for his solid insight and input on PLE. We also appreciate Kaela Loffler and Ramiro Nobre for providing an overview on how micro-drops can influence overall service performance. Similarly, we would like to express our gratitude to Carmine Scarpitta and Ahmed Abdelsalam for their guidance on FRRouting's SRv6 implementation.

The authors had the pleasure of collaborating with Swisscom, a leading service provider based in Switzerland, on several aspects covered in this book. The insights gained from Swisscom's exposure to engineering, migrations, and operations have enriched the content, providing field perspectives that are invaluable for readers.

This book wouldn't have been possible without the support of many people on the Cisco Press team. Brett Bartow, Product Line Manager of the Pearson IT professional Group, was instrumental in sponsoring the book and driving it to execution. Sandra Schroeder, Managing Editor, was masterful with book graphics. Ellie Bru, Development Editor, has done a wonderful job in the technical review cycle; it has been a pleasure working with you. Mandie Frank, Senior Project Editor, thank you for leading the book to success through the production cycle. Kitty Wilson, Copy Editor, thank you for polishing up the book and making the content more shiny. Also, many thanks to the numerous Cisco Press unknown soldiers working behind the scenes to make this book happen.

We would like to express our deepest gratitude to our Cisco management for supporting and encouraging us in creating this book. Thank you to everyone who has contributed to this book. Your support and expertise have made this project possible.

Finally, we would like to extend our heartfelt thanks to our families. Your unwavering support, patience, and understanding have been our pillars of strength throughout the writing process. The countless hours spent away from you to work on this book have not gone unnoticed, and we are deeply grateful for your encouragement and understanding.



## Contents at a Glance

Introduction xx

### **Part I Introduction**

Chapter 1 MPLS in a Nutshell 1

Chapter 2 What Is Segment Routing over MPLS (SR-MPLS)? 33

Chapter 3 What Is Segment Routing over IPv6 (SRv6)? 103

### **Part II Segment Routing**

Chapter 4 Segment Routing in Detail 219

Chapter 5 Migrating to Segment Routing 353

### **Part III Service Design**

Chapter 6 L2VPN Service Deployment: Configuration and Verification Techniques 439

Chapter 7 L3VPN Service Deployment: Configuration and Verification Techniques 605

Chapter 8 Service Assurance 783

Chapter 9 High Availability and Fast Convergence 857

### **Part IV Business and Operational Considerations**

Chapter 10 Business Opportunities 997

Chapter 11 Organizational Considerations 1043

Appendix A Reference Diagrams and Information 1109

Index 1115

### **Online Element:**

Chapter 12 SRv6 Ecosystem Deployment Use Cases 1

## Reader Services

**Register your copy** at [www.ciscopress.com/title/ISBN](http://www.ciscopress.com/title/ISBN) for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to [ciscopress.com/register](http://ciscopress.com/register) and log in or create an account\*. Enter the product ISBN 9780138230937 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

\*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

For access to any available bonus content associated with this title, visit [ciscopress.com/sr](http://ciscopress.com/sr), sign in or create a new account, and register ISBN 9780138230937 by December 31, 2027.

## Contents

Introduction xx

### Part I Introduction

#### Chapter 1 MPLS in a Nutshell 1

How MPLS Operates 4

MPLS Label Structure 6

Control Plane and Data Plane 10

Label Distribution Protocol (LDP) 11

Label Allocation Mechanism 12

MPLS Label Operations 14

Traffic Forwarding Using Labels 15

MPLS VPN Services Overview 16

MPLS Traffic Protection 18

Challenges and Shortcomings of MPLS 19

MPLS Label Space Limitation 20

LSP and Summarization 21

Inter-AS Limitations 21

*Lack of End-to-end QoS Control 21*

*Configuration and Operational Complexity of MPLS/VPN  
and BGP 22*

RSVP-Based Traffic Engineering 22

LDP-IGP Synchronization 24

Load Balancing and Hashing 25

Beyond MPLS 28

Summary 30

References and Additional Reading 32

#### Chapter 2 What Is Segment Routing over MPLS (SR-MPLS)? 33

Problem Description and Requirements 40

Segment Routing over MPLS (SR-MPLS) 41

Data Plane 41

Segment Identifier (SID) 42

*SID Allocation 43*

*IGP Prefix Segment (Prefix SID) 45*

*IGP Adjacency Segment (Adjacency SID) 47*

*BGP Prefix Segment (BGP Prefix SID) 49*

	<i>BGP Peering Segments (BGP Peering SIDs)</i>	50
	<i>Binding Segment (Binding SID)</i>	52
	IGP Extensions	53
	<i>IS-IS Extensions for Segment Routing (RFC 8667)</i>	53
	<i>OSPF Extensions for Segment Routing (RFC 8665)</i>	64
	<i>IGP Flexible Algorithm (Flex Algo) (RFC 9350)</i>	73
	MP-BGP Extensions	83
	<i>SR Prefix SID Extensions for BGP (RFC 8669)</i>	85
	<i>BGP Link-State Extensions for SR (RFC 9085)</i>	87
	<i>BGP Link-State Extensions for SR BGP Egress Peer Engineering (RFC 9086)</i>	95
	Summary	99
	References and Additional Reading	100
<b>Chapter 3</b>	<b>What Is Segment Routing over IPv6 (SRv6)?</b>	<b>103</b>
	Introduction	103
	Segment Routing over IPv6 (SRv6)	103
	IPv6 for SRv6 Recap	104
	SRv6 Network Programming (RFC 8986)	107
	<i>SRv6 Segment Identifier (SID)</i>	107
	IPv6 Segment Routing Header (SRH) (RFC 8754)	123
	<i>Penultimate Segment Pop of the SRH</i>	127
	<i>Ultimate Segment Pop of the SRH</i>	129
	<i>Ultimate Segment Decapsulation</i>	130
	<i>SRv6 Policy Headend Behaviors</i>	132
	<i>SRv6 Policy Endpoint Behaviors</i>	141
	<i>SRv6 Headend and Endpoint Behavior Overview</i>	146
	SRv6 Network Programming Extension: SRv6 uSID Instruction	147
	<i>uN Endpoint Variants</i>	152
	<i>uA Endpoint Variants</i>	156
	<i>SID Compression</i>	161
	Addressing Considerations	162
	<i>IPv6 Addressing</i>	162
	<i>SRv6 Locator Addressing Scheme</i>	165
	<i>Summarization</i>	172
	IGP Extensions	175
	<i>IS-IS Extensions for Segment Routing over IPv6 (RFC 9352)</i>	175

MP-BGP Extensions	186
<i>BGP Overlay Services on SRv6 (RFC 9252)</i>	186
<i>BGP Link-State Extensions for SRv6 (RFC 9514)</i>	201
SR-Powered Network Evolution	205
MPLS Network Architecture: Control and Data Plane Overview	205
SR-MPLS Network Architecture: Control and Data Plane Overview	207
SRv6 Network Architecture: Control and Data Plane Overview	209
Network Evolution at a Glance	210
SR-MPLS or SRv6	212
Benefits of Deploying Segment Routing	212
Hardware and Software Support	214
Feature Support	215
Summary	215
References and Additional Reading	217

## **Part II Segment Routing**

### **Chapter 4 Segment Routing in Detail 219**

Link-State IGPs	221
IS-IS	221
<i>IS-IS Levels</i>	222
<i>IS-IS Areas</i>	222
<i>IS-IS Router Types</i>	222
<i>IS-IS Routing</i>	222
<i>IS-IS Route Propagation and Leaking</i>	223
<i>IS-IS Overload Bit</i>	223
OSPF	224
<i>OSPF SPF Algorithm</i>	224
OSPFv3	225
<i>OSPFv3 Route Summarization</i>	225
<i>OSPFv3 Route Filtering</i>	225
Segment Routing Baseline	225
SR-MPLS Baseline	226
<i>Segment Routing Global Block (SRGB)</i>	226
<i>Segment Routing Local Block (SRLB)</i>	227
SR-MPLS Addressing	227
SR-MPLS Configuration	228

<i>SR-MPLS Verification</i>	231
SRv6 Baseline	236
<i>SRv6 uSID</i>	236
<i>SRv6 Addressing</i>	237
<i>SRv6 uSID Configuration</i>	239
<i>SRv6 uSID Verification</i>	241
Segment Routing Control Plane (IGP)	243
SR-MPLS Control Plane	243
<i>SR-MPLS IS-IS</i>	244
<i>SR-MPLS OSPF</i>	250
<i>SR-MPLS Anycast SID</i>	254
SRv6 Control Plane	257
<i>SRv6 IS-IS</i>	257
<i>SRv6 OSPF</i>	301
Multiplane Topologies with Flex Algos	301
<i>Components of SR Flex Algos</i>	302
<i>Flex Algo Use Cases Scenarios</i>	304
<i>SR-MPLS Configuration for Flex Algo Use Cases</i>	322
Segment Routing Control Plane (BGP)	324
BGP Prefix SID	324
<i>BGP Prefix SID Configuration</i>	324
<i>BGP Prefix SID Verification</i>	327
Intra-AS BGP-LU with a BGP Prefix SID	328
<i>Intra-AS BGP-LU Design</i>	330
<i>BGP Additional Path</i>	331
<i>Intra-AS BGP-LU Configuration</i>	332
<i>Intra-AS BGP-LU Verification</i>	337
<i>Data Forwarding from PE-1 to PE-3</i>	341
Inter-AS BGP-LU	343
<i>Inter-AS BGP-LU Design</i>	343
<i>Inter-AS BGP-LU Configuration</i>	344
<i>Inter-AS BGP-LU Verification</i>	345
<i>Data Forwarding from PE-1 to PE-5</i>	349
Summary	350
References and Additional Reading	352

## **Chapter 5 Migrating to Segment Routing 353**

Deployment Models	354
Migration Strategies	355
SR-MPLS Migration	358
SR-MPLS Reference Network Topology	358
Enabling SR-MPLS in an Existing Network (Coexistence)	358
<i>Enabling SR-MPLS on P2, P3, and PE-3</i>	360
<i>Enabling and Preferring SR-MPLS on P1, P2, and PE-1</i>	363
Building a New SR-MPLS Network	365
<i>Enabling SRMS</i>	365
<i>Enabling LDP on the Border Node</i>	372
Enabling the BGP Prefix SID in an SR-MPLS Network	376
<i>BGP Proxy Prefix SID</i>	383
SRv6 Migration	387
Building a New SRv6 Network Using an SRv6 IWG	389
<i>Migration Use Case</i>	391
Building a New SRv6 Network Using Inter-AS Option A	401
<i>Migration Use Case</i>	403
Building a New SRv6 Network Using Dual-Connected PE Devices	413
<i>Migration Use Case</i>	415
High Availability	425
<i>Active-Active</i>	425
<i>Active-Backup</i>	426
<i>Load Sharing</i>	426
Migration Paths from MPLS to SRv6	427
Flat MPLS Network	429
Unified MPLS Network	430
MPLS Network with Inter-AS Option C	431
Carrier Supporting Carrier MPLS Network	434
Summary	435
References and Additional Reading	437

## **Part III Service Design**

### **Chapter 6 L2VPN Service Deployment: Configuration and Verification Techniques 439**

L2VPN (EVPN)	440
EVPN in Detail	442

	<i>EVPN Instance (EVI)</i>	445
	<i>Ethernet Segment (ES)</i>	447
	<i>Ethernet Tag ID</i>	450
	<i>EVPN BGP Routes</i>	452
	<i>EVPN E-LAN</i>	472
	<i>SRv6 EVPN E-LAN Service Configuration and Verification</i>	474
	<i>EVPN E-Tree</i>	551
	<i>SRv6 EVPN E-Tree Service Configuration</i>	552
	<i>EVPN E-Line</i>	569
	<i>SRv6 EVPN E-Line (VPWS) Service Configuration</i>	571
	Summary	602
	References and Additional Reading	602
<b>Chapter 7</b>	<b>L3VPN Service Deployment: Configuration and Verification Techniques</b>	<b>605</b>
	L3VPN	606
	SRv6 L3VPN Overlay Service	608
	SRv6 L3VPN Full-Mesh Service	610
	SRv6 L3VPN Hub-and-Spoke Service	636
	SRv6 L3VPN Extranet Service	657
	SR-MPLS L3VPN Overlay Service	677
	SR-MPLS L3VPN Full-Mesh Service	688
	SR-MPLS L3VPN Hub-and-Spoke Service	719
	SR-MPLS L3VPN Extranet Service	752
	Route Target Constraint	767
	Route Target Constraint Configuration and Verification	771
	Summary	780
	References and Additional Reading	781
<b>Chapter 8</b>	<b>Service Assurance</b>	<b>783</b>
	Transport	784
	Segment Routing Data Plane Monitoring (SR-DPM)	788
	SR-DPM Configuration and Verification	789
	Path Tracing (PT)	798
	Services	806
	L2VPN Service Assurance	806
	Ethernet Connectivity Fault Management (CFM)	808
	ITU-T Y.1731 Performance Measurement	816



	L3VPN Service Assurance	834
	<i>IPSLA and TWAMP</i>	834
	Summary	855
	References and Additional Reading	855
<b>Chapter 9</b>	<b>High Availability and Fast Convergence</b>	<b>857</b>
	BFD Failure Detection Mechanism	858
	BFD BoB Configuration	865
	BFD BoB Verification	868
	BFD BLB Configuration	872
	BFD BLB Verification	874
	Topology-Independent Loop-Free Alternate (TI-LFA)	878
	Link Protection Configuration	883
	Link Protection Verification: SR-MPLS	885
	Link Protection Verification: SRv6	902
	Node Protection Configuration	914
	Node Protection Verification: SR-MPLS	917
	Node Protection Verification: SRv6	919
	SRLG Protection Configuration	921
	SRLG Protection Verification: SR-MPLS	923
	SRLG Protection Verification: SRv6	931
	Microloop Avoidance	935
	BGP PIC Edge	943
	BGP PIC Edge Configuration: SR-MPLS	948
	BGP PIC Edge Verification: SR-MPLS	951
	BGP PIC Edge Configuration: SRv6	962
	BGP PIC Edge Unipath Verification: SRv6	970
	BGP PIC Edge Multipath Verification: SRv6	981
	Summary	995
	References and Additional Reading	995
<b>Part IV</b>	<b>Business and Operational Considerations</b>	
<b>Chapter 10</b>	<b>Business Opportunities</b>	<b>997</b>
	Technological Opportunities and Benefits	999
	Fewer Protocols	999
	More QoS Options	1001
	SR from the Access Network to the Data Center Network	1003
	Traffic Engineering and Network Slicing	1005

Scale	1007
Routed Optical Networks	1008
<i>Benefit 1: Simplified Long-Distance Connectivity</i>	1008
<i>Benefit 2: Easier and Cost-Effective Scaling</i>	1009
<i>Benefit 3: Simplified Redundancy</i>	1009
Private Line Emulation	1011
Integrated Visibility	1017
<i>Intent-Driven Configuration of Visibility Features</i>	1018
<i>Intent-/Model-Based Assurance</i>	1019
<i>High-Precision Probing</i>	1020
<i>Path Tracing</i>	1023
New Hardware Generation	1025
CapEx Savings	1026
OpEx Savings	1030
Business Case Guidance	1032
Summary	1039
References and Additional Reading	1040
<b>Chapter 11 Organizational Considerations</b>	<b>1043</b>
Scenario 1: Replacing or Enhancing a Legacy Network with SR	1046
Knowledge	1046
Migration Strategy	1049
IT Evolution and Gap Awareness	1051
Scenario 2: Consolidating Networks and Services	1056
Domain Definitions	1056
<i>Domain Architecture Blueprint</i>	1059
<i>Domain Responsibilities and Their Architectural Implications</i>	1067
Domain Organization and Transformation	1074
Existing and New Processes	1083
Service Portfolio Consolidation	1083
Development and Release Methodology	1084
<i>Process with a Clear Flow</i>	1086
<i>Environments</i>	1089
<i>Domain Releases: A Symphony of Component Builds and Release Candidates</i>	1096
<i>Tooling: Embracing Automation for Environment and Process Efficiency</i>	1097
<i>Tooling: Source of Truth</i>	1098

<i>Tooling: Binary Repository</i>	1100
<i>Tooling: Pipelines</i>	1101
Change Management Across Domains	1104
Summary	1106
References and Additional Reading	1107

## **Appendix A Reference Diagrams and Information 1109**

SR-MPLS Reference Network	1109
SRv6 Reference Network	1111
SR Migration Reference Network	1112


## **Index 1115**

### **Online Element:**

## **Chapter 12 SRv6 Ecosystem Deployment Use Cases 1**

SRv6 Open-Source Implementations	2
Linux Kernel	2
Free Range Routing (FRR)	7
Vector Packet Processor (VPP)	9
Software for Open Networking in the Cloud (SONiC)	13
SRv6 Open-Source Lab Deployment Examples	17
Linux SRv6 Deployment	18
Containerlab	20
Linux Underlay Connectivity	24
Linux IPv4 L3VPN Service	27
Linux IPv6 L3VPN Service	32
Linux IPv4/IPv6 L3VPN Service	34
Linux Point-to-Point L2VPN Service	36
VPP SRv6 Deployment	39
Basic VPP Setup	40
VPP Underlay Connectivity	43
VPP IPv4 L3VPN Service	45
VPP IPv6 L3VPN Service	54
VPP Point-to-Point L2VPN Service	57
SRv6 L3VPN Interoperability	61
Free Range Routing IPv4 L3VPN Service	62
Cisco Catalyst 8000V Edge IPv4 L3VPN Service	75
Summary	82
References and Additional Reading	83

## Icons Used in This Book



## Command Syntax Conventions

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

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Braces { } indicate a required choice.

## Introduction

Welcome to the future of MPLS and the realm of advanced networking technologies, where efficiency, scalability, and reliability are paramount. This book is your gateway to mastering segment routing (SR), a revolutionary technology that transforms IP data transport and network operations. From the foundational principles of MPLS to state-of-the-art implementations of SR over MPLS (SR-MPLS) and SR over IPv6 (SRv6), this book offers a comprehensive guide that bridges the gap between theory and practice.

The chapters cover the entire spectrum of SR, providing a holistic understanding of the technology. They feature practical examples for SR on both IOS XR and IOS XE platforms, ensuring that you have the knowledge to implement SR in different network environments.

This book also goes beyond technical details. It delves into the business opportunities and organizational implications of adopting SR, offering valuable insights into how SR can drive growth, improve customer experience, and streamline operations. Dedicated sections on the SRv6 ecosystem in data centers and cloud environments showcasing network functions virtualization (NFV) prepare you for the next wave of networking innovations.

Available online content enables you to gain hands-on practice and reinforce the theories covered. A business case template provides a tool to legitimize investments in SR technologies and calculate potential returns.

By the end of this book, you'll be equipped with the knowledge and tools to implement and manage SR technologies effectively, helping you stay ahead in this ever-evolving field.

## What Sets This Book Apart

What distinguishes our book is its unique blend of content, offering readers an unparalleled experience that combines theoretical aspects of segment routing with hands-on experience from actual deployments.

- **Practical insights:** Drawing from real-world experiences, particularly our collaboration with Swisscom, this book provides real-world insights that bridge the gap between theory and practice. In this book you will find detailed configurations, design guidelines, and troubleshooting tips that are directly applicable to your work environment.
- **Step-by-step approach:** The content of this book is structured to guide you through a logical progression of information, from basic concepts to advanced implementations, making it suitable for both beginners and seasoned professionals. Each chapter builds on the previous one, providing a smooth learning curve.
- **Future-proofing:** With a dedicated section on the future of SRv6, featuring open-source SRv6 NFV implementations that can fit in data centers and cloud environments, this book prepares you for the next wave of networking innovations.

- **Interactive learning:** Downloadable content provides lab topology definitions, configurations, scripts, and templates to help you set up your own labs and experiment with them. Additionally, a lab guide is available, offering the option to run your lab in the cloud.
- **Beyond technology:** In addition to technical content, the business-oriented chapters outline the benefits of adopting SR and offer guidance in various areas to help leadership and teams smoothly transition to SR.

## Goals and Methods

The primary goal of this book is to provide a thorough understanding of SR by offering detailed explanations, configuration examples, verification hints, and packet captures. It walks you through foundational concepts and practical implementations.

Two reference lab topologies, one for SR-MPLS and one for SRv6, are consistently referenced throughout the technical chapters. To make the learning process interactive and engaging, the downloadable lab support material enables you to set up your own lab so you can replicate and apply the described theory.

## Who Should Read This Book?

This book is designed for network engineers, architects, and operators who are involved in the design, deployment, and management of modern networking infrastructures. It is also valuable for IT professionals, students, and researchers who wish to deepen their understanding of SR technologies. Business leaders and decision makers will find the chapters on new service opportunities and what to consider within an organization on the journey to SR particularly insightful.

## How This Book Is Organized

The chapters in this book guide you from basic concepts to advanced implementations, ensuring a logical progression of information. A recap of MPLS is followed by an in-depth exploration of SR, detailed configurations, migration strategies, service assurance, and high availability. The final chapters provide a glimpse into the future of SRv6 in data centers and cloud environments, explore business opportunities that justify investment in this new technology, and offer non-technical thoughts to streamline the transition to SR.

