



Segment Routing Configuration Guide for Cisco NCS 560 Series Routers, IOS XR Release 25.1.x, 25.2.x

First Published: 2025-01-31

Last Modified: 2025-02-05

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at 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: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. 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. (1721R)

© 2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	YANG Data Models for Segment Routing Features	1
	Using YANG Data Models	1
CHAPTER 2	About Segment Routing	3
	Scope	3
	Need	4
	Benefits	4
	Workflow for Deploying Segment Routing	5
CHAPTER 3	Configure Segment Routing over IPv6 (SRv6)	7
	Segment Routing over IPv6 Overview	8
	SRv6 Micro-Segment (uSID)	17
	SRv6 uSID Terminology	18
	SRv6 uSID Carrier Format	19
	SRv6 uSID Allocation Within a uSID Block	19
	SRv6 Endpoint Behaviors Associated with uSID	20
	SRv6 uSID in Action - Example	20
	Usage Guidelines and Limitations	25
	Configuring SRv6	27
	Configuring SRv6 under IS-IS	35
	Configuring SRv6 Flexible Algorithm under IS-IS	36
	Configuring SRv6 Locator Prefix Summarization	38
	Configuring TI-LFA with SRv6 IS-IS	38
	Configuring SRv6 IS-IS Microloop Avoidance	41
	Configuring SRv6 BGP-Based Services	42
	SRv6 Services: IPv4 L3VPN	44

SRv6 Services: IPv6 L3VPN	56
SRv6 Services: IPv4 BGP Global	68
SRv6 Services: IPv6 BGP Global	71
SRv6 Services: IPv4 L3VPN Active-Standby Redundancy using Port-Active Mode	77
SRv6 Services for L3VPN Active-Standby Redundancy using Port-Active Mode: Operation	78
Configure SRv6 Services L3VPN Active-Standby Redundancy using Port-Active Mode	78
Configuration Example	78
Running Configuration	79
Verification	79
SRv6 Services: IPv4 L3VPN Active-Active Redundancy	82
SRv6 Services: EVPN VPWS — All-Active Multi-Homing	83
SRv6-Services: EVPN ELAN Layer 2 Gateway With Automated Steering To Flexible Algorithm Paths	87
SRv6 ESI Filtering	92
SRv6 Services: L3 EVPN	94
SRv6 Services: IPv6 L3VPN	97
SRv6 Services: L2 and L3 Services with Remote SIDs from W-LIB	109
SRv6/MPLS L3 Service Interworking Gateway	112
L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway	117
SRv6/MPLS Dual-Connected PE	120
Static SRv6 pseudo-wire	122
Configure static SRv6 pseudo-wire	123
SRv6 SID Information in BGP-LS Reporting	126
DHCPv4 Relay Agent and Proxy Support over SRv6	126
DHCPv6 Relay Agent Support over SRv6	127
Full-Replace Migration to SRv6 Micro-SID	128
VRF-to-VRF route leaking in SRv6 core	132
Usage guidelines and limitations for VRF-to-VRF route leaking in SRv6	133
Configure VRF-to-VRF route leaking in SRv6	133

CHAPTER 4
Configure Segment Routing Global Block and Segment Routing Local Block 137

About the Segment Routing Global Block	137
About the Segment Routing Local Block	139
Understanding Segment Routing Label Allocation	140

Setup a Non-Default Segment Routing Global Block Range	143
Setup a Non-Default Segment Routing Local Block Range	144

CHAPTER 5

Configure Segment Routing for IS-IS Protocol	147
Enabling Segment Routing for IS-IS Protocol	147
Configuring a Prefix-SID on the IS-IS Enabled Loopback Interface	149
Weighted Anycast SID-Aware Path Computation	152
Configuring an Adjacency SID	157
Manually Configure a Layer 2 Adjacency SID	160
Configuring Bandwidth-Based Local UCMP	163
IS-IS Unreachable Prefix Announcement	164
Configuration Steps	165
IS-IS Partition Detection and Leakage of Specific Route Advertisement	167
Conditional Prefix Advertisement	171
Segment Routing ECMP-FEC Optimization	173
Verify the SR ECMP-FEC Optimization	176

CHAPTER 6

Configure Segment Routing for OSPF Protocol	179
Enabling Segment Routing for OSPF Protocol	179
Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface	181
Configuring an Adjacency SID	183
Protected Adjacency SID Backup Timer	185
Segment Routing ECMP-FEC Optimization	186
Multi-area Loopback Interface Support for OSPF	186
Usage Guidelines and Limitations for Multi-area Loopback	188
Configurations	188
Display Neighbor Router Capabilities in OSPF Networks	190
Usage Guidelines and Limitations to Display Neighbor Router Capabilities	192
Verify the Neighbor Router Capabilities	192

CHAPTER 7

Configure Segment Routing for BGP	195
Segment Routing for BGP	195
Configure BGP Prefix Segment Identifiers	196
Segment Routing Egress Peer Engineering	197

Usage Guidelines and Limitations	198
Configure Segment Routing Egress Peer Engineering	198
Configuring Manual BGP-EPE Peering SIDs	200
Configure BGP Link-State	202
Configurable Filters for IS-IS Advertisements to BGP-Link State	207
Configure Filters for IS-IS Advertisements to BGP-LS	207
Use Case: Configuring SR-EPE and BGP-LS	208
Configure BGP Proxy Prefix SID	211
BGP-LU Inter-AS Option-C Interworking with LDP and IGP SR-MPLS using Proxy BGP-SR	213
SRv6 double recursion for multilayer BGP underlay	218
Usage guidelines and limitations for SRv6 double recursion	219
Configure SRv6 double recursion for multilayer BGP underlay	220

CHAPTER 8

Configure SR-TE Policies	225
SR-TE Policy Overview	225
Usage Guidelines and Limitations	226
Instantiation of an SR Policy	226
On-Demand SR Policy – SR On-Demand Next-Hop	227
SR-ODN Configuration Steps	228
Configuring SR-ODN: Examples	231
Manually Provisioned SR Policy	260
PCE-Initiated SR Policy	260
Cumulative Metric Bounds (Delay-Bound Use-Case)	260
SR-TE BGP Soft Next-Hop Validation For ODN Policies	263
SR-TE Policy Path Types	265
Dynamic Paths	265
Optimization Objectives	266
Constraints	267
Configure SR Policy with Dynamic Path	270
Explicit Paths	272
SR-TE Policy with Explicit Path	272
Explicit Path with a BGP Prefix SID as First Segment	275
Configuring Explicit Path with Affinity Constraint Validation	279
Configure Explicit Path with Segment Protection-Type Constraint	281

Protocols	282
Path Computation Element Protocol	282
Configure the Head-End Router as PCEP PCC	282
Configure SR-TE PCE Groups	288
BGP SR-TE	293
Configure BGP SR Policy Address Family at SR-TE Head-End	293
Traffic Steering	295
Automated Steering	295
Color-Only Automated Steering	296
Setting CO Flag	297
Address-Family Agnostic Automated Steering	298
Per-Flow Automated Steering	298
Using Binding Segments	304
L2VPN Preferred Path	304
Static Route over Segment Routing Policy	305
Autoroute Include	307
Policy-Based Tunnel Selection for SR-TE Policy	308
SR-TE Automated Steering Without BGP Prefix Path Label	310
Use Case: Centralized BGP EPE for 6PE in an SR-MPLS Network	312
Miscellaneous	318
SR Policy Liveness Monitoring	318
Programming Non-Active Candidate Paths of an SR Policy	319
LDP over Segment Routing Policy	325
SR-TE MPLS Label Imposition Enhancement	328
Path invalidation drop	330
SR-TE Reoptimization Timers	332
Circuit-Style SR-TE Policies	334
SRv6 policy counters POL.CP.SL.INT.E	347
Reporting of SR-TE Policies Using BGP- Link State	348
Configure Reporting of SRTE Policies using BGP-LS	350
SR-TE Policy Path Protection	358
SR Policy Path Computation for IPv6	362

Bud Node Support	367
Configure Static Segment Routing Tree-SID via CLI at SR-PCE	368
Running Config	370
Multicast VPN: Dynamic Tree-SID MVPN (with TI-LFA)	372
Multicast VPN IPv6: Dynamic Tree-SID Multicast VPN IPv6	387
Multicast: Cisco Nonstop Forwarding for Tree-SID	395

CHAPTER 10

Configure Segment Routing Path Computation Element	397
About SR-PCE	397
Usage Guidelines and Limitations	399
Configure SR-PCE	399
Configure the Disjoint Policy (Optional)	401
Global Maximum-delay Constraint	402
PCE override of SR-TE policies	403
Usage guidelines	404
Limitations	404
Prerequisites	404
PCE override of PCC-initiated SR-TE policies	405
Configure PCE override of PCC-initiated SR-TE policies	405
PCE override of PCE-initiated SR-TE policies	409
Configure PCE override of PCE-initiated SR-TE policies	409
PCE-Initiated SR Policies	412
SR-PCE Flexible Algorithm Multi-Domain Path Computation	413
Example: SR-PCE Flexible Algorithm Multi-Domain Path Computation Use Case	414
Exclude Network Resources during Path Computation over SR-TE Policies	418
Configuration to exclude network resources during Path Computation	420
Configure the Shortest Path for Disjoint Candidate Paths	423
Configure the shortest path for disjoint candidate paths	423
Enable Strict Disjointness for SR-TE policies in the PCE	425
Configure Strict Disjointness in the PCE	426
ACL Support for PCEP Connection	427
Inter-Domain Path Computation Using Redistributed SID	427
Example: Inter-Domain Path Computation Using Redistributed SID	429
PCE Support for MPLS-TE LSPs	430

Configuring the North-Bound API on SR-PCE 433

CHAPTER 11

Configure Performance Measurement 437

Liveness Monitoring 438

IP Endpoint Liveness Monitoring 438

IP Endpoint Liveness Detection in an SR MPLS Network 442

IP Endpoint Liveness in an SRv6 Network 443

SR Policy Liveness Monitoring 446

Configure SR Policy Liveness Monitoring in an MPLS Network 449

Configure Segment Lists to Activate Candidate Paths in SRv6 for PM Liveness 454

Configure Flow Labels in SRv6 Header for PM Liveness 457

Hardware offload of SR policy liveness monitoring 460

Usage Guidelines and Limitations 461

Configuration Example 461

Delay Measurement 463

Measurement Modes 464

Link Delay Measurement 467

Delay Normalization 480

Link Anomaly Detection with IGP Penalty 483

Delay Measurement for IP Endpoint 484

IP Endpoint Delay Measurement over MPLS Network Usecases 486

IP Endpoint Delay Measurement over SRv6 Network Usecase 495

SR Policy End-to-End Delay Measurement 497

Fallback delay advertisement 504

Configure interface fallback delay value 505

Two-Way Active Measurement Protocol Light Source Address Filtering 507

Usage Guidelines and Limitations 508

Configure IP address on querier and responder nodes 508

Synthetic Loss Measurement 511

Configure Synthetic Loss measurement 512

Delay and synthetic loss measurement for GRE tunnel interfaces 517

Usage guidelines and limitations for delay and synthetic loss measurement for GRE tunnels 518

Configure delay and synthetic loss measurement for GRE tunnel interfaces 518

Delay Measurement Using Software Timestamp 520

Key Concepts	521
Benefits of Software Timestamp	522
How Delay is Measured Using Software Timestamp?	522
Guidelines and Limitations for Software Timestamp	522
Configure Software Timestamp for Delay Measurement	522

CHAPTER 12
Configure Topology-Independent Loop-Free Alternate (TI-LFA) 527

Limitations	530
Usage guidelines and limitations for TI-LFA	530
Configuring TI-LFA for IS-IS	531
Configuring TI-LFA for OSPF	533
TI-LFA Node and SRLG Protection: Examples	534
Configuring Global Weighted SRLG Protection	535
SR-MPLS over GRE as TI-LFA Backup Path	538
Limitations	541
Example: SR-MPLS over GRE as TI-LFA Backup Path	542
Unlabeled IPv6 Traffic Protection	549

CHAPTER 13
Configure Segment Routing Microloop Avoidance 553

About Segment Routing Microloop Avoidance	553
Usage Guidelines and Limitations	556
Configure Segment Routing Microloop Avoidance for IS-IS	556
Microloop Avoidance for IS-IS with Per-Prefix Filtering	557
Configure Segment Routing Microloop Avoidance for OSPF	561

CHAPTER 14
Configure Segment Routing Mapping Server 563

Segment Routing Mapping Server	563
Usage Guidelines and Restrictions	564
Segment Routing and LDP Interoperability	565
Example: Segment Routing LDP Interoperability	565
Configuring Mapping Server	568
Enable Mapping Advertisement	570
Configure Mapping Advertisement for IS-IS	570
Configure Mapping Advertisement for OSPF	571

Enable Mapping Client 572

CHAPTER 15

Enabling Segment Routing Flexible Algorithm 573

Prerequisites for Flexible Algorithm 573

Building Blocks of Segment Routing Flexible Algorithm 573

Flexible Algorithm Definition 573

Flexible Algorithm Membership 574

Flexible Algorithm Definition Advertisement 574

Flexible Algorithm Link Attribute Advertisement 574

Flexible Algorithm Prefix-SID Advertisement 575

Calculation of Flexible Algorithm Path 575

Installation of Forwarding Entries for Flexible Algorithm Paths 578

Flexible Algorithm Prefix-SID Redistribution 579

Flexible Algorithm Prefix Metric 581

Configuring Flexible Algorithm 582

Flexible Algorithm Link Attribute Advertisement Behavior 584

Strict IS-IS ASLA Link Attribute 587

Flexible Algorithm-Specific TE Metric 587

Flexible Algorithm with Exclude SRLG Constraint 588

Flexible Algorithm with Exclude Minimum Bandwidth Constraint 592

Flexible Algorithm with Exclude Maximum Delay Constraint 594

Maximum Paths Per IS-IS Flexible Algorithm 595

Maximum Paths Per IS-IS Flexible Algorithm Per Prefix 597

Example: Configuring IS-IS Flexible Algorithm 599

Example: Configuring OSPF Flexible Algorithm 599

User-Defined Generic Metric Support for IS-IS Flex Algo 600

Usage Guidelines and Limitations for User-Defined Generic Metrics 603

Configure User-Defined Generic Metrics 603

Example: Traffic Steering to Flexible Algorithm Paths 605

BGP Routes on PE – Color Based Steering 605

Delay Normalization 608

CHAPTER 16

Using Segment Routing OAM 613

MPLS Ping and Traceroute for BGP and IGP Prefix-SID 613

Examples: MPLS Ping, Traceroute, and Tree Trace for Prefix-SID	614
MPLS LSP Ping and Traceroute Nil FEC Target	616
Examples: LSP Ping and Traceroute for Nil_FEC Target	616
Segment Routing Ping and Traceroute	618
Segment Routing Ping	618
Segment Routing Traceroute	620
Segment Routing Ping and Traceroute for Flexible Algorithm	623
Segment Routing Ping for Flexible Algorithm	623
Segment Routing Traceroute for Flexible Algorithm	624
Segment Routing over IPv6 OAM	624
Data Plane Validation Support for SR-MPLS IPv6-based LSPs	626
Examples: SR-MPLS Data Plane Validation over IPv6-based LSPs	628
MPLS OAM support for SR-TE Policies using MPLS IPv6-based LSPs	629
Usage Guidelines and Limitations	630
Examples: MPLS OAM support for SR-TE Policies with IPv6-based LSPs	631



CHAPTER 1

YANG Data Models for Segment Routing Features

This chapter provides information about the YANG data models for Segment Routing features.

- [Using YANG Data Models, on page 1](#)

Using YANG Data Models

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the [Github](#) repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPath. To view a comprehensive list of the data models supported in a release, navigate to the **Available-Content.md** file in the repository.

You can also view the data model definitions using the [YANG Data Models Navigator](#) tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.



CHAPTER 2

About Segment Routing

- [Scope, on page 3](#)
- [Need, on page 4](#)
- [Benefits, on page 4](#)
- [Workflow for Deploying Segment Routing, on page 5](#)

Scope

Segment routing is a method of forwarding packets on the network based on the source routing paradigm. The source chooses a path and encodes it in the packet header as an ordered list of segments. Segments are an identifier for any type of instruction. For example, topology segments identify the next hop toward a destination. Each segment is identified by the segment ID (SID) consisting of a flat unsigned 20-bit integer.

Segments

Interior gateway protocol (IGP) distributes two types of segments: prefix segments and adjacency segments. Each router (node) and each link (adjacency) has an associated segment identifier (SID).

- A prefix SID is associated with an IP prefix. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels, and is distributed by IS-IS or OSPF. The prefix segment steers the traffic along the shortest path to its destination. A node SID is a special type of prefix SID that identifies a specific node. It is configured under the loopback interface with the loopback address of the node as the prefix.

A prefix segment is a global segment, so a prefix SID is globally unique within the segment routing domain.

- An adjacency segment is identified by a label called an adjacency SID, which represents a specific adjacency, such as egress interface, to a neighboring router. An adjacency SID can be allocated dynamically from the dynamic label range or configured manually from the segment routing local block (SRLB) range of labels. The adjacency SID is distributed by IS-IS or OSPF. The adjacency segment steers the traffic to a specific adjacency.

An adjacency segment is a local segment, so the adjacency SID is locally unique relative to a specific router.

By combining prefix (node) and adjacency segment IDs in an ordered list, any path within a network can be constructed. At each hop, the top segment is used to identify the next hop. Segments are stacked in order at the top of the packet header. When the top segment contains the identity of another node, the receiving node

uses equal cost multipaths (ECMP) to move the packet to the next hop. When the identity is that of the receiving node, the node pops the top segment and performs the task required by the next segment.

Dataplane

Segment routing can be directly applied to the Multiprotocol Label Switching (MPLS) architecture with no change in the forwarding plane. A segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. The related label is popped from the stack, after the completion of a segment.

Services

Segment Routing integrates with the rich multi-service capabilities of MPLS, including Layer 3 VPN (L3VPN), Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), and Ethernet VPN (EVPN).

Segment Routing for Traffic Engineering

Segment routing for traffic engineering (SR-TE) takes place through a policy between a source and destination pair. Segment routing for traffic engineering uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the provider core network to follow the specified path instead of the shortest path calculated by the IGP. The destination is unaware of the presence of the policy.

Need

With segment routing for traffic engineering (SR-TE), the network no longer needs to maintain a per-application and per-flow state. Instead, it simply obeys the forwarding instructions provided in the packet.

SR-TE utilizes network bandwidth more effectively than traditional MPLS-TE networks by using ECMP at every segment level. It uses a single intelligent source and relieves remaining routers from the task of calculating the required path through the network.

Benefits

- **Ready for SDN:** Segment routing was built for SDN and is the foundation for Application Engineered Routing (AER). SR prepares networks for business models, where applications can direct network behavior. SR provides the right balance between distributed intelligence and centralized optimization and programming.
- **Minimal configuration:** Segment routing for TE requires minimal configuration on the source router.
- **Load balancing:** Unlike in RSVP-TE, load balancing for segment routing can take place in the presence of equal cost multiple paths (ECMPs).
- **Supports Fast Reroute (FRR):** Fast reroute enables the activation of a pre-configured backup path within 50 milliseconds of path failure.
- **Plug-and-Play deployment:** Segment routing policies are interoperable with existing MPLS control and data planes and can be implemented in an existing deployment.

Workflow for Deploying Segment Routing

Follow this workflow to deploy segment routing.

1. Configure the Segment Routing Global Block (SRGB)
2. Enable Segment Routing and Node SID for the IGP
3. Configure Segment Routing for BGP
4. Configure the SR-TE Policy
5. Configure the SR-PCE
6. Configure TI-LFA and Microloop Avoidance
7. Configure the Segment Routing Mapping Server



CHAPTER 3

Configure Segment Routing over IPv6 (SRv6)

Segment Routing for IPv6 (SRv6) is the implementation of Segment Routing over the IPv6 dataplane.

- [Segment Routing over IPv6 Overview, on page 8](#)
- [SRv6 Micro-Segment \(uSID\), on page 17](#)
- [Usage Guidelines and Limitations, on page 25](#)
- [Configuring SRv6, on page 27](#)
- [Configuring SRv6 under IS-IS, on page 35](#)
- [Configuring SRv6 Flexible Algorithm under IS-IS, on page 36](#)
- [Configuring SRv6 Locator Prefix Summarization, on page 38](#)
- [Configuring TI-LFA with SRv6 IS-IS, on page 38](#)
- [Configuring SRv6 IS-IS Microloop Avoidance, on page 41](#)
- [Configuring SRv6 BGP-Based Services, on page 42](#)
- [SRv6 Services: IPv4 L3VPN Active-Standby Redundancy using Port-Active Mode, on page 77](#)
- [SRv6 Services: IPv4 L3VPN Active-Active Redundancy , on page 82](#)
- [SRv6 Services: EVPN VPWS — All-Active Multi-Homing , on page 83](#)
- [SRv6-Services: EVPN ELAN Layer 2 Gateway With Automated Steering To Flexible Algorithm Paths , on page 87](#)
- [SRv6 Services: L3 EVPN, on page 94](#)
- [SRv6 Services: IPv6 L3VPN, on page 97](#)
- [SRv6 Services: L2 and L3 Services with Remote SIDs from W-LIB, on page 109](#)
- [SRv6/MPLS L3 Service Interworking Gateway, on page 112](#)
- [L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway, on page 117](#)
- [SRv6/MPLS Dual-Connected PE, on page 120](#)
- [Static SRv6 pseudo-wire, on page 122](#)
- [SRv6 SID Information in BGP-LS Reporting, on page 126](#)
- [DHCPv4 Relay Agent and Proxy Support over SRv6, on page 126](#)
- [DHCPv6 Relay Agent Support over SRv6, on page 127](#)
- [Full-Replace Migration to SRv6 Micro-SID, on page 128](#)
- [VRF-to-VRF route leaking in SRv6 core, on page 132](#)

Segment Routing over IPv6 Overview

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Network Instructions	Release 7.5.1	This feature is now supported on the routers and operate in the native mode.


Segment Routing (SR) can be applied on both MPLS and IPv6 data planes. Segment Routing over IPv6 (SRv6) extends Segment Routing support with IPv6 data plane.

In an SR-MPLS enabled network, an MPLS label represents an instruction. The source nodes programs the path to a destination in the packet header as a stack of labels.

SRv6 introduces the Network Programming framework that enables a network operator or an application to specify a packet processing program by encoding a sequence of instructions in the IPv6 packet header. Each instruction is implemented on one or several nodes in the network and identified by an SRv6 Segment Identifier (SID) in the packet. The SRv6 Network Programming framework is defined in [IETF RFC 8986 SRv6 Network Programming](#).

In SRv6, an IPv6 address represents an instruction. SRv6 uses a new type of IPv6 Routing Extension Header, called the Segment Routing Header (SRH), in order to encode an ordered list of instructions. The active segment is indicated by the destination address of the packet, and the next segment is indicated by a pointer in the SRH.

Figure 1: Network Program in the Packet Header



The SRv6 SRH is documented in IETF RFC [IPv6 Segment Routing Header \(SRH\)](#).

The SRH is defined as follows:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Next Header  | Hdr Ext Len | Routing Type | Segments Left |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Last Entry   | Flags      | Tag          |                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

```

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|           Segment List[0] (128-bit IPv6 address)
|
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|           ...
|
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|           Segment List[n] (128-bit IPv6 address)
|
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//
//           Optional Type Length Value objects (variable)
//
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The following list explains the fields in SRH:

- Next header—Identifies the type of header immediately following the SRH.
- Hdr Ext Len (header extension length)—The length of the SRH in 8-octet units, not including the first 8 octets.
- Segments left—Specifies the number of route segments remaining. That means, the number of explicitly listed intermediate nodes still to be visited before reaching the final destination.
- Last Entry—Contains the index (zero based) of the last element of the segment list.
- Flags—Contains 8 bits of flags.
- Tag—Tag a packet as part of a class or group of packets like packets sharing the same set of properties.
- Segment list—128-bit IPv6 addresses representing the *n*th segment in the segment list. The segment list encoding starts from the last segment of the SR policy (path). That means the first element of the segment list (Segment list [0]) contains the last segment of the SR policy, the second element contains the penultimate segment of the SR policy and so on.

In SRv6, a SID represents a 128-bit value, consisting of the following three parts:

- Locator: This is the first part of the SID with most significant bits and represents an address of a specific SRv6 node.
- Function: This is the portion of the SID that is local to the owner node and designates a specific SRv6 function (network instruction) that is executed locally on a particular node, specified by the locator bits.
- Args: This field is optional and represents optional arguments to the function.

The locator part can be further divided into two parts:

- SID Block: This field is the SRv6 network designator and is a fixed or known address space for an SRv6 domain. This is the most significant bit (MSB) portion of a locator subnet.
- Node Id: This field is the node designator in an SRv6 network and is the least significant bit (LSB) portion of a locator subnet.

SRv6 Node Roles

Each node along the SRv6 packet path has a different functionality:

- Source node—A node that can generate an IPv6 packet with an SRH (an SRv6 packet), or an ingress node that can impose an SRH on an IPv6 packet.
- Transit node—A node along the path of the SRv6 packet (IPv6 packet and SRH). The transit node does not inspect the SRH. The destination address of the IPv6 packet does not correspond to the transit node.
- Endpoint node—A node in the SRv6 domain where the SRv6 segment is terminated. The destination address of the IPv6 packet with an SRH corresponds to the end point node. The segment endpoint node executes the function bound to the SID

SRv6 Head-End Behaviors

The SR Headend with Encapsulation behaviors are documented in the [IETF RFC 8986 SRv6 Network Programming](#).

The SR Headend with Insertion head-end behaviors are documented in the following IETF draft:

<https://datatracker.ietf.org/doc/draft-filsfils-spring-srv6-net-pgm-insertion/>

This section describes a set of SR Policy headend behaviors. The following list summarizes them:

- H.Encaps—SR Headend Behavior with Encapsulation in an SRv6 Policy
- H.Encaps.Red—H.Encaps with Reduced Encapsulation
- H.Insert—SR Headend with insertion of an SRv6 Policy
- H.Insert.Red—H.Insert with reduced insertion

SRv6 Endpoint Behaviors

The SRv6 endpoint behaviors are documented in the [IETF RFC 8986 SRv6 Network Programming](#).

The following is a subset of defined SRv6 endpoint behaviors that can be associated with a SID.

- End—Endpoint function. The SRv6 instantiation of a Prefix SID [[RFC8402](#)].
- End.X—Endpoint with Layer-3 cross-connect. The SRv6 instantiation of an Adj SID [[RFC8402](#)].
- End.DX6—Endpoint with decapsulation and IPv6 cross-connect (IPv6-L3VPN - equivalent to per-CE VPN label).
- End.DX4—Endpoint with decapsulation and IPv4 cross-connect (IPv4-L3VPN - equivalent to per-CE VPN label).
- End.DT6—Endpoint with decapsulation and IPv6 table lookup (IPv6-L3VPN - equivalent to per-VRF VPN label).
- End.DT4—Endpoint with decapsulation and IPv4 table lookup (IPv4-L3VPN - equivalent to per-VRF VPN label).
- End.DT46—Endpoint with decapsulation and specific IP table lookup (IP-L3VPN - equivalent to per-VRF VPN label).
- End.DX2—Endpoint with decapsulation and L2 cross-connect (L2VPN use-case).

- End.B6.Encaps—Endpoint bound to an SRv6 policy with encapsulation. SRv6 instantiation of a Binding SID.
- End.B6.Encaps.RED—End.B6.Encaps with reduced SRH. SRv6 instantiation of a Binding SID.

SRv6 Endpoint Behavior Variants

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
SRv6: Ultimate Segment Decapsulation (USD) on Full-length SIDs	Release 7.5.2	The Ultimate Segment Decapsulation (USD) variant is supported on SRv6 endpoint nodes using full-length SIDs. One of the USD variant applications is the case of TI-LFA in P routers with encapsulation with H.Encaps. The USD variant allows the last Segment Endpoint Node in the repair path list to decapsulate the IPv6 header added at the TI-LFA Point of Local Repair and forward the inner packet. In earlier releases, the USD variant was supported on SRv6 endpoint nodes using Micro SIDs (uSIDs).

Depending on how the SRH is handled, different behavior variants are defined for the End and End.X behaviors. The End and End.X behaviors can support these variants, either individually or in combinations.

- **Penultimate Segment Pop (PSP) of the SRH variant**—An SR Segment Endpoint Nodes receive the IPv6 packet with the Destination Address field of the IPv6 Header equal to its SID address.

A penultimate SR Segment Endpoint Node is one that, as part of the SID processing, copies the last SID from the SRH into the IPv6 Destination Address and decrements the Segments Left value from one to zero.

The PSP operation takes place only at a penultimate SR Segment Endpoint Node and does not happen at non-penultimate endpoint nodes. When a SID of PSP-flavor is processed at a non-penultimate SR Segment Endpoint Node, the PSP behavior is not performed since Segments Left would not be zero.

The SR Segment Endpoint Nodes advertise the SIDs instantiated on them via control plane protocols. A PSP-flavored SID is used by the Source SR Node when it needs to instruct the penultimate SR Segment Endpoint Node listed in the SRH to remove the SRH from the IPv6 header.

- **Ultimate Segment Pop (USP) of the SRH variant**—The SRH processing of the End and End.X behaviors are modified as follows:

If Segments Left is 0, then:

1. Update the Next Header field in the preceding header to the Next Header value of the SRH
2. Decrease the IPv6 header Payload Length by 8*(Hdr Ext Len+1)

3. Remove the SRH from the IPv6 extension header chain
4. Proceed to process the next header in the packet

One of the applications of the USP flavor is when a packet with an SRH is destined to an application on hosts with smartNICs implementing SRv6. The USP flavor is used to remove the consumed SRH from the extension header chain before sending the packet to the host.

- **Ultimate Segment Decapsulation (USD) variant**—The Upper-layer header processing of the End and End.X behaviors are modified as follows:

- **End** behavior: If the Upper-layer Header type is 41 (IPv6), then:
 1. Remove the outer IPv6 Header with all its extension headers
 2. Submit the packet to the egress IPv6 FIB lookup and transmission to the new destination
 3. Else, if the Upper-layer Header type is 4 (IPv4)
 4. Remove the outer IPv6 Header with all its extension headers
 5. Submit the packet to the egress IPv4 FIB lookup and transmission to the new destination
 6. Else, process as per Section 4.1.1 (Upper-Layer Header) of [IETF RFC 8986 SRv6 Network Programming](#)
- **End.X** behavior: If the Upper-layer Header type is 41 (IPv6) or 4 (IPv4), then:
 1. Remove the outer IPv6 Header with all its extension headers
 2. Forward the exposed IP packet to the L3 adjacency J
 3. Else, process as per Section 4.1.1 (Upper-Layer Header) of [IETF RFC 8986 SRv6 Network Programming](#)

One of the applications of the USD flavor is the case of TI-LFA in P routers with encapsulation with H.Encaps. The USD flavor allows the last Segment Endpoint Node in the repair path list to decapsulate the IPv6 header added at the TI-LFA Point of Local Repair and forward the inner packet.

Usage Guidelines and Limitations

General Guidelines and Limitations

- Cisco IOS XR Release 7.5.2 and later supports the following SRv6 SID behaviors and variants:
 - END with PSP/USD
 - END.X with PSP/USD
 - END.DT4
 - END.DT6
- SRv6 Underlay support includes:
 - IGP redistribution/leaking between levels
 - Prefix Summarization on ABR routers

- IS-IS TI-LFA
- Microloop Avoidance
- Flex-algo

Configuring SRv6

To enable SRv6 globally, you should first configure a locator with its prefix. The IS-IS protocol announces the locator prefix in IPv6 network and SRv6 applications (like ISIS, BGP) use it to allocate SIDs.

The following usage guidelines and restrictions apply while configuring SRv6.

- All routers in the SRv6 domain should have the same SID block (network designator) in their locator.
- The locator length should be 64-bits long.
 - The SID block portion (MSBs) cannot exceed 40 bits. If this value is less than 40 bits, user should use a pattern of zeros as a filler.
 - The Node Id portion (LSBs) cannot exceed 24 bits.
- You can configure up to 8 locators to support SRv6 Flexible Algorithm. All locators prefix must share the same SID block (first 40-bits).

Enabling SRv6 on the Platform

Before configuring SRv6 on , you must first use the following command in config mode:

- **hw-module profile segment-routing srv6 mode base**

You must reload the router after enabling this command.



Note Starting from Release 7.10.1, the SRv6 mode is automatically set to **base-and-micro-segment-f3216 mode** (dual mode) even if you configure the base mode. The running configuration would still continue to reflect the user configured mode only.

You can verify the change using the following console log:

```
fia_driver[238]: %FABRIC-FIA_DRV-6-HW_MOD_PROFILE_AUTO_CONVERTED
:
Auto-converting SRv6 hw-module base profile to
base-and-micro-segment-f3216 profile
```

Enabling SRv6 with Locator

This example shows how to globally enable SRv6 and configure locator.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLoc1
Router(config-srv6-locator)# prefix 2001:db8:0:a2::/64
```

(Optional) Configuring SRv6 Anycast Locator

An SRv6 Anycast locator is a type of locator that identifies a set of nodes (END SIDs). SRv6 Anycast Locators and their associated END SIDs may be provisioned at multiple places in a topology.

The set of nodes (Anycast group) is configured to advertise a shared Anycast locator and END SID. Anycast routing enables the steering of traffic toward multiple advertising nodes. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.

One use case is to advertise Anycast END SIDs at exit points from an SRv6 network. Any of the nodes that advertise the common END SID could be used to forward traffic out of the SRv6 portion of the network to the topologically nearest node.

Unlike a normal locator, IS-IS does not program or advertise END.X SIDs associated with an anycast locator.



Note END SIDs allocated from Anycast locators will not be used in constructing TI-LFA backup paths or Microloop Avoidance primary paths. TI-LFA backup and Microloop Avoidance paths for an Anycast locator prefix may terminate on any node advertising that locator, which may be different from the node terminating the original primary path.



Note SRv6 anycast locators may have non-zero algorithm (Flexible Algorithm) values.

The following example shows how to globally enable SRv6 and configure Anycast locator.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLoc1 anycast
Router(config-srv6-locator)# prefix 2001:db8:0:a2::/64
```

Optional: Configuring Encapsulation Parameters

This example shows how to configure encapsulation parameters when configuring SRv6. These optional parameters include:

- **segment-routing srv6 encapsulation source-address *ipv6-addr***—Source Address of outer encapsulating IPv6 header. The default source address for encapsulation is one of the loopback addresses.
- **segment-routing srv6 encapsulation hop-limit {*count* | **propagate**}**—The hop limit of outer-encapsulating IPv6 header. The range for *count* is from 1 to 254; the default value for hop-limit is 254. Use **propagate** to set the hop-limit value by propagation (from incoming packet/frame).

```
Router(config)# segment-routing srv6
Router(config-srv6)# encapsulation source-address 1::1
Router(config-srv6)# hop-limit 60
```

Optional: Enabling Syslog Logging for Locator Status Changes

This example shows how to enable the logging of locator status.

```
Router(config)# segment-routing srv6
Router(config-srv6)# logging locator status
```

Verifying SRv6 Manager

This example shows how to verify the overall SRv6 state from SRv6 Manager point of view. The output displays parameters in use, summary information, and platform specific capabilities.

```
Router# show segment-routing srv6 manager
Parameters:
  Parameters:
    SRv6 Enabled: Yes
    SRv6 Operational Mode:
      Base:
        SID Base Block: 2001:db8::/40
  Encapsulation:
    Source Address:
      Configured: 1::1
      Default: 5::5
    Hop-Limit: Default
    Traffic-class: Default
Summary:
  Number of Locators: 1 (1 operational)
  Number of SIDs: 4 (0 stale)
  Max SIDs: 64000
  OOR:
    Thresholds: Green 3200, Warning 1920
    Status: Resource Available
      History: (0 cleared, 0 warnings, 0 full)
    Block 2001:db8:0:a2::/64:
      Number of SIDs free: 65470
      Max SIDs: 65470
      Thresholds: Green 3274, Warning 1965
      Status: Resource Available
        History: (0 cleared, 0 warnings, 0 full)
Platform Capabilities:
  SRv6: Yes
  TILFA: Yes
  Microloop-Avoidance: Yes
  Endpoint behaviors:
    End (PSP)
    End.X (PSP)
    End.DX6
    End.DX4
    End.DT6
    End.DT4
    End.DX2
    uN (PSP/USD)
    uA (PSP/USD)
    uDT6
    uDT4
    uDX2
    uB6 (Insert.Red)
  Headend behaviors:
    T
    H.Insert.Red
    H.Encaps.Red
  Security rules:
    SEC-1
    SEC-2
    SEC-3
  Counters:
    CNT-1
    CNT-3
  Signaled parameters:
    Max-SL : 3
```



```

Max-End-Pop-SRH : 3
Max-H-Insert    : 3 sids
Max-H-Encap     : 3 sids
Max-End-D       : 4
Configurable parameters (under srv6):
Encapsulation:
  Source Address: Yes
  Hop-Limit     : value=Yes, propagate=No
  Traffic-class  : value=Yes, propagate=Yes
Max SIDs: 64000
SID Holdtime: 3 mins

```

Verifying SRv6 Locator

This example shows how to verify the locator configuration and its operational status.

```

Router# show segment-routing srv6 locator myLoc1 detail

```

Name	ID	Prefix	Status
myLoc1*	5	2001:db8:0:a2::/64	Up

```

(*) : is-default
Interface:
  Name: srv6-myLoc1
  IFH : 0x00000170
  IPv6 address: 2001:db8:0:a2::/64
  Chkpt Obj ID: 0x2fc8
  Created: Apr 25 06:21:57.077 (00:03:37 ago)

```

Verifying SRv6 Local SIDs

This example shows how to verify the allocation of SRv6 local SIDs off locator(s).

```

Router# show segment-routing srv6 locator myLoc1 sid

```

SID	State	RW	Function	Context	Owner
2001:db8:0:a2:1::	InUse	Y	End (PSP)	'default':1	sidmgr
2001:db8:0:a2:40::	InUse	Y	End.DT4	'VRF1'	bgp-100
2001:db8:0:a2:41::	InUse	Y	End.X (PSP)	[Hu0/1/0/1, Link-Local]	isis-srv6

The following example shows how to display detail information regarding an allocated SRv6 local SID.

```

Router# show segment-routing srv6 locator myLoc1 sid 2001:db8:0:a2:40:: detail

```

SID	State	RW	Function	Context	Owner
2001:db8:0:a2:40::	InUse	Y	End.DT4	'VRF1'	bgp-100

```

  SID context: { table-id=0xe0000011 ('VRF1':IPv4/Unicast) }
  Locator: myLoc1'
  Allocation type: Dynamic

```

Created: Feb 1 14:04:02.901 (3d00h ago)

Similarly, you can display SID information across locators by using the **show segment-routing sid** command.

show Commands

You can use the following **show** commands to verify the SRv6 global and locator configuration:

Command	Description
show segment-routing srv6 manager	Displays the summary information from SRv6 manager, including platform capabilities.
show segment-routing srv6 locator <i>locator-name</i> [detail]	Displays the SRv6 locator information on the router.
show segment-routing srv6 locator <i>locator-name</i> sid [<i>sid-ipv6-address</i> [detail]]	Displays the information regarding SRv6 local SID(s) allocated from a given locator.
show segment-routing srv6 sid [<i>sid-ipv6-address</i> all stale] [detail]	Displays SID information across locators. By default, only “active” (i.e. non-stale) SIDs are displayed.
show route ipv6 local-srv6	Displays all SRv6 local-SID prefixes in IPv6 RIB.

SRv6 Micro-Segment (uSID)

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Micro-Segment (uSID)	Release 7.3.1	This feature is an extension of the SRv6 architecture. It leverages the existing SRv6 Network Programming architecture to encode up to six SRv6 Micro-SID (uSID) instructions within a single 128-bit SID address. Such a SID address is called a uSID Carrier. In addition, this feature leverages the existing SRv6 data plane and control plane with no changes. It also provides low MTU overhead; for example, 6 uSIDs per uSID carrier results in 18 source-routing waypoints in only 40 bytes of overhead (in SRH).

The SRv6 micro-segment (uSID) is an extension of the SRv6 architecture. It leverages the SRv6 Network Programming architecture to encode several SRv6 Micro-SID (uSID) instructions within a single 128-bit SID address. Such a SID address is called a uSID Carrier.

SRv6 uSID is documented in the IETF drafts [Network Programming extension: SRv6 uSID instruction](#) and [Compressed SRv6 Segment List Encoding in SRH](#).

Throughout this chapter, we will refer to SRv6 micro-segment as “uSID”.

The SRv6 uSID provides the following benefits:

- Leverages the SRv6 Network Programming with no change. SRv6 uSID is a new pseudo code in the existing SRv6 network programming framework.
- Leverages the SRv6 data plane (SRH) with no change. Any SID in the destination address or SRH can be an SRv6 uSID carrier.
- Leverages the SRv6 control plane with no change.
- Ultra-Scale—Scalable number of globally unique nodes in the domain, for example:
 - 16-bit uSID ID size: 65k uSIDs per domain block
 - 32-bit uSID ID size: 4.3M uSIDs per domain block
- Lowest MTU overhead
 - 6 uSIDs per uSID carrier
 - For example, 18 source-routing waypoints in only 40 bytes of overhead
- Hardware-friendliness:
 - Leverages mature hardware capabilities (inline IP Destination Address edit, IP Destination Address longest match).
 - Avoids any extra lookup in indexed mapping tables.
 - A micro-program with 6 or fewer uSIDs requires only legacy IP-in-IP encapsulation behavior.
- Scalable Control Plane:
 - Summarization at area/domain boundary provides massive scaling advantage.
 - No routing extension is required, a simple prefix advertisement suffices.
- Seamless Deployment:
 - A uSID may be used as a SID (the carrier holds a single uSID).
 - The inner structure of an SR Policy can stay opaque to the source. A carrier with uSIDs is just seen as a SID by the policy headend Security.
 - Leverages SRv6's native SR domain security.

SRv6 uSID Terminology

The SRv6 Network Programming is extended with the following terms:

- **uSID**—An identifier that specifies a micro-segment.
A uSID has an associated behavior that is the SRv6 function (for example, a node SID or Adjacency SID) associated with the given ID. The node at which an uSID is instantiated is called the “Parent” node.
- **uSID Carrier**—A 128-bit IPv6 address (carried in either in the packet destination address or in the SRH) in the following format:

```
<uSID-Block><Active-uSID><Next-uSID>...<Last-uSID><End-of-Carrier>...<End-of-Carrier>
```

where:

- **uSID Block**—An IPv6 prefix that defines a block of SRv6 uSIDs.
- **Active uSID**—The first uSID that follows the uSID block.
- **Next uSID**—The next uSID after the Active uSID.
- **Last uSID**—The last uSID in the carrier before the End-of-Carrier uSID.
- **End-of-Carrier**—A globally reserved uSID that marks the end of a uSID carrier. The End-of-Carrier ID is **0000**. All empty uSID carrier positions must be filled with the End-of-Carrier ID; therefore, a uSID carrier can have more than one End-of-Carrier.

The following is an example of an SRH with 3 Micro-SID carriers for a total of up to 18 micro-instructions:

Micro-SID Carrier1: {uInstruction1, uInstruction2... uInstruction6}
Micro-SID Carrier2: {uInstruction7, uInstruction8... uInstruction12}
Micro-SID Carrier3: {uInstruction13, uInstruction14... uInstruction18}

SRv6 uSID Carrier Format

The uSID carrier format specifies the type of uSID carrier supported in an SRv6 network. The format specification includes Block size and ID size.

• uSID Block

The uSID block is an IPv6 prefix that defines a block of SRv6 uSIDs. This can be an IPv6 prefix allocated to the provider (for example, /22, /24, and so on.), or it can be any well-known IPv6 address block generally available for private use, such as the ULA space FC/8, as defined in IETF draft [RFC4193](#).

An SRv6 network may support more than a single uSID block.

The length of block [prefix] is defined in bits. From a hardware-friendliness perspective, it is expected to use sizes on byte boundaries (16, 24, 32, and so on).

• uSID ID

The length of uSID ID is defined in bits. From a hardware-friendliness perspective, it is expected to use sizes on byte boundaries (8, 16, 24, 32, and so on).

The uSID carrier format is specified using the notation "Fbbuu", where "bb" is size of block and "uu" is size of ID. For example, "F3216" is a format with a 32-bit uSID block and 16-bit uSID IDs.



Note F3216 is the default format, and the only format that is supported in IOS XR 7.3.1 release.

SRv6 uSID Allocation Within a uSID Block

The architecture for uSID specifies both globally scoped and locally scoped uSIDs, where a globally scoped uSID is the type of uSID that provides reachability to the node.

On the other hand, a locally scoped uSID is associated to a local behavior, and therefore *must* be preceded by a globally scoped uSID of the parent node when relying on routing to forward the packet.

The Global ID block (GIB) is the set of IDs available for globally scoped uSID allocation. The Local ID block (LIB) is the set of IDs available for locally scoped uSID allocation.

A globally scoped uSID is a uSID from the GIB. A globally scoped uSID typically identifies a shortest path to a node in the SR domain. An IP route (for example, /48) is advertised by the parent node to each of its globally scoped uSIDs, under the associated uSID block. The parent node executes a variant of the END behavior.

The “Nodal” uSID (uN) is an example of a globally scoped behavior defined in uSID architecture.

A node can have multiple globally scoped uSIDs under the same uSID blocks (for example, one per IGP flex-algorithm). Multiple nodes may share the same globally scoped uSID (Anycast).

A locally scoped uSID is a uSID from the LIB. A locally scoped uSID identifies a local micro-instruction on the parent node; for example, it may identify a cross-connect to a direct neighbor over a specific interface or a VPN context. Locally scoped uSIDs are not routeable.

For example, if N1 and N2 are two different physical nodes of the uSID domain and *L* is a locally scoped uSID value, then N1 and N2 may bind two different behaviors to *L*.

The uSIDs are allocated in one of following ways: auto, dynamic, or explicit.

- The request to allocate locally scoped uSIDs comes from SRv6 clients (such as IS-IS or BGP). The request can be to allocate any available ID (dynamic allocation) or to allocate a specific ID (explicit allocation).

SRv6 Endpoint Behaviors Associated with uSID

The SRv6 Network Programming is extended with new types of SRv6 SID endpoint behaviors:

- uN—A short notation for the NEXT-CSID (Compressed SID) End behavior with a pseudocode of shift-and-lookup, and PSP/USD flavors
- uA—A short notation for the NEXT-CSID End.X behavior with a pseudocode of shift-and-xconnect, and PSP/USD flavors
- uDT—A short notation for the NEXT-CSID End.DT behavior with the same pseudocode as End.DT4/End.DT6/End.DT46/End.DT2U/End.DT2M
- uDX—A short notation for the NEXT-CSID End.DX behavior with the same pseudocode as End.DX4/End.DX6/End.DX2

SRv6 uSID in Action - Example

This example highlights an integrated VPN and Traffic Engineering use-case leveraging SRv6 uSID.

VPNv4 site A connected to Node 1 sends packets to VPNv4 site B connected to Node 2 alongside a traffic engineered path via Node 8 and Node 7 using a single 128-bit SRv6 SID.

Node 1 is the ingress PE; Node 2 is the egress PE.

Nodes 3, 4, 5, and 6 are classic IPv6 nodes. Traffic received on these nodes use classic IP forwarding without changing the outer DA.

Nodes 1, 8, 7 and 2 are SRv6 capable configured with:

- 32-bit SRv6 block = fcbb:bb01
- 16-bit SRv6 ID


For example:

- Node 7 uN = fcbb:bb01:0700::/48
- Node 8 uN = fcbb:bb01:0800::/48

The following IGP routes are advertised:

- Node 8 advertises the IGP route fcbb:bb01:**0800**::/48
- Node 7 advertises the IGP route fcbb:bb01:**0700**::/48
- Node 2 advertises the IGP route fcbb:bb01:**0200**::/48

Figure 2: Integrated VPN and Traffic Engineering SRv6 uSID Use-case




- Node 1 encapsulates IPv4 packet from Site A and sends an IPv6 packet with DA = fcbb:bb01:0800:0700:0200:f001::
- Traffic engineered path via 8 and 7 using a single 128-bit SRv6 SID
- One single micro-program in the DA is enough

521410

Node 1 encapsulates an IPv4 packet from VPN Site A and sends an IPv6 packet with destination address fcbb:bb01:0800:0700:0200:f001::. This is a uSID carrier, with a list of micro-instructions (uSIDs) (0800, 0700, 0200, f001, and 0000 – indicating the end of the instruction).

uSIDs (uNs) 0800, 0700, 0200 are used to realize the traffic engineering path to Node 2 with way points at Nodes 8 and 7. uSID f001 is the BGP-signalled instruction (uDT4) advertized by Node 2 for the VPNv4 service

Figure 3: Node 1: End.B6.Encaps Behavior




DA = fcbb:bb01:0800:0700:0200:f001:0000:0000

521411

Nodes 4 and 5 simply forward the packet along the shortest path to Node 8, providing seamless deployment through classic IPv6 nodes.

Figure 4: Node 4 and Node 5: Classic IPv6 Nodes




DA = fcbb:bb01:0800:0700:0200:f001:0000:0000

521412

When Node 8 receives the packet, it performs SRv6 uN behavior (shift-and-lookup with PSP/USD). It removes its outer DA (0800) and advances the micro program to the next micro instruction by doing the following:

1. Pops its own uSID (0800)
2. **Shifts** the remaining DA by 16-bits to the left
3. Fills the remaining bits with 0000 (End-of-Carrier)
4. Performs a **lookup** for the shortest path to the next DA (fcbb:bb01:**0700::/48**)
5. Forwards it using the new DA fcbb:bb01:**0700**:0200:f001:0000:0000:0000


Figure 5: Node 8: SRv6 uN Behavior (Shift and Forward)



When Node 7 receives the packet, it performs the same SRv6 uN behavior (shift-and-lookup with PSP/USD), forwarding it using the new DA fcbb:bb01:**0200**:f001:0000:0000:0000:0000

521413

Figure 6: Node 7: SRv6 uN Behavior (Shift and Forward)




FIB Longest-Match fcbb:bb01:0700::/48 → SRv6 Instruction:

- Shift micro-program by one micro-instruction
- Set last micro-instruction to "End-of-Carrier"
- Lookup the updated DA and forward

Nodes 6 and 3 simply forward the packet along the shortest path to Node 2, providing seamless deployment through classic IPv6 nodes.


Figure 7: Node 6 and Node 3: Classic IPv6 Nodes



DA = fcbb:bb01:0200:f001:0000:0000:0000:0000

When Node 2 receives the packet, it performs an SRv6 uDT4 behavior (End.DT4—Endpoint with decapsulation and IPv4 table lookup) to VPNv4 Site B.

Figure 8: Node 2: SRv6 uDT4 Behavior



FIB Longest-Match fcbb:bb01:0200:f001::/64 → SRv6 Instruction:

- Decapsulate and Lookup of inner IPv4 packet

521416

To recap, this example showed an integrated VPN and Traffic Engineering use-case, where VPNv4 site A connected to Node 1 sent packets to VPNv4 site B connected to Node 2 alongside a traffic engineered path via Node 8 and Node 7 using a single 128-bit SRv6 SID:

- @1: inner packet P encapsulated with outer DA fcbb:bb01:**0800:0700:0200**:f001:0000:0000
- @4 & @5: classic IP forwarding, outer DA unchanged
- @8: SRv6 uN behavior: shift and lookup, outer DA becomes fcbb:bb01:**0700:0200**:f001:0000:0000:0000
- @7: SRv6 uN behavior: shift and lookup, outer DA becomes fcbb:bb01:**0200**:f001:0000:0000:0000:0000
- @6 & @3: classic IP forwarding, outer DA unchanged
- @2: SRv6 End.DT4: Decapsulate and IPv4 table lookup

Usage Guidelines and Limitations

General Guidelines and Limitations

- Cisco IOS XR supports uSIDs with 32-bit uSID block and 16-bit uSID IDs (3216).
A single UCF format must be used for uSID locators in a SRv6 uSID domain.
- Cisco IOS XR supports up to 8 uSID locator prefixes.
Multiple locator prefixes are used when configuring Anycast locators or SRv6 Flexible Algorithm instances, for example.
- Cisco IOS XR supports uSID locator prefixes from different uSID blocks.
Up to 64 uSID blocks can be used across all uSID locators in the network.
- Cisco IOS XR Release 7.3.1 and later supports the following SRv6 uSID behaviors and variants:

- uN with PSP/USD
- uA with PSP/USD
- uDT4
- uDT6
- SRv6 Underlay support includes:
 - IGP redistribution/leaking between levels
 - Prefix Summarization on ABR routers
 - IS-IS TI-LFA
 - Microloop Avoidance
 - Flex-algo
- SRv6 over GRE interface is not supported
- SRv6 over channelized port is not supported

uSID Allocation Recommendation

We recommend that the uSID block allocation is made from the IPv6 Unique Local Address (ULA) range.



Note Allocation from the public Global Unicast Addresses (GUA) range is also supported.

- Use ULA /24 base from FC00::/8 space
 - FCBB:BB/24, with *B* indicating a nibble value picked by operator
- 256 uSID blocks possible from this allocation
 - In this release, 64 uSID blocks are supported
 - FCBB:BBVV/32, with *VV* two variable nibbles. The supported values for *VV* in Cisco IOS XR Release 7.3.1 are 0x00 to 0x3F.

For example:

- ULA /24 base = FC00:01/24
- uSID block space = 64 uSID blocks (from FC00:01**00**/32 to FC00:01**3F**/32)

Configuring SRv6

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Traffic Class DSCP Marking Enhancement for QoS	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.
SRv6 Traffic Class DSCP Marking Enhancement for QoS	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS This feature allows you to set the traffic class values to use the full range of DSCP (6 bits) for SRv6 packets that are encapsulated. The feature extends the existing policy-map setting of 3-bit traffic class DSCP marking to 6 bits using the new policy-map-extend option in the configuration for L2 and L3 traffic for SRv6. This is a mutually inclusive feature for L2 and L3 traffic. Therefore, this feature cannot be used in the individual configuration for L2 and L3 traffic. This feature introduces this change: CLI: The policy-map-extend keyword is introduced in the hw-module profile segment-routing srv6 mode microsegment f3216 encapsulation traffic-class command.

Feature Name	Release Information	Feature Description
SRv6 Traffic Class QoS Enhancement	Release 7.7.1	The modified hw-module profile segment-routing srv6 mode command option provides you with better flexibility to customize the optional SRv6 encapsulation parameters. The updated command supports both L2 and L3 traffic types of SRv6 parameters. Encapsulation is a sub-mode from Release 7.7.1. The I3-traffic configuration supports the additional policy-map option that sets SRv6 traffic-class DSCP based on qos-group selected by input policy-map. The following commands are updated: • hw-module profile segment-routing srv6 mode: Mode is a mandatory parameter The following commands are introduced: • encapsulation l2-traffic • encapsulation l3-traffic

Enabling SRv6 involves the following high-level configuration steps:

- Enable SRv6 on the platform
- Configure SRv6 locator(s)
- Enable SRv6 under IS-IS
- Enable SRv6 Services under BGP

Enable SRv6 on the Platform

Before configuring SRv6 on the router, you must first use the following command:

- **hw-module profile segment-routing srv6 mode micro-segment format f3216**



Note You must reload the router after enabling this command.

(Optional) Configure Network Role

By default, after enabling SRv6 on the platform, the node can serve as both edge (services) and core roles.

Optionally, you can customize the node role as "core-only" using the following command:

- **hw-module profile network-role core-only**



Note You must reload the router after enabling this command.

Given that there is different budget for underlay SID encap based on the node role in the network (P-only vs Edge), an operator can use this configuration to provide a hint to the platform and control plane to use a larger SID encap budget when operating as a P-only node.

(Optional) Configure Merge Overlay/Underlay SID Mode

One of the main benefits of SRv6 uSID is compression (or packing) of multiple uSIDs into a uSID carrier. This is possible when they share the same uSID block and when there is enough space in the carrier.

The underlay SIDs are always programmed in compressed form, if possible. The overlay SID is programmed separately.

The **segment-routing srv6 micro-segment merge-overlay-underlay-sids** command is used to enable the platform to merge overlay/underlay SIDs.



Note From Cisco IOS XR Release 7.7.1 and later, the compression/merging of uSID lists is automatically done in the dataplane during imposition. This configuration option is ignored by the dataplane.

When there is a need to send overlay traffic, the data path implementation attempts to merge the underlay SIDs and overlay SIDs into a single carrier, if possible. With H.Encaps.Red encapsulation, this yields a packet with no SRH.



Note If the overlay and underlay use different uSID blocks, this merge is not possible.

By default, the Cisco NCS platform does not automatically merge the overlay/underlay SIDs.

To enable the platform to merge overlay/underlay SIDs, use the following command:

- **segment-routing srv6 micro-segment merge-overlay-underlay-sids**



Caution This command should only be enabled when a single block is required.

After you enable this command, this CLI will modify the behavior for all new overlay routes being programmed afterwards.

If you enable this command after SRv6 overlay routes are already programmed, we recommend that you clear the SRv6 overlay routes (using the **clear route [vrf WORD]** command) in order to trigger the re-programming in the “merge” mode.

If you do not to clear the overlay routes, those routes would continue to be programmed in the “non-merge” mode.

(Optional) Customize SRv6 Encapsulation Parameters

From Cisco IOS XR Release 7.7.1, the **hw-module profile segment-routing srv6 mode** command has been updated to support both L2 and L3 traffic types of SRv6 parameters. The existing traffic-class CLI (common

to L2 and L3) is still accepted in the ASCII format, but will be converted into the new **I2-traffic** and **I3-traffic** submode format after the config is committed.

From Cisco IOS XR Release 7.7.1, ensure that you use mandatory keyword **mode**. Note that **encapsulation** is now a sub-mode and not a configuration option.

The following example shows how to configure the traffic-class field settings on the IPv6 header:

```
Router(config)#hw-module profile segment-routing srv6 mode micro-segment format f3216
Router(config-srv6)#encapsulation
Router(config-srv6-encap)#I2-traffic
Router(config-srv6-encap-I2)#traffic-class propagate
```

```
Router(config)#hw-module profile segment-routing srv6 mode micro-segment format f3216
Router(config-srv6)#encapsulation
Router(config-srv6-encap)#I3-traffic
Router(config-srv6-encap-I3)#traffic-class policy-map
```

The **I2-traffic** and **I3-traffic** are SRv6 parameters for L2 traffic type and L3 traffic type respectively. The **traffic-class** parameter controls the traffic-class field of SRv6 IPv6 header for inner L2 and L3 traffic. Use **propagate** to set the traffic-class value by propagation (from incoming packet/frame). Use **policy-map** to set traffic-class DSCP to qos-group selected by input policy-map.

The **I3-traffic** configuration supports the additional **policy-map** option that sets SRv6 traffic-class DSCP based on qos-group selected by input **policy-map**.

The **I2-traffic** configuration continues to support the **propagate** option that sets the SRv6 traffic-class based on either the inner packet (default) or input **policy-map** (if match action found).

Starting with Cisco IOS XR Software Release 24.2.1, both the L2 and L3 traffic configurations support the **policy-map-extend** option to set the SRv6 traffic-class DSCP marking to 6 bits. This is a mutually inclusive feature for L2 and L3 traffic. Therefore, this feature cannot be used in the individual configuration for L2 and L3 traffic.

When you use the **policy-map-extend** keyword in the **hw-module profile segment-routing srv6 mode microsegment f3216 encapsulation traffic-class** CLI, the existing 3-bit DSCP marking set by the **hw-module profile segment-routing srv6 mode microsegment f3216 encapsulation traffic-class I3-traffic policy-map** CLI in the L3 traffic configuration is modified to 6-bits.

Go through the following topics for more information on the CLI:

- [hw-module profile segment-routing srv6 mode](#)
- [encapsulation I2-traffic](#)
- [encapsulation I3-traffic](#)
- [traffic-class policy-map-extend](#)

Specify the *value* (as 2 hexadecimal nibbles) for traffic class; valid values are from 0x0 to 0xff.

Configure SRv6 Locator Name, Prefix, and uSID-Related Parameters

This section shows how to globally enable SRv6 and configure locator.

- **segment-routing srv6 locators locator *locator***—Globally enable SRv6 and configure the locator.
- **segment-routing srv6 locators locator *locator* prefix *ipv6_prefix/length***—Configure the locator prefix value.

- **segment-routing srv6 locators locator** *locator* **micro-segment behavior unode psp-usd**—Specifies the locator as a micro-segment (uSID) locator as well as specifies that IGP underlay uSID (uN/uA) variant is PSP-USD for this locator.

(Optional) Configure Algorithm Associated with Locator

- **segment-routing srv6 locators locator** *locator* **algorithm** *algo*—(Optional) Configure Algorithm associated with the locator. Valid values for *algo* are from 128 to 255.

For additional information about SRv6 Flexible Algorithm, see [Configuring SRv6 Flexible Algorithm under IS-IS, on page 36](#).

For detailed information about Flexible Algorithm, see [Enabling Segment Routing Flexible Algorithm, on page 573](#).

(Optional) Configure Anycast Locator

An SRv6 Anycast locator is a type of locator that identifies a set of nodes (uN SIDs). SRv6 Anycast Locators and their associated uN SIDs may be provisioned at multiple places in a topology.

The set of nodes (Anycast group) is configured to advertise a shared Anycast locator and uN SID. Anycast routing enables the steering of traffic toward multiple advertising nodes. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.

One use case is to advertise Anycast uN SIDs at exit points from an SRv6 network. Any of the nodes that advertise the common uN SID could be used to forward traffic out of the SRv6 portion of the network to the topologically nearest node.

The following behaviors apply to Anycast Locator:

- Unlike a normal locator, IS-IS does not program or advertise uA SIDs associated with an Anycast locator.
- uN SIDs allocated from Anycast locators will not be used in constructing TI-LFA backup paths or Microloop Avoidance primary paths. TI-LFA backup and Microloop Avoidance paths for an Anycast locator prefix may terminate on any node advertising that locator, which may be different from the node terminating the original primary path.
- SRv6 anycast locators may have non-zero algorithm (Flexible Algorithm) values.

Use the following commands to configure the Anycast locator and advertise Anycast prefixes associated with an interface.

- **segment-routing srv6 locators locator** *locator* **anycast**—Configure the Anycast locator
- **router isis** *instance-id* **interface** *Loopback* *instance* **prefix-attributes anycast level** *level*—Advertise the Anycast prefixes associated with an interface.

Example 1:

The following example shows how to globally enable SRv6 and configure a locator.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLoc1
Router(config-srv6-locator)# micro-segment behavior unode psp-usd
Router(config-srv6-locator)# prefix 2001:0:8::/48
```

Example 2:

The following example shows how to configure Flexible Algorithm associated with locator.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLocAlgo128
Router(config-srv6-locator)# algorithm 128
Router(config-srv6-locator)# micro-segment behavior unode psp-usd
Router(config-srv6-locator)# prefix 2001:0:88::/48
```

Example 3:

The following example shows how to configure Anycast locator.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLocAnycast
Router(config-srv6-locator)# anycast
Router(config-srv6-locator)# micro-segment behavior unode psp-usd
Router(config-srv6-locator)# prefix 2001:0:100::/48
```

The following example shows how to advertise the Anycast prefixes associated with an interface.

```
Router(config)# router isis core
Router(config-isis)# interface Loopback100
Router(config-isis-if)# prefix-attributes anycast level 1
```

(Optional) Customize SRv6 Encapsulation Parameters

This section describes the configurable SRv6 encapsulation parameters. These optional parameters include:

- **segment-routing srv6 encapsulation source-address *ipv6-addr***—Source Address of outer encapsulating IPv6 header. The default source address for encapsulation is one of the loopback addresses.
- **segment-routing srv6 encapsulation hop-limit {*count* | **propagate**}**—The hop limit of outer-encapsulating IPv6 header. The range for *count* is from 1 to 254; the default value for hop-limit is 254. Use **propagate** to set the hop-limit value by propagation (from incoming packet/frame).
- **segment-routing srv6 encapsulation evpn next-header *protocol-number***—The protocol number to use in the Next-header field of the IPv6 or SRH header. The range for *protocol-number* is from 59 to 252.

(Optional) Customize SRv6 Logging for Locator Status Changes

- **segment-routing srv6 logging locator status**—Enable the logging of locator status.

(Optional) Customize SRv6 SID Parameters

- **segment-routing srv6 sid holdtime *minutes***—The holdtime for a stale or freed SID. The range of *minutes* is from 0 (disabled) to 60 minutes.

Example 4:

The following example shows how to configure optional SRv6 parameters:

```
RP/0/RSP0/CPU0:Node1(config)# segment-routing srv6 encapsulation
RP/0/RSP0/CPU0:Node1(config-srv6-encap)# source-address 1::1
RP/0/RSP0/CPU0:Node1(config-srv6-encap)# hop-limit 60
```

```
RP/0/RSP0/CPU0:Node1(config-srv6-encap)# evpn next-header 65
RP/0/RSP0/CPU0:Node1(config-srv6-encap)# exit
RP/0/RSP0/CPU0:Node1(config-srv6)# logging locator status
RP/0/RSP0/CPU0:Node1(config-srv6)# sid holdtime 10
RP/0/RSP0/CPU0:Node1(config-srv6)# micro-segment merge-overlay-underlay-sids
```

This config applies to only new SRv6 micro-segment overlay routes and does not update already programmed routes.

Please flap any existing SRv6 micro-segment overlay routes after making this configuration change.

```
RP/0/RSP0/CPU0:Node1(config-srv6)#
```

Verifying SRv6 Manager

This example shows how to verify the overall SRv6 state from SRv6 Manager point of view. The output displays parameters in use, summary information, and platform specific capabilities.

```
Router# show segment-routing srv6 manager
Parameters:
  SRv6 Enabled: Yes
  SRv6 Operational Mode:
    Micro-segment:
      SID Base Block: 2001::/24
  Encapsulation:
    Source Address:
      Configured: ::
      Default: ::
    Hop-Limit: Default
    Traffic-class: Default
Summary:
  Number of Locators: 3 (3 operational)
  Number of SIDs: 3 (0 stale)
  Max SIDs: 64000
  OOR:
    Thresholds: Green 3200, Warning 1920
    Status: Resource Available
    History: (0 cleared, 0 warnings, 0 full)
  Block 2001::/32:
    Number of SIDs free: 7680
    Max SIDs: 7680
    Thresholds: Green 384, Warning 231
    Status: Resource Available
    History: (0 cleared, 0 warnings, 0 full)
Platform Capabilities:
  SRv6: Yes
  TILFA: Yes
  Microloop-Avoidance: Yes
  Endpoint behaviors:
    End (PSP)
    End.X (PSP)
    End.DX6
    End.DX4
    End.DT6
    End.DT4
    End.OP
    uN (PSP/USD)
    uA (PSP/USD)
    uDT6
    uDT4
    uDX2
    uB6 (Insert.Red)
  Headend behaviors:
```

```

T
H.Insert.Red
H.Encaps.Red
Security rules:
SEC-1
SEC-2
SEC-3
Counters:
CNT-1
CNT-3
Signaled parameters:
Max-SL      : 3
Max-End-Pop-SRH : 3
Max-H-Insert : 3 sids
Max-H-Encap  : 3 sids
Max-End-D    : 4
Configurable parameters (under srv6):
Encapsulation:
Source Address: Yes
Hop-Limit     : value=Yes, propagate=No
Traffic-class  : value=Yes, propagate=Yes
Max SIDs: 64000
SID Holdtime: 3 mins

```

Verifying SRv6 Locator

This example shows how to verify the locator configuration and its operational status.

```
Router# show segment-routing srv6 locator myLoc1 detail
```

Name	ID	Algo	Prefix	Status	Flags
myLoc1	3	0	2001:0:8::/48	Up	U

(U): Micro-segment (behavior: uN (PSP/USD))

Interface:

Name: srv6-myLoc1

IFH : 0x02000120

IPv6 address: 2001:0:8::/48

Number of SIDs: 1

Created: Dec 10 21:26:54.407 (02:52:26 ago)

Verifying SRv6 SIDs

This example shows how to verify the allocation of SRv6 local SIDs off locator(s).

```
Router# show segment-routing srv6 locator myLoc1 sid
```

SID	State	RW	Behavior	Context	Owner
2001:0:8::	InUse	Y	uN (PSP/USD)	'default':1	sidmgr

The following example shows how to display detail information regarding an allocated SRv6 local SID.

```
Router# show segment-routing srv6 locator myLoc1 sid 2001:0:8:: detail
```

SID	State	RW	Behavior	Context	Owner
2001:0:8::	InUse	Y	uN (PSP/USD)	'default':1	sidmgr

```

2001:0:8::                                uN (PSP/USD)      'default':8      sidmgr
      InUse Y
SID Function: 0x8
SID context: { table-id=0xe0800000 ('default':IPv6/Unicast), opaque-id=8 }
Locator: 'myLoc1'
Allocation type: Dynamic
Created: Dec 10 22:10:51.596 (02:10:05 ago)

```

Similarly, you can display SID information across locators by using the **show segment-routing srv6 sid** command.

Configuring SRv6 under IS-IS

Intermediate System-to-Intermediate System (IS-IS) protocol already supports segment routing with MPLS dataplane (SR-MPLS). This feature enables extensions in IS-IS to support Segment Routing with IPv6 data plane (SRv6). The extensions include advertising the SRv6 capabilities of nodes and node and adjacency segments as SRv6 SIDs.

SRv6 IS-IS performs the following functionalities:

1. Interacts with SID Manager to learn local locator prefixes and announces the locator prefixes in the IGP domain.
2. Learns remote locator prefixes from other IS-IS neighbor routers and installs the learned remote locator IPv6 prefix in RIB or FIB.
3. Allocate or learn prefix SID and adjacency SIDs, create local SID entries, and advertise them in the IGP domain.

Usage Guidelines and Restrictions

The following usage guidelines and restrictions apply for SRv6 IS-IS:

- An IS-IS address-family can support either SR-MPLS or SRv6, but both at the same time is not supported.

Configuring SRv6 under IS-IS

To configure SRv6 IS-IS, use the **router isis** command. Enable SRv6 under the IS-IS IPv6 address-family and assign SRv6 locator(s) to it. Use the **level {1 | 2}** keywords to advertise the locator only in the specified IS-IS level.



Note If no level is specified, local locators will be advertised into all configured ISIS levels. Ensure that locators are included in the redistribution or propagation policy to prevent potential loops when redistributing between multiple instances or propagating between Level 2 and Level 1.

The following example shows how to configure SRv6 under IS-IS.

```

Router(config)# router isis core
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing srv6
Router(config-isis-srv6)# locator myLoc1 level 1

```

```
Router(config-isis-srv6-loc)# exit
```

For more information about configuring IS-IS, refer to the ["Implementing IS-IS"](#) chapter in the *Routing Configuration Guide*.

Configuring SRv6 Flexible Algorithm under IS-IS

This feature introduces support for implementing Flexible Algorithm using IS-IS SRv6.

SRv6 Flexible Algorithm allows operators to customize IGP shortest path computation according to their own needs. An operator can assign custom SR prefix-SIDs to realize forwarding beyond link-cost-based SPF. As a result, Flexible Algorithm provides a traffic-engineered path automatically computed by the IGP to any destination reachable by the IGP.

For detailed information about Flexible Algorithm, see [Enabling Segment Routing Flexible Algorithm, on page 573](#).

Usage Guidelines and Restrictions

Observe the following usage guidelines and restrictions:

- You can configure up to 8 locators to support SRv6 Flexible Algorithm.
- The Flexible Algorithm locator prefix follows the same usage guidelines and restrictions of algo-0 locator prefixes. See [Usage Guidelines and Limitations, on page 25](#).
- The Locator Algorithm value range is 128 to 255.

Configuring SRv6 Flexible Algorithm under IS-IS

The following sections show you the steps to enable SRv6 Flexible Algorithm. The example highlights a delay-based Flexible Algorithm instance.

1. Configure SRv6 locators
2. Assign SRv6 locators under IS-IS
3. Configure Flexible Algorithm definition and associated metric (for example, delay)
4. Configure the delay probe under the interface. For more information on SR performance measurement, see [Configure performance measurement](#).

The following section shows how to configure two SRv6 locators: one associated with Algo 0, and the other associated with Algo 128.

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLocBestEffort // best-effort locator
Router(config-srv6-locator)# micro-segment behavior unode psp-usd
Router(config-srv6-locator)# prefix 2001:0:1::/48
Router(config-srv6-locator)# exit

Router(config-srv6-locators)# locator myLocLowLat // low-latency (flex algo 128) locator
Router(config-srv6-locator)# micro-segment behavior unode psp-usd
Router(config-srv6-locator)# prefix 2001:0:2::/48
```

```

Router(config-srv6-locator)# algorithm 128
Router(config-srv6-locator)# exit
Router(config-srv6)# exit

```

The following section shows how to assign multiple SRv6 locators under IS-IS.

```

Router(config)# router isis core
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing srv6
Router(config-isis-srv6)# locator myLocBestEffort
Router(config-isis-srv6-loc)# exit
Router(config-isis-srv6)# locator myLocLowLat
Router(config-isis-srv6-loc)# exit

```

The following section shows how to configure the Flexible Algorithm definition.

```

Router(config)# router isis core
Router(config-isis)# flex-algo 128
Router(config-isis-flex-algo)# metric-type delay
Router(config-isis-flex-algo)# exit
Router(config-isis)# interface GigabitEthernet0/0/0/0
Router(config-isis-if)# address-family ipv6 unicast

```

The following section shows how to configure the delay probe under the interface.

```

Router(config)# performance-measurement
Router(config-perf-meas)# interface GigabitEthernet0/0/0/0
Router(config-pm-intf)# delay-measurement
Router(config-pm-intf-dm)# commit

```

Verification

```
Router# show segment-routing srv6 locator
```

Name	ID	Algo	Prefix	Status	Flags
myLoc1	3	0	2001:0:8::/48	Up	U
myLocBestEffort	5	0	2001:0:1::/48	Up	U
myLocLowLat	4	128	2001:0:2::/48	Up	U

```
Router# show isis flex-algo 128
```

```
IS-IS core Flex-Algo Database
```

```
Flex-Algo 128:
```

```
Level-2:
```

```

Definition Priority: 128
Definition Source: Router.00, (Local)
Definition Equal to Local: Yes
Disabled: No

```

```
Level-1:
```

```

Definition Priority: 128
Definition Source: Router.00, (Local)
Definition Equal to Local: Yes

```

```

Disabled: No

Local Priority: 128
FRR Disabled: No
Microloop Avoidance Disabled: No

```

Configuring SRv6 Locator Prefix Summarization

SRv6 leverages longest-prefix-match IP forwarding. Massive-scale reachability can be achieved by summarizing locators at ABRs and ASBRs.

Use the **summary-prefix locator** [**algorithm algo**] [**explicit**] command in IS-IS address-family configuration mode to specify that only locators from the specified algorithm contribute to the summary. The **explicit** keyword limits the contributing prefixes to only those belonging to the same algorithm.

The following example shows how to configure SRv6 IS-IS Algorithm Summarization for regular algorithm and Flexible Algorithm (128).

```

Router(config)# router isis core
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# summary-prefix 2001:0:1::/48
Router(config-isis-af)# summary-prefix 2001:0:2::/48 algorithm 128 explicit

```

Configuring TI-LFA with SRv6 IS-IS

This feature introduces support for implementing Topology-Independent Loop-Free Alternate (TI-LFA) using SRv6 IS-IS.

TI-LFA provides link protection in topologies where other fast reroute techniques cannot provide protection. The goal of TI-LFA is to reduce the packet loss that results while routers converge after a topology change due to a link failure. TI-LFA leverages the post-convergence path which is planned to carry the traffic and ensures link and node protection within 50 milliseconds. TI-LFA with IS-IS SR-MPLS is already supported.

TI-LFA provides link, node, and Shared Risk Link Groups (SRLG) protection in any topology.

For more information, see [Configure Topology-Independent Loop-Free Alternate \(TI-LFA\), on page 527](#).

Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- TI-LFA provides link protection by default. Additional tiebreaker configuration is required to enable node or SRLG protection.
- Usage guidelines for node and SRLG protection:
 - TI-LFA node protection functionality provides protection from node failures. The neighbor node is excluded during the post convergence backup path calculation.
 - Shared Risk Link Groups (SRLG) refer to situations in which links in a network share a common fiber (or a common physical attribute). These links have a shared risk: when one link fails, other links in the group might also fail. TI-LFA SRLG protection attempts to find the post-convergence

backup path that excludes the SRLG of the protected link. All local links that share any SRLG with the protecting link are excluded.

- When you enable link protection, you can also enable node protection, SRLG protection, or both, and specify a tiebreaker priority in case there are multiple LFAs.
- Valid priority values are from 1 to 255. The lower the priority value, the higher the priority of the rule. Link protection always has a lower priority than node or SRLG protection.

Configuring SRv6 IS-IS TI-LFA

The following example shows how to configure different types of TI-LFA protection for SRv6 IS-IS.