This book is broken into 12 chapters and an appendix:

- **Chapter 1, “MPLS in a Nutshell”:** This chapter provides an introduction to MPLS technology and its significance in modern networking as well as an overview of MPLS mechanisms and benefits.
- **Chapter 2, “What Is Segment Routing over MPLS (SR-MPLS)?”:** This chapter provides a basic introduction to the general concepts of segment routing (SR) and explores SR-MPLS in the control plane and data plane.

- **Chapter 3, “What Is Segment Routing over IPv6 (SRv6)?”**: This chapter explores SRv6 in the control plane and data plane and provides an overview of the evolution and simplification of SR-driven networks.
- **Chapter 4, “Segment Routing in Detail”**: This chapter includes detailed information on configuration and verification of SR on Cisco devices and describes design guidelines and advanced features for SR networks.
- **Chapter 5, “Migrating to Segment Routing”**: This chapter is a practical roadmap for migrating from MPLS to SR and presents strategies for greenfield and brownfield deployments.
- **Chapter 6, “L2VPN Service Deployment: Configuration and Verification Techniques”**: This chapter provides an overview of customer-related L2VPN services implemented on an SRv6 underlay network. It includes basic configuration methodologies and service verifications.
- **Chapter 7, “L3VPN Service Deployment: Configuration and Verification Techniques”**: This chapter provides an overview of customer-related L3VPN services implemented on SRv6 and SR-MPLS underlay networks. It includes basic configuration methodologies and service verifications.
- **Chapter 8, “Service Assurance”**: This chapter presents procedures and processes for improving customer experience and satisfaction. It includes a discussion of tools and protocols for SLA monitoring and fault management in the transport network layer, as well as in the L2VPN and L3VPN service overlays.
- **Chapter 9, “High Availability and Fast Convergence”**: This chapter introduces technologies and features for high availability and fast convergence in SR networks. It includes a detailed discussion of failure detection, path computation, and network convergence.
- **Chapter 10, “Business Opportunities”**: This chapter describes the benefits of investing in SR and why network service providers should transition to this technology. It offers insights into how SRv6 can offer substantial opportunities and advantages over SR-MPLS as well as business case calculation guidance.
- **Chapter 11, “Organizational Considerations”**: This chapter examines the impact of SR on organizational structures and processes and provides a guide to managing the transformation to a programmable SR network.
- **Chapter 12, “SRv6 Ecosystem Deployment Use Cases”**: This online only chapter discusses the potential of SRv6 in data centers and cloud environments. It provides examples and interoperability scenarios involving open-source software.
- **Appendix A, “Reference Diagrams and Information”**: This appendix describes the reference diagrams and information of the SR-MPLS, SRv6, and SR migration network topologies used throughout the book in a single location for the reader's convenience.

## Downloadable Content

Readers can access downloadable content using the companion website as per the instructions below:

1. The user enters [ciscopress.com/sr](http://ciscopress.com/sr) in his browser or clicks the hyperlink in the online book version.
2. The user completes the registration/login process.
3. The user confirms the already prepopulated ISBN number of the book and answers a proof-of-purchase challenge question, to access additional content.
4. The user clicks on the desired attachment.

The following attachments can be downloaded for use with this book:

- **SR-MPLS-Reference-Configuration-Lab.zip:** These files, which can be used across the entire book, include a reference diagram and configurations to deploy an SR-MPLS and services lab.
- **SRv6-Reference-Configurations-Lab.zip:** These files, which can be used across the entire book, include a reference diagram and configurations to deploy an SRv6 and services lab.
- **SRv6-Online-Lab-Guide.pdf:** This guide provides clear instructions on how to use an online lab, offering you the flexibility to run your SRv6 lab in the cloud.
- **SRv6-Migration-Lab.zip:** These files, which are meant to be used with Chapter 5, include a reference diagram and configuration for migration use cases.
- **SRv6-Linux-Lab.zip:** These files, which are meant to be used with Chapter 12, include a container-based lab topology definition and initialization script required to spin up the SRv6 Linux lab topology.
- **SRv6-VPP-Lab.zip:** These files, which are meant to be used with Chapter 12, include a bash script to spin up the SRv6 VPP topology, VPP instances, and startup configurations.
- **SRv6-Interop-Lab.zip:** These files, which are meant to be used with Chapter 12, include a Cisco CML topology definition and a running configuration of PE1 (IOS XR), P (IOS XE), and PE3 (IOS XE), as well as FRR settings and configuration for PE2 (FRR).
- **Segment-Routing-Business-Case-Template.xlsx:** This file, which is meant to be used with Chapter 10, includes a Microsoft Excel-based business case template to help you create your case for SR in your organization. Please review the provided sample data and update all cells that have a yellow background to reflect your specific information.

**Note** This book contains references to the companion website in later chapters which leverage the previously listed downloadable content.



*This page intentionally left blank*

## What Is Segment Routing over MPLS (SR-MPLS)?

We took a brief look at MPLS and its shortcomings in Chapter 1, “MPLS in a Nutshell.” Now it is time to build a solid understanding of segment routing (SR) so you will be ready for the upcoming chapters, which cover high-level design, configuration, and verification of various transport- and service-related aspects of SR-enabled networks. This chapter introduces basic segment routing concepts by using an analogy and then goes into the theory behind the MPLS data plane encapsulation implementation. This chapter covers Segment Routing for MPLS (SR-MPLS), and Chapter 3, “What Is Segment Routing over IPv6 (SRv6)?” covers IPv6 (SRv6) data plane encapsulations. The terms, abbreviations, and acronyms introduced in this chapter are consistently used throughout the remainder of this book.

Before delving into the more technical specifications of segment routing, let’s consider a simplified high-level analogy that serves as an example to explain underlying key concepts. The central processing unit (CPU) installed in an everyday device, such as a mobile phone, smart TV, laptop, or router is the brain of the system that controls other components, such as memory, hard disk, and a network interface card. The main task of the CPU is to execute program instructions in the form of machine code. Machine code is platform-specific binary code consisting of zeros and ones that is not human readable. Machine code for a given program is not portable between processor architectures; for example, ARM64 architecture-based machine code is not compatible with and cannot be run on x64 architecture-based devices and vice versa. You could think of it as two gingerbread recipes, one written in English and one in Bahasa Indonesian, each providing a list of instructions. While the Indonesian alphabet uses the same 26 letters as the English alphabet, a native English speaker will not be able to read or follow the recipe written in Bahasa Indonesian.

High-level programming languages such as Python, Java, C++, and Go allow programmers to write code that is independent of the underlying hardware architecture and human readable and that provides an abstraction layer to hide low-level hardware details. For instance, Example 2-1 shows a simple computer program that allocates a few variables,

stores the sum of  $a + b$  in a variable, and sends the result to the standard output (that is, the user's screen in the terminal).

**Example 2-1** *High-Level C++ Source Code*

```
#include <stdio.h>

int main(void){
    int a,b,c;
    a=1;
    b=2;
    c=a+b;
    printf("%d + %d = %d\n",a,b,c);
}
```

A compiler is a special program that translates high-level programming language source code into machine code that can be executed on a CPU. As an intermediate step, a compiler creates assembler code, which is one step away from machine code. Unlike machine code, assembler code is human readable and nicely shows the order of instructions that must be executed by the CPU to achieve the specified outcome of the high-level source code. Example 2-2 shows the same program from Example 2-1 but in assembler code.

**Example 2-2** *Low-Level Assembler Source Code*

```
.LC0:
.string "%d + %d = %d\n"
main:
push rbp
mov rbp, rsp
sub rsp, 16
mov DWORD PTR [rbp-4], 1
mov DWORD PTR [rbp-8], 2
mov edx, DWORD PTR [rbp-4]
mov eax, DWORD PTR [rbp-8]
add eax, edx
mov DWORD PTR [rbp-12], eax
mov ecx, DWORD PTR [rbp-12]
mov edx, DWORD PTR [rbp-8]
mov eax, DWORD PTR [rbp-4]
mov esi, eax
mov edi, OFFSET FLAT:.LC0
mov eax, 0
call printf
mov eax, 0
leave
ret
```

The assembler program consists of a list of instructions whose machine code counterparts will be executed one by one by the CPU at runtime. A special CPU register, generally referred to as the program counter, stores the memory address of the current instruction. Upon completion, the program counter is incremented, and the next instruction is fetched from the updated memory address to be executed. In other words, the program counter keeps track of where the CPU is in the program execution—that is, where it is in the sequence of instructions.

Don't worry if you don't understand the assembler program. The details are not relevant. What is relevant is the fact that there are different instructions, such as **push**, **mov**, **sub**, and **add**, that seem to accept one or more parameters. The supported instructions vary between hardware architectures and CPU models. The instruction set architecture (ISA) defines which instructions can be used by a software program to control the CPU. Reading such a manual reveals that instructions have the following format:

```
label: mnemonic argument1, argument2, argument3
```

where:

- *label* is an identifier (not related to MPLS labels).
- *mnemonic* is a name for a class of instructions that have the same function.
- Arguments are mandatory or optional, depending on the mnemonic.

Example 2-2 shows a label called **main**, followed by **push** (*mnemonic*) and **rbp** (*argument1*). This instruction tells the CPU to store a special register on the stack, whereas the **add eax, edx** instruction takes two arguments to perform the addition of  $a + b$  in the source code. This simple program uses common instructions, but applications in the field of artificial intelligence (AI) and machine learning (ML) use more complex and specialized instructions. In principle, there are no limits on what kind of instructions a CPU can execute, as long as it is implemented in hardware and there is a practical benefit of implementing it. The length of an instruction may vary within an ISA, depending on the underlying hardware architecture.

Finally, executing the binary yields the output shown in Example 2-3.

### Example 2-3 Output of Program Execution

```
cisco@ubuntu-server:~/Code$ ./program
1 + 2 = 3
cisco@ubuntu-server:~/Code$
```


At this point, you might wonder about the relevance of CPU instructions, program counters, and instruction formats in a segment routing book. The coming paragraphs shed light on the analogy and emphasize similarities between computer and segment routing architectures.

Segment routing (RFC 8402) leverages source routing, which allows the source node (ingress PE node) to steer a packet flow through the SR domain. This ability is a key

difference from traditional MPLS-based networks, where ingress PE nodes lack such fine-grained control over the traffic path through the network when relying on LDP labels. Traffic engineering (TE) techniques enable the optimization of traffic distribution in MPLS networks at the cost of additional protocols such as Resource Reservation Protocol (RSVP) and network state information (TE tunnels) in the network, which is challenging to operate and negatively impacts the overall network scale. Segment routing significantly simplifies the network protocol stack by superseding signaling protocols like LDP or RSVP-TE.

Instead, SR extensions elevate the underlying link-state routing protocol, providing a comprehensive view of the network topology across the entire domain, to provide the same functionality that relied on multiple protocols in the past. The interior gateway protocol (IGP) advertises *segments*, which are essentially network instructions, throughout the network, which guarantees that every node within the domain has the same view. The flooding of segments enables the IGP to replace the previously mentioned signaling protocols and facilitates moving any tunnel state information from the network to the packet headers. A segment can have global significance within the network, such as instructing nodes in the SR domain to steer traffic to a specific node, or local significance, such as instructing a specific node to steer traffic across a specific interface.

Figure 2-1 shows the two supported data planes of the segment routing architecture. SR-MPLS reuses the MPLS data plane, whereas SR IPv6 (SRv6) relies on the IPv6 data plane.



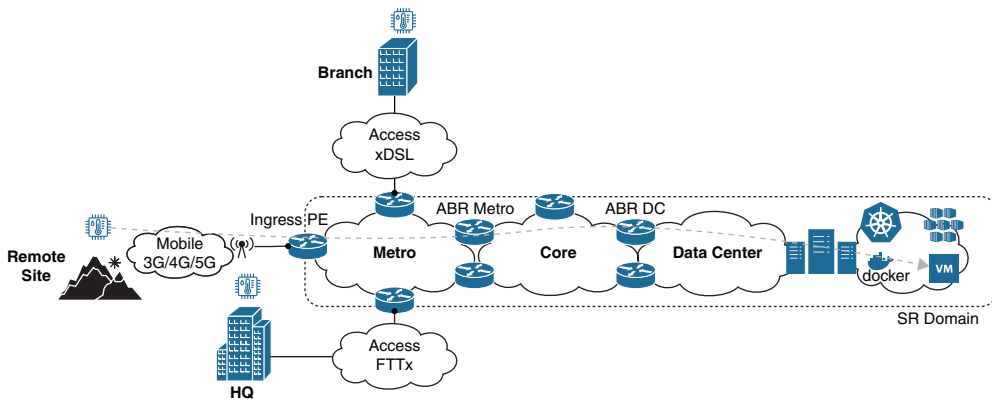
**Figure 2-1** Segment Routing Data Planes

As previously mentioned, a segment represents a single instruction identified by a segment identifier (SID). The length of a SID depends on the underlying data plane. For SR-MPLS, the SID is 20 bits long and is written in the Label field of the MPLS header. In contrast, an SRv6 SID is a 128-bit identifier in the Destination Address field of the IPv6 header. As with the assembler program shown in Example 2-2, multiple ordered instructions can be expressed as a list of segments. A list of segments can be realized using multiple SIDs, which in the MPLS data plane results in a label stack. In the SRv6 data plane, a list of segments may be encoded using the segment routing header (SRH), a micro-SID

(uSID) carrier, or a combination of both, depending on the SRv6 flavor and the number of segments. The fundamental terminology of segment routing is agnostic to the underlying data plane; the concept of a segment, SID, and list of segments applies to both encapsulation types.

**Note** SRv6 terms and concepts, such as SRH and uSID, are explained in Chapter 3. The different segments in SR-MPLS are presented in more detail later in this chapter, in the section “Segment Routing for MPLS (SR-MPLS).”

You may have come across the term *network as a computer* in the context of segment routing, in reference to the network as a large distributed system where several devices work together to execute a network program consisting of a list of instructions or segments. All nodes within an SR domain must speak a common language to be able to interpret the segments correctly. SR can be applied to both the MPLS and IPv6 architectures, which means that nodes within an SR domain are not limited to networking devices if they understand the underlying data plane. This is especially true for IPv6, which is widely supported across a range of different networking nodes from the Internet of Things (IoT) in the industry to containers in the data center. Figure 2-2 shows an imaginary local weather station with sensors in three different locations and some services running in a data center (DC).



**Figure 2-2** Weather Station Network Topology

The sensors connect to a local service provider (SP) using different access technologies. Each sensor measures temperature, humidity, and barometric pressure on a regular basis and transmits the data to a microservice hosted in a remote data center (on the right-hand side of the figure). The collected data is processed, stored, and evaluated every 24 hours to provide the weather forecast for the next seven days. It goes without saying that meteorology is far more complex than presented here, but this illustration will suffice for our example.

The SR domain in our example includes metro, core, and data center, up to and including the virtual machine or container, which means that segments could be executed by any of the nodes belonging to the SR domain. Within an SR domain, different roles can be distinguished:

- **Source/ingress node:** Handles the traffic as it enters the SR domain.
- **Transit node:** Handles the forwarding of traffic within the SR domain.
- **Endpoint/egress node:** Handles the traffic as it leaves the SR domain.

In traditional Layer 3 virtual private network (VPN) services, service provider and customer networks are isolated logically using virtual routing and forwarding (VRF) instances or access lists on the service edge to protect the SP infrastructure. Consequently, VRF instances or access lists are used to enforce the demarcation point of the SR domain. In our example, all three weather station sensors are isolated from the SP through VRF instances on the PE node, which means a network program can only be initiated by the ingress PE device receiving customer traffic.

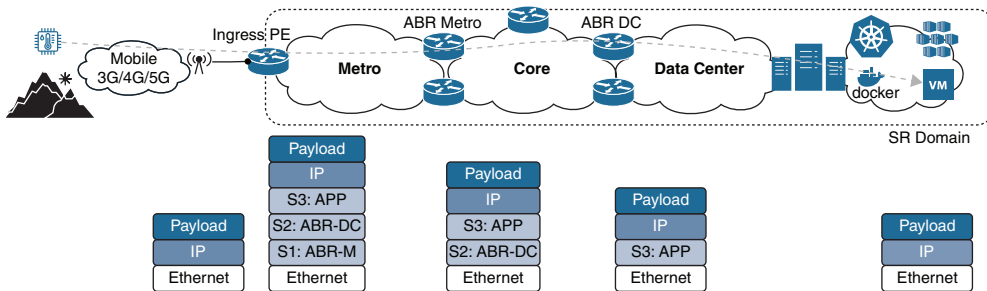
Unlike in traditional software development, with segment routing there are no high-level network programming languages available. Instead, a network program is defined as an ordered list of segments, also known as an *SR policy*, that steers a packet flow along a desired path in the network. SR policies are source-routed policies identified through the tuple, such as headend, color, or endpoint. Headend and endpoint should be self-explanatory; the 32-bit color value identifies the intent or objective of the policy. The endpoint and color are used as identifiers to steer traffic into the corresponding SR policy. Examples of such an intent are low latency or MACsec encrypted paths from the headend to the endpoint. The source routing is crucial in moving the traffic engineering tunnel state from intermediate routers to the packet headers imposed by the ingress node through an SR policy.

Complementary information on how to implement such traffic engineering capabilities using the IGP is provided in the section “IGP Flexible Algorithm (Flex Algo) (RFC 9350),” later in this chapter. Example 2-4 shows an imaginary SR policy that defines a loose path from the ingress PE node (source node) to the container (endpoint node) hosting the weather application via two transit nodes. Note that there are two area border routers (ABRs) in the metro and the data center, which may result in equal-cost multipath routing (ECMP). If desired, a more restrictive path could be defined, such as using a specific ABR or only traversing the core over MACsec-encrypted links.

#### **Example 2-4** *Network Program Pseudocode*

```
policy weather-app-policy
1 goto ABR Metro
2 goto ABR DC
3 goto container weather-app
```

Figure 2-3 shows the ordered list of segments expressed in this pseudocode.



**Figure 2-3** Network Program Segment Routing Policy (SR-MPLS)

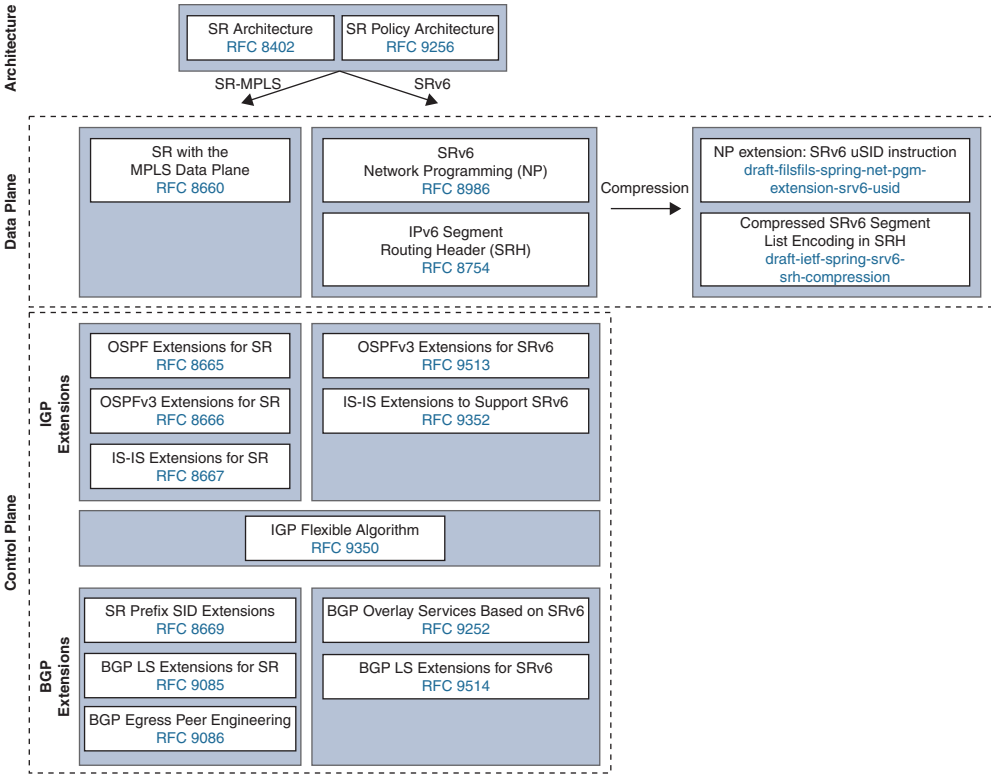
The ingress PE node imposes one or more additional headers to encapsulate the original customer packet. Note that the exact headers depend on the underlying data plane, as discussed in detail in the section “Segment Routing for MPLS (SR-MPLS),” later in this chapter, and in Chapter 3. The additional encapsulation overhead is negligible in most cases and justified by the significant scalability gains in the backbone network achieved by transferring the tunnel state information from the network to the packet. In the case of SR-MPLS, the length of the list of segments decreases as the network program is executed. The first segment is executed by one of the ABR metro nodes. The metro ABR pops its own instruction from the stack and forwards the packet toward an ABR DC, which pops its own instruction and forwards the packet toward the weather-app container. Eventually, the packet reaches its destination, which in our example is the SR-aware weather-app container that decapsulates and processes the inner IP packet. Note that this example excludes a few details, such as penultimate hop popping (PHP) and the BGP service label for simplicity.