```
Router(config)# router isis core
Router(config-isis)# interface bundle-ether 1201
Router(config-isis-if)# address-family ipv6 unicast
Router(config-isis-if-af)# fast-reroute per-prefix
Router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
Router(config-isis-if-af)# exit
Router(config-isis-if)# exit
Router(config-isis)# interface bundle-ether 1301
Router(config-isis-if)# address-family ipv6 unicast
Router(config-isis-if-af)# fast-reroute per-prefix
Router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
Router(config-isis-if-af)# fast-reroute per-prefix tiebreaker node-protecting index 100
Router(config-isis-if-af)# fast-reroute per-prefix tiebreaker srlg-disjoint index 200
Router(config-isis-if-af)# exit
```

Configuring SRv6 IS-IS TI-LFA with Flexible Algorithm

TI-LFA backup paths for particular Flexible Algorithm are computed using the same constraints as the calculation of the primary paths for such Flexible Algorithm. These paths use the locator prefix advertised specifically for such Flexible Algorithm in order to enforce a backup path.

By default, LFA/TI-LFA for SRv6 Flexible Algorithm uses the LFA/TI-LFA configuration of Algo 0.

Use the **fast-reroute disable** command to disable the LFA/TI-LFA calculation on a per-algorithm basis:

```
Router(config)# router isis core
Router(config-isis)# flex-algo 128
Router(config-isis-flex-algo)# fast-reroute disable
```

Verification

This example shows how to verify the SRv6 IS-IS TI-LFA configuration using the **show isis ipv6 fast-reroute ipv6-prefix detail** command.

```
Router# show isis ipv6 fast-reroute cafe:0:2::2/128 detail

L2 cafe:0:2::2/128 [20/115] Label: None, medium priority
   via fe80::e00:ff:fe3a:c700, HundredGigE0/0/0/0, Node2, Weight: 0
   Backup path: TI-LFA (link), via fe80::1600:ff:feec:fe00, HundredGigE0/0/0/1 Node3,
Weight: 0, Metric: 40
      P node: Node4.00 [cafe:0:4::4], SRv6 SID: cafe:0:4:: uN (PSP/USD)
      Backup-src: Node2.00
      P: No, TM: 40, LC: No, NP: No, D: No, SRLG: Yes
      src Node2.00-00, cafe:0:2::2
```


This example shows how to verify the SRv6 IS-IS TI-LFA configuration using the **show route ipv6 ipv6-prefix detail** command.

```
Router# show route ipv6 cafe:0:2::2/128 detail
Tue Feb 23 23:08:48.151 UTC

Routing entry for cafe:0:2::2/128
  Known via "isis 1", distance 115, metric 20, type level-2
  Installed Feb 23 22:57:38.900 for 00:11:09
  Routing Descriptor Blocks
    fe80::1600:ff:feec:fe00, from cafe:0:2::2, via HundredGigE0/0/0/1, Backup (TI-LFA)
      Repair Node(s): cafe:0:4::4
      Route metric is 40
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:65          Path ref count:1
      NHID:0x20002(Ref:19)
      SRv6 Headend: H.Insert.Red [f3216], SID-list {cafe:0:4::}
    fe80::e00:ff:fe3a:c700, from cafe:0:2::2, via HundredGigE0/0/0/0, Protected
      Route metric is 20
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:1          Path ref count:0
      NHID:0x20001(Ref:19)
      Backup path id:65
  Route version is 0x4 (4)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Flow-tag: Not Set
  Fwd-class: Not Set
  Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type RIB_SVD_TYPE_LOCAL
  Download Priority 1, Download Version 66
  No advertising protos.
```

This example shows how to verify the SRv6 IS-IS TI-LFA configuration using the **show cef ipv6 ipv6-prefix detail location location** command.

```
Router# show cef ipv6 cafe:0:2::2/128 detail location 0/0/cpu0
Tue Feb 23 23:09:07.719 UTC
cafe:0:2::2/128, version 66, SRv6 Headend, internal 0x1000001 0x210 (ptr 0x8e96fd2c) [1],
0x0 (0x8e93fae0), 0x0 (0x8f7510a8)
Updated Feb 23 22:57:38.904
local adjacency to HundredGigE0/0/0/0

Prefix Len 128, traffic index 0, precedence n/a, priority 1
gateway array (0x8e7b5c78) reference count 1, flags 0x500000, source rib (7), 0 backups
[2 type 3 flags 0x8401 (0x8e86ea40) ext 0x0 (0x0)]
LW-LDI[type=3, refc=1, ptr=0x8e93fae0, sh-ldi=0x8e86ea40]
gateway array update type-time 1 Feb 23 22:57:38.904
LDI Update time Feb 23 22:57:38.913
LW-LDI-TS Feb 23 22:57:38.913
  via fe80::1600:ff:feec:fe00/128, HundredGigE0/0/0/1, 9 dependencies, weight 0, class 0,
  backup (TI-LFA) [flags 0xb00]
    path-idx 0 NHID 0x20002 [0x8f5850b0 0x0]
    next hop fe80::1600:ff:feec:fe00/128, Repair Node(s): cafe:0:4::4
    local adjacency
    SRv6 H.Insert.Red SID-list {cafe:0:4::}
  via fe80::e00:ff:fe3a:c700/128, HundredGigE0/0/0/0, 6 dependencies, weight 0, class 0,
```

```
protected [flags 0x400]
  path-idx 1 bkup-idx 0 NHID 0x20001 [0x8f8420b0 0x0]
  next hop fe80::e00:ff:fe3a:c700/128

Load distribution: 0 (refcount 2)

Hash  OK  Interface          Address
0      Y   HundredGigE0/0/0/0  fe80::e00:ff:fe3a:c700
```

Configuring SRv6 IS-IS Microloop Avoidance

This feature introduces support for implementing microloop avoidance using IS-IS SRv6.

Microloops are brief packet loops that occur in the network following a topology change (link down, link up, or metric change events). Microloops are caused by the non-simultaneous convergence of different nodes in the network. If nodes converge and send traffic to a neighbor node that has not converged yet, traffic may be looped between these two nodes, resulting in packet loss, jitter, and out-of-order packets.

The SRv6 Microloop Avoidance feature detects if microloops are possible following a topology change. If a node computes that a microloop could occur on the new topology, the node creates a loop-free SR-TE policy path to the destination using a list of segments. After the RIB update delay timer expires, the SR-TE policy is replaced with regular forwarding paths.

Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- The Routing Information Base (RIB) update delay value specifies the amount of time the node uses the microloop avoidance policy before updating its forwarding table. The *delay-time* range is from 1 to 60000 milliseconds; the default value is 5000.

Configuring SRv6 IS-IS Microloop Avoidance

The following example shows how to configure SRv6 IS-IS Microloop Avoidance and set the Routing Information Base (RIB) update delay value.



Note Complete the [Configuring SRv6, on page 27](#) before performing these steps.

```
Router(config)# router isis test-igp
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# microloop avoidance segment-routing
Router(config-isis-af)# microloop avoidance rib-update-delay 2000
Router(config-isis-af)# commit
```

Configuring SRv6 IS-IS Microloop Avoidance with Flexible Algorithm

Microloop Avoidance paths for particular Flexible Algorithm are computed using the same constraints as the calculation of the primary paths for such Flexible Algorithm. These paths use the Locator prefix advertised specifically for such Flexible Algorithm in order to enforce a microloop avoidance path.

By default, Microloop Avoidance for SRv6 Flexible Algorithm uses the Microloop Avoidance configuration of Algo 0.

Use the **microloop avoidance disable** command to disable the microloop calculation on a per-algorithm basis:

```
Router(config)# router isis test-tilfa
Router(config-isis)# flex-algo 128
Router(config-isis-flex-algo)# microloop avoidance disable
```

Configuring SRv6 BGP-Based Services

Building on the messages and procedures defined in IETF draft "[BGP/MPLS IP Virtual Private Networks \(VPNs\)](#)", BGP has been extended to provide services over an SRv6 network, such as:

- IPv4 Layer-3 VPNs
- IPv6 Layer-3 VPNs
- IPv4 BGP global
- IPv6 BGP global
- Layer-2 VPNs - Ethernet VPNs (EVPN)

For more information about BGP, refer to the BGP Configuration Guide.

In SRv6-based services, the egress PE signals an SRv6 Service SID with the BGP service route. The ingress PE encapsulates the payload in an outer IPv6 header where the destination address is the SRv6 Service SID advertised by the egress PE. BGP messages between PEs carry SRv6 Service SIDs as a means to interconnect PEs and form VPNs. SRv6 Service SID refers to a segment identifier associated with one of the SRv6 service-specific behaviors advertised by the egress PE router, such as:

- uDT4 (Endpoint with decapsulation and IPv4 table lookup)
- uDT6 (Endpoint with decapsulation and IPv6 table lookup)
- uDX4 (Endpoint with decapsulation and IPv4 cross-connect)
- uDX6 (Endpoint with decapsulation and IPv6 cross-connect)

Based on the messages and procedures defined in IETF draft "[SRv6 BGP based Overlay services](#)", BGP encodes the SRv6 Service SID in the prefix-SID attribute of the corresponding BGP Network Layer Reachability Information (NLRI) and advertises it to its IPv6 BGP peers.

Usage Guidelines and Restrictions

- The following SRv6 BGP-based services are supported:
 - [IPv4 Layer-3 VPNs](#)
 - [IPv6 Layer-3 VPNs](#)
 - [IPv4 BGP global](#)
 - [IPv6 BGP global](#)

- uDT4 and uDT6 for L3VPN and BGP global are supported.
- Dual-Stack L3 Services (IPv4 L3VPN, IPv6 L3VPN, IPv4 BGP global, IPv6 BGP global) are supported.

SRv6 Locator Inheritance Rules

SRv6 locators can be assigned at different levels inside the BGP routing process. BGP allocates SRv6 Service SIDs from configured locator spaces according to the following inheritance rules:

1. Use the locator as defined under the service.
If not defined under the specific service, then:
2. Use the locator as defined under the corresponding address-family.
If not defined under the corresponding address-family, then:
3. Use the locator as defined globally under BGP.

Enabling SRv6 Globally under BGP

Use the **router bgp *as-number* segment-routing srv6** command to enable SRv6 globally under the BGP routing process. The *as-number* is from 1-65535.

```
RP/0/0/CPU0:Node1(config)# router bgp 100 segment-routing srv6
```

Assigning SRv6 Locator Globally under BGP

Use the **router bgp *as-number* segment-routing srv6 locator *WORD*** command to assign an SRv6 locator globally under the BGP routing process. The *as-number* is from 1-65535.

This example shows how to assign a locator:

```
RP/0/0/CPU0:Node1(config)# router bgp 100 segment-routing srv6 locator Node1-locator
```

For more information on how to configure an SRv6 locator, see [Configuring SRv6, on page 27](#).

For more information on how to assign an SRv6 locator under the BGP service or BGP address-family, see the following SRv6 Services sections.

SRv6 Services: IPv4 L3VPN

Table 5: Feature History Table

Feature Name	Release	Description
Per-Prefix SRv6 Locator Assignment	Release 7.5.1	This feature provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes (IPv4/IPv6 GRT, IPv4/IPv6 VPN). The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.
Support for iBGP as PE-CE protocol	Release 7.5.1	This feature introduces support for iBGP as PE-CE protocol.
BGP Route Leaking	Release 7.5.1	This feature adds support for importing routes from default-VRF to non-default VRF and routes from non-default VRF to default VRF.

This feature provides IPv4 L3VPNs (VPNv4) over an SRv6 network.

Usage Guidelines and Limitations

- SRv6 locator can be assigned globally, for all VRFs, or for an individual VRF.
SRv6 locator can be assigned globally, for all VRFs, for an individual VRF, or per-prefix.
- Per-VRF allocation mode is supported (uDT4 behavior)
- Dual-Stack L3VPN Services (IPv4, IPv6) are supported
- Equal-Cost Multi-path (ECMP) and Unequal Cost Multipath (UCMP) are supported.
- eBGP, OSPF, Static are supported as PE-CE protocol.
BGP (iBGP, eBGP), OSPF, Static are supported as PE-CE protocol.
- BGP route leaking between BGP Global and L3VPN is supported.
- Per-CE allocation mode is not supported (uDX4 behavior)
- iBGP is not supported as PE-CE protocol
- BGP route leaking is not supported

Configuring SRv6 based IPv4 L3VPN

To enable SRv6-based L3VPN, you need to enable SRv6 under BGP, specify the locator, and configure the SID allocation mode. The assignment of the locator can be done in different places under the **router bgp** configuration. See [SRv6 Locator Inheritance Rules, on page 43](#).

Use Case 1: Assigning SRv6 Locator Globally

This example shows how to enable SRv6 and configure the SRv6 locator name under BGP Global:

```
Node1(config)# router bgp 100
Node1(config-bgp)# segment-routing srv6
Node1(config-bgp-gbl-srv6)# locator Node1-locator
Node1(config-bgp-gbl-srv6)# exit
Node1(config-bgp)# address-family vpnv4 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::1:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv4 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust1
Node1(config-bgp-vrf)# rd 100:1
Node1(config-bgp-vrf)# address-family ipv4 unicast
Node1(config-bgp-vrf-af)# commit
```

Running Config

```
router bgp 100
  segment-routing srv6
    locator Node1-locator
  !
  address-family vpnv4 unicast
  !
  neighbor 3001::1:1:1:4
    remote-as 100
    address-family vpnv4 unicast
  !
  !
  vrf vrf_cust1
    rd 100:1
    address-family ipv4 unicast
  !
  !
  !
end
```

Use Case 2: Assigning SRv6 Locator for All VRFs

To configure the SRv6 locator for all VRFs under VPNv4 Address Family and specify the allocation mode, use the following commands:

- **router bgp *as-number* address-family vpnv4 unicast vrf all segment-routing srv6:** Enable SRv6
- **router bgp *as-number* address-family vpnv4 unicast vrf all segment-routing srv6 alloc mode {*per-vrf*}:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT4 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp *as-number* address-family vpnv4 unicast vrf all segment-routing srv6 locator *WORD*:** Specify the locator

This example shows how to enable SRv6 and configure the SRv6 locator for all VRFs under VPNv4 Address Family, with per-VRF label allocation mode:

```
Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv4 unicast
```

```

Node1(config-bgp-af)# vrf all
Node1(config-bgp-af-vrfall)# segment-routing srv6
Node1(config-bgp-af-vrfall-srv6)# locator Node1-locator
Node1(config-bgp-af-vrfall-srv6)# alloc mode per-vrf
Node1(config-bgp-af-vrfall-srv6)# exit
Node1(config-bgp-af-vrfall)# exit
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::1:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv4 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust1
Node1(config-bgp-vrf)# rd 100:1
Node1(config-bgp-vrf)# address-family ipv4 unicast
Node1(config-bgp-vrf-af)# commit

```

Running Config

```

router bgp 100
  address-family vpnv4 unicast
    vrf all
      segment-routing srv6
        locator Node1-locator
        alloc mode per-vrf
    !
  !
  neighbor 3001::1:1:1:4
    remote-as 100
    address-family vpnv4 unicast
  !
  !
  vrf vrf_cust1
    rd 100:1
    address-family ipv4 unicast
  !
  !
end

```

Use Case 3: Assigning SRv6 Locator for a specific VRF

To configure the SRv6 locator for a specific VRF under IPv4 Address Family and specify the allocation mode, use the following commands:

- **router bgp as-number vrf WORD address-family ipv4 unicast segment-routing srv6:** Enable SRv6
- **router bgp as-number vrf WORD address-family ipv4 unicast segment-routing srv6 alloc mode { per-vrf }:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT4 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp as-number vrf WORD address-family ipv4 unicast segment-routing srv6 locator WORD:** Specify the locator

This example shows how to configure the SRv6 locator for an individual VRF, with per-VRF label allocation mode:

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv4 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::1:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv4 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust1
Node1(config-bgp-vrf)# rd 100:1
Node1(config-bgp-vrf)# address-family ipv4 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6
Node1(config-bgp-vrf-af-srv6)# locator Node1-locator
Node1(config-bgp-vrf-af-srv6)# alloc mode per-vrf
Node1(config-bgp-vrf-af-srv6)# commit

```

Running Config

```

router bgp 100
 address-family vpnv4 unicast
 !
 neighbor 3001::1:1:1:4
  remote-as 100
  address-family vpnv4 unicast
  !
 !
vrf vrf_cust1
 rd 100:1
  address-family ipv4 unicast
   segment-routing srv6
   locator Node1-locator
   alloc mode per-vrf
  !
 !
 !
end

```

Use Case 4: Assigning SRv6 Locator for a Specific Prefix

This use case provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes. The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.

To assign an SRv6 locator for a specific prefix, configure a route policy to specify the SID allocation mode based on match criteria. Examples of match criteria are destination-based match or community-based match.

- Supported SID allocation modes are per-VRF and per-CE.
- For per-VRF allocation mode, you can also specify the SRv6 locator.
 - If an SRv6 locator is specified in the route policy, BGP will use that to allocate per-VRF SID. If the specified locator is invalid, the SID will not be allocated.
 - If an SRv6 locator is not specified in the route policy, the default locator configured under BGP is used to allocate the SID. If the default locator is not configured, then the SID will not be allocated.
- Per-CE allocation mode always uses the default locator configured under BGP to allocate the SID.

For more information on configuring routing policies, refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following example shows a route policy specifying the SID allocation mode with destination-based match:

```

Node1(config)# route-policy set_per_prefix_locator_rpl
Node1(config-rpl)# if destination in (10.1.1.0/24) then
Node1(config-rpl-if)# set srv6-alloc-mode per-vrf locator locator1
Node1(config-rpl-if)# elseif destination in (2.2.2.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf locator locator2
Node1(config-rpl-elseif)# elseif destination in (3.3.3.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf
Node1(config-rpl-elseif)# elseif destination in (4.4.4.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-ce
Node1(config-rpl-elseif)# else
Node1(config-rpl-elseif)# drop
Node1(config-rpl-elseif)# endif
Node1(config-rpl)# end-policy
Node1(config)#

```

To specify per-prefix allocation mode for a specific VRF under IPv4 address family, use the following command:

- **router bgp** *as-number* **vrf** *WORD* **address-family** **ipv4** **unicast** **segment-routing** **srv6** **alloc mode** **route-policy** *policy_name*

This example shows how to configure per-prefix allocation mode for a specific VRF (vrf_cust1) under IPv4 address family

```

Node1(config)# router bgp 100
Node1(config-bgp)# vrf vrf_cust1
Node1(config-bgp-vrf)# address-family ipv4 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6
Node1(config-bgp-vrf-af-srv6)# alloc mode route-policy set_per_prefix_locator_rpl

```

Running Configuration

```

route-policy set_per_prefix_locator_rpl
  if destination in (10.1.1.0/24) then
    set srv6-alloc-mode per-vrf locator locator1
  elseif destination in (2.2.2.0/24) then
    set srv6-alloc-mode per-vrf locator locator2
  elseif destination in (3.3.3.0/24) then
    set srv6-alloc-mode per-vrf
  elseif destination in (4.4.4.0/24) then
    set srv6-alloc-mode per-ce
  else
    drop
  endif
end-policy
!
router bgp 100
  vrf vrf_cust1
    address-family ipv4 unicast
      segment-routing srv6
        alloc mode route-policy set_per_prefix_locator_rpl
    !
  !
!

```

Verify that the local and received SIDs have been correctly allocated under VPNv4 and specific VRF (vrf_cust1):

```

Node1# show bgp vpnv4 unicast local-sids
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Local Sid	Alloc mode	Locator
Route Distinguisher: 8:8			
*>i8.8.8.8/32	NO SRv6 Sid	-	-
* i	NO SRv6 Sid	-	-
Route Distinguisher: 10.1.1.1:0 (default for vrf vrf_cust1)			
*> 10.1.1.0/24	fc00:0:1:40::	per-vrf	locator1
*> 2.2.2.0/24	fc00:8:1:40::	per-vrf	locator2
*> 3.3.3.0/24	fc00:9:1:40::	per-vrf	locator4
*> 4.4.4.0/24	fc00:9:1:41::	per-ce	locator4
*> 10.1.1.5/32	NO SRv6 Sid	-	-
*> 3.3.3.3/32	NO SRv6 Sid	-	-
*>i8.8.8.8/32	NO SRv6 Sid	-	-

```

Node1# show bgp vpnv4 unicast received-sids
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Received Sid
Route Distinguisher: 8:8		
*>i8.8.8.8/32	10.1.1.2	fc00:0:2:42::
* i	2400:2020:42:2fff::1	fc00:0:2:42::
Route Distinguisher: 10.1.1.1:0 (default for vrf vrf_cust1)		
*> 10.1.1.0/24	11.1.1.2	NO SRv6 Sid
*> 2.2.2.0/24	11.1.1.2	NO SRv6 Sid
*> 3.3.3.0/24	11.1.1.2	NO SRv6 Sid
*> 4.4.4.0/24	11.1.1.2	NO SRv6 Sid
*> 10.1.1.5/32	11.1.1.2	NO SRv6 Sid
*> 3.3.3.3/32	13.2.2.2	NO SRv6 Sid
*>i8.8.8.8/32	10.1.1.2	fc00:0:2:42::

```

Node1# show bgp vrf vrf_cust1 local-sids
BGP VRF vrf_cust1, state: Active
BGP Route Distinguisher: 10.1.1.1:0
VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013   RD version: 37

```

```

BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Local Sid          Alloc mode  Locator
Route Distinguisher: 10.1.1.1:0 (default for vrf vrf_cust1)
*> 10.1.1.0/24      fc00:0:1:40::          per-vrf      locator1
*> 2.2.2.0/24       fc00:8:1:40::          per-vrf      locator2
*> 3.3.3.0/24       fc00:9:1:40::          per-vrf      locator4
*> 4.4.4.0/24       fc00:9:1:41::          per-ce       locator4
*> 10.1.1.5/32      NO SRv6 Sid            -            -
*> 3.3.3.3/32       NO SRv6 Sid            -            -
*>i8.8.8.8/32       NO SRv6 Sid            -            -


Node1# show bgp vrf vrf_cust1 received-sids
BGP VRF vrf_cust1, state: Active
BGP Route Distinguisher: 10.1.1.1:0
VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013  RD version: 37
BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Received Sid
Route Distinguisher: 10.1.1.1:0 (default for vrf vrf_cust1)
*> 10.1.1.0/24      11.1.1.2          NO SRv6 Sid
*> 2.2.2.0/24       11.1.1.2          NO SRv6 Sid
*> 3.3.3.0/24       11.1.1.2          NO SRv6 Sid
*> 4.4.4.0/24       11.1.1.2          NO SRv6 Sid
*> 10.1.1.5/32      11.1.1.2          NO SRv6 Sid
*> 3.3.3.3/32       13.2.2.2          NO SRv6 Sid
*>i8.8.8.8/32       10.1.1.2          fc00:0:2:42::

```

Verification

The following figure shows a VPNv4 scenario. The sequence of commands included correspond to router Node1 acting as Ingress PE, and routers Node4 and Node5 acting as Egress PEs.



The following example shows how to verify the SRv6 based L3VPN configuration using the **show segment-routing srv6 sid** command.

In this example, we can observe the uDT4 SIDs associated with the IPv4 L3VPN; where uDT4 behavior represents Endpoint with decapsulation and IPv4 table lookup.

```
Node1# show segment-routing srv6 sid
```

```
*** Locator: 'Node1-locator' ***
```

SID	State	RW	Behavior	Context	Owner
cafe:0:1::	InUse	Y	uN (PSP/USD)	'default':1	sidmgr
cafe:0:1:e000::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/0, Link-Local]:0	isis-1
cafe:0:1:e001::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/1, Link-Local]:0	isis-1
cafe:0:1:e002::	InUse	Y	uDT4	'vrf_cust1'	bgp-100
cafe:0:1:e003::	InUse	Y	uDT4	'vrf_cust2'	bgp-100
cafe:0:1:e004::	InUse	Y	uDT4	'vrf_cust3'	bgp-100
cafe:0:1:e005::	InUse	Y	uDT4	'vrf_cust4'	bgp-100
cafe:0:1:e006::	InUse	Y	uDT4	'vrf_cust5'	bgp-100

The following example shows how to verify the SRv6 based L3VPN configuration using the **show segment-routing srv6SID-prefixdetail** command.

```
Node1# show segment-routing srv6 sid cafe:0:1:e002:: detail
Tue Feb 9 17:50:40.621 UTC
```

```
*** Locator: 'Node1-locator' ***
```

SID	State	RW	Behavior	Context	Owner
cafe:0:1:e002::	InUse	Y	uDT4	'vrf_cust1'	bgp-100
SID Function: 0xe002 SID context: { table-id=0xe0000011 ('vrf_cust1':IPv4/Unicast) } Locator: 'Node1-locator' Allocation type: Dynamic Created: Feb 9 17:41:07.475 (00:09:33 ago)					

The following example shows how to verify the SRv6 based L3VPN configuration using the **show bgp vpnv4 unicast** commands on Egress PE.

```
Node1# show bgp vpnv4 unicast summary
```

```
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 36
BGP NSR Initial initsync version 16 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	36	36	36	36	36	0

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
cafe:0:4::4	0	100	47	48	36	0	0	00:40:05	5
cafe:0:5::5	0	100	47	47	36	0	0	00:39:56	5

```
Node1# show bgp vpnv4 unicast rd 100:1
```

```
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 36
BGP NSR Initial initsync version 16 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network      Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf vrf_cust1)
*> 12.1.1.1/32   0.0.0.0                        0         32768 ?
*>i12.4.4.4/32   cafe:0:4::4                     0        100    0 ?
*>i12.5.5.5/32   cafe:0:5::5                     0        100    0 ?
```

Processed 3 prefixes, 3 paths

```

Node1# show bgp vpnv4 unicast rd 100:1 12.4.4.4/32

BGP routing table entry for 12.4.4.4/32, Route Distinguisher: 100:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          22        22
Last Modified: Feb 23 22:57:56.756 for 00:40:08
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local, (received & used)
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
      Received Label 0xe00400
      Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 22
      Extended community: RT:1:1 RT:100:1
      PSID-Type:L3, SubTLV Count:1
      SubTLV:
        T:1(Sid information), Sid:cafe:0:4::, Behavior:63, SS-TLV Count:1
        SubSubTLV:
          T:1(Sid structure):
            Source AFI: VPNv4 Unicast, Source VRF: vrf_cust1, Source Route Distinguisher: 100:1

```

The following examples show how to verify the BGP prefix information for VRF instances using the **show bgp vrf** commands:

```

Node1# show bgp vrf vrf_cust1 ipv4 unicast

BGP VRF vrf_cust1, state: Active
BGP Route Distinguisher: 100:1
VRF ID: 0x60000002
BGP router identifier 10.1.1.1, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000011  RD version: 32
BGP main routing table version 36
BGP NSR Initial initsync version 16 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf vrf_cust1)
*> 12.1.1.1/32      0.0.0.0              0         32768 ?
*>i12.4.4.4/32      cafe:0:4::4          0         100    0 ?
*>i12.5.5.5/32      cafe:0:5::5          0         100    0 ?

Processed 3 prefixes, 3 paths

```

```

Node1# show bgp vrf vrf_cust1 ipv4 unicast 12.4.4.4/32
Tue Feb 23 23:39:57.499 UTC
BGP routing table entry for 12.4.4.4/32, Route Distinguisher: 100:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          22        22
Last Modified: Feb 23 22:57:56.756 for 00:42:01
Paths: (1 available, best #1)

```

```

Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
Local, (received & used)
  cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
    Received Label 0xe00400
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
    Received Path ID 0, Local Path ID 1, version 22
    Extended community: RT:1:1 RT:100:1
    PSID-Type:L3, SubTLV Count:1
    SubTLV:
      T:1(Sid information), Sid:cafe:0:4::, Behavior:63, SS-TLV Count:1
    SubSubTLV:
      T:1(Sid structure):
        Source AFI: VPNv4 Unicast, Source VRF: vrf_cust1, Source Route Distinguisher: 100:1

```

The following example shows how to verify the SRv6 based L3VPN configuration using the **show route vrf** commands.

Node1# **show route vrf vrf_cust1**

```

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
A - access/subscriber, a - Application route
M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup path

```

Gateway of last resort is not set

```

L   12.1.1.1/32 is directly connected, 00:44:43, Loopback100
B   12.4.4.4/32 [200/0] via cafe:0:4::4 (nexthop in vrf default), 00:42:45
B   12.5.5.5/32 [200/0] via cafe:0:5::5 (nexthop in vrf default), 00:42:45

```

Node1# **show route vrf vrf_cust1 12.4.4.4/32**

```

Routing entry for 12.4.4.4/32
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Feb 23 22:57:56.746 for 00:43:12
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
      Route metric is 0
  No advertising protos.

```

Node1# **show route vrf vrf_cust1 12.4.4.4/32 detail**

```

Routing entry for 12.4.4.4/32
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Feb 23 22:57:56.746 for 00:43:37
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
      Route metric is 0
      Label: None
      Tunnel ID: None
      Binding Label: None

```

```

Extended communities count: 0
Source RD attributes: 0x0000:100:1
NHID:0x0 (Ref:0)
SRv6 Headend: H.Encaps.Red [f3216], SID-list {cafe:0:4:e004::}
Route version is 0x1 (1)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_REMOTE
Download Priority 3, Download Version 3
No advertising protos.

```

The following example shows how to verify the SRv6 based L3VPN configuration using the **show cef vrf** commands.

Node1# **show cef vrf vrf_cust1**

Prefix	Next Hop	Interface
0.0.0.0/0	drop	default handler
0.0.0.0/32	broadcast	
12.1.1.1/32	receive	Loopback100
12.4.4.4/32	cafe:0:4::/128	<recursive>
12.5.5.5/32	cafe:0:5::/128	<recursive>
224.0.0.0/4	0.0.0.0/32	
224.0.0.0/24	receive	
255.255.255.255/32	broadcast	

Node1# **show cef vrf vrf_cust1 12.4.4.4/32**

```

12.4.4.4/32, version 3, SRv6 Headend, internal 0x5000001 0x30 (ptr 0x78b9a61c) [1], 0x0
(0x0), 0x0 (0x88873720)
Updated Feb 23 22:57:56.749
Prefix Len 32, traffic index 0, precedence n/a, priority 3
via cafe:0:4::/128, 3 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x78e2da14 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop cafe:0:4::/128 via cafe:0:4::/48
SRv6 H.Encaps.Red SID-list {cafe:0:4:e004::}

```

Node1# **show cef vrf vrf_cust1 12.4.4.4/32 detail**

```

12.4.4.4/32, version 3, SRv6 Headend, internal 0x5000001 0x30 (ptr 0x78b9a61c) [1], 0x0
(0x0), 0x0 (0x88873720)
Updated Feb 23 22:57:56.749
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0x88a740a8) reference count 5, flags 0x2010, source rib (7), 0 backups
[1 type 3 flags 0x48441 (0x789cbcc8) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Feb 23 22:57:56.749
LDI Update time Feb 23 22:57:56.754

Level 1 - Load distribution: 0
[0] via cafe:0:4::/128, recursive

via cafe:0:4::/128, 3 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x78e2da14 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop cafe:0:4::/128 via cafe:0:4::/48

```



```
SRv6 H.Encaps.Red SID-list {cafe:0:4:e004::}
```

```
Load distribution: 0 1 (refcount 1)
```

```
Hash  OK  Interface      Address
0     Y   HundredGigE0/0/0/1  remote
1     Y   HundredGigE0/0/0/0  remote
```

SRv6 Services: IPv6 L3VPN

Table 6: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Services: IPv6 L3VPN	Release 7.3.1	With this feature, the egress PE can signal an SRv6 Service SID with the BGP overlay service route. The ingress PE encapsulates the IPv4/IPv6 payload in an outer IPv6 header where the destination address is the SRv6 Service SID provided by the egress PE. BGP messages between PEs carry SRv6 Service SIDs as a means to interconnect PEs and form VPNs.

This feature provides IPv6 L3VPNs (VPNv6) over an SRv6 network.

Usage Guidelines and Limitations

- SRv6 locator can be assigned globally, for all VRFs, or for an individual VRF.
SRv6 locator can be assigned globally, for all VRFs, for an individual VRF, or per-prefix.
- Per-VRF allocation mode is supported (uDT6 behavior)
- Dual-Stack L3VPN Services (IPv4, IPv6) are supported
- Equal-Cost Multi-path (ECMP) and Unequal Cost Multipath (UCMP) are supported.
- eBGP, OSPF, Static are supported as PE-CE protocol.
BGP (iBGP, eBGP), OSPF, Static are supported as PE-CE protocol.
- BGP route leaking between BGP Global and L3VPN is supported.
- Per-CE allocation mode is not supported (uDX6 behavior)
- iBGP is not supported as PE-CE protocol
- BGP route leaking is not supported

Configuring SRv6-based IPv6 L3VPN

To enable SRv6-based L3VPN, you need to enable SRv6 under BGP, specify the locator, and configure the SID allocation mode. The assignment of the locator can be done in different places under the **router bgp** configuration. See [SRv6 Locator Inheritance Rules](#), on page 43.

Use Case 1: Assigning SRv6 Locator Globally

This example shows how to configure the SRv6 locator name under BGP Global:

```
Node1(config)# router bgp 100
Node1(config-bgp)# segment-routing srv6
Node1(config-bgp-gbl-srv6)# locator Node1-locator
Node1(config-bgp-gbl-srv6)# exit
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# commit
```

Running Configuration

```
router bgp 100
  segment-routing srv6
    locator Node1-locator
  !
  address-family vpnv6 unicast
  !
  neighbor 3001::12:1:1:4
    remote-as 100
    address-family vpnv6 unicast
  !
  !
  vrf vrf_cust6
    rd 100:6
    address-family ipv6 unicast
  !
  !
  !
end
```

Use Case 2: Assigning SRv6 Locator for All VRFs

To configure the SRv6 locator for all VRFs under VPNv6 Address Family and specify the allocation mode, use the following commands:

- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6:** Enable SRv6
- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6 alloc mode {*per-vrf*}:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT6 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6 locator *WORD*:** Specify the locator

This example shows how to configure the SRv6 locator for all VRFs under VPNv6 Address Family, with per-VRF label allocation mode:

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# vrf all
Node1(config-bgp-af-vrfall)# segment-routing srv6
Node1(config-bgp-af-vrfall-srv6)# locator Node1-locator
Node1(config-bgp-af-vrfall-srv6)# alloc mode per-vrf
Node1(config-bgp-af-vrfall-srv6)# exit
Node1(config-bgp-af-vrfall)# exit
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# commit

```

Running Configuration

```

router bgp 100
 address-family vpnv6 unicast
   vrf all
     segment-routing srv6
       locator Node1-locator
       alloc mode per-vrf
   !
!
!
neighbor 3001::12:1:1:4
 remote-as 100
 address-family vpnv6 unicast
!
!
vrf vrf_cust6
 rd 100:6
 address-family ipv6 unicast
!
!
!
end

```

Use Case 3: Assigning SRv6 Locator for a specific VRF

To configure the SRv6 locator for a specific VRF under IPv6 Address Family and specify the allocation mode, use the following commands:

- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6:** Enable SRv6
- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 alloc mode { per-vrf }:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT6 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 locator WORD:** Specify the locator

This example shows how to configure the SRv6 locator for an individual VRF, with per-VRF label allocation mode:

```
Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6
Node1(config-bgp-vrf-af-srv6)# locator Node1-locator
Node1(config-bgp-vrf-af-srv6)# alloc mode per-vrf
Node1(config-bgp-vrf-af-srv6)# commit
```

Running Configuration

```
router bgp 100
 address-family vpnv6 unicast
 !
 neighbor 3001::12:1:1:4
  remote-as 100
  address-family vpnv6 unicast
  !
 !
vrf vrf_cust6
 rd 100:6
 address-family ipv6 unicast
  segment-routing srv6
  locator Node1-locator
  alloc mode per-vrf
 !
 !
 !
end
```

Use Case 4: Assigning SRv6 Locator for a Specific Prefix

This use case provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes. The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.

To assign an SRv6 locator for a specific prefix, configure a route policy to specify the SID allocation mode based on match criteria. Examples of match criteria are destination-based match or community-based match.

- Supported SID allocation modes are per-VRF and per-CE.
- For per-VRF allocation mode, you can also specify the SRv6 locator.
 - If an SRv6 locator is specified in the route policy, BGP will use that to allocate per-VRF SID. If the specified locator is invalid, the SID will not be allocated.
 - If an SRv6 locator is not specified in the route policy, the default locator is used to allocate the SID. If the default locator is not configured in BGP, then the SID will not be allocated.
- Per-CE allocation mode always uses the default locator configured under BGP to allocate the SID.

For more information on configuring routing policies, refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following example shows a route policy specifying the SID allocation mode with destination-based match:

```

Node1(config)# route-policy set_per_prefix_locator_rpl
Node1(config-rpl)# if destination in (3001::1:1:1:1/128) then
Node1(config-rpl-if)# set srv6-alloc-mode per-vrf locator locator1
Node1(config-rpl-if)# elseif destination in (3001::2:2:2:2/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf locator locator2
Node1(config-rpl-elseif)# elseif destination in (3001::3:3:3:3/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf
Node1(config-rpl-elseif)# elseif destination in (3001::4:4:4:4/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-ce
Node1(config-rpl-elseif)# else
Node1(config-rpl-else)# drop
Node1(config-rpl-else)# endif
Node1(config-rpl)# end-policy

```

To specify per-prefix allocation mode for a specific VRF under IPv6 Address Family, use the following command:

- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 alloc mode route-policy policy_name**

This example shows how to specify per-prefix allocation mode for a specific VRF (vrf_cust6) under the IPv6 address family:

```

Node1(config)# router bgp 100
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6
Node1(config-bgp-vrf-af-srv6)# alloc mode route-policy set_per_prefix_locator_rpl

```

Running Configuration

```

route-policy set_per_prefix_locator_rpl
  if destination in (3001::1:1:1:1/128) then
    set srv6-alloc-mode per-vrf locator locator1
  elseif destination in (3001::2:2:2:2/128) then
    set srv6-alloc-mode per-vrf locator locator2
  elseif destination in (3001::3:3:3:3/128) then
    set srv6-alloc-mode per-vrf
  elseif destination in (3001::4:4:4:4/128) then
    set srv6-alloc-mode per-ce
  else
    drop
  endif
end-policy
!
router bgp 100
  vrf vrf_cust6
    address-family ipv6 unicast
      segment-routing srv6
        alloc mode route-policy set_per_prefix_locator_rpl
    !
  !
!
!
!

```

Verify that the local and received SIDs have been correctly allocated under VPNv6 and specific VRF (vrf_cust6):

Node1# **show bgp vpnv6 unicast local-sids**

```
BGP router identifier 10.1.1.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Local Sid	Alloc mode	Locator
Route Distinguisher: 8:8			
*>i3008::8:8:8:8/128	NO SRv6 Sid	-	-
* i	NO SRv6 Sid	-	-
Route Distinguisher: 100:6 (default for vrf vrf_cust6)			
*> 3001::1:1:1:1/128	fc00:0:1:40::	per-vrf	locator1
*> 3001::2:2:2:2/128	fc00:8:1:40::	per-vrf	locator2
*> 3001::3:3:3:3/128	fc00:9:1:40::	per-vrf	locator4
*> 3001::4:4:4:4/128	fc00:9:1:41::	per-ce	locator4
*> 3001::5:5:5:5/128	NO SRv6 Sid	-	-
*> 3001::12:1:1:5/128	NO SRv6 Sid	-	-
*>i3008::8:8:8:8/128	NO SRv6 Sid	-	-

Node1# **show bgp vpnv6 unicast received-sids**

```
BGP router identifier 10.1.1.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Received Sid
Route Distinguisher: 8:8		
*>i3008::8:8:8:8/128	10.1.1.2	fc00:0:2:42::
* i	2400:2020:42:2fff::1	fc00:0:2:42::
Route Distinguisher: 100:6 (default for vrf vrf_cust6)		
*> 3001::1:1:1:1/128	11.1.1.2	NO SRv6 Sid
*> 3001::2:2:2:2/128	11.1.1.2	NO SRv6 Sid
*> 3001::3:3:3:3/128	11.1.1.2	NO SRv6 Sid
*> 3001::4:4:4:4/128	11.1.1.2	NO SRv6 Sid
*> 3001::5:5:5:5/128	11.1.1.2	NO SRv6 Sid
*> 3001::12:1:1:5/128	13.2.2.2	NO SRv6 Sid
*>i3008::8:8:8:8/128	10.1.1.2	fc00:0:2:42::

Node1# **show bgp vrf vrf_cust6 local-sids**

```
BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 10.1.1.1:0
```

```

VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013   RD version: 37
BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Local Sid          Alloc mode   Locator
Route Distinguisher: 8:8
*>i3008::8:8:8:8/128  NO SRv6 Sid          -             -
* i                  NO SRv6 Sid          -             -
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1:1/128  fc00:0:1:40::        per-vrf        locator1
*> 3001::2:2:2:2/128  fc00:8:1:40::        per-vrf        locator2
*> 3001::3:3:3:3/128  fc00:9:1:40::        per-vrf        locator4
*> 3001::4:4:4:4/128  fc00:9:1:41::        per-ce        locator4
*> 3001::5:5:5:5/128  NO SRv6 Sid          -             -
*> 3001::12:1:1:5/128 NO SRv6 Sid          -             -
*>i3008::8:8:8:8/128  NO SRv6 Sid          -             -


Node1# show bgp vrf vrf_cust6 received-sids
BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 100:6
VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013   RD version: 37
BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Received Sid
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1:1/128  11.1.1.2          NO SRv6 Sid
*> 3001::2:2:2:2/128  11.1.1.2          NO SRv6 Sid
*> 3001::3:3:3:3/128  11.1.1.2          NO SRv6 Sid
*> 3001::4:4:4:4/128  11.1.1.2          NO SRv6 Sid
*> 3001::5:5:5:5/128  11.1.1.2          NO SRv6 Sid
*> 3001::12:1:1:5/128 13.2.2.2          NO SRv6 Sid
*>i3008::8:8:8:8/128  10.1.1.2          fc00:0:2:42::

```

Verification

The following figure shows a VPNv6 scenario. The sequence of commands included correspond to router Node1 acting as Ingress PE, and routers Node4 and Node5 acting as Egress PEs.



The following examples shows how to verify the SRv6 based L3VPN configurations for an Individual VRF with per VRF label allocation mode.

In this example, we can observe the uDT6 SID associated with the IPv6 L3VPN, where uDT6 behavior represents Endpoint with decapsulation and IPv6 table lookup.

```
Node1# show segment-routing srv6 sid
Fri Jan 29 19:31:53.293 UTC
```

```
*** Locator: 'Node1-locator' ***
```

SID	State	RW	Behavior	Context	Owner
cafe:0:1::	InUse	Y	uN (PSP/USD)	'default':1	sidmgr
cafe:0:1:e000::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/0, Link-Local]:0	isis-1
cafe:0:1:e001::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/1, Link-Local]:0	isis-1
cafe:0:1:e002::	InUse	Y	uDT4	'vrf_cust1'	bgp-100
cafe:0:1:e003::	InUse	Y	uDT4	'vrf_cust2'	bgp-100
cafe:0:1:e004::	InUse	Y	uDT4	'vrf_cust3'	bgp-100
cafe:0:1:e005::	InUse	Y	uDT4	'vrf_cust4'	bgp-100
cafe:0:1:e006::	InUse	Y	uDT4	'vrf_cust5'	bgp-100
cafe:0:1:e007::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/0, Link-Local]:0:P	isis-1
cafe:0:1:e008::	InUse	Y	uA (PSP/USD)	[Hu0/0/0/1, Link-Local]:0:P	isis-1
cafe:0:1:e009::	InUse	Y	uDT6	'default'	bgp-100
cafe:0:1:e00a::	InUse	Y	uDT6	'vrf_cust6'	bgp-100


```
InUse Y
```

The following examples show how to verify the SRv6 based L3VPN configuration using the **show bgp vpnv6 unicast** commands on the Ingress PE.

```
Node1# show bgp vpnv6 unicast summary
```

```
Fri Jan 29 19:33:01.177 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 6
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	6	6	6	6	6	0

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
cafe:0:4::4	0	100	122	123	6	0	0	00:20:05	1
cafe:0:5::5	0	100	111	111	0	0	0	00:49:46	1

```
Node1# show bgp vpnv6 unicast rd 100:6
```

```
Fri Jan 29 19:41:01.334 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 8
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::12:1:1:1/128 :: 0 32768 ?
*>i3001::12:1:1:4/128 cafe:0:4::4 0 100 0 ?
*>i3001::12:1:1:5/128 cafe:0:5::5 0 100 0 ?
```

Processed 3 prefixes, 3 paths

```
Node1# show bgp vpnv6 unicast rd 100:6 3001::12:1:1:4/128
```

```
Fri Jan 29 19:41:42.008 UTC
BGP routing table entry for 3001::12:1:1:4/128, Route Distinguisher: 100:6
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          6         6
Last Modified: Jan 29 19:29:35.858 for 00:12:06
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local, (received & used)
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
```

```

Received Label 0xe00a00
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
Received Path ID 0, Local Path ID 1, version 6
Extended community: RT:100:6
PSID-Type:L3, SubTLV Count:1
SubTLV:
  T:1(Sid information), Sid:cafe:0:4::, Behavior:62, SS-TLV Count:1
SubSubTLV:
  T:1(Sid structure):
    Source AFI: VPNv6 Unicast, Source VRF: vrf_cust6, Source Route Distinguisher: 100:6

```

The following examples show how to verify the BGP prefix information for VRF instances:

```

Node1# show bgp vrf vrf_cust6 ipv6 unicast
Fri Jan 29 19:42:05.675 UTC
BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 100:6
VRF ID: 0x60000007
BGP router identifier 10.1.1.1, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0800016 RD version: 8
BGP main routing table version 8
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::12:1:1:1/128 ::                0          32768 ?
*>i3001::12:1:1:4/128 cafe:0:4::4        0          100    0 ?
*>i3001::12:1:1:5/128 cafe:0:5::5        0          100    0 ?

Processed 3 prefixes, 3 paths

Node1# show bgp vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128

BGP routing table entry for 3001::12:1:1:4/128, Route Distinguisher: 100:6
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          17        17
Last Modified: Jan 15 16:50:44.032 for 01:48:21
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local, (received & used)
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
    Received Label 0xe00a00
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
    Received Path ID 0, Local Path ID 1, version 17
    Extended community: RT:100:6
    PSID-Type:L3, SubTLV Count:1
    SubTLV:
      T:1(Sid information), Sid:cafe:0:4::, Behavior:62, SS-TLV Count:1
    SubSubTLV:
      T:1(Sid structure):
        Source AFI: VPNv6 Unicast, Source VRF: vrf_cust6, Source Route Distinguisher: 100:6

```

The following examples show how to verify the current routes in the Routing Information Base (RIB):

```
Node1# show route vrf vrf_cust6 ipv6 unicast
```

```
Fri Jan 29 19:43:28.067 UTC
```

```
Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
A - access/subscriber, a - Application route
M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup path
```

```
Gateway of last resort is not set
```

```
L 3001::12:1:1:1/128 is directly connected,
   01:01:23, Loopback105
B 3001::12:1:1:4/128
   [200/0] via cafe:0:4::4 (nexthop in vrf default), 00:13:52
B 3001::12:1:1:5/128
   [200/0] via cafe:0:5::5 (nexthop in vrf default), 00:05:53
```

```
Node1# show route vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128
```

```
Fri Jan 29 19:43:55.645 UTC
```

```
Routing entry for 3001::12:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:29:35.696 for 00:14:20
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
      Route metric is 0
  No advertising protos.
```

```
Node1# show route vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128 detail
```

```
Fri Jan 29 19:44:17.914 UTC
```

```
Routing entry for 3001::12:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:29:35.696 for 00:14:42
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
      Route metric is 0
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Source RD attributes: 0x0000:100:6
      NHID:0x0(Ref:0)
      SRv6 Headend: H.Encaps.Red [f3216], SID-list {cafe:0:4:e00a::}
  Route version is 0x1 (1)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Flow-tag: Not Set
  Fwd-class: Not Set
  Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_REMOTE
  Download Priority 3, Download Version 3
  No advertising protos.
```

The following examples show how to verify the current IPv6 Cisco Express Forwarding (CEF) table:

```

Node1# show cef vrf vrf_cust6 ipv6
Fri Jan 29 19:44:56.888 UTC

::/0
  drop      default handler
3001::12:1:1:1/128
  receive   Loopback105
3001::12:1:1:4/128
  recursive cafe:0:4::/128
3001::12:1:1:5/128
  recursive cafe:0:5::/128
fe80::/10
  receive
ff02::/16
  receive
ff02::2/128
  receive
ff02::1:ff00:0/104
  receive
ff05::/16
  receive
ff12::/16
  receive

Node1# show cef vrf vrf_cust6 ipv6 3001::12:1:1:4/128
Fri Jan 29 19:45:23.607 UTC
3001::12:1:1:4/128, version 3, SRv6 Headend, internal 0x5000001 0x30 (ptr 0x78f2e0e0) [1],
  0x0 (0x0), 0x0 (0x888a3ac8)
Updated Jan 29 19:29:35.700
Prefix Len 128, traffic index 0, precedence n/a, priority 3
  via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
    next hop VRF - 'default', table - 0xe0800000
    next hop cafe:0:4::/128 via cafe:0:4::/48
    SRv6 H.Encaps.Red SID-list {cafe:0:4:e00a::}

Node1# show cef vrf vrf_cust6 ipv6 3001::12:1:1:4/128 detail
Fri Jan 29 19:45:55.847 UTC
3001::12:1:1:4/128, version 3, SRv6 Headend, internal 0x5000001 0x30 (ptr 0x78f2e0e0) [1],
  0x0 (0x0), 0x0 (0x888a3ac8)
Updated Jan 29 19:29:35.700
Prefix Len 128, traffic index 0, precedence n/a, priority 3
  gateway array (0x78afe238) reference count 1, flags 0x2010, source rib (7), 0 backups
    [1 type 3 flags 0x48441 (0x78ba9a60) ext 0x0 (0x0)]
  LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 1 Jan 29 19:29:35.699
  LDI Update time Jan 29 19:29:35.701

Level 1 - Load distribution: 0
[0] via cafe:0:4::/128, recursive

  via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
    next hop VRF - 'default', table - 0xe0800000
    next hop cafe:0:4::/128 via cafe:0:4::/48
    SRv6 H.Encaps.Red SID-list {cafe:0:4:e00a::}

Load distribution: 0 1 (refcount 1)

Hash  OK  Interface          Address
0      Y   HundredGigE0/0/0/0      remote

```

```
1      Y      HundredGigE0/0/0/1      remote
```

SRv6 Services: IPv4 BGP Global

This feature extends support of SRv6-based BGP services to include IPv4 global BGP by implementing uDT4 SRv6 functions at the PE node ([draft-ietf-bess-srv6-services](#)).

Usage Guidelines and Limitations

- SRv6 locator can be assigned globally or under IPv4 unicast address family
- Equal-Cost Multi-path (ECMP) and Unequal Cost Multipath (UCMP) are supported.
- BGP, OSPF, Static are supported as PE-CE protocol.
- BGP route leaking between BGP Global and L3VPN is supported.

BGP Global IPv4 Over SRv6 with Per-AFI SID Allocation Mode (uDT4)

To configure BGP global IPv4 over SRv6, use the following commands:

- **router bgp *as-number* address-family ipv4 unicast segment-routing srv6**: Enable SRv6
- **router bgp *as-number* address-family ipv4 unicast segment-routing srv6 alloc mode per-vrf**: Specify the SID behavior (allocation mode).
The **per-vrf** keyword specifies that the same label is be used for all the routes advertised from a unique VRF.
- **router bgp *as-number* address-family ipv4 unicast segment-routing srv6 alloc mode {per-vrf | route-policy *policy_name*}**: Specify the SID behavior (allocation mode).
 - **per-vrf**: Specifies that the same label is be used for all the routes advertised from a unique VRF.
 - **route-policy *policy_name***: Uses a route policy to determine the SID allocation mode and locator (if provided) for given prefix.
- **router bgp *as-number* address-family ipv4 unicast segment-routing srv6 locator *WORD***: Specify the locator
- **router bgp *as-number* {af-group *WORD* | neighbor-group *WORD* | neighbor *ipv6-addr*} address-family ipv4 unicast encapsulation-type srv6**: Specify the encapsulation type for SRv6.
 - Use **af-group *WORD*** to apply the SRv6 encapsulation type to the address family group for BGP neighbors.
 - Use **neighbor-group *WORD*** to apply the SRv6 encapsulation type to the neighbor group for BGP neighbors.
 - Use **neighbor *ipv6-addr*** to apply the SRv6 encapsulation type to the specific BGP neighbor.

Use Case 1: BGP Global IPv4 over SRv6 with Per-AFI SID Allocation

The following example shows how to configure BGP global IPv4 over SRv6 with per-AFI SID allocation.

```

Node1(config)# router bgp 1
Node1(config-bgp)# bgp router-id 10.1.0.1
Node1(config-bgp)# address-family ipv4 unicast
Node1(config-bgp-af)# segment-routing srv6
Node1(config-bgp-af-srv6)# locator Node1
Node1(config-bgp-af-srv6)# alloc mode per-vrf
Node1(config-bgp-af-srv6)# exit
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 60::2
Node1(config-bgp-nbr)# remote-as 1
Node1(config-bgp-nbr)# update-source Loopback1
Node1(config-bgp-nbr)# address-family ipv4 unicast
Node1(config-bgp-nbr-af)# encapsulation-type srv6
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# neighbor 52.52.52.1
Node1(config-bgp-nbr)# remote-as 3
Node1(config-bgp-nbr)# address-family ipv4 unicast
Node1(config-bgp-nbr-af)# route-policy passall in
Node1(config-bgp-nbr-af)# route-policy passall out
Node1(config-bgp-nbr-af)# commit

```

Running Configuration

```

router bgp 1
  bgp router-id 10.1.0.1
  address-family ipv4 unicast
    segment-routing srv6
      locator Node1
      alloc mode per-vrf
  !
  !
  neighbor 60::2
    remote-as 1
    update-source Loopback1
    address-family ipv4 unicast
      encapsulation-type srv6
  !
  !
  neighbor 52.52.52.1
    remote-as 3
    address-family ipv4 unicast
      route-policy passall in
      route-policy passall out
  !
  !
  !

```

Use Case 2: BGP Global IPv4 over SRv6 with Per-Prefix SID Allocation

This use case provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes. The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.

To assign an SRv6 locator for a specific prefix, configure a route policy to specify the SID allocation mode based on match criteria. Examples of match criteria are destination-based match or community-based match.

- Supported SID allocation modes are per-VRF and per-CE.
- For per-VRF allocation mode, you can also specify the SRv6 locator.

- If an SRv6 locator is specified in the route policy, BGP will use that to allocate per-VRF SID. If the specified locator is invalid, the SID will not be allocated.
- If an SRv6 locator is not specified in the route policy, the default locator is used to allocate the SID. If the default locator is not configured in BGP, then the SID will not be allocated.
- Per-CE allocation mode always uses the default locator to allocate the SID.

For more information on configuring routing policies, refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following example shows a route policy specifying the SID allocation mode with destination-based match:

```

Node1(config)# route-policy set_per_prefix_locator_rpl
Node1(config-rpl)# if destination in (10.1.1.0/24) then
Node1(config-rpl-if)# set srv6-alloc-mode per-vrf locator locator1
Node1(config-rpl-if)# elseif destination in (2.2.2.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf locator locator2
Node1(config-rpl-elseif)# elseif destination in (3.3.3.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf
Node1(config-rpl-elseif)# elseif destination in (4.4.4.0/24) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-ce
Node1(config-rpl-elseif)# else
Node1(config-rpl-else)# drop
Node1(config-rpl-else)# endif
Node1(config-rpl)# end-policy
Node1(config)#

```

The following example shows how to configure BGP global IPv4 over SRv6 with a route policy to determine the SID allocation mode for given prefix.

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family ipv4 unicast
Node1(config-bgp-af)# segment-routing srv6
Node1(config-bgp-af-srv6)# alloc mode route-policy set_per_prefix_locator_rpl

```

Running Configuration

```

route-policy set_per_prefix_locator_rpl
  if destination in (10.1.1.0/24) then
    set srv6-alloc-mode per-vrf locator locator1
  elseif destination in (2.2.2.0/24) then
    set srv6-alloc-mode per-vrf locator locator2
  elseif destination in (3.3.3.0/24) then
    set srv6-alloc-mode per-vrf
  elseif destination in (4.4.4.0/24) then
    set srv6-alloc-mode per-ce
  else
    drop
  endif
end-policy
!
router bgp 100
  address-family ipv4 unicast
    segment-routing srv6
      alloc mode route-policy set_per_prefix_locator_rpl
  !
!
!

```

Verify that the local and received SIDs have been correctly allocated under BGP IPv4 address family:

```

Node1# show bgp ipv4 unicast local-sids
...
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Local Sid                               Alloc mode   Locator
*> 10.1.1.0/24       fc00:8:1:41::                               per-vrf      locator2
*> 2.2.2.0/24        fc00:0:1:41::                               per-vrf      locator1
*> 3.3.3.0/24        fc00:9:1:42::                               per-vrf      locator4
*> 4.4.4.0/24        fc00:9:1:43::                               per-ce       locator4
*> 10.1.1.5/32       NO SRv6 Sid                                 -            -
* i8.8.8.8/32       NO SRv6 Sid                                 -            -

Node1# show bgp ipv4 unicast received-sids
...
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop                               Received Sid
*> 10.1.1.0/24       66.2.2.2                               NO SRv6 Sid
*> 2.2.2.0/24       66.2.2.2                               NO SRv6 Sid
*> 3.3.3.0/24       66.2.2.2                               NO SRv6 Sid
*> 4.4.4.0/24       66.2.2.2                               NO SRv6 Sid
*> 10.1.1.5/32      66.2.2.2                               NO SRv6 Sid
* i8.8.8.8/32      77.1.1.2                               fc00:0:2:41::

```

SRv6 Services: IPv6 BGP Global

Table 7: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Services: BGP Global IPv6	Release 7.3.1	With this feature, the egress PE can signal an SRv6 Service SID with the BGP global route. The ingress PE encapsulates the IPv4/IPv6 payload in an outer IPv6 header where the destination address is the SRv6 Service SID provided by the egress PE. BGP messages between PEs carry SRv6 Service SIDs as a means to interconnect PEs.

This feature extends support of SRv6-based BGP services to include IPv6 global BGP by implementing uDT6 SRv6 functions at the PE node ([draft-ietf-bess-srv6-services](#)).

Usage Guidelines and Limitations

- SRv6 locator can be assigned globally or under IPv6 unicast address family
- Equal-Cost Multi-path (ECMP) and Unequal Cost Multipath (UCMP) are supported.
- BGP, OSPF, Static are supported as PE-CE protocol.

- BGP route leaking between BGP Global and L3VPN is supported.

BGP Global IPv6 Over SRv6 with Per-AFI SID Allocation Mode (uDT6)

To configure BGP global IPv6 over SRv6, use the following commands:

- **router bgp *as-number* address-family ipv6 unicast segment-routing srv6**: Enable SRv6
- **router bgp *as-number* address-family ipv6 unicast segment-routing srv6 alloc mode per-vrf**: Specify the SID behavior (allocation mode).
The **per-vrf** keyword specifies that the same label is be used for all the routes advertised from a unique VRF.
- **router bgp *as-number* address-family ipv6 unicast segment-routing srv6 alloc mode {per-vrf | route-policy *policy_name*}**: Specify the SID behavior (allocation mode).
 - **per-vrf**: Specifies that the same label is be used for all the routes advertised from a unique VRF.
 - **route-policy *policy_name***: Uses a route policy to determine the SID allocation mode and locator (if provided) for given prefix.
- **router bgp *as-number* address-family ipv6 unicast segment-routing srv6 locator *WORD***: Specify the locator
- **router bgp *as-number* {af-group *WORD* | neighbor-group *WORD* | neighbor *ipv6-addr*} address-family ipv6 unicast encapsulation-type srv6**: Specify the encapsulation type for SRv6.
 - Use **af-group *WORD*** to apply the SRv6 encapsulation type to the address family group for BGP neighbors.
 - Use **neighbor-group *WORD*** to apply the SRv6 encapsulation type to the neighbor group for Border Gateway Protocol (BGP) neighbors.
 - Use **neighbor *ipv6-addr*** to apply the SRv6 encapsulation type to the specific BGP neighbor.

Use Case 1: BGP Global IPv6 over SRv6 with Per-AFI SID Allocation

The following example shows how to configure BGP global IPv6 over SRv6 with per-AFI SID allocation.

```

Node1(config)# router bgp 100
Node1(config-bgp)# bgp router-id 10.1.1.1
Node1(config-bgp)# segment-routing srv6
Node1(config-bgp-gbl-srv6)# locator Node1
Node1(config-bgp-gbl-srv6)# exit
Node1(config-bgp)# address-family ipv6 unicast
Node1(config-bgp-af)# segment-routing srv6
Node1(config-bgp-af-srv6)# locator Node1
Node1(config-bgp-af-srv6)# alloc mode per-vrf
Node1(config-bgp-af-srv6)# exit
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor cafe:0:4::4
Node1(config-bgp-nbr)# address-family ipv6 unicast
Node1(config-bgp-nbr-af)# encapsulation-type srv6
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# neighbor cafe:0:5::5
Node1(config-bgp-nbr)# address-family ipv6 unicast
Node1(config-bgp-nbr-af)# encapsulation-type srv6

```

```
Node1(config-bgp-nbr-af)# commit
```

Running Configuration

```
router bgp 100
  bgp router-id 10.1.1.1
  segment-routing srv6
    locator Node1
  !
  address-family ipv6 unicast
    segment-routing srv6
      locator Node1
      alloc mode per-vrf
  !
  !
  neighbor cafe:0:4::4
    address-family ipv6 unicast
      encapsulation-type srv6
  !
  !
  neighbor cafe:0:5::5
    address-family ipv6 unicast
      encapsulation-type srv6
```

Use Case 2: BGP Global IPv6 over SRv6 with Per-Prefix SID Allocation

This use case provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes. The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.

To assign an SRv6 locator for a specific prefix, configure a route policy to specify the SID allocation mode based on match criteria. Examples of match criteria are destination-based match or community-based match.

- Supported SID allocation modes are per-VRF and per-CE.
- For per-VRF allocation mode, you can also specify the SRv6 locator.
 - If an SRv6 locator is specified in the route policy, BGP will use that to allocate per-VRF SID. If the specified locator is invalid, the SID will not be allocated.
 - If an SRv6 locator is not specified in the route policy, the default locator is used to allocate the SID. If the default locator is not configured in BGP, then the SID will not be allocated.
- Per-CE allocation mode always uses the default locator to allocate the SID.

For more information on configuring routing policies, refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following example shows a route policy specifying the SID allocation mode with destination-based match:

```
Node1(config)# route-policy set_per_prefix_locator_rpl
Node1(config-rpl)# if destination in (3001::1:1:1:1/128) then
Node1(config-rpl-if)# set srv6-alloc-mode per-vrf locator locator1
Node1(config-rpl-if)# elseif destination in (3001::2:2:2:2/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf locator locator2
Node1(config-rpl-elseif)# elseif destination in (3001::3:3:3:3/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf
Node1(config-rpl-elseif)# elseif destination in (3001::4:4:4:4/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-ce
Node1(config-rpl-elseif)# else
```

```

Node1(config-rpl-else)# drop
Node1(config-rpl-else)# endif
Node1(config-rpl)# end-policy

```

The following example shows how to configure BGP global IPv6 over SRv6 with a route policy to determine the SID allocation mode for given prefix.

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family ipv6 unicast
Node1(config-bgp-af)# segment-routing srv6
Node1(config-bgp-af-srv6)# alloc mode route-policy set_per_prefix_locator_rpl

```

Running Configuration

```

route-policy set_per_prefix_locator_rpl
  if destination in (3001::1:1:1:1/128) then
    set srv6-alloc-mode per-vrf locator locator1
  elseif destination in (3001::2:2:2:2/128) then
    set srv6-alloc-mode per-vrf locator locator2
  elseif destination in (3001::3:3:3:3/128) then
    set srv6-alloc-mode per-vrf
  elseif destination in (3001::4:4:4:4/128) then
    set srv6-alloc-mode per-ce
  else
    drop
  endif
end-policy
!
router bgp 100
  address-family ipv6 unicast
    segment-routing srv6
    alloc mode route-policy set_per_prefix_locator_rpl
  !
!

```

Verify that the local and received SIDs have been correctly allocated under BGP IPv6 address family:

```

Node1# show bgp ipv6 unicast local-sids
...
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Local Sid	Alloc mode	Locator
*> 3001::1:1:1:1/128	fc00:8:1:41::	per-vrf	locator2
*> 3001::2:2:2:2/128	fc00:0:1:41::	per-vrf	locator1
*> 3001::3:3:3:3/128	fc00:9:1:42::	per-vrf	locator4
*> 3001::4:4:4:4/128	fc00:9:1:43::	per-ce	locator4
*> 3001::5:5:5:5/128	NO SRv6 Sid	-	-
* i3008::8:8:8:8/128	NO SRv6 Sid	-	-

```

Node1# show bgp ipv6 unicast received-sids
...
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```


Network	Next Hop	Received Sid
*> 3001::1:1:1:1/128	66.2.2.2	NO SRv6 Sid
*> 3001::2:2:2:2/128	66.2.2.2	NO SRv6 Sid
*> 3001::3:3:3:3/128	66.2.2.2	NO SRv6 Sid
*> 3001::4:4:4:4/128	66.2.2.2	NO SRv6 Sid
*> 3001::5:5:5:5/128	66.2.2.2	NO SRv6 Sid

* i3008::8:8:8:8/128 77.1.1.2

fc00:0:2:41::

Verification

The following figure shows a IPv6 BGP global scenario. The sequence of commands included correspond to router Node1 acting as Ingress PE, and routers Node4 and Node5 acting as Egress PEs.



The following examples show how to verify the BGP global IPv6 configuration using the **show bgp ipv6 unicast** commands.

```
Node1# show bgp ipv6 unicast summary
Fri Jan 29 19:48:23.255 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0800000 RD version: 4
BGP main routing table version 4
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	4	4	4	4	4	0

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
cafe:0:4::4	0	100	137	138	4	0	0	00:35:27	1
cafe:0:5::5	0	100	138	137	4	0	0	00:10:54	1

```
Node1# show bgp ipv6 unicast
Fri Jan 29 19:49:05.688 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
```

```

BGP table state: Active
Table ID: 0xe0800000   RD version: 4
BGP main routing table version 4
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network             Next Hop              Metric LocPrf Weight Path
*> 3001::13:1:1:1/128 ::                        0           32768 i
*>i3001::13:1:1:4/128 cafe:0:4::4                0           100    0 i
*>i3001::13:1:1:5/128 cafe:0:5::5                0           100    0 i

Processed 3 prefixes, 3 paths

Node1# show bgp ipv6 unicast 3001::13:1:1:4/128
Fri Jan 29 19:49:22.067 UTC
BGP routing table entry for 3001::13:1:1:4/128
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          3         3
Last Modified: Jan 29 19:14:13.858 for 00:35:08
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 1, version 3
      PSID-Type:L3, SubTLV Count:1
      SubTLV:
        T:1(Sid information), Sid:cafe:0:4:e009::, Behavior:62, SS-TLV Count:1
        SubSubTLV:
          T:1(Sid structure):

```

The following examples show how to verify the current routes in the Routing Information Base (RIB):

```

Node1# show route ipv6 3001::13:1:1:4/128
Fri Jan 29 19:53:26.839 UTC

Routing entry for 3001::13:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:14:13.397 for 00:35:28
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Route metric is 0
  No advertising protos.

Node1# show route ipv6 3001::13:1:1:4/128 detail
Fri Jan 29 19:50:08.601 UTC

Routing entry for 3001::13:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:14:13.397 for 00:35:55
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Route metric is 0
      Label: None
      Tunnel ID: None
      Binding Label: None

```

```

Extended communities count: 0
NHID:0x0(Ref:0)
SRv6 Headend: H.Encaps.Red [f3216], SID-list {cafe:0:4:e009::}
Route version is 0x1 (1)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 4, Download Version 106
No advertising protos.

```

The following examples show how to verify the current IPv6 Cisco Express Forwarding (CEF) table:

```

Node1# show cef ipv6 3001::13:1:1:4/128
Fri Jan 29 19:50:29.149 UTC
3001::13:1:1:4/128, version 106, SRv6 Headend, internal 0x5000001 0x40 (ptr 0x78 cd3944)
[1], 0x0 (0x0), 0x0 (0x888a3a80)
Updated Jan 29 19:14:13.401
Prefix Len 128, traffic index 0, precedence n/a, priority 4
via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
next hop cafe:0:4::/128 via cafe:0:4::/48
SRv6 H.Encaps.Red SID-list {cafe:0:4:e009::}

Node1# show cef ipv6 3001::13:1:1:4/128 detail
Fri Jan 29 19:51:00.920 UTC
3001::13:1:1:4/128, version 106, SRv6 Headend, internal 0x5000001 0x40 (ptr 0x78cd3944)
[1], 0x0 (0x0), 0x0 (0x888a3a80)
Updated Jan 29 19:14:13.401
Prefix Len 128, traffic index 0, precedence n/a, priority 4
gateway array (0x78afef150) reference count 1, flags 0x2010, source rib (7), 0 backups
[1 type 3 flags 0x48441 (0x78ba99e8) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Jan 29 19:14:13.401
LDI Update time Jan 29 19:14:13.401

Level 1 - Load distribution: 0
[0] via cafe:0:4::/128, recursive

via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
next hop cafe:0:4::/128 via cafe:0:4::/48
SRv6 H.Encaps.Red SID-list {cafe:0:4:e009::}

Load distribution: 0 1 (refcount 1)

Hash OK Interface Address
0 Y HundredGigE0/0/0/0 remote
1 Y HundredGigE0/0/0/1 remote

```

SRv6 Services: IPv4 L3VPN Active-Standby Redundancy using Port-Active Mode

The Segment Routing IPv6 (SRv6) Services: IPv4 L3VPN Active-Standby Redundancy using Port-Active Mode feature provides all-active per-port load balancing for multihoming. The forwarding of traffic is determined based on a specific interface rather than per-flow across multiple Provider Edge routers. This

feature enables efficient load-balancing and provides faster convergence. In an active-standby scenario, the active PE router is detected using designated forwarder (DF) election by modulo calculation and the interface of the standby PE router brought down. For Modulo calculation, byte 10 of the Ethernet Segment Identifier (ESI) is used.

Usage Guidelines and Restrictions

- This feature can only be configured for bundle interfaces.
- When an EVPN Ethernet Segment (ES) is configured with port-active load-balancing mode, you cannot configure ACs of that bundle on bridge-domains with a configured EVPN instance (EVI). EVPN Layer 2 bridging service is not compatible with port-active load-balancing.

SRv6 Services for L3VPN Active-Standby Redundancy using Port-Active Mode: Operation

Under port-active operational mode, EVPN Ethernet Segment (ES) routes are exchanged across BGP for the routers servicing the multihomed ES. Each PE router then builds an ordered list of the IP addresses of all PEs connected to the ES, including itself, and assigns itself an ordinal for its position in the list. The ordinals are used with the modulo calculation to determine which PE will be the Designated Forwarder (DF) for a given ES. All non-DF PEs will take the respective bundles out of service.

In the case of link or port failure, the active DF PE withdraws its ES route. This re-triggers DF election for all PEs that service the ES and a new PE is elected as DF.

Configure SRv6 Services L3VPN Active-Standby Redundancy using Port-Active Mode

This section describes how you can configure SRv6 services L3VPN active-standby redundancy using port-active mode under an Ethernet Segment (ES).

Configuration Example

```
/* Configure Ethernet Link Bundles */
Router# configure
Router(config)# interface Bundle-Ether10
Router(config-if)# ipv4 address 10.0.0.2 255.255.255.0
Router(config-if)# ipv6 address 2001:DB8::1
Router(config-if)# lacp period short
Router(config-if)# mac-address 1.2.3
Router(config-if)# bundle wait-while 0
Router(config-if)# exit
Router(config)# interface GigabitEthernet 0/2/0/5
Router(config-if)# bundle id 14 mode active
Router(config-if)# commit

/* Configure load balancing. */
Router# configure
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether10
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 11.11.11.11.11.11.11.14
```

```

Router(config-evpn-ac-es)# load-balancing-mode port-active
Router(config-evpn-ac-es)# commit
!
/* Configure address family session in BGP. */
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# bgp router-id 192.168.0.2
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp)# neighbor 192.168.0.3
Router(config-bgp-nbr)# remote-as 200
Router(config-bgp-nbr)# update-source Loopback 0
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr)# commit

```

Running Configuration

```

interface Bundle-Ether14
  ipv4 address 14.0.0.2 255.255.255.0
  ipv6 address 14::2/64
  lacp period short
  mac-address 1.2.3
  bundle wait-while 0
!
interface GigabitEthernet0/2/0/5
  bundle id 14 mode active
!
evpn
  interface Bundle-Ether14
    ethernet-segment
      identifier type 0 11.11.11.11.11.11.11.14
      load-balancing-mode port-active
    !
  !
!
router bgp 100
  bgp router-id 192.168.0.2
  address-family l2vpn evpn
  !
  neighbor 192.168.0.3
    remote-as 100
    update-source Loopback0
    address-family l2vpn evpn
  !
!
!

```

Verification

Verify the SRv6 services L3VPN active-standby redundancy using port-active mode configuration.

```

/* Verify ethernet-segment details on active DF router */
Router# show evpn ethernet-segment interface Bundle-Ether14 detail

```

Ethernet Segment Id	Interface	Nexthops
0011.1111.1111.1111.1114	BE14	192.168.0.2 192.168.0.3

```

      ES to BGP Gates      : Ready
      ES to L2FIB Gates   : Ready
      Main port           :
        Interface name    : Bundle-Ether14
        Interface MAC     : 0001.0002.0003

```



```

    IfHandle       : 0x000041d0
    State          : Up
    Redundancy     : Not Defined
    ESI type       : 0
    Value          : 11.1111.1111.1111.1114
    ES Import RT   : 1111.1111.1111 (from ESI)
    Source MAC     : 0000.0000.0000 (N/A)
    Topology       :
      Operational  : MH
      Configured   : Port-Active
    Service Carving : Auto-selection
    Multicast      : Disabled
    Peering Details :
      192.168.0.2 [MOD:P:00]
      192.168.0.3 [MOD:P:00]

```

```

Service Carving Results:
  Forwarders      : 0
  Permanent       : 0
  Elected        : 0
  Not Elected    : 0
MAC Flushing mode : STP-TCN
Peering timer     : 3 sec [not running]
Recovery timer    : 30 sec [not running]
Carving timer     : 0 sec [not running]
Local SHG label   : None
Remote SHG labels : 0

```

/* Verify bundle Ethernet configuration on active DF router */

Router# **show bundle bundle-ether 14**

Bundle-Ether14

```

Status: Up
Local links <active/standby/configured>: 1 / 0 / 1
Local bandwidth <effective/available>: 1000000 (1000000) kbps
MAC address (source): 0001.0002.0003 (Configured)
Inter-chassis link: No
Minimum active links / bandwidth: 1 / 1 kbps
Maximum active links: 64
Wait while timer: Off
Load balancing:
  Link order signaling: Not configured
  Hash type: Default
  Locality threshold: None
LACP: Operational
  Flap suppression timer: Off
  Cisco extensions: Disabled
  Non-revertive: Disabled
mLACP: Not configured
IPv4 BFD: Not configured
IPv6 BFD: Not configured

```

Port	Device	State	Port ID	B/W, kbps
Gi0/2/0/5	Local	Active	0x8000, 0x0003	1000000
Link is Active				

/* Verify ethernet-segment details on standby DF router */

Router# **show evpn ethernet-segment interface bundle-ether 10 detail**

```

Router# show evpn ethernet-segment interface Bundle-Ether24 detail
Ethernet Segment Id      Interface      Nexthops

```

```

-----
0011.1111.1111.1111.1114 BE24                                     192.168.0.2
                                                                192.168.0.3

ES to BGP Gates      : Ready
ES to L2FIB Gates    : Ready
Main port            :
    Interface name    : Bundle-Ether24
    Interface MAC     : 0001.0002.0003
    IfHandle          : 0x000041b0
    State             : Standby
    Redundancy        : Not Defined
ESI type             : 0
    Value             : 11.1111.1111.1111.1114
ES Import RT         : 1111.1111.1111 (from ESI)
Source MAC           : 0000.0000.0000 (N/A)
Topology             :
    Operational       : MH
    Configured        : Port-Active
Service Carving      : Auto-selection
    Multicast         : Disabled
Peering Details      :
    192.168.0.2 [MOD:P:00]
    192.168.0.3 [MOD:P:00]

Service Carving Results:
    Forwarders        : 0
    Permanent         : 0
    Elected          : 0
    Not Elected       : 0
MAC Flushing mode    : STP-TCN
Peering timer        : 3 sec [not running]
Recovery timer       : 30 sec [not running]
Carving timer        : 0 sec [not running]
Local SHG label      : None
Remote SHG labels    : 0

/* Verify bundle configuration on standby DF router */
Router# show bundle bundle-ether 24

Bundle-Ether24
Status: LACP OOS (out of service)
Local links <active/standby/configured>: 0 / 1 / 1
Local bandwidth <effective/available>: 0 (0) kbps
MAC address (source): 0001.0002.0003 (Configured)
Inter-chassis link: No
Minimum active links / bandwidth: 1 / 1 kbps
Maximum active links: 64
Wait while timer: Off
Load balancing:
    Link order signaling: Not configured
    Hash type: Default
    Locality threshold: None
LACP: Operational
    Flap suppression timer: Off
    Cisco extensions: Disabled
    Non-revertive: Disabled
mLACP: Not configured
IPv4 BFD: Not configured
IPv6 BFD: Not configured

Port          Device          State          Port ID          B/W, kbps
-----
Gi0/0/0/4     Local          Standby       0x8000, 0x0002   1000000
Link is in standby due to bundle out of service state

```

SRv6 Services: IPv4 L3VPN Active-Active Redundancy

This feature provides active-active connectivity to a CE device in a L3VPN deployment. The CE device can be Layer-2 or Layer-3 device connecting to the redundant PEs over a single LACP LAG port.

Depending on the bundle hashing, an ARP or IPv6 Network Discovery (ND) packet can be sent to any of the redundant routers. As a result, not all entries will exist on a given PE. In order to provide complete awareness, Layer-3 local route learning is augmented with remote route-synchronization programming.