It should be becoming clear now that the execution of segments in a segment routing domain and the execution of instructions in computer architectures share several fundamental principles. In fact, those similarities are even more prominent with SRv6, as you will see in Chapter 3, which covers the Segments Left field of the SRH and the SRv6 SID format that are comparable to the program counter and instruction format, respectively.

The segment routing ecosystem encompasses a wide variety of Internet Engineering Task Force (IETF) standards and drafts across numerous working groups. The standardization process for segment routing has been progressing at an impressive pace, and most key drafts have become proposed standards. One exception worth highlighting here is the SRv6 compression drafts that are in the later stages of the standardization process. The successful mass-scale rollout of SR lead operators shows that there is no reason to delay the SR adoption.

Figure 2-4 displays a selection of the most important building blocks that make up segment routing (RFC 8402) and the segment routing policy architecture (RFC 9256).





**Figure 2-4** *Segment Routing Architecture*

**Note** Some of the official standards and draft titles have been shortened or modified in Figure 2-4 for better readability. This book covers many of the drafts and proposed standards in more detail, accompanying the somewhat dry theory with visual illustrations and packet captures.

## What Is Segment Routing over IPv6 (SRv6)?

### Introduction

This chapter covers the theory behind the Segment Routing over IPv6 (SRv6) data plane encapsulation implementation. It provides valuable decision-making process inputs and outlines potential pitfalls when evaluating the network architecture evolution to Segment Routing over IPv6. In addition, the section “SR-Powered Network Evolution” describes the network evolution journey that began with the introduction of Segment Routing for MPLS (SR-MPLS) and ends with a converged SDN transport network based on SRv6.


### Segment Routing over IPv6 (SRv6)

This section introduces SRv6, which shares many fundamental concepts with SR-MPLS. Although some operators might view SR-MPLS as a transitional step toward SRv6, this chapter shows that SRv6 is a superior solution that effectively addresses the challenges associated with MPLS discussed in the section “Challenges and Shortcomings of MPLS,” in Chapter 1, “MPLS in a Nutshell.” While SR-MPLS is already well established, a select number of compression-related SRv6 extensions are still in the process of being standardized as of this writing. The IETF has been advancing at an impressive pace, and all the major key drafts have been successfully standardized, achieving RFC status. This standardization marks a significant milestone in the evolution of SRv6, showcasing its readiness for widespread deployment and the promise of enhanced network efficiency.

Since SRv6 relies on the IPv6 data plane, it is crucial to have a solid understanding of IPv6 encapsulation and the IPv6 header. In fact, as you will see in this chapter, the vast majority of SRv6 use cases rely on IPv6 routing using an IPv6 header without any extension headers.

### IPv6 for SRv6 Recap

Figure 3-1 shows the format of the IPv6 header, as specified in RFC 8200.



**Figure 3-1** *IPv6 Header*

The IPv6 header consists of the following fields:

- **Version (4 bits):** Specifies the version of IP; set to 6 for IPv6.
- **Traffic Class (8 bits):** Used for traffic management (QoS) based on DSCP (6-bit) and ECN (2-bit).
- **Flow Label (20 bits):** Used for encoding entropy of the payload and subsequent flow hashing (load balancing).
- **Payload Length (16 bits):** Specifies the length of the payload following the IPv6 header.
- **Next Header (8 bits):** Identifies the header following the IPv6 header (for example, IPv4, IPv6, Ethernet, ICMP, TCP).
- **Hop Limit (8 bits):** Equivalent to the Time to Live (TTL) field of the IPv4 header.
- **Source Address (128 bits):** Identifies the source of the packet.
- **Destination Address (128 bits):** Identifies the destination of the packet.

Most of the fields are self-explanatory or easy to grasp. However, special attention should be paid to the Traffic Class, Flow Label, and Next Header fields.

The Traffic Class field is used for quality of service (QoS) marking, which involves Differentiated Services Codepoint (DSCP) and Explicit Congestion Notification (ECN). The 6-bit value of DSCP covers the decimal range from 0 to 63, which makes it possible to distinguish more than eight traffic classes. The 3-bit value of MPLS EXP in SR-MPLS is a significant limitation with SR-MPLS.

The Flow Label field facilitates efficient flow classification in combination with other IPv6 header fields, such as the Source Address and Destination Address fields. A

sequence of packets belonging to the same Layer 3 flow are generally classified based on the 5-tuple of network addresses, transport protocol, and transport ports. Note that not all of those identifiers may exist in a flow, depending on the payload (for example ICMP), or they may be unavailable due to encryption or fragmentation. Layer 2 traffic flows usually take into account data link layer information for classification and may or may not include some of the higher-layer protocol information. Often, flow classification is not only vendor dependent but also platform dependent, with some devices supporting 7 or more-tuple flow classification, taking into account one or more MPLS labels.

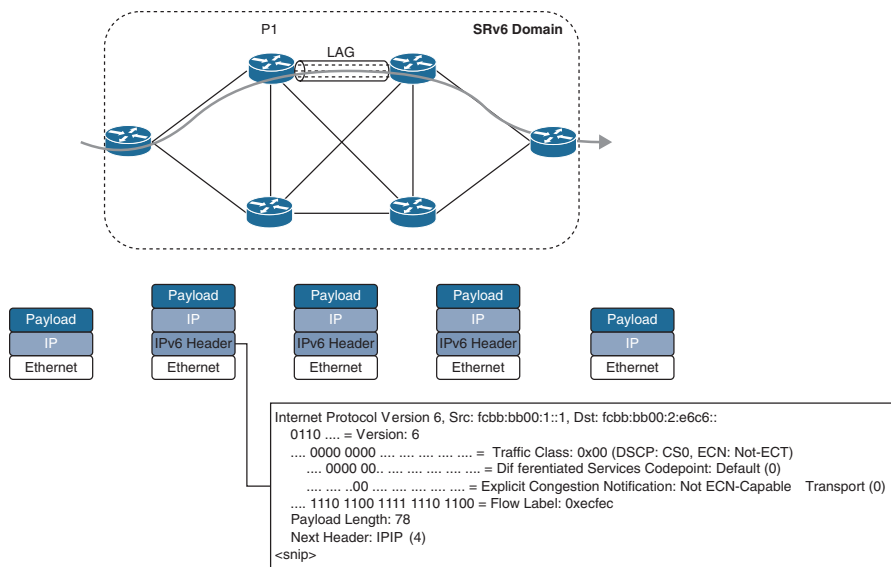
The Flow Label field is a radical simplification for IPv6 compared to IPv4 or MPLS. Instead of cumbersome inspecting a packet and trying to figure out where the relevant fields are located within the packet to extract the 5+-tuple, IPv6 uses a 3-tuple consisting of Source Address, Destination Address and Flow Label fields. Having all those fields at fixed positions within the IPv6 header simplifies the extraction of flow identifiers and consequently the hardware implementation of this process.

The Flow Label value is computed by the source node and not changed by transit nodes along the path. The source node in an SRv6 domain is usually an ingress PE device, which encapsulates the received packet coming from the edge into an outer IPv6 header with an optional segment routing header (SRH) extension header. The exact algorithm to compute the hash for the Flow Label value is implementation specific and may differ between vendors or platforms. However, depending on the type of service, different fields are considered, such as the following common identifiers:

- **Layer 2 VPN service:** Source and destination MAC addresses and source and destination IP addresses (IP payload only)
- **Layer 3 VPN service:** Source and destination IP addresses, transport protocol, and source and destination ports

This list is an example, and different implementations may consider additional fields. It is important to understand that the Flow Label value is computed only once in the network, at the source node, which is service aware; that is, Layer 2 or Layer 3 VPN services can be easily distinguished, and the tuple used for hashing can be extracted before additional encapsulation takes place. After the hash has been computed, it is written to the Flow Label field, which is, in turn, used by all transit nodes. In essence, the hard work of computing a proper hash needs to be performed only once by the source node, and all other nodes along the path can take advantage of this hash, which greatly reduces the complexity of flow classification to achieve proper load balancing for ECMP routing or LAG hashing.

Figure 3-2 shows an example of a traffic flow entering an SRv6 domain. The network is highly symmetric, with one core link relying on a link aggregation group (LAG). The ingress PE device encapsulates the received packet from the edge into an IPv6 header and populates the Flow Label field with the computed 20-bit hash (0xecfec). The packet entering the SRv6 domain is an IPv4 packet, as you can see from the Next Header field of the IPv6 header.



**Figure 3-2** IPv6 Flow Label and Capture

**Note** Many of the network figures related to MPLS/SR-MPLS in this book include an inner IP header with payload to emphasize the fact that MPLS sits between Layer 2 (the data link layer) and Layer 3 (the network layer) in the OSI model, which is why it is sometimes referred to as a Layer 2.5 networking protocol. SRv6 goes back to the roots of the OSI model and no longer relies on this shim layer. SRv6 network figures in this book are generally drawn with an inner payload only, unless the context asks for more detailed inner header information, such as IPv4/IPv6 or Ethernet.

The flow label must not change en route, and P1 classifies the flow based on the IPv6 source address, destination address, and flow label. P1 chooses the ECMP path and physical link of the LAG based on the hash of the IPv6 header, as shown in Figure 3-2.

The Next Header field identifies the upper layer protocol, which follows immediately after the IPv6 header. A key difference between IPv4 and IPv6 is the flexible support for extensions and options in IPv6, where *extension headers* are placed between the IPv6 header and upper layer protocols (for example, TCP or UDP).

**Note** IP protocol numbers and IPv6 extension headers registered with IANA are listed at <https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml> and <https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml#extension-header>.

The routing header for an IPv6 extension is a central puzzle piece of the SRv6 solution as it allows the insertion of an optional segment routing header (SRH) after the IPv6 header.

The details of the SRH are introduced later in this chapter, in the section “IPv6 Segment Routing Header (SRH) (RFC 8754).”

## SRv6 Network Programming (RFC 8986)

RFC 8986 lays the foundation of the segment routing architecture in the IPv6 data plane. Network instructions have to be encoded into the IPv6 header, which differs fundamentally from MPLS, where each instruction is represented by a label. Many fundamental concepts covered in Chapter 2, “What Is Segment Routing over MPLS (SR-MPLS)?” are the same for SRv6, though. An SRv6 SID is still associated with a segment, but instead of using an MPLS label, it is now represented as an IPv6 address. The IPv6 destination address in the outer IPv6 header is set to the SRv6 SID, which represents a network program, including a single instruction or an SR policy with a single segment. SRv6-unaware transit nodes forward an SRv6 SID based on the longest-prefix-match lookup on the IPv6 destination address. An IPv6 address associated with an SRv6 SID has a special format.

### SRv6 Segment Identifier (SID)

SRv6 SIDs are 128-bit long IPv6 addresses that follow the format shown in Figure 3-3:



**Figure 3-3** SRv6 SID Format

- **Locator:**
  - Most significant bits
  - Routable part, which points to the parent node that instantiated the SID
  - Advertised through a link-state IGP (IS-IS or OSPF)
  - Should be unique within the SRv6 domain, except for SRv6 anycast locators
- **Function:**
  - Identifies a locally significant behavior of the parent node
- **Arguments:**
  - Least significant bits
  - One or more input arguments to the function (for example, service or flow information)
  - Optional

The length of the Locator (L), Function (F), and Arguments (A) fields are flexible as long as the total length is less than or equal to 128 bits. If the total length is less than 128 bits, the SID should be padded to 128 bits with zeros. As mentioned previously, the Arguments field is optional.

As shown in Figure 3-4, the Locator field may be expressed as two different fields: SID block (B) and Node ID (N). A common format for early SRv6 deployments was to allocate B::/48 for the SID block and B:N::/64 for the locator. In Cisco documentation, this is commonly referred to as *base format*. The section “SRv6 Locator Addressing Scheme,” later in this chapter, discusses alternative locator assignments suited for large-scale deployments. For now, this basic split into SID Block, Node ID, and Function fields will suffice as an introduction to SRv6 SIDs.




Figure 3-4 SRv6 SID Format (Simplified)

Let us look at the example presented in Figure 3-5:




Figure 3-5 SRv6 SID Format (Example)

- A service provider allocates the SRv6 SID block—for example fcbb:bbbb:bb00::/48—from the unique local address (ULA) space for SRv6 locators in the network shared by all SRv6 nodes.
- Router 10 would be assigned the locator fcbb:bbbb:bb00:000a::/64 (L = 64 bits).
- Router 10 locally allocates function 0x0001 (F = 16 bits) without any arguments (A = 0 bits) for its SRv6 SID. The sum of L + F + A equals 80, which means that the remaining 48 bits must be padded with zero since SRv6 SIDs are 128-bit addresses.
- The resulting SRv6 SID associated with the segment to Router 10 equals fcbb:bbbb:bb00:a:1::.

**Note** IPv6 addresses are expressed in hexadecimal, where 10 translates to 0xA and not 0x10, which would be 16.

For reachability and backward compatibility between SRv6-capable and IPv6-capable nodes, SRv6 nodes advertise the locator (for example, /64) as an IPv6 prefix in the link-state IGP, as shown in Figure 3-6. The locator prefixes act as aggregate prefixes for source and transit nodes to perform longest-prefix-match lookups on SRv6 SIDs and forward the packets accordingly.

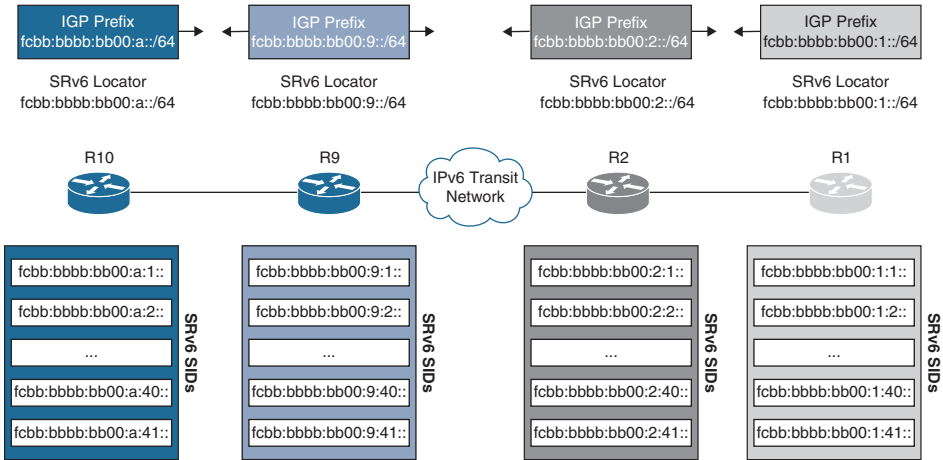


Figure 3-6 SRv6 Locator and IGP Prefix

Therefore, it is possible to provision SRv6 services over a native IPv6 network as long as the service edge is SRv6 capable. Obviously, certain SRv6 functionality, such as Topology Independent Loop-Free Alternate (TI-LFA) or traffic engineering, will not be available on native IPv6 transit nodes. This is discussed in more detail later in this chapter, in the section “IPv6 Segment Routing Header (SRH) (RFC 8754).” At this point, it may still be confusing how SRv6 SIDs can be used to provision services. A simplified analogy using familiar concepts from SR-MPLS will hopefully shed some light. Figure 3-7 shows the data plane encapsulation for both SR-MPLS and SRv6. It should be clear by now that SRv6 does not use labels, so how are transport and service identifiers encoded using a single IPv6 header?

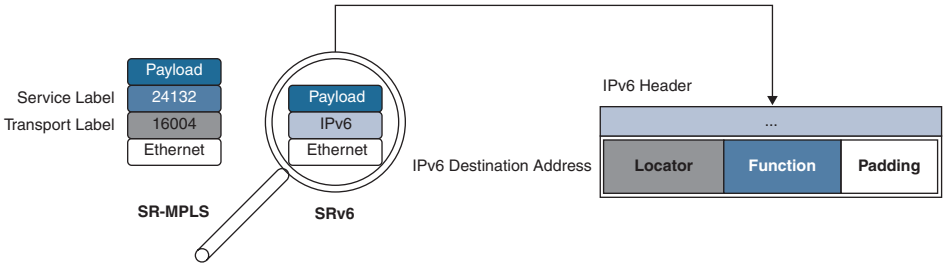


Figure 3-7 SR-MPLS and SRv6 SID Analogy



## Migrating to Segment Routing

Network operators face the ongoing challenge of staying competitive, which often requires proactively adopting new technologies. This brings us to the concept of *migration*, which involves shifting to newer technological landscapes. But what does migration really entail? It is a concept that is deceptively simple to state yet complex to execute.

At its core, migration involves moving from an existing network configuration to an enhanced, desired state. This journey unfolds a tapestry of complexity as it demands adaptations in hardware, software, configurations, IT systems, personnel, and business processes to accommodate the new technology. For network operators, especially those managing large-scale or diverse networks, the transformation is multifaceted and challenging. It can range from deploying entirely new infrastructures and overhauling networks to replacing outdated equipment, upgrading existing hardware, implementing new software solutions, and embracing novel protocols.

With the introduction of two segment routing (SR) technologies and their advantages in the previous chapters, this chapter provides a practical roadmap for implementing SR, whether you're starting afresh with a greenfield approach, or using a brownfield method, which involves integration with existing systems. This chapter presents simple steps to mitigate the risks associated with transitioning services from MPLS to SR.

This chapter covers the following topics:

- Deployment models and strategies for achieving connectivity between MPLS and SR networks during the migration phase, which is a critical period during which a network incrementally adopts SR
- Specifics of migrating from an LDP network to SR-MPLS and details of both deployment models and MPLS and SR interconnectivity options
- Three distinct strategies for migrating from MPLS to SRv6
- A roadmap for migration to an SRv6 network for four different MPLS networks

**Note** In this chapter, *MPLS* refers to both LDP and SR-MPLS, and *SR* denotes SR-MPLS and SRv6.

The migration process typically concerns the transition of customer edge (CE) or provider edge (PE) devices from one network environment to another. Although this chapter uses IPv4 to illustrate service migrations, the concepts are equally applicable to IPv6 services.

**Note** This chapter does not cover the decommissioning of legacy networks after a migration, although that is an important step.

## Deployment Models

*Greenfield network deployment* refers to the installation of a network where previously there was none. This term is derived from the construction industry, where new development on previously undeveloped land is termed a *greenfield development*. An important advantage of greenfield deployment is the opportunity to implement cutting-edge technology solutions from the ground up, free from the constraints or dependencies of existing infrastructure, software, biases, or business processes. In the context of SR, a greenfield deployment refers to building a separate SR network and then migrating the services from the MPLS network to the SR network. During the migration phase, services may connect with both the SR and MPLS networks, and so interworking is essential to maintain seamless service connectivity across the networks. Implementing interworking might necessitate additional hardware and software, which can be phased out after the completion of the service migration.


In contrast to a greenfield deployment, a *brownfield deployment* involves an upgrade or expansion of an existing network. This type of deployment involves installation and configuration of new hardware or network technology that is designed to coexist with the legacy network. One benefit of brownfield development is the ability to enhance existing technology solutions within established business processes; in addition, an organization can avoid extra capital expenditure on new infrastructure by undertaking a brownfield development. However, brownfield projects also come with their own challenges, including the need for a comprehensive and accurate understanding of the existing network's limitations and issues with legacy infrastructure that can potentially slow the development process and inflate overall costs.

A brownfield deployment involves activating SR in the MPLS network. This process can be conducted in phases during several maintenance windows. Throughout these periods, SR and MPLS coexist within the network, which introduces increased complexity in terms of features and configurations. However, it also reduces the effort and mitigates risks typically associated with maintenance windows.

## Migration Strategies


Transitioning the services to SR-MPLS or SRv6 in a network can be achieved through a multitude of viable methodologies, and this section illustrates several strategies. When embarking on a greenfield deployment and establishing a new SR core alongside metro and access layers, there are two principal strategies that stand out for migrating services from MPLS to SR network. The first strategy relies on the ability of the PE routers to simultaneously support MPLS and SR during the migration phase, and the second strategy depends on interworking.

Figure 5-1 illustrates the first strategy, which leverages an interworking gateway to sequentially transition the MPLS PE devices to the SR network in multiple maintenance windows. This strategy assumes that all the MPLS PE devices support SR.



**Figure 5-1** *Interworking Greenfield Strategy: Migrating PE Devices to a New SR Network*

In cases where the MPLS PE devices do not support SR, the CE devices can be migrated from the MPLS PE devices to the SR PE devices, as shown in Figure 5-2.




**Figure 5-2** *Interworking Greenfield Strategy: Migrating CE Devices to the New SR Network*

Later sections of this chapter covering the interworking strategy involve the migration of PE devices or CE devices:

- The section “Building a New SR-MPLS Network” covers migration from LDP to SR-MPLS, where a border node is the gateway between the LDP and SR-MPLS networks
- The section “Building a New SRv6 Network Using an SRv6 IWG” covers migration from MPLS to SRv6 where an SRv6 interworking gateway (IWG) acts as the IWG between the MPLS and SRv6 networks, and the section “Building a New SRv6 Network Using Inter-AS Option A” presents a solution using ASBRs as IWGs between the two networks. The information on the migration of CE devices is also valid for the migration of PE devices.

Figure 5-3 shows the second strategy, where PE devices can be concurrently connected to both the MPLS and SR networks, thus removing the necessity for an interworking gateway between the two networks.




**Figure 5-3** *Dual-Homed Greenfield Strategy : PE Devices Connected to the MPLS and SR Networks*

The PE devices must have sufficient port capacity, CPU, and memory to connect to the MPLS and SR networks simultaneously. The dual-homed strategy for migration from MPLS to SR is described in the section “SRv6 Network Using Dual-Connected PE Devices.” The MPLS network can be decommissioned once all the MPLS PE devices have been migrated to the SR network.

A brownfield deployment uses a coexistence strategy, incorporating SR into the existing network as illustrated in Figure 5-4. This strategy entails activating SR in specific parts of the MPLS network and progressively expanding its scope in multiple migration windows.

For this approach to work, the existing network must be able to support SR. Throughout the transition phase, PE and P devices are configured to support MPLS and SR concurrently. When all the devices in the network are migrated to SR, the MPLS-related configuration can be removed from the network.




**Figure 5-4** *Coexistence Brownfield Strategy: Enabling SR in an MPLS Network*

The coexistence strategy is discussed in the following sections:

- LDP to SR-MPLS migration is presented in the section "Enabling SR-MPLS in an Existing Network (Coexistence)."
- MPLS to SRv6 migration is explained in the section "SRv6 Network Using Dual-Connected PE Devices." Although this section shows the migration strategy for a greenfield deployment, it is also valid for a brownfield deployment.

Figure 5-5 shows another SRv6-specific brownfield strategy, which involves expanding the coverage of SRv6-based network services between PE devices connected to an existing IPv6 network.



**Figure 5-5** *IPv6 Backhaul Brownfield Strategy: SRv6-Based Services over an IPv6 Network*


Since the data packets from an SRv6 PE device are IPv6 packets with optional extension headers, they can be transported across any IPv6 network to facilitate SRv6-based service connectivity. However, this strategy comes with certain limitations, like the lack of TI-LFA and SRv6 traffic engineering in the native IPv6 network.

**Note** Migration of SRv6-based services using an existing IPv6 network is not covered in this chapter.

# L2VPN Service Deployment: Configuration and Verification Techniques

In this chapter, we explore L2VPN overlay services established across the SRv6 transport underlay frameworks we've already examined. This chapter outlines fundamental approaches for configuration and methods for confirming the integrity of L2VPN service structures. Keep in mind that the content provided here does not delve extensively into L2VPN services. Instead, this chapter serves as a primer on the transition and deployment of well-established L2VPN technologies within SRv6 transport infrastructures.

If there is a requirement for Layer 2 connectivity, a service provider must set up an L2VPN service using technologies such as Virtual Private LAN Service (VPLS), Virtual Private WAN Service (VPWS), or the more recent advancement in L2VPN technology, Ethernet VPN (EVPN). These overlay services are then transported across a unified core transport network. Transport networks capable of handling L2VPN can offer Ethernet LAN (E-LAN), Ethernet Tree (E-Tree), or Ethernet Line (E-Line) services. In this setup, the service provider maps the incoming customer Layer 2 traffic into bridge domains or establishes point-to-point circuits. These bridge domains or point-to-point circuits are then interconnected across the transport network with other customer site Layer 2 bridge domains or point-to-point circuits. In this way, L2VPN services facilitate the extension of subnets from one end to the other, enabling the provision of managed services such as point-to-point, Internet connectivity, intranet, and extranet services to end customers. Figure 6-1 provides a high-level illustration of how L2VPN overlay services are transported over an SRv6 core network, which will be the primary focus of this chapter.”



**Figure 6-1** L2VPN Connectivity Across an SRv6 Core Network

## L2VPN (EVPN)

The introduction of Multiprotocol Label Switching (MPLS) in RFC 3031, with its highly efficient and flexible data transportation capabilities through the use of MPLS labels, allowed for fast and efficient forwarding decisions without the need for complex IP lookups at each hop. MPLS is protocol agnostic, meaning it can transport packets of various network protocols, such as IP and Ethernet packets. This flexibility is a key part of MPLS's utility in creating sophisticated and scalable network services, including Layer 2 virtual private networks (L2VPNs).

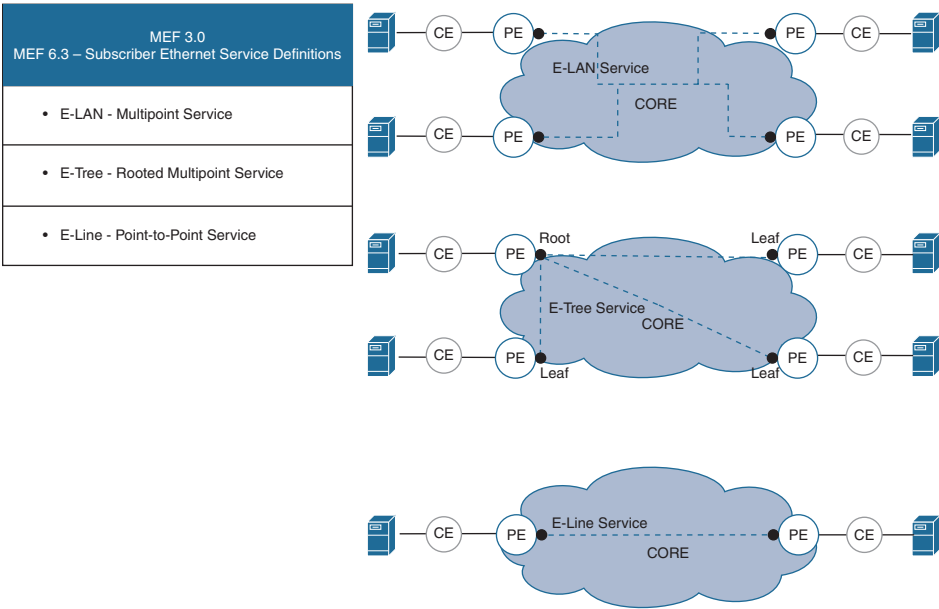
MPLS L2VPN services provide a transport mechanism for Layer 2 frames between multiple customer sites across an MPLS backbone. These services essentially allow the extension of customer Layer 2 networks across geographically dispersed locations, making it possible to create a VPN that emulates a single LAN segment to the customer. L2VPNs are built using two main architectures: Virtual Private LAN Service (VPLS) and Virtual Private Wire Service (VPWS). VPLS provides a multipoint-to-multipoint service, emulating an Ethernet LAN, while VPWS offers point-to-point connectivity, similar to a traditional leased line. Both VPLS and VPWS are essentially transported across a core network through the use of MPLS pseudowires (PW). A pseudowire is simply an emulated point-to-point connection established across a packet-switched network (PSN) that uses Label Distribution Protocol (LDP) for setting up the pseudowire circuits.

Over the past 15 years, Ethernet VPN (EVPN) has begun to gain traction with service providers and large enterprises, and it is now seen as the logical evolution from the VPLS architecture for Layer 2 provisioning. Limitations that are inherent to VPLS lack of multipathing and multihoming capabilities, lack of multicast optimization and redundancy, among others are resolved through EVPN.

EVPN uses Border Gateway Protocol (BGP) to address these limitations via the control plane, whereas VPLS is inherently a data plane learning and forwarding solution. BGP as its control plane protocol provides more scalability than the flooding and learning mechanism used in VPLS. EVPN can handle a larger number of endpoints without suffering from the same level of complexity and resource consumption that VPLS might encounter due to network expansion. EVPN is able to multicast traffic more efficiently

through the use of inclusive multicast Ethernet routes, eliminating the requirement to flood multicast traffic to all endpoints. VPLS, in contrast, typically floods multicast traffic across the entire Layer 2 domain. The multihoming capabilities inherent with EVPN enable single customer edge (CE) routers to connect to multiple provider edge (PE) devices for increased redundancy and load balancing. More granular traffic isolation through unique route targets (RTs) and route distinguishers (RDs) for different services is an additional benefit of EVPN. This is an improvement over VPLS, which typically relies on a single broadcast domain for all connected sites. EVPN uses BGP to advertise MAC addresses, leading to more optimal forwarding paths and faster convergence in the event of network failures. VPLS, on the other hand, relies on traditional MAC learning, which can be slower to converge. EVPN supports both Layer 2 VPN and Layer 3 VPN services, allowing for integrated routing and bridging in the same service instance. This provides greater flexibility in network design compared to VPLS. EVPN is therefore the superior option for Layer 2 VPN services, offering enhanced robustness, scalability, and flexibility over VPLS.

Metro Ethernet Forum (MEF) is an industry consortium that defines standards for carrier Ethernet services. Within the Metro Ethernet Forum 3.0 (MEF 3.0) umbrella standard, several subscriber and operator services standards are defined. One of them, MEF 6.3, is a specification that outlines various Ethernet services and attributes from the perspective of the subscriber. It defines the characteristics and types of Ethernet services that service providers can offer to their subscribers. Figure 6-2 provides a diagrammatic overview of these subscriber services.

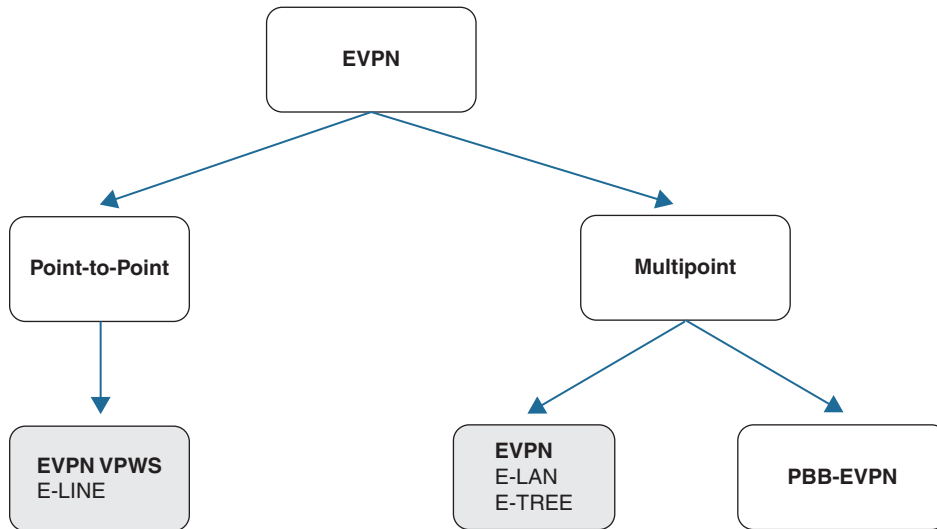


**Figure 6-2** MEF 6.3: Subscriber Ethernet Service Definitions



Figure 6-3 illustrates how EVPN services map to the MEF 6.3 subscriber service definitions:

- **E-LAN service (multipoint-to-multipoint connection):** EVPN E-LAN interconnects multiple endpoints in a multipoint-to-multipoint fashion that allows for any-to-any connectivity between customer sites, similar to traditional Ethernet LAN services.
- **E-Tree service (rooted multipoint connection):** EVPN E-tree involves designating certain sites as “root” or “hub” sites and others as “leaf” sites, with traffic flowing from leaf to root and from root to leaf but not directly between leaf sites.
- **E-Line service (point-to-point connection):** EVPN VPWS can be used to create a virtual point-to-point Ethernet service. This is typically achieved by setting up an EVPN instance with only two endpoints to provide a dedicated Ethernet connection between two customer sites



**Figure 6-3** *EVPN Family: Next-Generation Ethernet Service Solutions*

## EVPN in Detail

Ethernet VPN (EVPN) is a technology that provides Layer 2 VPN services over IP/MPLS transport networks and supports both unicast and multicast traffic, using BGP to distribute MAC/IP address reachability information. EVPN provides extensible and flexible multihoming VPN solutions for intra-subnet connectivity among tenant systems and end devices, which may be physical or virtual. EVPN provides E-LAN services and also allows for the provisioning of E-LINE services, with either port-active, single-active, or all-active multihoming with flow-based load balancing. EVPN VPWS simplifies pseudowire (PW) signaling and provides fast convergence upon node or link failure. The RFCs that define standards for EVPN include RFC 7432, RFC 8214, RFC 8365, and RFC 9135

## Business Opportunities

While the previous chapters provide detailed technical information about segment routing (SR), this chapter explores the business opportunities that emerge from adopting this innovative technology. No company will invest in a technology if its business cannot benefit in some way. This chapter discusses various topics to help build a bridge between engineering, operations, marketing, finance, and product management. It can help stakeholders and leadership identify relevant benefits in their areas, effectively communicate, and justify the investment in a transformation where silos are torn down and barriers are identified and addressed in a timely manner for a successful transition to SR.

Networks and associated IT systems are implemented differently across network service providers and large enterprises, tailored to offer the desired services to internal and external customers. Likewise, network service providers adhere to different standards and frameworks, and each organization's history and decisions influence to what extent standards and frameworks are adopted. We appreciate that organizational structures, processes, terminologies, services, and other aspects vary significantly across network service providers, and this chapter attempts to be descriptive without using any specific standard, framework, or terminology.

Echoing the insights from earlier technical chapters—such as Chapter 2, “What Is Segment Routing over MPLS (SR-MPLS)?” and Chapter 3, “What Is Segment Routing over IPv6 (SRv6)?” —this chapter highlights substantial opportunities and advantages of SRv6 compared to SR-MPLS.

**Note** Given that not all organizations will prioritize SRv6 transformation work items or perceive the benefits the same way, there is no universal blueprint that fits all possible scenarios. Consequently, the topics in this chapter are presented with different perspectives in mind.

The introduction of SRv6 technology presents a unique opportunity for network service providers. By harnessing its full potential, providers can gain a significant market edge or at least ensure that they remain competitive with their network service offerings over the coming decade. The improved performance, greater scale, network simplification, and new service options associated with SRv6 allow for the convergence of multiple networks into one. This consolidation ideally results in just one network to purchase, build, operate, support, power, cool, and host, thereby leading to substantial reductions in capital expenditures (CapEx) and operational expenditures (OpEx). The benefits extend to potential organizational optimizations, avoided costs for redundancy and scale-related spare capacity in multiple networks, simpler and improved service-level management (SLM), fewer integration points, and the convergence of operation support systems (OSS), business support systems (BSS), and IT systems in general. These factors can multiply the business opportunity related to the introduction of SRv6.

However, the advantages of SRv6 are not limited to consolidation and optimization. You may recall the transition from Asynchronous Transfer Mode (ATM)-based services to IP technology and its profound market impact two to three decades ago. Similarly, SR protocol options enable the transportation of new services over IP. Leased line and optical point-to-point services can now be offered over IP, reducing the need for optical network-based services to be exposed to customers. This shift simplifies the optical network stack, reducing the requirements for its OSS, BSS, and related IT systems to implement, test, support, and maintain. Instead, the IP services stack can incorporate these as additional service flavors, potentially reducing the cost of offering traditional leased line or optical services and enabling providers to offer these services at more competitive prices. Network service providers currently relying on third-party leased line or optical point-to-point services may even consider offering such services themselves, using their SR network. SRv6 allows for simplified chaining of network connectivity services with additional services such as Network Address Translation (NAT), firewalls, deep packet inspection (DPI), intrusion prevention systems (IPSs), and services offered on virtual machines or containers within data centers or clouds, reducing complexity and costs.

Network service providers operating single IP networks and requiring relatively simple IP or VPN connectivity services can also benefit from the introduction of SRv6. Although traditional MPLS VPN transport networks have been around for two decades, the networking industry is less likely to invest in significant developments or address known limitations due to its maturity, and so the transition to SRv6 is the next logical step. With device generations increasingly focusing their feature support around SRv6 and reducing legacy features to remove complexity and costs, it's worth evaluating investment in SRv6 for the coming decade.

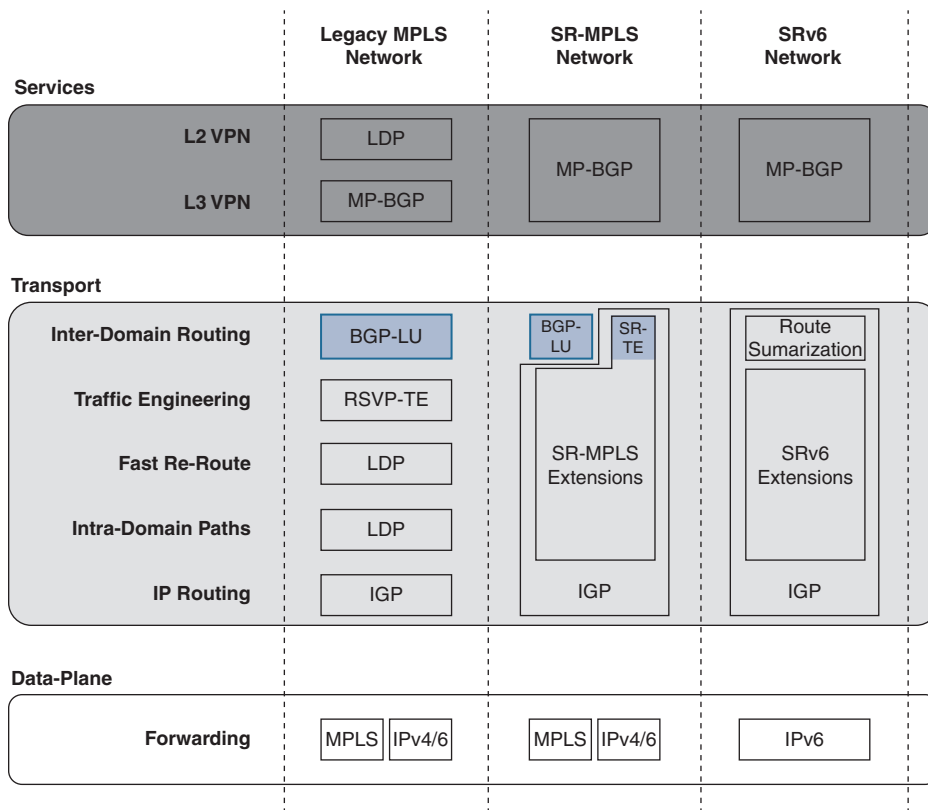
The introduction of SRv6 represents a tremendous opportunity for network service providers. Those who thoroughly analyze and leverage its full potential can undergo a true transformation, benefiting their business and customers for years to come. A detailed analysis may even reveal the business benefits associated with SRv6 as justification for an early network lifecycle. The following sections delve into how simplification, convergence, and standardization can lead to new business and enrich existing services.

## Technological Opportunities and Benefits

Before we dive into CapEx and OpEx opportunities, this section provides background information on selected opportunities. Some of these opportunities are not directly related to SR as a technology but rather to the fact that a new network may be built, or a new technology may be introduced.

### Fewer Protocols

Figure 10-1 shows network technology–specific protocol stacks (though it omits protocols that are in common across all technologies).



**Figure 10-1** Comparison of MPLS, SR-MPLS, and SRv6 Protocols

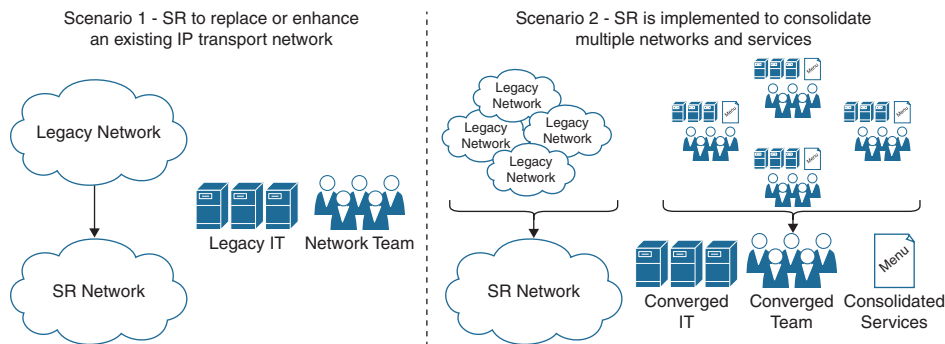
## Organizational Considerations

After reviewing the promising business case for segment routing (SR) in Chapter 10, “Business Opportunities,” you are now ready to evaluate how the introduction of SR will affect your organization. The potential impacts can vary greatly, depending on the number of networks converging and migrating to a new SR-based IP transport network, as well as the complexity of traditional non-IP services that may migrate to SR using new capabilities such as private line emulation (PLE). Just as the transition from Time-Division Multiplexing (TDM)-based services to IP/MPLS networks posed challenges in the mid-2000s, SR may impact more than just network engineering or operation teams. Departments that take care of marketing, sales, customer relationship management, and product and service portfolio management, along with any business partners and resellers will need to adapt to varying degrees.

Network service providers adhere to a variety of standards and frameworks, each shaped by an organization’s unique history and set of decisions. This chapter serves as a guide, outlining important considerations and potential pitfalls for those leading the transformation to a programmable SR network and offering strategies to circumvent those pitfalls before they hinder progress. Given the considerable diversity in organizational structures, processes, terminologies, services, and other aspects across network service providers, this chapter is descriptive and deliberately avoids adherence to any specific standard, framework, or terminology, allowing for broad applicability and flexibility.

Throughout this book, the term *domain* refers to a segment of a network. This chapter expands on that concept, discussing how forming an SR domain can affect various areas, such as personnel, network infrastructure, IT frameworks, processes, service offerings, and development activities.