Route synchronization between service PEs is required in order to provide minimum interruption to unicast and multicast services after failure on a redundant service PE. The following EVPN route-types are used for Layer-3 route synchronization:

- EVPN route-type 2 for synchronizing ARP tables
- EVPN route-type 7/8 for synchronizing IGMP JOINS/LEAVES

In a Layer-3 CE scenario, the router that connects to the redundant PEs may establish an IGP adjacency on the bundle port. In this case, the adjacency will be formed to one of the redundant PEs, and IGP customer routes will only be present on that PE. To synchronize Layer-3 customer subnet routes (IP Prefixes), the EVPN route-type 5 is used to carry the ESI and ETAG as well as the gateway address (prefix next-hop address).



Note Gratuitous ARP (GARP) or IPv6 Network Advertisement (NA) replay is not needed for CEs connected to the redundant PEs over a single LAG port.

The below configuration enables Layer-3 route synchronization for routes learned on the Ethernet-segment sub-interfaces.

```
evpn
 route-sync vrf default
 !
 vrf RED
  evi route-sync 10
 !
 vrf BLUE
  evi route-sync 20
 !
```



Note EVPN does not support untagged interfaces.

SRv6 Services: EVPN VPWS — All-Active Multi-Homing

Table 8: Feature History Table

Feature Name	Release	Description
SRv6 Services: EVPN VPWS — All-Active Multi-Homing (SRv6 Micro SID)	Release 7.3.2	This feature provides an ELINE (P2P) service with all-active multihoming capability over an SRv6 network. All-Active Multi-Homing enables an operator to connect a customer edge (CE) device to two or more provider edge (PE) devices to provide load balancing and redundant connectivity. With All-Active Multi-Homing, all the PEs can forward traffic to and from the multi-homed device.

EVPN VPWS All-Active Multi-Homing over SRv6 provides an ELINE (P2P) service with all-active multihoming capability over an SRv6 network.

All-Active Multi-Homing enables an operator to connect a customer edge (CE) device to two or more provider edge (PE) devices to provide load balancing and redundant connectivity. With All-Active Multi-Homing, all the PEs can forward traffic to and from the multi-homed device.



Note For information about EVPN VPWS, refer to the "EVPN Virtual Private Wire Service (VPWS)" chapter in the *L2VPN and Ethernet Services Configuration Guide*.

Configuring EVPN VPWS over SRv6



Note Complete the steps in [Configuring SRv6, on page 27](#) before performing these steps.

An SRv6 Locator for an EVPN VPWS service can be configured at 3 different levels independently:

- `global_locator` is the default locator for all EVPN-VPWS services
- `evi_locator` is applied to all EVPN-VPWS services for the specific EVI
- `evi_service_locator` is applied to an individual EVI service

When locators are configured at different levels at the same time, the following priority is implemented:

1. `evi_service_locator`
2. `evi_locator`

3. global_locator

This example shows how to configure an EVPN VPWS over SRv6 using a global locator for EVPN:

```
evpn
  segment-routing srv6
    locator sample_global_loc

l2vpn
  xconnect group sample_xcg
  p2p sample-vpws-12001-2002
    interface Bundle-Ether12001.2002
    neighbor evpn evi 12001 service 2002 segment-routing srv6
```

This example shows how to configure EVPN VPWS over SRv6 using specific EVI locator:

```
evpn
  evi 11001 segment-routing srv6
    locator sample_evi_loc

l2vpn
  xconnect group sample_xcg
  p2p sample-vpws-11001-2002
    interface Bundle-Ether11001.2002
    neighbor evpn evi 11001 service 2002 segment-routing srv6
```

This example shows how to configure an EVPN VPWS over SRv6 using a locator for an individual EVI service:

```
l2vpn
  xconnect group sample_xcg
  p2p sample-vpws-11001-2001
    interface Bundle-Ether11001.2001
    neighbor evpn evi 11001 service 2001 segment-routing srv6
    locator sample_evi_service_loc
```

Verification

Router# **show segment-routing srv6 locator**

Name	ID	Algo	Prefix	Status	Flags
sample_evi_loc	1	128	2001:0:8::/48	Up	U
sample_global_loc	2	0	2001:0:1::/48	Up	U

Router# **show segment-routing srv6 sid**

*** Locator: 'sample_evi_loc' ***

SID	State	RW	Behavior	Context	Owner
2001:0:8::	InUse	Y	uN (PSP/USD)	'default':8	sidmgr
2001:0:8:e000::	InUse	Y	uDX2	11001:2002	l2vpn_srv6
2001:0:8:e002::	InUse	Y	uA (PSP/USD)	[BE11, Link-Local]:128	isis-20
2001:0:8:e004::	InUse	Y	uA (PSP/USD)	[BE60, Link-Local]:128	isis-20

```

                InUse Y
2001:0:8:e006::          uA (PSP/USD)      [BE30, Link-Local]:128      isis-20
                InUse Y

*** Locator: 'sample_global_loc' ***

2001:0:1::              uN (PSP/USD)      'default':1                sidmgr
                InUse Y
2001:0:1:e001::         uDX2              12001:2002                 12vpn_srv6
                InUse Y
2001:0:1:e003::         uA (PSP/USD)      [BE11, Link-Local]:0        isis-20
                InUse Y
2001:0:1:e005::         uA (PSP/USD)      [BE60, Link-Local]:0        isis-20
                InUse Y
2001:0:1:e007::         uA (PSP/USD)      [BE30, Link-Local]:0        isis-20
                InUse Y

```

Router# **show evpn segment-routing srv6 detail**

```

Configured default locator: sample_global_loc
EVIs with unknown locator config: 0
VPWS with unknown locator config: 0

```

Locator name	Prefix	OOB	Service count	SID count
-----	-----	---	-----	-----
sample_evi_loc	2001:0:8::/48	False	1	1
Configured on EVIs <evi>: 11001				
sample_global_loc	2001:0:1::/48	False	1	1
Default locator				

Router# **show 12vpn xconnect group sample_xcg detail**

Thu Sep 2 14:39:22.575 UTC

Group sample_xcg, XC sample-vpws-11001-2002, state is up; Interworking none

```

AC: Bundle-Ether11001.2002, state is up
Type VLAN; Num Ranges: 1
Rewrite Tags: []
VLAN ranges: [2002, 2002]
MTU 1504; XC ID 0xc0002ee8; interworking none
Statistics:
  packets: received 0, sent 0
  bytes: received 0, sent 0
  drops: illegal VLAN 0, illegal length 0

```

EVPN: neighbor ::ffff:10.0.0.1, PW ID: evi 11001, ac-id 2002, state is up (established

)

```

XC ID 0xa0001f47
Encapsulation SRv6
Encap type Ethernet
Ignore MTU mismatch: Enabled
Transmit MTU zero: Disabled
Reachability: Up

```

SRv6	Local	Remote
-----	-----	-----
uDX2	2001:0:8:e000::	2001:0:3:e000::
AC ID	2002	2002
MTU	1518	1518
Locator	sample_evi_loc	N/A
Locator Resolved	Yes	N/A
SRv6 Headend	H.Encaps.L2.Red	N/A

```

Statistics:
  packets: received 0, sent 0
  bytes: received 0, sent 0

Group sample_xcg, XC sample-vpws-12001-2002, state is up; Interworking none
AC: Bundle-Ether12001.2002, state is up
  Type VLAN; Num Ranges: 1
  Rewrite Tags: []
  VLAN ranges: [2002, 2002]
  MTU 1504; XC ID 0xc0002eea; interworking none
  Statistics:
    packets: received 0, sent 0
    bytes: received 0, sent 0
    drops: illegal VLAN 0, illegal length 0
  EVPN: neighbor ::ffff:10.0.0.2, PW ID: evi 12001, ac-id 2002, state is up ( established
)
  XC ID 0xa0001f49
  Encapsulation SRv6
  Encap type Ethernet
  Ignore MTU mismatch: Enabled
  Transmit MTU zero: Disabled
  Reachability: Up

      SRv6              Local              Remote
      -----
uDX2      2001:0:1:e001::      2001:0:2:e001::
AC ID      2002                  2002
MTU        1518                  1518
Locator    sample_global_loc    N/A
Locator Resolved Yes            N/A
SRv6 Headend H.Encaps.L2.Red              N/A

Statistics:
  packets: received 0, sent 0
  bytes: received 0, sent 0

```

SRv6-Services: EVPN ELAN Layer 2 Gateway With Automated Steering To Flexible Algorithm Paths

Table 9: Feature History Table

Feature Name	Release Information	Feature Description
SRv6-Services: EVPN ELAN Layer 2 Gateway With Automated Steering To Flexible Algorithm Paths	Release 7.5.2	This feature builds upon EVPN BGP signaling to provide Emulated Local Area Network (ELAN) multipoint-to-multipoint Ethernet services over an SRv6-based network. You can enable automated steering of EVPN ELAN traffic into the path associated with a best-effort or Flex- Algorithm locator. This feature combines the benefits of EVPN ELAN service and SRv6 Micro-SIDs. For this feature, the segment-routing srv6 option was added to the evi command: evi (bridge-domain)


You can transport EVPN ELAN bridged unicast and broadcast, unknown unicast, and multicast (BUM) traffic over an SRv6 network in the Micro-SID format. Relevant SRv6 headend and endpoint definitions are noted below:

- **H.Encaps.L2.Red**: This headend router operation involves reduced encapsulation of Layer 2 or Ether frames using an SRv6 Policy.
- **uDT2U**: This endpoint router operation involves traffic decapsulation and unicast MAC L2 table lookup. This is used for the EVPN bridging unicast traffic use case.
- **uDT2M**: This endpoint router operation involves traffic decapsulation and L2 table flooding. This is used for the EVPN bridging BUM traffic with ESI filtering use case.



Note For more information on SRv6 headend and endpoint behaviors, refer to Segment Routing over IPv6 Overview.

The following topology is used to explain this feature.



Topology pointers:

- Customer edge (CE) devices send traffic between each other over the SRv6 network. The CE devices are CE-A, CE-B, and CE-C.
- The SRv6 network devices transport customer traffic, and they are R1, R2 .. till R8.
- The provider edge (PE) devices, R1, R3 and R8, are displayed in green. The SRv6 EVPN configurations must be enabled on the PE devices since they participate in the EVPN EVI.

This is a high-level overview of the traffic flow from CE-A to CE-C:

1. CE-A sets the source and destination addresses of the L2 frame and sends it to the connected PE device, R1.
2. R1 looks up the destination MAC address in the frame. Based on its forwarding table, R1 performs an H.Encaps.L2.Red operation and adds the destination DT2U SRv6 SID (say, fccc:ccc1:a1:e000::) to the packet.
3. From R1, traffic is sent over the SRv6 network to destination PE device R8.
4. When R8 receives the traffic, it performs the uDT2U function - It decapsulates the packet, performs a destination MAC address lookup in its forwarding table, and sends the frame through the local interface to CE-C.

Guidelines and Limitations

- For transporting BUM traffic, the BGP Route Reflector device should have an IOS XR release version 7.5.2 or later.

Configure SRv6 EVPN Bridging

Enable the following configurations on the PE routers R1, R3 and R8 since they participate in the EVPN EVI.



Note Complete the steps in Segment Routing over IPv6 Overview before performing these steps.

Associate SRv6 with EVPN

```
Router# configure terminal
Router(config)# evpn
```

Enable SRv6 under the EVPN mode and associate a global locator (**sample**, in this case) with EVPN.

```
Router(config-evpn)# segment-routing srv6
Router(config-evpn-srv6)# locator sample
Router(config-evpn-srv6)# exit
```

Associate an EVI-specific locator (**sample_evi_loc**) with EVI 1.

```
Router(config-evpn)# evi 1 segment-routing srv6
Router(config-evpn-instance)# locator sample_evi_loc
Router(config-evpn-instance)# commit
```

Associate SRv6 with L2VPN

Associate the sub-interface to the bridge domain:

```
Router(config)# l2vpn
Router(config-l2vpn)# bridge group bg1
Router(config-l2vpn-bg)# bridge-domain bd1
Router(config-l2vpn-bg-bd)# interface Hu0/0/0/0.1
Router(config-l2vpn-bg-bd-ac)# exit
```

Enable the **evi 1 segment-routing srv6** command under L2VPN bridge domain **bd1**.

```
Router(config-l2vpn-bg-bd)# evi 1 segment-routing srv6
Router(config-l2vpn-bg-bd-evi-srv6)# commit
```

Verification

In this sample output, SRv6 EVPN ELAN traffic unicast and multicast SID information is displayed.

```
Router# show evpn evi vpn-id 1 detail
```

VPN-ID	Encap	Bridge Domain	Type
1	SRv6	bd1	EVPN

```
..
  Stitching: Regular
  Unicast SID: fccc:cccl:a1:e000::
  Multicast SID: fccc:cccl:a1:e001::
..
```

In this sample output, EVI 1 details, including the corresponding SID and EVPN MAC address details are displayed.

```
Router# show evpn evi vpn-id 1 mac
```

VPN-ID	Encap	MAC address	IP address	Nexthop	Label	SID
1	SRv6	0010.3000.01d0	::	Hu0/0/0/0.1	0	fccc:cccl:a1:e000::

In this sample output, for the specified EVI and EVPN MAC address, SRv6 EVPN ELAN traffic details are displayed.

```
Router# show evpn evi vpn-id 1 mac 0010.3000.01d0 detail
```

VPN-ID	Encap	MAC	IP	Nexthop	Label	SID

```

----- address address -----
1      SRv6 ee03.0500.0130 :: 192.168.0.3 IMP=NULL fccc:ccc1:a3:e000::

Ethernet Tag : 0
Multi-paths Resolved : True
Multi-paths Internal label : None
Local Static : No
Remote Static : Yes
Local Ethernet Segment : N/A
Remote Ethernet Segment : 0100.0205.acce.5500.0500
Local Sequence Number : N/A
Remote Sequence Number : 0
Local Encapsulation : N/A
Remote Encapsulation : SRv6
Local E-Tree : Root
Remote E-Tree : Root
Remote matching E-Tree RT : No
Local AC-ID : 0x0
Remote AC-ID : 0x13

```

In this sample output, for the specified EVI, multicast SID details are displayed.

```
Router# show evpn evi vpn-id 1 inclusive-multicast detail
```

```

VPN-ID      Encap      EtherTag      Originating IP
----      -
1          SRv6        0             192.168.0.1
..
      TEPid: 0xffffffff
      PMSI Type: 6
      Nexthop: ::
      SR-TE Info: N/A
      SID: fccc:ccc1:a1:e001::
      Source: Local
      E-Tree: Root
..

```

In this sample output, for the specified MAC address, bridge domain information is displayed.

```
Router# show l2route evpn mac all | i ee03.0500.0130
```

```

Topo ID      Mac Address      Producer      Next Hop(s)
-----      -
1          ee03.0500.0130      L2VPN        ::ffff:10.0.0.10/IID/v6, N/A

```

In this sample output, SRv6 network locator and corresponding SID information are displayed.

uDT2U and **uDT2M** refer to SRv6 network endpoint operations. **uDT2U** indicates SRv6 traffic decapsulation, wherein EVPN bridged unicast traffic is forwarded out of the SR network. **uDT2M** indicates SRv6 traffic decapsulation, wherein EVPN bridged multicast traffic is forwarded out of the SR network.

```
Router# show segment-routing srv6 sid
```

```
*** Locator: 'sample_evi_loc' ***
```

```

SID          Behavior      Context      Owner      State      RW
---          -
fccc:ccc1:a1:: uN(PSP/USD) 'default':161 sidmgr      InUse      Y
fccc:ccc1:a1:e000:: uDT2U 7:0      l2vpn_srv6  InUse      Y
fccc:ccc1:a1:e001:: uDT2M 7:0      l2vpn_srv6  InUse      Y

```

In this sample output, CEF information is displayed, including SRv6 network endpoint details. **uDT2U** is an SRv6 network endpoint operation wherein SRv6 traffic is decapsulated and EVPN bridged unicast traffic is forwarded out of the SR network.

```

Router# show cef ipv6 fccc:cccl:a1:e000:: detail

fccc:cccl:a1:e000::, version 14, SRv6 Endpoint uDT2U, internal 0x1000001
0x0 (ptr 0x8ba26050) [1], 0x400 (0x8bbf7b58), 0x0 (0x92396138)

Prefix Len 64, traffic index 0, precedence n/a, priority 0
gateway array (0x8ba33e90) reference count 4, flags 0x0, source rib (7), 0 backups
[5 type 3 flags 0x8401 (0x8baf8ca8) ext 0x0 (0x0)]

LW-LDI[type=3, refc=1, ptr=0x8bbf7b58, sh-ldi=0x8baf8ca8]
gateway array update type-time 1 Sep  8 11:46:51.242

LDI Update time Sep  8 11:46:51.303
LW-LDI-TS Sep  8 11:46:51.380
via ::/128, 0 dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 NHID 0x0 [0x8afdf120 0x0]
next hop ::/128
XConnect ID: 0x80000003
Bridge ID: 0x1
Shg ID: 0x1

Load distribution: 0 (refcount 5)

Hash OK Interface Address
0 Y recursive Lookup in table

```

In the following examples, SRv6 EVPN ELAN traffic-related IID information is displayed.

```

Router# show evpn internal-id vpn-id 3001 detail

VPN-ID  Encap  Ethernet Segment Id  EtherTag  Internal ID
-----  -----
1        SRv6   0001.0001.0001.1501.0015  0          ::ffff:10.0.0.4

```

Summary pathlist:

```

0x05000002 (P) 192.168.0.3 fccc:cccl:a3:e000::
0x05000002 (P) 192.168.0.3 fccc:cccl:a4:e000::

```

```

Router# show cef vrf **iid ipv6 ::ffff:10.0.0.4

::ffff:10.0.0.4/128, version 39, SRv6 Headend, IID (EVPN-MH), internal 0x1000001 0x0 (ptr
0x8ba21798) [3], 0x0 (0x0), 0x0 (0x923967b0)
Updated Sep  8 18:01:06.495
Prefix Len 128, traffic index 0, precedence n/a, priority 0
gateway array (0x8ba36018) reference count 1, flags 0x2010, source rib (7), 0 backups
[1 type 3 flags 0x48441 (0x8baf9a28) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Sep  8 18:01:06.495
LDI Update time Sep  8 18:01:06.495

Level 1 - Load distribution: 0
[0] via fccc:cccl:a3::/128, recursive

via fccc:cccl:a3:e000::/128, 10 dependencies, recursive [flags 0x0]
path-idx 0 NHID 0x0 [0x8ba24e78 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop fccc:cccl:a3::/128 via fccc:cccl:a3::/48
SRv6 H.Encaps.L2.Red SID-list { fccc:cccl:a3:e000::}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Hu0/0/0/0 remote

```

```

via fccc:cccl:a4::/128, 10 dependencies, recursive [flags 0x100]
path-idx 0 NHID 0x0 [0x8ba24e78 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop fccc:cccl:a4::/128 via fccc:cccl:a4::/48
SRv6 H.Encaps.L2.Red SID-list {fccc:cccl:a4:e000::}

```

In this sample output, SRv6 EVPN ELAN traffic-related IID information is displayed.

Router# **show rib ipv6 iid**


IID	Prefix	Context	Owner	State	RW
0xa000001	::ffff:10.0.0.1	[EVPN-ELAN:evi=7:esi=8300.fccc.cccl.00a4.0000:nh=fccc:cccl:a4::eth_tag=0:type=0:encap=255:opaque=0]			
			l2vpn_iid	InUse	Y
0xa000002	::ffff:10.0.0.2	[EVPN-ELAN:evi=8:esi=8300.fccc.cccl.00a4.0000:nh=fccc:cccl:a4::eth_tag=0:type=0:encap=255:opaque=0]			
			l2vpn_iid	InUse	Y
0xa000003	::ffff:10.0.0.3	[EVPN-ELAN:evi=9:esi=8300.fccc.cccl.00a4.0000:nh=fccc:cccl:a4::eth_tag=0:type=0:encap=255:opaque=0]			
			l2vpn_iid	InUse	Y

SRv6 ESI Filtering

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 ESI Filtering	Release 7.11.1	Split Horizon Group (SHG) labels and Ethernet Segment Identifier (ESI) filtering functionalities exist on MPLS underlay networks. This feature introduces ESI filtering functionality to SRv6 underlay networks, using the End.DT2M SRv6 endpoint behavior. This behavior uses the "Arg.FE2" argument for SRv6, which is similar to the SHG label for MPLS. This feature allows nodes to identify BUM traffic based on the advertised ESI and prevent a loop by avoiding re-broadcasting the same traffic back towards the access node. This functionality is enabled by default.

Consider the below network topology where CE1 is attached to PE1 and PE2.



When broadcast, unknown unicast, and multicast (BUM) traffic is received by PE1 from CE1, PE1 floods this traffic on the core network. This results in a copy of the BUM traffic being sent to PE2. PE2 in turn floods it to the core and access network, and the BUM traffic would be received at CE1. This is undesirable as we are flooding the same traffic to the source from which it originated, thus causing a loop.

SRv6 ESI Filtering

A unique 16-bit ID is generated by all nodes in an Ethernet Segment and is transmitted to all nodes in the same EVPN instance via ES/EAD Route Type 1 NLRI (see ["EVPN Route Types"](#) for description of Route Type 1: Ethernet Auto-Discovery (AD) Route). This unique label functions in a similar way as the Split Horizon group label (SHG label). Ethernet Segment Identifier (ESI) filtering is the mechanism that controls which nodes to re-broadcast BUM traffic.



Note For more information, see ["Split Horizon Groups"](#) in the *L2VPN and Ethernet Services Configuration Guide*.

When PE1 receives BUM traffic from an attachment circuit (AC), it floods the traffic to all nodes in its core network, where PE2 is also member node. The traffic flooded to PE2 is sent with its previously advertised ID. When PE2 decodes the packet and examines the label, it recognizes this to be the label that it had generated. In this mechanism, PE2 will not flood the BUM traffic back towards the access. PE2 has successfully prevented a loop by avoiding re-broadcasting the same traffic back towards the source (CE1).

This feature introduces ESI filtering to SRv6 underlay networks, using the End.DT2M SRv6 endpoint behavior. This behavior uses the "Arg.FE2" argument, as defined in [IETF RFC 8986 SRv6 Network Programming](#) and [IETF draft SRv6 Argument Signaling for BGP Services](#). This argument provides a local mapping to ESI for split-horizon filtering of the received traffic to exclude a specific outgoing L2 interface (OIF), or a set of OIFs, from L2 table T flooding. The allocation of the argument values is local to the SR Segment Endpoint Node instantiating this behavior, and the signaling of the argument to other nodes for the EVPN functionality occurs via the control plane.

Usage Guidelines and Limitations

This feature is supported on NCS540 and NCS560 routers.



Note This feature is not supported on the N540-24Q8L2DD-SYS large-density router.

Verification

The output of the **show evpn ethernet-segment interface interface-name** now displays both the SHG label information for MPLS EVPN instances (EVIs) and the Arg.FE2 information for SRv6 EVIs:

```
Router# show evpn ethernet-segment interface Bundle-Ether 1
```

Ethernet Segment Id	Interface	Nexthops
0011.1111.1111.1111.1114	BE1	192.168.0.2 192.168.0.3
. . .		

```

Peering timer : 3 sec [not running]
Recovery timer : 30 sec [not running]
Local SHG label : 34251
Remote SHG labels : 2
38216 : nexthop 192.168.0.1
Arg.FE2 1:16 : nexthop 192.168.0.2

```

SRv6 Services: L3 EVPN

EVPN Route Type 5 (RT5) is used for the advertisement of EVPN routes using IP prefixes (refer to IETF [RFC 9136 - IP Prefix Advertisement in Ethernet VPN \(EVPN\)](#)) to provide end-to-end L3 connectivity

This feature adds support for carrying L3VPN routes in L2VPN EVPN EVPN RT5 address family instead of VPNv4 unicast and/or VPNv6 unicast address-family across SRv6 core (EVPN over SRv6 underlay).

Usage Guidelines and Limitations

BGP does not support dual VPNv4/v6 address family and EVPN RT5 address family on the same BGP session. For the route reflector (RR) to receive both Type-5 EVPN route and VPNv4/v6 address family, we recommend that you configure two pairs of loopback interfaces and configure two BGP loopback sessions between the RR and the PE: one session for VPNv4/v6 address family and one session for EVPN address family.

BGP sends all VRF routes via either VPNv4/v6 or EVPN address family. We recommend that you mark the VRF route via export route-policy and use neighbor out policy to either drop or pass the route for an address family to achieve the same net effect.


The following behaviors are supported:

- IPv4, IPv6, and IPv4/IPv6 (dual stack) L3 EVPN over SRv6
- uDT4
- uDT6
- uDT46
- Automated Steering to Flex-Algo (BGP per-VRF locator Flex-Algo (per-prefix))
- Automated Steering to SRv6 Policy (ODN/AS)

Configuring SRv6-based L3 EVPN

To enable SRv6-based L3 EVPN, you must enable SRv6 under BGP, specify the locator, and configure the SID allocation mode. The assignment of the locator can be done in multiple ways under the **router bgp** configuration. See [SRv6 Locator Inheritance Rules](#).

Figure 9: Configuration Example: Dual Stack L3 EVPN over SRv6



Configure the VRF (Dual-Stack IPv4/IPv6)

```

Router(config)# vrf VPN1
Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target
Router(config-vrf-import-rt)# 1:1
Router(config-vrf-import-rt)# exit
Router(config-vrf-af)# export route-target
Router(config-vrf-export-rt)# 1:1
Router(config-vrf-export-rt)# exit
Router(config-vrf)# address-family ipv6 unicast
Router(config-vrf-af)# import route-target
Router(config-vrf-import-rt)# 1:1
Router(config-vrf-import-rt)# exit
Router(config-vrf-af)# export route-target
Router(config-vrf-export-rt)# 1:1
Router(config-vrf-export-rt)# exit
Router(config-vrf-af)#

```

Configure the SRv6 Locator for an Individual VRF, with Per-VRF Label Allocation Mode

```

Router(config)# router bgp 100
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# additional-paths receive
Router(config-bgp-af)# additional-paths send
Router(config-bgp-af)# additional-paths selection route-policy add-path
Router(config-bgp-af)# exit
Router(config-bgp)# address-family vpnv6 unicast
Router(config-bgp-af)# additional-paths receive
Router(config-bgp-af)# additional-paths send
Router(config-bgp-af)# additional-paths selection route-policy add-path
Router(config-bgp-af)# exit
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp-af)# additional-paths receive
Router(config-bgp-af)# additional-paths send
Router(config-bgp-af)# additional-paths selection route-policy add-path
Router(config-bgp-af)# exit

Router(config-bgp)# neighbor 1111::1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# advertise vpnv4 unicast
Router(config-bgp-nbr-af)# advertise vpnv6 unicast
Router(config-bgp-nbr-af)# exit

```



```

Router(config-bgp-nbr) # exit

Router(config-bgp) # vrf VPN1
Router(config-bgp-vrf) # rd 100:1
Router(config-bgp-vrf) # address-family ipv4 unicast
Router(config-bgp-vrf-af) # segment-routing srv6
Router(config-bgp-vrf-af-srv6) # locator LOC1
Router(config-bgp-vrf-af-srv6) # alloc mode per-vrf
Router(config-bgp-vrf-af-srv6) # exit
Router(config-bgp-vrf-af) # exit
Router(config-bgp-vrf) # address-family ipv6 unicast
Router(config-bgp-vrf-af) # segment-routing srv6
Router(config-bgp-vrf-af-srv6) # locator LOC1
Router(config-bgp-vrf-af-srv6) # alloc mode per-vrf
Router(config-bgp-vrf-af-srv6) # exit
Router(config-bgp-vrf-af) # exit

Router(config-bgp-vrf) # neighbor 1.1.1.1
Router(config-bgp-vrf-nbr) # remote-as 200
Router(config-bgp-vrf-nbr) # address-family ipv4 unicast
Router(config-bgp-vrf-nbr-af) # exit
Router(config-bgp-vrf-nbr) # exit
Router(config-bgp-vrf) # neighbor 3333::3
Router(config-bgp-vrf-nbr) # remote-as 200
Router(config-bgp-vrf-nbr) # address-family ipv6 unicast

```

Running Configuration

```

vrf VPN1
  address-family ipv4 unicast
    import route-target
      1:1
    !
    export route-target
      1:1
    !
  !
  address-family ipv6 unicast
    import route-target
      1:1
    !
    export route-target
      1:1
    !
  !
!
router bgp 100
  address-family vpnv4 unicast
    additional-paths receive
    additional-paths send
    additional-paths selection route-policy add-path
  !
  address-family vpnv6 unicast
    additional-paths receive
    additional-paths send
    additional-paths selection route-policy add-path
  !
  address-family l2vpn evpn
    additional-paths receive
    additional-paths send
    additional-paths selection route-policy add-path
  !

```

```

neighbor 1111::1
  remote-as 100
  address-family l2vpn evpn
    advertise vpnv4 unicast
    advertise vpnv6 unicast
  !
!
vrf VPN1
  rd 100:1
  address-family ipv4 unicast
    segment-routing srv6
      locator LOC1
      alloc mode per-vrf
  !
  !
  address-family ipv6 unicast
    segment-routing srv6
      locator LOC1
      alloc mode per-vrf
  !
  !
neighbor 1.1.1.1
  remote-as 200
  address-family ipv4 unicast
  !
!
neighbor 3333::3
  remote-as 200
  address-family ipv6 unicast
  !
!
!
!

```

SRv6 Services: IPv6 L3VPN

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Services: IPv6 L3VPN	Release 7.3.1	With this feature, the egress PE can signal an SRv6 Service SID with the BGP overlay service route. The ingress PE encapsulates the IPv4/IPv6 payload in an outer IPv6 header where the destination address is the SRv6 Service SID provided by the egress PE. BGP messages between PEs carry SRv6 Service SIDs as a means to interconnect PEs and form VPNs.

This feature provides IPv6 L3VPNs (VPNv6) over an SRv6 network.

Usage Guidelines and Limitations

- SRv6 locator can be assigned globally, for all VRFs, or for an individual VRF.
SRv6 locator can be assigned globally, for all VRFs, for an individual VRF, or per-prefix.
- Per-VRF allocation mode is supported (uDT6 behavior)
- Dual-Stack L3VPN Services (IPv4, IPv6) are supported
- Equal-Cost Multi-path (ECMP) and Unequal Cost Multipath (UCMP) are supported.
- eBGP, OSPF, Static are supported as PE-CE protocol.
BGP (iBGP, eBGP), OSPF, Static are supported as PE-CE protocol.
- BGP route leaking between BGP Global and L3VPN is supported.
- Per-CE allocation mode is not supported (uDX6 behavior)
- iBGP is not supported as PE-CE protocol
- BGP route leaking is not supported

Configuring SRv6-based IPv6 L3VPN

To enable SRv6-based L3VPN, you need to enable SRv6 under BGP, specify the locator, and configure the SID allocation mode. The assignment of the locator can be done in different places under the **router bgp** configuration. See [SRv6 Locator Inheritance Rules, on page 43](#).

Use Case 1: Assigning SRv6 Locator Globally

This example shows how to configure the SRv6 locator name under BGP Global:

```
Node1(config)# router bgp 100
Node1(config-bgp)# segment-routing srv6
Node1(config-bgp-gbl-srv6)# locator Node1-locator
Node1(config-bgp-gbl-srv6)# exit
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# commit
```

Running Configuration

```
router bgp 100
  segment-routing srv6
    locator Node1-locator
  !
  address-family vpnv6 unicast
  !
  neighbor 3001::12:1:1:4
    remote-as 100
    address-family vpnv6 unicast
  !
```

```

!
vrf vrf_cust6
  rd 100:6
  address-family ipv6 unicast
!
!
end

```

Use Case 2: Assigning SRv6 Locator for All VRFs

To configure the SRv6 locator for all VRFs under VPNv6 Address Family and specify the allocation mode, use the following commands:

- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6:** Enable SRv6
- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6 alloc mode {*per-vrf*}:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT6 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp *as-number* address-family vpnv6 unicast vrf all segment-routing srv6 locator *WORD*:** Specify the locator

This example shows how to configure the SRv6 locator for all VRFs under VPNv6 Address Family, with per-VRF label allocation mode:

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# vrf all
Node1(config-bgp-af-vrfall)# segment-routing srv6
Node1(config-bgp-af-vrfall-srv6)# locator Node1-locator
Node1(config-bgp-af-vrfall-srv6)# alloc mode per-vrf
Node1(config-bgp-af-vrfall-srv6)# exit
Node1(config-bgp-af-vrfall)# exit
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# commit

```

Running Configuration

```

router bgp 100
  address-family vpnv6 unicast
    vrf all
      segment-routing srv6
        locator Node1-locator
        alloc mode per-vrf
    !
  !
  neighbor 3001::12:1:1:4
    remote-as 100
    address-family vpnv6 unicast

```

```

!
!
vrf vrf_cust6
  rd 100:6
  address-family ipv6 unicast
!
!
end

```

Use Case 3: Assigning SRv6 Locator for a specific VRF

To configure the SRv6 locator for a specific VRF under IPv6 Address Family and specify the allocation mode, use the following commands:

- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6:** Enable SRv6
- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 alloc mode { per-vrf }:** Specify the SID behavior (allocation mode)
 - Use the **per-vrf** keyword to specify that the same service SID (uDT6 behavior) be used for all the routes advertised from a unique VRF.
- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 locator WORD:** Specify the locator

This example shows how to configure the SRv6 locator for an individual VRF, with per-VRF label allocation mode:

```

Node1(config)# router bgp 100
Node1(config-bgp)# address-family vpnv6 unicast
Node1(config-bgp-af)# exit
Node1(config-bgp)# neighbor 3001::12:1:1:4
Node1(config-bgp-nbr)# remote-as 100
Node1(config-bgp-nbr)# address-family vpnv6 unicast
Node1(config-bgp-nbr-af)# exit
Node1(config-bgp-nbr)# exit
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# rd 100:6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6
Node1(config-bgp-vrf-af-srv6)# locator Node1-locator
Node1(config-bgp-vrf-af-srv6)# alloc mode per-vrf
Node1(config-bgp-vrf-af-srv6)# commit

```

Running Configuration

```

router bgp 100
  address-family vpnv6 unicast
  !
  neighbor 3001::12:1:1:4
  remote-as 100
  address-family vpnv6 unicast
  !
  !
vrf vrf_cust6
  rd 100:6
  address-family ipv6 unicast
  segment-routing srv6
    locator Node1-locator
    alloc mode per-vrf

```

```

!
!
!
!
end

```

Use Case 4: Assigning SRv6 Locator for a Specific Prefix

This use case provides the ability to assign a specific SRv6 locator for a given prefix or a set of prefixes. The egress PE advertises the prefix with the specified locator. This allows for per-prefix steering into desired transport behaviors, such as Flex Algo.

To assign an SRv6 locator for a specific prefix, configure a route policy to specify the SID allocation mode based on match criteria. Examples of match criteria are destination-based match or community-based match.

- Supported SID allocation modes are per-VRF and per-CE.
- For per-VRF allocation mode, you can also specify the SRv6 locator.
 - If an SRv6 locator is specified in the route policy, BGP will use that to allocate per-VRF SID. If the specified locator is invalid, the SID will not be allocated.
 - If an SRv6 locator is not specified in the route policy, the default locator is used to allocate the SID. If the default locator is not configured in BGP, then the SID will not be allocated.
- Per-CE allocation mode always uses the default locator configured under BGP to allocate the SID.

For more information on configuring routing policies, refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following example shows a route policy specifying the SID allocation mode with destination-based match:

```

Node1(config)# route-policy set_per_prefix_locator_rpl
Node1(config-rpl)# if destination in (3001::1:1:1:1/128) then
Node1(config-rpl-if)# set srv6-alloc-mode per-vrf locator locator1
Node1(config-rpl-if)# elseif destination in (3001::2:2:2:2/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf locator locator2
Node1(config-rpl-elseif)# elseif destination in (3001::3:3:3:3/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-vrf
Node1(config-rpl-elseif)# elseif destination in (3001::4:4:4:4/128) then
Node1(config-rpl-elseif)# set srv6-alloc-mode per-ce
Node1(config-rpl-elseif)# else
Node1(config-rpl-else)# drop
Node1(config-rpl-else)# endif
Node1(config-rpl)# end-policy

```

To specify per-prefix allocation mode for a specific VRF under IPv6 Address Family, use the following command:

- **router bgp as-number vrf WORD address-family ipv6 unicast segment-routing srv6 alloc mode route-policy policy_name**

This example shows how to specify per-prefix allocation mode for a specific VRF (vrf_cust6) under the IPv6 address family:

```

Node1(config)# router bgp 100
Node1(config-bgp)# vrf vrf_cust6
Node1(config-bgp-vrf)# address-family ipv6 unicast
Node1(config-bgp-vrf-af)# segment-routing srv6

```

```
Node1(config-bgp-vrf-af-srv6)# alloc mode route-policy set_per_prefix_locator_rpl
```

Running Configuration

```
route-policy set_per_prefix_locator_rpl
  if destination in (3001::1:1:1:1/128) then
    set srv6-alloc-mode per-vrf locator locator1
  elseif destination in (3001::2:2:2:2/128) then
    set srv6-alloc-mode per-vrf locator locator2
  elseif destination in (3001::3:3:3:3/128) then
    set srv6-alloc-mode per-vrf
  elseif destination in (3001::4:4:4:4/128) then
    set srv6-alloc-mode per-ce
  else
    drop
  endif
end-policy
!
router bgp 100
  vrf vrf_cust6
    address-family ipv6 unicast
      segment-routing srv6
        alloc mode route-policy set_per_prefix_locator_rpl
    !
  !
!
!
```

Verify that the local and received SIDs have been correctly allocated under VPNv6 and specific VRF (vrf_cust6):

```
Node1# show bgp vpnv6 unicast local-sids
```

```
BGP router identifier 10.1.1.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Local Sid              Alloc mode   Locator
Route Distinguisher: 8:8
*>i3008::8:8:8:8/128 NO SRv6 Sid              -            -
* i                 NO SRv6 Sid              -            -
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1:1/128 fc00:0:1:40::          per-vrf       locator1
*> 3001::2:2:2:2/128 fc00:8:1:40::          per-vrf       locator2
*> 3001::3:3:3:3/128 fc00:9:1:40::          per-vrf       locator4
*> 3001::4:4:4:4/128 fc00:9:1:41::          per-ce        locator4
*> 3001::5:5:5:5/128 NO SRv6 Sid              -            -
*> 3001::12:1:1:5/128 NO SRv6 Sid             -            -
*>i3008::8:8:8:8/128 NO SRv6 Sid              -            -
```

```
Node1# show bgp vpnv6 unicast received-sids
```

```
BGP router identifier 10.1.1.1, local AS number 1
BGP generic scan interval 60 secs
```

```

Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 50
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network                Next Hop                Received Sid
Route Distinguisher: 8:8
*>i3008::8:8:8/128        10.1.1.2                fc00:0:2:42::
* i                        2400:2020:42:2fff::1    fc00:0:2:42::

Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1/128        11.1.1.2                NO SRv6 Sid
*> 3001::2:2:2/128        11.1.1.2                NO SRv6 Sid
*> 3001::3:3:3/128        11.1.1.2                NO SRv6 Sid
*> 3001::4:4:4/128        11.1.1.2                NO SRv6 Sid
*> 3001::5:5:5/128        11.1.1.2                NO SRv6 Sid
*> 3001::12:1:1:5/128     13.2.2.2                NO SRv6 Sid
*>i3008::8:8:8/128        10.1.1.2                fc00:0:2:42::

```

Node1# show bgp vrf vrf_cust6 local-sids

```

BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 10.1.1.1:0
VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013   RD version: 37
BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network                Local Sid                Alloc mode   Locator
Route Distinguisher: 8:8
*>i3008::8:8:8/128        NO SRv6 Sid                -            -
* i                        NO SRv6 Sid                -            -
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1/128        fc00:0:1:40::              per-vrf      locator1
*> 3001::2:2:2/128        fc00:8:1:40::              per-vrf      locator2
*> 3001::3:3:3/128        fc00:9:1:40::              per-vrf      locator4
*> 3001::4:4:4/128        fc00:9:1:41::              per-ce       locator4
*> 3001::5:5:5/128        NO SRv6 Sid                -            -
*> 3001::12:1:1:5/128     NO SRv6 Sid                -            -
*>i3008::8:8:8/128        NO SRv6 Sid                -            -

```

Node1# show bgp vrf vrf_cust6 received-sids

```

BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 100:6
VRF ID: 0x60000004
BGP router identifier 10.1.1.1, local AS number 1
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000013   RD version: 37
BGP main routing table version 37
BGP NSR Initial initsync version 18 (Reached)

```



```
BGP NSR/ISSU Sync-Group versions 0/0
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
```

```
i - internal, r RIB-failure, S stale, N Nexthop-discard
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```


```

Network          Next Hop          Received Sid
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::1:1:1:1/128 11.1.1.2          NO SRv6 Sid
*> 3001::2:2:2:2/128 11.1.1.2          NO SRv6 Sid
*> 3001::3:3:3:3/128 11.1.1.2          NO SRv6 Sid
*> 3001::4:4:4:4/128 11.1.1.2          NO SRv6 Sid
*> 3001::5:5:5:5/128 11.1.1.2          NO SRv6 Sid
*> 3001::12:1:1:5/128 13.2.2.2         NO SRv6 Sid
*>i3008::8:8:8:8/128 10.1.1.2         fc00:0:2:42::

```

Verification

The following figure shows a VPNv6 scenario. The sequence of commands included correspond to router Node1 acting as Ingress PE, and routers Node4 and Node5 acting as Egress PEs.



The following examples shows how to verify the SRv6 based L3VPN configurations for an Individual VRF with per VRF label allocation mode.

In this example, we can observe the uDT6 SID associated with the IPv6 L3VPN, where uDT6 behavior represents Endpoint with decapsulation and IPv6 table lookup.

```
Node1# show segment-routing srv6 sid
```

```
Fri Jan 29 19:31:53.293 UTC
```

```
*** Locator: 'Node1-locator' ***
```

SID	State	RW	Behavior	Context	Owner
cafe:0:1::			uN (PSP/USD)	'default':1	sidmgr

```

InUse Y
cafe:0:1:e000::      uA (PSP/USD)      [Hu0/0/0/0, Link-Local]:0      isis-1
InUse Y
cafe:0:1:e001::      uA (PSP/USD)      [Hu0/0/0/1, Link-Local]:0      isis-1
InUse Y
cafe:0:1:e002::      uDT4              'vrf_cust1'                  bgp-100
InUse Y
cafe:0:1:e003::      uDT4              'vrf_cust2'                  bgp-100
InUse Y
cafe:0:1:e004::      uDT4              'vrf_cust3'                  bgp-100
InUse Y
cafe:0:1:e005::      uDT4              'vrf_cust4'                  bgp-100
InUse Y
cafe:0:1:e006::      uDT4              'vrf_cust5'                  bgp-100
InUse Y
cafe:0:1:e007::      uA (PSP/USD)      [Hu0/0/0/0, Link-Local]:0:P    isis-1
InUse Y
cafe:0:1:e008::      uA (PSP/USD)      [Hu0/0/0/1, Link-Local]:0:P    isis-1
InUse Y
cafe:0:1:e009::      uDT6              'default'                    bgp-100
InUse Y
cafe:0:1:e00a::      uDT6            'vrf_cust6'                  bgp-100
InUse Y

```

The following examples show how to verify the SRv6 based L3VPN configuration using the **show bgp vpnv6 unicast** commands on the Ingress PE.

```

Node1# show bgp vpnv6 unicast summary
Fri Jan 29 19:33:01.177 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 6
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	6	6	6	6	6	0

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
cafe:0:4::4	0	100	122	123	6	0	0	00:20:05	1
cafe:0:5::5	0	100	111	111	0	0	0	00:49:46	1

```

Node1# show bgp vpnv6 unicast rd 100:6
Fri Jan 29 19:41:01.334 UTC
BGP router identifier 10.1.1.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 8
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

```

Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard

```

Origin codes: i - IGP, e - EGP, ? - incomplete
  Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::12:1:1:1/128 ::                0          32768 ?
*>i3001::12:1:1:4/128 cafe:0:4::4          0      100      0 ?
*>i3001::12:1:1:5/128 cafe:0:5::5          0      100      0 ?

Processed 3 prefixes, 3 paths

Node1# show bgp vpv6 unicast rd 100:6 3001::12:1:1:4/128
Fri Jan 29 19:41:42.008 UTC
BGP routing table entry for 3001::12:1:1:4/128, Route Distinguisher: 100:6
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          6         6
Last Modified: Jan 29 19:29:35.858 for 00:12:06
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local, (received & used)
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
    Received Label 0xe00a00
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
    Received Path ID 0, Local Path ID 1, version 6
    Extended community: RT:100:6
PSID-Type:L3, SubTLV Count:1
SubTLV:
T:1(Sid information), Sid:cafe:0:4::, Behavior:62, SS-TLV Count:1
SubSubTLV:
T:1(Sid structure):
    Source AFI: VPNv6 Unicast, Source VRF: vrf_cust6, Source Route Distinguisher: 100:6

```

The following examples show how to verify the BGP prefix information for VRF instances:

```

Node1# show bgp vrf vrf_cust6 ipv6 unicast
Fri Jan 29 19:42:05.675 UTC
BGP VRF vrf_cust6, state: Active
BGP Route Distinguisher: 100:6
VRF ID: 0x60000007
BGP router identifier 10.1.1.1, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0800016 RD version: 8
BGP main routing table version 8
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
  Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:6 (default for vrf vrf_cust6)
*> 3001::12:1:1:1/128 ::                0          32768 ?
*>i3001::12:1:1:4/128 cafe:0:4::4          0      100      0 ?
*>i3001::12:1:1:5/128 cafe:0:5::5          0      100      0 ?

Processed 3 prefixes, 3 paths

Node1# show bgp vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128

BGP routing table entry for 3001::12:1:1:4/128, Route Distinguisher: 100:6

```

```

Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          17         17
Last Modified: Jan 15 16:50:44.032 for 01:48:21
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local, (received & used)
    cafe:0:4::4 (metric 30) from cafe:0:4::4 (10.1.1.4)
      Received Label 0xe00a00
      Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 17
      Extended community: RT:100:6
      PSID-Type:L3, SubTLV Count:1
      SubTLV:
        T:1(Sid information), Sid:cafe:0:4::, Behavior:62, SS-TLV Count:1
        SubSubTLV:
          T:1(Sid structure):
            Source AFI: VPNv6 Unicast, Source VRF: vrf_cust6, Source Route Distinguisher: 100:6

```

The following examples show how to verify the current routes in the Routing Information Base (RIB):

```

Node1# show route vrf vrf_cust6 ipv6 unicast
Fri Jan 29 19:43:28.067 UTC

```

```

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G - DAGR, l - LISp
       A - access/subscriber, a - Application route
       M - mobile route, r - RPL, t - Traffic Engineering, (!) - FRR Backup path

```

Gateway of last resort is not set

```

L   3001::12:1:1:1/128 is directly connected,
    01:01:23, Loopback105
B   3001::12:1:1:4/128
    [200/0] via cafe:0:4::4 (nexthop in vrf default), 00:13:52
B   3001::12:1:1:5/128
    [200/0] via cafe:0:5::5 (nexthop in vrf default), 00:05:53

```

```

Node1# show route vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128
Fri Jan 29 19:43:55.645 UTC

```

```

Routing entry for 3001::12:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:29:35.696 for 00:14:20
  Routing Descriptor Blocks
    cafe:0:4::4, from cafe:0:4::4
      Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
      Route metric is 0
  No advertising protos.

```

```

Node1# show route vrf vrf_cust6 ipv6 unicast 3001::12:1:1:4/128 detail
Fri Jan 29 19:44:17.914 UTC

```

```

Routing entry for 3001::12:1:1:4/128
  Known via "bgp 100", distance 200, metric 0, type internal
  Installed Jan 29 19:29:35.696 for 00:14:42

```

```

Routing Descriptor Blocks
  cafe:0:4::4, from cafe:0:4::4
    Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
    Route metric is 0
    Label: None
    Tunnel ID: None
    Binding Label: None
    Extended communities count: 0
    Source RD attributes: 0x0000:100:6
    NHID:0x0(Ref:0)
    SRv6 Headend: H.Encaps.Red [f3216], SID-list {cafe:0:4:e00a::}
Route version is 0x1 (1)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_REMOTE
Download Priority 3, Download Version 3
No advertising protos.

```

The following examples show how to verify the current IPv6 Cisco Express Forwarding (CEF) table:

```

Node1# show cef vrf vrf_cust6 ipv6
Fri Jan 29 19:44:56.888 UTC

::/0
  drop      default handler
3001::12:1:1:1/128
  receive   Loopback105
3001::12:1:1:4/128
  recursive  cafe:0:4::/128
3001::12:1:1:5/128
  recursive  cafe:0:5::/128
fe80::/10
  receive
ff02::/16
  receive
ff02::2/128
  receive
ff02::1:ff00:0/104
  receive
ff05::/16
  receive
ff12::/16
  receive

Node1# show cef vrf vrf_cust6 ipv6 3001::12:1:1:4/128
Fri Jan 29 19:45:23.607 UTC
3001::12:1:1:4/128, version 3, SRv6 Headend, internal 0x50000001 0x30 (ptr 0x78f2e0e0) [1],
0x0 (0x0), 0x0 (0x888a3ac8)
Updated Jan 29 19:29:35.700
Prefix Len 128, traffic index 0, precedence n/a, priority 3
  via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
  path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
  next hop VRF - 'default', table - 0xe0800000
  next hop cafe:0:4::/128 via cafe:0:4::/48
  SRv6 H.Encaps.Red SID-list {cafe:0:4:e00a::}

Node1# show cef vrf vrf_cust6 ipv6 3001::12:1:1:4/128 detail
Fri Jan 29 19:45:55.847 UTC
3001::12:1:1:4/128, version 3, SRv6 Headend, internal 0x50000001 0x30 (ptr 0x78f2e0e0) [1],
0x0 (0x0), 0x0 (0x888a3ac8)
Updated Jan 29 19:29:35.700

```

```

Prefix Len 128, traffic index 0, precedence n/a, priority 3
gateway array (0x78afe238) reference count 1, flags 0x2010, source rib (7), 0 backups
[1 type 3 flags 0x48441 (0x78ba9a60) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Jan 29 19:29:35.699
LDI Update time Jan 29 19:29:35.701

Level 1 - Load distribution: 0
[0] via cafe:0:4::/128, recursive

via cafe:0:4::/128, 7 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x78cd2a14 0x0]
next hop VRF - 'default', table - 0xe0800000
next hop cafe:0:4::/128 via cafe:0:4::/48
SRv6 H.Encaps.Red SID-list {cafe:0:4:e00a::}

Load distribution: 0 1 (refcount 1)

Hash  OK  Interface                Address
0      Y   HundredGigE0/0/0/0        remote
1      Y   HundredGigE0/0/0/1        remote

```

SRv6 Services: L2 and L3 Services with Remote SIDs from W-LIB

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Services: L2 and L3 Services with Remote SIDs from Wide Local ID Block	Release 7.9.1	This feature enables an SRv6 headend node to receive and install remote SIDs with Wide (32-bit) functions (Remote W-LIB). The Remote W-LIB is supported for Layer 3 (VPN/BGP global) and Layer 2 EVPN services (ELINE/ELAN). This capability is enabled by default.

This capability is enabled by default; there is no CLI to configure this capability at the ingress PE.

An SRv6 Service SID is used to identify a specific service function. This Service SID inserted into the packet header by the source node is used to steer the packet along a specific path that includes the service function.

The Service SID signaled by transposing a variable part of the SRv6 SID value (function, argument, or both) and carrying them in the existing label fields to achieve more efficient compression of those service prefix NLRIs in BGP update messages. The SRv6 SID Structure Sub-Sub-TLV (SSTLV) contains appropriate length fields when the SRv6 Service SID is signaled in split parts to enable the receiver to put together the SID accurately.

The Transposition Offset indicates the bit position. The Transposition Length indicates the number of bits that are being taken out of the SRv6 SID value and put into high order bits of label field.

For example, a remote W-LIB uSID **fcbb:bb00:0200:fff0:0001::** with a SRv6 SID SSTLV of **BL=32; NL=16; FL=32; AL=0, TPOS len/offset=16/64** is defined as follows:

- Block length (BL) of 32 bits = fcbb:bb00

- Node length (NL) of 16 bits = 0200
- Function length (FL) of 32 bits = fff0:0001
- Argument length (AL) of 0
- Transposition length (TPOS len) of 16 bits = 0001
- Transposition offset (TPOS offset) of 64 bits = fcbb:bb00:0200:fff0:

This results in a SID value of **fcbb:bb00:0200:fff0::** and Label value of **0x0001**.

Example

The following example shows output of a BGP route table for a VPNv4 prefix learned from three egress PEs:

- BGP Path 1 from next-hop 7::1 and a 32-bit uDT4 function (0xffff0 4002) allocated from W-LIB
- BGP Path 2 from next-hop 9::1 and a 16-bit uDT4 function (0x4002) allocated from LIB
- BGP Path 3 from next-hop 8::1 and a 16-bit uDT4 function (0x4002) allocated from LIB

Note the following fields in the output:

- Function length of 16 bits for LIB and 32 bits for W-LIB
- Transposition offset value of 48 bits for LIB and 64 bits for W-LIB
- Transposition length value of 16 bits for LIB/W-LIB

```
Router# show bgp vpnv4 un rd 100:2 2.2.0.1/32 detail
```

```
BGP routing table entry for 2.2.0.1/32, Route Distinguisher: 100:2
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          5314      5314
```

```
Flags: 0x20061292+0x00060000; multipath; backup available;
```

```
Last Modified: Jan 20 14:37:59.189 for 00:00:19
```

```
Paths: (3 available, best #1)
```

```
Not advertised to any peer
```

```
Path #1: Received by speaker 0
```

```
Flags: 0x2000000085070005+0x00, import: 0x39f
```

```
Not advertised to any peer
```

```
Local
```

```
7::1 (metric 20) from 2::1 (192.0.0.1), if-handle 0x00000000
```

```
Received Label 0x40020
```

```
Origin IGP, localpref 150, valid, internal, best, group-best, multipath,
```

```
import-candidate, imported
```

```
Received Path ID 1, Local Path ID 1, version 5314
```

```
Extended community: RT:100:2
```

```
Originator: 192.0.0.1, Cluster list: 2.0.0.1
```

```
PSID-Type:L3, SubTLV Count:1, R:0x00,
```

```
SubTLV:
```

```
T:1(Sid information), Sid:fcbb:cc00:7001:fff0::, F:0x00, R2:0x00, Behavior:63,  
R3:0x00, SS-TLV Count:1
```

```
SubSubTLV:
```

```
T:1(Sid structure):
```

```
Length [Loc-blk,Loc-node,Func,Arg]:[32,16,32,0], Tpose-len:16, Tpose-offset:64
```

```
Source AFI: VPNv4 Unicast, Source VRF: VRF_2, Source Route Distinguisher: 100:2
```

```
Path #2: Received by speaker 0
```

```
Flags: 0x2000000084060005+0x00, import: 0x096
```

```
Not advertised to any peer
```

```

Local
  9::1 (metric 20) from 2::1 (192.0.0.3), if-handle 0x00000000
    Received Label 0x40020
    Origin IGP, localpref 100, valid, internal, backup(protect multipath), add-path,
import-candidate, imported
    Received Path ID 2, Local Path ID 5, version 5314
    Extended community: RT:100:2
    Originator: 192.0.0.3, Cluster list: 2.0.0.1
    PSID-Type:L3, SubTLV Count:1, R:0x00,
    SubTLV:
      T:1(Sid information), Sid:fccc:cc00:9001::, F:0x00, R2:0x00, Behavior:63, R3:0x00,
SS-TLV Count:1
    SubSubTLV:
      T:1(Sid structure):
        Length [Loc-blk,Loc-node,Func,Arg]:[32,16,16,0], Tpose-len:16, Tpose-offset:48
        Source AFI: VPNv4 Unicast, Source VRF: VRF_2, Source Route Distinguisher: 100:2
    Path #3: Received by speaker 0
    Flags: 0x2000000084070005+0x00, import: 0x296
    Not advertised to any peer
Local
  8::1 (metric 20) from 2::1 (192.0.0.2), if-handle 0x00000000
    Received Label 0x40020
    Origin IGP, localpref 150, valid, internal, multipath, backup, add-path,
import-candidate, imported
    Received Path ID 3, Local Path ID 4, version 5314
    Extended community: RT:100:2
    Originator: 192.0.0.2, Cluster list: 2.0.0.1
    PSID-Type:L3, SubTLV Count:1, R:0x00,
    SubTLV:
      T:1(Sid information), Sid:fccc:cc00:8001::, F:0x00, R2:0x00, Behavior:63, R3:0x00,
SS-TLV Count:1
    SubSubTLV:
      T:1(Sid structure):
        Length [Loc-blk,Loc-node,Func,Arg]:[32,16,16,0], Tpose-len:16, Tpose-offset:48
        Source AFI: VPNv4 Unicast, Source VRF: VRF_2, Source Route Distinguisher: 100:2

```


SRv6/MPLS L3 Service Interworking Gateway

Table 13: Feature History Table

Feature Name	Release	Description
Identical Route Distinguisher (RD) for Interworking Gateways between MPLS and SRv6 Domains	Release 24.1.1	You can now configure the same Route Distinguisher (RD) for interworking gateways catering to both MPLS and SRv6 domains that help conserve hardware resources, reduce the BGP table scale and minimize the processing load on routers. At the same time, it ensures seamless connectivity across SRv6 and MPLS L3 EVPN domains, thus promoting interoperability and efficiency in modern network environments. Previously, a unique RD was required to extend L3 services between MPLS and SRv6 domains resulting in higher router load and resource consumption, which could have affected performance.
SRv6/MPLS L3 Service Interworking Gateway (SRv6 Micro-SID)	Release 7.3.2	This feature enables you to extend L3 services between MPLS and SRv6 domains by providing service continuity on the control plane and data plane. This feature allows for SRv6 L3VPN domains to interwork with existing MPLS L3VPN domains. The feature also allows a way to migrate from MPLS L3VPN to SRv6 L3VPN.

SRv6/MPLS L3 Service Interworking Gateway enables you to extend L3 services between MPLS and SRv6 domains by providing service continuity on the control plane and data plane.

This feature allows for SRv6 L3VPN domains to interwork with existing MPLS L3VPN domains. The feature also allows a way to migrate from MPLS L3VPN to SRv6 L3VPN.

The SRv6/MPLS L3 Service Interworking Gateway provides both transport and service termination at the gateway node. The gateway generates both SRv6 VPN SIDs and MPLS VPN labels for all prefixes under the VRF configured for re-origination. The gateway supports traffic forwarding from MPLS domain to SRv6 domain by popping the MPLS VPN label, looking up the destination prefix, and pushing the appropriate SRv6 encapsulation. From SRv6 domain to MPLS domain, the gateway removes the outer IPv6 header, looks up the destination prefix, and pushes the VPN and next-hop MPLS labels.

VRFs on the gateway node are configured with 2 sets of route targets (RTs):

- MPLS L3VPN RTs
- SRv6 L3VPN RTs (called *stitching RTs*)


The gateway performs the following actions:

- Imports service routes received from one domain (MPLS or SRv6)
- Re-advertises exported service routes to the other domain (next-hop-self)
- Stitches the service on the data plane (uDT4/H.Encaps.Red ↔ service label)

SRv6/MPLS L3 Service Interworking Gateway Scenarios


The following scenario is used to describe the gateway functionality:

- Node 1 is an L3VPN PE in the MPLS domain with an SR prefix SID label of 16001 for its Loopback interface 10.1.1.1/32.
- Node 2 is the SRv6/MPLS L3 Service Interworking Gateway. In the MPLS domain, it has an SR prefix SID label of 16002 for its Loopback interface 10.1.1.2/32. In the SRv6 domain, it has an SRv6 locator of B:0:2::/48 and Loopback interface B:0:2::2/128.
- Node 3 is an L3VPN PE in the SRv6 domain with SRv6 locator of B:0:3::/48 and Loopback interface B:0:3::3/128.



Scenario 1: SRv6-to-MPLS Control-Plane Direction/MPLS-to-SRv6 Data-Plane Direction

The figure below describes the associated control-plane behaviors in the SRv6-to-MPLS direction for traffic in the MPLS-to-SRv6 data-plane direction.



A. Node 3 advertises a BGP L3VPN update for prefix B.0.0.0/8 with RD corresponding to VRFA, including the SRv6 VPN SID (B:0:3:V9::) assigned to this VRF, in the SRv6 domain.



Note SRv6 uDT4 function value "V9" is not a valid hex number, however it is used for illustration purposes to remind you of its connection to a VRF.

B. Node 2 (gateway) imports the BGP L3VPN update and programs its FIB:

- MPLS label 24010 is allocated for VRFA
- Prefix B.0.0.0/8 is programmed with an "SR Headend Behavior with Reduced Encapsulation in an SR Policy" function (H.Encaps.Red) of B:0:3:V9::



Note The gateway follows per-VRF label and per-VRF SID allocation methods.

C. Node 2 re-originates a BGP L3VPN update for the same prefix, including the MPLS VPN label (24010) allocated for the VRF, in the MPLS domain.

D. Site A sends traffic to an IPv4 prefix (B.B.B.B) of Site B

E. Node 1 encapsulates incoming traffic with the MPLS VPN label (24010) and the prefix SID MPLS label (16002) of the BGP next-hop (Node 2).


F. Node 2 performs the following actions:

- Pops the MPLS VPN label and looks up the destination prefix
- Encapsulates the payload in an outer IPv6 header with destination address (DA) equal to the H.Encaps.Red function (B:0:3:V9::)

G. Node 3 removes the outer IPv6 header, looks up the payload destination address (B.B.B.B), and forwards to Site B.

Scenario 2: MPLS-to-SRv6 Control-Plane Direction/SRv6-to-MPLS Data-Plane Direction

The figure below describes the associated control-plane behaviors in the MPLS-to-SRv6 direction for traffic in the SRv6-to-MPLS data-plane direction.



A. Node 1 advertises a BGP L3VPN update for prefix A.0.0.0/8 with RD corresponding to VRFA, including the MPLS VPN label (24055) assigned to this VRF, in the MPLS domain.

B. Node 2 (gateway) imports the BGP L3VPN update and programs its FIB:

- Prefix A.0.0.0/8 is programmed to impose an MPLS VPN label (24055) and the prefix SID MPLS label (16001) of the BGP next-hop (Node 1)
- "Endpoint with decapsulation and IPv4 table lookup" function (uDT4) of B:0:2:V8:: is allocated to VRFA



Note SRv6 uDT4 function value "V8" is not a valid hex number, however it is used for illustration purposes to remind you of its connection to a VRF.



Note The gateway follows per-VRF label and per-VRF SID allocation methods.

C. Node 2 re-originates a BGP L3VPN update for the same prefix, including the uDT4 function (B:0:2:V8::) allocated for the VRF, in the SRv6 domain.

D. Site B sends traffic to an IPv4 prefix (A.A.A.A) of Site A.

E. Node 3 Encapsulates the payload in an outer IPv6 header with destination address (DA) equal to the uDT4 function (B:0:2:V8::).

F. Node 2 performs the following actions:

- Removes the outer IPv6 header and looks up the destination prefix
- Pushes the MPLS VPN label (24055) and the prefix SID MPLS label (16001) of the BGP next-hop (Node 1)

G. Node 1 pops the MPLS VPN label, looks up the payload destination address (A.A.A.A), and forwards to Site A.

Example

Leveraging the topology described in the above use-case, this example shows the SRv6/MPLS L3 Service Interworking Gateway configuration required at Node 2.

The following configuration shows how to enable SRv6 with locator and configure encapsulation parameters:

```
segment-routing
srv6
  encapsulation
    source-address B:0:2::2
  !
  locators
    locator LOC1
    prefix B:0:2::/48
  !
  !
  !
  !
```

The following configuration shows how to configure a VPNv4 VRF with the following route targets (RTs):

- 1111:1, RT used for MPLS L3VPN
- 2222:1, RT used for SRv6 L3VPN (stitching RT)

```
vrf ACME
address-family ipv4 unicast
import route-target
  1111:1
  2222:1 stitching
!
export route-target
  1111:1
  2222:1 stitching
!
!
!
```

The following configuration shows how to configure SRv6/SRv6 VPNs under BGP:

```
router bgp 100
segment-routing srv6
  locator LOC1
!
neighbor 10.1.1.1
address-family vpnv4 unicast
  import re-originate stitching-rt
  route-reflector-client
  advertise vpnv4 unicast re-originated
!
neighbor B:0:3::1
address-family vpnv4 unicast
```

```

import stitching-rt re-originate
route-reflector-client
encapsulation-type srv6
advertise vpnv4 unicast re-originated stitching-rt
!
vrf ACME
address-family ipv4 unicast
enable label-mode
segment-routing srv6

```

You can configure same route distinguisher (RD) on the Node 1, Node 2 and GW. This example shows how to configure same route distinguisher (RD) on the Node 1, Node 2 and GW. In this example, **rd 5000:2** is used on Node 1, Node 2 and GW.

```

/* Configuration on Node 1*/
vrf ACMErd 5000:2
address-family ipv4 unicast
import route-target
1111:1
!
export route-target
1111:1
!
!
!

/* Configuration on Node 2*/
vrf ACMErd 5000:2
address-family ipv4 unicast
import route-target
2222:1
!
export route-target
2222:1
!
!
!

/* Configuration on GW*/
vrf ACMErd 5000:2
address-family ipv4 unicast
import route-target
1111:1
2222:1 stitching
!
export route-target
1111:1
2222:1 stitching
!
!
!


```

L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway

This feature adds support for L3 EVPN interworking between SRv6 and MPLS.

L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway enables you to extend L3 EVPN services between MPLS and SRv6 domains by providing service continuity on the control plane and data plane.

This feature allows for SRv6 L3 EVPN domains to interwork with existing MPLS L3 EVPN domains. The feature also allows a way to migrate from MPLS L3 EVPN to SRv6 L3 EVPN.



The L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway provides both transport and service termination at the gateway node.

VRFs on the gateway node are configured with 2 sets of route targets (RTs):

- L3 EVPN/MPLS RTs
- L3 EVPN/SRv6 RTs (called *stitching RTs*)

The gateway performs the following actions:

- Imports service routes received from one domain (L3 EVPN/MPLS or L3 EVPN/SRv6)
- Re-originates exported service routes to the other domain and setting next-hop-self
- Stitches the service routes in the data plane (uDT4/H.Encaps.Red ↔ MPLS service label)

The gateway generates both L3 EVPN/SRv6 SIDs and L3 EVPN/MPLS labels for all prefixes under the VRF configured for re-origination:

- MPLS-to-SRv6 Control Plane Direction

The gateway imports routes received from the MPLS side (via EVPN RT5) and re-originates them in L3VPN VRF with a per-VRF SRv6 SID.

- SRv6-to-MPLS Control Plane Direction

The gateway imports routes received from the SRv6 side (via EVPN RT5) and re-originates them in L3VPN VRF with a per-VRF label.

In the data plane, the gateway forwards traffic from the MPLS domain to the SRv6 domain by popping the MPLS L3 EVPN label, looking up the destination prefix, and pushing the appropriate SRv6 encapsulation. In the opposite direction, the gateway removes the outer IPv6 header, looks up the destination prefix, and pushes the L3 EVPN and next-hop MPLS labels.

Usage Guidelines and Limitations

L3 EVPN/SRv6 and L3 EVPN/MPLS Interworking Gateway is supported for IPv4 and IPv6.

Configuration Example

Leveraging the topology described above, this example shows the SRv6/MPLS L3 EVPN Service Interworking Gateway configuration required at Node 2.

The following configuration shows how to enable SRv6 with locator and configure encapsulation parameters.

```
segment-routing
  srv6
    encapsulation
      source-address b:0:2::2
    !
    locators
      locator LOC1
      prefix b:0:2::/48
    !
  !
!
!
```

The following configuration shows how to configure a VPNv4/VPNv6 VRF with the following route targets (RTs):

- 1111:1, RT used for MPLS L3 EVPN
- 2222:1, RT used for SRv6 L3 EVPN (stitching RT)

```
vrf VPN1
  address-family ipv4 unicast
    import route-target
      1111:1
      2222:1 stitching
    !
    export route-target
      1111:1
      2222:1 stitching
    !
  !
  address-family ipv6 unicast
    import route-target
      1111:1
      2222:1 stitching
    !
    export route-target
      1111:1
      2222:1 stitching
    !
  !
!
```

The following configuration shows how to configure SRv6/SRv6 VPNs under BGP:

```
router bgp 100
  segment-routing srv6
    locator LOC1
  !
  address-family vpnv4 unicast
  !
  address-family vpnv6 unicast
  !
  address-family l2vpn evpn
  !
  neighbor 2222::2
    remote-as 100
    description SRv6 side peering
    address-family l2vpn evpn
      import reoriginate stitching-rt (Imports NLRI's that match normal route target
        identifier and exports re-originated NLRI's assigned with the stitching
```



```

        route target identifier)
    route-reflector-client
    encapsulation-type srv6
    advertise vpnv4 unicast re-originated (Specifies advertisement of re-originated
        VPNv4 unicast routes)
    advertise vpnv6 unicast re-originated (Specifies advertisement of re-originated
        VPNv6 unicast routes)
!
!
neighbor 3.3.3.3
remote-as 100
description MPLS side peering stitching side
address-family l2vpn evpn
    import stitching-rt reoriginate (Imports NLRIs that match stitching route target
        identifier and exports re-originated NLRIs assigned with the normal route
        target identifier)
    advertise vpnv4 unicast re-originated stitching-rt (Advertise local VPNv4 unicast
        routes assigned with stitching route target identifier)
    advertise vpnv6 unicast re-originated stitching-rt (Advertise local VPNv6 unicast
        routes assigned with stitching route target identifier)
!
!
vrf VPN1
rd 100:2
address-family ipv4 unicast
    mpls alloc enable
!
address-family ipv6 unicast
    mpls alloc enable
!
!
!
```

SRv6/MPLS Dual-Connected PE


Table 14: Feature History Table

Feature Name	Release	Description
SRv6/MPLS Dual-Connected PE (SRv6 Micro SID)	Release 7.3.2	This feature allows a PE router to support IPv4 L3VPN services for a given VRF with both MPLS and SRv6. This is MPLS and SRv6 L3VPNv4 co-existence scenario and is sometimes referred to as dual-connected PE.

A PE router can support IPv4 L3VPN service for a given VRF with both MPLS and SRv6. This is MPLS and SRv6 L3VPNv4 co-existence scenario and is sometimes referred to as dual-connected PE.

In the figure below, node 2 is a dual-connected PE to Site C, providing:

- MPLS/IPv4 L3VPN between Site A and Site C
- SRv6/IPv4 L3VPN between Site B and Site C



522080

Configure BGP to Support Dual-Mode

Enable MPLS Label Allocation

Use the **router bgp as-number vrf WORD address-family ipv4 unicast mpls alloc enable** command under the VRF address-family to enable per-prefix mode for MPLS labels. Additionally, use the **router bgp as-number vrf WORD address-family ipv4 unicast label mode {per-ce | per-vrf}** command to choose the type of label allocation.

```
Router(config)# router bgp 100
Router(config-bgp)# vrf blue
Router(config-bgp-vrf)# rd 1:10
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# mpls alloc enable
Router(config-bgp-vrf-af)# label mode per-ce
Router(config-bgp-vrf-af)# segment-routing srv6
Router(config-bgp-vrf-af-srv6)# alloc mode per-ce
Router(config-bgp-vrf-af-srv6)# exit
Router(config-bgp-vrf-af)# exit
Router(config-bgp-vrf)# exit
Router(config-bgp)#
```

Configure Encaps on Neighbor to Send the SRv6 SID Toward the SRv6 Dataplane

By default, if a VRF prefix has both an MPLS label and an SRv6 SID, the MPLS label is sent when advertising the prefix to the PE. To advertise a VRF prefix with an SRv6 SID to an SRv6 session, use the **encapsulation-type srv6** command under the neighbor VPN address-family.