Although each network service provider follows its own unique path to SR, this chapter categorizes the various paths into the two scenarios, shown in Figure 11-1, to examine impacts and assist in navigating potential challenges.



**Figure 11-1** *Two Scenarios for Implementing SR*

The following areas are relevant in Scenario 1, where SR is introduced as a new technology to replace or enhance an existing IP transport network:

- **Knowledge:** The network architecture, engineering, and operation teams need to familiarize themselves with the SR technology.
- **Migration strategy:** A review of the current infrastructure for feature support, scalability, and anticipated remaining lifetime supports the choice between migrating to SR-MPLS versus SRv6.
- **IT evolution and gap awareness:** Applications involved in automating network resource and service configurations, along with monitoring and assurance systems, need adjustments to handle SR technology specifics. An assessment can provide an overview on the adaptations required to manage a new SR network. In the absence of viable options, establishing a new SR IT stack could be seen as a strategic move to modernize and phase out legacy systems.

Each of these points is discussed in greater detail later in this chapter.

When implementing SR to merge various networks and services, as in Scenario 2, the impact is potentially even greater. It's not just about combining networks; it's about fusing teams, processes, IT systems, and operational domains into a single entity. Alongside the still-relevant focus areas just listed for the first scenario, the more complex Scenario 2 calls for thoughtful evaluation of several additional considerations.

- **IT evolution and gap awareness (extension):** In Scenario 2, IT systems from various departments need to be consolidated. Separate workflow automation, fulfillment, inventory, IP address management, backup, and other systems must be converged to maximize simplicity, efficiency, and business benefits.
- **Domain definitions:** Consolidating teams and their operational domains demands a strategic approach to guarantee that the newly formed entity overseeing the SR network domain operates efficiently and effectively. This process includes evaluating and, if needed, redefining domains, roles, authorities, and responsibilities, as well as


potentially consolidating physical network locations. Thorough preparation paves the way for a smoother transition for all affected teams and domains.

- **Team organization and transformation:** The merging of teams involves not only the blending of different skill sets but also the integration of diverse backgrounds and varying approaches to work, communication, and decision making. It also introduces significant uncertainty, raising questions about team composition. Developing a clear strategy for the team's evolution and maintaining open communication are essential to preserve motivation and ensure continuity throughout the transformation.
- **Existing and new processes:** Reviewing existing processes to pinpoint those affected by the transition to the new SR network domain is crucial for defining the overall transition scope. By simultaneously capturing the efficiency and effectiveness of existing processes, it is possible to identify potential templates for any new processes required in the SR domain.
- **Network services portfolio consolidation:** Merging network service portfolios is pivotal in consolidating multiple networks. Services often vary widely across networks, with some potentially offering numerous manual configuration options. Developing a service model that consolidates all services from the affected legacy networks demands considerable effort and will help determine which variants should be phased out to establish a standardized service definition. Although unlikely, it may be possible for the service modeling process to reveal a comprehensive model that encapsulates all service options from the merging networks. Regardless of the details, harmonizing the service portfolio is essential in order to streamline automation, assurance, testing, migration, and operations of the converged SR network.
- **Development and release methodology:** Individuals forming the new SR domain team will bring a variety of experiences from their previous roles, where they might have used Agile, Waterfall, DevOps, or a blend of these and other methodologies. They are likely comfortable with a variety of practices, artifacts, lab environments, processes, and tools. The integration of network services does more than just merge these different professional experiences; it also consolidates engineering and operational responsibilities, risks, and accountability within a unified domain. To effectively navigate the complexities of the SR domain and ensure both superior quality and efficient operations, a well-defined and robust development and release methodology is crucial.
- **Change management across domains:** When converging organizations, there is a need for a comprehensive change management strategy that addresses all levels of the affected domains in the organization. This strategy should include communication plans to keep all stakeholders informed, a common overall roadmap to align domains, training programs to upskill employees where necessary, and feedback mechanisms to address any concerns and challenges that may arise. Such a central change management strategy is critical to help domains and their employees transition to new ways of working, to foster acceptance of the new organizational structure, and to ensure that the combined entity can achieve its desired synergies and performance objectives.

The subsequent sections of this chapter delve into all these aspects in detail.

## Scenario 1: Replacing or Enhancing a Legacy Network with SR

Every journey to SR needs to consider at least the three areas shown in Figure 11-2.



**Figure 11-2** *Scenario 1: Replacing or Enhancing a Legacy Network with SR*

The following subsections explore these areas in greater detail and offer ideas to simplify the transition to SR.

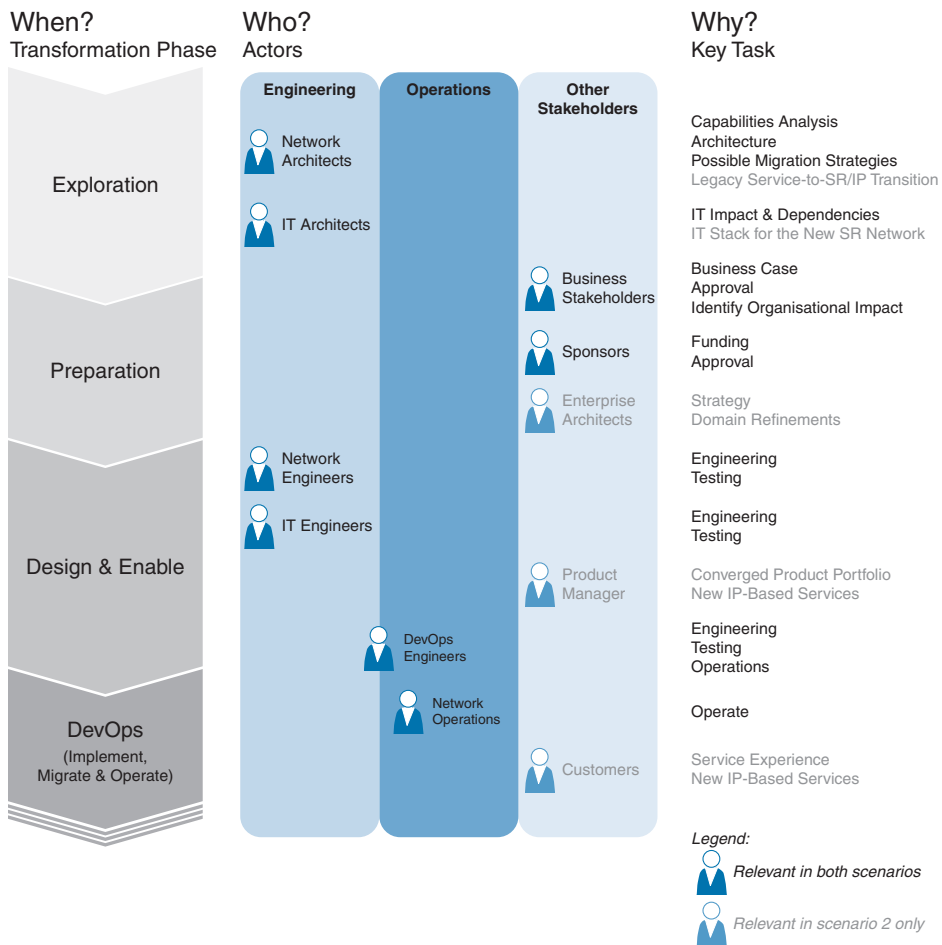
### Knowledge

What type of SR knowledge is essential for various roles within the organization? Who is tasked with making investment decisions? Who will fill which roles in the upcoming months, and who must possess substantial SR knowledge? This section aims to address these questions by identifying the key actors responsible for introducing SR within a service provider's network and discussing potential sources of knowledge. The actors and the sequence in which they require SR knowledge have been greatly simplified to serve as an introductory guide. To prevent redundancy, Scenario 2, which involves a more comprehensive transition to SR, is visually distinguished in Figure 11-3 by gray highlighting on the Scenario 2–relevant stakeholders and key tasks.

Let's now look at the reasoning for the required knowledge and the type of expertise needed by the actors at each chronological phase:

- **Network architects:** Faced with lifecycle challenges such as software or hardware nearing end-of-life or scaling issues, these professionals are tasked with finding suitable alternatives or successor technologies. They evaluate SR capabilities, identify network elements that lack SR support (such as load balancers and NAT), assess the impact on existing infrastructure, explore new service opportunities or enhancements using SR, and develop strategies for seamless SR migration to maintain service continuity. All these elements are then integrated into a target architecture, which serves as a baseline to evaluate the transition's impact on interfacing networks and existing IT systems.





**Figure 11-3** SR Knowledge: Who Needs It, When, and Why

# Index

## Symbols

---

\* wildcard, 1112

## A

---

ABR (area border router), 5, 38, 39, 224  
acceptance test, 1085, 1087  
active-active mode, 425  
active-backup mode, 426  
address family, BGP, 84  
adjacency, IS-IS, verification, 265–266  
adjacency segment, 226  
Adjacency SID, 47–49, 226  
Adj-SID sub-TLV, 59–61, 68–69  
algorithm. *See also* Flex Algo  
    operator-defined, 74–75  
    SPF (Shortest Path First), 56, 73–74, 220  
Angstrom era, ONLINE  
anycast, set, 219  
anycast SID, 45, 47, 219  
    SR-MPLS, 254–255  
        *configuration*, 254–256  
        *verification*, 256–257  
SRv6, 270–271

*configuration*, 271

*use case*, 272–282

*verification*, 272

architecture, SPRING, 40–41

area

    IS-IS, 222

    OSPF, 224

ASBR (autonomous system border router),  
    5, 20

assembler code, 34–35

assurance, intent-/model-based,  
    1019–1020. *See also* service assurance

automation, 1065, 1097–1098

    pipeline, 1098, 1101–1104

    solution lifecycle, 1067

    test, 1094

## B

---

backup path, 19

base format, 108

BD (bridge domain), 447

BFD (Bidirectional Forwarding Detection),  
    857–859. *See also* BLB (BFD over  
    Logical Bundle); BoB (BFD over  
    Bundle)

- async mode, 859–860
- BLB, 861
  - configuration*, 872–874
  - packet capture*, 863–865
  - verification*, 874–883
- BoB, 861
  - configuration*, 865–868
  - packet capture*, 863
  - verification*, 868–871
- demand mode, 859
- echo, 859–860
- over LAG interfaces, 860
- packet capture, 861–862
- SH (single-hop) sessions, 860–861
- BG (bridge group)**, 447
- BGP. *See also* MP-BGP (Multiprotocol Border Gateway Protocol)**
  - address families, 84
  - free core, 1
  - L3VPN, Flex Algo, 318–322
  - Link-State Extensions for SRv6, 201–205
  - MP\_REACH\_NLRI path attribute, 608–609
  - multiprotocol, 17
  - NHT (next hop tracking), 220
  - operational complexity, 22
  - PIC (Prefix Independent Convergence), 292–293
  - PIC Edge, 945–948
    - multipath verification*, SRv6, 981–995
    - unipath configuration*, SR-MPLS, 948–951
    - unipath configuration*, SRv6, 962–970
    - unipath verification*, SR-MPLS, 951–962
    - unipath verification*, SRv6, 970–981
  - Prefix SID, 324, 609–610
    - configuration*, 324–326
    - enabling in an SR-MPLS network*, 376–383
    - verification*, 327–328
  - protocol ID, 96
  - proxy Prefix SID, 383–387
  - RR (route reflector), 6
  - UPDATE message, 83–84, 86–87, 96–98, 186, 187–188, 190, 191, 203–204
    - for SRv6 L2 services*, 192–193
    - for SRv6 L3 services*, 194–195, 196–198
- BGP Link-State extensions for SR**, 87–95
- BGP Peering SID**, 50–52
- BGP Prefix SID**, 49–50, 85–87
- BGP-LU (BGP Labeled Unicast)**, 220
  - inter-AS, 343
    - configuration*, 344–345
    - data forwarding*, 349–350
    - design*, 343
    - verification*, 345–348
  - intra-AS
    - design*, 330
    - with Prefix SID*, 328–330
- binary repository**, 1100–1101
- Binding Segment SID**, 52–53
- BLB (BFD over Logical Bundle)**, 860, 861
  - configuration*, 872–874
  - packet capture*, 863–865
  - verification*, 874–883
- BoB (BFD over Bundle)**, 860, 861
  - configuration*, 865–868
  - packet capture*, 863
  - verification*, 868–871
- broadcast domain, EVPN**, 444
- brownfield deployment**, 354
  - coexistence migration strategy, 356–357
  - IPv6 backhaul strategy, 356–357
  - SR-MPLS, 358–360
    - enabling and preferring on P1, P2, and PE-1*, 363–365

*enabling on P2, P3, and PE-3,*  
360–363

**BVLAN (BFD over VLAN over Bundle),**  
860

## C

**CCM (Continuity Check Message),** 809–  
810

**CE (customer edge) router,** 5, 6

**CFM (Ethernet Connectivity Fault  
Management),** 808

CCM (Continuity Check Message),  
809–810

configuration, 811–812

Linktrace Protocol, 811

Loopback Protocol, 810

MA (maintenance association), 808

MD (maintenance domain), 808

MEP (maintenance endpoint), 808–809

MIP (maintenance intermediate point),  
809

PDU's, 809

verification, 813–816

**change management,** 1104–1106

**CI/CD/CT (continuous integration/  
continuous delivery/continuous  
testing),** 1101

**Cilium,** ONLINE

**CML (Cisco Modeling Lab),** ONLINE

code, assembler, 34–35

coexistence brownfield strategy, 356–357

collision, label, 43

**command/s**

containerlab deploy, ONLINE

containerlab inspect, ONLINE-ONLINE

debug lacp packets, 483

help sr localsid, ONLINE-ONLINE

hw-module, 240

hw-module bfd-hw-offload enable  
location, 867

microloop avoidance rib-update-delay  
5000, 942–943

pcap trace, ONLINE-ONLINE

ping, 249, 254, 270, 349, 423, 635, 676,  
767, ONLINE-ONLINE, ONLINE-  
ONLINE, ONLINE-ONLINE,  
ONLINE-ONLINE, ONLINE-  
ONLINE

ping ethernet cfm domain service, 815

segment-routing mpls sr-prefer, 361

service vpp status, ONLINE-ONLINE

service vpp stop, ONLINE-ONLINE

show bfd ipv6 session, 868–869

show bfd ipv6 session interface bundle-  
Ether 1 detail, 869

show bfd ipv6 session interface TF0/0/0/0  
detail, 870–871

show bfd neighbors, 877

show bfd neighbors interface port-  
channel 1 details, 877–878

show bfd session, 875

show bfd session interface bundle-ether 1  
detail, 875–876

show bgp ipv4 labeled-unicast, 347, 348,  
377–378, 379–381, 713, 740–741

show bgp ipv4 rt-filter, 771, 774

show bgp ipv4 rt-filter neighbors,  
774–775

show bgp ipv4 unicast, 328, 348,  
713–714

show bgp ipv4 vpn, ONLINE-ONLINE

show bgp l2vpn evpn bridge-domain,  
446, 521–522, 525–527, 537, 540,  
549–550, 568–569

show bgp l2vpn evpn bridge-domain  
200-BD, 527–529

show bgp l2vpn evpn bridge-domain  
200-BD received-sids wide, 523–524

show bgp l2vpn evpn bridge-domain  
VPWS:300, 594–599

show bgp l2vpn evpn rd, 458–461, 464,  
467–469, 470, 523, 601

show bgp l2vpn evpn route-type, 530

show bgp l2vpn evpn route-type ethernet-segment, 530–532, 600  
 show bgp l2vpn evpn summary, 518–519  
 show bgp segment-routing srv6, ONLINE-ONLINE  
 show bgp summary, ONLINE-ONLINE  
 show bgp vpnv4 uni vrf, ONLINE-ONLINE  
 show bgp vpnv4 unicast, 741–742  
 show bgp vpnv4 unicast advertised summary, 408–409  
 show bgp vpnv4 unicast vrf, 626, 628–630, 646–647, 649–650, 651–652, 655, 659–661, 663, 667–668, 670, 673–674, 702–703, 704–705, 705–706, 733–735, 758–759, 760, 951–953, 954, 958–960, 971–973, 976–977, 982–984, 989–990  
 show bgp vpnv4 unicast vrf local-sids, 627, 647–648, 663–664, 671  
 show bgp vpnv4 unicast vrf received-sids, 627–628, 648, 664, 671  
 show bgp vrf, 276, 320–321, 409–410, 420–421  
 show bgp vrf nexthop-set, 634, 650, 667  
 show bgp vrf-db table, 566–567, 656  
 show bgp vrf-db table all, 519–521, 594  
 show bpg vpnv4 unicast vrf, 665–666  
 show bundle bundle-Ether 200, 478, 484  
 show bundle bundle-Ether 250, 481, 482  
 show bundle bundle-Ether 300, 575–576  
 show cef, 887–888, 892–893, 898–899, 926  
 show cef detail, 714–715, 716, 717–718, 743, 744–745, 747, 748, 749–750  
 show cef ipv6, 904–905, 908–909, 912–913, 937–940, 942  
 show cef ipv6 detail, 980–981, 993–994  
 show cef vrf, 276, 279  
 show cef vrf detail, 632–633, 654, 669–670, 675, 709–711, 738, 764, 766, 955–957, 960–962, 974–975, 979, 986–988, 991–992  
 show ethernet cfm local meps domain service, 813  
 show ethernet cfm peer meps domain service, 814–815  
 show ethernet sla operations detail profile, 820–821, 827–828  
 show ethernet sla statistics brief profile, 821–823, 829–830  
 show ethernet sla statistics detail profile, 823–825, 831–834  
 show evpn ethernet-segment, 494, 580  
 show evpn ethernet-segment interface BE200 carving detail, 496–498  
 show evpn ethernet-segment interface BE300 carving detail, 581–582  
 show evpn ethernet-segment interface Bundle-Ether260 detail, 538–539  
 show evpn evi, 498, 583–584  
 show evpn evi ead, 584–585  
 show evpn evi inclusive-multicast detail, 502–503  
 show evpn evi vpn 200 mac, 501–502  
 show evpn evi vpn 205 mac, 560–561  
 show evpn evi vpn-id 200 detail, 499–500  
 show evpn summary, 493–494  
 show interfaces brief, 477–478, 483, 574–575, ONLINE-ONLINE  
 show ip cef detail, 715, 717, 718, 744, 745–746, 749  
 show ip cef internal, 889–890, 894–895, 900–901, 930–931  
 show ip cef vrf, ONLINE-ONLINE  
 show ip cef vrf detail, 711, 739, 765, 957  
 show ip ospf database, 234  
 show ip route repair-paths, 888–889, 893, 899–900, 928, 954–955  
 show ip route vrf, 737, 762  
 show ip route vrf repair-paths, 709  
 show ipsla statistics 11, 843–844

show ipsla statistics 21, 844–846  
 show ipsla statistics aggregated 11, 846–847  
 show ipsla statistics aggregated 21, 847–850  
 show ipsla twamp session, 854  
 show ipv6 cef, ONLINE-ONLINE  
 show ipv6 route isis, ONLINE-ONLINE  
 show isis adjacency, 265–266  
 show isis database, 233, 247, 265, 272, 288, 291, 299, 309–310  
 show isis fast-reroute summary, 884  
 show isis fast-reroute ti-lfa tunnel, 894, 900, 929  
 show isis flex-algo, 308  
 show isis ipv4 fast-reroute detail, 887, 892, 898, 904, 918, 925, 927–928, 933–934, 934  
 show isis ipv6 fast-reroute detail, 908, 912  
 show isis neighbor, ONLINE-ONLINE  
 show isis rib, 929–930  
 show isis srv6 locators det, ONLINE-ONLINE  
 show l2vpn bridge-domain bd-name 200-BD detail, 508–510  
 show l2vpn bridge-domain bd-name detail, 563–565  
 show l2vpn bridge-domain brief, 508, 511, 544  
 show l2vpn forwarding bridge-domain ELAN-BG:200-BD mac-address location 0/RP0/CPU0, 510–511  
 show l2vpn forwarding bridge-domain ELAN-BG:210-BD mac-address location 0/RP0/CPU0, 545  
 show l2vpn forwarding xconnect detail location, 591  
 show l2vpn mac-learning, 512  
 show l2vpn mac-learning mac all location, 511  
 show l2vpn xconnect group VPWS-XC, 589–591  
 show lacp, 479–480, 576–577  
 show lacp bundle-Ether 250, 481, 482  
 show memif, ONLINE-ONLINE  
 show mpls forwarding, 362, 364, 374  
 show mpls oam dpm adjacency, 797  
 show mpls oam dpm prefix, 797  
 show mpls oam dpm summary, 795–796  
 show route, 891, 896–897, 917–918, 924–925  
 show route 10.0.1.8/32, 886–887  
 show route ipv6, 280–281, 282, 289, 290, 291–292, 300, 311–312, 903–904  
 show route ipv6 detail, 906–907, 911–912  
 show route vrf, 278  
 show route vrf detail, 631–632, 653–654, 668–669, 674–675, 708, 736, 761, 762–763, 973–974, 977–978, 985–986  
 show segment-routing srv6 capabilities-parameters, ONLINE-ONLINE  
 show segment-routing srv6 locator, ONLINE-ONLINE, ONLINE-ONLINE  
 show segment-routing srv6 locator MAIN detail, ONLINE-ONLINE  
 show segment-routing srv6 locator MAIN sid, 513–514, 546–547  
 show segment-routing srv6 sid, 503–504, 586, ONLINE-ONLINE, ONLINE-ONLINE  
 show segment-routing srv6 sid detail, 633, 634–635, 656  
 show segment-routing traffic-eng policy, 274–275  
 show slrg name, 932  
 show sr localsids, ONLINE-ONLINE, ONLINE-ONLINE  
 show sr policies, ONLINE-ONLINE  
 show sr steering-policies, ONLINE-ONLINE