```
Router(config-bgp)# neighbor 192::6
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# encapsulation-type srv6
Router(config-bgp-nbr-af)# exit
```

Running Config

```
router bgp 100
neighbor 192::6
remote-as 1
address-family ipv4 unicast
encapsulation-type srv6
```

```

!
!
vrf blue
rd 1:10
address-family ipv4 unicast
  mpls alloc enable
  label mode per-ce
  segment-routing srv6
  alloc mode per-ce
!
!
!
!
!

```

Static SRv6 pseudo-wire

Static SRv6 pseudo-wire (PW) gives the flexibility to

- configure single-homing static SRv6 pseudo-wire between two Provider Edge (PE) routers, in an SRv6 core plane, and
- extends Virtual Private Wire Service (VPWS) capabilities by incorporating SRv6.

Table 15: Feature History Table

Feature Name	Release Information	Feature Description
Static SRv6 pseudowire	Release 25.1.1	This feature introduces support for Static Segment Routing over IPv6 (SRv6) Pseudo-wires in the IOSXR platform. This enhancement aims to extend the existing Virtual Private Wire Service (VPWS) capabilities by incorporating SRv6, providing improved flexibility and scalability for service providers.

Key aspects of SRv6 static pseudo-wire

This feature enhances the existing Virtual Private Wire Service (VPWS) by incorporating Segment Routing over IPv6 (SRv6) capabilities, which previously supported only MPLS.

This enhancement enables the configuration and management of static SRv6 pseudo-wires, thereby providing enhanced flexibility and scalability for service providers. Additionally, it addresses the specific requirements of customers who require static configuration for their network services.

- **Static Configuration** : Unlike dynamic pseudo-wires that rely on signaling protocols like LDP or BGP, static pseudo-wires are manually configured. This allows for more control and customization in specific network scenarios.
- **Segment Routing over IPv6 (SRv6)**: SRv6 is a network architecture that uses IPv6 addresses to encode the path that packets should follow through the network. Each segment in the path is represented by an SRv6 Segment Identifier (SID).
- **Pseudo-wire (PW)**: A pseudo-wire is a mechanism that emulates the properties of a traditional telecommunications circuit over a packet-switched network. It allows for the transport of Layer 2 frames over an IP/MPLS network.

Benefits of Static SRv6 pseudo-wire

The benefits of having an option to configure pseudo-wire in a static SRv6 core plane gives businesses an edge over their competitors, as listed.

- **Enhanced Flexibility:** By using SRv6, service providers can define more granular and flexible paths for their pseudo-wires, improving network efficiency and performance.
- **Scalability:** SRv6 allows for a scalable solution that can handle a large number of pseudo-wires without the need for complex BGP signaling protocols.
- **Interoperability:** The feature is designed to work across various platforms and planes making it versatile for different network environments.

Configure static SRv6 pseudo-wire

The purpose of this task is to configure static SRv6 pseudo-wires. To configure this, the network operators must define the local and remote SRv6 SIDs for the pseudo-wire endpoints. This involves specifying the segment routing configuration and ensuring that the pseudo-wire is correctly mapped to the desired network path.

Follow these steps to configure static SRv6 pseudo-wire

Procedure

Step 1 Define the Locator

Example:

```
Router(config)#segment-routing
Router(config-sr)#srv6
Router(config-srv6)#formats
Router(config-srv6-fmts)#format usid-f3216
Router(config-srv6-fmts)#usid local-id-block explicit start 0xee00 (default start is 0xfe00
- optional required if you if more than 256 )
Router(config-srv6-fmts)#exit

Router(config-srv6)encapsulation
Router(config-srv6)source-address 1::1
Router(config-srv6)traffic-class propagate

Router(config-srv6)locators
Router(config-srv6-locators)locator locator0
Router(config-srv6-locators)micro-segment behavior unode psp-usd
Router(config-srv6-locators)prefix fccc:cc00:1::/48
```

Step 2 Define the Sub-interface

Example:

```
Router#sh run int be101.40010001

interface Bundle-Ether101.40010001 l2transport
 encapsulation dot1q 1
 rewrite ingress tag pop 1 symmetric
!
```

Step 3 Define the xconnect group**Example:**

```
Router(config)#l2vpn
Router(config-l2vpn)#xconnect group xg4001
Router(config-l2vpn-xc)#p2p vpws-400100150
Router(config-l2vpn-xc-p2p)#interface Bundle-Ether101
Router(config-l2vpn-xc-p2p)#neighbor segment-routing srv6 static local fccc:cc00:1:ee96::
remote fccc:cc00:2:ee96::
```

Note

The **remote fccc:cc00:2:ee96::** SID is the remote device SID. You must configure this on both ends of the PE. However the local and remote SIDs interchange

Step 4 Run the following show commands to view the defined locator, evpn, L2vpn, xconnect and evpn internal ID details.**Example:**

```
Router#show segment-routing srv6 locator locator0 sid fccc:cc00:1:ee01:: detail
```

SID	State	RW	Behavior	Context	Owner
fccc:cc00:1:ee01::	InUse	Y	uDX2	pw_id:3221225475	l2vpn_srv6

SID Function: 0xee01
 SID context: { static pw_id=3221225475 }
 Locator: 'locator0'
 Allocation type: Explicit
 Created: Dec 16 15:38:33.852 (00:05:19 ago)

```
Router# show evpn segment-routing srv6 detail
```

```
Configured default locator: None
Configured default SID Function Length: 16 bits
EVI's with unknown locator config: 0
VPWS with unknown locator config: 0
Global SID Function Length: 16 bits
No SRv6 locators in use
```

```
Router# show l2vpn xconnect group xg4001 xc-name vpws-40010001 det
```

```
Group xg4001, XC vpws-40010001, state is up; Interworking none
Decoupled mode: Disabled
AC: Bundle-Ether101.40010001, state is up
Type VLAN; Num Ranges: 1
Rewrite Tags: []
VLAN ranges: [1, 1]
MTU 9186; XC ID 0xc0000003; interworking none
Statistics:
  packets: received 7109476, sent 5069183
  bytes: received 910012144, sent 648854934
  drops: illegal VLAN 0, illegal length 0
PW: neighbor ::ffff:10.0.0.4, PW ID 2684354565, state is up
PW class not set, XC ID 0xc0000003
Encapsulation SRv6, protocol none
PW type Ethernet, control word unknown, interworking none
PW backup disable delay 0 sec
```

```

Ignore MTU mismatch: Disabled
Transmit MTU zero: Disabled
Reachability: Up
Nexthop type: Internal ID ::ffff:10.0.0.4

```

SRv6	Local	Remote
-----	-----	-----
uDX2	fccc:cc00:1:ee01::	fccc:cc00:2:ee01::
AC ID	0	0
MTU	9200	9200
Locator	locator0	N/A
Locator Resolved	Yes	N/A
SRv6 Headend	H.Encaps.L2.Red	N/A

```

Statistics:
  packets: received 5069183, sent 7109476
  bytes: received 648854934, sent 910012144

```

```
Router#show evpn internal-id
```

VPN-ID	Encap	Ethernet Segment Id	EtherTag	Internal ID
-----	-----	-----	-----	-----
0	SRv6	fccc:cc00:2:ee01::	13421820	::ffff:10.0.0.4
	Summary pathlist (ID 0x000000000000001a):			
0x05000005		fccc:cc00:2:ee01::		fccc:cc00:2:ee01::
0	SRv6	fccc:cc00:2:ee02::	13421820	::ffff:10.0.0.2
	Summary pathlist (ID 0x000000000000000a):			
0x05000002		fccc:cc00:2:ee02::		fccc:cc00:2:ee02::

```
Router#show evpn internal-id detail
```

VPN-ID	Encap	Ethernet Segment Id	EtherTag	Internal ID
-----	-----	-----	-----	-----
0	SRv6	fccc:cc00:2:ee01::	13421820	::ffff:10.0.0.4
	Path resolved: TRUE			
	Path Internal ID: ::ffff:10.0.0.4			
	IP-Tunnel (ID:0x0000000000000019)			
		fccc:cc00:2:ee01::		fccc:cc00:2:ee01::
	Path Version:1, Originating PE::			
	Summary pathlist (ID 0x000000000000001a):			
0x05000005		fccc:cc00:2:ee01::		fccc:cc00:2:ee01::
0	SRv6	fccc:cc00:2:ee02::	13421820	::ffff:10.0.0.2
	Path resolved: TRUE			
	Path Internal ID: ::ffff:10.0.0.2			
	IP-Tunnel (ID:0x0000000000000009)			
		fccc:cc00:2:ee02::		fccc:cc00:2:ee02::
	Path Version:1, Originating PE::			
	Summary pathlist (ID 0x000000000000000a):			
0x05000002		fccc:cc00:2:ee02::		fccc:cc00:2:ee02::

SRv6 SID Information in BGP-LS Reporting

BGP Link-State (BGP-LS) is used to report the topology of the domain using nodes, links, and prefixes. This feature adds the capability to report SRv6 Segment Identifier (SID) Network Layer Reachability Information (NLRI).

The following NLRI has been added to the BGP-LS protocol to support SRv6:

- Node NLRI: SRv6 Capabilities, SRv6 MSD types
- Link NLRI: End.X, LAN End.X, and SRv6 MSD types
- Prefix NLRI: SRv6 Locator
- SRv6 SID NLRI (for SIDs associated with the node): Endpoint Function, BGP-EPE Peer Node/Set

This example shows how to distribute IS-IS SRv6 link-state data using BGP-LS:

```
Router(config)# router isis 200
Router(config-isis)# distribute link-state instance-id 200
```

DHCPv4 Relay Agent and Proxy Support over SRv6

This feature introduces support for DHCPv4 Relay Agent and Proxy over SRv6.

An IOS XR router can act as a DHCPv4 relay agent/proxy with a DHCPv4 server connected over an SRv6 network.

The following functionality is supported:

- DHCPv4 relay agent/proxy over SRv6 with DHCPv4 server (helper-address) located in default VRF (global)
- DHCPv4 relay agent/proxy over SRv6 with DHCPv4 server (helper-address) located in non-default VRF
- DHCPv4 relay agent/proxy on interfaces associated with a default VRF (global)
- DHCPv4 relay agent/proxy on interfaces associated with a non-default VRF
- DHCPv4 relay agent/proxy on Ethernet physical interfaces
- DHCPv4 relay agent/proxy on Ethernet bundle interfaces

For information on configuring DHCPv4 relay agent and proxy, refer to the “Implementing the Dynamic Host Configuration Protocol” chapter in the *IP Addresses and Services Configuration Guide*.

DHCPv6 Relay Agent Support over SRv6

Table 16: Feature History Table

Feature Name	Release Information	Feature Description
DHCPv6 Relay Agent Support on SRv6	Release 7.2.2	An IOS XR router can act as a DHCPv6 relay agent with a DHCPv6 server connected over an SRv6 network. A DHCP relay agent is a host that forwards DHCP packets between clients and servers that do not reside on a shared physical subnet.

This feature introduces support for DHCPv6 Relay Agent over SRv6.

An IOS XR router can act as a DHCPv6 relay agent with a DHCPv6 server connected over an SRv6 network.

The following functionality is supported:

- DHCPv6 relay agent over SRv6 with DHCPv6 server (helper-address) located in default VRF (global)
- DHCPv6 relay agent over SRv6 with DHCPv6 server (helper-address) located in non-default VRF
- DHCPv6 relay agent on interfaces associated with a default VRF (global)
- DHCPv6 relay agent on interfaces associated with a non-default VRF
- DHCPv6 relay agent on Ethernet physical interfaces
- DHCPv6 relay agent on Ethernet bundle interfaces

For information on configuring DHCPv6 relay agent, refer to the “Implementing the Dynamic Host Configuration Protocol” chapter in the *IP Addresses and Services Configuration Guide*.

Full-Replace Migration to SRv6 Micro-SID

Table 17: Feature History Table

Feature Name	Release	Description
Full-Replace Migration to SRv6 Micro-SID	Release 7.8.1	This feature enables migration of existing SRv6 SID format1 to SRv6 Micro-SIDs (f3216) formats. Earlier, only one format was supported at a time, and you had to choose either format1 or Micro-SID format for the deployment of services. Migration from Full-length SIDs to SRv6 Micro-SIDs was not possible. The hw-module profile segment-routing srv6 mode base-and-micro-segment-f3216 subcommand is introduced under hw-module profile segment-routing srv6 mode.

During the Full-Replace migration, both underlay and services are migrated from format1 to f3216. The underlay migration is done using the *Ship in the night* strategy, where updates into your environment are incremental, thereby phasing out your existing transport protocols when ready. This method minimizes the service disruption, and is recommended for seamless migration. The services migration is done using *swap* procedures, where the incoming transport label is swapped with an outgoing transport label.

The format1 to f3216 migration is seamless, requires minimal configurations, and no IETF signaling extensions. The migration enables preference of Micro-SID f3216 over format1, and minimizes traffic drop with faster convergence.

EVPN supports migration of the following services from format1 to f3216:

- IS-IS underlay (TILFA, uLoop, FlexAlgo)
- L3 overlay (VPNv4/VPNv6 and IPv4/IPv6)
- L2 overlay (EVPN VPWS - All-Active Multi-Homing)
- SRv6-MPLS IW Gateway, dual-connected PE

The following modes are supported in the context of migration:

- **Base:** SRv6 classic with format1 only.
- **Dual:** SRv6 classic with format1 and SRv6 Micro-SID with f3216 will both coexist.
- **f3216:** Micro-segment format. f3216 represents the format 3216, which is 32-bit block and 16-bit IDs.

The migration process involves the following steps:

1. **Prepare for migration:** Upgrade the network nodes to an image that is Micro-SID f3216 capable, and allows the coexistence of format1 and f3216.

Use the **hw-module profile segment-routing srv6 mode base-and-micro-segment-f3216** command to change the hardware module profile to SRv6 dual-mode.

2. **Migrate the underlay to f3216:** Enable IS-IS as an underlay protocol on PE nodes. The IS-IS configuration adds f3216 locators to format1 locators. Both format1 and f3216 endpoint SIDs are allocated, installed, and announced during this stage. f3216 is the preferred option over format1 for underlay paths.

The IS-IS SR headends provide faster convergence to f3216. Faster convergence to f3216 is done on the per-prefix per-path level, does not need any new CLI, and avoids packet drops. The format1 locators are removed after underlay traffic convergence to f3216 on all nodes. The format1 locators are unconfigured from IS-IS, and deleted from SRv6.

Change the hardware profile to SRv6 f3216 mode.

At the end of this step, the migration status of the following P Nodes are:

- Locator reachability: f3216 only
- Underlay endpoint/headends: f3216 only
- The mode is f3216

At the end of this step, the migration status of the following PE Nodes are:

- Locator reachability: format1 and f3216
- Underlay endpoint/headends: f3216 only
- Overlay endpoint/headends: format1
- The mode is base-and-f3216

3. **Migrate the overlay to f3216:** Enables overlay f3216 under BGP and EVPN on all PE nodes. The BGP and EVPN configuration replaces format1 by f3216 locators. During this stage, the f3216 Micro-SIDs are allocated, installed, and announced, while the format1 SIDs are deallocated, uninstalled, and withdrawn.

The format1 locators are removed after overlay traffic convergence to f3216 on all nodes. The format1 locators are unconfigured from BGP and EVPN, and deleted from SRv6. For a transient period, BGP and EVPN might have some paths with format1 and some with f3216.

At the end of this step, the migration status of the following is:

- For P/PE Nodes:
 - Locator reachability: f3216 only
 - Underlay endpoint/headends: f3216 only
 - Overlay endpoint/headends: f3216 only
 - The mode is f3216

The migration starts with SRv6 base format1, and ends with SRv6 Micro-SID f3216. The migration states are:

1. **Initial state:** This is the early migration state of a deployment, for the supported features. This state comprises SRv6 base with format1.

You must use the **hw-module profile srv6 mode base** command for the initial state.

This example shows the initial state of migration with SRv6 and configure locator:

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLoc0
Router(config-srv6-locators)# prefix flbb:bbbb:bb00:0001::/64
```

This example shows the initial state of migration with SRv6 and IS-IS:

```
Router(config)# router isis 100
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing srv6
Router(config-isis-srv6)# locator myLoc0
```

This example shows the initial state of migration with SRv6 and BGP/EVPN:

```
Router(config)# router bgp 100
Router(config-bgp)# bgp router-id 10
Router(config-bgp)# segment-routing srv6
Router(config-bgp-srv6)# locator myLoc0
```

```
Router(config)# evpn
Router(config-evpn)# segment-routing srv6
Router(config-evpn-srv6)# locator myLoc0
```

2. **In-migration state:** The migration procedures are initiated, and are in progress. This state comprises SRv6 in dual mode (base with format1, and Micro-SID with f3216).

You must use the **hw-module profile segment-routing srv6 mode base-and-micro-segment-f3216** command for the in-migration state.

This example shows the in-migration state with SRv6 and configure locator:

```
Router(config)# segment-routing srv6
Router(config-srv6)# locators
Router(config-srv6-locators)# locator myLoc0
Router(config-srv6-locators)# prefix flbb:bbbb:bb00:0001::/64
Router(config-srv6-locators)# delayed-delete
Router(config-srv6-locators)# locator myuLoc0
Router(config-srv6-locators)# micro-segment behavior unode psp-usd
Router(config-srv6-locators)# prefix fcbb:bb00:0001::/48
```

This example shows the in-migration state with SRv6 and IS-IS:

```
Router(config)# router isis 100
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing srv6
Router(config-isis-srv6)# locator myLoc0
Router(config-isis-srv6)# locator myuLoc0
```

This example shows the in-migration state with SRv6 and BGP/EVPN:

```
Router(config)# router bgp 100
Router(config-bgp)# bgp router-id 10
```

```
Router(config-bgp) # segment-routing srv6
Router(config-bgp-srv6) # locator myuLoc0
```

```
Router(config) # evpn
Router(config-evpn) # segment-routing srv6
Router(config-evpn-srv6) # locator myuLoc0
```

3. **End state:** This is the state of deployment at the end of the migration. At the end state, you can update the network and add new features. The Full-Replace migration end state can be of two modes:
- **Full-Replace:** Both underlay and overlay are migrated to Micro-SID f3216. Full-Replace is the Cisco recommended migration type.
 - **uF1:** Underlay migrated to Micro-SID f3216, overlay remains format1. The uF1 migration is a transient state of the Full-Replace migration type.

You need to use the **hw-module profile srv6 mode f3216** command for the end migration state.

This example shows the end state with SRv6 and configure locator:

```
Router(config) # segment-routing srv6
Router(config-srv6) # locators
Router(config-srv6-locators) # locator myuLoc0
Router(config-srv6-locators) # micro-segment behavior unode psp-usd
Router(config-srv6-locators) # prefix fcbb:bb00:0001::/48
```

This example shows the end state with SRv6 and IS-IS:

```
Router(config) # router isis 100
Router(config-isis) # address-family ipv6 unicast
Router(config-isis-af) # segment-routing srv6
Router(config-isis-srv6) # locator myuLoc0
```

This example shows the end state with SRv6 and BGP/EVPN:

```
Router(config) # router bgp 100
Router(config-bgp) # bgp router-id 10
Router(config-bgp) # segment-routing srv6
Router(config-bgp-srv6) # locator myuLoc0
```

```
Router(config) # evpn
Router(config-evpn) # segment-routing srv6
Router(config-evpn-srv6) # locator myuLoc0
```

Run the following command to check the result of migration, as shown in the example:

```
RP/0/RSP0/CPU0:Router# sh route ipv6 fc00:cc30:600:e004:: detail
Wed Nov 10 18:57:56.645 UTC

Routing entry for fc00:cc30:600::/48
  Known via "isis 2", distance 115, metric 141, SRv6-locator, type level-2
  Installed Nov 2 18:56:55.718 for 00:01:01
  Routing Descriptor Blocks
    fe80::232:17ff:fec3:58c0, from 7511::1, via TenGigE0/0/0/16.1, Protected
    Route metric is 141
    Label: None
    Tunnel ID: None
    Binding Label: None
```

```

Extended communities count: 0
Path id:1 Path ref count:0
NHID:0x20006(Ref:193)
Backup path id:65
fe80::226:80ff:fe36:7c01, from 7511::1, via TenGigE1/0/9/1.1, Backup (TI-LFA)
Repair Node(s): 3888::1
Route metric is 251
Label: None
Tunnel ID: None
Binding Label: None
Extended communities count: 0
Path id:65 Path ref count:1
NHID:0x20007(Ref:163)
SRv6 Headend:H.Insert.Red [f3216], SID-list {fc00:cc30:700::}
Route version is 0x0 (8)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-Class: Not Set
Route Priority:RIB_PRIORITY_NON_RECURSIVE_LOW (8) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 2, Download Version 261731
No advertising protos.

```

VRF-to-VRF route leaking in SRv6 core

VRF-to-VRF route leaking is an SRv6 core feature that allows communication between separate VRFs by sharing specific routes while keeping others isolated. This is achieved by configuring import and export route targets in each VRF, ensuring that only the selected routes are exchanged.

Table 18: Feature History Table

Feature Name	Release Information	Feature Description
VRF-to-VRF route leaking in SRv6 core	Release 24.4.1	VRF-to-VRF route leaking enables sharing of routes between VRFs while maintaining their isolation. This feature allows the source VRF to send leaked routes to remote PEs or Route Reflectors (RRs) across an SRv6 core network, similar to an MPLS core network, enabling communication between different service tenants or administrative domains without compromising VRF isolation.

Route leaking between VRFs in an SRv6 core network extends the MPLS-based VRF route leaking functionalities to SRv6, by taking advantage of the SRv6 flexibility for path selection and optimization. You can configure the destination VRF to send leaked routes to a remote PE or RR across an SRv6 core. The feature enables VRF-to-VRF communication in an SRv6-enabled environment while maintaining control over routing and traffic engineering decisions.

Benefits of VRF-to-VRF route leaking in SRv6

The key benefits of the feature are:

- **Improved Traffic Management:** The feature allows routes in the source VRF to use SRv6 SIDs from a best-effort locator and routes in the destination VRF to use SIDs from a low-latency locator. This setup enables differentiated traffic treatment in the SRv6 core.
- **Enhanced Flexibility:** The feature leaks routes between VRFs and advertises them as VPNv4 or VPNv6 or EVPN RT5 prefixes to remote PE routers, providing better flexibility in managing network traffic and inter-VRF communication.
- **Scalability:** The feature dynamically leaks routes that help to scale the network by automating the redistribution process between VRFs.
- **Security and Isolation:** The feature uses route targets and policies to control route leaking, ensuring that it only occurs between intended VRFs, maintaining both security and isolation.

Usage guidelines and limitations for VRF-to-VRF route leaking in SRv6

The usage guidelines and limitations that are listed apply:

- VRF-to-VRF route leaking does not support multicast routes.
- The feature supports both SRv6 Full-length SID and Micro-SIDs.
- Depending on the destination VRF configuration, the PE router assigns SIDs to the leaked route based on the SID allocation mode, which can be per-VRF or per-VRF-46.

Configure VRF-to-VRF route leaking in SRv6

Before you begin

Enable Segment Routing over IPv6 under BGP in the source and destination VRFs.

Procedure

Step 1 Run the **export route-policy** command to configure and attach route leaking in the source VRF.

- Configure the static export Route Target to leak all prefixes to the destination VRF. In the below configuration, the leaked Route Target is 1:12.

```
Router(config)#vrf vrf-be
Router(config-vrf)#address-family ipv4 unicast
Router(config-vrf-af)#import route-target 1:10
Router(config-vrf-af)#export route-target 1:10
Router(config-vrf-af)#export route-target 1:12
Router(config-vrf-af)#commit
```

- Configure a route policy that attaches appropriate Route Target to the leaked prefixes to leak specific prefixes to the destination VRF. Apply the IF condition in the Route Policy Language (RPL) to leak specific prefixes.

```

Router(config)#prefix-set allowed-leaked-route
Router(config-pfx)#192.168.1.0/32
Router(config-pfx)#end-set
Router(config)#route-policy export-policy
Router(config-rpl)#if destination in allowed-leaked-route then
Router(config-rpl-if)#set extcommunity rt 1:12
Router(config-rpl-if)#endif
Router(config-rpl)#end-policy
Router(config)#commit
Router(config)#vrf vrf-be
Router(config-vrf)#address-family ipv4 unicast
Router(config-vrf-af)#import route-target 1:10
Router(config-vrf-af)#export route-target 1:10
Router(config-vrf-af)#export route-policy export-policy
Router(config-vrf-af)#commit

```

Step 2 Configure the destination VRF to import routes from the source VRF.

Example:

```

Router(config)#vrf vrf-ef
Router(config-vrf)#address-family ipv4 unicast
Router(config-vrf-af)#import route-target 2:2
Router(config-vrf-af)#import route-target 1:12
Router(config-vrf-af)#export route-target 2:2

```

Step 3 Run the **show running-config** command to verify the running configuration.

Example:

```

vrf vrf-be
  address-family ipv4 unicast
    import route-target 1:10
    export route-target 1:10
    export route-target 1:12
  !
!
vrf vrf-ef
  address-family ipv4 unicast
    import route-target 2:2
    import route-target 1:12
    export route-target 2:2
  !
!

```

Step 4 Run the **import from vrf advertise-as-vpn** command to forward the imported routes to a remote PE or VPN RR peer through configuration.

Example:

```

Router(config)#vrf vrf-ef
Router(config-vrf)#address-family ipv4 unicast
Router(config-vrf-af)#import from vrf advertise-as-vpn
Router(config-vrf-af)#commit

```

Step 5 Run the **show bgp vrf** command to verify the route leaking from the source VRF vrf-be.

In the below show output, the source VRF vrf-be leaks the Route Target 1:12.

Example:

```

Router#show bgp vrf vrf-be 192.168.1.0/32 detail
Mon Aug 19 14:06:22.668 UTC
BGP routing table entry for 192.168.1.0/32, Route Distinguisher: 1.1.1.1:11

```

```

Versions:
  Process          bRIB/RIB    SendTblVer
  Speaker          3434714    3434714
    SRv6-VPN SID: fc00:1:4:fff0:7d1::/80
    Format: f3216
    Alloc Mode/Locator ID: per-vrf/1
    Flags: 0x00143001+0x01000000+0x00000000
Last Modified: Aug 19 09:53:33.351 for 04:12:49
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Flags: 0x3000000005040003+0x00, import: 0x31f
  Advertised to update-groups (with more than one peer):
    0.2
4
  100.4.0.1 from 100.4.0.1 (193.0.0.1), if-handle 0x00000000
    Origin IGP, localpref 100, valid, external, best, group-best, import-candidate
    Received Path ID 0, Local Path ID 1, version 3434714
    Extended community: RT:1:10 RT:1:12

```

The described show output indicates that the destination VRF vrf-ef imports the prefix from the source VRF vrf-be, as shown by the **imported** flag. The output also includes details of the source VRF. A non-zero value in the **Flags** field confirms that the prefix is imported.

Example:

```

Router#show bgp vrf vrf-ef 192.168.1.0/24 detail
Mon Aug 19 14:08:07.102 UTC
BGP routing table entry for 192.168.1.0/24, Route Distinguisher: 1.1.1.1:21
Versions:
  Process          bRIB/RIB    SendTblVer
  Speaker          3440133    3440133
    SRv6-VPN SID: fc00:2:4:fff0:7d1::/80
    Format: f3216
    Alloc Mode/Locator ID: per-vrf/2
    Flags: 0x00103001+0x01000000+0x00000000
Last Modified: Aug 19 10:48:24.351 for 03:19:42
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
Flags: 0x3100000005040003+0x00, import: 0x080
  Advertised to update-groups (with more than one peer):
    0.2
4
  100.4.0.1 from 100.4.0.1 (193.0.0.1), if-handle 0x00000000
    Origin IGP, localpref 100, valid, external, best, group-best, import-candidate,
imported
    Received Path ID 0, Local Path ID 1, version 3440133
    Extended community: RT:1:10 RT:1:12 RT:1:20
    Source AFI: VPNv4 Unicast, Source VRF: vrf-be, Source Route Distinguisher: 1.1.1.1:11

```

The below show output is an example of the output on remote PE:

```

Router#show bgp vpnv4 unicast rd 1.1.1.1:21 192.168.1.0/32 detail
Fri Dec 13 13:21:29.136 PST
BGP routing table entry for 192.168.1.0/32, Route Distinguisher: 1.1.1.1:21
Versions:
  Process          bRIB/RIB    SendTblVer
  Speaker          690        690
    Flags: 0x00040028+0x00000000+0x00000000
Last Modified: Dec 13 13:14:37.000 for 00:06:52

```



```

Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x2000000025060005+0x00, import: 0x31f
  Not advertised to any peer
  10
    192::1 (metric 30) from 192::4 (192.168.0.1), if-handle 0x00000000
      Received Label 0x7d10
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate, not-in-vrf
      Received Path ID 1, Local Path ID 1, version 690
      Extended community:RT:2:2
      Originator: 192.168.0.1, Cluster list: 192.168.1.4
      PSID-Type:L3, SubTLV Count:1, R:0x00,
      SubTLV:
        T:1(Sid information), Sid:fcc00:2:4:fff0::(Transposed), F:0x00, R2:0x00, Behavior:61,
R3:0x00, SS-TLV Count:1
        SubSubTLV:
          T:1(Sid structure):
            Length [Loc-blk,Loc-node,Func,Arg]:[32,16,32,0], Tpose-len:16, Tpose-offset:64

```



CHAPTER 4


Configure Segment Routing Global Block and Segment Routing Local Block

Local label allocation is managed by the label switching database (LSD). The Segment Routing Global Block (SRGB) and Segment Routing Local Block (SRLB) are label values preserved for segment routing in the LSD.

- [About the Segment Routing Global Block, on page 137](#)
- [About the Segment Routing Local Block, on page 139](#)
- [Understanding Segment Routing Label Allocation, on page 140](#)
- [Setup a Non-Default Segment Routing Global Block Range, on page 143](#)
- [Setup a Non-Default Segment Routing Local Block Range, on page 144](#)

About the Segment Routing Global Block

The Segment Routing Global Block (SRGB) is a range of labels reserved for Segment Routing global segments. A prefix-SID is advertised as a domain-wide unique index. The prefix-SID index points to a unique label within the SRGB range. The index is zero-based, meaning that the first index is 0. The MPLS label assigned to a prefix is derived from the Prefix-SID index plus the SRGB base. For example, considering an SRGB range of 16,000 to 23,999, a prefix 10.1.1.65/32 with prefix-SID index of **65** is assigned the label value of **16065**.



To keep the configuration simple and straightforward, we strongly recommended that you use a homogenous SRGB (meaning, the same SRGB range across all nodes). Using a heterogenous SRGB (meaning, a different SRGB range of the same size across nodes) is also supported but is not recommended.

Behaviors and Limitations

- The default SRGB in IOS XR has a size of 8000 starting from label value 16000. The default range is 16000 to 23,999. With this size, and assuming one loopback prefix per router, an operator can assign prefix SIDs to a network with 8000 routers.
- There are instances when you might need to define a different SRGB range. For example:
 - Non-IOS XR nodes with a SRGB range that is different than the default IOS XR SRGB range.
 - The default SRGB range is not large enough to accommodate all required prefix SIDs.
- A non-default SRGB can be configured following these guidelines:
 - The SRGB starting value can be configured anywhere in the dynamic label range space (16,000 to 1,048,575).
 - In Cisco IOS XR release earlier than 6.6.3, the SRGB can have a maximum configurable size of 262,143.
 - In Cisco IOS XR release 6.6.3 and later, the SRGB can be configured to any size value that fits within the dynamic label range space.
- Allocating an SRGB label range does not mean that all the labels in this range are programmed in the forwarding table. The label range is just reserved for SR and not available for other purposes. Furthermore, a platform may limit the number of local labels that can be programmed.
- We recommend that the non-default SRGB be configured under the **segment-routing** global configuration mode. By default, all IGP instances and BGP use this SRGB.

- You can also configure a non-default SRGB under the IGP, but it is not recommended.

SRGB Label Conflicts

When you define a non-default SRGB range, there might be a label conflict (for example, if labels are already allocated, statically or dynamically, in the new SRGB range). The following system log message indicates a label conflict:

```
%ROUTING-ISIS-4-SRGB_ALLOC_FAIL : SRGB allocation failed: 'SRGB reservation not
successful for [16000,80000], SRGB (16000 80000, SRGB_ALLOC_CONFIG_PENDING, 0x2)
(So far 16 attempts). Make sure label range is free'
```

To remove this conflict, you must reload the router to release the currently allocated labels and to allocate the new SRGB.

After the system reloads, LSD does not accept any dynamic label allocation before IS-IS/OSPF/BGP have registered with LSD. Upon IS-IS/OSPF/BGP registration, LSD allocates the requested SRGB (either the default range or the customized range).

After IS-IS/OSPF/BGP have registered and their SRGB is allocated, LSD starts serving dynamic label requests from other clients.



Note To avoid a potential router reload due to label conflicts, and assuming that the default SRGB size is large enough, we recommend that you use the default IOS XR SRGB range.



Note Allocating a non-default SRGB in the upper part of the MPLS label space increases the chance that the labels are available and a reload can be avoided.



Caution Modifying a SRGB configuration is disruptive for traffic and may require a reboot if the new SRGB is not available entirely.

About the Segment Routing Local Block

A local segment is automatically assigned an MPLS label from the dynamic label range. In most cases, such as TI-LFA backup paths and SR-TE explicit paths defined with IP addresses, this dynamic label allocation is sufficient. However, in some scenarios, it could be beneficial to allocate manually local segment label values to maintain label persistency. For example, an SR-TE policy with a manual binding SID that is performing traffic steering based on incoming label traffic with the binding SID.

The Segment Routing Local Block (SRLB) is a range of label values preserved for the manual allocation of local segments, such as adjacency segment identifiers (adj-SIDs), Layer 2 adj-SIDs, binding SIDs (BSIDs), and BGP peering SIDs. These labels are locally significant and are only valid on the nodes that allocate the labels.

Behaviors and Limitations

- The default SRLB has a size of 1000 starting from label value 15000; therefore, the default SRLB range goes from 15000 to 15,999.
- A non-default SRLB can be configured following these guidelines:
 - The SRLB starting value can be configured anywhere in the dynamic label range space (16,000 to 1,048,575).
 - In Cisco IOS XR release earlier than 6.6.3, the SRLB can have a maximum configurable size of 262,143.
 - In Cisco IOS XR release 6.6.3 and later, the SRLB can be configured to any size value that fits within the dynamic label range space.

SRLB Label Conflicts

When you define a non-default SRLB range, there might be a label conflict (for example, if labels are already allocated, statically or dynamically, in the new SRLB range). In this case, the new SRLB range will be accepted, but not applied (pending state). The previous SRLB range (active) will continue to be in use.

To remove this conflict, you must reload the router to release the currently allocated labels and to allocate the new SRLB.



Caution

You can use the **clear segment-routing local-block discrepancy all** command to clear label conflicts. However, using this command is disruptive for traffic since it forces all other MPLS applications with conflicting labels to allocate new labels.



Note

To avoid a potential router reload due to label conflicts, and assuming that the default SRGB size is large enough, we recommend that you use the default IOS XR SRLB range.



Note

Allocating a non-default SRLB in the upper part of the MPLS label space increases the chance that the labels are available and a reload can be avoided.

Understanding Segment Routing Label Allocation

In IOS XR, local label allocation is managed by the Label Switching Database (LSD). MPLS applications must register as a client with the LSD to allocate labels. Most MPLS applications (for example: LDP, RSVP, L2VPN, BGP [LU, VPN], IS-IS and OSPF [Adj-SID], SR-TE [Binding-SID]) use labels allocated dynamically by LSD.

With Segment Routing-capable IOS XR software releases, the LSD *preserves* the default SRLB label range (15,000 to 15,999) and default SRGB label range (16,000 to 23,999), even if Segment Routing is not enabled.

This preservation of the default SRLB/SRGB label range makes future Segment Routing activation possible without a reboot. No labels are allocated from this preserved range. When you enable Segment Routing with the default SRLB/SRGB in the future, these label ranges will be available and ready for use.


The LSD allocates dynamic labels starting from 24,000.



Note If an MPLS label range is configured and it overlaps with the default SRLB/SRGB label ranges (for example, **mpls label range 15000 1048575**), then the default SRLB/SRGB preservation is disabled.

Example 1: LSD Label Allocation When SR is not Configured

- Special use: 0-15
- MPLS static: 16 to 14,999
- SRLB (preserved): 15,000 to 15,999
- SRGB (preserved): 16,000 to 23,999
- Dynamic: 24,000 to max



Example 2: LSD Label Allocation When SR is Configured with Default SRGB and Default SRLB

- Special use: 0-15
- MPLS static: 16 to 14,999
- SRLB (reserved): 15,000 to 15,999
- SRGB (reserved): 16,000 to 23,999
- Dynamic: 24,000 to max

0	Special-purpose and MPLS static labels
...	
14,999	
15,000 - 15,999	Reserved range (SRLB)
16,000	Reserved range (SRGB)
...	
23,999	
24,000	Dynamic label range
...	
...	
1,048,575	

521228

Example 3: LSD Label Allocation When SR is Configured with Non-default SRGB and Non-default SRLB

- Special use: 0-15
- MPLS static: 16 to 14,999
- SRLB (preserved): 15,000 to 15,999
- SRGB (preserved): 16,000 to 23,999
- Dynamic: 24000 to 28,999
- SRLB (reserved): 29,000 to 29,999
- SRGB (reserved): 30,000 to 39,999
- Dynamic: 40,000 to max

0	
...	Special-purpose and MPLS static labels
14,999	
15,000 - 15,999	Preserved range (SRLB)
16,000	
...	Preserved range (SRGB)
23,999	
24,000	
...	Dynamic label range
28,999	
29,000 - 29,999	Reserved range (SRLB)
30,000	
...	Reserved range (SRGB)
39,999	
40,000	
...	Dynamic label range
...	
1,048,575	521,229

Setup a Non-Default Segment Routing Global Block Range

This task explains how to configure a non-default SRGB range.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	segment-routing global-block <i>starting_value</i> <i>ending_value</i> Example: RP/0/RP0/CPU0:router(config)# segment-routing global-block 16000 80000	Enter the lowest value that you want the SRGB range to include as the starting value. Enter the highest value that you want the SRGB range to include as the ending value.
Step 3	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session.

	Command or Action	Purpose
		• No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

Use the **show mpls label table [label label-value]** command to verify the SRGB configuration:

```
Router# show mpls label table label 16000 detail
Table Label      Owner                               State Rewrite
-----
0      16000      ISIS (A) :1                               InUse  No
      (Lbl-blk SRGB, vers:0, (start_label=16000, size=64001))
```

What to do next

Configure prefix SIDs and enable segment routing.

Setup a Non-Default Segment Routing Local Block Range

This task explains how to configure a non-default SRLB range.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	segment-routing local-block starting_value ending_value Example: RP/0/RP0/CPU0:router(config)# segment-routing local-block 30000 30999	Enter the lowest value that you want the SRLB range to include as the starting value. Enter the highest value that you want the SRLB range to include as the ending value.
Step 3	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session.

	Command or Action	Purpose
		• No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

Use the **show mpls label table [label label-value] [detail]** command to verify the SRLB configuration:

```
Router# show mpls label table label 30000 detail

Table Label   Owner                               State Rewrite
-----
0      30000   LSD(A)                               InUse  No
(Lbl-blk SRLB, vers:0, (start_label=30000, size=1000, app_notify=0)
```

```
Router# show segment-routing local-block inconsistencies
```

No inconsistencies

The following example shows an SRLB label conflict in the range of 30000 and 30999. Note that the default SRLB is active and the configured SRLB is pending:

```
Router(config)# segment-routing local-block 30000 30999

%ROUTING-MPLS_LSD-3-ERR_SRLB_RANGE : SRLB allocation failed: 'SRLB reservation not successful
for [30000,30999]. Use with caution 'clear segment-routing local-block discrepancy all'
command
to force srlb allocation'
```



Caution You can use the **clear segment-routing local-block discrepancy all** command to clear label conflicts. However, using this command is disruptive for traffic since it forces all other MPLS applications with conflicting labels to allocate new labels.

```
Router# show mpls label table label 30000 detail

Table Label   Owner                               State Rewrite
-----
0      30000   LSD(A)                               InUse  No
(Lbl-blk SRLB, vers:0, (start_label=30000, size=1000, app_notify=0)

Router# show segment-routing local-block inconsistencies
SRLB inconsistencies range: Start/End: 30000/30999

Router# show mpls lsd private | i SRLB

SRLB Lbl Mgr:
Current Active SRLB block      = [15000, 15999]
Configured Pending SRLB block = [30000, 30999]
```

Reload the router to release the currently allocated labels and to allocate the new SRLB:

```
Router# reload
Proceed with reload? [confirm]yes
```

After the system is brought back up, verify that there are no label conflicts with the SRLB configuration:

```
Router# show mpls lsd private | i SRLB

SRLB Lbl Mgr:
  Current Active SRLB block      = [30000, 30999]
  Configured Pending SRLB block = [0, 0]

Router# show segment-routing local-block inconsistencies

No inconsistencies
```

What to do next

Configure adjacency SIDs and enable segment routing.



CHAPTER 5

Configure Segment Routing for IS-IS Protocol

Integrated Intermediate System-to-Intermediate System (IS-IS), Internet Protocol Version 4 (IPv4), is a standards-based Interior Gateway Protocol (IGP). The Cisco IOS XR software implements the IP routing capabilities described in International Organization for Standardization (ISO)/International Engineering Consortium (IEC) 10589 and RFC 1995, and adds the standard extensions for single topology and multitopology IS-IS for IP Version 6 (IPv6).

This module provides the configuration information used to enable segment routing for IS-IS.

- [Enabling Segment Routing for IS-IS Protocol, on page 147](#)
- [Configuring a Prefix-SID on the IS-IS Enabled Loopback Interface, on page 149](#)
- [Weighted Anycast SID-Aware Path Computation, on page 152](#)
- [Configuring an Adjacency SID, on page 157](#)
- [Configuring Bandwidth-Based Local UCMP, on page 163](#)
- [IS-IS Unreachable Prefix Announcement, on page 164](#)
- [IS-IS Partition Detection and Leakage of Specific Route Advertisement, on page 167](#)
- [Conditional Prefix Advertisement, on page 171](#)
- [Segment Routing ECMP-FEC Optimization, on page 173](#)

Enabling Segment Routing for IS-IS Protocol

Segment routing on the IS-IS control plane supports the following:

- IPv4 and IPv6 control plane
- Level 1, level 2, and multi-level routing
- Prefix SIDs for host prefixes on loopback interfaces
- Adjacency SIDs for adjacencies
- MPLS penultimate hop popping (PHP) and explicit-null signaling

This task explains how to enable segment routing for IS-IS.

Before you begin

Your network must support the MPLS Cisco IOS XR software feature before you enable segment routing for IS-IS on your router.



Note You must enter the commands in the following task list on every IS-IS router in the traffic-engineered portion of your network.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router isis instance-id Example: RP/0/RP0/CPU0:router(config)# router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. Note You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.
Step 3	address-family { ipv4 ipv6 } [unicast] Example: RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 4	metric-style wide [level { 1 2 }] Example: RP/0/RP0/CPU0:router(config-isis-af)# metric-style wide level 1	Configures a router to generate and accept only wide link metrics in the Level 1 area.
Step 5	router-id loopback loopback interface used for prefix-sid Example: RP/0/RP0/CPU0:router(config-isis-af)# router-id loopback0	Configures router ID for each address-family (IPv4/IPv6). IS-IS advertises the router ID in TLVs 134 (for IPv4 address family) and 140 (for IPv6 address family). Required when traffic engineering is used.
Step 6	segment-routing mpls [sr-prefer] Example: RP/0/RP0/CPU0:router(config-isis-af)# segment-routing mpls	Segment routing is enabled by the following actions: • MPLS forwarding is enabled on all interfaces where IS-IS is active. • All known prefix-SIDs in the forwarding plain are programmed, with the

	Command or Action	Purpose
		prefix-SIDs advertised by remote routers or learned through local or remote mapping server. The prefix-SIDs locally configured are advertised. Use the sr-prefer keyword to set the preference of segment routing (SR) labels over label distribution protocol (LDP) labels.
Step 7	exit Example: RP/0/RP0/CPU0:router(config-isis-af)# exit RP/0/RP0/CPU0:router(config-isis)# exit	
Step 8	Use the commit or end command.	commit—Saves the configuration changes and remains within the configuration session. end—Prompts user to take one of these actions: Yes — Saves configuration changes and exits the configuration session. No —Exits the configuration session without committing the configuration changes. Cancel —Remains in the configuration session, without committing the configuration changes.

What to do next

Configure the prefix SID.

Configuring a Prefix-SID on the IS-IS Enabled Loopback Interface

A prefix segment identifier (SID) is associated with an IP prefix. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels. A prefix SID is configured under the loopback interface with the loopback address of the node as the prefix. The prefix segment steers the traffic along the shortest path to its destination.

A prefix SID can be a node SID or an Anycast SID. A node SID is a type of prefix SID that identifies a specific node. An Anycast SID is a type of prefix SID that identifies a set of nodes, and is configured with n-flag clear. The set of nodes (Anycast group) is configured to advertise a shared prefix address and prefix SID. Anycast

routing enables the steering of traffic toward multiple advertising nodes. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.

The prefix SID is globally unique within the segment routing domain.

This task explains how to configure prefix segment identifier (SID) index or absolute value on the IS-IS enabled Loopback interface.

Before you begin

Ensure that segment routing is enabled on the corresponding address family.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router isis instance-id Example: RP/0/RP0/CPU0:router(config)# router isis 1	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.
Step 3	interface Loopback instance Example: RP/0/RP0/CPU0:router(config-isis)# interface Loopback0	Specifies the loopback interface and instance.
Step 4	address-family { ipv4 ipv6 } [unicast] Example: The following is an example for ipv4 address family: RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 5	prefix-sid [algorithm algorithm-number] {index SID-index absolute SID-value} [n-flag-clear] [explicit-null] Example: RP/0/RP0/CPU0:router(config-isis-if-af)# prefix-sid index 1001	Configures the prefix-SID index or absolute value for the interface. Specify algorithm algorithm-number to configure SR Flexible Algorithm. Specify index SID-index for each node to create a prefix SID based on the lower boundary of the SRGB + the index.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config-isis-if-af)# prefix-sid absolute 17001	Specify absolute SID-value for each node to create a specific prefix SID within the SRGB. By default, the n-flag is set on the prefix-SID, indicating that it is a node SID. For specific prefix-SID (for example, Anycast prefix-SID), enter the n-flag-clear keyword. IS-IS does not set the N flag in the prefix-SID sub Type Length Value (TLV). To disable penultimate-hop-popping (PHP) and add explicit-Null label, enter explicit-null keyword. IS-IS sets the E flag in the prefix-SID sub TLV.
Step 6	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

Verify the prefix-SID configuration:

```
RP/0/RP0/CPU0:router# show isis database verbose
```

```
IS-IS 1 (Level-2) Link State Database
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
router.00-00         * 0x0000039b  0xfc27        1079          0/0/0
  Area Address: 49.0001
  NLPID:             0xcc
  NLPID:             0x8e
  MT:                 Standard (IPv4 Unicast)
  MT:                 IPv6 Unicast                                0/0/0
  Hostname:          router
  IP Address:         10.0.0.1
  IPv6 Address:       2001:0db8:1234::0a00:0001
  Router Cap:         10.0.0.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
    SR Algorithm:
      Algorithm: 0

<...>
Metric: 0             IP-Extended 10.0.0.1/32
  Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0

<...>
```


Weighted Anycast SID-Aware Path Computation

Table 19: Feature History Table

Feature Name	Release Information	Feature Description
Weighted Anycast SID-Aware Path Computation	Release 7.3.1	This feature extends Anycast SIDs with weighted nodes. Weighted Anycast nodes advertise a cost (weight) along with the Anycast SID. Traffic is then distributed according to the weights. Weighted Anycast SIDs allow for highly available paths with node redundancy and path optimality that provide Fast ReRoute (FRR) for node failure of service provider edge (PE) routers and ABR/ASBRs nodes in multi-domain networks.

The Weighted Anycast SID feature extends Anycast SIDs with weighted nodes.

Anycast routing enables the steering of traffic toward multiple advertising nodes, providing load-balancing and redundancy. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes. With the default (unweighted) behavior, the traffic is load-balanced across each node in the group evenly.

Weighted Anycast nodes advertise a cost along with the Anycast SID. This cost serves as a weight. Traffic to the SID is then distributed according to the weights.


Weighted Anycast SIDs allow for highly available paths with node redundancy and path optimality that provide FRR for node failure of service provider edge (PE) routers and ABR/ASBR nodes in multi-domain networks.

In addition, Weighted Anycast SIDs allow for scaled computation at the PCE of multi-domain paths.

The native SR path computation algorithms are augmented to compute optimum paths relying on Weighted Anycast SIDs during path encoding.

Consider the example depicted below. Nodes A and B are part of the same Anycast groups, represented by different SIDs (100, 200, 300).

- SID 100 sends traffic preferentially to node A
- SID 200 sends traffic preferentially to node B
- SID 300 sends traffic equally to both nodes




The Anycast replacement algorithm runs after an SR-TE path has been computed. It examines the prefix SIDs in the path and swaps them with Anycast SIDs that contain the same node. The new paths are checked against the original constraints and kept if suitable.

If a node is part of multiple Anycast groups, the algorithm considers them according to their weights.

Example

The following figure shows 3 isolated IGP domains without redistribution and without BGP 3107. Each Area Border Router (ABR) 1 through 4 is configured with a node SID. The link delays are also shown.



ABRs 1 and 2 share the following Anycast SIDs:

- 16012 – sends traffic to either Node 1 or 2 (the topologically nearest node)
- 17012 – sends traffic preferentially to Node 1


- 18012 – sends traffic preferentially to Node 2

ABRs 3 and 4 share the following Anycast SIDs:

- 16034 – sends traffic either Node 3 or 4 (the topologically nearest node)
- 17034 – sends traffic preferentially to Node 3
- 18034 – sends traffic preferentially to Node 4


Consider the case where routers A and Z are provider edge (PE) routers in the same VPN. Router A receives a VPN route with BGP next-hop to router Z. Router A resolves the SR path to router Z using SR-ODN or SR-PCE.

Before considering Anycast SIDs, the head-end router or SR-PCE computes the SID list.



In this case, the optimized computed path from router A to router Z is $16002 > 16003 > 1600Z$.

Using the weighted Anycast-encoded SID list, the optimized computed path from router A to router Z is $18012 > 17034 > 1600Z$. This path has a cumulative delay of 31.



521469

Using node SIDs, failures inside each domain (for example, links) benefit from fast TI-LFA convergence. However, failures of the ABR nodes would be dependent on SR-PCE reoptimization.

Using weighted Anycast SIDs, failures of the ABR nodes and failures inside each domain benefit from fast TI-LFA convergence.

Configuration

Based on the topology in Figure NN, this example shows the Weighted Anycast SID configuration of ABRs 1 and 2.

ABR 1 Configuration

```
RP/0/RSP0/CPU0:ios(config)# router isis 1
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback0
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 16001
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback1
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 16012
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback2
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# metric 1
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 17012
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback3
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# metric 100000
```

```
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 18012
```

Running Config

```
router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001 // Node SID
  !
 !
 interface Loopback1
  prefix-attributes anycast
  address-family ipv4 unicast
   prefix-sid absolute 16012 //Anycast SID - (prefer node 1 or 2)
  !
 !
 interface Loopback2
  prefix-attributes anycast
  address-family ipv4 unicast
   metric 1
   prefix-sid absolute 17012 // Weighted Anycast SID (prefer node 1)
  !
 !
 interface Loopback3
  prefix-attributes anycast
  address-family ipv4 unicast
   metric 100000
   prefix-sid absolute 18012 // Weighted Anycast SID (prefer node 2)
  !
 !
end
```

ABR 2 Configuration

```
RP/0/RSP0/CPU0:ios(config)# router isis 1
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback0
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 16002
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback1
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 16012
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback2
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# metric 100000
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 17012
RP/0/RSP0/CPU0:ios(config-isis-if-af)# exit
RP/0/RSP0/CPU0:ios(config-isis-if)# exit
RP/0/RSP0/CPU0:ios(config-isis)# interface Loopback3
RP/0/RSP0/CPU0:ios(config-isis-if)# prefix-attributes anycast
RP/0/RSP0/CPU0:ios(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:ios(config-isis-if-af)# metric 1
RP/0/RSP0/CPU0:ios(config-isis-if-af)# prefix-sid absolute 18012
```

Running Config

```

router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16002 // Node SID
  !
 !
 interface Loopback1
  prefix-attributes anycast
  address-family ipv4 unicast
   prefix-sid absolute 16012 // Anycast SID (prefer any)
  !
 !
 interface Loopback2
  prefix-attributes anycast
  address-family ipv4 unicast
   metric 100000
   prefix-sid absolute 17012 // Weighted Anycast SID (prefer node 1)
  !
 !
 interface Loopback3
  prefix-attributes anycast
  address-family ipv4 unicast
   metric 1
   prefix-sid absolute 18012 // Weighted Anycast SID (prefer node 2)
  !
 !
end

```

Configuring an Adjacency SID

An adjacency SID (Adj-SID) is associated with an adjacency to a neighboring node. The adjacency SID steers the traffic to a specific adjacency. Adjacency SIDs have local significance and are only valid on the node that allocates them.

An adjacency SID can be allocated dynamically from the dynamic label range or configured manually from the segment routing local block (SRLB) range of labels.

Adjacency SIDs that are dynamically allocated do not require any special configuration, however there are some limitations:

- A dynamically allocated Adj-SID value is not known until it has been allocated, and a controller will not know the Adj-SID value until the information is flooded by the IGP.
- Dynamically allocated Adj-SIDs are not persistent and can be reallocated after a reload or a process restart.
- Each link is allocated a unique Adj-SID, so the same Adj-SID cannot be shared by multiple links.

Manually allocated Adj-SIDs are persistent over reloads and restarts. They can be provisioned for multiple adjacencies to the same neighbor or to different neighbors. You can specify that the Adj-SID is protected. If the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.

Adjacency SIDs are advertised using the existing IS-IS Adj-SID sub-TLV. The S and P flags are defined for manually allocated Adj-SIDs.

0 1 2 3 4 5 6 7

```

+---+---+---+---+---+---+
| F | B | V | L | S | P |   |
+---+---+---+---+---+---+

```

Table 20: Adjacency Segment Identifier (Adj-SID) Flags Sub-TLV Fields

Field	Description
S (Set)	This flag is set if the same Adj-SID value has been provisioned on multiple interfaces.
P (Persistent)	This flag is set if the Adj-SID is persistent (manually allocated).

Manually allocated Adj-SIDs are supported on point-to-point (P2P) interfaces.

This task explains how to configure an Adj-SID on an interface.

Before you begin

Ensure that segment routing is enabled on the corresponding address family.

Use the **show mpls label table detail** command to verify the SRLB range.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router isis instance-id Example: RP/0/RP0/CPU0:router(config)# router isis 1	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.
Step 3	interface type interface-path-id Example: RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet0/0/0/7	Specifies the interface and enters interface configuration mode.
Step 4	point-to-point Example: RP/0/RP0/CPU0:router(config-isis-if)# point-to-point	Specifies the interface is a point-to-point interface.

	Command or Action	Purpose
Step 5	address-family { ipv4 ipv6 } [unicast] Example: The following is an example for ipv4 address family: <pre>RP/0/RP0/CPU0:router(config-isis-if) # address-family ipv4 unicast</pre>	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 6	adjacency-sid {index adj-SID-index absolute adj-SID-value } [protected] Example: <pre>RP/0/RP0/CPU0:router(config-isis-if-af) # adjacency-sid index 10 RP/0/RP0/CPU0:router(config-isis-if-af) # adjacency-sid absolute 15010</pre>	Configures the Adj-SID index or absolute value for the interface. Specify index adj-SID-index for each link to create an Adj-SID based on the lower boundary of the SRLB + the index. Specify absolute adj-SID-value for each link to create a specific Adj-SID within the SRLB. Specify if the Adj-SID is protected . For each primary path, if the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.
Step 7	Use the commit or end command.	commit — Saves the configuration changes and remains within the configuration session. end — Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No — Exits the configuration session without committing the configuration changes. • Cancel — Remains in the configuration session, without committing the configuration changes.

Verify the Adj-SID configuration:

```
RP/0/RP0/CPU0:router# show isis segment-routing label adjacency persistent
Mon Jun 12 02:44:07.085 PDT
```

IS-IS 1 Manual Adjacency SID Table

```
15010 AF IPv4
    GigabitEthernet0/0/0/3: IPv4, Protected 1/65/N, Active
    GigabitEthernet0/0/0/7: IPv4, Protected 2/66/N, Active

15100 AF IPv6
    GigabitEthernet0/0/0/3: IPv6, Not protected 255/255/N, Active
```


Verify the labels are added to the MPLS Forwarding Information Base (LFIB):

```
RP/0/RP0/CPU0:router# show mpls forwarding labels 15010
Mon Jun 12 02:50:12.172 PDT
Local   Outgoing   Prefix      Outgoing    Next Hop    Bytes
Label   Label      or ID       Interface   Interface   Switched
-----
15010   Pop        SRLB (idx 10)  Gi0/0/0/3   10.0.3.3    0
        Pop        SRLB (idx 10)  Gi0/0/0/7   10.1.0.5    0
        16004     SRLB (idx 10)  Gi0/0/0/7   10.1.0.5    0           (!)
        16004     SRLB (idx 10)  Gi0/0/0/3   10.0.3.3    0           (!)
```

Manually Configure a Layer 2 Adjacency SID

Typically, an adjacency SID (Adj-SID) is associated with a Layer 3 adjacency to a neighboring node, to steer the traffic to a specific adjacency. If you have Layer 3 bundle interfaces, where multiple physical interfaces form a bundle interface, the individual Layer 2 bundle members are not visible to IGP; only the bundle interface is visible.

You can configure a Layer 2 Adj-SID for the individual Layer 2 bundle interfaces. This configuration allows you to track the availability of individual bundle member links and to verify the segment routing forwarding over the individual bundle member links, for Operational Administration and Maintenance (OAM) purposes.

A Layer 2 Adj-SID can be allocated dynamically or configured manually.

- IGP dynamically allocates Layer 2 Adj-SIDs from the dynamic label range for each Layer 2 bundle member. A dynamic Layer 2 Adj-SID is not persistent and can be reallocated as the Layer 3 bundle link goes up and down.
- Manually configured Layer 2 Adj-SIDs are persistent if the Layer 3 bundle link goes up and down. Layer 2 Adj-SIDs are allocated from the Segment Routing Local Block (SRLB) range of labels. However, if the configured value of Layer 2 Adj-SID does not fall within the available SRLB, a Layer 2 Adj-SID will not be programmed into forwarding information base (FIB).

Restrictions

- Adj-SID forwarding requires a next-hop, which can be either an IPv4 address or an IPv6 address, but not both. Therefore, manually configured Layer 2 Adj-SIDs are configured per address-family.
- Manually configured Layer 2 Adj-SID can be associated with only one Layer 2 bundle member link.
- A SID value used for Layer 2 Adj-SID cannot be shared with Layer 3 Adj-SID.
- SR-TE using Layer 2 Adj-SID is not supported.

This task explains how to configure a Layer 2 Adj-SID on an interface.

Before you begin

Ensure that segment routing is enabled on the corresponding address family.

Use the **show mpls label table detail** command to verify the SRLB range.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	segment-routing Example: RP/0/RP0/CPU0:Router(config)# segment-routing	Enters segment routing configuration mode.
Step 3	adjacency-sid Example: RP/0/RP0/CPU0:Router(config-sr)# adjacency-sid	Enters adjacency SID configuration mode.
Step 4	interface type interface-path-id Example: RP/0/RP0/CPU0:Router(config-sr-adj)# interface GigabitEthernet0/0/0/3	Specifies the interface and enters interface configuration mode.
Step 5	address-family { ipv4 ipv6 } [unicast] Example: RP/0/RP0/CPU0:Router(config-sr-adj-intf)# address-family ipv4 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 6	l2-adjacency sid {index adj-SID-index absolute adj-SID-value } [next-hop {ipv4_address ipv6_address }] Example: RP/0/RP0/CPU0:Router(config-sr-adj-intf-af)# l2-adjacency sid absolute 15015 next-hop 10.1.1.4	Configures the Adj-SID index or absolute value for the interface. Specify index adj-SID-index for each link to create an Adj-SID based on the lower boundary of the SRLB + the index. Specify absolute adj-SID-value for each link to create a specific Adj-SID within the SRLB. For point-to-point interfaces, you are not required to specify a next-hop. However, if you do specify the next-hop, the Layer 2 Adj-SID will be used only if the specified next-hop matches the neighbor address. For LAN interfaces, you must configure the next-hop IPv4 or IPv6 address. If you do not configure the next-hop, the Layer 2 Adj-SID will not be used for LAN interface.

	Command or Action	Purpose
Step 7	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.
Step 8	end	
Step 9	router isis <i>instance-id</i> Example: RP/0/RP0/CPU0:Router(config) # router isis isp	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.
Step 10	address-family { ipv4 ipv6 } [unicast] Example: RP/0/RP0/CPU0:Router(config-isis) # address-family ipv4 unicast	Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.
Step 11	segment-routing bundle-member-adj-sid Example: RP/0/RP0/CPU0:Router(config-isis-af) # segment-routing bundle-member-adj-sid	Programs the dynamic Layer 2 Adj-SIDs, and advertises both manual and dynamic Layer 2 Adj-SIDs. Note This command is not required to program manual L2 Adj-SID, but is required to program the dynamic Layer 2 Adj-SIDs and to advertise both manual and dynamic Layer 2 Adj-SIDs.

Verify the configuration:

```
Router# show mpls forwarding detail | i "Pop|Outgoing Interface|Physical Interface"
Tue Jun 20 06:53:51.876 PDT
. . .
15001 Pop          SRLB (idx 1)      BE1          10.1.1.4        0
    Outgoing Interface: Bundle-Ether1 (ifhandle 0x000000b0)
    Physical Interface: GigabitEthernet0/0/0/3 (ifhandle 0x000000b0)
```

```

Router# show running-config segment-routing
Tue Jun 20 07:14:25.815 PDT
segment-routing
adjacency-sid
interface GigabitEthernet0/0/0/3
address-family ipv4 unicast
  12-adjacency-sid absolute 15015 next-hop 10.1.1.4
!
!
!
!

```

Configuring Bandwidth-Based Local UCMP

Bandwidth-based local Unequal Cost Multipath (UCMP) allows you to enable UCMP functionality locally between Equal Cost Multipath (ECMP) paths based on the bandwidth of the local links.

Bandwidth-based local UCMP is performed for prefixes, segment routing Adjacency SIDs, and Segment Routing label cross-connects installed by IS-IS, and is supported on any physical or virtual interface that has a valid bandwidth.

For example, if the capacity of a bundle interface changes due to the link or line card up/down event, traffic continues to use the affected bundle interface regardless of the available provisioned bundle members. If some bundle members were not available due to the failure, this behavior could cause the traffic to overload the bundle interface. To address the bundle capacity changes, bandwidth-based local UCMP uses the bandwidth of the local links to load balance traffic when bundle capacity changes.

Before you begin

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router isis <i>instance-id</i> Example: RP/0/RP0/CPU0:router(config)# router isis 1	Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode. You can change the level of routing to be performed by a particular routing instance by using the is-type router configuration command.
Step 3	address-family { ipv4 ipv6 } [unicast] Example: The following is an example for ipv4 address family:	Specifies the IPv4 or IPv6 address family, and enters IS-IS address family configuration mode.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast	
Step 4	apply-weight ecmp-only bandwidth Example: RP/0/RP0/CPU0:router(config-isis-af)# apply-weight ecmp-only bandwidth	Enables UCMP functionality locally between ECMP paths based on the bandwidth of the local links.
Step 5	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

IS-IS Unreachable Prefix Announcement

Table 21: Feature History Table

Feature Name	Release	Description
IS-IS Unreachable Prefix Announcement	Release 7.8.1	The Unreachable Prefix Announcement (UPA) notifies the loss of prefix reachability between areas or domains, for prefixes that are covered by the summary address range during inter-area or inter-domain summarization. This feature helps in identifying the routers that are facing prefix unreachability issues faster and fix it. The new commands introduced for this feature are: • summary-prefix • prefix-unreachable

The organization of networks into levels or areas and/or IGP domains helps to limit the scope of link-state information within certain boundaries. However, the state that is related to prefix reachability often requires propagation across these areas (Level1/Level2) or domains (Autonomous System Boundary Router (ASBR)). An Autonomous System Boundary Router (ASBR) is a router that is running multiple protocols and serves

as a gateway to routers outside the Open Shortest Path First (OSPF) domain and those operating with different protocols.

Route summarization, also known as route aggregation, is a method to minimize the number of routing tables in an IP network. It consolidates selected multiple routes into a single route advertisement.

Techniques such as summarization address the scale challenges associated with the advertisement of the individual prefix state outside of local area/domain. MPLS architecture did not allow for the effective use of the summarization due to its end-to-end Label Switched Path (LSP) requirement. With the introduction of the SRv6, which does not have such requirement, the use of summarization has become important again.

Summarization results in suppression of the individual prefix state that is useful for triggering fast-convergence mechanisms outside of the Interior Gateway Routing Protocols (IGPs (for example - Border Gateway Protocol - Prefix Independent Convergence (BGP PIC) Edge).

This feature enables the notification of the individual prefixes becoming unreachable in its area/domain, when the summarization is used between areas/domains to advertise the reachability for these prefixes.

There are existing SRv6 deployments that use summarization and require fast detection of the egress Provider Edge (PE) going down. To address these deployments in timely manner, we use the existing Protocol Data Units (PDUs) and Tag-Length-Values (TLVs), which is based on the Prefix Unreachability Advertisement (UPA).

Configuration Steps

The configuration steps that are required to set up the Unreachable Prefix Announcement (UPA) feature are as follows:

- **UPA Advertisement**

An existing IS-IS address-family submode **summary-prefix** command was extended for UPA advertisement.

```
Router(config)#router isis 1
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#summary-prefix beef:10::/32 level 2 adv-unreachable
Router(config-isis-af)#summary-prefix beef:11::/32 level 2 algorithm 128 adv-unreachable
  unreachable-component-tag 777
Router(config-isis-af)#commit
```

- **Prefix Unreachable**

The new **prefix-unreachable** command includes new commands that control the UPA advertisement such as, lifetime, metric, limit the maximum number if UPAs and UPA processing. For more details see, [prefix-unreachable](#)

```
Router(config)#router isis 1
Router(config-isis)#address-family ipv6
Router(config-isis-af)#prefix-unreachable
Router(config-isis-prefix-unreachable)#adv-lifetime 500
Router(config-isis-prefix-unreachable)#adv-metric 4261412866
Router(config-isis-prefix-unreachable)#adv-maximum 77
Router(config-isis-prefix-unreachable)#rx-process-enable
Router(config-isis-prefix-unreachable)#commit
```

Running Configuration

Execute the following show commands to review the L1/L2 (area) or ASBR (domain) running configuration:

Run the **show run router isis 1 address-family ipv6 unicast** command to view the summary prefix under as well as UPA parameters under it.

```
Router#sh run router isis 1 address-family ipv6 unicast
router isis 1
address-family ipv6 unicast
  advertise application lfa link-attributes srlg
  advertise link attributes
  prefix-unreachable
    adv-lifetime 300
  !
  summary-prefix 10::/64
  summary-prefix beef:10::/32 adv-unreachable
  summary-prefix beef:11::/32 algorithm 128 adv-unreachable
  summary-prefix ceef:10::/32 adv-unreachable
  propagate level 2 into level 1 route-policy L2_TO_L1
  segment-routing srv6
    locator USID_ALG0
  !
    locator USID_ALG128
  !
!
!
```


IS-IS Partition Detection and Leakage of Specific Route Advertisement

Table 22: Feature History Table

Feature Name	Release	Description
IS-IS Partition Detection and Leakage of Specific Route Advertisements	Release 7.10.1	In an open ring topology, a single fiber cut may partition the area or domain into two pieces. With summarization enabled, the area (domain) partition may result in traffic drops. Depending on the configuration in the Area Border Routers (ABRs) or Autonomous System Boundary Routers (ASBRs) that is picked as an entry point to the partitioned area (domain), the traffic is delivered to its destination or dropped as unreachable at ABR or ASBR. IS-IS partition detection and leakage of specific route advertisements features are introduced to retain connectivity for the partitioned area (domain) when summarization is used. The ABRs or ASBRs detect a network partition within an area (domain) and upon detection, ensure that the summary route is replaced with specific route advertisements in IS-IS. The feature introduces these changes: New Command: • partition-detect Modified Command: • The partition-repair keyword is introduced in the summary-prefix command. YANG Data Model: • New XPath for <code>Cisco-IOS-XR-um-router-isis-cfg.yang</code> (see GitHub, YANG Data Models Navigator)

In service provider networks, the Layer 1 (L1) area is often represented by a set of routers connected in a ring. Sometimes the ring is not closed (for example, Area Border Routers (ABRs) are not connected directly inside the L1 area). In such cases, a single fiber cut partitions the area into two pieces. Route Summarization is basically advertising many routes into one route, also called route aggregation. When the partition is detected, summarization is suppressed, and all previously summarized prefixes are advertised in IS-IS. The individual prefix advertisements preserve the connectivity end-to-end.

Figure 10: Interarea Topology with L1 Area in an Open Ring



This feature addresses the following summarization problems:

- **Area Partition Detection and Avoidance:** Area partition detection only works for summarization from L1 to L2. It is not supported for summarization from L2 to L1. Ensure you have router-id that is configured for the address-family (IPv4 or IPv6) for which you are enabling the partition detection and avoidance.



Note

Router-id must be enabled for this feature to work. Area partition is tracked for each algorithm, algo 0, and any enabled flex-algo, independently. For flex-algo, the Area Border Router must participate in the algo for the tracking to work for such flex-algo

- **Domain Partition Detection and Avoidance:** Networks use multi-domain design, where they split their network into multiple IGP domains. They redistribute between domains and summarize during the redistribution. An IGP domain may represent an open ring and a single link cut may split it into two parts.

Configuration Steps

Configure IS-IS Partition Detection and Leakage of Specific as follows:

- **Area Partition Detection and Avoidance:** To configure the area partition detection and avoidance there are two configuration steps:

1. For each summary prefix that you want the area partition and avoidance to work, enable it with the summary-prefix command:

```
Router(config)#router isis 1
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#commit
```

2. Configure the tracking of the Area Border Router (ABR) reachability. Here, the example is for two ABRs, but you can enable the partition for as many ABRs in the area.

```
Router(config)#router isis 1
Router(config-isis)#address-family ipv6 unicast
```

```

Router(config-isis-af)#router-id 2001:DB8:1::1
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#partition-detect
Router(config-isis-af)#track 2001:DB8:4::4
Router(config-isis-af)#commit

Router(config)#router isis 1
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#router-id 2001:DB8:4::4
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#partition-detect
Router(config-isis-af)#track 2001:DB8:1::1
Router(config-isis-af)#commit

```

- **Domain Partition Detection and Avoidance:** It is similar to Area Partition and requires two configuration steps:

1. For each summary prefix that you want the area partition and avoidance to work, enable it with the summary-prefix command:

```

Router(config)#router isis 1
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#commit

```

2. Configure the tracking of the Area Border Router reachability. The example is for two ABRs, but you can enable the partition for as many ABRs in the area.

To track ASBR, two identifiers are required:

- First is the internal router-id in the instance under which the configuration is done (similar to area partition).
- Second is the address of the ASBR in the other domain that is redistributed to the instance where the configuration is done.



Note

The implementation of the feature ensures that the reachability of the external-address is only tracked in algorithm 0. The loss of the external-address is used for algo 0 and all flex-algos. The reachability of the internal address is kept per flex-algo.

```

Router(config)#router isis 2
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#router-id 2001:DB8:1::1
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#partition-detect

```

```

Router(config-isis-af)#track 2001:DB8:4::4 external-address 2001:DB8:10::4
Router(config-isis-af)#commit

Router(config)#router isis 2
Router(config-isis)#address-family ipv6 unicast
Router(config-isis-af)#router-id 2001:DB8:4::4
Router(config-isis-af)#summary-prefix 2001:DB8::/32 level 2 partition-repair
Router(config-isis-af)#summary-prefix 2001:DB9::/32 level 2 algorithm 128
partition-repair
Router(config-isis-af)#partition-detect
Router(config-isis-af)#track 2001:DB8:1::1 external-address 2001:DB8:10::1
Router(config-isis-af)#commit

```

Verification

• For Area Partition Detection and Avoidance

Use the show command **show isis instance 1 flex-algo 128** to check if the area partition is detected:

```

Router# show isis instance 1 flex-algo 128
IS-IS 1 Flex-Algo Database
Flex-Algo 128:
Level-2:
Definition Priority: 128
Definition Source: plzen.00, (Local)
Definition Equal to Local: Yes
Definition Metric Type: IGP
Definition Flex-Algo Prefix Metric: No
Exclude Any Affinity Bit Positions:
Include Any Affinity Bit Positions:
Include All Affinity Bit Positions:
Reverse Exclude Any Affinity Bit Positions:
Reverse Include Any Affinity Bit Positions:
Reverse Include All Affinity Bit Positions:
Exclude SRLGs:
Disabled: No

Level-1:
Definition Priority: 128
Definition Source: plzen.00, (Local)
Definition Equal to Local: Yes
Definition Metric Type: IGP
Definition Flex-Algo Prefix Metric: No
Exclude Any Affinity Bit Positions:
Include Any Affinity Bit Positions:
Include All Affinity Bit Positions:
Reverse Exclude Any Affinity Bit Positions:
Reverse Include Any Affinity Bit Positions:
Reverse Include All Affinity Bit Positions:
Exclude SRLGs:
Disabled: No

Topologies supported:
IPv4 Unicast
Partition-Detect:
ABR: Internal-Address: 10.4.4.4
ASBR: Internal-Address: 10.4.4.4 External-Address: 10.10.10.4
ABR: Internal-Address: 10.5.5.5
IPv6 Unicast
Partition-Detect:
ABR: Internal-Address: 2001:DB8:4::4

```

```

Local Priority: 128
FRR Disabled: No
Microloop Avoidance Disabled: No
Data Plane Segment Routing: Yes
Data Plane IP: No

```

• Domain Partition Detection and Avoidance

Use the show command **show isis instance 1 flex-algo 128** to check if the domain partition is detected:

```

Router# show isis instance 1 flex-algo 128
IS-IS 1 Flex-Algo Database
Flex-Algo 128:

Level-2:
Definition Priority: 128
Definition Source: plzen.00, (Local)
Definition Equal to Local: Yes
Definition Metric Type: IGP
Definition Flex-Algo Prefix Metric: No
Exclude Any Affinity Bit Positions:
Include Any Affinity Bit Positions:
Include All Affinity Bit Positions:
Reverse Exclude Any Affinity Bit Positions:
Reverse Include Any Affinity Bit Positions:
Reverse Include All Affinity Bit Positions:
Exclude SRLGs:
Disabled: No

Level-1:
Definition Priority: 128
Definition Source: plzen.00, (Local)
Definition Equal to Local: Yes
Definition Metric Type: IGP
Definition Flex-Algo Prefix Metric: No
Exclude Any Affinity Bit Positions:
Include Any Affinity Bit Positions:
Include All Affinity Bit Positions:
Reverse Exclude Any Affinity Bit Positions:
Reverse Include Any Affinity Bit Positions:
Reverse Include All Affinity Bit Positions:
Exclude SRLGs:
Disabled: No

Topologies supported:
IPv4 Unicast
Partition-Detect:
ABR: Internal-Address: 10.4.4.4 (Active)
ASBR: Internal-Address: 10.4.4.4 External-Address: 10.10.10.4 (Active)
ABR: Internal-Address: 10.5.5.5
IPv6 Unicast
Partition-Detect:
ABR: Internal-Address: 2001:DB8:4::4 (Active)
ASBR: Internal-Address: 2001:DB8:4::4 External-Address: 2001:DB8:10::4 (Active)

```

Conditional Prefix Advertisement

In some situations, it's beneficial to make the IS-IS prefix advertisement conditional. For example, an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) that has lost its connection to one of the areas or autonomous systems (AS) might keep advertising a prefix. If an ABR or ASBR advertises the

Segment Routing (SR) SID with this prefix, the label stack of the traffic routed toward the disconnected area or AS might use this SID, which would result in dropped traffic at the ABR or ASBR.

ABRs or ASBRs are often deployed in pairs for redundancy and advertise a shared Anycast prefix SID. Conditional Prefix Advertisement allows an ABR or an ASBR to advertise its Anycast SID only when connected to a specific area or domain. If an ABR or ASBR becomes disconnected from the particular area or AS, it stops advertising the address for a specified interface (for example, Loopback).

Configure the conditional prefix advertisement under a specific interface. The prefix advertisement on this interface is associated with the route-policy that tracks the presence of a set of prefixes (prefix-set) in the Routing Information Base (RIB).

For faster convergence, the route-policy used for conditional prefix advertisement uses the new event-based **rib-has-route async** condition to notify IS-IS of the following situations:

- When the last prefix from the prefix-set is removed from the RIB.
- When the first prefix from the prefix-set is added to the RIB.

Configuration

To use the conditional prefix advertisement in IS-IS, create a prefix-set to be tracked. Then create a route policy that uses the prefix-set.

```
Router(config)# prefix-set prefix-set-name
Router(config-pfx)# prefix-address-1/length[, prefix-address-2/length,,
prefix-address-16/length]
Router(config-pfx)# end-set

Router(config)# route-policy rpl-name
Router(config-rpl)# if rib-has-route async prefix-set-name then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
```

To advertise the loopback address in IS-IS conditionally, use the **advertise prefix route-policy** command under IS-IS interface address-family configuration sub-mode.

```
Router(config)# router isis 1
Router(config-isis)# interface Loopback0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# advertise prefix route-policy rpl-name
Router(config-isis-if-af)# commit
```

Example

```
Router(config)# prefix-set domain_2
Router(config-pfx)# 2.3.3.3/32, 2.4.4.4/32
Router(config-pfx)# end-set
Router(config)# route-policy track_domain_2
Router(config-rpl)# if rib-has-route async domain_2 then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
Router(config)# router isis 1
Router(config-isis)# interface Loopback0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# advertise prefix route-policy track_domain_2
```

```
Router(config-isis-if-af)# commit
```

Running Configuration

```
prefix-set domain_2
  2.3.3.3/32,
  2.4.4.4/32
end-set
!
route-policy track_domain_2
  if rib-has-route async domain_2 then
    pass
  endif
end-policy
!
router isis 1
  interface Loopback0
    address-family ipv4 unicast
    advertise prefix route-policy track_domain_2
  !
!
```

Segment Routing ECMP-FEC Optimization

Table 23: Feature History Table

Feature Name	Release Information	Feature Description
SR ECMP-FEC Optimization for IPv6 Prefixes	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.

Feature Name	Release Information	Feature Description
SR ECMP-FEC Optimization for IPv6 Prefixes	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS You can now minimize ECMP-FEC resource consumption and duplication of resources during underlay programming for an SR-MPLS network for IPv6 prefixes. The feature supports sharing the same ECMP-FEC, regular FEC, and Egress Encapsulation DB (EEDB) entries for all paths with the same outgoing MPLS label among IPv6 prefixes. This improvement is now possible by using the hw-module fib mpls label lsr-optimized command. In earlier releases, ECMP-FEC optimization for IPv4 prefixes was available.
SR ECMP-FEC Optimization L2 and L3 Recursive Services	Release 7.4.1	This feature adds support for L2VPN service Label Edge Router (LER) and BGP PIC for Layer 3 BGP services when SR ECMP-FEC Optimization is enabled.

ECMP-FECs are used for any ECMP programming on the system, such as MPLS LSP ECMP, VPN multipath, and [EVPN multi-homing](#) EVPN multi-homing.

The SR ECMP-FEC optimization solution minimizes ECMP-FEC resource consumption and duplication, during underlay programming for an SR-MPLS network. This feature supports sharing the same ECMP-FEC, regular FEC, and Egress Encapsulation DB (EEDB) entries among Segment Routing prefixes with the same set of next hops.

ECMP-FEC optimization is triggered when all the out_labels associated with the ECMP paths for a given prefix have the same value. If this rule is not met, then the prefix is programmed with a dedicated ECMP-FEC.

Segment Routing Label Edge Router (LER) ECMP-FEC Optimization enables ECMP-FEC optimization originally developed for Label Switched Router (LSR) nodes (MPLS P) to be enabled on LER (Layer 3 MPLS PE) routers.

Usage Guidelines and Limitations

- SR ECMP-FEC Optimization is not supported on Cisco NCS 5500 series routers that have the Cisco NC57 line cards installed and operating in the native or compatible modes.
- For the labeled prefixes with ECMP across a combination of labeled and unlabeled (PHP) paths, the SR ECMP-FEC Optimization cannot be triggered since the paths associated with the prefix do not have the same outgoing label and/or label action.
- For prefixes with LFA backup paths, the SR ECMP-FEC Optimization is possible since these backup paths do not require an extra label to be pushed; all paths associated with the prefix (primary and backup) have the same outgoing label value.
- For prefixes with TI-LFA backup paths requiring extra labels to be pushed on to the backup, the SR ECMP-FEC Optimization is not possible since all the paths associated with the prefix do not have the same outgoing label value.
- For the duration of time that prefixes are programmed to avoid microloops (when SR MicroLoop Avoidance is triggered), SR ECMP-FEC Optimization is not possible since all the paths associated with the prefix do not have the same outgoing label value. After removal of the microloop-avoidance programming, the SR ECMP-FEC Optimization might be possible again.
- For scenarios with prefixes where the SR ECMP-FEC Optimization is not possible, dedicated ECMP-FEC is allocated per prefix. This could potentially lead to ECMP FEC out-of-resource (OOR) considering the baseline usage of ECMP FEC resources at steady state. During ECMP-FEC OOR, prefixes with multiple paths are programmed with a single path in order to avoid traffic disruption.
- SR ECMP-FEC optimization is applicable in the following instances:
 - Label Switched Router (LSR) nodes (MPLS P)
 - L3VPN Label Edge Router (LER) nodes
 - L2VPN LER nodes
 - ASBR node with BGP-LU swap
- BGP PIC is supported
- SR ECMP-FEC optimization should not be enabled in the following instances:
 - L2VPN/L3VPN LER nodes with VPN over BGP-LU over SR
- For the labeled prefixes, transitioning from TI-LFA to SR ECMP-FEC optimization can cause ECMP-FEC OOR due to different output labels (ECMP label vs backup path's label) at make-before-break. This results in a few second traffic loss depending on route scale

Enable SR ECMP-FEC Optimization

To enable SR ECMP-FEC optimization, use the **hw-module fib mpls label lsr-optimized** command in global configuration mode. After enabling this feature, reload the line card. For more information about the command, see the *MPLS Label Distribution Protocol Commands* chapter in the *MPLS Command Reference Guide*.

```
Router(config)# hw-module fib mpls label lsr-optimized
Router(config)# commit
```

```
LC/0/0/CPU0:Oct 11 20:19:12.540 UTC: fia_driver[185]:
%FABRIC-FIA_DRV-4-MPLS_HW_PROFILE_MISMATCH :
```



```
Mismatch found, reload LC to activate the new mpls profile
```

```
Router# reload location 0/0/CPU0
```

```
Proceed with reload? [confirm]
```

```
Reloading node 0/0/CPU0
```

Verify the SR ECMP-FEC Optimization

From Cisco IOS XR Release 24.2.1, in addition to IPv4 prefixes, you can view the ECMP-FEC optimization details for IPv6 prefixes using the **show controllers npu resources** command. In earlier releases, ECMP-FEC optimization for IPv4 prefixes was available.