- show trace, [ONLINE-ONLINE](#),  
[ONLINE-ONLINE](#), [ONLINE-ONLINE](#)
- systemctl restart frr, [ONLINE-ONLINE](#)
- traceroute, 249, 312, 342, 362–363,  
364–365, 375–376, 382, 383, 412, 423,  
645, 677, 718–719, 750–751, 767
- traceroute ethernet cfm domain service,  
816
- compiler, 34**
- Containerlab, [ONLINE-ONLINE](#)**
  - Linux SRv6 lab deployment, [ONLINE-ONLINE](#)
  - topology definition, [ONLINE](#), [ONLINE-ONLINE](#)
- containerlab deploy command, [ONLINE](#)**
- containerlab inspect command, [ONLINE-ONLINE](#)**
- CONTINUE operation, 41**
- control plane, 10, 87–88**
  - MPLS, 205–207
  - show isis database, 257
  - SR-MPLS, 207–209, 243–244
  - SRv6, 209–210, 257
  - traceroute, 254
  - VPP (Vector Packet Processor), [ONLINE-ONLINE](#)
- converged SR domain, 1083**
  - development and release process,  
1084–1086
    - environments, 1089–1096*
    - key phases, 1086–1089*
  - service portfolio consolidation,  
1083–1084
- convergence, 878–879. *See also* Microloop Avoidance**
  - calculating, 858
  - failure event detection time, 858
  - IPFRR (IP Fast Reroute), 878–881
  - microloops, 935
  - network event information propagation  
time, 858

- RIB/FIB update time, 858
- RLFA (Remote Loop-Free Alternate),  
880–882
- topology update and repair path  
computation time, 858
- CPU (central processing unit), 33, 35, [ONLINE](#)**
  - ISA (instruction set architecture), 35
  - program counter, 35
- CSC MPLS network migration, 434–435**

## D

---

- DARPA (Defense Advanced Research Projects Agency), 28**
- data plane, 10**
  - MPLS, 205–207
  - segment routing, 36–37
  - SONiC (Software for Open Networking in the Cloud), [ONLINE](#)
  - SR-MPLS, 41–42, 207–209
  - SRv6, 209–210
  - VPP (Vector Packet Processor), [ONLINE-ONLINE](#)
- database, label switching, 43, 44. *See also* FIB (Forwarding Information Base); RIB (Routing Information Base)**
- debug lacp packets command, 483**
- deployment model**
  - brownfield, 354
  - greenfield, 354
- dev pipeline, 1098**
- development and release process, 1084**
  - automation, 1097–1098
  - domain releases, 1096–1097
  - key phases, 1086–1089
  - lab environments, 1089–1096
  - testing, 1097
- development test, 1091**
- Dijkstra SPF algorithm, 224**

disaggregation, SONiC (Software for Open Networking in the Cloud),  
ONLINE-ONLINE

domain, 1043. *See also* converged SR domain

architecture blueprint, 1059–1061

*API gateway and graphical user interface, 1061–1062*

*apps and microservices, 1065*

*error and event handling, 1066*

*inventories, 1063–1064*

*messaging, 1066*

*network resources, 1065*

*resource fulfillment, 1064*

*rollout and migration automation, 1064–1065*

*security measures and events, 1066*

*service and resource assurance, 1062–1063*

*service catalog, 1063*

*service orchestration, 1064*

*solution lifecycle automation, 1067*

*workflow automation engine, 1065*

*workforce and customer apps, 1066*

availability of components as shared services by other domains, 1070

boundaries, 1056–1058, 1059

component criticality, 1069

component sourcing, 1071–1074

converged SR, 1083

*development and release methodology, 1084–1086*

*service portfolio consolidation, 1083–1084*

key roles, 1075–1082

organization and transformation, 1074–1075

recommendations for components, 1070–1071

responsibilities, 1067–1069

SAP (service access point), 1058–1059

segment routing, 38

squads, 1082

teams, 1082–1083

downstream on-demand, 11

DPDK, ONLINE-ONLINE

dual-homed greenfield strategy, 356

## E

---

ECMP (equal-cost multipath), 25, 38, 798–800

EFP (Ethernet flow point), 447

E-LAN, 472–473

port-active multihomed service

*access circuit configuration, 474–477*

*access circuit verification, 477–484*

*BGP configuration, 514–518*

*BGP verification, 518–533*

*EVPN configuration, 484–491*

*EVPN verification, 491–504*

*L2VPN configuration, 504–507*

*L2VPN verification, 507–514*

route types and usage, 473

single-homed service, 533–534

*BGP configuration, 520–547*

*BGP verification, 547–551*

*EFP configuration, 534–535*

*EFP verification, 535*

*EVPN configuration, 535–536*

*EVPN verification, 536–542*

*L2VPN configuration, 542–543*

*L2VPN verification, 543–547*

E-Line

all-active BGP



- configuration*, 591–593
  - verification*, 593–601
- all-active EFP (access circuit)
  - configuration*, 571–574
  - verification*, 574–577
- all-active EVPN
  - configuration*, 577–579
  - verification*, 579–586
- all-active L2VPN
  - configuration*, 586–588
  - verification*, 588–591
- BGP route types and usage, 570–571
- End behavior, SRv6, 113–114
- End.B6.Encaps behavior, 142–143
- End.B6.Encaps.Red behavior, 143–144
- End.B6.Insert behavior, 144–145
- End.B6.Insert.Red behavior, 145–146
- End.DT2M behavior, SRv6, 122–123
- End.DT2U behavior, SRv6, 120–121
- End.DT4 behavior, SRv6, 116–117
- End.DX2 behavior, SRv6, 119
- End.DX4 behavior, SRv6, 117–118
- endpoint behaviors, 141
  - SID (segment identifier), 1113
  - SRv6 policy, 141, 146
  - uSID (micro SID), 149–150
- endpoint node, 126
- endpoint/egress node, 38
- end-to-end QoS, MPLS (Multiprotocol Label Switching), 21–22
- End.X behavior, SRv6, 114–115
- EPE (egress peer engineering), 41
- ES (Ethernet segment), 445, 447–449
- ESI (Ethernet segment identifier), 445, 449–450
- Ethernet CFM (Connectivity Fault Management). *See* CFM (Ethernet Connectivity Fault Management)
- Ethernet OAM (Operations, Administration, and Maintenance), 806–807
- Ethernet tag ID, 450–452
- EVC (Ethernet virtual circuit), 447
- EVI (EVPN instance), 444, 445–447
- EVPN (Ethernet VPN), 212, 215, 440–441
  - aliasing, 444
  - benefits, 443–444
  - BGP routes, 452–454
  - broadcast domain, 444
  - E-LAN, 472–473
    - port-active MHD BGP configuration*, 514–518
    - port-active MHD BGP verification*, 518–533
    - port-active MHD EVPN configuration*, 484–491
    - port-active MHD EVPN verification*, 491–504
    - port-active MHD L2VPN configuration*, 504–507
    - port-active MHD L2VPN verification*, 507–514
    - port-active multihomed service*, 474–484
    - route types and usage*, 473
    - SHD BGP configuration*, 520–547
    - SHD BGP verification*, 547–551
    - SHD EFP configuration*, 534–535
    - SHD EFP verification*, 535
    - SHD EVPN configuration*, 535–536
    - SHD EVPN verification*, 536–542
    - SHD L2VPN configuration*, 542–543
    - SHD L2VPN verification*, 543–547
    - single-homed service*, 533–534

- all-active BGP configuration*, 591–593
  - all-active BGP verification*, 593–601
  - all-active EFP configuration*, 571–574
  - all-active EFP verification*, 574–577
  - all-active EVPN configuration*, 577–579
  - all-active EVPN verification*, 579–586
  - all-active L2VPN configuration*, 586–588
  - all-active L2VPN verification*, 588–591
  - BGP route types and usage*, 570–571
  - ES (Ethernet segment), 445, 447–449
  - ESI (Ethernet segment identifier), 445, 449–450
  - Ethernet tag ID, 445, 450–452
  - E-Tree, 551–552
    - BGP configuration*, 565
    - BGP verification*, 565–569
    - EFP (access circuit) configuration*, 553–555
    - EFP (access circuit) verification*, 555–556
    - EVPN configuration*, 556–559
    - EVPN verification*, 559–561
    - L2VPN configuration*, 561–563
    - L2VPN verification*, 563–565
  - EVI (EVPN instance), 445–447
  - MAC VRF instance, 444
  - MEF 6.3, 441–442
  - Route Type 1
    - per-ESI Ethernet A-D route*, 454–458
    - per-EVI Ethernet A-D route*, 458–463
  - Route Type 2, 463–466
  - Route Type 3, 467–469, 512
  - Route Type 4, 470–472
  - VLAN services interfaces, 451–452
  - EVPN VPWS**, 191–192, 199–200
  - extranet L3VPN**
    - over SR-MPLS
      - configuration*, 752–756
      - verification*, 756–767
    - over SRv6
      - configuration*, 657
      - verification*, 662–677
- ## F
- 
- F3216 format**, 148
  - failure event detection time**, 858
  - FEC (forwarding equivalence class)**, 11
  - FIB (Forwarding Information Base)**, 10
    - pre- and post-failure entries, 936–942
    - update time, 858
  - flat MPLS network migration**, 429–430
  - Flex Algo**, 73–78, 1005–1007
    - assigning to a BGP L3VPN service, 318–322
    - calculation of a path, 303
    - definition, 302
    - definition advertisement, 302
    - disjoint paths, 312–313
      - configuration*, 313–315
      - verification*, 315–318
    - installation of forwarding entries, 303
    - link attribute advertisement, 302–303
    - low-latency path, 304–305
      - configuration*, 305–308
      - verification*, 308–312
    - metrics, 78
    - prefix metric, 303
    - SID advertisement, 303
    - SR-MPLS configuration, 322–324

sub-sub-TLVs, 79–83

use cases, 304

Flexible Algorithm Definition sub-TLV,  
78–79

Flow Label field, IPv6 header, 104–106  
forwarding plane, 10

FRR (Fast Reroute), 18–19, 41, 858

FRR (Free Range Routing), ONLINE-  
ONLINE

daemons, ONLINE-ONLINE

IPv4 L3VPN service, ONLINE-ONLINE  
release 9.1, ONLINE-ONLINE

SRv6 headend and endpoint support,  
ONLINE-ONLINE

system architecture, ONLINE

vytish, ONLINE-ONLINE

full-mesh L3VPN

over SR-MPLS

*BGP label allocation modes*, 701

*configuration*, 689–700

*verification*, 701–719

over SRv6

*BGP configuration on PE-2*,  
624–625

*configuration*, 610–622

*full-mesh verification*, 625–636

*SRv6 configuration*, 623–624

*VRF and interface configuration*,  
622

## G

---

GRE over L3VPN, 26

greenfield deployment, 354

dual-homed migration strategy, 356

interworking migration strategy, 355–356

SR-MPLS, 365

*BGP proxy Prefix SID*, 383–387

*enabling LDP on the border node*,  
372–376

*enabling SRMS*, 365–372

*enabling the BGP Prefix SID*,  
376–383

SRv6

*building a new network using an  
IWG*, 389–401

*building a new network using  
dual-connected PE devices*,  
413–424

*building a new network using  
inter-AS Option A*, 401–413

GUA (global unicast address), 162, 164

## H

---

hash, n-tuple, 25

headend behaviors, 132

header

IPv6, 104

*Flow Label field*, 104–106

*Next Header field*, 106

*Traffic Class field*, 104

MPLS, 6–7

segment routing, 36–37, 106–107,  
125–127. *See also* SRH (segment  
routing header)

*fields*, 124–125

*Penultimate Segment Pop*,  
127–128

*Ultimate Segment Pop*, 129

*USD (Ultimate Segment  
Decapsulation)*, 130–132

help sr localsid command, ONLINE-  
ONLINE

H.Encaps behavior, 132–133

H.Encaps.L2 behavior, 135

H.Encaps.L2.Red behavior, 136–137

H.Encaps.Red behavior, 133–135

high-availability, 425, 857–858. *See  
also* BFD (Bidirectional Forwarding  
Detection); FRR (Fast Reroute);  
LFA (Loop-Free Alternate); TI-LFA

- (Topology-Independent Loop-Free Alternate)
  - active-active, 425
  - active-backup, 426
  - BFD (Bidirectional Forwarding Detection), 857–859
    - async mode*, 859–860
    - BLB*, 861
    - BoB*, 861, 863
    - demand mode*, 859
    - echo*, 859–860
    - over LAG interfaces*, 860
    - packet capture*, 861–862
    - SH (single-hop) sessions*, 860–861
  - FRR (Fast Reroute), 858
  - LFA (Loop-Free Alternate), 879–880
  - load sharing, 426
  - TI-LFA (Topology-Independent Loop-Free Alternate), 883
  - high-level programming language, 33–34
  - high-precision probing, 1020–1023
  - H.Insert behavior, 138–140
  - hub-and-spoke L3VPN
    - over SR-MPLS
      - configuration*, 719–732
      - verification*, 732–751
    - over SRv6
      - configuration*, 636–645
      - verification*, 645–657
  - hw-module bfd-hw-offload enable location command, 867
  - hw-module command, 240
- 
- IGP (interior gateway protocol), 10
    - link-state. *See* IS-IS; OSPF (Open Shortest Path First)
    - PCE (path computation element), 211
    - Prefix SID, 45–47
      - segments, 36
      - summarization, 220
  - IMIX (Internet Mix), ONLINE
  - integration test, 1080, 1091
  - intent-driven configuration, 1018
  - intent-/model-based assurance, 1019–1020
  - inter-AS BGP-LU, 341–342, 343
    - configuration, 344–345
    - data forwarding, 349–350
    - design, 343
    - verification, 345–348
  - inter-AS routing, 21
    - Option A, 401–413
    - Option C, 431–433
  - interface, loopback, 163–164
  - interface loopback address, SRv6, 237–238
  - interworking greenfield strategy, 355–356
  - intra-AS BGP-LU
    - BGP Additional Path feature, 331–332
    - configuration, 332–337
    - design, 330
    - with Prefix SID, 328–330
    - verification, 337–341
  - IOS XE, 251–253
    - advertising the BGP Prefix SID, 326
    - assigning Prefix SID to Loopback0, 231
    - BGP-LU session on ASBR-3, 344
    - configuring anycast SID, 256
    - enabling segment routing, 230
    - enabling segment routing for IS-IS, 246
    - enabling segment routing for OSPF, 250
    - SR-MPLS IS-IS verification, 247–248
    - verifying Prefix SID assignment, 236
    - verifying the advertisement of the anycast SID, 257
  - IOS XR
    - advertise the SRv6 locator, 307

- advertising the BGP Prefix SID, 326
- assigning an SRv6 locator to flex algo 128, 307
- assigning link delay, 307
- assigning Prefix SID to Loopback0, 231
- BGP-LU session on ASBR-1, 344
- Cisco EVC framework, 446–447
- configuring anycast SID, 255–256
- configuring anycast SRv6 locator on P-3, 271
- configuring IS-IS on ABR-1, 263
- configuring IS-IS on P-1, 262–263
- configuring IS-IS on PE-1, 263–264
- configuring IS-IS on RR-1, 264–265
- enabling segment routing, 230
- enabling segment routing for IS-IS, 246
- enabling segment routing for OSPF, 250
- enabling SRv6, 240
- enabling UPA processing, 298
- increasing the maximum number of UPAs, 297
- modifying the UPA lifetime, 296
- prefix tagging for UPA generation, 297
- propagating UPAs from IS-IS to Level 1, 298
- selectively generating UPAs, 297
- SR-MPLS IS-IS verification, 247
- SR-MPLS OSPF verification, 250–251
- SRv6-TE policy and traffic steering using color, 273–274
- summary route advertisement, 286
- summary route for flex algo 128, 308
- suppressing default route generation, 286
- UPA announcement, 296
- verifying Prefix SID assignment, 235
- verifying the advertisement of the anycast SID, 257
- IoT (Internet of things), 37**
- IP (Internet Protocol), 28**
- IPFRR (IP Fast Reroute), 878–881**
- iproute2, ONLINE-ONLINE, ONLINE, ONLINE, ONLINE-ONLINE**
- IPSLA, 834**
  - configuration, 836–843
  - verification, 843–851
- IPv6**
  - backhaul brownfield strategy, 356–357
  - GUA (global unicast address), 162, 164
  - header
    - Flow Label field, 104–106*
    - Next Header field, 106*
    - Traffic Class field, 104*
  - LLA (link local address), 162, 164–165
  - ULA (unique local address), 162
- ISA (instruction set architecture), 35**
- IS-IS. *See also* SR-MPLS, IS-IS; SRv6, IS-IS**
  - areas, 222
  - extensions for segment routing (RFC 8667), 53–54
    - Adj-SID sub-TLV, 59–61*
    - Flex Algo sub-sub-TLVs, 79–83*
    - LAN-Adj-SID sub-TLV, 61*
    - Prefix-SID sub-TLV, 57–59*
    - Router Capability TLV, 54–57*
    - Segment Routing Algorithm sub-TLV, 55–56*
    - SID/Label Binding TLV, 61–62*
    - SID/Label sub-TLV, 63–64*
    - SRLB sub-TLV, 56–57*
  - extensions for SRv6, 175–176
    - segment routing over IPv6 capabilities, 180–183*
    - segment routing over IPv6 locator, 176–180*
    - SIDs, 183–186*
  - levels, 222
  - LSP (label switched path), 179–180
  - overload bit, 223
  - route leaking, 223

- route propagation, 223
- router types, 222
- routing, 222
- verifying Prefix SID assignment, 235

## ITU-T Y.1731 Performance Measurement

- delay and jitter measurement
  - configuration*, 818–820
  - verification*, 820–825
- SLM (synthetic loss measurement), 816
  - configuration*, 826–827
  - verification*, 827–834

# L

---

## L2VPN, 3, 27. *See also* EVPN (Ethernet VPN)

- pseudowire, 11–12
- service assurance, 806–807
  - CFM (Ethernet Connectivity Fault Management)*, 808–816. *See also* CFM (Ethernet Connectivity Fault Management)
  - ITU-T Y.1731 Performance Measurement*. *See* ITU-T Y.1731 Performance Measurement

## L3VPN, 2, 605–606

- extranet
  - over SR-MPLS, configuration*, 752–756
  - over SR-MPLS, verification*, 756–767
  - over SRv6*, 656
  - over SRv6, configuration*, 657–662
  - over SRv6, verification*, 662–677
- Flex Algo, 318–322
- full-mesh, over SR-MPLS
  - BGP label allocation modes*, 701
  - configuration*, 689–700
  - verification*, 701–719
- full-mesh, over SRv6

- BGP configuration on PE-2*, 624–625
- configuration*, 610–622
- SRv6 configuration*, 623–624
- verification*, 625–636
- VRF and interface configuration*, 622

- GRE over, 26
- hub-and-spoke, over SR-MPLS
  - configuration*, 719–732
  - verification*, 732–751
- hub-and-spoke, over SRv6, 636
  - configuration*, 636–645
  - verification*, 645–657
- interoperability, ONLINE-ONLINE
  - Cisco Catalyst 8000V Edge IPv4 L3VPN service*, ONLINE-ONLINE
  - FRR IPv4 L3VPN service*, ONLINE-ONLINE
- over SR-MPLS, 677–679
  - configuration*, 679–685
  - verification*, 686–688
- over SRv6, 608
  - BGP Prefix SID path attribute*, 609–610
  - MP\_REACH\_NLRI path attribute*, 608
  - NLRI*, 608
- service assurance, 834, 847–850
  - IPSLA*, 834. *See also* IPSLA
  - TWAMP*, 835–836. *See also* TWAMP (Two-Way Active Measurement Protocol); TWAMP Light
- lab environment, 1089–1093
  - development, 1096–1097
  - preparation and construction, 1093–1096
- label/s, 1, 7–9
  - allocation, 12–13
  - based forwarding, 15–16