The following example shows NPU usage before enabling SR ECMP-FEC optimization for IPv4 and IPv6 prefixes:

```
Router#show controllers npu resources ecmpfec location 0/0/CPU0
```

```
Tue May 14 17:17:40.306 UTC
```

```
HW Resource Information
```

```
Name : ecmp_fec
Asic Type : JerichoPlus
```

```
NPU-0
```

```
OOR Summary
```

```
Estimated Max Entries : 4096
Red Threshold : 95 %
Yellow Threshold : 80 %
OOR State : Green
High Water Mark : 216
High Water Mark Time : 2024.May.14 17:10:46 UTC
Bank Info : ECMP
```

```
OFA Table Information
```

```
(May not match HW usage)
```

```
ipnhgroup : 111
ip6nhgroup : 105
```

```
Current Hardware Usage
```

```
Name: ecmp_fec
```

```
Estimated Max Entries : 4096
Total In-Use : 216 (5 %)
OOR State : Green
High Water Mark : 216
High Water Mark Time : 2024.May.14 17:10:46 UTC
Bank Info : ECMP
```

```
Name: hier_0
```

```
Estimated Max Entries : 4096
Total In-Use : 216
OOR State : Green
High Water Mark : 216
High Water Mark Time : 2024.May.14 17:10:46 UTC
Bank Info : ECMP
```

The following example shows NPU usage after enabling SR ECMP-FEC optimization for IPv4 and IPv6 prefixes. After enabling ECMP-FEC optimization, the overall hardware usage has dropped to 16%, and the group of next-hop addresses used for forwarding IPv4 and IPv6 is significantly reduced.

```
Router#show controllers npu resources ecmpfec location 0/0/CPU0
```

```
Tue May 14 17:07:56.266 UTC
```

```
HW Resource Information
```

```

  Name           : ecmp_fec
  Asic Type      : JerichoPlus

```

```
NPU-0
```

```
  OOR Summary
```

```

    Estimated Max Entries : 4096
    Red Threshold         : 95 %
    Yellow Threshold      : 80 %
    OOR State             : Green
    High Water Mark       : 16
    High Water Mark Time  : 2024.May.14 17:04:47 UTC
    Bank Info             : ECMP

```

```
OFA Table Information
```

```
(May not match HW usage)
```

```

  ipnhgroup      : 11
  ip6nhgroup     : 5

```

```
Current Hardware Usage
```

```

  Name: ecmp_fec
    Estimated Max Entries : 4096
    Total In-Use         : 16      (0 %)
    OOR State             : Green
    High Water Mark       : 16
    High Water Mark Time  : 2024.May.14 17:04:47 UTC
    Bank Info             : ECMP

```

```
  Name: hier_0
```

```

    Estimated Max Entries : 4096
    Total In-Use          : 16
    OOR State             : Green
    High Water Mark       : 16
    High Water Mark Time  : 2024.May.14 17:04:47 UTC
    Bank Info             : ECMP

```




CHAPTER 6

Configure Segment Routing for OSPF Protocol

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) developed by the OSPF working group of the Internet Engineering Task Force (IETF). Designed expressly for IP networks, OSPF supports IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets.

This module provides the configuration information to enable segment routing for OSPF.



Note For additional information on implementing OSPF on your , see the *Implementing OSPF* module in the .

- [Enabling Segment Routing for OSPF Protocol, on page 179](#)
- [Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface, on page 181](#)
- [Configuring an Adjacency SID, on page 183](#)
- [Segment Routing ECMP-FEC Optimization, on page 186](#)
- [Multi-area Loopback Interface Support for OSPF, on page 186](#)
- [Display Neighbor Router Capabilities in OSPF Networks, on page 190](#)

Enabling Segment Routing for OSPF Protocol

Segment routing on the OSPF control plane supports the following:

- OSPFv2 control plane
- Multi-area
- IPv4 prefix SIDs for host prefixes on loopback interfaces
- Adjacency SIDs for adjacencies
- MPLS penultimate hop popping (PHP) and explicit-null signaling

This section describes how to enable segment routing MPLS and MPLS forwarding in OSPF. Segment routing can be configured at the instance, area, or interface level.

Before you begin

Your network must support the MPLS Cisco IOS XR software feature before you enable segment routing for OSPF on your router.



Note You must enter the commands in the following task list on every OSPF router in the traffic-engineered portion of your network.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router ospf process-name Example: RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process and places the router in router configuration mode.
Step 3	segment-routing mpls Example: RP/0/RP0/CPU0:router(config-ospf)# segment-routing mpls	Enables segment routing using the MPLS data plane on the routing process and all areas and interfaces in the routing process. Enables segment routing forwarding on all interfaces in the routing process and installs the SIDs received by OSPF in the forwarding table.
Step 4	segment-routing sr-prefer Example: RP/0/RP0/CPU0:router(config-ospf)# segment-routing sr-prefer	Sets the preference of segment routing (SR) labels over label distribution protocol (LDP) labels.
Step 5	area area Example: RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode.
Step 6	segment-routing mpls Example: RP/0/RP0/CPU0:router(config-ospf-ar)# segment-routing mpls	(Optional) Enables segment routing using the MPLS data plane on the area and all interfaces in the area. Enables segment routing forwarding on all interfaces in the area and installs the SIDs received by OSPF in the forwarding table.
Step 7	exit Example:	

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config-ospf-ar) # exit RP/0/RP0/CPU0:router(config-ospf) # exit	
Step 8	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

What to do next

Configure the prefix SID.

Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface

A prefix segment identifier (SID) is associated with an IP prefix. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels. A prefix SID is configured under the loopback interface with the loopback address of the node as the prefix. The prefix segment steers the traffic along the shortest path to its destination.

A prefix SID can be a node SID or an Anycast SID. A node SID is a type of prefix SID that identifies a specific node. An Anycast SID is a type of prefix SID that identifies a set of nodes, and is configured with n-flag clear. The set of nodes (Anycast group) is configured to advertise a shared prefix address and prefix SID. Anycast routing enables the steering of traffic toward multiple advertising nodes. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.

The prefix SID is globally unique within the segment routing domain.

This task describes how to configure prefix segment identifier (SID) index or absolute value on the OSPF-enabled Loopback interface.

Before you begin

Ensure that segment routing is enabled on an instance, area, or interface.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router ospf <i>process-name</i> Example: RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing process, and places the router in router configuration mode.
Step 3	area <i>value</i> Example: RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode.
Step 4	interface Loopback <i>interface-instance</i> Example: RP/0/RP0/CPU0:router(config-ospf-ar)# interface loopback 0	Specifies the loopback interface and instance.
Step 5	prefix-sid { index <i>SID-index</i> absolute <i>SID-value</i> } [n-flag-clear] [explicit-null] Example: RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid index 1001 RP/0/RP0/CPU0:router(config-ospf-ar)# prefix-sid absolute 17001	Configures the prefix-SID index or absolute value for the interface. Specify index <i>SID-index</i> for each node to create a prefix SID based on the lower boundary of the SRGB + the index. Specify absolute <i>SID-value</i> for each node to create a specific prefix SID within the SRGB. By default, the n-flag is set on the prefix-SID, indicating that it is a node SID. For specific prefix-SID (for example, Anycast prefix-SID), enter the n-flag-clear keyword. OSPF does not set the N flag in the prefix-SID sub Type Length Value (TLV). To disable penultimate-hop-popping (PHP) and add an explicit-Null label, enter the explicit-null keyword. OSPF sets the E flag in the prefix-SID sub TLV.
Step 6	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions:

	Command or Action	Purpose
		• Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

Verify the prefix-SID configuration:

```
RP/0/RP0/CPU0:router# show ospf database opaque-area 7.0.0.1 self-originate
OSPF Router with ID (10.0.0.1) (Process ID 1)
Type-10 Opaque Link Area Link States (Area 0)
<...>
Extended Prefix TLV: Length: 20
Route-type: 1
AF      : 0
Flags   : 0x40
Prefix  : 10.0.0.1/32

SID sub-TLV: Length: 8
Flags   : 0x0
MTID    : 0
Algo    : 0
SID Index : 1001
```

Configuring an Adjacency SID

An adjacency SID (Adj-SID) is associated with an adjacency to a neighboring node. The adjacency SID steers the traffic to a specific adjacency. Adjacency SIDs have local significance and are only valid on the node that allocates them.

An adjacency SID can be allocated dynamically from the dynamic label range or configured manually from the segment routing local block (SRLB) range of labels.

Adjacency SIDs that are dynamically allocated do not require any special configuration, however there are some limitations:

- A dynamically allocated Adj-SID value is not known until it has been allocated, and a controller will not know the Adj-SID value until the information is flooded by the IGP.
- Dynamically allocated Adj-SIDs are not persistent and can be reallocated after a reload or a process restart.
- Each link is allocated a unique Adj-SID, so the same Adj-SID cannot be shared by multiple links.

Manually allocated Adj-SIDs are persistent over reloads and restarts. They can be provisioned for multiple adjacencies to the same neighbor or to different neighbors. You can specify that the Adj-SID is protected. If

the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.

Adjacency SIDs are advertised using the existing OSPF Adj-SID sub-TLV. The P-flag is defined for manually allocated Adj-SIDs.

```

 0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
|B|V|L|G|P|   |
+---+---+---+---+---+---+

```

Table 24: Adjacency Segment Identifier (Adj-SID) Flags Sub-TLV Fields

Field	Description
P (Persistent)	This flag is set if the Adj-SID is persistent (manually allocated).

This task explains how to configure an Adj-SID on an interface.

Before you begin

Ensure that segment routing is enabled on the corresponding address family.

Use the **show mpls label table detail** command to verify the SRLB range.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters mode.
Step 2	router ospf process-name Example: RP/0/RP0/CPU0:router(config)# router ospf 1	Enables OSPF routing for the specified routing instance, and places the router in router configuration mode.
Step 3	area area Example: RP/0/RP0/CPU0:router(config-ospf)# area 0	Enters area configuration mode.
Step 4	interface type interface-path-id Example: RP/0/RP0/CPU0:router(config-ospf-ar)# interface HundredGigE0/0/0/1	Specifies the interface and enters interface configuration mode.

	Command or Action	Purpose
Step 5	adjacency-sid {index <i>adj-SID-index</i> absolute <i>adj-SID-value</i>} [protected] Example: <pre>RP/0/RP0/CPU0:router(config-config-ospf-ar-if)# adjacency-sid index 10</pre> <pre>RP/0/RP0/CPU0:router(config-config-ospf-ar-if)# adjacency-sid absolute 15010</pre>	Configures the Adj-SID index or absolute value for the interface. Specify index <i>adj-SID-index</i> for each link to create an Adj-SID based on the lower boundary of the SRLB + the index. Specify absolute <i>adj-SID-value</i> for each link to create a specific Adj-SID within the SRLB. Specify if the Adj-SID is protected. For each primary path, if the Adj-SID is protected on the primary interface and a backup path is available, a backup path is installed. By default, manual Adj-SIDs are not protected.
Step 6	Use the commit or end command.	commit — Saves the configuration changes and remains within the configuration session. end — Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No — Exits the configuration session without committing the configuration changes. • Cancel — Remains in the configuration session, without committing the configuration changes.

What to do next

Configure the SR-TE policy.

Protected Adjacency SID Backup Timer

OSPF advertises a protected adjacency SID for an adjacency when a backup path is available. Primary and backup paths are programmed into the label switching database (LSD) as rewrites.

When an adjacency goes down, OSPF stops advertising the protected adjacency SID immediately, and the backup path is promoted and installed as LSD rewrite. After a specified amount of time, the LSD rewrite is deleted. If the installed path fails again, the protection ends there and traffic through the original protected adjacency SID is permanently lost.

The Protected Adjacency SID Backup Timer provides a configurable maintenance time period. During this time period, OSPF updates the LSD rewrite with primary and backup (if available) paths to the neighbor upon topology changes.

Configuration

Use the **segment-routing protected-adjacency-sid-delay** command in OSPF configuration mode. The range is from 30 to 3600 seconds; the default is 900 seconds (15 min).

```
Router(config)# router ospf 1
Router(config-ospf)# segment-routing protected-adjacency-sid-delay 360
```

Running Configuration

```
router ospf 1
 segment-routing protected-adjacency-sid-delay 360
 area 1
   interface HundredGigE0/0/0/1
     fast-reroute per-prefix
     fast-reroute per-prefix ti-lfa enable
   !
 !
 !
```

Segment Routing ECMP-FEC Optimization

ECMP-FECs are used for any ECMP programming on the system, such as MPLS LSP ECMP, VPN multipath, and [EVPN multi-homing](#) EVPN multi-homing.

The SR ECMP-FEC optimization solution minimizes ECMP-FEC resource consumption during underlay programming for an SR-MPLS network. This feature supports sharing the same ECMP-FEC, regular FEC, and Egress Encapsulation DB (EEDB) entries for all IPv4 and IPv6 Segment Routing prefixes with the same set of next hops. ECMP-FEC optimization is triggered when all the out_labels associated with the ECMP paths for a given prefix have the same value. If this rule is not met, then the prefix is programmed with a dedicated ECMP-FEC. Other prefixes that meet the rule are candidates for optimization.

Segment Routing Label Edge Router (LER) ECMP-FEC Optimization enables ECMP-FEC optimization originally developed for Label Switched Router (LSR) nodes (MPLS P) to be enabled on LER (Layer 3 MPLS PE) routers.

For usage guidelines, limitations, and configuration options, see [Segment Routing ECMP-FEC Optimization, on page 173](#).

Multi-area Loopback Interface Support for OSPF

Table 25: Feature History Table

Feature Name	Release Information	Feature Description
Multi-area Loopback Interface for OSPF	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.

Feature Name	Release Information	Feature Description
Multi-area Loopback Interface for OSPF	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS You can save IP addresses and resources, prevent the use of multiple node SIDs for labels associated with loopback interfaces, and save time configuring multiple loopback interfaces for an Area Border Router (ABR) in a network. These improvements are possible as you can now configure a single loopback interface for multiple areas. With this feature, an ABR can use a single loopback interface for all areas it connects to, eliminating the need for separate loopback interfaces for each area. Previously, each loopback interface was linked to only one area. The feature introduces these changes: CLI: The multi-area-interface command is extended to support loopback interfaces.

Loopback Interface Optimization for ABRs

With this feature, you can reuse a single loopback interface for multiple areas, with no specific limit on the number of areas that can be configured for an ABR. This means that an ABR requires only one loopback interface for all areas it connects to instead of requiring a separate loopback interface for each area.

In earlier releases, configuring separate loopback interfaces for each area was necessary, resulting in configuring as many loopback interfaces as areas on the ABR.

Key Benefits of Multi-area Loopback Interface

The key benefits of the feature are:

- The feature offers an efficient networking solution, eliminating the need to configure multiple loopback interfaces.

- The feature improves scaling by enabling a 1:n ratio of loopback interfaces to multiple areas.

Usage Guidelines and Limitations for Multi-area Loopback

The following usage guidelines and limitations apply:

- The feature supports only IPv4 addresses.
- Associate the loopback interface to an area of an ABR, and then reuse it as a multi-area interface in other areas.
- The feature is limited to a single OSPF instance and does not extend to multiple OSPF instances.
- The feature is limited to a single VRF and does not extend to multiple VRFs.

Configurations

Configure the loopback interface

Perform the following task to configure the loopback interface IPv4 or IPv6 address, which is the primary loopback address that you want to reuse for all areas:

```
Router(config)#interface Loopback 0
Router(config-if)#ipv4 address 10.10.10.1/8
Router(config-if)#commit
```

Running Configuration

```
!
interface Loopback0
  ipv4 address 10.10.10.1 255.0.0.0
!
```

Associate the loopback interface to an area

Perform the following task to associate the primary loopback interface to area 1:

```
Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 1
Router(config-ospf-ar)#interface Loopback 0
Router(config-ospf-ar-if)#commit
```

Running Configuration

```
!
router ospf 100
  router-id 10.10.10.1
  area 1
    interface Loopback 0
  !
!
```

Reuse the loopback interface

Perform the following task to reuse the primary loopback interface in another OSPF area:

```
Router(config)#router ospf 100
Router(config-ospf)#router-id 10.10.10.1
Router(config-ospf)#area 0
Router(config-ospf-ar)#multi-area-interface Loopback 0
Router(config-ospf-ar-mif)#commit
```

Running Configuration

```
!
router ospf 100
  router-id 10.10.10.1
  area 0
    multi-area-interface Loopback 0
  !
!
```

Verification

The following example shows how the loopback 0 interface is identified by area 1 in an ABR:

```
Router#show ospf routes 10.10.10.1/8 backup-path detail
Thu Mar 21 18:54:55.729 UTC

OSPF Route entry for 10.10.10.1/8
  Route type: Intra-area
  Last updated: Mar 11 02:43:17.876
  Metric: 1
  SPF priority: 8,   SPF version: 38
  RIB version: 0,   Source: Unknown
  Label Info: Default 16004 SSPF 18004 Type SR
               10.10.10.1, directly connected, via Loopback0
               LSA: 1/10.10.10.1/10.10.10.1, Area: 1
```

The following example shows that summary LSA is not advertised by ABR in area 0 after you enable the loopback interface in area 1:

```
Router#show ospf database summary 10.10.10.1
Thu Mar 21 18:30:57.294 UTC

OSPF Router with ID (10.10.10.1) (Process ID 1)

Summary Net Link States (Area 0)

LS age: 1760
Options: (No TOS-capability, DC)
LS Type: Summary Links (Network)
Link State ID: 10.10.10.1 (Summary Network Number)
Advertising Router: 10.0.0.1
LS Seq Number: 8000006e
Checksum: 0x5550
Length: 28
Network Mask: /32
TOS: 0 Metric: 2
```

The following example indicates the router link-state advertisement (LSA) in area 0:

```
Router#show ospf database opaque-area ext-prefix 10.10.10.1/8
Thu Mar 21 18:57:40.914 UTC

OSPF Router with ID (10.10.10.1) (Process ID 1)

Type-10 Opaque Link Area Link States (Area 0)

LS age: 345
```

```
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 7.0.0.3
Opaque Type: 7
Opaque ID: 3
Advertising Router: 10.10.10.1
LS Seq Number: 80000001
Checksum: 0xa3d6
Length: 56
```

```
Extended Prefix TLV: Length: 8
```

```
Route-type: 1
```

```
AF : 0
```

```
Flags : 0x40
```

```
Prefix : 10.10.10.1/32
```

```
SID sub-TLV: Length: 8
```

```
Flags : 0x0
```

```
MTID : 0
```

```
Algo : 0
```

```
SID Index : 4
```

```
SID sub-TLV: Length: 8
```

```
Flags : 0x0
```

```
MTID : 0
```

```
Algo : 1
```

```
SID Index : 2004
```

Display Neighbor Router Capabilities in OSPF Networks

Table 26: Feature History Table

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks	Release 24.2.11	This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.

Feature Name	Release Information	Feature Description
Display Neighbor Router Capabilities in OSPF Networks	Release 24.2.1	Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR: • N540-ACC-SYS • N540X-ACC-SYS • N540-24Z8Q2C-SYS You can now improve the user experience, understand the features and functionalities supported by neighboring routers, reduce operational work hours and human errors by gaining better visibility into the routing capabilities of directly connected neighboring routers in an OSPF network. These improvements are now possible when you enable Segment Routing on a neighboring connected router. Previously, there was no straightforward method to display the segment routing and other capabilities of neighboring routers in an OSPF network. The feature introduces these changes: CLI: • The capabilities keyword is introduced in the show ospf neighbor command. YANG Data Models: • Cisco-IOS-XR-ipv4-ospf-oper.yang See (GitHub, Yang Data Models Navigator)

The feature enables you to access detailed information of a neighboring router capability from the Router Information Link State Advertisement (RI LSA). This includes details such as segment routing feature, graceful restart and helper operation modes, hostname, and traffic engineering capabilities.

For more information about graceful restart and helper operation modes, see *Implementing OSPF* chapter in the *Routing Configuration Guide*.

Key Benefits of Displaying Neighbor Router Capabilities

- The feature ensures optimal path selection for faster packet forwarding.
- Reduces operational complexity, and enables you to make informed decisions for network upgrades.
- Perform compatibility checks with neighboring routers.

Usage Guidelines and Limitations to Display Neighbor Router Capabilities

The following usage guidelines and limitations apply:

- The RI LSA includes the capabilities of a router, only when Segment Routing capability is enabled for that router. This means that if Segment Routing is not enabled, the router does not advertise any capability in the RI LSA.
- The feature is area-oriented, meaning that it is specific to an OSPF area. This suggests that the capabilities advertised by a neighboring router are relevant and visible only within the same OSPF area.
- You must enable traffic engineering at the OSPF or at the area level.
- You must configure graceful restart and helper capabilities at the OSPF level.

Verify the Neighbor Router Capabilities

You can use the **show ospf neighbor capabilities** command to display the traffic engineering support, hostname, and graceful restart and helper capabilities of neighboring connected routers, which are enabled for Segment Routing.

```
Router#show ospf neighbor capabilities
Mon Apr 29 08:27:07.289 UTC
```

Capabilities advertised in RI LSA

```
TE - Traffic Engineering
GR - Graceful Restart
GH - Graceful Restart Helper
```

NA - Capabilities TLV Not Advertised

Neighbors for OSPF 1

Area 1

Neighbor ID	Address	Interface	SR-Enabled	Capabilities
192.168.199.137	209.165.200.225/27	Gi100/0/0/0	Yes	TE, GR, GH
HostName: Device_3				
192.168.48.1	209.165.200.230/27	Gi100/0/0/4	No	NA
192.168.48.200	209.165.200.235/27	Gi10/0/0/0	Yes	TE, GR, GH
HostName: Device_0				

Total neighbor count: 3

In this scenario, Area 1 consists of three connected routers. Routers 192.168.199.137 and 192.168.48.200 support Segment Routing, so their RI LSA includes the TE, GH, and GR capabilities. When the TE and GR are not configured, only the GH capability is advertised by default. Router 192.168.48.1, which does not support Segment Routing, does not include any capabilities in its RI LSA.

When you use the **show ospf neighbor** *neighbor-id* command to view the summary information of a specific neighboring router, the summary information includes the capabilities of that neighbor router if Segment Routing is enabled for it.

The following example shows the capabilities of neighbor 192.168.199.130 listed in the summary information:

```
Router#show ospf neighbor 192.168.199.130
Mon Nov 6 19:13:08.861 IST
* Indicates MADJ interface
# Indicates Neighbor awaiting BFD session up

Neighbors for OSPF 1

Neighbor 192.168.199.130, interface address 172.31.80.37
  In the area 0 via interface GigabitEthernet0/2/0/4
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x52
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:32
  Neighbor is up for 00:08:50
  Number of DBD retrans during last exchange 0
  Index 3/6, retransmission queue length 0, number of retransmission o
  First 0(0) /0(0) Next 0(0) /0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 1 msec, maximum is 8 msec
  LS Ack list: NSR-sync pending 0, high water mark o
  Adjacency SIDs:
    Label: 24010,      Dynamic, Unprotected
    Label: 24011,      Dynamic,  Protected (Has backup)
  Neighbor Interface ID: 4
Segment-Routing Enabled
  Strict SPF capable
  Start SID: 16000 Range Size: 8000
  Flex algo Enabled
Capabilities advertised in RI LSA
  Graceful Restart capable
  Graceful Restart Helper capable
HostName: r5

Total neighbor count: 1
```




CHAPTER 7

Configure Segment Routing for BGP

Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) that allows you to create loop-free inter-domain routing between autonomous systems. An autonomous system is a set of routers under a single technical administration. Routers in an autonomous system can use multiple Interior Gateway Protocols (IGPs) to exchange routing information inside the autonomous system and an EGP to route packets outside the autonomous system.

This module provides the configuration information used to enable Segment Routing for BGP.



Note For additional information on implementing BGP on your router, see the *Implementing BGP* module in the *Routing Configuration Guide*.

- [Segment Routing for BGP, on page 195](#)
- [Configure BGP Prefix Segment Identifiers, on page 196](#)
- [Segment Routing Egress Peer Engineering, on page 197](#)
- [Configure BGP Link-State, on page 202](#)
- [Configurable Filters for IS-IS Advertisements to BGP-Link State, on page 207](#)
- [Use Case: Configuring SR-EPE and BGP-LS, on page 208](#)
- [Configure BGP Proxy Prefix SID, on page 211](#)
- [SRv6 double recursion for multilayer BGP underlay, on page 218](#)

Segment Routing for BGP

In a traditional BGP-based data center (DC) fabric, packets are forwarded hop-by-hop to each node in the autonomous system. Traffic is directed only along the external BGP (eBGP) multipath ECMP. No traffic engineering is possible.

In an MPLS-based DC fabric, the eBGP sessions between the nodes exchange BGP labeled unicast (BGP-LU) network layer reachability information (NLRI). An MPLS-based DC fabric allows any leaf (top-of-rack or border router) in the fabric to communicate with any other leaf using a single label, which results in higher packet forwarding performance and lower encapsulation overhead than traditional BGP-based DC fabric. However, since each label value might be different for each hop, an MPLS-based DC fabric is more difficult to troubleshoot and more complex to configure.

BGP has been extended to carry segment routing prefix-SID index. BGP-LU helps each node learn BGP prefix SIDs of other leaf nodes and can use ECMP between source and destination. Segment routing for BGP

simplifies the configuration, operation, and troubleshooting of the fabric. With segment routing for BGP, you can enable traffic steering capabilities in the data center using a BGP prefix SID.



Note BGP flowspec support with SRv6 - Limitations

List of BGP address families interacts with SRv6. There are some supported and unsupported BGP address family for the interaction with SRv6.

Supported address family:

- address-families ipv6.

Unsupported address families:

- address-families ipv4
 - vpnv4
 - vpnv6
-

Configure BGP Prefix Segment Identifiers

Segments associated with a BGP prefix are known as BGP prefix SIDs. The BGP prefix SID is global within a segment routing or BGP domain. It identifies an instruction to forward the packet over the ECMP-aware best-path computed by BGP to the related prefix. The BGP prefix SID is manually configured from the segment routing global block (SRGB) range of labels.

Each BGP speaker must be configured with an SRGB using the **segment-routing global-block** command. See the [About the Segment Routing Global Block](#) section for information about the SRGB.



Note You must enable SR and explicitly configure the SRGB before configuring SR BGP. The SRGB must be explicitly configured, even if you are using the default range (16000 – 23999). BGP uses the SRGB and the index in the BGP prefix-SID attribute of a learned BGP-LU advertisement to allocate a local label for a given destination.

If SR and the SRGB are enabled after configuring BGP, then BGP is not aware of the SRGB, and therefore it allocates BGP-LU local labels from the dynamic label range instead of from the SRGB. In this case, restart the BGP process in order to allocate BGP-LU local labels from the SRGB.



Note Because the values assigned from the range have domain-wide significance, we recommend that all routers within the domain be configured with the same range of values.

To assign a BGP prefix SID, first create a routing policy using the **set label-index index** attribute, then associate the index to the node.

Example

The following example shows how to configure the SRGB, create a BGP route policy using a \$SID parameter and **set label-index** attribute, and then associate the prefix-SID index to the node.

```
RP/0/RP0/CPU0:router(config)# segment-routing global-block 16000 23999

RP/0/RP0/CPU0:router(config)# route-policy SID($SID)
RP/0/RP0/CPU0:router(config-rpl)# set label-index $SID
RP/0/RP0/CPU0:router(config-rpl)# end policy

RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# bgp router-id 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# network 10.1.1.3/32 route-policy SID(3)
RP/0/RP0/CPU0:router(config-bgp-af)# allocate-label all
RP/0/RP0/CPU0:router(config-bgp-af)# commit
RP/0/RP0/CPU0:router(config-bgp-af)# end

RP/0/RP0/CPU0:router# show bgp 10.1.1.3/32
BGP routing table entry for 10.1.1.3/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker              74         74
Local Label: 16003
Last Modified: Sep 29 19:52:18.155 for 00:07:22
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.2
  3
    99.3.21.3 from 99.3.21.3 (10.1.1.3)
Received Label 3
  Origin IGP, metric 0, localpref 100, valid, external, best, group-best
  Received Path ID 0, Local Path ID 1, version 74
  Origin-AS validity: not-found
Label Index: 3
```

Segment Routing Egress Peer Engineering

Segment routing egress peer engineering (EPE) uses a controller to instruct an ingress provider edge, or a content source (node) within the segment routing domain, to use a specific egress provider edge (node) and a specific external interface to reach a destination. BGP peer SIDs are used to express source-routed inter-domain paths.

Below are the BGP-EPE peering SID types:

- **PeerNode SID**—To an eBGP peer. Pops the label and forwards the traffic on any interface to the peer.
- **PeerAdjacency SID**—To an eBGP peer via interface. Pops the label and forwards the traffic on the related interface.
- **PeerSet SID**—To a set of eBGP peers. Pops the label and forwards the traffic on any interface to the set of peers. All the peers in a set might not be in the same AS.

Multiple PeerSet SIDs can be associated with any combination of PeerNode SIDs or PeerAdjacency SIDs.

The controller learns the BGP peer SIDs and the external topology of the egress border router through BGP-LS EPE routes. The controller can program an ingress node to steer traffic to a destination through the egress node and peer node using BGP labeled unicast (BGP-LU).

EPE functionality is only required at the EPE egress border router and the EPE controller.

Usage Guidelines and Limitations

- When enabling BGP EPE, you must enable MPLS encapsulation on the egress interface connecting to the eBGP peer. This can be done by enabling either BGP labeled unicast (BGP-LU) address family or MPLS static for the eBGP peer.

For information about BGP-LU, refer to the “[Implementing BGP](#)” chapter in the *BGP Configuration Guide*.

For information about MPLS static, refer to the “[Implementing MPLS Static Labeling](#)” chapter in the *MPLS Configuration Guide*.

Configure Segment Routing Egress Peer Engineering

This task explains how to configure segment routing EPE on the EPE egress node.

Procedure

	Command or Action	Purpose
Step 1	router bgp as-number Example: <pre>RP/0/RP0/CPU0:router(config)# router bgp 1</pre>	Specifies the BGP AS number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 2	neighbor ip-address Example: <pre>RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.1.3</pre>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
Step 3	remote-as as-number Example: <pre>RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 3</pre>	Creates a neighbor and assigns a remote autonomous system number to it.

	Command or Action	Purpose
Step 4	egress-engineering Example: <pre>RP/0/RP0/CPU0:router(config-bgp-nbr)# egress-engineering</pre>	Configures the egress node with EPE for the eBGP peer.
Step 5	exit Example: <pre>RP/0/RP0/CPU0:router(config-bgp-nbr)# exit RP/0/RP0/CPU0:router(config-bgp)# exit RP/0/RP0/CPU0:router(config)#</pre>	
Step 6	mpls static Example: <pre>RP/0/RP0/CPU0:router(config)# mpls static</pre>	Configure MPLS static on the egress interface connecting to the eBGP peer.
Step 7	interface type interface-path-id Example: <pre>RP/0/RP0/CPU0:router(config-mpls-static)# interface GigabitEthernet0/0/1/2</pre>	Specifies the egress interface connecting to the eBGP peer.
Step 8	Use the commit or end command.	commit —Saves the configuration changes and remains within the configuration session. end —Prompts user to take one of these actions: • Yes — Saves configuration changes and exits the configuration session. • No —Exits the configuration session without committing the configuration changes. • Cancel —Remains in the configuration session, without committing the configuration changes.

Example

Running Config:


```

router bgp 1
 neighbor 192.168.1.3
  remote-as 3
  egress-engineering
  !
!
mpls static
 interface GigabitEthernet0/0/1/2
  !
!

```

Configuring Manual BGP-EPE Peering SIDs

Configuring manual BGP-EPE Peer SIDs allows for persistent EPE label values. Manual BGP-EPE SIDs are advertised through BGP-LS and are allocated from the Segment Routing Local Block (SRLB). See [Configure Segment Routing Global Block and Segment Routing Local Block, on page 137](#) for information about the SRLB.

Each PeerNode SID, PeerAdjacency SID, and PeerSet SID is configured with an index value. This index serves as an offset from the configured SRLB start value and the resulting MPLS label (SRLB start label + index) is assigned to these SIDs. This label is used by CEF to perform load balancing across the individual BGP PeerSet SIDs, BGP PeerNode SID, or ultimately across each first-hop adjacency associated with that BGP PeerNode SID or BGP PeerSet SID.

Configuring Manual PeerNode SID

Each eBGP peer will be associated with a PeerNode SID index that is configuration driven.

```

RP/0/0/CPU0:PE1(config)# router bgp 10
RP/0/0/CPU0:PE1(config-bgp)# neighbor 10.10.10.2
RP/0/0/CPU0:PE1(config-bgp-nbr)# remote-as 20
RP/0/0/CPU0:PE1(config-bgp-nbr)# egress-engineering
RP/0/0/CPU0:PE1(config-bgp-nbr)# peer-node-sid index 600

```

Configuring Manual PeerAdjacency SID

Any first-hop for which an adjacency SID is configured needs to be in the resolution chain of at least one eBGP peer that is configured for egress-peer engineering. Otherwise such a kind of “orphan” first-hop with regards to BGP has no effect on this feature. This is because BGP only understands next-hops learnt by the BGP protocol itself and in addition only the resolving IGP next-hops for those BGP next-hops.

```

RP/0/0/CPU0:PE1(config)# router bgp 10
RP/0/0/CPU0:PE1(config-bgp)# adjacencies
RP/0/0/CPU0:PE1(config-bgp-adj)# 10.1.1.2
RP/0/0/CPU0:PE1(config-bgp-adj)# adjacency-sid index 500

```

Configuring Manual PeerSet SID

The PeerSet SID is configured under global Address Family. This configuration results in the creation of a Peer-Set SID EPE object.

```

RP/0/0/CPU0:PE1(config)# router bgp 10
RP/0/0/CPU0:PE1(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:PE1(config-bgp-afi)# peer-set-id 1


```

```
RP/0/0/CPU0:PE1(config-bgp-peer-set)# peer-set-sid 300
```

Example

Topology

The example in this section uses the following topology.



In this example, BGP-EPE peer SIDs are allocated from the default SRLB label range (15000 – 15999). The BGP-EPE peer SIDs are configured as follows:

- PeerNode SIDs to 10.10.10.2 with index 600 (label 15600), and for 20.10.10.2 with index 700 (label 15700)
- PeerAdj SID to link 10.1.1.2 with index 500 (label 15500)
- PeerSet SID 1 to load balance over BGP neighbors 10.10.10.1 and 20.10.10.2 with SID index 300 (label 15300)
- PeerSet SID 2 to load balance over BGP neighbor 20.10.10.2 and link 10.1.1.2 with SID index 400 (label 15400)

Configuration on R1

```
router bgp 10
 address-family ipv4 unicast
  peer-set-id 1
  peer-set-sid index 300
  !
  peer-set-id 2
  peer-set-sid index 400
  !
  !
  adjacencies
  10.1.1.2
  adjacency-sid index 500
  peer-set 2
  !
  !
  neighbor 10.10.10.2
```

```
remote-as 20
egress-engineering
peer-node-sid index 600
peer-set 1
!
neighbor 20.10.10.2
egress-engineering
peer-node-sid index 700
peer-set 1
peer-set 2
!
```

To further show the load balancing of this example:

- 15600 is load balanced over {10.1.1.1 and 2.1.1.1}
- 15700 is load balanced over {3.1.1.1 and 4.1.1.1}
- 15500 is load balanced over {10.1.1.1}
- 15300 is load balanced over {10.1.1.1, 2.1.1.1, 3.1.1.1 and 4.1.1.1}
- 15400 is load balanced over {10.1.1.1, 3.1.1.1 and 4.1.1.1}


Configure BGP Link-State

BGP Link-State (LS) is an Address Family Identifier (AFI) and Sub-address Family Identifier (SAFI) originally defined to carry interior gateway protocol (IGP) link-state information through BGP. The BGP Network Layer Reachability Information (NLRI) encoding format for BGP-LS and a new BGP Path Attribute called the BGP-LS attribute are defined in [RFC7752](#). The identifying key of each Link-State object, namely a node, link, or prefix, is encoded in the NLRI and the properties of the object are encoded in the BGP-LS attribute.

The BGP-LS Extensions for Segment Routing are documented in [RFC9085](#).

BGP-LS applications like an SR Path Computation Engine (SR-PCE) can learn the SR capabilities of the nodes in the topology and the mapping of SR segments to those nodes. This can enable the SR-PCE to perform path computations based on SR-TE and to steer traffic on paths different from the underlying IGP-based distributed best-path computation.

The following figure shows a typical deployment scenario. In each IGP area, one or more nodes (BGP speakers) are configured with BGP-LS. These BGP speakers form an iBGP mesh by connecting to one or more route-reflectors. This way, all BGP speakers (specifically the route-reflectors) obtain Link-State information from all IGP areas (and from other ASes from eBGP peers).



Usage Guidelines and Limitations

- BGP-LS supports IS-IS and OSPFv2.
- The identifier field of BGP-LS (referred to as the Instance-ID) identifies the IGP routing domain where the NLRI belongs. The NLRIs representing link-state objects (nodes, links, or prefixes) from the same IGP routing instance must use the same Instance-ID value.
- When there is only a single protocol instance in the network where BGP-LS is operational, we recommend configuring the Instance-ID value to **0**.
- Assign consistent BGP-LS Instance-ID values on all BGP-LS Producers within a given IGP domain.
- NLRIs with different Instance-ID values are considered to be from different IGP routing instances.
- Unique Instance-ID values must be assigned to routing protocol instances operating in different IGP domains. This allows the BGP-LS Consumer (for example, SR-PCE) to build an accurate segregated multi-domain topology based on the Instance-ID values, even when the topology is advertised via BGP-LS by multiple BGP-LS Producers in the network.
- If the BGP-LS Instance-ID configuration guidelines are not followed, a BGP-LS Consumer may see duplicate link-state objects for the same node, link, or prefix when there are multiple BGP-LS Producers deployed. This may also result in the BGP-LS Consumers getting an inaccurate network-wide topology.

- The following table defines the supported extensions to the BGP-LS address family for carrying IGP topology information (including SR information) via BGP. For more information on the BGP-LS TLVs, refer to [Border Gateway Protocol - Link State \(BGP-LS\) Parameters](#).

Table 27: IOS XR Supported BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs

TLV Code Point	Description	Produced by IS-IS	Produced by OSPFv2	Produced by BGP
256	Local Node Descriptors	X	X	—
257	Remote Node Descriptors	X	X	—
258	Link Local/Remote Identifiers	X	X	—
259	IPv4 interface address	X	X	—
260	IPv4 neighbor address	X		
261	IPv6 interface address	X	—	—
262	IPv6 neighbor address	X	—	—
263	Multi-Topology ID	X	—	—
264	OSPF Route Type	—	X	—
265	IP Reachability Information	X	X	—
266	Node MSD TLV	X	X	—
267	Link MSD TLV	X	X	—
512	Autonomous System	—	—	X
513	BGP-LS Identifier	—	—	X
514	OSPF Area-ID	—	X	—
515	IGP Router-ID	X	X	—
516	BGP Router-ID TLV	—	—	X
517	BGP Confederation Member TLV	—	—	X
1024	Node Flag Bits	X	X	—
1026	Node Name	X	X	—
1027	IS-IS Area Identifier	X	—	—
1028	IPv4 Router-ID of Local Node	X	X	—
1029	IPv6 Router-ID of Local Node	X	—	—
1030	IPv4 Router-ID of Remote Node	X	X	—
1031	IPv6 Router-ID of Remote Node	X	—	—
1034	SR Capabilities TLV	X	X	—
1035	SR Algorithm TLV	X	X	—
1036	SR Local Block TLV	X	X	—

TLV Code Point	Description	Produced by IS-IS	Produced by OSPFv2	Produced by BGP
1039	Flex Algo Definition (FAD) TLV	X	X	—
1044	Flex Algorithm Prefix Metric (FAPM) TLV	X	X	—
1088	Administrative group (color)	X	X	—
1089	Maximum link bandwidth	X	X	—
1090	Max. reservable link bandwidth	X	X	—
1091	Unreserved bandwidth	X	X	—
1092	TE Default Metric	X	X	—
1093	Link Protection Type	X	X	—
1094	MPLS Protocol Mask	X	X	—
1095	IGP Metric	X	X	—
1096	Shared Risk Link Group	X	X	—
1099	Adjacency SID TLV	X	X	—
1100	LAN Adjacency SID TLV	X	X	—
1101	PeerNode SID TLV	—	—	X
1102	PeerAdj SID TLV	—	—	X
1103	PeerSet SID TLV	—	—	X
1114	Unidirectional Link Delay TLV	X	X	—
1115	Min/Max Unidirectional Link Delay TLV	X	X	—
1116	Unidirectional Delay Variation TLV	X	X	—
1117	Unidirectional Link Loss	X	X	—
1118	Unidirectional Residual Bandwidth	X	X	—
1119	Unidirectional Available Bandwidth	X	X	—
1120	Unidirectional Utilized Bandwidth	X	X	—
1122	Application-Specific Link Attribute TLV	X	X	—
1152	IGP Flags	X	X	—
1153	IGP Route Tag	X	X	—
1154	IGP Extended Route Tag	X	—	—
1155	Prefix Metric	X	X	—
1156	OSPF Forwarding Address	—	X	—
1158	Prefix-SID	X	X	—
1159	Range	X	X	—

TLV Code Point	Description	Produced by IS-IS	Produced by OSPFv2	Produced by BGP
1161	SID/Label TLV	X	X	—
1170	Prefix Attribute Flags	X	X	—
1171	Source Router Identifier	X	—	—
1172	L2 Bundle Member Attributes TLV	X	—	—
1173	Extended Administrative Group	X	X	—

Exchange Link State Information with BGP Neighbor

The following example shows how to exchange link-state information with a BGP neighbor:

```
Router# configure
Router(config)# router bgp 1
Router(config-bgp)# neighbor 10.0.0.2
Router(config-bgp-nbr)# remote-as 1
Router(config-bgp-nbr)# address-family link-state link-state
Router(config-bgp-nbr-af)# exit
```

IGP Link-State Database Distribution

A given BGP node may have connections to multiple, independent routing domains. IGP link-state database distribution into BGP-LS is supported for both OSPF and IS-IS protocols in order to distribute this information on to controllers or applications that desire to build paths spanning or including these multiple domains.

To distribute IS-IS link-state data using BGP-LS, use the **distribute link-state** command in router configuration mode.

```
Router# configure
Router(config)# router isis isp
Router(config-isis)# distribute link-state instance-id 32
```

To distribute OSPFv2 link-state data using BGP-LS, use the **distribute link-state** command in router configuration mode.

```
Router# configure
Router(config)# router ospf 100
Router(config-ospf)# distribute link-state instance-id 32
```

Configurable Filters for IS-IS Advertisements to BGP-Link State

Table 28: Feature History Table

Feature Name	Release Information	Feature Description
Configurable Filters for IS-IS Advertisements to BGP-Link State	Release 7.10.1	This feature allows you to configure a route map to filter IS-IS route advertisements to BGP-Link State (LS). It also provides a per-area configuration knob to disable IS-IS advertisements for external and propagated prefixes. This configuration of filters hence reduces the amount of redundant data for external and interarea prefixes sent to the BGP - LS clients. The feature introduces exclude-external, exclude-interarea, and route-policy name optional keywords in the distribute link-state command.

In a large IS-IS network, there are multiple routers in different areas distributing their link-state databases through BGP-LS. In addition, other protocols, such as OSPF do their own BGP-LS reporting and have routes that are redistributed into IS-IS. This can result in substantial amounts of redundant data for external and interarea prefixes which are sent to the BGP-LS clients only to be discarded.

Rather than sending redundant information, this feature provides the option of limiting the prefixes for which IS-IS TLV information is sent to BGP-LS.

There are three options to filter prefix Type-Length-Values (TLVs) that are reported in BGP-LS and the operators can specify these options on a per-level basis:

- **exclude-external**: Omits information for external prefixes that are redistributed from a different protocol or instance. These are identified by the “X” bit set in its Extended Reachability Attribute Flags or the ‘X’ bit of TLVs 236 and 237.
- **exclude-interarea**: Omits information for interarea prefixes and summaries. These are identified by the ‘R’ bit set in their Extended Reachability Attribute Flags or the ‘up or down’ bit set in TLVs 135, 235, 236, and 237.
- **route-policyname**: Allows specification of a route-policy to provide filtering based on a set of destination prefixes.

The filtering is implemented at the point where the individual prefix TLVs are read from a label-switched path to generate updates to BGP-LS. It does not affect the advertisement of a node or the link information.

Configure Filters for IS-IS Advertisements to BGP-LS

Configuration Example

You can configure any of these filters for IS-IS advertisements to BGP-LS:

```
Router#config
Router(config)#router isis 1
```



```

Router(config-isis)#distribute link-state exclude-external
Router(config-isis)#commit

Router#config
Router(config)#router isis 1
Router(config-isis)#distribute link-state exclude-interarea
Router(config-isis)#commit

Router# config
Router(config)# router isis 1
Router(config-isis)#distribute link-state route-policy isis-rp-1
Router(config-isis)#commit

```



Note This feature does not introduce any new failure modes to IS-IS.

Running Configuration

To check the filter for IS-IS advertisements to BGP-LS, you can run the following command:

```

Router# show running-config
router isis 1
  distribute link-state exclude-external
  commit
!
!

router isis 1
  distribute link-state exclude-interarea
  commit
!
!


router isis 1
  distribute link-state route-policy isis-rp-1
  commit
!
!

```

Use Case: Configuring SR-EPE and BGP-LS

In the following figure, segment routing is enabled on autonomous system AS1 with ingress node A and egress nodes B and C. In this example, we configure EPE on egress node C.

Figure 11: Topology



Procedure

Step 1 Configure node C with EPE for eBGP peers D and E.

Example:

```
RP/0/RP0/CPU0:router_C(config)# router bgp 1
RP/0/RP0/CPU0:router_C(config-bgp)# neighbor 192.168.1.3
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# remote-as 3
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# description to E
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# egress-engineering
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# route-policy bgp_in in
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# route-policy bgp_out out
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# exit
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router_C(config-bgp)# neighbor 192.168.1.2
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# description to D
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# egress-engineering
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# route-policy bgp_in in
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# route-policy bgp_out out
RP/0/RP0/CPU0:router_C(config-bgp-nbr-af)# exit
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# exit
```

Step 2 Configure node C to advertise peer node SIDs to the controller using BGP-LS.

Example:

```
RP/0/RP0/CPU0:router_C(config-bgp)# neighbor 172.29.50.71
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# description to EPE_controller
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# address-family link-state link-state
RP/0/RP0/CPU0:router_C(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router_C(config-bgp)# exit
```

Step 3 Configure MPLS static on the egress interfaces connecting to the eBGP peer.

Example:

```
RP/0/RP0/CPU0:router_C(config)# mpls static
RP/0/RP0/CPU0:router_C(config-mpls-static)# interface TenGigE 0/3/0/0
```

```
RP/0/RP0/CPU0:router_C(config-mpls-static)# interface TenGigE 0/1/0/0
RP/0/RP0/CPU0:router_C(config-mpls-static)# exit
```

Step 4 Commit the configuration.

Example:

```
RP/0/RP0/CPU0:router_C(config)# commit
```

Step 5 Verify the configuration.

Example:

```
RP/0/RP0/CPU0:router_C# show bgp egress-engineering

Egress Engineering Peer Set: 192.168.1.2/32 (10b87210)
  Nexthop: 192.168.1.2
  Version: 2, rn_version: 2
  Flags: 0x00000002
  Local ASN: 1
  Remote ASN: 2
  Local RID: 10.1.1.3
  Remote RID: 10.1.1.4
  First Hop: 192.168.1.2
  NHID: 3
  Label: 24002, Refcount: 3
  rpc_set: 10b9d408

Egress Engineering Peer Set: 192.168.1.3/32 (10be61d4)
  Nexthop: 192.168.1.3
  Version: 3, rn_version: 3
  Flags: 0x00000002
  Local ASN: 1
  Remote ASN: 3
  Local RID: 10.1.1.3
  Remote RID: 10.1.1.5
  First Hop: 192.168.1.3
  NHID: 4
  Label: 24003, Refcount: 3
  rpc_set: 10be6250
```

The output shows that node C has allocated peer SIDs for each eBGP peer.

Example:

```
RP/0/RP0/CPU0:router_C# show mpls forwarding labels 24002 24003
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
24002	Pop	No ID	Te0/0/0/1	192.168.1.2	0
24003	Pop	No ID	Te0/0/0/2	192.168.1.3	0

The output shows that node C installed peer node SIDs in the Forwarding Information Base (FIB).

Configure BGP Proxy Prefix SID

To support segment routing, Border Gateway Protocol (BGP) requires the ability to advertise a segment identifier (SID) for a BGP prefix. A BGP-Prefix-SID is the segment identifier of the BGP prefix segment in a segment routing network. BGP prefix SID attribute is a BGP extension to signal BGP prefix-SIDs. However, there may be routers which do not support BGP extension for segment routing. Hence, those routers also do not support BGP prefix SID attribute and an alternate approach is required.

BGP proxy prefix SID feature allows you to attach BGP prefix SID attributes for remote prefixes learnt from BGP labeled unicast (LU) neighbours which are not SR-capable and propagate them as SR prefixes. This allows an LSP towards non SR endpoints to use segment routing global block in a SR domain. Since BGP proxy prefix SID uses global label values it minimizes the use of limited resources such as ECMP-FEC and provides more scalability for the networks.

BGP proxy prefix SID feature is implemented using the segment routing mapping server (SRMS). SRMS allows the user to configure SID mapping entries to specify the prefix-SIDs for the prefixes. The mapping server advertises the local SID-mapping policy to the mapping clients. BGP acts as a client of the SRMS and uses the mapping policy to calculate the prefix-SIDs.

Configuration Example:

This example shows how to configure the BGP proxy prefix SID feature for the segment routing mapping server.

```
RP/0/RSP0/CPU0:router(config)# segment-routing
RP/0/RSP0/CPU0:router(config-sr)# mapping-server
RP/0/RSP0/CPU0:router(config-sr-ms)# prefix-sid-map
RP/0/RSP0/CPU0:router(config-sr-ms-map)# address-family ipv4
RP/0/RSP0/CPU0:router(config-sr-ms-map-af)# 10.1.1.1/32 10 range 200
RP/0/RSP0/CPU0:router(config-sr-ms-map-af)# 192.168.64.1/32 400 range 300
```

This example shows how to configure the BGP proxy prefix SID feature for the segment-routing mapping client.

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# address-family ip4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# segment-routing prefix-sid-map
```

Verification

These examples show how to verify the BGP proxy prefix SID feature.

```
RP/0/RSP0/CPU0:router# show segment-routing mapping-server prefix-sid-map ipv4 detail
Prefix
10.1.1.1/32
  SID Index:      10
  Range:          200
  Last Prefix:    10.1.1.200/32
  Last SID Index: 209
  Flags:
Number of mapping entries: 1
```

```
RP/0/RSP0/CPU0:router# show bgp ipv4 labeled-unicast 192.168.64.1/32
```

```
BGP routing table entry for 192.168.64.1/32
```

```

Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          117        117
  Local Label: 16400
Last Modified: Oct 25 01:02:28.562 for 00:11:45Paths: (2 available, best #1)
Advertised to peers (in unique update groups):
  201.1.1.1
Path #1: Received by speaker 0  Advertised to peers (in unique update groups):
  201.1.1.1
Local
  20.0.101.1 from 20.0.101.1 (20.0.101.1)      Received Label 61
  Origin IGP, localpref 100, valid, internal, best, group-best, multipath, labeled-unicast

  Received Path ID 0, Local Path ID 0, version 117
Prefix SID Attribute Size: 7
Label Index: 1

RP/0/RSP0/CPU0:router# show route ipv4 unicast 192.68.64.1/32 detail

```

```

Routing entry for 192.168.64.1/32
  Known via "bgp 65000", distance 200, metric 0, [ei]-bgp, labeled SR, type internal
  Installed Oct 25 01:02:28.583 for 00:20:09
  Routing Descriptor Blocks
    20.0.101.1, from 20.0.101.1, BGP multi path
    Route metric is 0
    Label: 0x3d (61)
    Tunnel ID: None
    Binding Label: None
    Extended communities count: 0
    NHID:0x0(Ref:0)
    Route version is 0x6 (6)
    Local Label: 0x3e81 (16400)
    IP Precedence: Not Set
    QoS Group ID: Not Set
    Flow-tag: Not Set
    Fwd-class: Not Set
    Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_LOCAL
    Download Priority 4, Download Version 242
    No advertising protos.

```

```

RP/0/RSP0/CPU0:router# show cef ipv4 192.168.64.1/32 detail
192.168.64.1/32, version 476, labeled SR, drop adjacency, internal 0x5000001 0x80 (ptr
0x71c42b40) [1], 0x0 (0x71c11590), 0x808 (0x722b91e0)
Updated Oct 31 23:23:48.733
Prefix Len 32, traffic index 0, precedence n/a, priority 4
Extensions: context-label:16400
  gateway array (0x71ae7e78) reference count 3, flags 0x7a, source rib (7), 0 backups
    [2 type 5 flags 0x88401 (0x722eb450) ext 0x0 (0x0)]
  LW-LDI[type=5, refc=3, ptr=0x71c11590, sh-ldi=0x722eb450]
  gateway array update type-time 3 Oct 31 23:49:11.720
  LDI Update time Oct 31 23:23:48.733
  LW-LDI-TS Oct 31 23:23:48.733
    via 20.0.101.1/32, 0 dependencies, recursive, bgp-ext [flags 0x6020]
    path-idx 0 NHID 0x0 [0x7129a294 0x0]
    recursion-via-/32
    unresolved
    local label 16400
    labels imposed {ExpNullv6}

```

```

RP/0/RSP0/CPU0:router# show bgp labels
BGP router identifier 2.1.1.1, local AS number 65000
BGP generic scan interval 60 secs

```

```

Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 245
BGP main routing table version 245
BGP NSR Initial initsync version 16 (Reached)
BGP NSR/ISSU Sync-Group versions 245/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network                Next Hop           Rcvd Label       Local Label
*>i10.1.1.1/32            10.1.1.1             3                16010
*> 2.1.1.1/32             0.0.0.0             no-label         3
*> 192.68.64.1/32        20.0.101.1           2                16400
*> 192.68.64.2/32        20.0.101.1           2                16401

```

BGP-LU Inter-AS Option-C Interworking with LDP and IGP SR-MPLS using Proxy BGP-SR

Table 29: Feature History Table

Feature Name	Release	Description
BGP-LU Inter-AS Option-C Interworking with LDP and IGP SR-MPLS using Proxy BGP-SR	Release 7.3.2	This feature extends the current Proxy BGP-SR functionality by allowing the BGP-LU ASBR router with Proxy BGP-SR configured to also interconnect attached LDP domains. The Proxy BGP-SR feature allows interconnection of IGP SR-MPLS domains and legacy domains via BGP-LU Inter-AS option-C. It provides a prefix-to-SID mapping for BGP-LU prefixes that are learned without a Prefix-SID.

The Proxy BGP-SR feature allows interconnection of IGP SR-MPLS domains and legacy domains via BGP-LU Inter-AS option-C. It provides a prefix-to-SID mapping for BGP-LU prefixes that are learned without a Prefix-SID. This new feature extends the current functionality by allowing the BGP-LU ASBR router (configured with Proxy BGP-SR) to also interconnect attached LDP domains.


With this enhancement, when performing redistribution from BGP into IGP, LDP would use the same local label assigned by BGP for a prefix learned by BGP-LU. The local label value is based on the SR mapping server configuration (Proxy-BGP SR feature). This behavior allows incoming LDP traffic destined to a redistributed prefix to be switched over to the BGP-LU Inter-AS LSP.

Use Case

In the following figure, Router A does the following:

- Is an ASBR for BGP AS 100 running BGP-LU with BGP AS 200

- Interconnects two IS-IS processes: one running LDP and another running Segment Routing
- Redistributes prefixes learned by BGP-LU from AS 200 into both IS-IS instances
- Runs SR Mapping Server (SRMS) in order to assign mappings to prefixes learned by BGP LU from AS 200 without a prefix SID (proxy BGP-SR) and prefixes learned from the LDP domain



Configuration on Router A - ASBR for AS100

```

prefix-set pfxset-bgplu
  100.1.1.8/32 // The Prefix under test
end-set
!
prefix-set LOOPBACKS
  192.168.0.1,
  192.168.0.2,
  192.168.0.3,
  192.168.0.4,
  192.168.0.8
end-set
!
route-policy Pass
  pass
end-policy
!
route-policy rpl-bgplu
  if destination in pfxset-bgplu then
    pass
  else
    drop
  endif
end-policy
!
route-policy MATCH_LOOPBACKS

```

```

        if destination in LOOPBACKS then
            pass
        else
            drop
        endif
    end-policy
!
router static
    address-family ipv4 unicast
        10.10.10.2/32 GigabitEthernet0/2/0/0
    !
!
router isis 50
    is-type level-2-only
    net 49.0001.0000.0000.0001.00
    address-family ipv4 unicast
        metric-style wide
    redistribute bgp 100 route-policy rp1-bgplu // Redistribute prefixes learned by BGP-LU
    into IS-IS LDP domain
    !
    interface Loopback0
        passive
        address-family ipv4 unicast
    !
    !
    interface GigabitEthernet0/2/0/1
        address-family ipv4 unicast
    !
    !
!
router isis 100
    is-type level-2-only
    net 49.0001.0000.0000.0011.00
    distribute link-state
    address-family ipv4 unicast
        metric-style wide
    mpls traffic-eng level-2-only
    mpls traffic-eng router-id Loopback0
    redistribute bgp 100 route-policy rp1-bgplu // Redistribute prefixes learned by BGP-LU
    into IS-IS SR domain
    segment-routing mpls
    segment-routing prefix-sid-map advertise-local
    !
    interface Loopback0
        passive
        address-family ipv4 unicast
        prefix-sid index 1
    !
    !
    interface GigabitEthernet0/2/0/2
        address-family ipv4 unicast
    !
    !
!
router bgp 100
    bgp router-id 192.168.0.1
    address-family ipv4 unicast
    segment-routing prefix-sid-map // SR Proxy SID Configuration
    network 192.168.0.1/32
    redistribute isis 50 route-policy MATCH_LOOPBACKS
    redistribute isis 100 route-policy MATCH_LOOPBACKS
    allocate-label all
    !
    neighbor 10.10.10.2

```



```

remote-as 200
address-family ipv4 labeled-unicast
  route-policy Pass in
  route-policy Pass out
!
!
!
mpls ldp
router-id 192.168.0.1
interface GigabitEthernet0/2/0/1
!
!
segment-routing
global-block 16000 23999
mapping-server
  prefix-sid-map // SRMS configuration
  address-family ipv4
    100.1.1.8/32 108 range 1 // SRMS mapping - LU prefix 100.1.1.8/32 assigned prefix index
108
    192.168.0.3/32 3 range 1 // SRMS mapping - LDP prefix Router C assigned prefix index
3
!
!
!
!

```

Configuration on Router B - ASBR for AS200

```

route-policy Pass
  pass
end-policy
!
router static
  address-family ipv4 unicast
    10.10.10.1/32 GigabitEthernet0/0/0/0
!
!
router bgp 200
  bgp router-id 192.168.0.2
  address-family ipv4 unicast
    network 100.1.1.8/32 // Import/Inject route into BGP
    network 192.168.0.2/32
    allocate-label all
!
  neighbor 10.10.10.1
  remote-as 100
  address-family ipv4 labeled-unicast
    route-policy Pass in
    route-policy Pass out
!
!
!

```

Configuration on Router C in the LDP Domain

```

router isis 50
  is-type level-2-only
  net 49.0001.0000.0000.0003.00
  address-family ipv4 unicast
    metric-style wide
!
interface Loopback0
  passive

```

```
    address-family ipv4 unicast
    !
  !
  interface GigabitEthernet0/0/0/0
    address-family ipv4 unicast
    !
  !
  !
  mpls ldp
    router-id 192.168.0.3
    interface GigabitEthernet0/0/0/0
    !
  !
```

Configuration on Router D in the SR IS-IS Domain

```
router isis 100
  is-type level-2-only
  net 49.0001.0000.0000.0004.00
  address-family ipv4 unicast
    metric-style wide
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
  segment-routing mpls
  !
  interface Loopback0
    passive
    address-family ipv4 unicast
    prefix-sid index 4
    !
  !
  interface GigabitEthernet0/0/0/0
    address-family ipv4 unicast
    !
  !
  !
  segment-routing
  !
```

SRv6 double recursion for multilayer BGP underlay

Table 30: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 double recursion for multilayer BGP underlay	Release 24.4.1	The feature introduces support for SRv6 double recursion, where network services such as BGP VPNs (Layer 2 and Layer 3) require multiple resolution layers. Specifically, one routing layer resolves over another before reaching the final destination. Double recursion is achieved by collapsing the underlay, typically involving protocols like IGP or BGP in the packet forwarding chain. This enables three-level load balancing and an even distribution of traffic across multiple layers of the network stack. The feature is supported on the ingress Provider Edge (PE) router. The feature introduces these changes: CLI: • <code>tag-map tag <value> map forwarding-hierarchy level-2-used-as-nexthop</code> • The <code>show cef ipv6</code> and <code>show cef ipv4</code> commands now include the Layer 2 prefix information, which resolve as nexthop Layer 3 prefixes. YANG Data Models: • <code>Cisco-IOS-XR-um-router-rib-cfg:router</code> (see GitHub, Yang Data Models Navigator) data model.

SRv6 double recursion and load balancing

SRv6 double recursion refers to scenarios where network service such as BGP VPN (Layer 2/Layer 3) requires multiple layers of resolution, specifically where one routing layer resolves over another before reaching its final destination.

The SRv6 double recursion feature is designed to support three-level load balancing by collapsing the underlay, which typically involves protocols like IGP or BGP in the forwarding chain. This involves specific configurations in RIB, BGP, and static routes to indicate IPv4 and IPv6 unicast routes with Layer 2 prefixes, which resolve as nexthop Layer 3 prefixes.

Without this feature, the traffic drops the Layer 3 prefixes leading to packet loss.

Key benefits of SRv6 double recursion

- Support for complex network scenarios: In traditional service provider (SP) networks, BGP VPN services typically resolve directly over IGP reachability, meaning there is a straightforward, single-level resolution path from the service to the IGP underlay. However, with the evolution of networking, especially in data centers, new architectures have emerged where BGP underlay networks are used. In these scenarios, BGP VPN services resolve over a BGP underlay, which in turn resolves over IGP or directly connected routes.
- Support for enhanced load balancing: SRv6 double recursion allows the routing platform to balance traffic more efficiently across all available paths, enhancing overall network performance and reducing congestion. Proper load balancing at each level of recursion helps in distributing traffic evenly, avoiding bottlenecks, and ensuring that packets take optimal paths even in complex topologies.
- Improved network flexibility: Double recursion allows advanced routing solutions in SP and data center networks, accommodating complex use cases beyond single-layer limitations.

Single and double recursion in BGP VPN services

To optimize the network routing efficiency, it is important to understand the resolution processes of BGP VPN services.

- Single Recursion: BGP VPN service → IGP reachability. In a single recursion, the BGP VPN service directly resolves to IGP (Interior Gateway Protocol) reachability. Single recursion is required when a BGP VPN service can directly resolve to an IGP route without any intermediary steps.
- Double Recursion: BGP VPN service → BGP underlay reachability → IGP reachability. In double recursion, the BGP VPN service first resolves over a BGP underlay, which then leads to IGP reachability. Double recursion is necessary in cases where a BGP VPN service cannot resolve directly to an IGP route but instead must traverse a BGP underlay first. This creates a layered resolution process that standard single-level load balancing cannot handle effectively.

Usage guidelines and limitations for SRv6 double recursion

Limitations

You cannot use a combination of BGP-SU (BGP Service Unicast) and BGP-IP paths at Layer 2.

Usage guidelines

The usage guidelines that are listed apply:

- Configure the locator prefix for IPv6 prefixes that are Layer 2 prefixes.
- If the collapsed BGP paths are a combination of IGP IPv6 and SRv6, the router filters out only the IPv6 paths.
- In a collapsed chain scenario with SRv6 as the IGP, the router configures the encapsulation (Encap) ID, provided it supports SRv6 and BGP extensions for SRv6.

Configure SRv6 double recursion for multilayer BGP underlay

Before you begin

- For Layer 2 IPv6 prefixes, set the locator prefix using the **prefix-set** command in RIB.
- Configure the set of IPv4 and IPv6 Layer 2 prefixes, which resolve as nexthop Layer 3 prefixes using the **prefix-set** command.

Procedure

-
- Step 1** Enable the hardware support for BGP-LU to allow the ingress PE router to advertise and forward the MPLS-labelled unicast routes.
- Example:**
- ```
Router#config
Router(config)#hw-module profile cef bgplu enable
```
- Step 2** Configure the IPv4 and IPv6 unicast routes in the RIB to ensure that the Layer 2 prefixes can be used as nexthop for Layer 3 prefixes.
- Example:**
- ```
Router#config
Router(config)#router bgp 100
Router(config-bgp)#address-family ipv4 unicast
Router(config-bgp-af)#table-policy level2-ipv4-policy
Router(config-bgp-af)#exit
Router(config-bgp)#address-family ipv6 unicast
Router(config-bgp-af)#table-policy level2-ipv6-policy
Router(config-bgp-af)#commit
```
- Step 3** Assign the required tag to the IPv4 and IPv6 unicast routes to forward the route information to the Forwarding Information Base (FIB).
- For the IPv4 and IPv6 routes, the RIB assigns the tags and forwards the route information to the FIB, indicating that these Layer 2 prefixes resolve as nexthop for Layer 3 prefixes.
- Example:**
- IPv4 Layer 2 prefixes:
- ```
Router(config)#route-policy level2-ipv4-policy
Router(config-rpl)#if destination in level2_prefixes-ipv4 then
```

```
Router(config-rpl-if)#set tag 100
Router(config-rpl-if)#else
Router(config-rpl-else)#pass
Router(config-rpl-else)#endif
Router(config-rpl)#end-policy
Router(config)#commit
```

**Example:**

IPv6 Layer 2 prefixes:

```
Router(config)#route-policy level2-ipv6-policy
Router(config-rpl)#if destination in level2_prefixes-ipv6 then
Router(config-rpl-if)#set tag 100
Router(config-rpl-if)#endif
Router(config-rpl)#if destination in level2_locators then
Router(config-rpl-if)#set locator-prefix
Router(config-rpl-if)#else
Router(config-rpl-else)#pass
Router(config-rpl-else)#endif
Router(config-rpl)#end-policy
```

- Step 4** Map the tags in the RIB for the IPv4 and IPv6 unicast routes to ensure that the routes resolve as nexthop for Layer 3 prefixes.

In the following example, the RIB maps all the IPv4 and IPv6 routes tagged with the value 100, which indicates that these routes resolve as nexthop for Layer 3 prefixes. The RIB adds the *FIB\_UPDATE\_ROUTE\_FLAG\_EXTN\_LVL2\_HAS\_DEPENDENT* flag when it sends the route update to the FIB.

**Example:**

```
Router#config
Router(config)#router rib
Router(config-rib)#tag-map tag 100 map forwarding-hierarchy level-2-used-as-nexthop
Router(config-rib)#commit
```

- Step 5** Run the **show running-config** command to verify the running configuration.

**Example:**

```
hw-module profile cef bgplu enable
!
router bgp 100
 address-family ipv4 unicast
 table-policy level2-ipv4-policy
 !
 address-family ipv6 unicast
 table-policy level2-ipv6-policy
 !
!
router rib
 tag-map tag 100 map forwarding-hierarchy level-2-used-as-nexthop
!
!
```

- Step 6** Run the **show cef ipv4** and **show cef ipv6** commands to verify the Layer 2 collapsed prefixes.

In the show output command example, the IPv6 Layer 2 prefixes, which resolve as nexthop for Layer 3 prefixes are collapsed. This is indicated by the **collapsed** keyword in the output. The SRv6 SID lists indicate the different encapsulation layers or hierarchy.

**Example:**

```

Router#show cef ipv6 2001:DB8:A:B::1/64
Thu Jun 6 12:48:52.399 EDT
2001:DB8:A:B::1/64, version 8, SRv6 Headend, internal 0x1000001 0x0 (ptr 0x63851c98) [1],
0x1400 (0x63851da0), 0x0 (0x638b2128)
Updated Jun 6 12:41:10.589
Prefix Len 64, traffic index 0, precedence n/a, priority 0, encap-id 0x11deadbeef
gateway array (0x61e1a798) reference count 1, flags 0x10, source rib (7), 0 backups
 [2 type 3 flags 0x40008501 (0x63853e38) ext 0x0 (0x0) (collapsed)]
 LW-LDI[type=3, refc=1, ptr=0x63851da0, sh-ldi=0x63853e38]
 gateway array update type-time 1 Jun 6 12:41:10.589
 LDI Update time Jun 6 12:41:10.629
 LW-LDI-TS Jun 6 12:41:10.629
Accounting: Disabled
 via 2001:DB8::1/128, 1 dependency, recursive [flags 0x3000000]
 path-idx 0 NHID 0x0 [0x63a2c098 0x0]
 next hop 2001:DB8::1/128 via 2001:DB8::1
 SRv6 H.Encaps.Red SID-list {2001:DB8:1:e002::}
 SRv6 H.Insert.Red SID-list {}
 SRv6 H.Insert.Red SID-list {bbbb:bbbb:3:: bbbb:bbbb:4::}
 via 2001:DB8::1/128, 1 dependency, recursive [flags 0x3000000]
 path-idx 1 NHID 0x0 [0x63a2c270 0x0]
 next hop 2001:DB8::1/128 via 2001:DB8::1
 SRv6 H.Encaps.Red SID-list {bbbb:bbbb:2:e002::}
 SRv6 H.Insert.Red SID-list {}

 Load distribution: 0 1 2 2 (refcount 2)

Hash OK Interface Address
0 Y UNKNOWN intf 0x00000013 10::2
1 Y UNKNOWN intf 0x00000014 20::2
2 Y UNKNOWN intf 0x00000013 10::2
3 Y UNKNOWN intf 0x00000013 10::2

```

The described show output indicates that the IPv4 Layer 2 prefixes, which resolve as nexthop for Layer 3 prefixes are collapsed.

#### Example:

```

Router#show cef 209.165.201.1 detail
Output received:
Mon Dec 2 08:31:43.765 UTC
209.165.201.1/27, version 47031, internal 0x5000001 0x40 (ptr 0x98246ad8) [1], 0x0 (0x0),
0x0 (0x0)
Updated Dec 2 08:27:35.523
Prefix Len 32, traffic index 0, precedence n/a, priority 4
gateway array (0x98099098) reference count 1, flags 0x2010, source rib (7), 0 backups
[1 type 3 flags 0x40441 (0x98134438) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Dec 2 08:27:35.523
LDI Update time Dec 2 08:31:14.951

Level 1 - Load distribution: 0 1
[0] via 209.165.200.225/27, recursive
[1] via 209.165.200.226/27, recursive

via 209.165.200.225/27, 3 dependencies, recursive, bgp-multipath [flags 0x6080]
path-idx 0 NHID 0x0 [0x982355c8 0x0]
next hop 209.165.200.225/27 via 209.165.200.225/27

Load distribution: 0 1 2 2 (refcount 1)

Hash OK Interface Address
0 Y HundredGigE0/0/0/0/2 fe80::2
1 Y Bundle-Ether1201 fe80::2

```

```

2 Y Bundle-Ether1301 fe80::3
3 Y Bundle-Ether1301 fe80::3

```

```

via 209.165.200.226/27, 3 dependencies, recursive, bgp-multipath [flags 0x6080]
path-idx 1 NHID 0x0 [0x98235678 0x0]
next hop 209.165.200.226/27 via 209.165.200.226/27

```

```
Load distribution: 0 1 2 2 (refcount 1)
```

| Hash | OK | Interface            | Address |
|------|----|----------------------|---------|
| 4    | Y  | HundredGigE0/0/0/0/2 | fe80::2 |
| 5    | Y  | Bundle-Ether1201     | fe80::2 |
| 6    | Y  | Bundle-Ether1301     | fe80::3 |
| 7    | Y  | Bundle-Ether1301     | fe80::3 |

**Step 7** Run the **router static** command to configure and map the tags in RIB for IPv6 or IPv4 static routes.

**Example:**

```

Router#config
Router(config)#router static
Router(config-static)#address-family ipv6 unicast
Router(config-static-afi)#2001:DB8:8::/48 4::4 tag 100
Router(config-static-afi)#commit

```

**Step 8** Run the **show route** command to view the configuration for IPv4 or IPv6 static routes.

**Example:**

The below show output displays the IPv4 static route configuration.

```
Router#show route 209.165.201.30
```

```
Tue Dec 3 18:22:39.579 UTC
```

```

Routing entry for 209.165.201.30/27
Known via "bgp 100", distance 200, metric 0, lvl2 has dependent
Tag 100, type internal
Installed Dec 3 18:22:35.820 for 00:00:03
Routing Descriptor Blocks
1::2, from 1::2, BGP multi path
Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
Route metric is 0
1::3, from 1::3, BGP multi path
Nexthop in Vrf: "default", Table: "default", IPv6 Unicast, Table Id: 0xe0800000
Route metric is 0
No advertising protos.

```

**Example:**

The below show output displays the IPv6 static route configuration.

```
Router#show route ipv6 2001:DB8:A:B::1 detail
```

```
Tue Dec 3 18:23:55.390 UTC
```

```

Routing entry for 2001:DB8:A:B::1/64
Known via "bgp 100", distance 200, metric 0, lvl2 has dependent
Tag 100, type internal
Installed Dec 3 18:22:45.835 for 00:01:09
Routing Descriptor Blocks
1::2, from 1::2
Route metric is 0
Label: None
Tunnel ID: None
Binding Label: None
Extended communities count: 0

```



```
NHID: 0x0 (Ref: 0)
Path Grouping ID: 100
SRv6 Headend: H.Encaps.Red [f3216], SID-list {fccc:bb01:2:e002::}
Route version is 0x1c (28)
No local label
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 4, Download Version 17265
No advertising protos.
```

---



## CHAPTER 8

# Configure SR-TE Policies

This module provides information about segment routing for traffic engineering (SR-TE) policies, how to configure SR-TE policies, and how to steer traffic into an SR-TE policy.

- [SR-TE Policy Overview, on page 225](#)
- [Usage Guidelines and Limitations, on page 226](#)
- [Instantiation of an SR Policy, on page 226](#)
- [SR-TE Policy Path Types, on page 265](#)
- [Protocols, on page 282](#)
- [Traffic Steering, on page 295](#)
- [Miscellaneous, on page 318](#)

## SR-TE Policy Overview

Segment routing for traffic engineering (SR-TE) uses a “policy” to steer traffic through the network. An SR-TE policy path is expressed as a list of segments that specifies the path, called a segment ID (SID) list. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the network to follow the specified path instead of following the shortest path calculated by the IGP. If a packet is steered into an SR-TE policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

An SR-TE policy is identified as an ordered list (head-end, color, end-point):

- Head-end – Where the SR-TE policy is instantiated
- Color – A numerical value that distinguishes between two or more policies to the same node pairs (Head-end – End point)
- End-point – The destination of the SR-TE policy

Every SR-TE policy has a color value. Every policy between the same node pairs requires a unique color value.

An SR-TE policy uses one or more candidate paths. A candidate path is a single segment list (SID-list) or a set of weighted SID-lists (for weighted equal cost multi-path [WECMP]). A candidate path is either dynamic or explicit. See *SR-TE Policy Path Types* section for more information.

# Usage Guidelines and Limitations

Observe the following guidelines and limitations for the platform.

- Broadcast links are not supported, configure IGP's interface as P2P (point-to-point).
- The ECMP path-set of an IGP route with a mix of SR-TE Policy paths (Autoroute Include) and unprotected native paths is supported.
- The ECMP path-set of an IGP route with a mix of SR-TE Policy paths (Autoroute Include) and protected (LFA/TI-LFA) native paths is not supported.
- Before configuring SR-TE policies, use the **distribute link-state** command under IS-IS or OSPF to distribute the link-state database to external services.
- L3VPN BGP PIC over SR-TE is supported.

BGP PIC over SR-TE is supported when both primary and backup paths each resolve into the BSID of an SR policy. BGP PIC over SR-TE is not supported when primary and backup paths are of different resolution types. For example, when a primary path resolves into the BSID of an SR policy, the backup path cannot point to a native LSP. When this happens, the backup path will not be programmed. For information about BGP PIC, refer to the BGP PIC chapter in the *BGP Configuration Guide*.

- SR-TE over BVI is not supported. An SR-TE policy cannot be resolved over an MPLS-enabled BVI interface.
- Counter implications when BVI and SR-TE co-exist in same NPU—Counters for a BVI's logical interface are not allocated when the same NPU hosts layer-2 (sub)interface(s) associated with the BVI alongside other port(s) used as egress interface(s) for an SR policy
- GRE tunnel as primary interface for an SR policy is not supported.
- GRE tunnel as backup interface for an SR policy with TI-LFA protection is not supported.
- Head-end computed inter-domain SR policy with Flex Algo constraint and IGP redistribution is not supported. This is supported with Flex Algo-aware path computation at SR-PCE, with or without IGP redistribution. See [SR-PCE Flexible Algorithm Multi-Domain Path Computation, on page 413](#).

## Instantiation of an SR Policy

An SR policy is instantiated, or implemented, at the head-end router.