- collisions, 43
- POP operation, 14
- PUSH operation, 14
- service, 14, 16
- space limitation, 20–21
- SWAP operation, 14
- LAG (link aggregation group), 25, 105**
- LAN-Adj-SID sub-TLV, 61**
- LDP (Label Distribution Protocol), 1, 12.**
  - See also mLDP*
  - downstream on-demand, 11
  - enabling on the border node, 372–376
  - ID, 11
  - IGP synchronization, 24
  - label/s
    - allocation, 12–13*
    - based forwarding, 15–16*
    - service, 14, 16*
    - space limitation, 20–21*
  - neighbor discovery, 11–12
  - unsolicited downstream, 11
- legacy network. *See also domain***
  - replacing or enhancing with SR
    - IT evolution and gap awareness, 1051–1056*
    - migration strategy, 1049–1051*
    - required knowledge and expertise, 1046–1049*
  - SR consolidation, 1056–1074
- LER (label edge router), 5**
- levels, IS-IS, 222
- LFA (Loop-Free Alternate), 878–880**
- LFIB (Label Forwarding Information Base), 10**
- Linktrace Protocol, 811**
- Linux SRv6/Linux kernel, ONLINE-ONLINE**
  - Containerlab topology definition, ONLINE-ONLINE
  - hardware offloading, ONLINE-ONLINE
  - lab deployment, ONLINE-ONLINE
    - Containerlab, ONLINE-ONLINE*
    - Layer 3 overlay services, ONLINE*
    - Linux IPv4 L3VPN service, ONLINE-ONLINE*
    - Linux IPv4/IPv6 L3VPN service, ONLINE-ONLINE*
    - Linux IPv6 L3VPN service, ONLINE-ONLINE*
    - Linux point-to-point L2VPN service, ONLINE-ONLINE*
    - overlay services VPN-1, ONLINE-ONLINE*
    - overlay services VPN-2, ONLINE-ONLINE*
    - overlay services VPN-3, ONLINE-ONLINE*
    - underlay connectivity, ONLINE-ONLINE*
    - underlay domain transport, ONLINE-ONLINE*
  - network stack, ONLINE-ONLINE
  - RPD (routing policy database), ONLINE-ONLINE, ONLINE-ONLINE
  - SRv6 endpoint support, ONLINE-ONLINE
  - SRv6 headend support, ONLINE-ONLINE
  - SRv6 overlay services configuration, ONLINE-ONLINE
  - uA/End.X (NEXT-CSID) endpoint behavior, ONLINE
- LLA (link local address), 162, 164–165**
- load sharing, 426**
- locator addressing scheme, SRv6, 165–169**
  - large-scale deployments, 170–172
  - small- and medium-scale deployments, 169–170
  - uSID, 238–239
- loopback interface, 163–164**
- Loopback Protocol, 810**

LSA (link-state advertisement), 224  
 LSD (label switching database), 43, 44  
 LSP (label switched path), 21  
   echo request packet, 790–791  
   IS-IS, 179–180  
   labels, 380  
 LSR (label switch router), 5

## M

---

MA (maintenance association), 808  
 MAC VRF instance, 444  
 MCD (Midpoint Compressed Data), 800–801  
 MD (maintenance domain), 808  
 MEF (Metro Ethernet Forum), 441–442  
 message/s  
   BGP Route Target Constraint, 770  
   BGP UPDATE, 83–84, 86–87, 96–98, 186, 187–188, 190, 191, 203–204  
     *for SRv6 L2 services, 192–193*  
     *for SRv6 L3 services, 194–195, 196–198*  
   SR-MPLS DPM syslog, 798  
 methodology, 1084  
 metrics, Flex Algo, 78  
 Microloop Avoidance, 935, 942–943  
   local protection, 935–940  
   remote protection, 940–941  
 microloop avoidance rib-update-delay 5000 command, 942–943  
 microloops, 935  
 migration test, 1096  
 migration to segment routing, 353–354.  
   *See also* brownfield deployment;  
   greenfield deployment  
   coexistence brownfield strategy, 356–357  
   deployment models  
     *brownfield, 354*  
     *greenfield, 354*  
   dual-homed greenfield strategy, 356  
   interworking greenfield strategy, 355–356  
   IPv6 backhaul brownfield strategy, 356–357  
   MPLS to SRv6, 427–429  
     *CSC MPLS network, 434–435*  
     *flat MPLS network, 429–430*  
     *MPLS network with inter-AS option C, 431–433*  
     *unified MPLS network, 430–432*  
 SR-MPLS, 358  
 SRv6, 387–389  
   *building a new network using an IWG, 389–401*  
   *building a new network using dual-connected PE devices, 413–424*  
   *building a new network using inter-AS Option A, 401–413*  
 MIP (maintenance intermediate point), 809  
 mLDP, 3  
 MP\_REACH\_NLRI path attribute, 608–609  
 MP-BGP (Multiprotocol Border Gateway Protocol), 6, 17  
   extension/s, 83–85  
     *BGP Link-State, 87–95*  
     *BGP Link-State extensions for SR BGP Egress Peer Engineering (RFC 9086), 95–98*  
     *BGP overlay services on SRv6, 186–201*  
     *BGP Prefix SID, 85–87*  
 MPLS (Multiprotocol Label Switching), 4, 6, 28, 440. *See also* LDP (Label Distribution Protocol)  
   backbone routes, 5  
   BGP-free core, 1  
   CE (customer edge) router, 6  
   control plane, 205–207



- data plane, 205–207
- FRR (fast reroute), 18–19
- header, 6–7
- L2VPN, 3, 27. *See also* L2VPN
- L3VPN, 2, 26. *See also* L3VPN
- label/s, 1, 7–9
  - based forwarding, 15–16
  - POP operation, 14
  - PUSH operation, 14
  - service, 16
  - space limitation, 20–21
  - SWAP operation, 14
- LDP (Label Distribution Protocol), 11, 12
  - downstream on-demand, 11
  - IGP synchronization, 24
  - unsolicited downstream, 11
- LER (label edge router), 5
- LSP (label switched path), 21
- LSR (label switch router), 5
- mLDP, 3
- operational complexity, 22
- pseudowire, 440
- QoS, end-to-end, 21–22
- RSVP-TE, 22–23
  - limitations, 23–24
  - tunnel, 22–23
- services, 2
- traffic protection, 18
- Unified, 210
- use cases, 3–4
- VPN, 16–18, 20
- VRF instance, 3, 5
- MPLS TE (Traffic Engineering), 3–4
- mVPN, 3

## N

---

- neighbor adjacency, LDP, 11
- neighbor discovery, LDP (Label Distribution Protocol), 11–12

- network, as a computer, 37
- network event information propagation time, 858
- Next Header field, IPv6 header, 106
- NEXT operation, 41
- NLRI, 201–203, 608
- Node SID, 45–46, 226
- nonrouted SID, 110–111
- n-tuple hash, 25

## O

---

- open-source SRv6, [ONLINE](#), [ONLINE-ONLINE](#). *See also* open-source SRv6 lab deployment
  - FRR (Free Range Routing), [ONLINE-ONLINE](#)
    - daemons, [ONLINE-ONLINE](#)
    - release 9.1, [ONLINE-ONLINE](#)
    - SRv6 headend and endpoint support, [ONLINE-ONLINE](#)
    - system architecture, [ONLINE](#)
    - vytish, [ONLINE-ONLINE](#)
- Linux kernel, [ONLINE-ONLINE](#)
  - hardware offloading, [ONLINE-ONLINE](#)
  - network stack, [ONLINE-ONLINE](#)
  - SRv6 endpoint support, [ONLINE-ONLINE](#)
  - SRv6 headend support, [ONLINE-ONLINE](#)
  - uA/End.X (NEXT-CSID) endpoint behavior, [ONLINE](#)
- SONiC (Software for Open Networking in the Cloud), [ONLINE-ONLINE](#)
  - cli, [ONLINE](#)
  - container modules, [ONLINE-ONLINE](#)
  - data plane, [ONLINE](#)
  - disaggregation, [ONLINE-ONLINE](#)
  - sonic-cfggen, [ONLINE](#)

- system architecture, ONLINE-ONLINE*
- VPP (Vector Packet Processor), ONLINE-ONLINE
  - data plane, ONLINE-ONLINE*
  - endpoint support, ONLINE-ONLINE*
  - headend support, ONLINE-ONLINE*
  - linux-cp control plane integration plug-in, ONLINE-ONLINE*
  - pps (packets per second), ONLINE-ONLINE*
  - vector packet processing graph, ONLINE-ONLINE*
- open-source SRv6 lab deployment
  - L3VPN interoperability, ONLINE-ONLINE
    - Cisco Catalyst 8000V Edge IPv4 L3VPN service, ONLINE-ONLINE*
    - FRR IPv4 L3VPN service, ONLINE-ONLINE*
  - Linux SRv6, ONLINE-ONLINE
    - Containerlab, ONLINE-ONLINE*
    - Layer 3 overlay services, ONLINE*
    - Linux IPv4 L3VPN service, ONLINE-ONLINE*
    - Linux IPv4/IPv6 L3VPN service, ONLINE-ONLINE*
    - Linux IPv6 L3VPN service, ONLINE-ONLINE*
    - Linux point-to-point L2VPN service, ONLINE-ONLINE*
    - overlay services VPN-1, ONLINE-ONLINE*
    - overlay services VPN-2, ONLINE-ONLINE*
    - overlay services VPN-3, ONLINE-ONLINE*
    - underlay connectivity, ONLINE-ONLINE*
    - underlay domain transport, ONLINE-ONLINE*
- VPP (Vector Packet Processor), ONLINE-ONLINE
  - basic setup, ONLINE-ONLINE*
  - IPv4 L3VPN service, ONLINE-ONLINE*
  - IPv6 L3VPN service, ONLINE-ONLINE*
  - point-to-point L2VPN service, ONLINE-ONLINE*
  - underlay connectivity, ONLINE-ONLINE*
- operator-defined algorithm, 74–75
- optical transport network (OTN), 73–74
- OSPF (Open Shortest Path First). *See also* SR-MPLS, OSPF; SRv6, OSPF
  - ABR (area border router), 224, 229
  - areas, 224
  - extensions for segment routing (RFC 8665), 64
    - Adj-SID sub-TLV, 68–69*
    - Prefix Range TLV, 72–73*
    - Prefix-SID sub-TLV, 70–71*
    - Segment Routing Algorithm TLV, 65*
    - SID/Label Range TLV, 65–66*
    - SRLB TLV, 66–67*
    - SRMS Preference TLV, 67*
  - LSA (link-state advertisement), 224
  - shortest path tree, 224
  - SPF (Shortest Path First) algorithm, 224
  - verifying Prefix SID assignment, 235
- OSPFv3, 225
  - route filtering, 225
  - route summarization, 225
- overlay, 1018
- overload bit, IS-IS, 223

# P

---

P router, 6

path attribute. *See also* BGP

BGP Prefix SID, 609–610

MP\_REACH\_NLRI, 608–609

path divergence, 788–789

pcap trace command, ONLINE-ONLINE

PCE (path computation element), 211

PE router, 5, 606, 767–768

RTC (route target constraint), 768–771

*configuration*, 771–774

*memberships*, 769

*verification*, 774–780

Penultimate Segment Pop, 127–128

per-ESI Ethernet A-D route, 454–458

per-EVI Ethernet A-D route, 458–463

PHP (penultimate hop popping), 39

PIC (Prefix Independent Convergence),  
943, 944–945

PIC Edge

multipath verification, SRv6, 981–995

unipath

*SR-MPLS, configuration*, 948–951

*SR-MPLS, verification*, 951–962

*SRv6, configuration*, 962–970

*SRv6, verification*, 970–981

ping command, 249, 254, 270, 292, 349,  
423, 635, 676, 767, ONLINE-ONLINE,  
ONLINE-ONLINE, ONLINE-ONLINE,  
ONLINE-ONLINE, ONLINE-ONLINE

ping ethernet cfm domain service  
command, 815

pipeline, 1101–1104

PLE (private line emulation), 1011–1017

PLR (point of local repair), 878–879

policy

segment routing, 38–40, 41, 52–53

SRv6, 126

*End.B6.Encaps behavior*, 142–143

*End.B6.Encaps.Red behavior*,  
143–144

*End.B6.Insert behavior*, 144–145

*End.B6.Insert.Red behavior*,  
145–146

*endpoint behaviors*, 141, 146

*headend behaviors*, 132, 140

*H.Encaps behavior*, 132–133

*H.Encaps.L2 behavior*, 135

*H.Encaps.L2.Red behavior*,  
136–137

*H.Encaps.Red behavior*, 133–135

*H.Insert behavior*, 138–140

POP operation, 14

pps (packets per second), ONLINE-  
ONLINE

Prefix Attribute Flags sub-TLV, 178–179

prefix metric, Flex Algo, 303

Prefix Range TLV, 72–73

Prefix SID, 76–77, 226

BGP, 324

*configuration*, 324–326

*enabling in an SR-MPLS network*,  
376–383

*proxy*, 383–387

*verification*, 327–328

verifying, 235–236

Prefix SID TLV, 95

Prefix-SID sub-TLV, 58–59, 70–71

primary path, 18

program counter, 35

protocol ID, BGP, 96

pseudowire, 3, 11–12, 27, 440

P-space, 881–882, 896–897

PT (Path Tracing), 784, 787, 798–801,  
1023–1025

MCD (Midpoint Compressed Data),  
800–801

probe packets, 801–806

PUSH operation, 14, 41

## Q-R

---

QoS, end-to-end, 21–22

Q-space, 881–882, 896–897

RD (route distinguisher), 17

reachability

SR-MPLS, verification, 248–249,  
253–254

SRv6, verification, 270

UPA (Unreachability Prefix  
Announcement), 293–295  
*configuration*, 295–298  
*verification*, 298–301

resource assurance, 1062–1063

RFC 4385, 26, 27

RFC 4760, 84

RFC 5036, 11

RFC 5357, 785

RFC 6391, 26

RFC 6513, 3

RFC 6790, 26, 26

RFC 7432, 449, 454

RFC 7855, 40–41

RFC 8317, 552

RFC 8402, 87–88

RFC 8667, IS-IS extensions for segment  
routing, 53–57, 256

RFC 8668, 48, 93

RFC 8754, 123–132

RFC 8762, 785

RFC 8986, 107, 111

RFC 9085, 87–95

RFC 9086, 95–98

RFC 9252, 186–201

RFC 9350, 74

RFC 9352, 175–186

RFC 9356, 48

RFC 9417, 1020

RIB (Routing Information Base), 10, 12,  
858

RLFA (Remote Loop-Free Alternate),  
880–882

RONs (routed optical networks),  
1008–1011

route filtering, OSPFv3, 225

route leaking, IS-IS, 223

route propagation, IS-IS, 223

route summarization. *See also*  
summarization

OSPFv3, 225

SRv6, 172–175

Route Type 1

per-ESI Ethernet A-D route, 454–458

per-EVI Ethernet A-D route, 458–463

Route Type 2, 463–466

Route Type 3, 467–469, 512

Route Type 4, 470–472

routed SID, 110–111

router

area border, 5, 38, 224

autonomous system border, 5, 20

CE (customer edge), 5, 6

IS-IS, 222

label edge, 5

label switch, 5

P (provider), 6

PE (provider edge), 5, 606, 767–768

PLR (point of local repair), 878–879

P-space, 881–882, 896–897

Q-space, 881–882, 896–897

Router Capability TLV, 54–57

routing, inter-AS, 21

RPD (routing policy database), ONLINE-  
ONLINE, ONLINE-ONLINE

RR (route reflector), 6

RTC (route target constraint), 768–771

*configuration*, 771–774

*memberships*, 769

*verification*, 774–780

## RSVP-TE, 22–23

limitations, 23–24

tunnel, 22–23

## RT (route target), 17–18, 389, 390

## RTC (route target constraint), 768–771

configuration, 771–774

memberships, 769

verification, 774–780

# S

---

## SAP (service access point), 1058–1059

## scalability, 220

## SDN (software-defined networking), 28–29

## segments, 36

## Segment Routing Algorithm sub-TLV, 55–56

## Segment Routing Algorithm TLV, 65

## Segment Routing Capability sub-TLV, 54–55

## segment-routing mpls sr-prefer command, 361

## sensors, IoT, 37

## service assurance, 783, 1062–1063

### L2VPN, 806–807

*CFM (Ethernet Connectivity Fault Management)*, 808–816. *See also* *CFM (Ethernet Connectivity Fault Management)*

*ITU-T Y.1731 Performance Measurement. See ITU-T Y.1731 Performance Measurement*

### L3VPN, 834

*IPSLA*, 834. *See also* *IPSLA*

*TWAMP*, 835–836. *See also* *TWAMP (Two-Way Active Measurement Protocol)*; *TWAMP Light*

transport-related. *See* transport-related service assurance

## service catalog, 1063

## service chaining, 213

## service label, 14, 16

## service orchestration, 1064

## service portfolio consolidation, 1083–1084

## service provider/s

benefits of SRv6 adoption, 998–999

network evolution, 210–212

PE (provider edge) routers, 606. *See also* PE router

technological opportunities and benefits of SR

*CapEx savings*, 1026–1029

*fewer protocols*, 999–1001

*integrated visibility*, 1017–1018

*more QoS options*, 1001–1003

*new hardware generation*, 1025

*OpEx savings*, 1030–1032

*PLE (private line emulation)*, 1011–1017

*RONs (routed optical networks)*, 1008–1011

*scale*, 1007–1008

*traffic engineering and network slicing*, 1005–1007

*unification across domains*, 1003–1005

## service vpp status command, ONLINE-ONLINE

## service vpp stop command, ONLINE-ONLINE

## show bfd ipv6 session command, 868–869

## show bfd ipv6 session interface bundle-Ether 1 detail command, 869

## show bfd ipv6 session interface TF0/0/0/0 detail command, 870–871

## show bfd neighbors command, 877

show bfd neighbors interface port-channel  
     1 details command, 877–878  
 show bfd session command, 875  
 show bfd session interface bundle-ether 1  
     detail command, 875–876  
 show bgp ipv4 labeled-unicast command,  
     347, 348, 377–378, 379–381, 713,  
     740–741  
 show bgp ipv4 rt-filter command, 771,  
     774  
 show bgp ipv4 rt-filter neighbors  
     command, 774–775  
 show bgp ipv4 unicast command, 328,  
     348, 713–714  
 show bgp ipv4 vpn command, ONLINE-  
     ONLINE  
 show bgp l2vpn evpn bridge-domain  
     200-BD command, 527–529  
 show bgp l2vpn evpn bridge-domain  
     200-BD received-sids wide command,  
     523–524  
 show bgp l2vpn evpn bridge-domain  
     command, 446, 521–522, 525–527,  
     537, 540, 549–550, 568–569  
 show bgp l2vpn evpn bridge-domain  
     VPWS:300 command, 594–599  
 show bgp l2vpn evpn rd command,  
     458–461, 464, 467–469, 470, 523,  
     601  
 show bgp l2vpn evpn route-type  
     command, 530  
 show bgp l2vpn evpn route-type ethernet-  
     segment command, 530–532, 600  
 show bgp l2vpn evpn summary command,  
     518–519  
 show bgp segment-routing srv6  
     command, ONLINE-ONLINE  
 show bgp summary command, ONLINE-  
     ONLINE  
 show bgp vpnv4 uni vrf command,  
     ONLINE-ONLINE  
 show bgp vpnv4 unicast command,  
     741–742  
 show bgp vpnv4 unicast vrf command,  
     626, 628–630, 646–647, 649–650,  
     651–652, 655, 659–661, 663,  
     667–668, 670, 673–674, 702–703,  
     704–705, 705–706, 733–735,  
     758–759, 760, 951–953, 954,  
     958–960, 971–973, 976–977,  
     982–984, 989–990  
 show bgp vpnv4 unicast vrf local-sids  
     command, 627, 647–648, 663–664,  
     671  
 show bgp vpnv4 unicast vrf received-sids  
     command, 627–628, 648, 664, 671  
 show bgp vrf command, 276, 320–321,  
     409–410, 420–421  
 show bgp vrf nexthop-set command, 634,  
     650, 667  
 show bgp vrf-db table all command,  
     519–521, 594  
 show bgp vrf-db table command,  
     566–567, 656  
 show bpg vpnv4 unicast vrf command,  
     665–666  
 show bundle bundle-Ether 200 command,  
     478, 484  
 show bundle bundle-Ether 250 command,  
     481, 482  
 show bundle bundle-Ether 300 command,  
     575–576  
 show cef command, 887–888, 892–893,  
     898–899, 919, 926  
 show cef detail command, 714–715, 716,  
     717–718, 743, 744–745, 747, 748,  
     749–750  
 show cef ipv6 command, 904–905,  
     908–909, 912–913, 936–937,  
     937–940, 942  
 show cef ipv6 detail command, 980–981,  
     993–994  
 show cef vrf command, 276, 279  
 show cef vrf detail command, 632–633,  
     654, 669–670, 675, 709–711, 738,  
     764, 766, 955–957, 960–962,  
     974–975, 979, 986–988, 991–992