The following sections provide details on the SR policy instantiation methods:


- [On-Demand SR Policy – SR On-Demand Next-Hop , on page 227](#)
- [Manually Provisioned SR Policy, on page 260](#)
- [PCE-Initiated SR Policy, on page 260](#)

## On-Demand SR Policy – SR On-Demand Next-Hop

Segment Routing On-Demand Next Hop (SR-ODN) allows a service head-end router to automatically instantiate an SR policy to a BGP next-hop when required (on-demand). Its key benefits include:

- **SLA-aware BGP service** – Provides per-destination steering behaviors where a prefix, a set of prefixes, or all prefixes from a service can be associated with a desired underlay SLA. The functionality applies equally to single-domain and multi-domain networks.
- **Simplicity** – No prior SR Policy configuration needs to be configured and maintained. Instead, operator simply configures a small set of common intent-based optimization templates throughout the network.
- **Scalability** – Device resources at the head-end router are used only when required, based on service or SLA connectivity needs.

The following example shows how SR-ODN works:



1. An egress PE (node H) advertises a BGP route for prefix T/t. This advertisement includes an SLA intent encoded with a BGP color extended community. In this example, the operator assigns color purple (example value = 100) to prefixes that should traverse the network over the delay-optimized path.
2. The route reflector receives the advertised route and advertises it to other PE nodes.
3. Ingress PEs in the network (such as node F) are pre-configured with an ODN template for color purple that provides the node with the steps to follow in case a route with the intended color appears, for example:
  - Contact SR-PCE and request computation for a path toward node H that does not share any nodes with another LSP in the same disjointness group.
  - At the head-end router, compute a path towards node H that minimizes cumulative delay.
4. In this example, the head-end router contacts the SR-PCE and requests computation for a path toward node H that minimizes cumulative delay.

5. After SR-PCE provides the compute path, an intent-driven SR policy is instantiated at the head-end router. Other prefixes with the same intent (color) and destined to the same egress PE can share the same on-demand SR policy. When the last prefix associated with a given [intent, egress PE] pair is withdrawn, the on-demand SR policy is deleted, and resources are freed from the head-end router.

An on-demand SR policy is created dynamically for BGP global or VPN (service) routes. The following services are supported with SR-ODN:

- EVPN-VPWS (single-homing)
- EVPN-VPWS (multi-homing)
- EVPN (single-homing/multi-homing)



**Note** For EVPN single-homing, you must configure an EVPN Ethernet Segment Identifier (ESI) with a non-zero value.



**Note** Colored per-ESI/per-EVI EVPN Ethernet Auto-Discovery route (route-type 1) and Inclusive Multicast Route (route-type 3) are used to trigger instantiation of ODN SR-TE policies.



**Note** The following scenarios involving virtual Ethernet Segments (vES) are also supported with EVPN ODN:

- VPLS VFI as vES for single-active Multi-Homing to EVPN
- Active/backup Pseudo-wire (PW) as vES for Single-Homing to EVPN
- Static Pseudo-wire (PW) as vES for active-active Multi-Homing to EVPN

## SR-ODN Configuration Steps



**Note** If you are on a release before Cisco IOS XR Release 7.4.1, you can configure SR-ODN with Flexible Algorithm constraints using the **segment-routing traffic-eng on-demand color color dynamic sid-algorithm algorithm-number** command.

Starting with Cisco IOS XR release 7.4.1, you can also configure SR-ODN with Flexible Algorithm constraints using the new **segment-routing traffic-eng on-demand color color constraints segments sid-algorithm algorithm-number** command.

From Cisco IOS XR Release 7.9.1, the **segment-routing traffic-eng on-demand color color dynamic sid-algorithm algorithm-number** command is deprecated. Previous configurations stored in NVRAM will be rejected at boot-up. Performing In-Service Software Upgrade (ISSU) to Cisco IOS XR Release 7.9.1 will also be rejected.

Hence, for Cisco IOS XR Release 7.9.1, you must reconfigure all SR-ODN configurations with Flexible Algorithm constraints that use the **on-demand dynamic sid-algorithm** with the **on-demand constraints** command.

To configure SR-ODN, complete the following configurations:

1. Define the SR-ODN template on the SR-TE head-end router.  
(Optional) If using Segment Routing Path Computation Element (SR-PCE) for path computation:
  - a. Configure SR-PCE. For detailed SR-PCE configuration information, see [Configure SR-PCE, on page 399](#).
  - b. Configure the head-end router as Path Computation Element Protocol (PCEP) Path Computation Client (PCC). For detailed PCEP PCC configuration information, see [Configure the Head-End Router as PCEP PCC](#).
2. Define BGP color extended communities. Refer to the "Implementing BGP" chapter in the *BGP Configuration Guide*.
3. Define routing policies (using routing policy language [RPL]) to set BGP color extended communities. Refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

The following RPL attach-points for setting/matching BGP color extended communities are supported:



**Note** The following table shows the supported RPL match operations; however, routing policies are required primarily to set BGP color extended community. Matching based on BGP color extended communities is performed automatically by ODN's on-demand color template.

| Attach Point      | Set | Match |
|-------------------|-----|-------|
| VRF export        | X   | X     |
| VRF import        | —   | X     |
| Neighbor-in       | X   | X     |
| Neighbor-out      | X   | X     |
| Inter-AFI export  | —   | X     |
| Inter-AFI import  | —   | X     |
| Default-originate | X   | —     |

4. Apply routing policies to a service. Refer to the "Implementing Routing Policy" chapter in the *Routing Configuration Guide*.

### Configure On-Demand Color Template

- Use the **on-demand color** *color* command to create an ODN template for the specified color value. The head-end router automatically follows the actions defined in the template upon arrival of BGP global or VPN routes with a BGP color extended community that matches the color value specified in the template.

The *color* range is from 1 to 4294967295.

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 10
```



**Note** Matching based on BGP color extended communities is performed automatically via ODN's on-demand color template. RPL routing policies are not required.

- Use the **on-demand color color dynamic** command to associate the template with on-demand SR policies with a locally computed dynamic path (by SR-TE head-end router utilizing its TE topology database) or centrally (by SR-PCE). The head-end router will first attempt to install the locally computed path; otherwise, it will use the path computed by the SR-PCE.

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 10 dynamic
```

- Use the **on-demand color color dynamic pcep** command to indicate that only the path computed by SR-PCE should be associated with the on-demand SR policy. With this configuration, local path computation is not attempted; instead the head-end router will only instantiate the path computed by the SR-PCE.

```
Router(config-sr-te)# on-demand color 10 dynamic pcep
```

### Configure Dynamic Path Optimization Objectives

- Use the **metric type {igp | te | latency}** command to configure the metric for use in path computation.

```
Router(config-sr-te-color-dyn)# metric type te
```

- Use the **metric margin {absolute value| relative percent}** command to configure the On-Demand dynamic path metric margin. The range for *value* and *percent* is from 0 to 2147483647.

```
Router(config-sr-te-color-dyn)# metric margin absolute 5
```

### Configure Dynamic Path Constraints

- Use the **disjoint-path group-id group-id type {link | node | srlg | srlg-node} [sub-id sub-id]** command to configure the disjoint-path constraints. The *group-id* and *sub-id* range is from 1 to 65535.

```
Router(config-sr-te-color-dyn)# disjoint-path group-id 775 type link
```

- Use the **affinity {include-any | include-all | exclude-any} {name WORD}** command to configure the affinity constraints.

```
Router(config-sr-te-color-dyn)# affinity exclude-any name CROSS
```

- Use the **maximum-sid-depth value** command to customize the maximum SID depth (MSD) constraints advertised by the router.

The default MSD *value* is equal to the maximum MSD supported by the platform (12).

```
Router(config-sr-te-color)# maximum-sid-depth 5
```

See [Customize MSD Value at PCC, on page 285](#) for information about SR-TE label imposition capabilities.

- Use the **constraints segments sid-algorithm** *algorithm-number* command to configure the SR Flexible Algorithm constraints. The *algorithm-number* range is from 128 to 255.

```
Router(config-sr-te-color)# constraints segments sid-algorithm 128
```

- Use the **constraints segments protection** {**protected-only** | **protected-preferred** | **unprotected-only** | **unprotected-preferred**} command to configure the Adj-SID protection behavior constraints.

```
Router(config-sr-te-color)# constraints segments protection protected-only
```

See [Segment Protection-Type Constraint, on page 269](#) for information about Adj-SID protection behavior.

## Configuring SR-ODN: Examples

### Configuring SR-ODN: Layer-3 Services Examples

The following examples show end-to-end configurations used in implementing SR-ODN on the head-end router.

#### Configuring ODN Color Templates: Example

Configure ODN color templates on routers acting as SR-TE head-end nodes. The following example shows various ODN color templates:

- color 10: minimization objective = te-metric
- color 20: minimization objective = igp-metric
- color 21: minimization objective = igp-metric; constraints = affinity
- color 22: minimization objective = te-metric; path computation at SR-PCE; constraints = affinity

```
segment-routing
traffic-eng
on-demand color 10
dynamic
metric
type te
!
!
on-demand color 20
dynamic
metric
type igp
!
!
on-demand color 21
dynamic
metric
type igp
!
affinity exclude-any
name CROSS
!
!
on-demand color 22
```



```

dynamic
 pcep
 !
 metric
 type te
 !
 affinity exclude-any
 name CROSS
 !
 !
 !
 on-demand color 128
 constraints
 segments
 sid-algorithm 128
 !
 !
 !
end

```

### Configuring BGP Color Extended Community Set: Example

The following example shows how to configure BGP color extended communities that are later applied to BGP service routes via route-policies.




---

**Note** In most common scenarios, egress PE routers that advertise BGP service routes apply (set) BGP color extended communities. However, color can also be set at the ingress PE router.

---

```

extcommunity-set opaque color10-te
 10
end-set
!
extcommunity-set opaque color20-igp
 20
end-set
!
extcommunity-set opaque color21-igp-excl-cross
 21
end-set
!

```

### Configuring RPL to Set BGP Color (Layer-3 Services): Examples

The following example shows various representative RPL definitions that set BGP color community.

The examples include the set color action only. The last RPL example performs the set color action for selected destinations based on a prefix-set.

```

route-policy SET_COLOR_LOW_LATENCY_TE
 set extcommunity color color10-te
 pass
end-policy
!
route-policy SET_COLOR_HI_BW
 set extcommunity color color20-igp
 pass
end-policy
!

```

```

prefix-set sample-set
 88.1.0.0/24
end-set
!
route-policy SET_COLOR_GLOBAL
 if destination in sample-set then
 set extcommunity color color10-te
 else
 pass
 endif
end-policy

```

### Applying RPL to BGP Services (Layer-3 Services): Example

The following example shows various RPLs that set BGP color community being applied to BGP Layer-3 VPN services (VPNv4/VPNv6) and BGP global.

- The L3VPN examples show the RPL applied at the VRF export attach-point.
- The BGP global example shows the RPL applied at the BGP neighbor-out attach-point.

```

vrf vrf_cust1
 address-family ipv4 unicast
 export route-policy SET_COLOR_LOW_LATENCY_TE
 !
 address-family ipv6 unicast
 export route-policy SET_COLOR_LOW_LATENCY_TE
 !
!
vrf vrf_cust2
 address-family ipv4 unicast
 export route-policy SET_COLOR_HI_BW
 !
 address-family ipv6 unicast
 export route-policy SET_COLOR_HI_BW
 !
!

router bgp 100
 neighbor-group BR-TO-RR
 address-family ipv4 unicast
 route-policy SET_COLOR_GLOBAL out
 !
!
end

```

### Verifying BGP VRF Information

Use the **show bgp vrf** command to display BGP prefix information for VRF instances. The following output shows the BGP VRF table including a prefix (88.1.1.0/24) with color 10 advertised by router 10.1.1.8.

```
RP/0/RP0/CPU0:R4# show bgp vrf vrf_cust1
```

```

BGP VRF vrf_cust1, state: Active
BGP Route Distinguisher: 10.1.1.4:101
VRF ID: 0x60000007
BGP router identifier 10.1.1.4, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000007 RD version: 282
BGP main routing table version 287

```

```

BGP NSR Initial initsync version 31 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.4:101 (default for vrf vrf_cust1)
*> 44.1.1.0/24 40.4.101.11 0 400 {1} i
*>i55.1.1.0/24 10.1.1.5 100 0 500 {1} i
*>i88.1.1.0/24 10.1.1.8 C:10 100 0 800 {1} i
*>i99.1.1.0/24 10.1.1.9 100 0 800 {1} i

Processed 4 prefixes, 4 paths

```

The following output displays the details for prefix 88.1.1.0/24. Note the presence of BGP extended color community 10, and that the prefix is associated with an SR policy with color 10 and BSID value of 24036.

```

RP/0/RP0/CPU0:R4# show bgp vrf vrf_cust1 88.1.1.0/24

BGP routing table entry for 88.1.1.0/24, Route Distinguisher: 10.1.1.4:101
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 282 282
Last Modified: May 20 09:23:34.112 for 00:06:03
Paths: (1 available, best #1)
 Advertised to CE peers (in unique update groups):
 40.4.101.11
 Path #1: Received by speaker 0
 Advertised to CE peers (in unique update groups):
 40.4.101.11
 800 {1}
10.1.1.8 C:10 (bsid:24036) (metric 20) from 10.1.1.55 (10.1.1.8)
 Received Label 24012
 Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
 Received Path ID 0, Local Path ID 1, version 273
Extended community: Color:10 RT:100:1
 Originator: 10.1.1.8, Cluster list: 10.1.1.55
SR policy color 10, up, registered, bsid 24036, if-handle 0x08000024

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 10.1.1.8:101

```

### Verifying Forwarding (CEF) Table

Use the **show cef vrf** command to display the contents of the CEF table for the VRF instance. Note that prefix 88.1.1.0/24 points to the BSID label corresponding to an SR policy. Other non-colored prefixes, such as 55.1.1.0/24, point to BGP next-hop.

```

RP/0/RP0/CPU0:R4# show cef vrf vrf_cust1

Prefix Next Hop Interface

0.0.0.0/0 drop default handler
0.0.0.0/32 broadcast
40.4.101.0/24 attached TenGigE0/0/0/0.101
40.4.101.0/32 broadcast TenGigE0/0/0/0.101
40.4.101.4/32 receive TenGigE0/0/0/0.101
40.4.101.11/32 40.4.101.11/32 TenGigE0/0/0/0.101
40.4.101.255/32 broadcast TenGigE0/0/0/0.101
44.1.1.0/24 40.4.101.11/32 <recursive>
55.1.1.0/24 10.1.1.5/32 <recursive>

```

```

88.1.1.0/24 24036 (via-label) <recursive>
99.1.1.0/24 10.1.1.9/32 <recursive>
224.0.0.0/4 0.0.0.0/32
224.0.0.0/24 receive
255.255.255.255/32 broadcast

```

The following output displays CEF details for prefix 88.1.1.0/24. Note that the prefix is associated with an SR policy with BSID value of 24036.

```

RP/0/RP0/CPU0:R4# show cef vrf vrf_cust1 88.1.1.0/24

88.1.1.0/24, version 51, internal 0x5000001 0x0 (ptr 0x98c60ddc) [1], 0x0 (0x0), 0x208
(0x98425268)
Updated May 20 09:23:34.216
Prefix Len 24, traffic index 0, precedence n/a, priority 3
via local-label 24036, 5 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x97091ec0 0x0]
recursion-via-label
next hop VRF - 'default', table - 0xe0000000
next hop via 24036/0/21
next hop srte_c_10_ep labels imposed {ImplNull 24012}

```

### Verifying SR Policy

Use the **show segment-routing traffic-eng policy** command to display SR policy information.

The following outputs show the details of an on-demand SR policy that was triggered by prefixes with color 10 advertised by node 10.1.1.8.

```

RP/0/RP0/CPU0:R4# show segment-routing traffic-eng policy color 10 tabular

Color Endpoint Admin Oper Binding

10 10.1.1.8 up up 24036

```

The following outputs show the details of the on-demand SR policy for BSID 24036.



**Note** There are 2 candidate paths associated with this SR policy: the path that is computed by the head-end router (with preference 200), and the path that is computed by the SR-PCE (with preference 100). The candidate path with the highest preference is the active candidate path (highlighted below) and is installed in forwarding.

```

RP/0/RP0/CPU0:R4# show segment-routing traffic-eng policy binding-sid 24036

SR-TE policy database

Color: 10, End-point: 10.1.1.8
Name: srte_c_10_ep_10.1.1.8
Status:
Admin: up Operational: up for 4d14h (since Jul 3 20:28:57.840)
Candidate-paths:
Preference: 200 (BGP ODN) (active)
Requested BSID: dynamic
PCC info:
Symbolic name: bgp_c_10_ep_10.1.1.8_discr_200
PLSP-ID: 12

```

```

Dynamic (valid)
 Metric Type: TE, Path Accumulated Metric: 30
 16009 [Prefix-SID, 10.1.1.9]
 16008 [Prefix-SID, 10.1.1.8]
Preference: 100 (BGP ODN)
Requested BSID: dynamic
PCC info:
 Symbolic name: bgp_c_10_ep_10.1.1.8_discr_100
 PLSP-ID: 11
Dynamic (pce 10.1.1.57) (valid)
 Metric Type: TE, Path Accumulated Metric: 30
 16009 [Prefix-SID, 10.1.1.9]
 16008 [Prefix-SID, 10.1.1.8]
Attributes:
 Binding SID: 24036
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```

### Verifying SR Policy Forwarding

Use the **show segment-routing traffic-eng forwarding policy** command to display the SR policy forwarding information.

The following outputs show the forwarding details for an on-demand SR policy that was triggered by prefixes with color 10 advertised by node 10.1.1.8.

```
RP/0/RP0/CPU0:R4# show segment-routing traffic-eng forwarding policy binding-sid 24036
tabular
```

| Color | Endpoint | Segment<br>List | Outgoing<br>Label | Outgoing<br>Interface | Next Hop | Bytes<br>Switched | Pure<br>Backup |
|-------|----------|-----------------|-------------------|-----------------------|----------|-------------------|----------------|
| 10    | 10.1.1.8 | dynamic         | 16009             | Gi0/0/0/4             | 10.4.5.5 | 0                 |                |
|       |          |                 | 16001             | Gi0/0/0/5             | 11.4.8.8 | 0                 | Yes            |

```
RP/0/RP0/CPU0:R4# show segment-routing traffic-eng forwarding policy binding-sid 24036
detail
```

```
Mon Jul 8 11:56:46.887 PST
```

```
SR-TE Policy Forwarding database
```

```

Color: 10, End-point: 10.1.1.8
Name: srte_c_10_ep_10.1.1.8
Binding SID: 24036
Segment Lists:
 SL[0]:
 Name: dynamic
 Paths:
 Path[0]:
 Outgoing Label: 16009
 Outgoing Interface: GigabitEthernet0/0/0/4
 Next Hop: 10.4.5.5
 Switched Packets/Bytes: 0/0
 FRR Pure Backup: No
 Label Stack (Top -> Bottom): { 16009, 16008 }
 Path-id: 1 (Protected), Backup-path-id: 2, Weight: 64
 Path[1]:
 Outgoing Label: 16001
 Outgoing Interface: GigabitEthernet0/0/0/5
 Next Hop: 11.4.8.8
 Switched Packets/Bytes: 0/0

```

```

FRR Pure Backup: Yes
Label Stack (Top -> Bottom): { 16001, 16009, 16008 }
Path-id: 2 (Pure-Backup), Weight: 64
Policy Packets/Bytes Switched: 0/0
Local label: 80013

```

## Configuring SR-ODN: EVPN Services Examples

### Configuring BGP Color Extended Community Set: Example

The following example shows how to configure BGP color extended communities that are later applied to BGP service routes via route-policies.

```

extcommunity-set opaque color-44
 44
end-set

extcommunity-set opaque color-55
 55
end-set

extcommunity-set opaque color-77
 77
end-set

extcommunity-set opaque color-88
 88
end-set

```

### Configuring RPL to Set BGP Color (EVPN Services): Examples

The following examples shows various representative RPL definitions that set BGP color community.

The following RPL examples match on EVPN route-types and then set the BGP color extended community.

```

route-policy sample-export-rpl
 if evpn-route-type is 1 then
 set extcommunity color color-44
 endif
 if evpn-route-type is 3 then
 set extcommunity color color-55
 endif
end-policy

route-policy sample-import-rpl
 if evpn-route-type is 1 then
 set extcommunity color color-77
 elseif evpn-route-type is 3 then
 set extcommunity color color-88
 else
 pass
 endif
end-policy

```

The following RPL example sets BGP color extended community while matching on the following:

- Route Distinguisher (RD)
- Ethernet Segment Identifier (ESI)
- Ethernet Tag (ETAG)
- EVPN route-types

```

route-policy sample-bgpneighbor-rpl
 if rd in (10.1.1.1:3504) then
 set extcommunity color color3504
 elseif rd in (10.1.1.1:3505) then
 set extcommunity color color3505
 elseif rd in (10.1.1.1:3506) then
 set extcommunity color color99996
 elseif esi in (0010.0000.0000.0000.1201) and rd in (10.1.1.1:3508) then
 set extcommunity color color3508
 elseif etag in (30509) and rd in (10.1.1.1:3509) then
 set extcommunity color color3509
 elseif etag in (0) and rd in (10.1.1.1:2001) and evpn-route-type is 1 then
 set extcommunity color color82001
 elseif etag in (0) and rd in (10.1.1.1:2001) and evpn-route-type is 3 then
 set extcommunity color color92001
 endif
 pass
end-policy

```

### Applying RPL to BGP Services (EVPN Services): Example

The following examples show various RPLs that set BGP color community being applied to EVPN services.

The following 2 examples show the RPL applied at the EVI export and import attach-points.



**Note** RPLs applied under EVI import or export attach-point also support matching on the following:

- Ethernet Segment Identifier (ESI)
- Ethernet Tag (ETAG)
- EVPN-Originator

```

evpn
 evi 101
 bgp
 route-target 101:1
 route-target import 100:1
 route-target export 101:1
 route-policy import sample-import-rpl
 !
 advertise-mac
 !
!
 evi 102
 bgp
 route-target 102:1
 route-target import 100:2
 route-target export 102:1
 route-policy export sample-export-rpl
 !
 advertise-mac
 !
!
!

```

The following example shows the RPL applied at the BGP neighbor-out attach-point.



**Note** RPLs defined under BGP neighbor-out attach-point also support matching on the following:

- EVPN-Originator

```
router bgp 100
 bgp router-id 10.1.1.1
 address-family l2vpn evpn
 !
 neighbor-group evpn-rr
 remote-as 100
 update-source Loopback0
 address-family l2vpn evpn
 !
 neighbor 10.10.10.10
 use neighbor-group evpn-rr
 address-family l2vpn evpn
 route-policy sample-bgpneighbor-rpl out
```

### Configuring SR-ODN for EVPN-VPWS: Use Case

This use case shows how to set up a pair of ELINE services using EVPN-VPWS between two sites. Services are carried over SR policies that must not share any common links along their paths (link-disjoint). The SR policies are triggered on-demand based on ODN principles. An SR-PCE computes the disjoint paths.

This use case uses the following topology with 2 sites: Site 1 with nodes A and B, and Site 2 with nodes C and D.

**Figure 12: Topology for Use Case: SR-ODN for EVPN-VPWS**






Table 31: Use Case Parameters

|                                                                                                         |                                                                                                                                            |                                                                                                                                            |
|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>IP Addresses of Loopback0 (Lo0) Interfaces</b>                                                       | SR-PCE Lo0: 10.1.1.207                                                                                                                     |                                                                                                                                            |
|                                                                                                         | Site 1: <ul style="list-style-type: none"> <li>• Node A Lo0: 10.1.1.5</li> <li>• Node B Lo0: 10.1.1.6</li> </ul>                           | Site 2: <ul style="list-style-type: none"> <li>• Node C Lo0: 10.1.1.2</li> <li>• Node D Lo0: 10.1.1.4</li> </ul>                           |
| <b>EVPN-VPWS Service Parameters</b>                                                                     | ELINE-1: <ul style="list-style-type: none"> <li>• EVPN-VPWS EVI 100</li> <li>• Node A: AC-ID = 11</li> <li>• Node C: AC-ID = 21</li> </ul> | ELINE-2: <ul style="list-style-type: none"> <li>• EVPN-VPWS EVI 101</li> <li>• Node B: AC-ID = 12</li> <li>• Node D: AC-ID = 22</li> </ul> |
| <b>ODN BGP Color Extended Communities</b>                                                               | Site 1 routers (Nodes A and B): <ul style="list-style-type: none"> <li>• set color 10000</li> <li>• match color 11000</li> </ul>           | Site 2 routers (Nodes C and D): <ul style="list-style-type: none"> <li>• set color 11000</li> <li>• match color 10000</li> </ul>           |
| <b>Note</b><br>These colors are associated with the EVPN route-type 1 routes of the EVPN-VPWS services. |                                                                                                                                            |                                                                                                                                            |
| <b>PCEP LSP Disjoint-Path Association Group ID</b>                                                      | Site 1 to Site 2 LSPs (from Node A to Node C/from Node B to Node D): <ul style="list-style-type: none"> <li>• group-id = 775</li> </ul>    | Site 2 to Site 1 LSPs (from Node C to Node A/from Node D to Node B): <ul style="list-style-type: none"> <li>• group-id = 776</li> </ul>    |

The use case provides configuration and verification outputs for all devices.

| Configuration                                             | Verification                                             |
|-----------------------------------------------------------|----------------------------------------------------------|
| <a href="#">Configuration: SR-PCE, on page 240</a>        | <a href="#">Verification: SR-PCE, on page 245</a>        |
| <a href="#">Configuration: Site 1 Node A, on page 241</a> | <a href="#">Verification: Site 1 Node A, on page 249</a> |
| <a href="#">Configuration: Site 1 Node B, on page 242</a> | <a href="#">Verification: Site 1 Node B, on page 251</a> |
| <a href="#">Configuration: Site 2 Node C, on page 243</a> | <a href="#">Verification: Site 2 Node C, on page 254</a> |
| <a href="#">Configuration: Site 2 Node D, on page 244</a> | <a href="#">Verification: Site 2 Node D, on page 257</a> |

### Configuration: SR-PCE

For cases when PCC nodes support, or signal, PCEP association-group object to indicate the pair of LSPs in a disjoint set, there is no extra configuration required at the SR-PCE to trigger disjoint-path computation.



**Note** SR-PCE also supports disjoint-path computation for cases when PCC nodes do not support PCEP association-group object. See [Configure the Disjoint Policy \(Optional\), on page 401](#) for more information.

**Configuration: Site 1 Node A**

This section depicts relevant configuration of Node A at Site 1. It includes service configuration, BGP color extended community, and RPL. It also includes the corresponding ODN template required to achieve the disjointness SLA.

Nodes in Site 1 are configured to set color 10000 on originating EVPN routes, while matching color 11000 on incoming EVPN routes from routers located at Site 2.

Since both nodes in Site 1 request path computation from SR-PCE using the same disjoint-path group-id (775), the PCE will attempt to compute disjointness for the pair of LSPs originating from Site 1 toward Site 2.

```
/* EVPN-VPWS configuration */

interface GigabitEthernet0/0/0/3.2500 l2transport
 encapsulation dot1q 2500
 rewrite ingress tag pop 1 symmetric
!
l2vpn
 xconnect group evpn_vpws_group
 p2p evpn_vpws_100
 interface GigabitEthernet0/0/0/3.2500
 neighbor evpn evi 100 target 21 source 11
 !
 !
!
!

/* BGP color community and RPL configuration */

extcommunity-set opaque color-10000
 10000
end-set
!
route-policy SET_COLOR_EVPN_VPWS
 if evpn-route-type is 1 and rd in (ios-regex '.*...*(100)') then
 set extcommunity color color-10000
 endif
 pass
end-policy
!
router bgp 65000
 neighbor 10.1.1.253
 address-family l2vpn evpn
 route-policy SET_COLOR_EVPN_VPWS out
 !
!
!

/* ODN template configuration */

segment-routing
 traffic-eng
 on-demand color 11000
 dynamic
 pcep
 !
 metric
 type igp
 !
 disjoint-path group-id 775 type link
!
```

```

!
!
!

```

### Configuration: Site 1 Node B

This section depicts relevant configuration of Node B at Site 1.

```

/* EVPN-VPWS configuration */

interface TenGigE0/3/0/0/8.2500 l2transport
 encapsulation dot1q 2500
 rewrite ingress tag pop 1 symmetric
!
l2vpn
 xconnect group evpn_vpws_group
 p2p evpn_vpws_101
 interface TenGigE0/3/0/0/8.2500
 neighbor evpn evi 101 target 22 source 12
 !
 !
 !
!

/* BGP color community and RPL configuration */

extcommunity-set opaque color-10000
 10000
end-set
!
route-policy SET_COLOR_EVPN_VPWS
 if evpn-route-type is 1 and rd in (ios-regex '.*...*...:(101)') then
 set extcommunity color color-10000
 endif
 pass
end-policy
!
router bgp 65000
 neighbor 10.1.1.253
 address-family l2vpn evpn
 route-policy SET_COLOR_EVPN_VPWS out
 !
 !
!

/* ODN template configuration */

segment-routing
 traffic-eng
 on-demand color 11000
 dynamic
 pcep
 !
 metric
 type igp
 !
 disjoint-path group-id 775 type link
 !
 !
!
!
!

```

### Configuration: Site 2 Node C

This section depicts relevant configuration of Node C at Site 2. It includes service configuration, BGP color extended community, and RPL. It also includes the corresponding ODN template required to achieve the disjointness SLA.

Nodes in Site 2 are configured to set color 11000 on originating EVPN routes, while matching color 10000 on incoming EVPN routes from routers located at Site 1.

Since both nodes on Site 2 request path computation from SR-PCE using the same disjoint-path group-id (776), the PCE will attempt to compute disjointness for the pair of LSPs originating from Site 2 toward Site 1.

```
/* EVPN-VPWS configuration */

interface GigabitEthernet0/0/0/3.2500 l2transport
 encapsulation dot1q 2500
 rewrite ingress tag pop 1 symmetric
!
l2vpn
 xconnect group evpn_vpws_group
 p2p evpn_vpws_100
 interface GigabitEthernet0/0/0/3.2500
 neighbor evpn evi 100 target 11 source 21
 !
 !
!
!

/* BGP color community and RPL configuration */

extcommunity-set opaque color-11000
 11000
end-set
!
route-policy SET_COLOR_EVPN_VPWS
 if evpn-route-type is 1 and rd in (ios-regex '.*..*..*..*(100)') then
 set extcommunity color color-11000
 endif
 pass
end-policy
!
router bgp 65000
 neighbor 10.1.1.253
 address-family l2vpn evpn
 route-policy SET_COLOR_EVPN_VPWS out
 !
!
!

/* ODN template configuration */

segment-routing
 traffic-eng
 on-demand color 10000
 dynamic
 pcep
 !
 metric
 type igp
 !
 disjoint-path group-id 776 type link
!
```

```

!
!
!

```

### Configuration: Site 2 Node D

This section depicts relevant configuration of Node D at Site 2.

```

/* EVPN-VPWS configuration */

interface GigabitEthernet0/0/0/1.2500 l2transport
 encapsulation dot1q 2500
 rewrite ingress tag pop 1 symmetric
!
l2vpn
 xconnect group evpn_vpws_group
 p2p evpn_vpws_101
 interface GigabitEthernet0/0/0/1.2500
 neighbor evpn evi 101 target 12 source 22
 !
 !
 !
!

/* BGP color community and RPL configuration */

extcommunity-set opaque color-11000
 11000
end-set
!
route-policy SET_COLOR_EVPN_VPWS
 if evpn-route-type is 1 and rd in (ios-regex '.*...*...:(101)') then
 set extcommunity color color-11000
 endif
 pass
end-policy
!
router bgp 65000
 neighbor 10.1.1.253
 address-family l2vpn evpn
 route-policy SET_COLOR_EVPN_VPWS out
 !
 !
!

/* ODN template configuration */

segment-routing
 traffic-eng
 on-demand color 10000
 dynamic
 pcep
 !
 metric
 type igp
 !
 disjoint-path group-id 776 type link
 !
 !
!
!
!

```

**Verification: SR-PCE**

Use the **show pce ipv4 peer** command to display the SR-PCE's PCEP peers and session status. SR-PCE performs path computation for the 4 nodes depicted in the use-case.

```
RP/0/0/CPU0:SR-PCE# show pce ipv4 peer
Mon Jul 15 19:41:43.622 UTC

PCE's peer database:

Peer address: 10.1.1.2
 State: Up
 Capabilities: Stateful, Segment-Routing, Update, Instantiation

Peer address: 10.1.1.4
 State: Up
 Capabilities: Stateful, Segment-Routing, Update, Instantiation

Peer address: 10.1.1.5
 State: Up
 Capabilities: Stateful, Segment-Routing, Update, Instantiation

Peer address: 10.1.1.6
 State: Up
 Capabilities: Stateful, Segment-Routing, Update, Instantiation
```

Use the **show pce association group-id** command to display information for the pair of LSPs assigned to a given association group-id value.

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. In particular, disjoint LSPs from site 1 to site 2 are identified by association group-id 775. The output includes high-level information for LSPs associated to this group-id:

- At Node A (10.1.1.5): LSP symbolic name = `bgp_c_11000_ep_10.1.1.2_discr_100`
- At Node B (10.1.1.6): LSP symbolic name = `bgp_c_11000_ep_10.1.1.4_discr_100`

In this case, the SR-PCE was able to achieve the desired disjointness level; therefore the Status is shown as "Satisfied".

```
RP/0/0/CPU0:SR-PCE# show pce association group-id 775
Thu Jul 11 03:52:20.770 UTC

PCE's association database:

Association: Type Link-Disjoint, Group 775, Not Strict
Associated LSPs:
 LSP[0]:
 PCC 10.1.1.6, tunnel name bgp_c_11000_ep_10.1.1.4_discr_100, PLSP ID 18, tunnel ID 17,
 LSP ID 3, Configured on PCC
 LSP[1]:
 PCC 10.1.1.5, tunnel name bgp_c_11000_ep_10.1.1.2_discr_100, PLSP ID 18, tunnel ID 18,
 LSP ID 3, Configured on PCC
Status: Satisfied
```

Use the **show pce lsp** command to display detailed information of an LSP present in the PCE's LSP database. This output shows details for the LSP at Node A (10.1.1.5) that is used to carry traffic of EVPN VPWS EVI 100 towards node C (10.1.1.2).

```
RP/0/0/CPU0:SR-PCE# show pce lsp pcc ipv4 10.1.1.5 name bgp_c_11000_ep_10.1.1.2_discr_100
Thu Jul 11 03:58:45.903 UTC

PCE's tunnel database:
```

```

PCC 10.1.1.5:

Tunnel Name: bgp_c_11000_ep_10.1.1.2_discr_100
Color: 11000
Interface Name: srte_c_11000_ep_10.1.1.2
LSPs:
 LSP[0]:
 source 10.1.1.5, destination 10.1.1.2, tunnel ID 18, LSP ID 3
 State: Admin up, Operation up
 Setup type: Segment Routing
 Binding SID: 80037
 Maximum SID Depth: 10
 Absolute Metric Margin: 0
 Relative Metric Margin: 0%
 Preference: 100
 Bandwidth: signaled 0 kbps, applied 0 kbps
 PCEP information:
 PLSP-ID 0x12, flags: D:1 S:0 R:0 A:1 O:1 C:0
 LSP Role: Exclude LSP
 State-sync PCE: None
 PCC: 10.1.1.5
 LSP is subdelegated to: None
 Reported path:
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Adj, Label 80003, Address: local 11.5.8.5 remote 11.5.8.8
 SID[1]: Node, Label 16007, Address 10.1.1.7
 SID[2]: Node, Label 16002, Address 10.1.1.2
 Computed path: (Local PCE)
 Computed Time: Thu Jul 11 03:49:48 UTC 2019 (00:08:58 ago)
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Adj, Label 80003, Address: local 11.5.8.5 remote 11.5.8.8
 SID[1]: Node, Label 16007, Address 10.1.1.7
 SID[2]: Node, Label 16002, Address 10.1.1.2
 Recorded path:
 None
 Disjoint Group Information:
 Type Link-Disjoint, Group 775

```

This output shows details for the LSP at Node B (10.1.1.6) that is used to carry traffic of EVPN VPWS EVI 101 towards node D (10.1.1.4).

```

RP/0/0/CPU0:SR-PCE# show pce lsp pcc ipv4 10.1.1.6 name bgp_c_11000_ep_10.1.1.4_discr_100
Thu Jul 11 03:58:56.812 UTC

```

PCE's tunnel database:

```

PCC 10.1.1.6:

Tunnel Name: bgp_c_11000_ep_10.1.1.4_discr_100
Color: 11000
Interface Name: srte_c_11000_ep_10.1.1.4
LSPs:
 LSP[0]:
 source 10.1.1.6, destination 10.1.1.4, tunnel ID 17, LSP ID 3
 State: Admin up, Operation up
 Setup type: Segment Routing
 Binding SID: 80061
 Maximum SID Depth: 10
 Absolute Metric Margin: 0
 Relative Metric Margin: 0%
 Preference: 100
 Bandwidth: signaled 0 kbps, applied 0 kbps
 PCEP information:

```

```

 PLSP-ID 0x12, flags: D:1 S:0 R:0 A:1 O:1 C:0
LSP Role: Disjoint LSP
State-sync PCE: None
PCC: 10.1.1.6
LSP is subdelegated to: None
Reported path:
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16001, Address 10.1.1.1
 SID[1]: Node, Label 16004, Address 10.1.1.4
Computed path: (Local PCE)
 Computed Time: Thu Jul 11 03:49:48 UTC 2019 (00:09:08 ago)
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16001, Address 10.1.1.1
 SID[1]: Node, Label 16004, Address 10.1.1.4
Recorded path:
 None
Disjoint Group Information:
 Type Link-Disjoint, Group 775

```

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. In particular, disjoint LSPs from site 2 to site 1 are identified by association group-id 776. The output includes high-level information for LSPs associated to this group-id:

- At Node C (10.1.1.2): LSP symbolic name = `bgp_c_10000_ep_10.1.1.5_discr_100`
- At Node D (10.1.1.4): LSP symbolic name = `bgp_c_10000_ep_10.1.1.6_discr_100`

In this case, the SR-PCE was able to achieve the desired disjointness level; therefore, the Status is shown as "Satisfied".

```

RP/0/0/CPU0:SR-PCE# show pce association group-id 776
Thu Jul 11 03:52:24.370 UTC

```

PCE's association database:

**Association: Type Link-Disjoint, Group 776, Not Strict**

Associated LSPs:

```

LSP[0]:
 PCC 10.1.1.4, tunnel name bgp_c_10000_ep_10.1.1.6_discr_100, PLSP ID 16, tunnel ID 14,
 LSP ID 1, Configured on PCC
LSP[1]:
 PCC 10.1.1.2, tunnel name bgp_c_10000_ep_10.1.1.5_discr_100, PLSP ID 6, tunnel ID 21,
 LSP ID 3, Configured on PCC

```

**Status: Satisfied**

Use the **show pce lsp** command to display detailed information of an LSP present in the PCE's LSP database. This output shows details for the LSP at Node C (10.1.1.2) that is used to carry traffic of EVPN VPWS EVI 100 towards node A (10.1.1.5).

```

RP/0/0/CPU0:SR-PCE# show pce lsp pcc ipv4 10.1.1.2 name bgp_c_10000_ep_10.1.1.5_discr_100
Thu Jul 11 03:55:21.706 UTC

```

PCE's tunnel database:

PCC 10.1.1.2:

Tunnel Name: `bgp_c_10000_ep_10.1.1.5_discr_100`

**Color: 10000**

**Interface Name: srte\_c\_10000\_ep\_10.1.1.5**

LSPs:

```

LSP[0]:
 source 10.1.1.2, destination 10.1.1.5, tunnel ID 21, LSP ID 3
 State: Admin up, Operation up

```



```

Setup type: Segment Routing
Binding SID: 80052
Maximum SID Depth: 10
Absolute Metric Margin: 0
Relative Metric Margin: 0%
Preference: 100
Bandwidth: signaled 0 kbps, applied 0 kbps
PCEP information:
 PLSP-ID 0x6, flags: D:1 S:0 R:0 A:1 O:1 C:0
LSP Role: Exclude LSP
State-sync PCE: None
PCC: 10.1.1.2
LSP is subdelegated to: None
Reported path:
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16007, Address 10.1.1.7
 SID[1]: Node, Label 16008, Address 10.1.1.8
 SID[2]: Adj, Label 80005, Address: local 11.5.8.8 remote 11.5.8.5
Computed path: (Local PCE)
 Computed Time: Thu Jul 11 03:50:03 UTC 2019 (00:05:18 ago)
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16007, Address 10.1.1.7
 SID[1]: Node, Label 16008, Address 10.1.1.8
 SID[2]: Adj, Label 80005, Address: local 11.5.8.8 remote 11.5.8.5
Recorded path:
 None
Disjoint Group Information:
 Type Link-Disjoint, Group 776

```

This output shows details for the LSP at Node D (10.1.1.4) used to carry traffic of EVPN VPWS EVI 101 towards node B (10.1.1.6).

```

RP/0/0/CPU0:SR-PCE# show pce lsp pcc ipv4 10.1.1.4 name bgp_c_10000_ep_10.1.1.6_discr_100
Thu Jul 11 03:55:23.296 UTC

```

PCE's tunnel database:

-----  
PCC 10.1.1.4:

Tunnel Name: bgp\_c\_10000\_ep\_10.1.1.6\_discr\_100

**Color: 10000**

**Interface Name: srte\_c\_10000\_ep\_10.1.1.6**

LSPs:

```

LSP[0]:
 source 10.1.1.4, destination 10.1.1.6, tunnel ID 14, LSP ID 1
 State: Admin up, Operation up
 Setup type: Segment Routing
 Binding SID: 80047
 Maximum SID Depth: 10
 Absolute Metric Margin: 0
 Relative Metric Margin: 0%
 Preference: 100
 Bandwidth: signaled 0 kbps, applied 0 kbps
 PCEP information:
 PLSP-ID 0x10, flags: D:1 S:0 R:0 A:1 O:1 C:0
 LSP Role: Disjoint LSP
 State-sync PCE: None
 PCC: 10.1.1.4
 LSP is subdelegated to: None
 Reported path:
 Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16001, Address 10.1.1.1
 SID[1]: Node, Label 16006, Address 10.1.1.6
 Computed path: (Local PCE)

```

```

Computed Time: Thu Jul 11 03:50:03 UTC 2019 (00:05:20 ago)
Metric type: IGP, Accumulated Metric 40
 SID[0]: Node, Label 16001, Address 10.1.1.1
 SID[1]: Node, Label 16006, Address 10.1.1.6
Recorded path:
 None
Disjoint Group Information:
 Type Link-Disjoint, Group 776

```

### Verification: Site 1 Node A

This section depicts verification steps at Node A.

Use the **show bgp l2vpn evpn** command to display BGP prefix information for EVPN-VPWS EVI 100 (rd 10.1.1.5:100). The output includes an EVPN route-type 1 route with color 11000 originated at Node C (10.1.1.2).

```

RP/0/RSP0/CPU0:Node-A# show bgp l2vpn evpn rd 10.1.1.5:100
Wed Jul 10 18:57:57.704 PST
BGP router identifier 10.1.1.5, local AS number 65000
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 360
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.5:100 (default for vrf VPWS:100)
*> [1][0000.0000.0000.0000.0000][11]/120
 0.0.0.0 0 i
*>i [1][0000.0000.0000.0000.0000][21]/120
 10.1.1.2 C:11000 100 0 i

```

The following output displays the details for the incoming EVPN RT1. Note the presence of BGP extended color community 11000, and that the prefix is associated with an SR policy with color 11000 and BSID value of 80044.

```

RP/0/RSP0/CPU0:Node-A# show bgp l2vpn evpn rd 10.1.1.5:100
[1][0000.0000.0000.0000.0000][21]/120
Wed Jul 10 18:57:58.107 PST
BGP routing table entry for [1][0000.0000.0000.0000.0000][21]/120, Route Distinguisher:
10.1.1.5:100
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 360 360
Last Modified: Jul 10 18:36:18.369 for 00:21:40
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local
 10.1.1.2 C:11000 (bsid:80044) (metric 40) from 10.1.1.253 (10.1.1.2)
 Received Label 80056
 Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, rib-install
 Received Path ID 0, Local Path ID 1, version 358

```

```

Extended community: Color:11000 RT:65000:100
Originator: 10.1.1.2, Cluster list: 10.1.1.253
SR policy color 11000, up, registered, bsid 80044, if-handle 0x00001b20

```

Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.1.1.2:100

Use the **show l2vpn xconnect** command to display the state associated with EVPN-VPWS EVI 100 service.

```

RP/0/RSP0/CPU0:Node-A# show l2vpn xconnect group evpn_vpws_group
Wed Jul 10 18:58:02.333 PST
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
 SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

```

| XConnect Group  |  | Name          | ST | Segment 1 Description | ST | Segment 2 Description | ST |
|-----------------|--|---------------|----|-----------------------|----|-----------------------|----|
| evpn_vpws_group |  | evpn_vpws_100 | UP | Gi0/0/0/3.2500        | UP | EVPN 100,21,10.1.1.2  | UP |

The following output shows the details for the service. Note that the service is associated with the on-demand SR policy with color 11000 and end-point 10.1.1.2 (node C).

```

RP/0/RSP0/CPU0:Node-A# show l2vpn xconnect group evpn_vpws_group xc-name evpn_vpws_100
detail
Wed Jul 10 18:58:02.755 PST

```

```

Group evpn_vpws_group, XC evpn_vpws_100, state is up; Interworking none
AC: GigabitEthernet0/0/0/3.2500, state is up
 Type VLAN; Num Ranges: 1
 Rewrite Tags: []
 VLAN ranges: [2500, 2500]
 MTU 1500; XC ID 0x120000c; interworking none
 Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0
 drops: illegal VLAN 0, illegal length 0
EVPN: neighbor 10.1.1.2, PW ID: evi 100, ac-id 21, state is up (established)
 XC ID 0xa0000007
 Encapsulation MPLS
 Source address 10.1.1.5
 Encap type Ethernet, control word enabled
 Sequencing not set
Preferred path Active : SR TE srte_c_11000_ep_10.1.1.2, On-Demand, fallback enabled
Tunnel : Up
Load Balance Hashing: src-dst-mac

```

| EVPN         | Local    | Remote   |
|--------------|----------|----------|
| Label        | 80040    | 80056    |
| MTU          | 1500     | 1500     |
| Control word | enabled  | enabled  |
| AC ID        | 11       | 21       |
| EVPN type    | Ethernet | Ethernet |

```

Create time: 10/07/2019 18:31:30 (1d17h ago)
Last time status changed: 10/07/2019 19:42:00 (1d16h ago)
Last time PW went down: 10/07/2019 19:40:55 (1d16h ago)
Statistics:
 packets: received 0, sent 0

```

```
bytes: received 0, sent 0
```

Use the **show segment-routing traffic-eng policy** command with **tabular** option to display SR policy summary information.

The following output shows the on-demand SR policy with BSID 80044 that was triggered by EVPN RT1 prefix with color 11000 advertised by node C (10.1.1.2).

```
RP/0/RSP0/CPU0:Node-A# show segment-routing traffic-eng policy color 11000 tabular
Wed Jul 10 18:58:00.732 PST
```

| Color | Endpoint | Admin<br>State | Oper<br>State | Binding<br>SID |
|-------|----------|----------------|---------------|----------------|
| 11000 | 10.1.1.2 | up             | up            | 80044          |

The following output shows the details for the on-demand SR policy. Note that the SR policy's active candidate path (preference 100) is computed by SR-PCE (10.1.1.207).

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. Specifically, from site 1 to site 2, LSP at Node A (srte\_c\_11000\_ep\_10.1.1.2) is link-disjoint from LSP at Node B (srte\_c\_11000\_ep\_10.1.1.4).

```
RP/0/RSP0/CPU0:Node-A# show segment-routing traffic-eng policy color 11000
Wed Jul 10 19:15:47.217 PST
```

```
SR-TE policy database

```

```
Color: 11000, End-point: 10.1.1.2
 Name: srte_c_11000_ep_10.1.1.2
 Status:
 Admin: up Operational: up for 00:39:31 (since Jul 10 18:36:00.471)
 Candidate-paths:
 Preference: 200 (BGP ODN) (shutdown)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_11000_ep_10.1.1.2_discr_200
 PLSP-ID: 19
 Dynamic (invalid)
 Preference: 100 (BGP ODN) (active)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_11000_ep_10.1.1.2_discr_100
 PLSP-ID: 18
 Dynamic (pce 10.1.1.207) (valid)
 Metric Type: IGP, Path Accumulated Metric: 40
 80003 [Adjacency-SID, 11.5.8.5 - 11.5.8.8]
 16007 [Prefix-SID, 10.1.1.7]
 16002 [Prefix-SID, 10.1.1.2]
 Attributes:
 Binding SID: 80044
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes
```

### Verification: Site 1 Node B

This section depicts verification steps at Node B.

Use the **show bgp l2vpn evpn** command to display BGP prefix information for EVPN-VPWS EVI 101 (rd 10.1.1.6:101). The output includes an EVPN route-type 1 route with color 11000 originated at Node D (10.1.1.4).

```
RP/0/RSP0/CPU0:Node-B# show bgp l2vpn evpn rd 10.1.1.6:101
Wed Jul 10 19:08:54.964 PST
BGP router identifier 10.1.1.6, local AS number 65000
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 322
BGP NSR Initial initsync version 7 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.6:101 (default for vrf VPWS:101)
*> [1][0000.0000.0000.0000.0000][12]/120
 0.0.0.0 0 i
*>i[1][0000.0000.0000.0000.0000][22]/120
 10.1.1.4 C:11000 100 0 i

Processed 2 prefixes, 2 paths
```

The following output displays the details for the incoming EVPN RT1. Note the presence of BGP extended color community 11000, and that the prefix is associated with an SR policy with color 11000 and BSID value of 80061.

```
RP/0/RSP0/CPU0:Node-B# show bgp l2vpn evpn rd 10.1.1.6:101
[1][0000.0000.0000.0000.0000][22]/120
Wed Jul 10 19:08:55.039 PST
BGP routing table entry for [1][0000.0000.0000.0000.0000][22]/120, Route Distinguisher:
10.1.1.6:101
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 322 322
Last Modified: Jul 10 18:42:10.408 for 00:26:44
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local
 10.1.1.4 C:11000 (bsid:80061) (metric 40) from 10.1.1.253 (10.1.1.4)
 Received Label 80045
 Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, rib-install
 Received Path ID 0, Local Path ID 1, version 319
 Extended community: Color:11000 RT:65000:101
 Originator: 10.1.1.4, Cluster list: 10.1.1.253
 SR policy color 11000, up, registered, bsid 80061, if-handle 0x00000560

 Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.1.1.4:101
```

Use the **show l2vpn xconnect** command to display the state associated with EVPN-VPWS EVI 101 service.

```
RP/0/RSP0/CPU0:Node-B# show l2vpn xconnect group evpn_vpws_group
Wed Jul 10 19:08:56.388 PST
```

Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,  
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

| XConnect        |                      |           | Segment 1        |           | Segment 2            |           |
|-----------------|----------------------|-----------|------------------|-----------|----------------------|-----------|
| Group           | Name                 | ST        | Description      | ST        | Description          | ST        |
| -----           |                      |           |                  |           |                      |           |
| evpn_vpws_group | <b>evpn_vpws_101</b> |           |                  |           |                      |           |
|                 |                      | <b>UP</b> | Te0/3/0/0/8.2500 | <b>UP</b> | EVPN 101,22,10.1.1.4 | <b>UP</b> |
| -----           |                      |           |                  |           |                      |           |

The following output shows the details for the service. Note that the service is associated with the on-demand SR policy with color 11000 and end-point 10.1.1.4 (node D).

```
RP/0/RSP0/CPU0:Node-B# show l2vpn xconnect group evpn_vpws_group xc-name evpn_vpws_101
Wed Jul 10 19:08:56.511 PST
```

```
Group evpn_vpws_group, XC evpn_vpws_101, state is up; Interworking none
AC: TenGigE0/3/0/0/8.2500, state is up
 Type VLAN; Num Ranges: 1
 Rewrite Tags: []
 VLAN ranges: [2500, 2500]
 MTU 1500; XC ID 0x2a0000e; interworking none
 Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0
 drops: illegal VLAN 0, illegal length 0
 EVPN: neighbor 10.1.1.4, PW ID: evi 101, ac-id 22, state is up (established)
 XC ID 0xa0000009
 Encapsulation MPLS
 Source address 10.1.1.6
 Encap type Ethernet, control word enabled
 Sequencing not set
Preferred path Active : SR TE srte_c_11000_ep_10.1.1.4, On-Demand, fallback enabled
 Tunnel : Up
 Load Balance Hashing: src-dst-mac
```

| EVPN         | Local    | Remote   |
|--------------|----------|----------|
| -----        |          |          |
| Label        | 80060    | 80045    |
| MTU          | 1500     | 1500     |
| Control word | enabled  | enabled  |
| AC ID        | 12       | 22       |
| EVPN type    | Ethernet | Ethernet |
| -----        |          |          |

```
Create time: 10/07/2019 18:32:49 (00:36:06 ago)
Last time status changed: 10/07/2019 18:42:07 (00:26:49 ago)
Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0
```

Use the **show segment-routing traffic-eng policy** command with **tabular** option to display SR policy summary information.

The following output shows the on-demand SR policy with BSID 80061 that was triggered by EVPN RT1 prefix with color 11000 advertised by node D (10.1.1.4).

```
RP/0/RSP0/CPU0:Node-B# show segment-routing traffic-eng policy color 11000 tabular
Wed Jul 10 19:08:56.146 PST
```

|       |          |       |      |         |
|-------|----------|-------|------|---------|
| Color | Endpoint | Admin | Oper | Binding |
|-------|----------|-------|------|---------|

|       |          | State | State | SID   |
|-------|----------|-------|-------|-------|
| 11000 | 10.1.1.4 | up    | up    | 80061 |

The following output shows the details for the on-demand SR policy. Note that the SR policy's active candidate path (preference 100) is computed by SR-PCE (10.1.1.207).

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. Specifically, from site 1 to site 2, LSP at Node B (srte\_c\_11000\_ep\_10.1.1.4) is link-disjoint from LSP at Node A (srte\_c\_11000\_ep\_10.1.1.2).

```
RP/0/RSP0/CPU0:Node-B# show segment-routing traffic-eng policy color 11000
Wed Jul 10 19:08:56.207 PST
```

```
SR-TE policy database

```

```
Color: 11000, End-point: 10.1.1.4
 Name: srte_c_11000_ep_10.1.1.4
 Status:
 Admin: up Operational: up for 00:26:47 (since Jul 10 18:40:05.868)
 Candidate-paths:
 Preference: 200 (BGP ODN) (shutdown)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_11000_ep_10.1.1.4_discr_200
 PLSP-ID: 19
 Dynamic (invalid)
 Preference: 100 (BGP ODN) (active)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_11000_ep_10.1.1.4_discr_100
 PLSP-ID: 18
 Dynamic (pce 10.1.1.207) (valid)
 Metric Type: IGP, Path Accumulated Metric: 40
 16001 [Prefix-SID, 10.1.1.1]
 16004 [Prefix-SID, 10.1.1.4]
 Attributes:
 Binding SID: 80061
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes
```

### Verification: Site 2 Node C

This section depicts verification steps at Node C.

Use the **show bgp l2vpn evpn** command to display BGP prefix information for EVPN-VPWS EVI 100 (rd 10.1.1.2:100). The output includes an EVPN route-type 1 route with color 10000 originated at Node A (10.1.1.5).

```
RP/0/RSP0/CPU0:Node-C# show bgp l2vpn evpn rd 10.1.1.2:100
BGP router identifier 10.1.1.2, local AS number 65000
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 21
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs
```

```

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.2:100 (default for vrf VPWS:100)
*>i [1] [0000.0000.0000.0000.0000] [11]/120
 10.1.1.5 C:10000 100 0 i
*> [1] [0000.0000.0000.0000.0000] [21]/120
 0.0.0.0 0 i

```

The following output displays the details for the incoming EVPN RT1. Note the presence of BGP extended color community 10000, and that the prefix is associated with an SR policy with color 10000 and BSID value of 80058.

```

RP/0/RSP0/CPU0:Node-C# show bgp l2vpn evpn rd 10.1.1.2:100
[1] [0000.0000.0000.0000.0000] [11]/120
BGP routing table entry for [1] [0000.0000.0000.0000.0000] [11]/120, Route Distinguisher:
10.1.1.2:100
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 20 20
Last Modified: Jul 10 18:36:20.503 for 00:45:21
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local
 10.1.1.5 C:10000 (bsid:80058) (metric 40) from 10.1.1.253 (10.1.1.5)
 Received Label 80040
 Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, rib-install
 Received Path ID 0, Local Path ID 1, version 18
 Extended community: Color:10000 RT:65000:100
 Originator: 10.1.1.5, Cluster list: 10.1.1.253
 SR policy color 10000, up, registered, bsid 80058, if-handle 0x000006a0

Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.1.1.5:100

```

Use the **show l2vpn xconnect** command to display the state associated with EVPN-VPWS EVI 100 service.

```

RP/0/RSP0/CPU0:Node-C# show l2vpn xconnect group evpn_vpws_group
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
 SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

```

| XConnect        |               |    | Segment 1      |    | Segment 2            |    |
|-----------------|---------------|----|----------------|----|----------------------|----|
| Group           | Name          | ST | Description    | ST | Description          | ST |
| -----           |               |    |                |    |                      |    |
| evpn_vpws_group |               |    |                |    |                      |    |
|                 | evpn_vpws_100 |    |                |    |                      |    |
|                 |               | UP | Gi0/0/0/3.2500 | UP | EVPN 100,11,10.1.1.5 | UP |
| -----           |               |    |                |    |                      |    |

The following output shows the details for the service. Note that the service is associated with the on-demand SR policy with color 10000 and end-point 10.1.1.5 (node A).

```

RP/0/RSP0/CPU0:Node-C# show l2vpn xconnect group evpn_vpws_group xc-name evpn_vpws_100

Group evpn_vpws_group, XC evpn_vpws_100, state is up; Interworking none
AC: GigabitEthernet0/0/0/3.2500, state is up
Type VLAN; Num Ranges: 1
Rewrite Tags: []

```



```

VLAN ranges: [2500, 2500]
MTU 1500; XC ID 0x1200008; interworking none
Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0
 drops: illegal VLAN 0, illegal length 0
EVPN: neighbor 10.1.1.5, PW ID: evi 100, ac-id 11, state is up (established)
XC ID 0xa0000003
Encapsulation MPLS
Source address 10.1.1.2
Encap type Ethernet, control word enabled
Sequencing not set
Preferred path Active : SR TE srte_c_10000_ep_10.1.1.5, On-Demand, fallback enabled
Tunnel : Up
Load Balance Hashing: src-dst-mac

```

| EVPN         | Local    | Remote   |
|--------------|----------|----------|
| Label        | 80056    | 80040    |
| MTU          | 1500     | 1500     |
| Control word | enabled  | enabled  |
| AC ID        | 21       | 11       |
| EVPN type    | Ethernet | Ethernet |

```

Create time: 10/07/2019 18:36:16 (1d19h ago)
Last time status changed: 10/07/2019 19:41:59 (1d18h ago)
Last time PW went down: 10/07/2019 19:40:54 (1d18h ago)
Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0

```

Use the **show segment-routing traffic-eng policy** command with **tabular** option to display SR policy summary information.

The following output shows the on-demand SR policy with BSID 80058 that was triggered by EVPN RT1 prefix with color 10000 advertised by node A (10.1.1.5).

```
RP/0/RSP0/CPU0:Node-C# show segment-routing traffic-eng policy color 10000 tabular
```

| Color | Endpoint | Admin State | Oper State | Binding SID |
|-------|----------|-------------|------------|-------------|
| 10000 | 10.1.1.5 | up          | up         | 80058       |

The following output shows the details for the on-demand SR policy. Note that the SR policy's active candidate path (preference 100) is computed by SR-PCE (10.1.1.207).

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. Specifically, from site 2 to site 1, LSP at Node C (srte\_c\_10000\_ep\_10.1.1.5) is link-disjoint from LSP at Node D (srte\_c\_10000\_ep\_10.1.1.6).

```
RP/0/RSP0/CPU0:Node-C# show segment-routing traffic-eng policy color 10000
```

```
SR-TE policy database
```

```

Color: 10000, End-point: 10.1.1.5
Name: srte_c_10000_ep_10.1.1.5
Status:
 Admin: up Operational: up for 00:12:35 (since Jul 10 19:49:21.890)
Candidate-paths:

```

```

Preference: 200 (BGP ODN) (shutdown)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_10000_ep_10.1.1.5_discr_200
 PLSP-ID: 7
 Dynamic (invalid)
Preference: 100 (BGP ODN) (active)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_10000_ep_10.1.1.5_discr_100
 PLSP-ID: 6
 Dynamic (pce 10.1.1.207) (valid)
 Metric Type: IGP, Path Accumulated Metric: 40
 16007 [Prefix-SID, 10.1.1.7]
 16008 [Prefix-SID, 10.1.1.8]
 80005 [Adjacency-SID, 11.5.8.8 - 11.5.8.5]
Attributes:
 Binding SID: 80058
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```

### Verification: Site 2 Node D

This section depicts verification steps at Node D.

Use the **show bgp l2vpn evpn** command to display BGP prefix information for EVPN-VPWS EVI 101 (rd 10.1.1.4:101). The output includes an EVPN route-type 1 route with color 10000 originated at Node B (10.1.1.6).

```

RP/0/RSP0/CPU0:Node-D# show bgp l2vpn evpn rd 10.1.1.4:101
BGP router identifier 10.1.1.4, local AS number 65000
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 570
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.4:101 (default for vrf VPWS:101)
*>i [1] [0000.0000.0000.0000.0000] [12]/120
 10.1.1.6 C:10000 100 0 i
*> [1] [0000.0000.0000.0000.0000] [22]/120
 0.0.0.0 0 i

Processed 2 prefixes, 2 paths

```

The following output displays the details for the incoming EVPN RT1. Note the presence of BGP extended color community 10000, and that the prefix is associated with an SR policy with color 10000 and BSID value of 80047.

```

RP/0/RSP0/CPU0:Node-D# show bgp l2vpn evpn rd 10.1.1.4:101
[1] [0000.0000.0000.0000.0000] [12]/120
BGP routing table entry for [1] [0000.0000.0000.0000.0000] [12]/120, Route Distinguisher:
10.1.1.4:101
Versions:

```

```

Process bRIB/RIB SendTblVer
Speaker 569 569
Last Modified: Jul 10 18:42:12.455 for 00:45:38
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
Local
 10.1.1.6 C:10000 (bsid:80047) (metric 40) from 10.1.1.253 (10.1.1.6)
 Received Label 80060
 Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, rib-install
 Received Path ID 0, Local Path ID 1, version 568
 Extended community: Color:10000 RT:65000:101
 Originator: 10.1.1.6, Cluster list: 10.1.1.253
 SR policy color 10000, up, registered, bsid 80047, if-handle 0x00001720

Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 10.1.1.6:101

```

Use the **show l2vpn xconnect** command to display the state associated with EVPN-VPWS EVI 101 service.

```

RP/0/RSP0/CPU0:Node-D# show l2vpn xconnect group evpn_vpws_group
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
 SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

```

| XConnect        |               | Segment 1   | ST | Segment 2            | ST |
|-----------------|---------------|-------------|----|----------------------|----|
| Group           | Name          | Description |    | Description          |    |
| evpn_vpws_group | evpn_vpws_101 | UP          | UP | EVPN 101,12,10.1.1.6 | UP |

The following output shows the details for the service. Note that the service is associated with the on-demand SR policy with color 10000 and end-point 10.1.1.6 (node B).

```

RP/0/RSP0/CPU0:Node-D# show l2vpn xconnect group evpn_vpws_group xc-name evpn_vpws_101

Group evpn_vpws_group, XC evpn_vpws_101, state is up; Interworking none
AC: GigabitEthernet0/0/0/1.2500, state is up
 Type VLAN; Num Ranges: 1
 Rewrite Tags: []
 VLAN ranges: [2500, 2500]
 MTU 1500; XC ID 0x120000c; interworking none
 Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0
 drops: illegal VLAN 0, illegal length 0
 EVPN: neighbor 10.1.1.6, PW ID: evi 101, ac-id 12, state is up (established)
 XC ID 0xa000000d
 Encapsulation MPLS
 Source address 10.1.1.4
 Encap type Ethernet, control word enabled
 Sequencing not set
 Preferred path Active : SR TE srte_c_10000_ep_10.1.1.6, On-Demand, fallback enabled
 Tunnel : Up
 Load Balance Hashing: src-dst-mac

```

| EVPN  | Local | Remote |
|-------|-------|--------|
| Label | 80045 | 80060  |
| MTU   | 1500  | 1500   |

```

Control word enabled enabled
AC ID 22 12
EVPN type Ethernet Ethernet

```

```

Create time: 10/07/2019 18:42:07 (00:45:49 ago)
Last time status changed: 10/07/2019 18:42:09 (00:45:47 ago)
Statistics:
 packets: received 0, sent 0
 bytes: received 0, sent 0

```

Use the **show segment-routing traffic-eng policy** command with **tabular** option to display SR policy summary information.

The following output shows the on-demand SR policy with BSID 80047 that was triggered by EVPN RT1 prefix with color 10000 advertised by node B (10.1.1.6).

```
RP/0/RSP0/CPU0:Node-D# show segment-routing traffic-eng policy color 10000 tabular
```

| Color | Endpoint | Admin State | Oper State | Binding SID |
|-------|----------|-------------|------------|-------------|
| 10000 | 10.1.1.6 | up          | up         | 80047       |

The following output shows the details for the on-demand SR policy. Note that the SR policy's active candidate path (preference 100) is computed by SR-PCE (10.1.1.207).

Based on the goals of this use case, SR-PCE computes link-disjoint paths for the SR policies associated with a pair of ELINE services between site 1 and site 2. Specifically, from site 2 to site 1, LSP at Node D (srte\_c\_10000\_ep\_10.1.1.6) is link-disjoint from LSP at Node C (srte\_c\_10000\_ep\_10.1.1.5).

```
RP/0/RSP0/CPU0:Node-D# show segment-routing traffic-eng policy color 10000
```

```
SR-TE policy database

```

```

Color: 10000, End-point: 10.1.1.6
Name: srte_c_10000_ep_10.1.1.6
Status:
 Admin: up Operational: up for 01:23:04 (since Jul 10 18:42:07.350)
Candidate-paths:
 Preference: 200 (BGP ODN) (shutdown)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_10000_ep_10.1.1.6_discr_200
 PLSP-ID: 17
 Dynamic (invalid)
 Preference: 100 (BGP ODN) (active)
 Requested BSID: dynamic
 PCC info:
 Symbolic name: bgp_c_10000_ep_10.1.1.6_discr_100
 PLSP-ID: 16
 Dynamic (pce 10.1.1.207) (valid)
 Metric Type: IGP, Path Accumulated Metric: 40
 16001 [Prefix-SID, 10.1.1.1]
 16006 [Prefix-SID, 10.1.1.6]
Attributes:
 Binding SID: 80047
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```

## Manually Provisioned SR Policy

Manually provisioned SR policies are configured on the head-end router. These policies can use dynamic paths or explicit paths. See the [SR-TE Policy Path Types, on page 265](#) section for information on manually provisioning an SR policy using dynamic or explicit paths.

## PCE-Initiated SR Policy

An SR-TE policy can be configured on the path computation element (PCE) to reduce link congestion or to minimize the number of network touch points.

The PCE collects network information, such as traffic demand and link utilization. When the PCE determines that a link is congested, it identifies one or more flows that are causing the congestion. The PCE finds a suitable path and deploys an SR-TE policy to divert those flows, without moving the congestion to another part of the network. When there is no more link congestion, the policy is removed.

To minimize the number of network touch points, an application, such as a Network Services Orchestrator (NSO), can request the PCE to create an SR-TE policy. PCE deploys the SR-TE policy using PCC-PCE communication protocol (PCEP).

For more information, see the [PCE-Initiated SR Policies, on page 412](#) section.

## Cumulative Metric Bounds (Delay-Bound Use-Case)

*Table 32: Feature History Table*

| Feature Name                                    | Release Information | Feature Description                                                                                                                                                                                                               |
|-------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cumulative Metric Bounds (Delay-Bound use-case) | Release 7.3.1       | With this feature, SRTE calculates a shortest path that satisfies multiple metric bounds.<br><br>This feature provides flexibility for finding paths within metric bounds, for parameters such as latency, hop count, IGP and TE. |

SRTE can calculate a shortest path with cumulative metric bounds. For example, consider these metric bounds:

- IGP metric  $\leq 10$
- TE metric  $\leq 60$
- Hop count  $\leq 4$
- Latency  $\leq 55$

When an SR policy is configured on a head-end node with these metric bounds, a path is finalized towards the specified destination only if it meets each of these criteria.

You can set the maximum number of attempts for computing a shortest path that satisfies the cumulative metric bounds criteria, by using the **kshortest-paths** command in SR-TE configuration mode.

## Restrictions

- PCE-based cumulative metric bounds computations are not supported. You must use non-PCE (SR-TE topology) based configuration for path calculation, for cumulative bounds.
- If you use PCE dynamic computation configuration with cumulative bounds, the PCE computes a path and validates against cumulative bounds. If it is valid, then the policy is created with this path on PCC. If the initial path doesn't respect the bounds, then the path is not considered, and no further K-shortest path algorithm is executed to find the path.

## Configuring SRTE Shortest Path Calculation For Cumulative Metric Bounds

You can enable this feature for SR, and ODN SR policy configurations, as shown below.

### SR Policy

**SR Policy** - A policy called **fromAtoB\_XTC** is created towards destination IP address 192.168.0.2. Also, the candidate-paths preference, and other attributes are enabled.

```
Router# configure terminal
Router(config)# segment-routing traffic-eng policy fromAtoB_XTC
Router(config-sr-te-policy)# color 2 end-point ipv4 192.168.0.2
Router(config-sr-te-policy)# candidate-paths preference 100
Router(config-sr-te-policy-path-pref)# dynamic metric type te
```

**Cumulative Metric bounds** – IGP, TE, hop count, and latency metric bounds are set. SRTE calculates paths only when each criterion is satisfied.

```
Router(config-sr-te-policy-path-pref)# constraints bounds cumulative
Router(config-sr-te-pref-const-bounds-type)# type igp 10
Router(config-sr-te-pref-const-bounds-type)# type te 60
Router(config-sr-te-pref-const-bounds-type)# type hopcount 4
Router(config-sr-te-pref-const-bounds-type)# type latency 55
Router(config-sr-te-pref-const-bounds-type)# commit
```

### ODN SR Policy

**SR ODN Policy** – An SR ODN policy with color 1000 is created. Also, the candidate-paths value is on-demand.

```
Router# configure terminal
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 1000 dynamic metric type te
Router(config-sr-te)# candidate-paths on-demand
Router(config-sr-te-candidate-path-type)# exit
Router(config-sr-te-candidate-path)# exit
```

**Cumulative Metric bounds** – IGP, TE, hop count, and latency metric bounds are set for the policy. SRTE calculates paths, only when each criterion is satisfied.

```
Router(config-sr-te)# on-demand color 1000 dynamic bounds cumulative
Router(config-sr-te-odc-bounds-type)# type igp 100
Router(config-sr-te-odc-bounds-type)# type te 60
Router(config-sr-te-odc-bounds-type)# type hopcount 6
Router(config-sr-te-odc-bounds-type)# type latency 1000
Router(config-sr-te-odc-bounds-type)# commit
```

To set the maximum number of attempts for computing paths that satisfy the cumulative metric bounds criteria, use the **kshortest-paths** command.

```
Router# configure terminal
Router(config)# segment-routing traffic-eng
```

```
Router(config-sr-te)# kshortest-paths 120
Router(config-sr-te)# commit
```

## Verification

Use this command to view SR policy configuration details. Pointers:

- The **Number of K-shortest-paths** field displays 4. It means that the K-shortest path algorithm took 4 computations to find the right path. The 4 shortest paths that are computed using K-shortest path algorithm did not respect the cumulative bounds. The fifth shortest path is valid against the bounds.
- The values for the metrics of the actual path (**TE, IGP, Cumulative Latency** and **Hop count** values in the **Dynamic** section) are within the configured cumulative metric bounds.

```
Router# show segment-routing traffic-eng policy color 2

Color: 2, End-point: 192.168.0.2
Name: srte_c_2_ep_192.168.0.2
Status:
 Admin: up Operational: up for 3d02h (since Dec 15 12:13:21.993)

Candidate-paths:

 Preference: 100 (configuration) (active)

 Name: fromAtoB_XTC
 Requested BSID: dynamic
 Constraints:
 Protection Type: protected-preferred
 Affinity:
 exclude-any:
 red
 Maximum SID Depth: 10
 IGP Metric Bound: 10
 TE Metric Bound: 60
 Latency Metric Bound: 55
 Hopcount Metric Bound: 4

 Dynamic (valid)

 Metric Type: TE, Path Accumulated Metric: 52
 Number of K-shortest-paths: 4
 TE Cumulative Metric: 52
 IGP Cumulative Metric: 3
 Cumulative Latency: 52
 Hop count: 3
 16004 [Prefix-SID, 192.168.0.4]
 24003 [Adjacency-SID, 16.16.16.2 - 16.16.16.5]
 24001 [Adjacency-SID, 14.14.14.5 - 14.14.14.4]

Attributes:

 Binding SID: 24011
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
```

## SR-TE BGP Soft Next-Hop Validation For ODN Policies

Table 33: Feature History Table

| Feature Name                                        | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE BGP Soft Next-Hop Validation For ODN Policies | Release 7.3.2       | <p>This feature addresses BGP Next-Hop reachability issues through BGP Next-Hop <i>soft</i> validation, and also enhances BGP best path selection.</p> <p>New commands:</p> <ul style="list-style-type: none"> <li>• <b>nexthop validation color-extcomm disable</b></li> <li>• <b>nexthop validation color-extcomm sr-policy</b></li> <li>• <b>bgp bestpath igp-metric sr-policy</b></li> </ul> |

Before a BGP router installs a route in the routing table, it checks its own reachability to the Next-Hop (NH) IP address of the route. In an SR-TE domain, a NH address may not be redistributed within the AS, or to a neighbor AS. So, BGP cannot reach the NH, and does not install the corresponding route into the routing table. The following workarounds are available, but they are tedious and might impact scalability:

1. Enable a non-default, static route to null0 covering the routes
2. Inject the routes into BGP using BGP-Labeled Unicast configuration
3. Redistribute routes between IGP domains

This feature introduces a more optimal design and solution - When you enable an SR policy on the SR-TE headend router, configure the `nexthop validation color-extcomm sr-policy` command in BGP configuration mode. It instructs BGP that, instead of NH reachability validation of BGP routes, the validation is done for SR policy-installed color NH addresses. When the NH address of such a route is reachable, the route is added to the routing table.

Also, this configuration on the ingress/headend PE router reduces the route scale for NH reachability, and service (VPN) routes automatically get NH reachability.

**RR configuration** – For intermediate router configuration, enable the RR with the `nexthop validation color-extcomm disable` command. When enabled, and L3VPN prefixes are associated with a color ID, BGP skips NH validation on the RR.

When the RR has no reachability to the color-extcomm NH, either enable this command, or use a legacy static route.

The following sequence occurs when the headend router receives L3VPN prefixes based on a color ID such as purple, green, etc.

1. The router checks/learns the local SR policy, or requests the ODN SR policy for color ID and NH



2. BGP does validation of the SR policy routes' NH addresses and applies the corresponding NH AD/metric. For a NH with a specific BGP-based color attribute, SR-PCE provides the AD/metric  
With BGP NH reachability, traffic is transported smoothly
3. On the RR, BGP does not validate NH reachability

### BGP Best Path Selection Based On SR Policy Effective Metric

BGP uses an algorithm to select the best path for installing the route in the RIB or for making a choice of which BGP path to propagate. At a certain point in the process, if there is IGP reachability to a BGP NH address, the algorithm chooses the path with the lowest IGP metric as the best path. The SR Policy path metric is not considered even if it has a better metric. This feature addresses the issue.

To ensure that BGP prefers the SR policy path metric over the IGP metric, enable `bgp bestpath igp-metric sr-policy` in BGP configuration mode.

## Configurations

### Configuring BGP Soft Next-Hop Validation (Headend Router)

```
Headend # configure
Headend (config) # router bgp 100
Headend (config-bgp) # nexthop validation color-extcomm sr-policy
Headend (config-bgp) # commit
Headend (config-bgp) # end
```

### Configuring BGP Soft Next-Hop Validation (Route Reflector)

```
RR # configure
RR (config) # router bgp 100
RR (config-bgp) # nexthop validation color-extcomm disable
RR (config-bgp) # commit
RR (config-bgp) # end
```

### Configuring BGP Best Path Selection Based on SR Policy Metric (Headend Router)

```
Headend # configure
Headend (config) # router bgp 100
Headend (config-bgp) # bgp bestpath igp-metric sr-policy
Headend (config-bgp) # commit
Headend (config-bgp) # end
```

## Verification

Use this command to view BGP Soft Next-Hop Validation details.

```
Headend # show bgp process detail | i Nexthop
Use SR-Policy admin/metric of color-extcomm Nexthop during path comparison: enabled ExtComm
Color Nexthop validation: SR-Policy then RIB
```

Use this command to view BGP Best Path Selection Based on SR Policy Metric.