- show ethernet cfm local meps domain service command, 813
- show ethernet cfm peer meps domain service command, 814–815
- show ethernet sla operations detail profile command, 820–821, 827–828
- show ethernet sla statistics brief profile command, 821–823, 829–830
- show ethernet sla statistics detail profile command, 823–825, 831–834
- show evpn ethernet-segment command, 494, 580
- show evpn ethernet-segment interface BE200 carving detail command, 496–498
- show evpn ethernet-segment interface BE300 carving detail command, 581–582
- show evpn ethernet-segment interface Bundle-Ether260 detail command, 538–539
- show evpn evi command, 498, 540–542, 583–584
- show evpn evi ead command, 501, 584–585
- show evpn evi inclusive-multicast detail command, 502–503
- show evpn evi vpn 200 mac command, 501–502
- show evpn evi vpn 205 mac command, 560–561
- show evpn evi vpn-id 200 detail command, 499–500
- show evpn summary command, 493–494
- show interfaces brief command, 477–478, 483, 574–575, ONLINE-ONLINE
- show ip cef detail command, 715, 717, 718, 744, 745–746, 749
- show ip cef internal command, 889–890, 894–895, 900–901, 930–931
- show ip cef vrf command, ONLINE-ONLINE
- show ip cef vrf detail command, 711, 739, 765, 957
- show ip ospf database command, 234
- show ip route repair-paths command, 888–889, 893, 899–900, 928, 954–955
- show ip route vrf command, 737, 762
- show ip route vrf repair-paths command, 709
- show ipsla statistics 11 command, 843–844
- show ipsla statistics 21 command, 844–846
- show ipsla statistics aggregated 11 command, 846–847
- show ipsla statistics aggregated 21 command, 847–850
- show ipsla twamp session command, 854
- show ipv6 cef command, ONLINE-ONLINE
- show ipv6 route isis command, ONLINE-ONLINE
- show isis adjacency command, 265–266
- show isis database command, 233, 247, 257, 265, 272, 288, 291, 299, 309–310
- show isis fast-reroute summary command, 884
- show isis fast-reroute ti-lfa tunnel command, 894, 900, 929
- show isis flex-algo command, 308
- show isis ipv4 fast-reroute detail command, 887, 892, 898, 904, 918, 920–921, 925, 927–928, 933–934
- show isis ipv6 fast-reroute detail command, 908, 912
- show isis neighbor command, ONLINE-ONLINE
- show isis rib command, 929–930
- show isis srv6 locators det command, ONLINE-ONLINE
- show l2vpn bridge-domain bd-name 200-BD detail command, 508–510
- show l2vpn bridge-domain bd-name detail command, 563–565



- show l2vpn bridge-domain brief  
command, 508, 511, 544
- show l2vpn forwarding bridge-domain  
ELAN-BG:200-BD mac-address  
location 0/RP0/CPU0 command,  
510–511
- show l2vpn forwarding bridge-domain  
ELAN-BG:210-BD mac-address  
location 0/RP0/CPU0 command, 545
- show l2vpn forwarding xconnect detail  
location command, 591
- show l2vpn mac-learning command, 512
- show l2vpn mac-learning mac all location  
command, 511
- show l2vpn xconnect group VPWS-XC  
command, 589–591
- show lacp bundle-Ether 250 command,  
481, 482
- show lacp command, 479–480, 576–577
- show memif command, ONLINE-ONLINE
- show mpls forwarding command, 362,  
364, 374
- show mpls oam dpm adjacency command,  
797
- show mpls oam dpm prefix command,  
797
- show mpls oam dpm summary command,  
795–796
- show route 10.0.1.8/32 command,  
886–887
- show route command, 891, 896–897,  
917–918, 924–925
- show route ipv6 command, 280–281,  
282, 289, 290, 291–292, 300,  
311–312, 903–904
- show route ipv6 detail command,  
906–907, 911–912
- show route vrf command, 278
- show route vrf detail command, 631–632,  
653–654, 668–669, 674–675, 708,  
736, 761, 762–763, 973–974,  
977–978, 985–986
- show segment-routing srv6 capabilities-  
parameters command, ONLINE-  
ONLINE
- show segment-routing srv6 locator  
command, ONLINE-ONLINE,  
ONLINE-ONLINE
- show segment-routing srv6 locator MAIN  
detail command, ONLINE-ONLINE
- show segment-routing srv6 locator MAIN  
sid command, 513–514, 546–547
- show segment-routing srv6 sid command,  
503–504, 586, ONLINE-ONLINE,  
ONLINE-ONLINE
- show segment-routing srv6 sid detail  
command, 633, 634–635, 656
- show segment-routing traffic-eng policy  
command, 274–275
- show slrg name command, 932
- show sr localsids command, ONLINE-  
ONLINE, ONLINE-ONLINE
- show sr policies command, ONLINE-  
ONLINE
- show sr steering-policies command,  
ONLINE-ONLINE
- show trace command, ONLINE-ONLINE,  
ONLINE-ONLINE, ONLINE-ONLINE
- SID (segment identifier), 36–37,  
42–43
  - Adjacency, 47–49, 226
  - allocation, 43–44, 115
  - anycast, 45, 47, 219
  - BGP Peering, 50–52
  - BGP Prefix, 49–50, 85–87
  - Binding Segment, 52–53
  - block addressing considerations,  
163
  - compression, 161–162
  - Node, 38–39
  - Prefix, 45–47, 76–77
  - SRv6, 107–108, 183–186
    - End behavior*, 113–114
    - End.DT2M behavior*,  
122–123
    - End.DT2U behavior*, 120–121
    - End.DT4 behavior*, 116–117



- End.DX2 behavior*, 119
- End.DX4 behavior*, 117–118
- endpoint behaviors*, 1113
- End.X behavior*, 114–115
- Locator field*, 108–110
- routed/nonrouted*, 110–111
- verification*, 266–270
- SID/Label Binding TLV, 61–62
- SID/Label Range TLV, 65–66
- SID/Label sub-TLV, 63–64
- SID/Label TLV, 91
- SLA (service-level agreement), 783–784, 806
- SLM (synthetic loss measurement), 816
  - configuration, 826–827
  - verification, 827–834
- SLRG (shared risk link group), 921. *See also* TI-LFA (Topology-Independent Loop-Free Alternate), SLRG protection
- SmartNICs, ONLINE
- software-defined networking, 28–29
- SONiC (Software for Open Networking in the Cloud), ONLINE-ONLINE
  - cli, ONLINE
  - container modules, ONLINE-ONLINE
  - data plane, ONLINE
  - disaggregation, ONLINE-ONLINE
  - sonic-cfggen, ONLINE
  - system architecture, ONLINE-ONLINE
- source code repository, 1098–1100
- source node, 126
- source routing, 35–36
- source/ingress node, 38
- SPF (Shortest Path First) algorithm, 56, 73–74, 220, 224
- SPRING architecture, 40–41
- SR (segment routing), 225. *See also* SR-MPLS
  - adjacency segment, 226
  - benefits, 212–214
  - BGP Link-State extensions, 87–95
  - business case guidance, 1032–1034, 1036–1037
    - opportunity analysis*, 1037–1038
    - refining known CapEx*, 1034–1035
    - refining known OpEx*, 1035–1036
  - data plane, 36–37
  - domain, 38
  - enabling
    - on IOS XE*, 230
    - on IOS XR*, 230
  - endpoint/egress node, 38
  - EPE (egress peer engineering), 41
  - feature support, 215
  - header, 106–107
  - IETF standards, 39, 40
  - IS-IS extensions (RFC 8667), 53–54
    - Adj-SID sub-TLV*, 59–61
    - Flex Algo sub-sub-TLVs*, 79–83
    - LAN-Adj-SID sub-TLV*, 61
    - Prefix-SID sub-TLV*, 57–59
    - Router Capability TLV*, 54–57
    - Segment Routing Algorithm sub-TLV*, 55–56
    - SID/Label Binding TLV*, 61–62
    - SID/Label sub-TLV*, 63–64
    - SRLB sub-TLV*, 56–57
  - network as a computer, 37
  - OSPF extensions (RFC 8665), 64
    - Adj-SID sub-TLV*, 68–69
    - Prefix Range TLV*, 72–73
    - Prefix-SID sub-TLV*, 70–71
    - Segment Routing Algorithm TLV*, 65
    - SID/Label Range TLV*, 65–66
    - SRLB TLV*, 66–67
    - SRMS Preference TLV*, 67
  - policy, 38–40
  - policy/ies, 41, 52–53

- replacing or enhancing a legacy network with. *See also* legacy network
  - IT evolution and gap awareness, 1051–1056*
  - migration strategy, 1049–1051*
  - required knowledge and expertise, 1046–1049*
- SID (segment identifier), 36–37
- source/ingress node, 38
- SPRING architecture, 70–77
- transit node, 38
- SR-DPM (Segment Routing Data Plane Monitoring), 784, 788–789**
  - adjacency verification and validation, 790–792
  - configuration, 789–790
  - prefix reachability verification, 793–798
- SRGB (segment routing global block), 44–45, 226–227**
  - reconfiguring, 229
  - verifying, 232–234
- SRH (segment routing header), 36–37, 123–124, 125–127, 147**
  - fields, 124–125
  - Penultimate Segment Pop, 127–128
  - Ultimate Segment Pop, 129
  - USD (Ultimate Segment Decapsulation), 130–132
  - uSID instruction, 147–148
- SRLB (segment routing local block), 45, 227–228**
  - reconfiguring, 229
  - verifying, 232–234
- SRLB sub-TLV, 56–57**
- SRLB TLV, 66–67**
- SR-MPLS, 41**
  - addressing, 227
  - anycast SID, 254–255
    - configuration, 254–256*
    - verification, 256–257*
  - building a new network, 365
    - BGP proxy Prefix SID, 383–387*
    - enabling LDP on the border node, 372–376*
    - enabling SRMS, 365–372*
    - enabling the BGP Prefix SID, 376–383*
  - configuration, 228–229
    - assign the prefix SID, 230–231*
    - enable segment routing, 230*
    - reconfiguring the SRGB/SRLB, 229*
  - control plane, 207–209, 243–244
  - data forwarding, 341–342
  - data plane, 41–42, 207–209
  - enabling in an existing network (coexistence), 358–360
    - enabling and preferring on P1, P2, and PE-1, 363–365*
    - enabling on P2, P3, and PE-3, 360–363*
  - Flex Algo, 322–324
  - IS-IS, 244–246
    - configuration, 246*
    - verification, 247–249*
  - L3VPN overlay service, 677–679
    - BGP label allocation modes, 701*
    - configuration, 679–685*
    - extranet, configuration, 752–756*
    - extranet, verification, 756–767*
    - full-mesh, configuration, 689–700*
    - full-mesh, verification, 701–719*
    - hub-and-spoke, configuration, 719–732*
    - hub-and-spoke, verification, 732–751*
    - verification, 686–688*
  - label, collisions, 43
  - LSD (label switching database), 43
  - migration to segment routing, 358
  - OSPF, 250

- configuration*, 250
- verification*, 250–253
- Prefix SID, 45–47
- reachability, verification, 248–249, 253–254
- SID (segment identifier), 42–43
  - Adjacency*, 47–49
  - allocation*, 43–44
  - anycast*, 47
  - BGP Peering*, 50–52
  - BGP Prefix*, 49–50
  - Binding Segment*, 52–53
  - Node*, 45–46
  - sub-TLVs*, 57–61
- SRGB (segment routing global block), 44–45
- SRLB (segment routing local block), 45
- verify the SRGB and SRLB, 232–234
- verifying the Prefix SID assignment, 235–236
- SRMS (segment routing mapping server)**, 365–372
- SRMS Preference TLV**, 67
- SR-PM (Segment Routing Performance Measurement)**, 784
  - end-to-end delay measurement of any endpoint, 786
  - end-to-end SR policy delay measurement, 785–786
  - end-to-end SR policy liveness detection, 787
  - link delay measurement, 785
  - PT (Path Tracing), 787
- SRv6**, 29–30, 103, 213–214, 601, **ONLINE-ONLINE**
  - BGP link-state extensions, 201–205
  - building a new network
    - using an IWG*, 389–401
    - using dual-connected PE devices*, 413–424
    - using inter-AS Option A*, 401–413
  - control plane, 209–210, 257
  - data plane, 209–210
  - endpoint node, 126
  - Flex Algo, 301. *See also* Flex Algo
  - hardware and software support, 214
  - interface loopback address, 237–238
  - IS-IS, 257–260
    - anycast SID*, 270–271
    - anycast SID configuration*, 271
    - anycast SID use case*, 272–282
    - anycast SID verification*, 272
    - configuration*, 260–265
    - reachability verification*, 270
    - summarization*, 282–285
    - summarization configuration*, 286–287
    - summarization verification*, 287–292
  - UPA (Unreachability Prefix Announcement)*, 293–295
  - UPA configuration*, 295–298
  - UPA verification*, 298–301
  - verification*, 265
  - verify the SIDs*, 266–270
  - verifying IS-IS adjacency*, 265–266
  - verifying the database*, 266
- IS-IS extensions (RFC 9352), 175–176
  - segment routing over IPv6 capabilities*, 180–183
  - segment routing over IPv6 locator*, 176–180
  - SIDs*, 183–186
- IWG (interworking gateway), 389–390
- L3VPN overlay service, 608. *See also* full-mesh L3VPN
  - BGP Prefix SID path attribute*, 609–610
  - extranet, configuration*, 657–662
  - extranet, verification*, 662–677
  - full-mesh, configuration*, 610–625

- full-mesh, verification*, 625–636
- hub-and-spoke, configuration*, 636–645
- hub-and-spoke, verification*, 645–657
- MP\_REACH\_NLRI path attribute, 608
- NLRI, 608
- locator addressing scheme, 165–169
  - large-scale deployments*, 170–172
  - small- and medium-scale deployments*, 169–170
- MP-BGP extensions, BGP overlay services, 186–201
- open-source. *See* open-source SRv6
- policy, 126
  - End.B6.Encaps behavior*, 142–143
  - End.B6.Encaps.Red behavior*, 143–144
  - End.B6.Insert behavior*, 144–145
  - End.B6.Insert.Red behavior*, 145–146
  - endpoint behaviors*, 141, 146
  - headend behaviors*, 132, 140
  - H.Encaps behavior*, 132–133
  - H.Encaps.L2 behavior*, 135
  - H.Encaps.L2.Red behavior*, 136–137
  - H.Encaps.Red behavior*, 133–135
  - H.Insert behavior*, 138–140
- SID (segment identifier), 107–108
  - block addressing considerations*, 163
  - compression*, 161–162
  - End behavior*, 113–114
  - End.DT2M behavior*, 122–123
  - End.DT2U behavior*, 120–121
  - End.DT4 behavior*, 116–117
  - End.DX2 behavior*, 119
  - End.DX4 behavior*, 117–118
  - endpoint behaviors*, 1113
  - End.X behavior*, 114–115
  - Locator field*, 108–110
  - routed/nonrouted*, 110–111
- source node, 126
- SRH, 125–127
  - fields*, 124–125
  - Penultimate Segment Pop*, 127–128
  - Ultimate Segment Pop*, 129
  - USD (Ultimate Segment Decapsulation)*, 130–132
- standards, 103
- summarization, 172–175
- transit node, 126
- uSID (micro SID), 236, 239
  - allocation*, 236–237
  - globally significant*, 150–151
  - instruction extension*, 147–150
  - locally significant*, 150–151
  - locator addressing scheme*, 238–239
  - routed/nonrouted*, 151
  - uA endpoint variants*, 156–160
  - uN endpoint variants*, 152–155
  - verification*, 241–242
- uSID configuration, 239
  - enable SRv6 uSIDs*, 239–240
  - modify SRv6 parameters*, 240
- VPNv6 services, 194–195
- SRv6 BGP PeerNode SID TLV, 205
- SRv6 End SID sub-TLV, 179
- SRv6 End.X SID sub-TLV, 183–185
- SRv6 Locator TLV, 177–178
- SRv6 Service Data sub-sub-TLV, 188–189
- SRv6 SID Information sub-TLV, 188
- stakeholder testing, 1088
- standards
  - segment routing, 39, 40
  - SRv6, 103

**stitching RT, 390**

**sub-sub-TLV, 79–80**

Flex Algo, 79–83

SRv6 IS-IS, 185–186

SRv6 Service Data, 188–189

**sub-TLV**

Adj-SID, 59–61, 68–69

Flexible Algorithm Definition, 78–79

IS-IS Flexible Algorithm Definition,  
77–79

LAN-Adj-SID, 61

Prefix Attribute Flags, 178–179

Prefix-SID, 58–59, 70–71

Segment Routing Algorithm, 55–56

Segment Routing Capability, 54–55

SID/Label, 63–64

SRLB, 56–57

SRv6 End SID, 179

SRv6 End.X SID, 183–185

SRv6 Locator, 176–177

SRv6 SID Information, 188

**summarization, 220**

IS-IS, 172–175, 282–285

OSPFv3, 225

**SWAP operation, 14**

**synchronization, LDP–IGP, 24**

**systemctl restart frr command, ONLINE-  
ONLINE**

## T

**tcpdump, ONLINE-ONLINE**

**TDP (Tag Distribution Protocol), 1. See  
also LDP (Label Distribution Protocol)**

**TE (traffic engineering), 35–36, 38, 41,  
210, 213**

**teams, 1082–1083. See also domain**

**test/ing, 1067, 1080, 1086, 1088, 1097**  
acceptance, 1085, 1087

automation, 1094

development, 1091

integration, 1080, 1091

migration, 1096

stakeholder, 1088

unit, 1079, 1087

**TI-LFA (Topology-Independent Loop-Free  
Alternate), 883**

combined SRLG and node protection,  
934

link protection

*configuration, 883–885*

*SR-MPLS verification, 885–902*

*SRv6 verification, 902–914*

node protection

*configuration, 914–917*

*SR-MPLS verification, 917–919*

*SRv6 verification, 919–921*

*tiebreakers, 916*

SRLG protection

*configuration, 921–923*

*SR-MPLS verification, 923–931*

*SRv6 verification, 931–948*

**time/timer, 857–858**

EVPN E-LAN port-active MHD EVPN,  
491

failure event detection, 858

network event information propagation,  
858

topology update and repair path  
computation, 858

**TLV (type length value), 53. See also ISIS;  
OSPF; sub-sub-TLV; sub-TLV**

L2 Bundle Member Attributes, 93

Prefix Range, 72–73

Prefix-SID, 95

Router Capability, 54–57

Segment Routing Algorithm, 65

SID/Label, 91

- SID/Label Binding, 61–62
- SID/Label Range, 65–66
- SRLB, 66–67
- SRMS Preference, 67
- SRv6 BGP PeerNode SID, 205
- SRv6 Locator, 177–178
- topology**
  - definition, ONLINE
  - update and repair path computation time, 858
- traceroute command**, 249, 254, 312, 342, 362–363, 364–365, 375–376, 382, 383, 412, 423, 645, 677, 718–719, 750–751, 767
- traceroute ethernet cfm domain service command**, 816
- traffic blackholing**, 788
- Traffic Class field, IPv6 header**, 104
- traffic load balancing**, 798–799
- transit node**, 38, 126
- transport-related service assurance**
  - PT (Path Tracing), 784, 787, 798–801, 1023–1025
    - MCD (Midpoint Compressed Data)*, 800–801
    - probe packets*, 801–806
  - SR-DPM (Segment Routing Data Plane Monitoring), 788–789
    - adjacency verification and validation*, 790–792
    - configuration*, 789–790
    - prefix reachability verification*, 793–798
- TTL (Time to Live)**, 7
- tunnel, RSVP-TE**, 22–23
- TWAMP (Two-Way Active Measurement Protocol)**, 306–307, 835–836
- TWAMP Light**, 785
  - configuration, 851–853
  - verification, 853–854

## U

---

- ULA (unique local address)**, 162
- Ultimate Segment Pop**, 129
- underlay connectivity**, 1018
  - Linux SRv6, ONLINE-ONLINE
  - VPP (Vector Packet Processor), ONLINE-ONLINE
- Unified MPLS**, 210
- unified MPLS network migration**, 430–432
- unit test**, 1079, 1087
- unsolicited downstream**, 11
- UPA (Unreachability Prefix Announcement)**, 293–295
  - configuration, 295–298
  - verification, 298–301
- update packing**, 190–191
- USD (Ultimate Segment Decapsulation)**, 130–132
- uSID (micro SID)**, 36–37, 147–149, 236, 239
  - allocation, 236–237
  - configuration, 239
  - endpoint behaviors, 149–150
  - F3216 format, 148
  - globally significant, 150–151
  - locally significant, 150–151
  - locator addressing scheme, 238–239
  - SRv6 configuration
    - enable SRv6 uSIDs*, 239–240
    - modify SRv6 parameters*, 240
    - verification*, 241–242
  - uA endpoint variants, 156–160
  - uN endpoint variants, 152–155

## V

---

- VPLS (Virtual Private LAN Service)**, 440
- VPN**, 20

extranet, 657

intranet, 657

MPLS, operational complexity, 22

VRF (virtual routing and forwarding),  
17–18

**VPP (Vector Packet Processor), ONLINE-  
ONLINE**

data plane, ONLINE-ONLINE

DPDK, ONLINE-ONLINE

endpoint support, ONLINE-ONLINE

headend support, ONLINE-ONLINE

lab deployment, ONLINE-ONLINE

*basic setup, ONLINE-ONLINE*

*IPv4 L3VPN service, ONLINE-  
ONLINE*

*IPv6 L3VPN service, ONLINE-  
ONLINE*

*point-to-point L2VPN service,  
ONLINE-ONLINE*

*underlay connectivity, ONLINE-  
ONLINE*

linux-cp control plane integration plug-  
in, ONLINE-ONLINE

pps (packets per second), ONLINE-  
ONLINE

vector packet processing graph,  
ONLINE-ONLINE

**VPWS (Virtual Private Wire Service),  
440**

**VRF (virtual routing and forwarding),  
17–18, 605, 606**

**VRF instance, 3, 38, 5**

**vytish, ONLINE-ONLINE**

## **W-X-Y-Z**

---

wildcard, \*, 1112