```
Headend # show bgp vrf VRF1002 ipv4 unicast 207.77.2.0

BGP routing table entry for 207.77.2.0/24, Route Distinguisher: 18522:1002 Versions:
Process bRIB/RIB SendTblVer
Speaker 5232243 5232243 Paths: (1 available, best #1)
Advertised to CE peers (in unique update groups): 10.11.2.11 101.15.2.2
Path #1: Received by speaker 0

Advertised to CE peers (in unique update groups): 10.11.2.11 101.15.2.2
```

```

16611 770
10.1.1.33 C:1129 (bsid:27163) (admin 20) (metric 25) from 10.1.1.100 (10.1.1.33)
Received Label 24007
Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
Received Path ID 1, Local Path ID 1, version 5232243
Extended community: Color:1129 RT:17933:1002 RT:18522:1002
Originator: 10.1.1.33, Cluster list: 10.1.1.100
SR policy color 1129, up, registered, bsid 27163, if-handle 0x200053dc
Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 18522:3002

```

#### Details

- **10.1.1.33 C:1129** - BGP path is selected based on the SR policy with color ID C:1129
- If no SR policy is up, or if the SR policy metric is not configured, only the RIB metric is displayed
- **admin 20** and **metric 25** are SR policy references

## SR-TE Policy Path Types

A **dynamic** path is based on an optimization objective and a set of constraints. The head-end computes a solution, resulting in a SID-list or a set of SID-lists. When the topology changes, a new path is computed. If the head-end does not have enough information about the topology, the head-end might delegate the computation to a Segment Routing Path Computation Element (SR-PCE). For information on configuring SR-PCE, see *Configure Segment Routing Path Computation Element* chapter.

An **explicit** path is a specified SID-list or set of SID-lists.

An SR-TE policy initiates a single (selected) path in RIB/FIB. This is the preferred valid candidate path. A path is selected when the path is valid and its preference is the best among all candidate paths for that policy.




---

**Note** The protocol of the source is not relevant in the path selection logic.

---

A candidate path has the following characteristics:

- It has a preference – If two policies have the same {color, endpoint} but different preferences, the policy with the highest preference is selected.
- It is associated with a single binding SID (BSID) – A BSID conflict occurs when there are different SR policies with the same BSID. In this case, the policy that is installed first gets the BSID and is selected.
- It is valid if it is usable.

## Dynamic Paths

#### Behaviors and Limitations

For a dynamic path that traverses a specific interface between nodes (segment), the algorithm may encode this segment using an Adj-SID. The SR-TE process prefers the protected Adj-SID of the link, if one is available. In addition, the SR-TE process prefers a manual protected Adj-SID over a dynamic protected Adj-SID.


You can configure the path to prefer the protected or unprotected Adj-SID, or to use only protected or unprotected Adj-SID. See [Segment Protection-Type Constraint](#), on page 269.

## Optimization Objectives

Optimization objectives allow the head-end router to compute a SID-list that expresses the shortest dynamic path according to the selected metric type:

- IGP metric — Refer to the "Implementing IS-IS" and "Implementing OSPF" chapters in the *Routing Configuration Guide for Series Routers*.
- TE metric — See the [Configure Interface TE Metrics](#), on page 266 section for information about configuring TE metrics.

This example shows a dynamic path from head-end router 1 to end-point router 3 that minimizes IGP or TE metric:



- The blue path uses the minimum IGP metric: Min-Metric (1 → 3, IGP) = SID-list <16003>; cumulative IGP metric: 20
- The green path uses the minimum TE metric: Min-Metric (1 → 3, TE) = SID-list <16005, 16004, 16003>; cumulative TE metric: 23

## Configure Interface TE Metrics

Use the **metric value** command in SR-TE interface submode to configure the TE metric for interfaces. The *value* range is from 0 to 2147483647.

```

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# interface type interface-path-id
Router(config-sr-te-if)# metric value

```

### Configuring TE Metric: Example

The following configuration example shows how to set the TE metric for various interfaces:

```

segment-routing
traffic-eng
interface TenGigE0/0/0/0
metric 100

```

```

!
interface TenGigE0/0/0/1
 metric 1000
!
interface TenGigE0/0/2/0
 metric 50
!
!
end

```

## Constraints

Constraints allow the head-end router to compute a dynamic path according to the selected metric type:

- **Affinity** — You can apply a color or name to links or interfaces by assigning affinity bit-maps to them. You can then specify an affinity (or relationship) between an SR policy path and link colors. SR-TE computes a path that includes or excludes links that have specific colors, or combinations of colors. See the [Named Interface Link Admin Groups and SR-TE Affinity Maps, on page 267](#) section for information on named interface link admin groups and SR-TE Affinity Maps.
- **Disjoint** — SR-TE computes a path that is disjoint from another path in the same disjoint-group. Disjoint paths do not share network resources. Path disjointness may be required for paths between the same pair of nodes, between different pairs of nodes, or a combination (only same head-end or only same end-point).
- **Flexible Algorithm** — Flexible Algorithm allows for user-defined algorithms where the IGP computes paths based on a user-defined combination of metric type and constraint.
- **Protection type** — For a dynamic path that traverses a specific interface between nodes (segment), or for an explicit path using IP addresses of intermediate links, the algorithm may encode this segment using an Adj-SID. You can specify the path to prefer protected or unprotected Adj-SIDs, or to use only protected or unprotected Adj-SIDs. See [Segment Protection-Type Constraint, on page 269](#) for information about configuring the protection type.

### Named Interface Link Admin Groups and SR-TE Affinity Maps

Named Interface Link Admin Groups and SR-TE Affinity Maps provide a simplified and more flexible means of configuring link attributes and path affinities to compute paths for SR-TE policies.

In the traditional TE scheme, links are configured with attribute-flags that are flooded with TE link-state parameters using Interior Gateway Protocols (IGPs), such as Open Shortest Path First (OSPF).

Named Interface Link Admin Groups and SR-TE Affinity Maps let you assign, or map, up to 256 color names for affinity and attribute-flag attributes instead of 32-bit hexadecimal numbers. After mappings are defined, the attributes can be referred to by the corresponding color name in the CLI. Furthermore, you can define constraints using *include-any*, *include-all*, and *exclude-any* arguments, where each statement can contain up to 10 colors.



#### Note

You can configure affinity constraints using attribute flags or the Flexible Name Based Policy Constraints scheme; however, when configurations for both schemes exist, only the configuration pertaining to the new scheme is applied.

*Configure Named Interface Link Admin Groups and SR-TE Affinity Maps*

Use the **affinity name NAME** command in SR-TE interface submode to assign affinity to interfaces. Configure this on routers with interfaces that have an associated admin group attribute.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# interface TenGigE0/0/1/2
Router(config-sr-if)# affinity
Router(config-sr-if-affinity)# name RED
```

Use the **affinity-map name NAME bit-position bit-position** command in SR-TE sub-mode to define affinity maps. The *bit-position* range is from 0 to 255.

Configure affinity maps on the following routers:

- Routers with interfaces that have an associated admin group attribute.
- Routers that act as SR-TE head-ends for SR policies that include affinity constraints.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# affinity-map
Router(config-sr-te-affinity-map)# name RED bit-position 23
```

**Configuring Link Admin Group: Example**

The following example shows how to assign affinity to interfaces and to define affinity maps. This configuration is applicable to any router (SR-TE head-end or transit node) with colored interfaces.

```
segment-routing
traffic-eng
interface TenGigE0/0/1/1
affinity
name CROSS
name RED
!
!
interface TenGigE0/0/1/2
affinity
name RED
!
!
interface TenGigE0/0/2/0
affinity
name BLUE
!
!
affinity-map
name RED bit-position 23
name BLUE bit-position 24
name CROSS bit-position 25
!
end
```

*Segment Protection-Type Constraint***Table 34: Feature History Table**

| Feature Name                       | Release Information | Feature Description                                                                                                                                                                                                                                                                  |
|------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Segment Protection-Type Constraint | Release 7.4.1       | <p>This feature introduces the ability to control whether protected or unprotected segments are used when encoding the SID-list of an SR policy candidate path.</p> <p>The types of segments that could be used when building a SID-list include prefix SIDs and adjacency SIDs.</p> |

This feature introduces the ability to control whether protected or unprotected segments are used when encoding the SID-list of an SR policy candidate path. The types of segments that could be used when building a SID-list include prefix SIDs and adjacency SIDs.

A prefix SID is a global segment representing a prefix that identifies a specific node. A prefix SID is programmed with a backup path computed by the IGP using TI-LFA.

An adjacency SID is a local segment representing an IGP adjacency. An adjacency SID can be programmed with or without protection. Protected adjacency SIDs are programmed with a link-protectant backup path computed by the IGP (TI-LFA) and are used if the link associated with the IGP adjacency fails.

Prefix SIDs and adjacency SIDs can be leveraged as segments in a SID-list in order to forward a packet along a path traversing specific nodes and/or over specific interfaces between nodes. The type of segment used when encoding the SID-list will determine whether failures along the path would be protected by TI-LFA. Depending on the offering, an operator may want to offer either unprotected or protected services over traffic engineered paths.

The following behaviors are available with the segment protection-type constraint:

- **protected-only** — The SID-list must be encoded using protected segments.
- **protected-preferred** — The SID-list should be encoded using protected segments if available; otherwise, the SID-list may be encoded using unprotected Adj-SIDs. This is the default behavior when no segment protection-type constraint is specified.
- **unprotected-only** — The SID-list must be encoded using unprotected Adj-SID.
- **unprotected-preferred** — The SID-list should be encoded using unprotected Adj-SID if available, otherwise SID-list may be encoded using protected segments.

**Usage Guidelines and Limitations**

Observe the following guidelines and limitations for the platform:

- This constraint applies to candidate-paths of manual SR policies with either dynamically computed paths or explicit paths.
- This constraint applies to On-Demand SR policy candidate-paths.

- PCEP has been augmented (vendor-specific object) to allow a PCC to indicate the segment protection-type constraint to the PCE.
- When the segment protection type constraint is protected-only or unprotected-only, the path computation must adhere to the constraint. If the constraint is not satisfied, the SR policy will not come up on such candidate path.
- When the segment protection-type constraint is unprotected-only, the entire SID-list must be encoded with unprotected Adj-SIDs.
- When the segment protection-type constraint is protected-only, the entire SID-list must be encoded with protected Adj-SIDs or Prefix SIDs.

### Configuring Segment Protection-Type Constraint

Use the **constraints segments protection** {**protected-only** | **protected-preferred** | **unprotected-only** | **unprotected-preferred**} command to configure the segment protection-type behavior.

The following example shows how to configure the policy with a SID-list that must be encoded using protected segments:

```
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 10 end-point ipv4 10.1.1.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# constraints
Router(config-sr-te-path-pref-const)# segments
Router(config-sr-te-path-pref-const-seg)# protection protected-only
```

## Configure SR Policy with Dynamic Path

To configure a SR-TE policy with a dynamic path, optimization objectives, and affinity constraints, complete the following configurations:

1. Define the optimization objectives. See the [Optimization Objectives, on page 266](#) section.
2. Define the constraints. See the [Constraints, on page 267](#) section.
3. Create the policy.

### Behaviors and Limitations

You can configure the path to prefer protected or unprotected segments, or to use only protected or unprotected segments.

### Examples

The following example shows a configuration of an SR policy at an SR-TE head-end router. The policy has a dynamic path with optimization objectives and affinity constraints computed by the head-end router.

```
segment-routing
traffic-eng
policy foo
color 100 end-point ipv4 10.1.1.2
candidate-paths
preference 100
dynamic
metric
```

```

 type te
 !
 !
 constraints
 affinity
 exclude-any
 name RED
 !
 !
 !
 !
 !
 !
 !
 !
 !

```

The following example shows a configuration of an SR policy at an SR-TE head-end router. The policy has a dynamic path with optimization objectives and affinity constraints computed by the SR-PCE.

```

segment-routing
traffic-eng
policy baa
color 101 end-point ipv4 10.1.1.2
candidate-paths
preference 100
dynamic
 pcep
 !
 metric
 type te
 !
 !
 constraints
 affinity
 exclude-any
 name BLUE
 !
 !
 !
 !
 !
 !
 !
 !

```

The following example shows a configuration of an SR policy at an SR-TE head-end router. The policy has a dynamic path with optimization objective and segment protection-type constraint computed by the head-end router.

```

segment-routing
traffic-eng
policy baa
color 101 end-point ipv4 10.1.1.2
candidate-paths
preference 100
dynamic
 metric
 type te
 !
 !
 constraints
 segments
 protection protected-only
 !
 !
 !
 !
 !

```



The following example shows a configuration of an SR policy at an SR-TE head-end router. The policy has a dynamic path with optimization objective and segment protection-type constraint computed by the SR-PCE.

```
segment-routing
traffic-eng
 policy baa
 color 101 end-point ipv4 10.1.1.2
 candidate-paths
 preference 100
 dynamic
 pcep
 !
 metric
 type te
 !
 !
 constraints
 segments
 protection protected-only
 !
 !
!
!
```

## Explicit Paths

## SR-TE Policy with Explicit Path

**Table 35: Feature History Table**

| Feature Name                                                    | Release Information | Feature Description                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE Explicit Segment Lists with Mix of IPv4 and IPv6 Segments | Release 7.9.1       | <p>This feature allows you to configure an explicit segment list with IPv4 addresses and include an IPv6 address as a non-first SID.</p> <p>You can thus deploy a centralized BGP EPE solution for 6PE in an SR-MPLS network where the last segment is associated with an EPE-enabled BGPv6 neighbor.</p> |

An explicit segment list is defined as a sequence of one or more segments. A segment can be configured as an IP address or an MPLS label representing a node or a link.

An explicit segment list can be configured with the following:

- IP-defined segments
- MPLS label-defined segments

- ## Usage Guidelines and Limitations

- An IP-defined segment can be associated with an IPv4 address (for example, a link or a Loopback address).
- When a segment of the segment list is defined as an MPLS label, subsequent segments can only be configured as MPLS labels.
- You can configure the path to prefer the protected or unprotected Adj-SID, or to use only protected or unprotected Adj-SID. See [Segment Protection-Type Constraint, on page 269](#).

To configure an SR-TE policy with an explicit path, complete the following configurations:

1. Create the segment list.
2. Create the SR-TE policy.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list name SIDLIST1
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.2
Router(config-sr-te-sl)# index 20 mpls adjacency 10.1.1.3
Router(config-sr-te-sl)# index 30 mpls adjacency 10.1.1.4
Router(config-sr-te-sl)# exit
```

```
Router(config-sr-te)# segment-list name SIDLIST2
Router(config-sr-te-sl)# index 10 mpls label 16002
Router(config-sr-te-sl)# index 20 mpls label 16003
Router(config-sr-te-sl)# index 30 mpls label 16004
Router(config-sr-te-sl)# exit
```

```
Router(config-sr-te)# segment-list name SIDLIST3
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.2
Router(config-sr-te-sl)# index 20 mpls label 16003
Router(config-sr-te-sl)# index 30 mpls label 16004
Router(config-sr-te-sl)# exit
```

```
Router(config-sr-te)# policy POLICY2
Router(config-sr-te-policy)# color 20 end-point ipv4 10.1.1.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 200
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST2
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path)# preference 100
```

```
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST1
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# exit
```

## Running Configuration

```
Router# show running-configuration

segment-routing
traffic-eng
segment-list SIDLIST1
index 10 mpls adjacency 10.1.1.2
index 20 mpls adjacency 10.1.1.3
index 30 mpls adjacency 10.1.1.4
!
segment-list SIDLIST2
index 10 mpls label 16002
index 20 mpls label 16003
index 30 mpls label 16004
!
segment-list SIDLIST3
index 10 mpls adjacency 10.1.1.2
index 20 mpls label 16003
index 30 mpls label 16004
!
policy POLICY2
color 20 end-point ipv4 10.1.1.4
candidate-paths
preference 200
explicit segment-list SIDLIST2
!
!
preference 100
explicit segment-list SIDLIST1
!
!
!!
!
```

## Verification

Verify the SR-TE policy configuration using:

```
Router# show segment-routing traffic-eng policy name srte_c_20_ep_10.1.1.4

SR-TE policy database

Color: 20, End-point: 10.1.1.4
Name: srte_c_20_ep_10.1.1.4
Status:
 Admin: up Operational: up for 00:00:15 (since Jul 14 00:53:10.615)
Candidate-paths:
 Preference: 200 (configuration) (active)
 Name: POLICY2
 Requested BSID: dynamic
 Protection Type: protected-preferred
 Maximum SID Depth: 8
 Explicit: segment-list SIDLIST2 (active)
 Weight: 1, Metric Type: TE
 16002
```

```

16003
16004

Preference: 100 (configuration) (inactive)
Name: POLICY2
Requested BSID: dynamic
Protection Type: protected-preferred
Maximum SID Depth: 8
Explicit: segment-list SIDLIST1 (inactive)
Weight: 1, Metric Type: TE
[Adjacency-SID, 10.1.1.2 - <None>]
[Adjacency-SID, 10.1.1.3 - <None>]
[Adjacency-SID, 10.1.1.4 - <None>]
Attributes:
Binding SID: 51301
Forward Class: Not Configured
Steering labeled-services disabled: no
Steering BGP disabled: no
IPv6 caps enable: yes
Invalidation drop enabled: no


```

## Explicit Path with a BGP Prefix SID as First Segment

**Table 36: Feature History Table**

| Feature Name                                               | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE Explicit Path with a BGP Prefix SID as First Segment | Release 7.11.1      | <p>This feature allows you to configure an SR-TE policy with an explicit path that uses a remote BGP prefix SID as its first segment. This path is achieved by leveraging the recursive resolution of the first SID, which is a BGP-Label Unicast (BGP-LU) SID. BGP-LU labels are used as the first SID in the SR policy to determine the egress paths for the traffic and program the SR-TE forwarding chain accordingly.</p> <p>This allows users to enable Segment Routing to leverage their existing BGP infrastructure and integrate it with the required Segment Routing functionalities.</p> |

The figure below shows a network setup where a PE router (PE1) is using BGP-SR as the routing protocol and is connected to a classic BGP-LU domain. The classic BGP-LU domain is also connected to other domains running BGP-SR




In BGP-SR, a BGP Prefix-SID is advertised along with a prefix in BGP Labeled Unicast (BGP-LU). This Prefix-SID attribute contains information about the label value and index for the route. This allows for interworking between the classic BGP-LU and BGP-SR domains.

The figure below illustrates a best-effort BGP-LU Label Switched Path (LSP) from PE1 to PE6, passing through transit Label Switching Router (LSR) nodes 2, 3, and 5.

With this setup, the operator can create an alternate path using an SR-TE policy at the ingress BGP-SR PE (PE1). The explicit path for this alternate path will follow a different transit node, using the BGP prefix SID for that transit node as the first segment.

For example, in the figure below, the SID-list shows that the explicit path uses the BGP prefix SID of transit LSR node 4 (16004) followed by the BGP prefix SID of PE6 (16006). The PE1 router resolves the first segment of this explicit path to the outgoing interface towards Node 2, using the outgoing label advertised by Node 2 for the BGP-LU prefix 10.1.1.4/32.



### Limitations for explicit Path with a BGP prefix SID as first segment

- SR-TE explicit path with a BGP prefix SID as the first segment does not support BGP peering on loopback interfaces. It is recommended to use interface peering for BGP sessions.

### Example

The following output shows the BGP-LU routes learned at the head-end (PE1 in the illustration above). Consider the IP prefix 1.1.1.4/32. A BGP-LU local label is assigned from the SRGB with a prefix SID index of 4, resulting in the value 16004. Additionally, the classic BGP-LU upstream neighbor (P2) advertises an outgoing label of 78000 for this IP prefix.

```
RP/0/RP0/CPU0:R1# show bgp ipv4 unicast labels
```

```
BGP router identifier 1.1.1.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
```

```

Table ID: 0xe0000000 RD version: 13
BGP main routing table version 13
BGP NSR Initial initsync version 6 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

| Network                 | Next Hop        | Rcvd Label   | Local Label  |
|-------------------------|-----------------|--------------|--------------|
| *> 1.1.1.1/32           | 0.0.0.0         | nolabel      | 3            |
| *> 1.1.1.2/32           | 10.1.2.2        | 3            | 24007        |
| *> 1.1.1.3/32           | 10.1.2.2        | 78005        | 24002        |
| <b>*&gt; 1.1.1.4/32</b> | <b>10.1.2.2</b> | <b>78000</b> | <b>16004</b> |
| *> 1.1.1.5/32           | 10.1.2.2        | 78002        | 16005        |
| *> 1.1.1.6/32           | 10.1.2.2        | 78001        | 16006        |

```

...

RP/0/RP0/CPU0:R1# show bgp ipv4 labeled-unicast 1.1.1.4/32

BGP routing table entry for 1.1.1.4/32
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 74 74
 Local Label: 16004
Last Modified: Sep 29 19:52:18.155 for 00:07:22
Paths: (1 available, best #1)
 Advertised to update-groups (with more than one peer):
 0.2
 Path #1: Received by speaker 0
 Advertised to update-groups (with more than one peer):
 0.2
 3
 10.1.2.2 from 10.1.2.2 (1.1.1.2)
 Received Label 78000
 Origin IGP, metric 0, localpref 100, valid, external, best, group-best
 Received Path ID 0, Local Path ID 1, version 74
 Origin-AS validity: not-found
 Label Index: 4

```

The following output shows the corresponding label forwarding entry for BGP prefix SID 16004 with corresponding outgoing label and outgoing interface.

```

RP/0/RP0/CPU0:R1# show mpls forwarding labels 16002

```

| Local Label | Outgoing Label | Prefix or ID   | Outgoing Interface | Next Hop | Bytes Switched |
|-------------|----------------|----------------|--------------------|----------|----------------|
| 16004       | 78000          | SR Pfx (idx 2) | Te0/0/0/0          | 10.1.2.2 | 44             |

The configuration below depicts an SR policy (color 100 and end-point 1.1.1.6) with an explicit path to PE6 using the BGP prefix SID of transit LSR node 4 (16004) as its first segment followed by the BGP prefix SID of PE6 (16006).

```

RP/0/RP0/CPU0:R1# configure
RP/0/RP0/CPU0:R1(config)# segment-routing
RP/0/RP0/CPU0:R1(config-sr)# traffic-eng
RP/0/RP0/CPU0:R1(config-sr-te)# segment-list SL-R4-R6

```

```

RP/0/RP0/CPU0:R1(config-sr-te-sl)# index 10 mpls label 16004
RP/0/RP0/CPU0:R1(config-sr-te-sl)# index 20 mpls label 16006
RP/0/RP0/CPU0:R1(config-sr-te-sl)# exit

RP/0/RP0/CPU0:R1(config-sr-te)# policy POL-to-R6
RP/0/RP0/CPU0:R1(config-sr-te-policy)# color 100 end-point ipv4 1.1.1.6
RP/0/RP0/CPU0:R1(config-sr-te-policy)# candidate-paths
RP/0/RP0/CPU0:R1(config-sr-te-policy-path)# preference 100
RP/0/RP0/CPU0:R1(config-sr-te-policy-path-pref)# explicit segment-list SL-R4-R6

RP/0/RP0/CPU0:R1# show running-configuration

segment-routing
traffic-eng
 segment-list SL-R4-R6
 index 10 mpls label 16004
 index 20 mpls label 16006
 !
 policy POL-to-R6
 color 100 end-point ipv4 1.1.1.6
 candidate-paths
 preference 100
 explicit segment-list SL-R4-R6
 !
 !
!
!
!
!
!
!

```

The following output depicts the forwarding information for the SR policy including the outgoing interface and outgoing label stack.

Observe how the first segment configured (MPLS label 16004) is replaced in the forwarding with the label advertised by the BGP-LU neighbor of the SRTE head-end (MPLS label value 78000).

```
RP/0/RP0/CPU0:R1# show segment-routing traffic-eng forwarding policy color 100
```

```

SR-TE Policy Forwarding database

Color: 100, End-point: 1.1.1.6
Name: srte_c_100_ep_1.1.1.6
Binding SID: 24006
Active LSP:
 Candidate path:
 Preference: 100 (configuration)
 Name: POL-to-R6
 Local label: 24005
 Segment lists:
 SL[0]:
 Name: SL-R4-R6
 Switched Packets/Bytes: 100/10000
 [MPLS -> MPLS]: 100/10000
 Paths:
 Path[0]:
 Outgoing Label: 78000
 Outgoing Interfaces: TenGigE0/0/0/0
 Next Hop: 10.1.2.2
 Switched Packets/Bytes: 100/10000
 [MPLS -> MPLS]: 100/10000
 FRR Pure Backup: No
 ECMP/LFA Backup: No
 Internal Recursive Label: Unlabelled (recursive)

```

```
Label Stack (Top -> Bottom): { 78000, 16006 }
```

```
Policy Packets/Bytes Switched: 100/9600
```

With the SR-TE policy installed, the head-end can apply SR-TE automated steering principles when programming BGP service overlay routes.

In the following output, BGP prefix 10.0.0.0/8 with next-hop 1.1.1.6 and color 100 is automatically steered over the SR policy configured above (color 100 and end-point 1.1.1.6).

```
RP/0/RP0/CPU0:R1# show bgp
```

```
BGP router identifier 1.1.1.1, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 13
BGP main routing table version 13
BGP NSR Initial initsync version 6 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network Next Hop Metric LocPrf Weight Path
...
*> 10.0.0.0/8 1.1.1.6 C:100 0 0 3 i
```

## Configuring Explicit Path with Affinity Constraint Validation

To fully configure SR-TE flexible name-based policy constraints, you must complete these high-level tasks in order:

1. Assign Color Names to Numeric Values
2. Associate Affinity-Names with SR-TE Links
3. Associate Affinity Constraints for SR-TE Policies

```
/* Enter the global configuration mode and assign color names to numeric values
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# affinity-map
Router(config-sr-te-affinity-map)# blue bit-position 0
Router(config-sr-te-affinity-map)# green bit-position 1
Router(config-sr-te-affinity-map)# red bit-position 2
Router(config-sr-te-affinity-map)# exit

/* Associate affinity-names with SR-TE links
Router(config-sr-te)# interface Gi0/0/0/0
Router(config-sr-te-if)# affinity
Router(config-sr-te-if-affinity)# blue
Router(config-sr-te-if-affinity)# exit
Router(config-sr-te-if)# exit
```



```

Router(config-sr-te)# interface Gi0/0/0/1
Router(config-sr-te-if)# affinity
Router(config-sr-te-if-affinity)# blue
Router(config-sr-te-if-affinity)# green
Router(config-sr-te-if-affinity)# exit
Router(config-sr-te-if)# exit
Router(config-sr-te)#

/* Associate affinity constraints for SR-TE policies
Router(config-sr-te)# segment-list name SIDLIST1
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.2
Router(config-sr-te-sl)# index 20 mpls adjacency 2.2.2.23
Router(config-sr-te-sl)# index 30 mpls adjacency 10.1.1.4
Router(config-sr-te-sl)# exit
Router(config-sr-te)# segment-list name SIDLIST2
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.2
Router(config-sr-te-sl)# index 30 mpls adjacency 10.1.1.4
Router(config-sr-te-sl)# exit
Router(config-sr-te)# segment-list name SIDLIST3
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.5
Router(config-sr-te-sl)# index 30 mpls adjacency 10.1.1.4
Router(config-sr-te-sl)# exit

Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 20 end-point ipv4 10.1.1.4
Router(config-sr-te-policy)# binding-sid mpls 1000
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 200
Router(config-sr-te-policy-path-pref)# constraints affinity exclude-any red
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST1
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST2
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST3

```

## Running Configuration

```

Router# show running-configuration
segment-routing
traffic-eng

interface GigabitEthernet0/0/0/0
 affinity
 blue
!
!
interface GigabitEthernet0/0/0/1
 affinity
 blue
 green
!
!

segment-list name SIDLIST1
 index 10 mpls adjacency 10.1.1.2
 index 20 mpls adjacency 2.2.2.23

```

```

 index 30 mpls adjacency 10.1.1.4
 !
 segment-list name SIDLIST2
 index 10 mpls adjacency 10.1.1.2
 index 30 mpls adjacency 10.1.1.4
 !
 segment-list name SIDLIST3
 index 10 mpls adjacency 10.1.1.5
 index 30 mpls adjacency 10.1.1.4
 !
 policy POLICY1
 binding-sid mpls 1000
 color 20 end-point ipv4 10.1.1.4
 candidate-paths
 preference 100
 explicit segment-list SIDLIST3
 !
 !
 preference 200
 explicit segment-list SIDLIST1
 !
 explicit segment-list SIDLIST2
 !
 constraints
 affinity
 exclude-any
 red
 !
 !
 !
 !
 !
 !
 !
 affinity-map
 blue bit-position 0
 green bit-position 1
 red bit-position 2
 !
 !
 !

```

## Configure Explicit Path with Segment Protection-Type Constraint

For an SR policy with an explicit path that includes IP addresses of links, the SR-TE process encodes these segments using the corresponding adjacency SID (Adj-SID) for each link. The type of Adj-SID used (protected or unprotected) is determined by the segment protection-type constraint configured under the SR policy. See the [Segment Protection-Type Constraint](#), on page 269.

### Configure Local SR-TE Policy Using Explicit Paths

Create a segment list with IP addresses:

```

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list name SIDLIST1
Router(config-sr-te-sl)# index 10 mpls adjacency 10.1.1.2
Router(config-sr-te-sl)# index 20 mpls adjacency 10.1.1.3
Router(config-sr-te-sl)# index 30 mpls adjacency 10.1.1.4
Router(config-sr-te-sl)# exit

```

Create the SR-TE policy with segment protection-type constraint:

```
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 20 end-point ipv4 10.1.1.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST1
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# constraints
Router(config-sr-te-path-pref-const)# segments
Router(config-sr-te-path-pref-const-seg)# protection protected-only
Router(config-sr-te-path-pref-const-seg)# commit
```

### Running Configuration

```
Router# show running-configuration
segment-routing
traffic-eng
segment-list SIDLIST1
index 10 mpls adjacency 10.1.1.2
index 20 mpls adjacency 10.1.1.3
index 30 mpls adjacency 10.1.1.4
!

policy POLICY1
color 10 end-point ipv4 10.1.1.4
candidate-paths
preference 100
explicit segment-list SIDLIST1
!
constraints
segments
protection protected-only
!
!
```

## Protocols

### Path Computation Element Protocol

The path computation element protocol (PCEP) describes a set of procedures by which a path computation client (PCC) can report and delegate control of head-end label switched paths (LSPs) sourced from the PCC to a PCE peer. The PCE can request the PCC to update and modify parameters of LSPs it controls. The stateful model also enables a PCC to allow the PCE to initiate computations allowing the PCE to perform network-wide orchestration.

### Configure the Head-End Router as PCEP PCC

Configure the head-end router as PCEP Path Computation Client (PCC) to establish a connection to the PCE. The PCC and PCE addresses must be routable so that TCP connection (to exchange PCEP messages) can be established between PCC and PCE.

### Configure the PCC to Establish a Connection to the PCE

Use the **segment-routing traffic-eng pcc** command to configure the PCC source address, the SR-PCE address, and SR-PCE options.

A PCE can be given an optional precedence. If a PCC is connected to multiple PCEs, the PCC selects a PCE with the lowest precedence value. If there is a tie, a PCE with the highest IP address is chosen for computing path. The precedence *value* range is from 0 to 255.

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc
Router(config-sr-te-pcc)# source-address ipv4 local-source-address
Router(config-sr-te-pcc)# pce address ipv4 PCE-address[precedence value]
Router(config-sr-te-pcc)# pce address ipv4 PCE-address[keychain WORD]
```

### Configure PCEP Authentication

TCP Message Digest 5 (MD5) authentication has been used for authenticating PCEP (TCP) sessions by using a clear text or encrypted password. This feature introduces support for TCP Authentication Option (TCP-AO), which replaces the TCP MD5 option.

TCP-AO uses Message Authentication Codes (MACs), which provides the following:

- Protection against replays for long-lived TCP connections
- More details on the security association with TCP connections than TCP MD5
- A larger set of MACs with minimal system and operational changes

TCP-AO is compatible with Master Key Tuple (MKT) configuration. TCP-AO also protects connections when using the same MKT across repeated instances of a connection. TCP-AO protects the connections by using traffic key that are derived from the MKT, and then coordinates changes between the endpoints.




---

**Note** TCP-AO and TCP MD5 are never permitted to be used simultaneously. TCP-AO supports IPv6, and is fully compatible with the proposed requirements for the replacement of TCP MD5.

---

### TCP Message Digest 5 (MD5) Authentication

Use the **password {clear | encrypted} LINE** command to enable TCP MD5 authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured password will be rejected. Specify if the password is encrypted or clear text

```
Router(config-sr-te-pcc)# pce address ipv4 PCE-address[password {clear | encrypted} LINE]
```

### TCP Authentication Option (TCP-AO)

Use the **tcp-ao key-chain [include-tcp-options]** command to enable TCP Authentication Option (TCP-AO) authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured key chain will be rejected. Use the **include-tcp-options** keyword to include other TCP options in the header for MAC calculation.

```
Router(config-sr-te-pcc)# pce address ipv4 PCE-address tcp-ao key-chain [include-tcp-options]
```

### Configure PCEP-Related Timers

Use the **timers keepalive** command to specify how often keepalive messages are sent from PCC to its peers. The range is from 0 to 255 seconds; the default value is 30.

```
Router(config-sr-te-pcc) # timers keepalive seconds
```

Use the **timers deadtimer** command to specify how long the remote peers wait before bringing down the PCEP session if no PCEP messages are received from this PCC. The range is from 1 to 255 seconds; the default value is 120.

```
Router(config-sr-te-pcc) # timers deadtimer seconds
```

Use the **timers delegation-timeout** command to specify how long a delegated SR policy can remain up without an active connection to a PCE. The range is from 0 to 3600 seconds; the default value is 60.

```
Router(config-sr-te-pcc) # timers delegation-timeout seconds
```

### PCE-Initiated SR Policy Timers

Use the **timers initiated orphans** command to specify the amount of time that a PCE-initiated SR policy will remain delegated to a PCE peer that is no longer reachable by the PCC. The range is from 10 to 180 seconds; the default value is 180.

```
Router(config-sr-te-pcc) # timers initiated orphans seconds
```

Use the **timers initiated state** command to specify the amount of time that a PCE-initiated SR policy will remain programmed while not being delegated to any PCE. The range is from 15 to 14440 seconds (24 hours); the default value is 600.

```
Router(config-sr-te-pcc) # timers initiated state seconds
```

To better understand how the PCE-initiated SR policy timers operate, consider the following example:

- PCE A instantiates SR policy P at head-end N.
- Head-end N delegates SR policy P to PCE A and programs it in forwarding.
- If head-end N detects that PCE A is no longer reachable, then head-end N starts the PCE-initiated **orphan** and **state** timers for SR policy P.
- If PCE A reconnects before the **orphan** timer expires, then SR policy P is automatically delegated back to its original PCE (PCE A).
- After the **orphan** timer expires, SR policy P will be eligible for delegation to any other surviving PCE(s).
- If SR policy P is not delegated to another PCE before the **state** timer expires, then head-end N will remove SR policy P from its forwarding.

### Enable SR-TE SYSLOG Alarms

Use the **logging policy status** command to enable SR-TE related SYSLOG alarms.

```
Router(config-sr-te) # logging policy status
```

### Enable PCEP Reports to SR-PCE

Use the **report-all** command to enable the PCC to report all SR policies in its database to the PCE.

```
Router(config-sr-te-pcc) # report-all
```

### Customize MSD Value at PCC

Use the **maximum-sid-depth** *value* command to customize the Maximum SID Depth (MSD) signaled by PCC during PCEP session establishment.

The default MSD *value* is equal to the maximum MSD supported by the platform (12).

```
Router(config-sr-te) # maximum-sid-depth value
```



**Note** The platform's SR-TE label imposition capabilities are as follows:

- Up to 12 transport labels when no service labels are imposed
- Up to 9 transport labels when service labels are imposed

For cases with path computation at PCE, a PCC can signal its MSD to the PCE in the following ways:

- During PCEP session establishment – The signaled MSD is treated as a node-wide property.
  - MSD is configured under **segment-routing traffic-eng maximum-sid-depth** *value* command
- During PCEP LSP path request – The signaled MSD is treated as an LSP property.
  - On-demand (ODN) SR Policy: MSD is configured using the **segment-routing traffic-eng on-demand color** *color* **maximum-sid-depth** *value* command
  - Local SR Policy: MSD is configured using the **segment-routing traffic-eng policy** *WORD* **candidate-paths preference** *preference* **dynamic metric sid-limit** *value* command.



**Note** If the configured MSD values are different, the per-LSP MSD takes precedence over the per-node MSD.

After path computation, the resulting label stack size is verified against the MSD requirement.

- If the label stack size is larger than the MSD and path computation is performed by PCE, then the PCE returns a "no path" response to the PCC.
- If the label stack size is larger than the MSD and path computation is performed by PCC, then the PCC will not install the path.



**Note** A sub-optimal path (if one exists) that satisfies the MSD constraint could be computed in the following cases:

- For a dynamic path with TE metric, when the PCE is configured with the **pce segment-routing te-latency** command or the PCC is configured with the **segment-routing traffic-eng te-latency** command.
- For a dynamic path with LATENCY metric
- For a dynamic path with affinity constraints

For example, if the PCC MSD is 4 and the optimal path (with an accumulated metric of 100) requires 5 labels, but a sub-optimal path exists (with accumulated metric of 110) requiring 4 labels, then the sub-optimal path is installed.

### Customize the SR-TE Path Calculation

Use the **te-latency** command to enable ECMP-aware path computation for TE metric.

```
Router(config-sr-te) # te-latency
```



**Note** ECMP-aware path computation is enabled by default for IGP and LATENCY metrics.

### Configure PCEP Redundancy Type

Use the **redundancy pcc-centric** command to enable PCC-centric high-availability model. The PCC-centric model changes the default PCC delegation behavior to the following:

- After LSP creation, LSP is automatically delegated to the PCE that computed it.
- If this PCE is disconnected, then the LSP is redelegated to another PCE.
- If the original PCE is reconnected, then the delegation fallback timer is started. When the timer expires, the LSP is redelegated back to the original PCE, even if it has worse preference than the current PCE.

```
Router(config-sr-te-pcc) # redundancy pcc-centric
```

### Configuring Head-End Router as PCEP PCC and Customizing SR-TE Related Options: Example

The following example shows how to configure an SR-TE head-end router with the following functionality:

- Enable the SR-TE head-end router as a PCEP client (PCC) with 3 PCEP servers (PCE) with different precedence values. The PCE with IP address 10.1.1.57 is selected as BEST.
- Enable SR-TE related syslogs.
- Set the Maximum SID Depth (MSD) signaled during PCEP session establishment to 5.
- Enable PCEP reporting for all policies in the node.

```
segment-routing
traffic-eng
pcc
```

```
source-address ipv4 10.1.1.2
pce address ipv4 10.1.1.57
precedence 150
password clear <password>
!
pce address ipv4 10.1.1.58
precedence 200
password clear <password>
!
pce address ipv4 10.1.1.59
precedence 250
password clear <password>
!
!
logging
policy status
!
maximum-sid-depth 5
pcc
report-all
!
!
end
```

## Verification

RP/0/RSP0/CPU0:Router# **show segment-routing traffic-eng pcc ipv4 peer**

PCC's peer database:

-----

**Peer address: 10.1.1.57, Precedence: 150, (best PCE)**

**State up**

Capabilities: Stateful, Update, Segment-Routing, Instantiation

**Peer address: 10.1.1.58, Precedence: 200**

**State up**

Capabilities: Stateful, Update, Segment-Routing, Instantiation

**Peer address: 10.1.1.59, Precedence: 250**

**State up**

Capabilities: Stateful, Update, Segment-Routing, Instantiation



## Configure SR-TE PCE Groups

**Table 37: Feature History Table**


| Feature Name     | Release       | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE PCE Groups | Release 7.3.2 | <p>This feature allows an SR policy to be delegated to a set of PCE servers configured under a PCE group. Multiple PCE groups can be configured to allow SR policies on the same head-end to be delegated to different sets of PCEs.</p> <p>With this functionality, an operator can designate sets of PCEs for various purposes, such as PCE-per-service-type or PCE-per-wholesale-customers.</p> |

This feature allows an SR policy to be delegated or reported to a set of PCE servers configured under a PCE group. Multiple PCE groups can be configured to allow different SR policies on the same head-end to be delegated or reported to different sets of PCEs.

With this functionality, an operator can designate sets of PCEs for various purposes, such as PCE-per-service-type or PCE-per-wholesale-customer.

In the figure below, Router A has a PCEP session with 5 PCEs. The PCEs are configured over 3 PCE groups. PCE1 is in the “default” group. PCE2 and PCE3 are in the RED group. PCE4 and PCE5 are in the BLUE group.

**Figure 13: Example: PCE Groups**



In case of PCE failure, each candidate path is re-delegated to the next-best PCE within the same PCE group. For example, if the best PCE in the RED group (PCE2) fails, then all candidate paths in the RED group fallback to the secondary PCE in the RED group (PCE3). If all the PCEs in the RED group fail, then all candidate paths in the RED group become undelegated; they are not delegated to the PCEs in the BLUE group. If there

are no more available PCEs in the given PCE group, then the outcome is the same as when there are no available PCEs.

### Configure PCE Groups

Use the **segment-routing traffic-eng pcc pce address {ipv4 *ipv4\_addr* | ipv6 *ipv6\_addr*} pce-group WORD** command to configure the PCE groups.

The following example shows how to configure the PCE groups

- PCE1 in the “default” group
- PCE2 and PCE3 in the “red” group
- PCE4 and PCE5 in the “blue” group

```
Router(config)# segment-routing traffic-eng pcc
Router(config-sr-te-pcc)# pce address ipv4 10.1.1.1
Router(config-pcc-pce)# precedence 10
Router(config-pcc-pce)# exit

Router(config-sr-te-pcc)# pce address ipv4 2.2.2.2
Router(config-pcc-pce)# precedence 20
Router(config-pcc-pce)# pce-group red
Router(config-pcc-pce)# exit

Router(config-sr-te-pcc)# pce address ipv4 3.3.3.3
Router(config-pcc-pce)# precedence 30
Router(config-pcc-pce)# pce-group red
Router(config-pcc-pce)# exit

Router(config-sr-te-pcc)# pce address ipv4 4.4.4.4
Router(config-pcc-pce)# precedence 40
Router(config-pcc-pce)# pce-group blue
Router(config-pcc-pce)# exit

Router(config-sr-te-pcc)# pce address ipv4 5.5.5.5
Router(config-pcc-pce)# precedence 50
Router(config-pcc-pce)# pce-group blue
Router(config-pcc-pce)# exit
```

### Verification

```
segment-routing
traffic-eng
pcc
pce address ipv4 10.1.1.1
precedence 10
!
pce address ipv4 2.2.2.2
precedence 20
pce-group red
!
pce address ipv4 3.3.3.3
precedence 30
pce-group red
!
pce address ipv4 4.4.4.4
precedence 40
pce-group blue
```

```

!
pce address ipv4 5.5.5.5
 precedence 50
 pce-group blue
!
!
!
!
!

```

### Assign PCE Group to a Candidate Path or ODN Template

Use the **segment-routing traffic-eng policy** *policy* **pce-group** *WORD* command to assign the PCE group to all candidate paths of an SR policy.

Use the **segment-routing traffic-eng policy** *policy* **candidate-paths preference** *pref* **pce-group** *WORD* command to assign the PCE group to a specific candidate path of an SR policy.

Use the **segment-routing traffic-eng on-demand color** *color* **pce-group** *WORD* command to assign the PCE group to on-demand candidate paths triggered by an ODN template.




---

**Note** Only one PCE group can be attached to a given SR policy candidate path.

---

The following example shows how to configure a policy with all candidate paths delegated/reported to PCEs in the default group:

```

Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy A
Router(config-sr-te-policy)# color 100 end-point ipv4 192.168.0.2
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# dynamic
Router(config-sr-te-pp-info)# pcep
Router(config-sr-te-path-pcep)# exit
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path)# exit
Router(config-sr-te-policy)# exit

```

The following example shows how to configure a policy with all candidate paths delegated/reported to PCEs in the red group:

```

Router(config-sr-te)# policy B
Router(config-sr-te-policy)# color 100 end-point ipv4 192.168.0.3
Router(config-sr-te-policy)# pce-group red
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# dynamic
Router(config-sr-te-pp-info)# pcep
Router(config-sr-te-path-pcep)# exit
Router(config-sr-te-pp-info)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path)# exit
Router(config-sr-te-policy)# exit

```

The following example shows how to configure a policy with a specific candidate path (explicit path) reported to PCEs in the blue group:

```

Router(config-sr-te)# policy C
Router(config-sr-te-policy)# color 100 end-point ipv4 192.168.0.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# pce-group blue
Router(config-sr-te-policy-path-pref)# explicit segment-list SLA
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path)# exit
Router(config-sr-te-policy)# exit

```

The following example shows how to configure an ODN template with on-demand candidate paths delegated/reported to PCEs in the blue group:

```

Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# pce-group blue
Router(config-sr-te-color)# dynamic
Router(config-sr-te-color-dyn)#pcep
Router(config-sr-te-color-dyn-pce)#

```

### Running Config

```

segment-routing
traffic-eng
 on-demand color 10
 dynamic
 pcep
 !
 !
 pce-group blue
 !
 policy A
 color 100 end-point ipv4 192.168.0.2
 candidate-paths
 preference 100
 dynamic
 pcep
 !
 !
 !
 !
 policy B
 color 100 end-point ipv4 192.168.0.3
 pce-group red
 candidate-paths
 preference 100
 dynamic
 pcep
 !
 !
 !
 !
 policy C
 color 100 end-point ipv4 192.168.0.4
 candidate-paths
 preference 100
 explicit segment-list SLA
 !
 pce-group blue
 !

```

```

!
!
!
!
end

```

## Verification

```
Router# show segment-routing traffic-eng pcc ipv4 peer
```

```
PCC's peer database:
```

```

```

```

Peer address: 10.1.1.1
 Precedence: 10 (best PCE)
 State up
 Capabilities: Stateful, Update, Segment-Routing, Instantiation

Peer address: 2.2.2.2
 Group: red, Precedence 20
 State up
 Capabilities: Stateful, Update, Segment-Routing, Instantiation

Peer address: 3.3.3.3
 Group: red, Precedence 30
 State up
 Capabilities: Stateful, Update, Segment-Routing, Instantiation

Peer address: 4.4.4.4
 Group: blue, Precedence 40
 State up
 Capabilities: Stateful, Update, Segment-Routing, Instantiation

Peer address: 5.5.5.5
 Group: blue, Precedence 50
 State up
 Capabilities: Stateful, Update, Segment-Routing, Instantiation

```

```
Router# show segment-routing traffic-eng policy name srte_c_100_ep_192.168.0.3
```

```
SR-TE policy database
```

```

```

```

Color: 100, End-point: 192.168.0.3
Name: srte_c_100_ep_192.168.0.3
Status:
 Admin: up Operational: up for 00:13:26 (since Sep 17 22:52:48.365)
Candidate-paths:
 Preference: 100 (configuration)
 Name: B
 Requested BSID: dynamic
 PCC info:
 Symbolic name: cfg_B_discr_100
 PLSP-ID: 2
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 PCE Group: red
 Dynamic (pce 192.168.1.4) (valid)
 Metric Type: TE, Path Accumulated Metric: 10
Attributes:
 Forward Class: 0
 Steering labeled-services disabled: no
 Steering BGP disabled: no

```

```
IPv6 caps enable: no
Invalidation drop enabled: no
```

## BGP SR-TE

BGP may be used to distribute SR Policy candidate paths to an SR-TE head-end. Dedicated BGP SAFI and NLRI have been defined to advertise a candidate path of an SR Policy. The advertisement of Segment Routing policies in BGP is documented in the IETF draft <https://datatracker.ietf.org/doc/draft-ietf-idr-segment-routing-te-policy/>

SR policies with IPv4 and IPv6 end-points can be advertised over BGPv4 or BGPv6 sessions between the SR-TE controller and the SR-TE headend.

The Cisco IOS-XR implementation supports the following combinations:

- IPv4 SR policy advertised over BGPv4 session
- IPv6 SR policy advertised over BGPv4 session
- IPv6 SR policy advertised over BGPv6 session

### Configure BGP SR Policy Address Family at SR-TE Head-End

Perform this task to configure BGP SR policy address family at SR-TE head-end:

#### Procedure

|               | Command or Action                                                                                                                                                     | Purpose                                                                                                               |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b>                                                                                                                                                      |                                                                                                                       |
| <b>Step 2</b> | <b>router bgp</b> <i>as-number</i><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config)# router bgp 65000</pre>                                                    | Specifies the BGP AS number and enters the BGP configuration mode, allowing you to configure the BGP routing process. |
| <b>Step 3</b> | <b>bgp router-id</b> <i>ip-address</i><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-bgp)# bgp router-id 10.1.1.1</pre>                                      | Configures the local router with a specified router ID.                                                               |
| <b>Step 4</b> | <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>sr-policy</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 sr-policy</pre> | Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.                     |
| <b>Step 5</b> | <b>exit</b>                                                                                                                                                           |                                                                                                                       |

|               | Command or Action                                                                                                                                                         | Purpose                                                                                                                |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Step 6</b> | <b>neighbor</b> <i>ip-address</i><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-bgp) #<br><b>neighbor 10.10.0.1</b>                                                   | Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer. |
| <b>Step 7</b> | <b>remote-as</b> <i>as-number</i><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-bgp-nbr) #<br><b>remote-as 1</b>                                                      | Creates a neighbor and assigns a remote autonomous system number to it.                                                |
| <b>Step 8</b> | <b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>sr-policy</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-bgp-nbr) #<br><b>address-family ipv4 sr-policy</b> | Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.                      |
| <b>Step 9</b> | <b>route-policy</b> <i>route-policy-name</i> { <b>in</b>   <b>out</b> }<br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-bgp-nbr-af) #<br><b>route-policy pass out</b>   | Applies the specified policy to IPv4 or IPv6 unicast routes.                                                           |

#### Example: BGP SR-TE with BGPv4 Neighbor to BGP SR-TE Controller

The following configuration shows the an SR-TE head-end with a BGPv4 session towards a BGP SR-TE controller. This BGP session is used to signal both IPv4 and IPv6 SR policies.

```

router bgp 65000
bgp router-id 10.1.1.1
!
address-family ipv4 sr-policy
!
address-family ipv6 sr-policy
!
neighbor 10.1.3.1
remote-as 10
description *** eBGP session to BGP SRTE controller ***
address-family ipv4 sr-policy
route-policy pass in
route-policy pass out
!
address-family ipv6 sr-policy
route-policy pass in
route-policy pass out
!
!
!

```

#### Example: BGP SR-TE with BGPv6 Neighbor to BGP SR-TE Controller

The following configuration shows an SR-TE head-end with a BGPv6 session towards a BGP SR-TE controller. This BGP session is used to signal both IPv4 and IPv6 SR policies.

```
router bgp 65000
 bgp router-id 10.1.1.1
 address-family ipv4 sr-policy
 !
 address-family ipv6 sr-policy
 !
 neighbor 3001::10:1:3:1
 remote-as 10
 description *** eBGP session to BGP SRTE controller ***
 address-family ipv4 sr-policy
 route-policy pass in
 route-policy pass out
 !
 address-family ipv6 sr-policy
 route-policy pass in
 route-policy pass out
 !
 !
 !
 !
```


## Traffic Steering

### Automated Steering

Automated steering (AS) allows service traffic to be automatically steered onto the required transport SLA path programmed by an SR policy.

With AS, BGP automatically steers traffic onto an SR Policy based on the next-hop and color of a BGP service route. The color of a BGP service route is specified by a color extended community attribute. This color is used as a transport SLA indicator, such as min-delay or min-cost.

When the next-hop and color of a BGP service route matches the end-point and color of an SR Policy, BGP automatically installs the route resolving onto the BSID of the matching SR Policy. Recall that an SR Policy on a head-end is uniquely identified by an end-point and color.



When a BGP route has multiple extended-color communities, each with a valid SR Policy, the BGP process installs the route on the SR Policy giving preference to the color with the highest numerical value.

The granularity of AS behaviors can be applied at multiple levels, for example:



- At a service level—When traffic destined to all prefixes in a given service is associated to the same transport path type. All prefixes share the same color.
- At a destination/prefix level—When traffic destined to a prefix in a given service is associated to a specific transport path type. Each prefix could be assigned a different color.
- At a flow level—When flows destined to the same prefix are associated with different transport path types

AS behaviors apply regardless of the instantiation method of the SR policy, including:

- On-demand SR policy
- Manually provisioned SR policy
- PCE-initiated SR policy

See the [Verifying BGP VRF Information, on page 233](#) and [Verifying Forwarding \(CEF\) Table, on page 234](#) sections for sample output that shows AS implementation.

## Color-Only Automated Steering

Color-only steering is a traffic steering mechanism where a policy is created with given color, regardless of the endpoint.

You can create an SR-TE policy for a specific color that uses a NULL end-point (0.0.0.0 for IPv4 NULL, and ::0 for IPv6 NULL end-point). This means that you can have a single policy that can steer traffic that is based on that color and a NULL endpoint for routes with a particular color extended community, but different destinations (next-hop).



**Note** Every SR-TE policy with a NULL end-point must have an explicit path-option. The policy cannot have a dynamic path-option (where the path is computed by the head-end or PCE) since there is no destination for the policy.

You can also specify a color-only (CO) flag in the color extended community for overlay routes. The CO flag allows the selection of an SR-policy with a matching color, regardless of endpoint Sub-address Family Identifier (SAFI) (IPv4 or IPv6). See [Setting CO Flag, on page 297](#).

### Configure Color-Only Steering

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy P1
Router(config-sr-te-policy)# color 1 end-point ipv4 0.0.0.0

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy P2
Router(config-sr-te-policy)# color 2 end-point ipv6 ::0
```

```

Router# show running-configuration
segment-routing
traffic-eng
policy P1
 color 1 end-point ipv4 0.0.0.0
!
policy P2
 color 2 end-point ipv6 ::
!
!
end

```

## Setting CO Flag

The BGP-based steering mechanism matches BGP color and next-hop with that of an SR-TE policy. If the policy does not exist, BGP requests SR-PCE to create an SR-TE policy with the associated color, end-point, and explicit paths. For color-only steering (NULL end-point), you can configure a color-only (CO) flag as part of the color extended community in BGP.



**Note** See [Color-Only Automated Steering, on page 296](#) for information about color-only steering (NULL end-point).

The behavior of the steering mechanism is based on the following values of the CO flags:

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>co-flag 00</b> | <ol style="list-style-type: none"> <li>1. The BGP next-hop and color &lt;N, C&gt; is matched with an SR-TE policy of same &lt;N, C&gt;.</li> <li>2. If a policy does not exist, then IGP path for the next-hop N is chosen.</li> </ol>                                                                                                                                                                                                                                                                                                   |
| <b>co-flag 01</b> | <ol style="list-style-type: none"> <li>1. The BGP next-hop and color &lt;N, C&gt; is matched with an SR-TE policy of same &lt;N, C&gt;.</li> <li>2. If a policy does not exist, then an SR-TE policy with NULL end-point with the same address-family as N and color C is chosen.</li> <li>3. If a policy with NULL end-point with same address-family as N does not exist, then an SR-TE policy with any NULL end-point and color C is chosen.</li> <li>4. If no match is found, then IGP path for the next-hop N is chosen.</li> </ol> |

### Configuration Example

```

Router(config)# extcommunity-set opaque overlay-color
Router(config-ext)# 1 co-flag 01
Router(config-ext)# end-set
Router(config)#
Router(config)# route-policy color
Router(config-rpl)# if destination in (5.5.5.1/32) then
Router(config-rpl-if)# set extcommunity color overlay-color
Router(config-rpl-if)# endif
Router(config-rpl)# pass
Router(config-rpl)# end-policy

```

```
Router(config)#
```

## Address-Family Agnostic Automated Steering

Address-family agnostic steering uses an SR-TE policy to steer both labeled and unlabeled IPv4 and IPv6 traffic. This feature requires support of IPv6 encapsulation (IPv6 caps) over IPV4 endpoint policy.

IPv6 caps for IPv4 NULL end-point is enabled automatically when the policy is created in Segment Routing Path Computation Element (SR-PCE). The binding SID (BSID) state notification for each policy contains an "ipv6\_caps" flag that notifies SR-PCE clients (PCC) of the status of IPv6 caps (enabled or disabled).

An SR-TE policy with a given color and IPv4 NULL end-point could have more than one candidate path. If any of the candidate paths has IPv6 caps enabled, then all of the remaining candidate paths need IPv6 caps enabled. If IPv6 caps is not enabled on all candidate paths of same color and end-point, traffic drops can occur.

You can disable IPv6 caps for a particular color and IPv4 NULL end-point using the **ipv6 disable** command on the local policy. This command disables IPv6 caps on all candidate paths that share the same color and IPv4 NULL end-point.

### Disable IPv6 Encapsulation

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy P1
Router(config-sr-te-policy)# color 1 end-point ipv4 0.0.0.0
Router(config-sr-te-policy)# ipv6 disable
```

## Per-Flow Automated Steering

*Table 38: Feature History Table*

| Feature Name                                                      | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Per-Flow Automated Steering: L3 / L2 BGP Services + BSID Steering | Release 7.4.1       | <p>This feature introduces support for BGP VPNv6 (6VPE) and BGP EVPN (single-home/multi-homed) over PFP, labeled traffic (Binding SID as top-most label in the stack) steering over per-flow policy (PFP).</p> <p>An ingress QoS policy applied to an input interface is used to classify flows and set corresponding MPLS experimental values.</p> |

The steering of traffic through a Segment Routing (SR) policy is based on the candidate paths of that policy. For a given policy, a candidate path specifies the path to be used to steer traffic to the policy's destination. The policy determines which candidate path to use based on the candidate path's preference and state. The candidate path that is valid and has the highest preference is used to steer all traffic using the given policy. This type of policy is called a Per-Destination Policy (PDP).

Per-Flow Automated Traffic Steering using SR-TE Policies introduces a way to steer traffic on an SR policy based on the attributes of the incoming packets, called a Per-Flow Policy (PFP).


A PFP provides up to 8 "ways" or options to the endpoint. With a PFP, packets are classified by a classification policy and marked using internal tags called forward classes (FCs). The FC setting of the packet selects the "way". For example, this "way" can be a traffic-engineered SR path, using a low-delay path to the endpoint. The FC is represented as a numeral with a value of 0 to 7.

A PFP defines an array of FC-to-PDP mappings. A PFP can then be used to steer traffic into a given PDP based on the FC assigned to a packet.

As with PDPs, PFPs are identified by a {headend, color, endpoint} tuple. The color associated with a given FC corresponds to a valid PDP policy of that color and same endpoint as the parent PFP. So PFP policies contain mappings of different FCs to valid PDP policies of different colors. Every PFP has an FC designated as its default FC. The default FC is associated to packets with a FC undefined under the PFP or for packets with a FC with no valid PDP policy.

The following example shows a per-flow policy from Node1 to Node4:


**Figure 14: PFP Example**



- FC=0 -> shortest path to Node4
  - IGP shortest path = 16004
- FC=1 -> Min-delay path to Node4
  - SID list = {16002,16004}

The same on-demand instantiation behaviors of PDPs apply to PFPs. For example, an edge node automatically (on demand) instantiates Per-Flow SR Policy paths to an endpoint by service route signaling. Automated Steering steers the service route in the matching SR Policy.

**Figure 15: PFP with ODN Example**



Like PDPs, PFPs have a binding SID (BSID). Existing SR-TE automated steering (AS) mechanisms for labeled traffic (via BSID) and unlabeled traffic (via BGP) onto a PFP is similar to that of a PDP. For example, a packet having the BSID of a PFP as the top label is steered onto that PFP. The classification policy on the

ingress interface marks the packet with an FC based on the configured class-map. The packet is then steered to the PDP that corresponds to that FC.

### Usage Guidelines and Limitations

The following guidelines and limitations apply to the platform when acting as a head-end of a PFP policy:

- BGP IPv4 unicast over PFP (steered via ODN/AS) is supported
- BGP IPv6 unicast (with IPv4 next-hop [6PE]) over PFP (steered via ODN/AS) is supported
- BGP IPv6 unicast (with IPv6 next-hop) over PFP (steered via ODN/AS) is supported
- BGP VPNv4 over PFP (steered via ODN/AS) is supported
- BGP VPNv6 (6VPE) over PFP (steered via ODN/AS) is supported
- BGP EVPN (single-home/multi-homed) over PFP (steered via ODN/AS) is supported
- Pseudowire and VPLS over PFP (steered with preferred-path) are supported
- BGP multipath is supported
- BGP PIC is not supported
- Labeled traffic (Binding SID as top-most label in the stack) steered over PFP is supported
- When not explicitly configured, FC 0 is the default FC.
- A PFP is considered valid as long as its default FC has a valid PDP.
- An ingress QoS policy applied to an input interface is used to classify flows and set corresponding MPLS experimental values.
- PFP implementation is accomplished with a double-pass through the ASIC (recirculation).
  - In the first pass, an ingress QoS policy applied to an input interface is used to classify flows and set MPLS EXP values, alongside push of service label and PFP Binding SID label.
  - In the absence of any ingress QoS policy, the default behavior is to copy PREC/DSCP/EXP to PFP BSID MPLS EXP.
  - In the second pass, a forwarding lookup based on PFP Binding SID label plus MPLS EXP is used to resolve to a given PFP's PDP.
- The PFP's BSID is allocated from a user-configured MPLS label block; see [Configuring PFP BSID Label Block, on page 301](#).
- The following counters are supported:
  - PFP's BSID counter (packet, bytes)
  - Per-FC counters (packet, byte)
    - Collected from the PDP's segment-list-per-path egress counters
  - If an SR policy is used for more than one purpose (as a regular policy as well as a PDP under one or more PFPs), then the collected counters will represent the aggregate of all contributions. To preserve independent counters, it is recommended that an SR policy be used only for one purpose.

- Inbound packet classification, based on the following fields, is supported:
  - IP precedence
  - IP DSCP
  - L3 ACL-based (L3 source/destination IP; L4 source/destination port)
  - MPLS EXP
  - Layer 2 CoS
  - MAC ACL
- A color associated with a PFP SR policy cannot be used by a non-PFP SR policy. For example, if a per-flow ODN template for color 100 is configured, then the system will reject the configuration of any non-PFP SR policy using the same color. You must assign different color value ranges for PFP and non-PFP SR policies.

### Configuring PFP BSID Label Block

Implementation on NCS platforms requires that the BSID assigned to a PFP be allocated from a preconfigured label block. The BSID is a local segment.

This label range cannot overlap with the existing SRLB or SRGB ranges allocated on the platform.

To configure the MPLS label block for PFP BSID allocation, use the **block name name type pfp start starting-value {end ending-value | size size } [client word]** command.

This example shows how to allocate a block of labels based on the size of the block:

```
Router(config)# mpls label blocks
Router(config-mpls-lbl-blks)# block name sample-pfp-bsid-block type pfp start 40000 size
1000 client any
```

This example shows how to allocate a block of labels based on specific starting and ending values:

```
Router(config)# mpls label blocks
Router(config-mpls-lbl-blks)# block name sample-pfp-bsid-block type pfp start 40000 end
41000 client any
```

### Configuring ODN Template for PFP Policies: Example

The following example depicts an ODN template for PFP policies that includes three FCs.

The example also includes the corresponding ODN templates for PDPs as follows:

- FC0 (default FC) mapped to color 10 = Min IGP path
- FC1 mapped to color 20 = Flex Algo 128 path
- FC2 mapped to color 30 = Flex Algo 129 path

```
segment-routing
traffic-eng
on-demand color 10
dynamic
```

```

metric
 type igp
!
!
!
on-demand color 20
constraints
 segments
 sid-algorithm 128
!
!
!
on-demand color 30
constraints
 segments
 sid-algorithm 129
!
!
!
on-demand color 1000
per-flow
 forward-class 0 color 10
 forward-class 1 color 20
 forward-class 2 color 30

```

### Manually Configuring a PFP and PDPs: Example

The following example depicts a manually defined PFP that includes three FCs and corresponding manually defined PDPs.

The example also includes the corresponding PDPs as follows:

- FC0 (default FC) mapped to color 10 = Min IGP path
- FC1 mapped to color 20 = Min TE path
- FC2 mapped to color 30 = Min delay path

```

mpls label blocks
block name sample-pfp-bsid-block type pfp start 400000 size 1000 client any
!
!
segment-routing
traffic-eng
policy MyPerFlow
 color 1000 end-point ipv4 10.1.1.4
 candidate-paths
 preference 100
 per-flow
 forward-class 0 color 10
 forward-class 1 color 20
 forward-class 2 color 30
!
policy MyLowIGP
 color 10 end-point ipv4 10.1.1.4
 candidate-paths
 preference 100
 dynamic
 metric type igp
!
policy MyLowTE
 color 20 end-point ipv4 10.1.1.4

```

```

candidate-paths
 preference 100
 dynamic
 metric type te
!
policy MyLowDelay
 color 30 end-point ipv4 10.1.1.4
 candidate-paths
 preference 100
 dynamic
 metric type delay

```

### Configuring Ingress Classification: Example

An MQC QoS policy is used to classify and mark traffic to a corresponding MPLS experimental value.

The following shows an example of such ingress classification policy:

```

class-map match-any MinDelay
 match dscp 46
end-class-map
!
class-map match-any PremiumHosts
 match access-group ipv4 PrioHosts
end-class-map
!
!
policy-map MyPerFlowClassificationPolicy
 class MinDelay
 set mpls experimental imposition 2
 !
 class PremiumHosts
 set mpls experimental imposition 1
 !
 class class-default
 !
end-policy-map
!
interface GigabitEthernet0/0/0/0
 description PE_Ingress_Interface
 service-policy input MyPerFlowClassificationPolicy
!

```

### Determining Per-Flow Policy State

A PFP is brought down for the following reasons:

- The PDP associated with the default FC is in a down state.
- All FCs are associated with PDPs in a down state.
- The FC assigned as the default FC is missing in the forward class mapping.

Scenario 1—FC 0 (default FC) is not configured in the FC mappings below:

```

policy foo
 color 1 end-point ipv4 10.1.1.1
 per-flow
 forward-class 1 color 10
 forward-class 2 color 20

```



Scenario 2—FC 1 is configured as the default FC, however it is not present in the FC mappings:

```
policy foo
 color 1 end-point ipv4 10.1.1.1
 per-flow
 forward-class 0 color 10
 forward-class 2 color 20
 forward-class default 1
```

## Using Binding Segments

The binding segment is a local segment identifying an SR-TE policy. Each SR-TE policy is associated with a binding segment ID (BSID). The BSID is a local label that is automatically allocated for each SR-TE policy when the SR-TE policy is instantiated.

BSID can be used to steer traffic into the SR-TE policy and across domain borders, creating seamless end-to-end inter-domain SR-TE policies. Each domain controls its local SR-TE policies; local SR-TE policies can be validated and rerouted if needed, independent from the remote domain's head-end. Using binding segments isolates the head-end from topology changes in the remote domain.

Packets received with a BSID as top label are steered into the SR-TE policy associated with the BSID. When the BSID label is popped, the SR-TE policy's SID list is pushed.

BSID can be used in the following cases:

- Multi-Domain (inter-domain, inter-autonomous system)—BSIDs can be used to steer traffic across domain borders, creating seamless end-to-end inter-domain SR-TE policies.
- Large-Scale within a single domain—The head-end can use hierarchical SR-TE policies by nesting the end-to-end (edge-to-edge) SR-TE policy within another layer of SR-TE policies (aggregation-to-aggregation). The SR-TE policies are nested within another layer of policies using the BSIDs, resulting in seamless end-to-end SR-TE policies.
- Label stack compression—If the label-stack size required for an SR-TE policy exceeds the platform capability, the SR-TE policy can be seamlessly stitched to, or nested within, other SR-TE policies using a binding segment.
- BGP SR-TE Dynamic—The head-end steers the packet into a BGP-based FIB entry whose next hop is a binding-SID.

## L2VPN Preferred Path

EVPN VPWS Preferred Path over SR-TE Policy feature allows you to set the preferred path between the two end-points for EVPN VPWS pseudowire (PW) using SR-TE policy.

L2VPN VPLS or VPWS Preferred Path over SR-TE Policy feature allows you to set the preferred path between the two end-points for L2VPN Virtual Private LAN Service (VPLS) or Virtual Private Wire Service (VPWS) using SR-TE policy.

Refer to the [EVPN VPWS Preferred Path over SR-TE Policy](#) and [L2VPN VPLS or VPWS Preferred Path over SR-TE Policy](#) sections in the "L2VPN Services over Segment Routing for Traffic Engineering Policy" chapter of the *L2VPN and Ethernet Services Configuration Guide*.

## Static Route over Segment Routing Policy

This feature allows you to specify a Segment Routing (SR) policy as an interface type when configuring static routes for MPLS data planes.

For information on configuring static routes, see the "Implementing Static Routes" chapter in the *Routing Configuration Guide*.

### Configuration Example

The following example depicts a configuration of a static route for an IPv4 destination over an SR policy.

```
Router(config)# router static
Router(config-static)# address-family ipv4 unicast
Router(config-static-afi)# 10.1.100.100/32 sr-policy sample-policy
```

### Running Configuration

```
Router# show run segment-routing traffic-eng

segment-routing
 traffic-eng
 segment-list sample-SL
 index 10 mpls adjacency 10.1.1.102
 index 20 mpls adjacency 10.1.1.103
 !
 policy sample-policy
 color 777 end-point ipv4 10.1.1.103
 candidate-paths
 preference 100
 explicit segment-list sample-SL

Router# show run segment-routing traffic-eng

router static
 address-family ipv4 unicast
 10.1.1.4/32 sr-policy srte_c_200_ep_10.1.1.4
 !
 !
```

### Verification

```
Router# show segment-routing traffic-eng policy candidate-path name sample-policy

SR-TE policy database

Color: 777, End-point: 10.1.1.103
Name: srte_c_777_ep_10.1.1.103
Status:
 Admin: up Operational: up for 00:06:35 (since Jan 17 14:34:35.120)
Candidate-paths:
 Preference: 100 (configuration) (active)
 Name: sample-policy
 Requested BSID: dynamic
 PCC info:
 Symbolic name: cfg_sample-policy_discr_100
 PLSP-ID: 5
 Constraints:
 Protection Type: protected-preferred
```

```

Maximum SID Depth: 9
Explicit: segment-list sample-SL (valid)
Weight: 1, Metric Type: TE
 SID[0]: 100102 [Prefix-SID, 10.1.1.102]
 SID[1]: 100103 [Prefix-SID, 10.1.1.103]
Attributes:
 Binding SID: 24006
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
 Max Install Standby Candidate Paths: 0

```

Router# **show static sr-policy sample-policy**

```

SR-Policy-Name State Binding-label Interface ifhandle VRF
Paths
sample-policy Up 24006 srte_c_777_ep_10.1.1.103 0x2000803c default
10.1.100.100/32
Reference count=1, Internal flags=0x0
Last Policy notification was Up at Jan 17 13:39:46.478

```

Router# **show route 10.1.100.100/32**

```

Routing entry for 10.1.100.100/32
 Known via "static", distance 1, metric 0
 Installed Jan 17 14:35:40.969 for 00:06:38
 Routing Descriptor Blocks
 directly connected, via srte_c_777_ep_10.1.1.103
 Route metric is 0
 No advertising protos.

```

Router# **show route 10.1.100.100/32 detail**

```

Routing entry for 10.1.100.100/32
 Known via "static", distance 1, metric 0
 Installed Jan 17 14:35:40.969 for 00:06:44
 Routing Descriptor Blocks
 directly connected, via srte_c_777_ep_10.1.1.103
 Route metric is 0
 Label: None
 Tunnel ID: None
 Binding Label: 0x5dc6 (24006)
 Extended communities count: 0
 NHID: 0x0 (Ref: 0)
 Route version is 0x1 (1)
 No local label
 IP Precedence: Not Set
 QoS Group ID: Not Set
 Flow-tag: Not Set
 Fwd-class: Not Set
 Route Priority: RIB_PRIORITY_STATIC (9) SVD Type RIB_SVD_TYPE_LOCAL
 Download Priority 3, Download Version 3169
 No advertising protos.

```

Router# **show cef 10.1.100.100/32**

```

10.1.100.100/32, version 3169, internal 0x1000001 0x30 (ptr 0x8b1b95d8) [1], 0x0 (0x0), 0x0
(0x0)
Updated Jan 17 14:35:40.971
Prefix Len 32, traffic index 0, precedence n/a, priority 3

```

```

gateway array (0x8a92f228) reference count 1, flags 0x2010, source rib (7), 0 backups
 [1 type 3 flags 0x48441 (0x8a9d1b68) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Jan 17 14:35:40.971
LDI Update time Jan 17 14:35:40.972
 via local-label 24006, 3 dependencies, recursive [flags 0x0]
 path-idx 0 NHID 0x0 [0x8ac59f30 0x0]
 recursion-via-label
 next hop via 24006/1/21

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y recursive 24006/1

```

## Autoroute Include

**Table 39: Feature History Table**

| Feature Name      | Release       | Description                                                                                                                                                                        |
|-------------------|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Autoroute Include | Release 7.3.2 | This feature allows you to steer specific IGP (IS-IS, OSPF) prefixes, or all prefixes, over non-shortest paths and to divert the traffic for those prefixes on to an SR-TE policy. |

You can configure SR-TE policies with Autoroute Include to steer specific IGP (IS-IS, OSPF) prefixes, or all prefixes, over non-shortest paths and to divert the traffic for those prefixes on to the SR-TE policy.

The **autoroute include all** option applies Autoroute Announce functionality for all destinations or prefixes.

The **autoroute include ipv4 address** option applies Autoroute Destination functionality for the specified destinations or prefixes. This option is supported for IS-IS only; it is not supported for OSPF.

The Autoroute SR-TE policy adds the prefixes into the IGP, which determines if the prefixes on the endpoint or downstream of the endpoint are eligible to use the SR-TE policy. If a prefix is eligible, then the IGP checks if the prefix is listed in the Autoroute Include configuration. If the prefix is included, then the IGP downloads the prefix route with the SR-TE policy as the outgoing path.

### Usage Guidelines and Limitations

- Autoroute Include supports three metric types:
  - Default (no metric): The path over the SR-TE policy inherits the shortest path metric.
  - Absolute (constant) metric: The shortest path metric to the policy endpoint is replaced with the configured absolute metric. The metric to any prefix that is Autoroute Included is modified to the absolute metric. Use the **autoroute metric constant** *constant-metric* command, where *constant-metric* is from 1 to 2147483647.
  - Relative metric: The shortest path metric to the policy endpoint is modified with the relative value configured (plus or minus). Use the **autoroute metric relative** *relative-metric* command, where *relative-metric* is from -10 to +10.



**Note** To prevent load-balancing over IGP paths, you can specify a metric that is lower than the value that IGP takes into account for autorouted destinations (for example, **autoroute metric relative -1**).

- LDP over SR-TE not supported.
- LDP to SR-TE interworking is not supported.
- Static route over SR-TE is not supported.

### Configuration Examples

The following example shows how to configure **autoroute include** for all prefixes:

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)#policy P1
Router(config-sr-te-policy)# color 20 end-point ipv4 10.1.1.2
Router(config-sr-te-policy)# autoroute include all
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-pp-index)# explicit segment-list Plist-1
```

The following example shows how to configure **autoroute include** for the specified IPv4 prefixes:



**Note** This option is supported for IS-IS only; it is not supported for OSPF.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)#policy P1
Router(config-sr-te-policy)# color 20 end-point ipv4 10.1.1.2
Router(config-sr-te-policy)# autoroute include ipv4 10.1.1.21/32
Router(config-sr-te-policy)# autoroute include ipv4 10.1.1.23/32
Router(config-sr-te-policy)# autoroute metric constant 1
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-pp-index)# explicit segment-list Plist-1
```

## Policy-Based Tunnel Selection for SR-TE Policy

Policy-Based Tunnel Selection (PBTS) is a mechanism that lets you direct traffic into specific SR-TE policies based on different classification criteria. PBTS benefits Internet service providers (ISPs) that carry voice and data traffic through their networks, who want to route this traffic to provide optimized voice service.

PBTS works by selecting SR-TE policies based on the classification criteria of the incoming packets, which are based on the IP precedence, experimental (EXP), differentiated services code point (DSCP), or type of

service (ToS) field in the packet. Default-class configured for paths is always zero (0). If there is no TE for a given forward-class, then the default-class (0) will be tried. If there is no default-class, then the packet is dropped. PBTS supports up to seven (exp 1 - 7) EXP values associated with a single SR-TE policy.

For more information about PBTS, refer to the "Policy-Based Tunnel Selection" section in the *MPLS Configuration Guide*.

### Configure Policy-Based Tunnel Selection for SR-TE Policies

The following section lists the steps to configure PBTS for an SR-TE policy.



**Note** Steps 1 through 4 are detailed in the "Implementing MPLS Traffic Engineering" chapter of the *MPLS Configuration Guide*.

1. Define a class-map based on a classification criteria.
2. Define a policy-map by creating rules for the classified traffic.
3. Associate a forward-class to each type of ingress traffic.
4. Enable PBTS on the ingress interface, by applying this service-policy.
5. Create one or more egress SR-TE policies (to carry packets based on priority) to the destination and associate the egress SR-TE policy to a forward-class.

### Configuration Example

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy POLICY-PBTS
Router(config-sr-te-policy)# color 1001 end-point ipv4 10.1.1.20
Router(config-sr-te-policy)# autoroute
Router(config-sr-te-policy-autoroute)# include all
Router(config-sr-te-policy-autoroute)# forward-class 1
Router(config-sr-te-policy-autoroute)# exit
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 1
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST1
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path-pref)# exit
Router(config-sr-te-policy-path-pref)# preference 2
Router(config-sr-te-policy-path-pref)# dynamic
Router(config-sr-te-policy-path-pref)# metric
Router(config-sr-te-policy-path-pref)# type te
Router(config-sr-te-policy-path-pref)# commit
```

### Running Configuration

```
segment-routing
traffic-eng
policy POLICY-PBTS
color 1001 end-point ipv4 10.1.1.20
autoroute
include all
forward-class 1
!
```

```

candidate-paths
 preference 1
 explicit segment-list SIDLIST1
 !
 preference 2
 dynamic
 metric
 type te

```

## SR-TE Automated Steering Without BGP Prefix Path Label

Table 40: Feature History Table

| Feature Name                                           | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE Automated Steering Without BGP Prefix Path Label | Release 7.9.1       | <p>This feature allows traffic to a BGP service route to be steered over an SR-TE policy using automated-steering principles without imposing the service route's prefix label.</p> <p>This feature allows you to deploy a centralized BGP EPE solution for 6PE in an SR-MPLS network.</p> <p>This feature introduces the <b>bgp prefix-path-label ignore</b> command.</p> |

This feature allows traffic to a BGP service route to be steered over an SR-TE policy using automated-steering principles without imposing the service route's prefix label (see [Automated Steering, on page 295](#)). BGP ignores the programming of the label associated with a prefix path (for example, 6PE/VPN label) when recursing onto the BSID of an SR-TE policy with this feature enabled.

This feature allows you to deploy a [centralized BGP EPE solution for 6PE in an SR-MPLS network](#).

### Usage Guidelines and Limitations

This functionality applies to local/manually configured SR-TE candidate-paths.

This functionality does not apply to on-demand SR-TE candidate-paths triggered by ODN.

This functionality does not apply to SR-TE candidate-paths instantiated via PCEP (PCE-initiated) or BGP-TE.

### Configuration

Use the **bgp prefix-path-label ignore** command in SR-TE policy steering config mode to indicate BGP ignores the programming of the label associated with a prefix path (for example, 6PE/VPN label) when recursing onto the BSID of an SR-TE policy with this feature enabled.

```

Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# steering

```

```

Router(config-sr-te-policy-steering)# bgp prefix-path-label ignore
Router(config-sr-te-policy-steering)# exit
Router(config-sr-te-policy)# color 100 end-point ipv4 0.0.0.0
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list sample-sl

```

## Verification

The following output displays the SR-TE policy (SR policy color 100, IPv4 null end-point) details showing the ignore prefix label steering behavior:

```
Router# show segment-routing traffic-eng policy candidate-path name FOO private
```

```

SR-TE policy database

Color: 100, End-point: 0.0.0.0 ID: 3
 Name: srte_c_100_ep_0.0.0.0
 Status:
 Admin: up Operational: up for 00:10:07 (since Feb 2 12:58:43.554)
 Candidate-paths:
 Preference: 100 (configuration) (active)
 Originator: ASN 0 node-address <None> discriminator: 100
 Name: FOO
 Requested BSID: dynamic
 Constraints:
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 ID: 1
 Source: 20.1.0.100
 Stale: no
 Checkpoint flags: 0x00000000
 Steering:
 Client: BGP
 Disabled: no
 Ignore prefix label: yes
 Explicit: segment-list sample-sl (valid)
 Weight: 1, Metric Type: TE
 IGP area: 2
 SID[0]: 16102 [Prefix-SID: 20.1.0.102, Algorithm: 0]
 SID[1]: 16103 [Prefix-SID: 20.1.0.103, Algorithm: 0]
 SID[2]: 24008 [Adjacency-SID, 15:15:15::4 - 15:15:15::5]

LSPs:

. . .

Attributes:
 Binding SID: 24030
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
 Max Install Standby Candidate Paths: 0
Notification to clients:
 Binding SID: 24030
 Bandwidth : 0 Kbps (0 Kbps)
 State: UP
 Flags: [add] [ipv6_caps] [ignore_prefix_label]
 Metric Type: NONE
 Metric Value: 2147483647
 Admin Distance: 100

```



```

ifhandle: 0x00000170
Source: 20.1.0.100
Transition count: 1
LSPs created count: 1
Reoptimizations completed count: 1
Retry Action Flags: 0x00000000, ()
Last Retry Timestamp: never (0 seconds ago)
Policy reference: 0x1f81e50

```

The following output shows that BGP received the ignore prefix label steering behavior for an SR policy color 100 and IPv4 null end-point:

```

Router# show bgp nexthops 0.0.0.0 color 100 | include "BGP prefix label"

 BGP prefix label: [No]

```

The following output shows the details for a IPv6 BGP global route (151:1::/64) learned from an IPv4 next-hop (6PE) that is steered over an SR policy (BSID 24030). BGP programs the prefix path ignoring its label.

```

Router# show bgp ipv6 labeled-unicast 151:1::/64 detail

BGP routing table entry for 151:1::/64
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 2003 2003
 Local Label: 81718 (no rewrite);
 Flags: 0x003e1001+0x30010000;
Last Modified: Nov 23 16:59:17.891 for 00:00:03
Paths: (400 available, best #1)
 Advertised IPv6 Unicast paths to update-groups (with more than one peer):
 0.2
 Advertised IPv6 Labeled-unicast paths to update-groups (with more than one peer):
 0.3
 Path #1: Received by speaker 0
 Flags: 0xa480000001060205+0x01, import: 0x020
 Advertised IPv6 Unicast paths to update-groups (with more than one peer):
 0.2
 Advertised IPv6 Labeled-unicast paths to update-groups (with more than one peer):
 0.3
 300, (Received from a RR-client)
 5.5.3.1 C:100 (bsid:24030) (admin 100) (metric 2147483647) from 4.4.4.1 (5.5.5.5),
if-handle 0x00000170
 Prefix Label not imposed due to SR policy config

```

## Use Case: Centralized BGP EPE for 6PE in an SR-MPLS Network

In this use case, an operator wants to control the egress peering router/egress transit autonomous system (AS) used by selected Internet IPv6 prefixes. To achieve this, SR policies with explicit paths are used to steer traffic to an intended egress peering router and intended egress transit AS. BGP-EPE SIDs are used in order to force traffic onto an intended egress transit AS. Traffic steering follows SR-TE automated-steering principles.

The following key features enable the use-case:

- [SR-TE Explicit Segment Lists with Mix of IPv4 and IPv6 Segments](#) — Allows the last segment of an SR policy's SID list to be associated with an EPE-enabled BGPv6 neighbor.
- [SR-TE Automated Steering Without BGP Prefix Path Label](#) — Allows traffic to an Internet IPv6 prefix to be steered over an SR-TE policy without imposing the 6PE label learned from the original advertising router.


## Topology

The below topology shows a single-stack IPv4 SR-MPLS and BGP-free network that delivers Internet IPv6 connectivity using 6PE.

Peering routers 2 and 3 learn IPv6 reachability through transit AS's (ASBR routers 4, 5, 6) via eBGPv6 neighbors.

BGP EPE SIDs are enabled on external BGPv6 neighbors at router 2 (for example, EPE label 24024) and router 3 (for example, EPE labels 24035 and 24036).


**Figure 16: Network Setup**



## BGP EPE Propagation via BGP-LS


Peering routers 2 and 3 advertise their EPE-enabled neighbors via BGP-LS. As a result, ingress router node 1 learns those EPE-enabled neighbors via BGP-LS. This allows the SR-TE database at the ingress router to include the external links.

Figure 17: BGP-LS

**Steady State (Non-Traffic-Engineered)**

At steady state, router 1 selects (as BGP best-path) the path from router 2 for IPv6 prefix 2001:db8:abcd::/48. Traffic to this prefix is sent over the SR-native LSP associated with router 2 (prefix SID 16002) along side the advertised 6PE label.

**Figure 18: Steady State**



### EPE Traffic-Engineered Path

To create a traffic-engineered path that steers traffic to an intended egress peering router/egress transit AS (for example, node 3/ASBR node 6), an SR policy can be configured at ingress border router 1 with the following:

- An IPv4 null (0.0.0.0) end-point, in order to perform color-only automated steering (see [Color-Only Automated Steering, on page 296](#)).
- An explicit segment list with SIDs corresponding to the intended egress node (for example, node 3) and the intended egress peering link (for example, ASBR node 6).
- The **bgp prefix-label ignore** steering command in order to indicate BGP to ignore the programming of the 6PE label associated with the prefix path.


```
segment-routing
traffic-eng
 segment-list sl-to_3-epe_36
 index 10 mpls adjacency 10.1.1.3
 index 20 mpls adjacency 2001:db8:20:3:6::3
 !
policy FOO
steering
 bgp prefix-path-label ignore
!
color 10 end-point ipv4 0.0.0.0
candidate-paths
 preference 100
```

When a given IPv6 Internet destination needs to be steered over an intended egress peering router/egress AS, the operator can perform one of the following:

- Advertise a new BGP prefix path from a Route Server that includes a color extended community value equal to the color of the SR-TE policy for the intended egress peering router/egress AS, or
- Apply a color extended community value equal to the color of the SR-TE policy for the intended egress peering router/egress AS at the peering router advertising the best path (for example, node 2), as shown below.



**Figure 19: EPE Traffic-Engineered Path**



The following output depicts the details of the SR-TE policy programmed at the ingress border router node 1 used to send traffic to the egress peering router node 3 and egress AS behind ASBR node 6:

```
Router1# show segment-routing traffic-eng policy candidate-path name F00 private
```

SR-TE policy database

Color: 10, End-point: 0.0.0.0 ID: 3

```

Name: srte_c_10_ep_0.0.0.0
Status:
 Admin: up Operational: up for 00:10:07 (since Feb 2 12:58:43.554)
Candidate-paths:
 Preference: 100 (configuration) (active)
 Originator: ASN 0 node-address <None> discriminator: 100
 Name: FOO
 Requested BSID: dynamic
 Constraints:
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 ID: 1
 Source: 10.1.1.1
 Stale: no
 Checkpoint flags: 0x00000000
 Steering:
 Client: BGP
 Disabled: no
 Ignore prefix label: yes
 Explicit: segment-list sl-to_3-epe_36 (valid)
 Weight: 1, Metric Type: TE
 IGP area: 2
 SID[0]: 16003 [Prefix-SID: 10.1.1.3, Algorithm: 0]
 SID[1]: 24036 [Adjacency-SID, 2001:db8:20:3:6::3 - 2001:db8:20:3:6::6]
LSPs:
. . .

Attributes:
 Binding SID: 24030
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
 Max Install Standby Candidate Paths: 0
Notification to clients:
 Binding SID: 24030
 Bandwidth : 0 Kbps (0 Kbps)
 State: UP
 Flags: [add] [ipv6_caps] [ignore_prefix_label]
 Metric Type: NONE
 Metric Value: 2147483647
 Admin Distance: 100
 ifhandle: 0x00000170
 Source: 10.1.1.1
 Transition count: 1
 LSPs created count: 1
 Reoptimizations completed count: 1
 Retry Action Flags: 0x00000000, ()
 Last Retry Timestamp: never (0 seconds ago)
 Policy reference: 0x1f81e50

```

The following output depicts the details of the IPv6 BGP global route (2001:db8:abcd::/48) being steered over the binding SID of the previously shown SR-TE policy (24030):

```

Router1# show bgp ipv6 labeled-unicast 2001:db8:abcd::/48 detail

BGP routing table entry for 2001:db8:abcd::/48
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 2003 2003
 Local Label: 81718 (no rewrite);

```

```

Flags: 0x003e1001+0x30010000;
Last Modified: Nov 23 16:59:17.891 for 00:00:03
Paths: (1 available, best #1)
 Advertised IPv6 Unicast paths to update-groups (with more than one peer):
 0.2
 Advertised IPv6 Labeled-unicast paths to update-groups (with more than one peer):
 0.3
 Path #1: Received by speaker 0
 Flags: 0xa480000001060205+0x01, import: 0x020
 Advertised IPv6 Unicast paths to update-groups (with more than one peer):
 0.2
 Advertised IPv6 Labeled-unicast paths to update-groups (with more than one peer):
 0.3
 300, (Received from a RR-client)
 10.1.1.2 C:10 (bsid:24030) (admin 100) (metric 2147483647) from 10.1.1.100 (10.1.1.2),
 if-handle 0x00000170
 Prefix Label not imposed due to SR policy config

```

## Miscellaneous

### SR Policy Liveness Monitoring

SR Policy liveness monitoring allows you to verify end-to-end traffic forwarding over an SR Policy candidate path by periodically sending performance monitoring (PM) packets. The head-end router sends PM packets to the SR policy's endpoint router, which sends them back to the head-end without any control-plane dependency on the endpoint router.

For more information about this feature, see [SR Policy Liveness Monitoring, on page 446](#).

## Programming Non-Active Candidate Paths of an SR Policy

Table 41: Feature History Table

| Feature Name                                           | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Programming Non-Active Candidate Paths of an SR Policy | Release 7.6.1       | <p>By programming non-active candidate paths (CPs) in the forwarding plane, you ensure that if the existing active CP is unavailable, the traffic switches quickly to the new CP, thus minimizing loss of traffic flow.</p> <p>In earlier releases, instantiating a non-active CP to the forwarding plane after the unavailability of the active CP could take a few seconds, resulting in potential loss of traffic flow.</p> <p>This feature introduces the following command:</p> <ul style="list-style-type: none"><li>• <a href="#">max-install-standby-cpaths</a></li></ul> |

An SR Policy is associated with one or more candidate paths (CP). A CP is selected as the active CP when it is valid and it has the highest preference value among all the valid CPs of the SR Policy. By default, only the active CP is programmed in the forwarding plane.

This feature allows the programming of multiple CPs of an SR policy in the forwarding plane. This minimizes traffic loss when a new CP is selected as active.

### Usage Guidelines and Limitations

Observe the following usage guidelines and limitations:

- Up to three non-active CPs can be programmed in the forwarding plane.
- Manually configured CPs are supported. This includes CPs with explicit paths or dynamic (head-end computed or PCE-delegated) paths.
- On-Demand instantiated CPs (ODN) are supported.
- BGP-initiated CPs are supported.
- PCE-initiated CPs via PCEP are not supported. This applies to policies created via CLI or via north-bound HTTP-based API.
- Programming of non-active CPs is not supported with SRv6-TE policies, Per-Flow Policies (PFP), or point-to-multipoint SR policies (Tree-SID)
- PCEP reporting of additional CPs is supported, but the PCEP reporting does not distinguish between active and non-active CPs.



- Programming of non-active CPs can be enabled for all SR policies (global), for a specific policy (local), or ODN template.

If enabled globally and also locally or on ODN template, the local or ODN configuration takes precedence over the global configuration.

- Programming of non-active CPs under global SR-TE and configuring policy path protection of an SR policy is supported. In this case, policy path protection takes precedence.
- Programming of non-active CPs for a specific SR policy and configuring policy path protection of an SR policy is not supported.
- The number of policies supported could be impacted by the number of non-active CPs per policy. Programming non-active CPs in the forwarding plane consumes hardware resources (such as local label and ECMP FEC) when more candidate paths are pre-programmed in forwarding than are actually carrying traffic.
- The active CP will be in programmed state. The remaining CPs will be in standby programmed state.
- We recommend that you create separate PM sessions for active and standby candidate paths to monitor the health of the paths end-to-end.

The recommended PM timers should be different for active and standby PM profiles. The PM timers should be less aggressive for the standby PM profile compared to the active PM profile. See [Configure Performance Measurement, on page 437](#) for information about configuring PM sessions.




---

**Note** PM sessions for BGP-TE policies are not supported. PM profiles can be configured only under configured policies at the head-end.

---

- The protected paths for each CP is programmed in the respective LSPs. The protected paths of active CPs are programmed in the active LSP, and the protected paths of standby CPs are programmed in the standby LSP.
- If a candidate path with higher preference becomes available, the traffic will switch to it in Make-Before-Break (MBB) behavior.

## Configuration

Programming of non-active CPs can be enabled for all SR policies (global), for a specific policy (local), or ODN template. If enabled globally, the local or ODN configuration takes precedence over the global configuration.

### Global SR-TE

Use the **max-install-standby-cpaths** *value* command to configure standby candidate paths for all SR policies, for a specific policy, or for an ODN template. The range for *value* is from 1 to 3. Use **no** **max-install-standby-cpaths** command to return to the default behavior.

The following example shows how to configure standby candidate paths globally:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# max-install-standby-cpaths 2
Router(config-sr-te)#
```

## Running Config

```
segment-routing
traffic-eng
max-install-standby-cpaths 2
```

## Local SR Policy

Use the **max-install-standby-cpaths** *value* command to configure standby candidate paths for a specific policy. The range for *value* is from 0 (disable) to 3.

If programming of non-active CPs is enabled for all SR policies (global), you can disable programming of non-active CPs for a specific policy using the **max-install-standby-cpaths 0** command.

The following example shows how to configure standby candidate paths for a specific SR policy:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy MyBackupPolicy
Router(config-sr-te-policy)# max-install-standby-cpaths 2
Router(config-sr-te-policy)#
```

## Running Config

```
segment-routing
traffic-eng
policy MyBackupPolicy
max-install-standby-cpaths 2
```

## SR ODN

When you create an ODN template, two CPs are created by default (PCE-delegated and head-end computed) with preference 100 and preference 200. You can use the **max-install-standby-cpaths 1** command to program the non-active CP in forwarding. If programming of non-active CPs is enabled for all SR policies (global), you can disable programming of non-active CPs on ODN template using the **max-install-standby-cpaths 0** command.

The following example shows how to configure standby candidate paths for an SR ODN template:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# max-install-standby-cpaths 1
Router(config-sr-te-color)#
```

## Running Config

```
segment-routing
traffic-eng
on-demand color 10
max-install-standby-cpaths 1
```

The following example shows how to enable three standby CPs globally and disable standby CPs on local SR policy and ODN template:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# max-install-standby-cpaths 3
Router(config-sr-te)# policy MyBackupPolicy
Router(config-sr-te-policy)# max-install-standby-cpaths 0
```

```
Router(config-sr-te-policy)# exit
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# max-install-standby-cpaths 0
Router(config-sr-te-color)#
```

## Verification

The following output shows the status of active and backup CPs:

```
Router# show segment-routing traffic-eng policy

SR-TE policy database

Color: 50, End-point: 1.1.1.4
 Name: srte_c_50_ep_1.1.1.4
 Status:
 Admin: up Operational: up for 08:17:32 (since Sep 9 13:16:02.818)
 Candidate-paths:
 Preference: 100 (configuration) (active)
 Name: NCP_STATIC
 Requested BSID: 5000
 PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_100
 PLSP-ID: 2
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Explicit: segment-list WORKING (valid)
 Reverse: segment-list REVERSE_WORKING
 Weight: 1, Metric Type: TE
 24010
 24012
 Preference: 80 (configuration) (standby)
 Name: NCP_STATIC
 Requested BSID: 5000
 PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_80
 PLSP-ID: 3
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Explicit: segment-list STANDBY1 (valid)
 Reverse: segment-list REVERSE_STANDBY1
 Weight: 1, Metric Type: TE
 24018
 24010
 Preference: 60 (configuration) (standby)
 Name: NCP_STATIC
 Requested BSID: 5000
 PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_60
 PLSP-ID: 4
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Explicit: segment-list STANDBY2 (valid)
 Reverse: segment-list REVERSE_STANDBY2
 Weight: 1, Metric Type: TE
 24014
 Preference: 40 (configuration)
 Name: NCP_STATIC
 Requested BSID: 5000
 PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_40
 PLSP-ID: 5
 Protection Type: protected-preferred
```

```

 Maximum SID Depth: 10
 Dynamic (valid)
 Metric Type: TE, Path Accumulated Metric: 10
 24005 [Adjacency-SID, 13.13.13.1 - 13.13.13.4]
Preference: 30 (configuration)
Name: NCP_STATIC
Requested BSID: 5000
PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_30
 PLSP-ID: 1
 Protection Type: protected-preferred
 Maximum SID Depth: 10
Dynamic (pce 1.1.1.4) (valid)
 Metric Type: TE, Path Accumulated Metric: 10
 24015 [Adjacency-SID, 13.13.13.1 - 13.13.13.4]
Preference: 20 (configuration)
Name: NCP_STATIC
Requested BSID: 5000
PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_20
 PLSP-ID: 6
 Protection Type: protected-preferred
 Maximum SID Depth: 10
Explicit: segment-list WORKING2 (valid)
Reverse: segment-list REVERSE_WORKING2
Weight: 1, Metric Type: TE
 24012
 24012
Preference: 10 (configuration)
Name: NCP_STATIC
Requested BSID: 5000
PCC info:
 Symbolic name: cfg_NCP_STATIC_discr_10
 PLSP-ID: 7
 Protection Type: protected-preferred
 Maximum SID Depth: 10
Explicit: segment-list WORKING3 (valid)
Reverse: segment-list REVERSE_WORKING3
Weight: 1, Metric Type: TE
 24010
 24014
 24010
Attributes:
 Binding SID: 5000
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
Max Install Standby CPaths: 2

```

Router# **show segment-routing traffic-eng forwarding policy**

SR-TE Policy Forwarding database

-----

```

Color: 50, End-point: 1.1.1.4
Name: srte_c_50_ep_1.1.1.4
Binding SID: 5000
Active LSP:
 Candidate path:
 Preference: 100 (configuration)
 Name: NCP_STATIC

```

```

Local label: 24021
Segment lists:
 SL[0]:
 Name: WORKING
 Switched Packets/Bytes: 0/0
 Paths:
 Path[0]:
 Outgoing Label: 24012
 Outgoing Interfaces: GigabitEthernet0/0/0/0
 Next Hop: 10.10.10.2
 Switched Packets/Bytes: 0/0
 FRR Pure Backup: No
 ECMP/LFA Backup: No
 Internal Recursive Label: Unlabelled (recursive)
 Label Stack (Top -> Bottom): { 24012 }
Standby LSP(s):
 LSP[0]:
 Candidate path:
 Preference: 80 (configuration)
 Name: NCP_STATIC
 Local label: 24024
 Segment lists:
 SL[0]:
 Name: STANDBY1
 Switched Packets/Bytes: 0/0
 Paths:
 Path[0]:
 Outgoing Label: 24010
 Outgoing Interfaces: GigabitEthernet0/0/0/2
 Next Hop: 12.12.12.3
 Switched Packets/Bytes: 0/0
 FRR Pure Backup: No
 ECMP/LFA Backup: No
 Internal Recursive Label: Unlabelled (recursive)
 Label Stack (Top -> Bottom): { 24010 }
 LSP[1]:
 Candidate path:
 Preference: 60 (configuration)
 Name: NCP_STATIC
 Local label: 24025
 Segment lists:
 SL[0]:
 Name: STANDBY2
 Switched Packets/Bytes: 0/0
 Paths:
 Path[0]:
 Outgoing Label: Pop
 Outgoing Interfaces: GigabitEthernet0/0/0/3
 Next Hop: 13.13.13.4
 Switched Packets/Bytes: 0/0
 FRR Pure Backup: No
 ECMP/LFA Backup: No
 Internal Recursive Label: Unlabelled (recursive)
 Label Stack (Top -> Bottom): { Pop }

Policy Packets/Bytes Switched: 2/136

```

## LDP over Segment Routing Policy

The LDP over Segment Routing Policy feature enables an LDP-targeted adjacency over a Segment Routing (SR) policy between two routers. This feature extends the existing MPLS LDP address family neighbor configuration to specify an SR policy as the targeted end-point.

LDP over SR policy is supported for locally configured SR policies with IPv4 end-points.

For more information about MPLS LDP, see the "Implementing MPLS Label Distribution Protocol" chapter in the *MPLS Configuration Guide*.

For more information about Autoroute, see the *Autoroute Announce for SR-TE* section.



**Note** Before you configure an LDP targeted adjacency over SR policy name, you need to create the SR policy under Segment Routing configuration. The SR policy interface names are created internally based on the color and endpoint of the policy. LDP is non-operational if SR policy name is unknown.

The following functionality applies:

1. Configure the SR policy – LDP receives the associated end-point address from the interface manager (IM) and stores it in the LDP interface database (IDB) for the configured SR policy.
2. Configure the SR policy name under LDP – LDP retrieves the stored end-point address from the IDB and uses it. Use the auto-generated SR policy name assigned by the router when creating an LDP targeted adjacency over an SR policy. Auto-generated SR policy names use the following naming convention: **srte\_c\_color\_val\_ep\_endpoint-address**. For example, **srte\_c\_1000\_ep\_10.1.1.2**

### Configuration Example

```
/* Enter the SR-TE configuration mode and create the SR policy. This example corresponds
to a local SR policy with an explicit path. */
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list sample-sid-list
Router(config-sr-te-sl)# index 10 address ipv4 10.1.1.7
Router(config-sr-te-sl)# index 20 address ipv4 10.1.1.2
Router(config-sr-te-sl)# exit
Router(config-sr-te)# policy sample_policy
Router(config-sr-te-policy)# color 1000 end-point ipv4 10.1.1.2
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list sample-sid-list
Router(config-sr-te-pp-info)# end

/* Configure LDP over an SR policy */
Router(config)# mpls ldp
Router(config-ldp)# address-family ipv4
Router(config-ldp-af)# neighbor sr-policy srte_c_1000_ep_10.1.1.2 targeted
Router(config-ldp-af)#
```



**Note** Do one of the following to configure LDP discovery for targeted hellos:

- Active targeted hellos (SR policy head end):

```
mpls ldp
 interface GigabitEthernet0/0/0/0
 !
 !
```

- Passive targeted hellos (SR policy end-point):

```
mpls ldp
 address-family ipv4
 discovery targeted-hello accept
 !
 !
```

### Running Configuration

```
segment-routing
traffic-eng
 segment-list sample-sid-list
 index 10 address ipv4 10.1.1.7
 index 20 address ipv4 10.1.1.2
 !
 policy sample_policy
 color 1000 end-point ipv4 10.1.1.2
 candidate-paths
 preference 100
 explicit segment-list sample-sid-list
 !
 !
 !
 !
 !
 !

mpls ldp
 address-family ipv4
 neighbor sr-policy srte_c_1000_ep_10.1.1.2 targeted
 discovery targeted-hello accept
 !
 !
```

### Verification

Router# **show mpls ldp interface brief**

| Interface           | VRF Name | Config   | Enabled  | IGP-Auto-Cfg | TE-Mesh-Grp cfg |
|---------------------|----------|----------|----------|--------------|-----------------|
| Te0/3/0/0/3         | default  | Y        | Y        | 0            | N/A             |
| Te0/3/0/0/6         | default  | Y        | Y        | 0            | N/A             |
| Te0/3/0/0/7         | default  | Y        | Y        | 0            | N/A             |
| Te0/3/0/0/8         | default  | N        | N        | 0            | N/A             |
| Te0/3/0/0/9         | default  | N        | N        | 0            | N/A             |
| <b>srte_c_1000_</b> | default  | <b>Y</b> | <b>Y</b> | 0            | N/A             |

Router# **show mpls ldp interface**

```
Interface TenGigE0/3/0/0/3 (0xa000340)
 VRF: 'default' (0x60000000)
```

```

 Enabled via config: LDP interface
Interface TenGigE0/3/0/0/6 (0xa000400)
 VRF: 'default' (0x60000000)
 Enabled via config: LDP interface
Interface TenGigE0/3/0/0/7 (0xa000440)
 VRF: 'default' (0x60000000)
 Enabled via config: LDP interface
Interface TenGigE0/3/0/0/8 (0xa000480)
 VRF: 'default' (0x60000000)
 Disabled:
Interface TenGigE0/3/0/0/9 (0xa0004c0)
 VRF: 'default' (0x60000000)
 Disabled:
Interface srte_c_1000_ep_10.1.1.2 (0x520)
 VRF: 'default' (0x60000000)
 Enabled via config: LDP interface

```

```
Router# show segment-routing traffic-eng policy color 1000
```

```
SR-TE policy database
```

```

Color: 1000, End-point: 10.1.1.2
Name: srte_c_1000_ep_10.1.1.2
Status:
 Admin: up Operational: up for 00:02:00 (since Jul 2 22:39:06.663)
Candidate-paths:
 Preference: 100 (configuration) (active)
 Name: sample_policy
 Requested BSID: dynamic
 PCC info:
 Symbolic name: cfg_sample_policy_discr_100
 PLSP-ID: 17
 Explicit: segment-list sample-sid-list (valid)
 Weight: 1, Metric Type: TE
 16007 [Prefix-SID, 10.1.1.7]
 16002 [Prefix-SID, 10.1.1.2]
Attributes:
 Binding SID: 80011
 Forward Class: 0
 Steering BGP disabled: no
 IPv6 caps enable: yes

```

```
Router# show mpls ldp neighbor 10.1.1.2 detail
```

```

Peer LDP Identifier: 10.1.1.2:0
TCP connection: 10.1.1.2:646 - 10.1.1.6:57473
Graceful Restart: No
Session Holdtime: 180 sec
State: Oper; Msgs sent/rcvd: 421/423; Downstream-Unsolicited
Up time: 05:22:02
LDP Discovery Sources:
 IPv4: (1)
 Targeted Hello (10.1.1.6 -> 10.1.1.2, active/passive)
 IPv6: (0)
Addresses bound to this peer:
 IPv4: (9)
 10.1.1.2 2.2.2.99 10.1.2.2 10.2.3.2
 10.2.4.2 10.2.22.2 10.2.222.2 10.30.110.132
 11.2.9.2
 IPv6: (0)
Peer holdtime: 180 sec; KA interval: 60 sec; Peer state: Estab

```



```

NSR: Disabled
Clients: LDP over SR Policy
Capabilities:
 Sent:
 0x508 (MP: Point-to-Multipoint (P2MP))
 0x509 (MP: Multipoint-to-Multipoint (MP2MP))
 0x50a (MP: Make-Before-Break (MBB))
 0x50b (Typed Wildcard FEC)
 Received:
 0x508 (MP: Point-to-Multipoint (P2MP))
 0x509 (MP: Multipoint-to-Multipoint (MP2MP))
 0x50a (MP: Make-Before-Break (MBB))
 0x50b (Typed Wildcard FEC)

```

## SR-TE MPLS Label Imposition Enhancement

The SR-TE MPLS Label Imposition Enhancement feature increases the maximum label imposition capabilities of the platform.

In previous releases, the platform supported:

- Up to 5 MPLS transport labels when no MPLS service labels are imposed
- Up to 3 MPLS transport labels when MPLS service labels are imposed

With the SR-TE MPLS Label Imposition Enhancement feature, the platform supports the following:

- Up to 12 MPLS transport labels when no MPLS service labels are imposed
- Up to 9 MPLS transport labels when MPLS service labels are imposed

This enhancement is enabled and disabled dynamically, as the label count changes. For example, if a path requires only 3 MPLS transport labels, the MPLS Label Imposition Enhancement feature is not enabled.

You can disable labeled services for SR-TE policies. The label switching database (LSD) needs to know if labeled services are disabled on top of an SR-TE policy to perform proper label stack splitting.

### Disable Labeled Services per Local Policy

Use the **labeled-services disable** command to disable steering for labeled services for a configured policy. This configuration applies per policy.

```

segment-routing
 traffic-eng
 policy policy name
 steering
 labeled-services disable

```

### Disable Labeled Services per ODN color

Use the **labeled-services disable** command to disable steering of labeled-services for on-demand color policies. This configuration applies for a specific ODN color.

```

segment-routing
 traffic-eng
 on-demand color color
 steering
 labeled-services disable

```

### Disable Labeled Services per Policy Type

Use the **labeled-services disable** command to disable steering of labeled services for all policies for the following policy types:

- **all** — all policies
- **local** — all locally configured policies
- **on-demand** — all BGP on-demand color policies
- **bgp-srte** — all controller-initiated BGP SR-TE policies
- **pcep** — all PCE-initiated policies



**Note** You can specify more than one policy type.

```
segment-routing
 traffic-eng
 steering
 labeled-services
 disable {all | local | on-demand | bgp-srte | pcep}
```

### Verification

Use the **show segment-routing traffic-eng policy** command to display SR policy information. The following output shows that steering of labeled services for the on-demand SR policy are disabled.

```
Router# show segment-routing traffic-eng policy color 10
Thu Jul 18 11:35:25.124 PDT

SR-TE policy database

Color: 10, End-point: 10.1.1.8
Name: srte_c_10_ep_10.1.1.8
Status:
 Admin: up Operational: up for 00:00:06 (since Jul 18 11:35:19.350)
Candidate-paths:
 Preference: 1 (configuration) (active)
 Name: test_pol_2
 Requested BSID: dynamic
 Dynamic (valid)
 Metric Type: TE, Path Accumulated Metric: 10
 24004 [Adjacency-SID, 10.1.1.1 - 10.1.1.2]
Attributes:
 Binding SID: 24011
 Forward Class: 0
 Steering labeled-services disabled: yes
 Steering BGP disabled: no
 IPv6 caps enable: yes
```

## Path invalidation drop

Table 42: Feature History Table

| Feature Name           | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Path invalidation drop | Release 7.4.1       | <p>By default, if an SR Policy becomes invalid (for example, if there is no valid candidate path available), traffic falls back to the native SR forwarding path. In some scenarios, a network operator may require that certain traffic be only carried over the path associated with an SR policy and never allow the native SR LSP to be used.</p> <p>This feature allows the SR policy to stay up in the control plane (to prevent prefixes mapped to the SR policy from falling back to the native SR LSP) but drop the traffic sent on the SR policy.</p> |

By default, if an SR Policy becomes invalid, traffic would fall back to the native SR forwarding path.

In some scenarios, a network operator may require that certain traffic be only carried over the path associated with an SR policy and never allow the native SR LSP to be used. The SR-TE path invalidation drop feature is introduced to meet this requirement.

With the path invalidation drop feature enabled, an SR policy that would become invalid (with no existing valid candidate path) is programmed to drop traffic. At the same time, the SR policy remains operationally UP to prevent prefixes steered over this SR policy from falling back to the native SR path.

Forwarding over the SR policy path resumes without dropping traffic, when the SR policy path becomes valid.



**Note** This feature applies to an SR policy that transitions from valid to invalid; it does not apply to an SR policy that has never been declared valid.

### Enable path invalidation drop for manual SR policy

Use the **segment-routing traffic-eng policy *name* steering path-invalidation drop** command to enable the dropping of traffic when all candidate paths of an SR Policy becomes invalid.

```
segment-routing
 traffic-eng
 policy foo
 steering
 path-invalidation drop
```

### Enable path Invalidation Drop for On-Demand SR Policy

Use the **segment-routing traffic-eng on-demand color *color* steering path-invalidation drop** command (where *color* is from 1 to 4294967295) to enable the dropping of traffic when an On-Demand SR Policy becomes invalid.

```
segment-routing
 traffic-eng
 on-demand color 10
 steering
 path-invalidation drop
```

### Enable Path invalidation drop for PCE-initiated SR policy

Use the **segment-routing traffic-eng pcc profile *profile* steering path-invalidation drop** command (where *profile* is from 1 to 65534) to enable the dropping of traffic when a PCE-Initiated SR Policy becomes invalid.

```
segment-routing
 traffic-eng
 pcc
 profile 7
 steering
 path-invalidation drop
```

### SR Policy output with path invalidation drop enabled and active

This output shows a detailed SR policy with path invalidation drop enabled and valid candidate path:

**Router# show segment-routing traffic-eng policy detail**

```
SR-TE policy database

Color: 1, End-point: 10.10.10.1
 Name: srte_c_1_ep_10.10.10.1
 Status:
 Admin: up Operational: up (path-invalidation drop) for 00:30:58 (since Feb 11
09:07:21.549)
 Candidate-paths:
 Preference: 1 (configuration) (inactive)
 Name: test
 Requested BSID: dynamic
 PCC info:
 Symbolic name: cfg_test_discr_1
 PLSP-ID: 1
 Invalidation drop enabled: yes
 Constraints:
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Performance-measurement:
 Reverse-path segment-list:
 Delay-measurement: Disabled
 Liveness-detection: Disabled
 Explicit: segment-list s11 (inactive)
 Last error: unresolved first label (12345)
 Weight: 1, Metric Type: TE (IANA PCEP/IGP: 2/2)
 SID[0]: 12345
 Preference: 0 (SR-TE) (active) (drop)
 Name: srte_c_1_ep_10.10.10.1_discr_0
 Requested BSID: dynamic
 LSPs:
 LSP[0]:
```

```

LSP-ID: 4 policy ID: 1 (active)
State: Programmed
Binding SID: 24003
Attributes:
 Binding SID: 24003
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: no
 Invalidation drop enabled: yes
 Max Install Standby Candidate Paths: 0
 Path Type: SRMPLSv4

```

### Configure path invalidation drop with performance measurement liveness detection

The Path Invalidation Drop feature can work alongside the **invalidation-action down** configuration in the Performance Measurement Liveness Detection feature. The Performance Measurement Liveness Detection feature enables end-to-end SR policy liveness detection for all segment lists of the active and standby candidate paths that are in the forwarding table. When **invalidation-action down** is configured and a candidate path becomes invalid, the candidate path is immediately operationally brought down and becomes invalid.

See [SR Policy Liveness Monitoring, on page 446](#) for information about configuring liveness detection and the invalidation action.

When both **path-invalidation drop** and **performance-measurement liveness-detection invalidation-action down** are enabled, the following behavior is observed:

1. If the PM liveness session goes down, the current active candidate path of the SR policy becomes invalid.
2. SR-TE path re-optimization occurs immediately to find the next valid candidate path in the SR policy.
3. If there is no valid candidate path in the SR policy, the policy remains operationally UP, but is marked with the invalidation-drop state.

## SR-TE Reoptimization Timers

SR-TE path re-optimization occurs when the head-end determines that there is a more optimal path available than the one currently used. For example, in case of a failure along the SR-TE LSP path, the head-end could detect and revert to a more optimal path by triggering re-optimization.

Re-optimization can occur due to the following events:

- The explicit path hops used by the primary SR-TE LSP explicit path are modified
- The head-end determines the currently used path-option are invalid due to either a topology path disconnect, or a missing SID in the SID database that is specified in the explicit-path
- A more favorable path-option (lower index) becomes available

For event-based re-optimization, you can specify various delay timers for path re-optimization. For example, you can specify how long to wait before switching to a reoptimized path

Additionally, you can configure a timer to specify how often to perform reoptimization of policies. You can also trigger an immediate reoptimization for a specific policy or for all policies.

## SR-TE Reoptimization

To trigger an immediate SR-TE reoptimization, use the **segment-routing traffic-eng reoptimization** command in Exec mode:

```
Router# segment-routing traffic-eng reoptimization {all | name policy}
```

Use the **all** option to trigger an immediate reoptimization for all policies. Use the **name policy** option to trigger an immediate reoptimization for a specific policy.

## Configuring SR-TE Reoptimization Timers

Use these commands in SR-TE configuration mode to configure SR-TE reoptimization timers:

- **timers candidate-path cleanup-delay seconds**—Specifies the delay before cleaning up candidate paths, in seconds. The range is from 0 (immediate clean-up) to 86400; the default value is 120
- **timers cleanup-delay seconds**—Specifies the delay before cleaning up previous path, in seconds. The range is from 0 (immediate clean-up) to 300; the default value is 10.
- **timers init-verify-restart seconds**—Specifies the delay for topology convergence after the topology starts populating due to a restart, in seconds. The range is from 10 to 10000; the default is 40.
- **timers init-verify-startup seconds**—Specifies the delay for topology convergence after topology starts populating for due to startup, in seconds. The range is from 10 to 10000; the default is 300
- **timers init-verify-switchover seconds**—Specifies the delay for topology convergence after topology starts populating due to a switchover, in seconds. The range is from 10 to 10000; the default is 60.
- **timers install-delay seconds**—Specifies the delay before switching to a reoptimized path, in seconds. The range is from 0 (immediate installation of new path) to 300; the default is 10.
- **timers periodic-reoptimization seconds**—Specifies how often to perform periodic reoptimization of policies, in seconds. The range is from 0 to 86400; the default is 600.

## Example Configuration

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# timers
Router(config-sr-te-timers)# candidate-path cleanup-delay 600
Router(config-sr-te-timers)# cleanup-delay 60
Router(config-sr-te-timers)# init-verify-restart 120
Router(config-sr-te-timers)# init-verify-startup 600
Router(config-sr-te-timers)# init-verify-switchover 30
Router(config-sr-te-timers)# install-delay 60
Router(config-sr-te-timers)# periodic-reoptimization 3000
```

## Running Config

```
segment-routing
 traffic-eng
 timers
 install-delay 60
 periodic-reoptimization 3000
 cleanup-delay 60
 candidate-path cleanup-delay 600
 init-verify-restart 120
 init-verify-startup 600
 init-verify-switchover 30
```

## Circuit-Style SR-TE Policies

**Table 43: Feature History Table**

| Feature Name                 | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Circuit-Style SR-TE Policies | Release 7.8.1       | <p>This solution allows Segment Routing to meet the requirements of a connection-oriented transport network, which was historically delivered over circuit-switched SONET/SDH networks.</p> <p>Circuit-style SR-TE policies allow a common network infrastructure to be used for both connection-oriented services and classic IP-based transport. This eliminates the need for multiple parallel networks, which greatly reduces both capital expenditures (CapEx) and operating expenditures (OpEx).</p> |

Segment Routing provides an architecture that caters to both connectionless transport (such as IP) as well as connection-oriented transport (such as TDM). IP-centric transport uses the benefits of ECMP and automated/optimum protection from TI-LFA. On the other hand, connection-oriented transport, which was historically delivered over circuit-switched SONET/SDH networks, requires the following:

- End-to-end bidirectional transport that provides congruent forward and reverse paths, predictable latency, and disjointness
- Bandwidth commitment to ensure there is no impact on the SLA due to network load from other services
- Monitoring and maintenance of path integrity with end-to-end 50-msec path protection
- Persistent end-to-end paths regardless of control-plane state


An SR network can satisfy these requirements by leveraging Circuit-Style SR-TE policies (CS-SR policies).

### Properties of Circuit-Style SR Policies

CS-SR policies have the following properties:

- **Guaranteed Latency over Non-ECMP Paths**

Consider the network below with three possible paths from node 1 to node 7. Of the three paths, the best end-to-end delay is provided by the blue path (1 -> 2 -> 3 -> 7). The chosen path is then encoded with Adj-SIDs corresponding to the traversed interfaces to avoid any ECMP, and therefore guarantee the latency over the path.




522677

#### • Control-Plane Independent Persistency

Adjacency SIDs can provide a persistent path independent from control-plane changes (such as IGP neighbor flaps), as well as network events (such as interface additions or interface flaps) and even the presence of IP on an interface. To achieve this, adjacency SIDs can be manually allocated to ensure persistent values, for example after a node reload event. In addition, adjacency SIDs can be programmed as non-protected to avoid any local TI-LFA protection.


With the Adj-SIDs depicted in the figure below, the path from node 1 to node 7 is encoded with the segment list of {24000, 24001, 24000}. By manually allocating the same Adj-SID values for other direction, the path from node 7 to node 1 is encoded with the same segment list of {24000, 24001, 24000}.



522676

#### • Co-Routed Bidirectional Path

Forward and return SR Policies with congruent paths are routed along the same nodes/interfaces.




522675



### • Liveness Monitoring with Path Protection Switching

Bi-directional liveness monitoring on the working and protect paths ensures fast and consistent switchover, while a protect path is pre-programmed over disjoint facilities.



### • Guaranteed Bandwidth


Most services carried over the CS-SR policy are constant-rate traffic streams. Any packet loss due to temporary congestion leads to bit errors at the service layer. Therefore, bandwidth must be managed very tightly and guaranteed to the services mapped to CS-SR policies.

A centralized controller manages the bandwidth reservation. The controller maintains the reserved bandwidth on each link based on the traffic usage:

- Monitors amount of traffic forwarded to each CS-SR policy in the network
- Uses knowledge of the active path used by the policy
- Computes the per-link reservable bandwidth accordingly

A per-hop behavior (as documented in [RFC3246](#) [Expedited Forwarding] or [RFC2597](#) [Assured Forwarding]) ensures that the specified bandwidth is available to CS-SR policies at all times independent of any other traffic.

Bandwidth is reserved on both the working and protect paths.



In addition, you can allocate one MPLS-EXP value for traffic steered over the CS SR-TE policies and use QoS (interface queueing) configuration to isolate the circuit traffic from the rest:

- QoS on headend nodes:
  - Define EXP value associated with CS services
  - Enforce rate limiting and perform EXP marking on service ingress interfaces
- QoS on transit nodes:
  - Classify incoming packets based on EXP value associated with CS services.
  - Enforce guaranteed bandwidth for the classified traffic on egress interfaces using bandwidth queues or priority queue with shaper.

### Components of the Circuit-Style Solution

CS-SR policy paths are computed and maintained by a stateful PCE. The stateful PCE has a centralized view of the network that it can leverage to compute the co-routed bidirectional end-to-end paths and perform bandwidth allocation control, as well as monitor capabilities to ensure SLA compliance for the life of the CS-SR Policy.

- Centralized Controller
  - Computes the path
  - Encodes the path in a list of Adj-SIDs
  - Monitors and controls bandwidth for SLA guarantee
- QoS configuration on every link to isolate guaranteed traffic

### Usage Guidelines and Limitations

Observe the following guidelines and limitations:


- The maximum SID depth (MSD) is 6.
- CS SR policy end-point IP address must be the router-ID of the intended node.

- SR policy path protection is required for both directions.
- SR policy with dynamic path bandwidth constraint is required for both directions and must have the same value for both directions.
- Candidate path (CP) behavior:
  - The working path is associated with the candidate path of the highest preference value.
  - The protect path is associated with the candidate path of the second-highest preference value.
  - The restore path is associated with the candidate path of the third-highest preference value and is configured as "backup ineligible".
  - Candidate paths with the same role in both directions (working, protect, restore) must have the same preference value.
- Bi-directional path behavior:
  - All paths must be configured as co-routed.
  - All paths with the same role in both directions (working, protect, restore) must have the same bi-directional association ID value.
  - The bi-directional association ID value must be globally unique.
- Disjointness constraint:
  - The working and protect paths under the CS SR policy must be configured with a disjointness constraint using the same disjoint association ID and disjointness type.
    - The disjointness association ID for a working and protect path pair in one direction must be globally unique from the corresponding working and protect path pair in the opposite direction.
  - Node and link disjoint constraint types are supported.
  - The disjoint type used in both directions must be the same.
  - The restore path must not be configured with a disjointness constraint.
- Path optimization objectives supported are TE, IGP, and latency.
- The path optimization objective must match across working, protect, and restore paths in both directions.
- Segment type constraint:
  - Working, protect, and restore paths must all be configured with unprotected-only segment type constraint.
  - Working, protect, and restore paths must all be configured with Adj-SID-only segment type constraint.
  - To ensure persistency throughout link failure events, manual adjacency SIDs allocated from the SRLB range should be created on all interfaces used by CS policies.
- Revert/recovery behavior:
  - When both working and protect paths are down, the restore path becomes active.

- The restore path remains active until the working or protect path recovers (partial recovery) and the lock duration timer expires.
- The lock duration timer is configured under the protect and restore CPs.
- The following functionalities are not supported:
  - Affinity constraint
  - Flex-Algo constraint
  - Metric-bounds constraint
  - SR-TE Path Invalidation Drop

### Configure Performance Measurement Liveness Profiles

Performance Measurement (PM) provides proper detection of candidate path liveness and effective path protection. See [SR Policy Liveness Monitoring](#), on page 446.




The following example shows how to create a liveness profile for the working and protect paths.

```
Router_1(config)# performance-measurement
Router_1(config-perf-meas)# liveness-profile name profile-WORKING
Router_1(config-pm-ld-profile)# liveness-detection
Router_1(config-pm-ld-profile-ld)# multiplier 3
Router_1(config-pm-ld-profile-ld)# exit
Router_1(config-pm-ld-profile)# probe
Router_1(config-pm-ld-probe)# tx-interval 30000
Router_1(config-pm-ld-probe)# exit
Router_1(config-pm-ld-profile)# exit
Router_1(config-perf-meas)# liveness-profile name profile-PROTECT
Router_1(config-pm-ld-profile)# liveness-detection
Router_1(config-pm-ld-profile-ld)# multiplier 3
Router_1(config-pm-ld-profile-ld)# exit
```

```
Router_1(config-pm-ls-profile)# probe
Router_1(config-pm-ls-probe)# tx-interval 100000
```

### Configuring CS SR-TE Policy

The following example shows how to configure a circuit-style SR policy from node 1 to node 7 with three candidate paths: working, protect, and restore.



### Create the SR-TE Policy

Configure the CS SR-TE policy.

Use the **bandwidth** *bandwidth* command in SR-TE policy configuration mode to configure the guaranteed reservable bandwidth for the policy. The range for *bandwidth* is from 1 to 4294967295 in kbps.

Use the **path-protection** command in SR-TE policy configuration mode to enable end-to-end path protection.

```
Router_1(config)# segment-routing
Router_1(config-sr)# traffic-eng
Router_1(config-sr-te)# policy cs-srte-to-node7
Router_1(config-sr-te-policy)# bandwidth 10000
Router_1(config-sr-te-policy)# color 10 end-point ipv4 10.1.1.7
Router_1(config-sr-te-policy)# path-protection
Router_1(config-sr-te-path-pref-protection)# exit
Router_1(config-sr-te-policy)#
```

### Enable Liveness Monitoring Under SR Policy

The following example shows how to enable liveness monitoring under SR Policy and associate the working and protect (backup) liveness-profiles.

```
Router_1(config)# segment-routing
Router_1(config-sr)# traffic-eng
Router_1(config-sr-te)# policy cs-srte-to-node7
Router_1(config-sr-te-policy)# performance-measurement
Router_1(config-sr-te-policy-perf-meas)# liveness-detection
Router_1(config-sr-te-policy-live-detect)# liveness-profile name profile-WORKING
```

```

Router_1(config-sr-te-policy-live-detect)# liveness-profile backup name profile-PROTECT
Router_1(config-sr-te-policy-live-detect)# exit
Router_1(config-sr-te-policy-perf-meas)# exit
Router_1(config-sr-te-policy)#

```

### Configure the Working Candidate Path

The working CP has the following characteristics:

- The working path is associated with the candidate path of the highest preference.
- The working CP uses unprotected-only Adj-SIDs in the segment list.
- The working CP is bidirectional and co-routed.
- The working CP in both directions must have the same bi-directional association ID value.
- The disjoint path constraint for the working CP must have the same group ID and disjoint type as the protect CP.

```

Router_1(config)# segment-routing
Router_1(config-sr)# traffic-eng
Router_1(config-sr-te)# policy cs-srte-to-node7
Router_1(config-sr-te-policy)# candidate-paths
Router_1(config-sr-te-policy-path)# preference 100
Router_1(config-sr-te-policy-path-pref)# dynamic
Router_1(config-sr-te-pp-info)# pcep
Router_1(config-sr-te-path-pcep)# exit
Router_1(config-sr-te-pp-info)# metric
Router_1(config-sr-te-path-metric)# type te
Router_1(config-sr-te-path-metric)# exit
Router_1(config-sr-te-pp-info)# exit
Router_1(config-sr-te-policy-path-pref)# constraints
Router_1(config-sr-te-path-pref-const)# segments
Router_1(config-sr-te-path-pref-const-seg)# protection unprotected-only
Router_1(config-sr-te-path-pref-const-seg)# adjacency-sid-only
Router_1(config-sr-te-path-pref-const-seg)# exit
Router_1(config-sr-te-path-pref-const)# disjoint-path group-id 3 type node
Router_1(config-sr-te-path-pref-const)# exit
Router_1(config-sr-te-policy-path-pref)# bidirectional
Router_1(config-sr-te-path-pref-bidir)# co-routed
Router_1(config-sr-te-path-pref-bidir)# association-id 1100
Router_1(config-sr-te-path-pref-bidir)# exit
Router_1(config-sr-te-policy-path-pref)# exit
Router_1(config-sr-te-policy-path)#

```

### Configure the Protect Candidate Path

The protect CP has the following characteristics:

- The protect path is associated with the candidate path of the second-highest preference.
- The protect CP uses unprotected-only Adj-SIDs in the segment list.
- The protect CP is bidirectional and co-routed.
- The protect CP in both directions must have the same bi-directional association ID value.
- The disjoint path constraint for the protect CP must have the same group ID and disjoint type as the working CP.

- When the working path is invalid, the protect path becomes active. After the working path has recovered, the protect path remains active until the default lock duration (300 seconds) expires. You can configure a different lock duration using the **lock duration** *duration* command. The *duration* range is 0 (disabled) to 3000 seconds. If the lock duration is 0 (disabled), then the working path becomes active as soon as it recovers. If *duration* is not specified, the protect path remains active.

```

Router_1(config-sr-te-policy-path)# preference 50
Router_1(config-sr-te-policy-path-pref)# dynamic
Router_1(config-sr-te-pp-info)# pcep
Router_1(config-sr-te-path-pcep)# exit
Router_1(config-sr-te-pp-info)# metric
Router_1(config-sr-te-path-metric)# type te
Router_1(config-sr-te-path-metric)# exit
Router_1(config-sr-te-pp-info)# exit
Router_1(config-sr-te-policy-path-pref)# lock duration 30
Router_1(config-sr-te-policy-path-pref)# constraints
Router_1(config-sr-te-path-pref-const)# segments
Router_1(config-sr-te-path-pref-const-seg)# protection unprotected-only
Router_1(config-sr-te-path-pref-const-seg)# adjacency-sid-only
Router_1(config-sr-te-path-pref-const-seg)# exit
Router_1(config-sr-te-path-pref-const)# disjoint-path group-id 3 type node
Router_1(config-sr-te-path-pref-const)# exit
Router_1(config-sr-te-policy-path-pref)# bidirectional
Router_1(config-sr-te-path-pref-bidir)# co-routed
Router_1(config-sr-te-path-pref-bidir)# association-id 1050
Router_1(config-sr-te-path-pref-bidir)# exit
Router_1(config-sr-te-policy-path-pref)# exit
Router_1(config-sr-te-policy-path)#

```

### Configure the Restore Candidate Path

The restore CP has the following characteristics:

- The restore path is associated with the candidate path of the the third-highest preference.
- The restore CP uses unprotected-only Adj-SIDs in the segment list.
- The restore CP is bidirectional and co-routed.
- The restore CP in both directions must have the same bidirectional association ID value.
- The restore CP must be configured with **backup-ineligible**. This configuration prevents the restore CP from being used as a fast reroute backup. The restore path is not computed until both working and protect paths become unavailable.
- Disjointness constraint is not configured on the restore CP.
- If both working and protect paths are unavailable, the restore path becomes active. After either the working or protect path has recovered, the restore path remains active until the default lock duration (300 seconds) expires. You can configure a different lock duration using the **lock duration** *duration* command. The *duration* range is 0 (disabled) to 3000 seconds. If the lock duration is 0 (disabled), then the working or protect path becomes active as soon as either recovers. If *duration* is not specified, the restore path remains active.

```

Router_1(config-sr-te-policy-path)# preference 10
Router_1(config-sr-te-policy-path-pref)# dynamic
Router_1(config-sr-te-pp-info)# pcep

```

```

Router_1(config-sr-te-path-pcep)# exit
Router_1(config-sr-te-pp-info)# metric
Router_1(config-sr-te-path-metric)# type te
Router_1(config-sr-te-path-metric)# exit
Router_1(config-sr-te-pp-info)# exit
Router_1(config-sr-te-policy-path-pref)# backup-ineligible
Router_1(config-sr-te-policy-path-pref)# lock duration 30
Router_1(config-sr-te-policy-path-pref)# constraints
Router_1(config-sr-te-path-pref-const)# segments
Router_1(config-sr-te-path-pref-const-seg)# protection unprotected-only
Router_1(config-sr-te-path-pref-const-seg)# adjacency-sid-only
Router_1(config-sr-te-path-pref-const-seg)# exit
Router_1(config-sr-te-path-pref-const)# exit
Router_1(config-sr-te-policy-path-pref)# bidirectional
Router_1(config-sr-te-path-pref-bidir)# co-routed
Router_1(config-sr-te-path-pref-bidir)# association-id 1010
Router_1(config-sr-te-path-pref-bidir)# exit
Router_1(config-sr-te-policy-path-pref)# exit
Router_1(config-sr-te-policy-path)#

```

## Running Configuration

```

Router_1# show running-config

. . .
segment-routing
 traffic-eng
 policy cs-srte-to-node7
 bandwidth 10000
 color 10 end-point ipv4 10.1.1.7
 path-protection
 !
 candidate-paths
 preference 10
 dynamic
 pcep
 !
 metric
 type te
 !
 !
 lock
 duration 30
 !
 backup-ineligible
 !
 constraints
 segments
 protection unprotected-only
 adjacency-sid-only
 !
 !
 bidirectional
 co-routed
 association-id 1010
 !
 !
 preference 50
 dynamic
 pcep
 !
 metric

```



```

 type te
 !
 !
 lock
 duration 30
 !
 constraints
 segments
 protection unprotected-only
 adjacency-sid-only
 !
 disjoint-path group-id 3 type node
 !
 bidirectional
 co-routed
 association-id 1050
 !
 !
 preference 100
 dynamic
 pcep
 !
 metric
 type te
 !
 !
 constraints
 segments
 protection unprotected-only
 adjacency-sid-only
 !
 disjoint-path group-id 3 type node
 !
 bidirectional
 co-routed
 association-id 1100
 !
 !
 !
 performance-measurement
 liveness-detection
 liveness-profile backup name profile-PROTECT
 liveness-profile name profile-WORKING
 invalidation-action down
 !
 !
 !
 !
 root
 performance-measurement
 liveness-profile name profile-PROTECT
 liveness-detection
 multiplier 3
 !
 probe
 tx-interval 100000
 !
 !
 liveness-profile name profile-WORKING
 liveness-detection
 multiplier 3
 !
 probe

```

```

 tx-interval 30000
 !
!
!

```

## Verification

Use the **show segment-routing traffic-eng policy detail** command to display the details of the CS SR policy on node 1:

```

Router_1# show segment-routing traffic-eng policy detail

SR-TE policy database

Color: 10, End-point: 10.1.1.7
Name: srte_c_10_ep_10.1.1.7
Status:
 Admin: up Operational: up for 00:02:24 (since Nov 30 08:03:36.588)
Candidate-paths:
 Preference: 100 (configuration) (active)
 Name: cs-srte-to-node7
 Requested BSID: 8000
 PCC info:
 Symbolic name: cfg_cs-srte-to-node7_discr_100
 PLSP-ID: 2
 Constraints:
 Protection Type: unprotected-only
 Maximum SID Depth: 6
 Adjacency SIDs Only: True
 Performance-measurement:
 Reverse-path Label: Not Configured
 Delay-measurement: Disabled
 Liveness-detection: Enabled
 Profile: profile-WORKING
 Invalidation Action: down
 Logging:
 Session State Change: No
 Statistics:
 Session Create : 1
 Session Update : 12
 Session Delete : 4
 Session Up : 8
 Session Down : 3
 Delay Notification: 0
 Session Error : 0
 Dynamic (pce 192.168.0.5) (valid)
 Metric Type: TE, Path Accumulated Metric: 10
 SID[0]: 24001 [Adjacency-SID, 10.10.10.1 - 10.10.10.2]
 Reverse path:
 SID[0]: 24000 [Adjacency-SID, 10.10.10.2 - 10.10.10.1]
 Protection Information:
 Role: WORKING
 Path Lock: Timed
 Lock Duration: 300(s)
 Preference: 50 (configuration) (protect)
 Name: cs-srte-to-node7
 Requested BSID: 8000
 PCC info:
 Symbolic name: cfg_cs-srte-to-node7_discr_50
 PLSP-ID: 1
 Constraints:

```

```

Protection Type: unprotected-only
Maximum SID Depth: 6
Adjacency SIDs Only: True
Performance-measurement:
Reverse-path Label: Not Configured
Delay-measurement: Disabled
Liveness-detection: Enabled
 Profile: profile-PROTECT
 Invalidation Action: down
 Logging:
 Session State Change: No
Statistics:
 Session Create : 0
 Session Update : 9
 Session Delete : 0
 Session Up : 1
 Session Down : 0
 Delay Notification: 0
 Session Error : 0
Dynamic (pce 192.168.0.5) (valid)
 Metric Type: TE, Path Accumulated Metric: 10
 SID[0]: 24002 [Adjacency-SID, 11.11.11.1 - 11.11.11.2]
 Reverse path:
 SID[0]: 24003 [Adjacency-SID, 11.11.11.2 - 11.11.11.1]
Protection Information:
 Role: PROTECT
 Path Lock: Timed
 Lock Duration: 30(s)
Preference: 10 (configuration) (inactive)
Name: cs-srte-to-node7
Requested BSID: 8000
Constraints:
 Protection Type: unprotected-only
 Maximum SID Depth: 6
 Adjacency SIDs Only: True
Performance-measurement:
Reverse-path Label: Not Configured
Delay-measurement: Disabled
Liveness-detection: Enabled
 Profile: working
 Invalidation Action: down
 Logging:
 Session State Change: No
Statistics:
 Session Create : 0
 Session Update : 0
 Session Delete : 0
 Session Up : 0
 Session Down : 0
 Delay Notification: 0
 Session Error : 0
Dynamic (pce) (inactive)
 Metric Type: TE, Path Accumulated Metric: 0
Protection Information:
 Role: RESTORE
 Path Lock: Timed
 Lock Duration: 30(s)
LSPs:
LSP[0]:
LSP-ID: 3 policy ID: 1 (standby)
Local label: 24037
State: Standby programmed state
Performance-measurement:
Reverse-path Label: Not Configured

```

```
Delay-measurement: Disabled
Liveness-detection: Enabled
Profile: profile-WORKING
Invalidation Action: down
Logging:
 Session State Change: No
 Session State: up, for 1d12h (since Nov 30 08:03:37.859)
LSP[1]:
 LSP-ID: 7 policy ID: 1 (active)
 Local label: 24036
 State: Programmed
 Binding SID: 8000
 Performance-measurement:
 Reverse-path Label: Not Configured
 Delay-measurement: Disabled
 Liveness-detection: Enabled
 Profile: profile-WORKING
 Invalidation Action: down
 Logging:
 Session State Change: No
 Session State: up, for 05:42:36 (since Dec 1 15:11:36.203)
Attributes:
 Binding SID: 8000
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Bandwidth Requested: 10000 kbps
 Bandwidth Current: 10000 kbps
 Invalidation drop enabled: no
 Max Install Standby Candidate Paths: 0
```

## SRv6 policy counters POL.CP.SL.INT.E

Policy counters are metrics used in networking to measure and report traffic statistics for specific policies. These counters help network administrators to monitor and manage network performance, capacity planning, and traffic engineering.

Table 44: Feature History Table

| Feature Name                           | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SRv6 policy counters (POL.CP.SL.INT.E) | Release 25.1.1      | <p>The network administrators can now monitor and manage network performance, capacity planning, and traffic engineering by reviewing the policy counters (POL.CP.SL.INT.E) in SRv6-TE.</p> <p>This feature is enabled by default.</p> <p><b>YANG Data Models:</b></p> <p>This Yang model allows operators to disable all the policy counters that are enabled by default.</p> <ul style="list-style-type: none"> <li>• Native model:<br/>Cisco-IOS-XR-accounting-cfg</li> <li>• UM model:<br/>Cisco-IOS-XR-um-cef-accounting-cfg</li> </ul> <p>(see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</p> |

**POL.CP.SL.INT.E** stands for **Per-SR-policy, per Candidate-Path, per-Segment-List, per-interface, egress traffic counter**. This counter in SRv6 is used to measure and report traffic statistics for specific paths and interfaces within a network policy. It provides granular insights into the traffic flow through different segments and interfaces of a network policy, enabling effective network management and optimization.

#### Benefits of SRv6 policy counters POL.CP.SL.INT.E

The benefits of SRv6 policy counters POL.CP.SL.INT.E, are as listed.

- **Network Optimization:** By analyzing the detailed traffic data, network operators can make informed decisions to optimize the network, such as rerouting traffic to avoid congestion.
- **Troubleshooting:** In case of network issues, these counters provide the necessary data to quickly identify and resolve problems.
- **Policy Validation:** Ensuring that the implemented SRv6 policies are functioning as intended and making adjustments as needed based on the counter data.

## Reporting of SR-TE Policies Using BGP- Link State

Table 45: Feature History Table

| Feature Name                                              | Release Information | Feature Description                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reporting of SR-TE Policies Using BGP-Link State for SRv6 | Release 24.1.1      | <p>You can gather the Traffic Engineering Policy information that is locally available in a node and advertise it in BGP-LS for SRv6. The configuration procedures from previous releases, which supported only SR-MPLS, are the same and apply to this feature.</p> |

| Feature Name                                                 | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reporting of SR-TE Policies Using BGP-Link State for SR-MPLS | Release 7.10.1      | <p>BGP- Link State (LS) is a mechanism by which LS and Traffic Engineering (TE) information can be collected from networks and shared with external components (such as, Segment Routing Path Computation Element (SR-PCE) or Crossword Optimization Engine (COE)) using the BGP routing protocol.</p> <p>This feature gathers the Traffic Engineering Policy information that is locally available in a node and advertises it in BGP-LS for SR-MPLS.</p> <p>The operators can now take informed decisions based on the information that is gathered on their network's path computation, reoptimization, service placement, network visualization, and so on.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• <b>distribute link-state</b></li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>• New XPaths for module <code>Cisco-IOS-XR-infra-xtc-agent-cfg.yang</code> (see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> |

This function is achieved using a BGP Network Layer Reachability Information (NLRI) encoding format. BGP-LS consumes structured IGP data (for example, router-id, remote-IP-address of a link, local-IP address of a link, link identifier, and so on). and creates BGP-LS (NLRI) or attributes that BGP or other components like Cisco IOS XR Traffic Controller (XTC) can consume. Current implementation of BGP-LS can report topology using Nodes, Links, and Prefixes.

**Note**

- Starting from Release 7.10.1 this feature supports MPLS.
- Starting from Release 24.1.1 this feature supports SRv6.

Modern Segment Routing (SR) networks often use SR Traffic Engineering (SR-TE) to influence the path that each specific traffic takes over the network. SR-TE tunnels can be provisioned manually on the tunnel head, but often they are calculated and provisioned by the central controller. Often operator of the network wants the ability to force the traffic over specific nodes and links.

Now the operators have the option to collect reports of the SR-TE and Policy information that is locally available in a node and advertise it into BGP-LS updates, which can be used by external components. Refer the [IEFT](#) for examples. This feature is implemented so that the operators have control over their network's path computation, reoptimization, service placement, network visualization, and so on.



**Note** Circuit Style (CS) SR policies are reported but without the CS policies' specific attributes, like the bidirectional constraints, per-hop behavior, and so on.

## Configure Reporting of SRTE Policies using BGP-LS

The reporting of policies to BGP-LS is disabled by default. Configuring the **distribute link-state** under the SR-PCE or SR-TE configuration enables this feature. Once enabled, all the existing SR policies or Candidate Path (CPs) are encoded to BGP-LS. You can disable this feature by removing the configuration and all the SR policies or CPs are withdrawn from BGP which deletes all of the previously encoded SR policies or CPs.



**Note** When SR policies that are reporting in BGP-LS are enabled by the operator, the Head End only reports the Active Candidate Path (CPs) (CPs installed in the forwarding). There are monitoring use cases that require reporting of inactive CPs. The following CLI is needed to report inactive CPs.

For both SR-TE and SR-PCE, you need to use the global CLI to enable the reporting of policies or Circuit Style (CS) to BGP-LS:



- Note**
- Starting from Release 7.10.1 this feature supports MPLS.
  - Starting from Release 24.1.1 this feature supports SRv6.

### Configuration Example

Configure the following command to enable reporting and syncing of SR policies in BGP-LS at the Head End:

```
Router# config
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# distribute link-state
Router(config-sr-te-distribute-ls)# report-candidate-path-inactive
Router(config-sr-te-distribute-ls)# commit
Router(config-sr-te-distribute-ls)# exit
```



**Note** The **report-candidate-path-inactive** command is an optional command to report inactive CPs.

### Running Configuration

This is a sample running configuration which shows that you have configured BGP-LS reporting feature.

```
segment-routing
 traffic-eng
 distribute link-state
 report-candidate-path-inactive
```

```

!
!
!

```

## Verification

Use the **show pce segment-routing traffic-eng policy private** and **show pce distributed-ls events** to verify the configuration.

**Router#show pce segment-routing traffic-eng policy private**

```

PCE's policy database:

PCC Address: 192.168.2.1
 Color: 100, Endpoint: 192.0.2.1
 Name: srte_c_100_ep_192.0.2.1
 Self Pointer: 0x317d7f0
 Active C-path Pointer: 0x317d9b0
 Candidate-paths:
 Symbolic-name: cfg_foo_discr_100 (Active)
 PLSP-ID: 1
 NB-API Notification pending flag: 0
 Self Pointer: 0x317d9b0
 Tunnel Pointer: 0x317d4e0
 Policy Pointer: 0x317d7f0
 Cached distribute LS element:
 Sense: TRUE
 Refcount: 1
 Is on queue: FALSE
 Node identifiers:
 Protocol: 9 router ID: 192.0.2.1
 Policy identifiers:
 Color: 100 Endpoint: 192.0.2.1
 Flags: 0x00
 Candidate path identifiers:
 Originator: 0.0.0.0 protocol: 3 ASN: 0
 Flags: 0x00
 Discriminator: 100
 Candidate path attributes:
 Name: foo
 Policy name: srte_c_100_ep_192.0.2.1
 State:
 Priority: 0 flags: 0x5800 (active, evaluated, valid-sid-list)
 Preference: 100
 BSID:
 Flags: 0x4000 (alloc)
 BSID: 24021
 Specified BSID: 0
 Constraints:
 Bitfield: 0x0020 flags: 0x4000 (protected-only) MT-ID: 0
 Algorithm: 0
 Bandwidth: 0 kbps
 Metric constraints[0]:
 Type: 0 flags: 0x80 (optimization)
 Margin: 0 bound: 0
 Metric constraints[1]:
 Type: 4 flags: 0x10 (bound)
 Margin: 0 bound: 10
 Segment lists:
 Segment list[0]:
 Flags: 0x3800 (computed, verified, first-seg-resolved) MT-ID: 0 algorithm:
0 weight: 0
 Metric[0]:

```



```

Type: 0 flags: 0x10 (value)
Margin: 0 bound: 0 value: 10
Segments:
Segment[0]:
 Bitfield: 0x0000 type: 3 flags: 0x8000 (sid-present)
 SID: 102000
 Descriptor:
 Algorithm: 0
 Local address: 192.0.2.1 [0] remote address: 0.0.0.0 [0]

```

#### Router#show pce distribute-ls events

```

Distribute LS events:

Event history (oldest first):
 Time Event
 Apr 11 01:50:48.985 [UID: 0] SR CP NLRI: node ID: SR RID 192.0.2.1, ls_id 0, asn 0,
bitf 0x00000200 policy ID: endpoint: 192.168.0.2 color: 100 CP ID: originator: config 0.0.0.0
asn: 0 discriminator: 100 (oper: add)
 Apr 11 01:50:50.046 [UID: 1] SR CP NLRI: node ID: SR RID 192.0.2.1, ls_id 0, asn 0,
bitf 0x00000200 policy ID: endpoint: 192.168.0.2 color: 100 CP ID: originator: config 0.0.0.0
asn: 0 discriminator: 100 (oper: add)
RP/0/0/CPU0:rtrX#show pce distribute-ls summary
Tue Apr 11 02:03:36.289 PDT
Distribution enabled: yes
Connected to LS-LIB: yes
Encode queue size: 0
Estimated encoding rate: 17280/s
NLRIs encoded to LS-LIB:
 SR candidate path: 2 added, 0 removed, 0 errored, 0 replaced
Element stats:
 SR candidate path: 1 (watermark: 2)
Throttle timer:
 Running: no

```

Below is the example show output for SRv6.

Router# **show segment-routing traffic-eng policy color 200 private**

```

SR-TE policy database

Color: 200, End-point: 192::2 ID: 2
Name: srte_c_200_ep_192::2
Status:
 Admin: up Operational: up for 00:01:07 (since Mar 6 11:17:42.580)
Candidate-paths:
 Preference: 100 (configuration) (active)
 Originator: ASN 0 node-address <None> discriminator: 100
 Name: bar
 Requested BSID: dynamic
 PCC info:
 Symbolic name: cfg_bar_discr_100
 PLSP-ID: 2
 Is orphan: no
 State timer:
 Running: no
 Constraints:
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 ID: 1
 Source: 192::1
 Stale: no
 Checkpoint flags: 0x00000000

```

```

Path Type: SRV6
Performance-measurement:
 Reverse-path segment-list:
 Delay-measurement: Disabled
 Liveness-detection: Disabled
Dynamic (pce 192.168.0.3) (valid)
Metric Type: IGP, Path Accumulated Metric: 10
IGP area: 0
 SID[0]: fccc:cccl:2::/48 Behavior: uN (PSP/USD) (48)
 Format: f3216
 LBL:32 LNL:16 FL:0 AL:80
 Address: 192::2
SRv6 Information:
 Locator: loc1Algo0
 Binding SID requested: Dynamic
 Binding SID behavior: uB6 (Insert.Red)
Cached distribute LS element:
 Sense: TRUE
 Refcount: 1
 Is on queue: FALSE
 Node identifiers:
 Protocol: 9 router ID: 192::1
 Policy identifiers:
 Color: 200 Endpoint: 192::2
 Flags: 0x80 (endpoint-v6)
 Candidate path identifiers:
 Originator: 0.0.0.0 protocol: 3 ASN: 0
 Flags: 0x00
 Discriminator: 100
 Candidate path attributes:
 Name: bar
 Policy name: srte_c_200_ep_192::2
 State:
 Priority: 0 flags: 0x5A00 (active, evaluated, valid-sid-list, delegated)
 Preference: 100
 SRv6 BSID:
 Flags: 0x8000 (alloc)
 BSID: fccc:cccl:1:e018::
 Specified BSID: ::
 Endpoint:
 Endpoint function: 71 flags: 0x00 algorithm: 0
 Structure:
 Locator block length: 32 locator node length: 16 function length: 16 arguments
length: 0
 Constraints:
 Bitfield: 0x0020 flags: 0xC000 (dataplane-v6, protected) MT-ID: 0
 Algorithm: 0
 Bandwidth: 0 kbps
 Metric constraints[0]:
 Type: 0 flags: 0x80 (optimization)
 Margin: 0 bound: 0
 Metric constraints[1]:
 Type: 4 flags: 0x10 (bound)
 Margin: 0 bound: 10
 Segment lists:
 Segment list[0]:
 Flags: 0xB800 (dataplane-v6, computed, verified, first-seg-resolved) MT-ID:
0 algorithm: 0 weight: 1
 Metric[0]:
 Type: 0 flags: 0x10 (value)
 Margin: 0 bound: 0 value: 10
 Segments:
 Segment[0]:

```

```

 Bitfield: 0x0003 type: 9 flags: 0x8000 (sid-present)
 SID: fccc:ccc1:2::
 Descriptor:
 Algorithm: 0
 Local address: 192::2 [0] remote address: :: [0]
 Endpoint:
 Endpoint function: 48 flags: 0x00 algorithm: 0
 Structure:
 Locator block length: 32 locator node length: 16 function length: 0
arguments length: 80
LSPs:
 LSP[0]:
 LSP-ID: 2 policy ID: 2 (active)
 State: Programmed
 Binding SID: fccc:ccc1:1:e018::
 Install timer:
 Running: no
 Cleanup timer:
 Running: no
 Delete timer:
 Running: no
 Revert timer:
 Running: no
 SM chain:
 Init -> Egress paths
 Egress paths pending -> BSID RW
 BSID rewrite pending -> Success
 Forwarding flags: 0x00000008
 Candidate path ID: 1
 Flags:
 SL-ID:
 Sent/Received/Transferred: 1/1/0
 SLs:
 SL[0]:
 Name: dynamic
 Type: Dynamic PCE
 Checkpoint id: 1
 NH SRV6 SID: fccc:ccc1:2::
 SL ID: 0xa000001
 Flags:
 Normalized Weight: 1
 ENS: 1
 Paths:
 Path[0]:
 Interface version: 1
 Flags:
 Outgoing interface: Gi0/2/0/0
 Weight: 1
 SID stack:
 Total SID count: 0
 Underlay Normalized Weight: 1
 Path[1]:
 Interface version: 1
 Flags:
 Outgoing interface: Gi0/2/0/1
 Weight: 1
 SID stack:
 Total SID count: 0
 Underlay Normalized Weight: 1
 Path[2]:
 Interface version: 1
 Flags:
 Outgoing interface: Gi0/2/0/2
 Weight: 1

```

```

 SID stack:
 Total SID count: 0
 Underlay Normalized Weight: 1
Attributes:
 Binding SID: fccc:cccl:1:e018::
 Forward Class: Not Configured
 Steering labeled-services disabled: no
 Steering BGP disabled: no
 IPv6 caps enable: yes
 Invalidation drop enabled: no
 Max Install Standby Candidate Paths: 0
 Path Type: SRV6
Notification to clients:
 Binding SID: fccc:cccl:1:e018::
 Bandwidth : 0 Kbps (0 Kbps)
 State: UP
 Flags: [add] [ipv6_caps]
 Metric Type: IGP
 Metric Value: 10
 Admin Distance: 30
ifhandle: 0x00000000
Source: 192::1
Transition count: 1
LSPs created count: 1
Reoptimizations completed count: 1
Retry Action Flags: 0x00000000, ()
Last Retry Timestamp: never (0 seconds ago)
Policy reference: 0x1222c10
Event history (oldest first):
 Time Event
 Mar 6 11:17:40.563 POLICY CREATE
 Mar 6 11:17:40.564 (x3) CP PCRPT: 1 hops:
 Mar 6 11:17:41.071 CP PCUPD: CP-ID: 1, path change: true, notify: true, hops:
fccc:cccl:2::
 Mar 6 11:17:41.072 CP PCRPT: 1 hops: fccc:cccl:2::
 Mar 6 11:17:41.072 LSP CREATE: 2, need BSID RW: true, BSID: ::
 Mar 6 11:17:41.072 CP CHANGE: PREF: 100 [PROTO: 30, ORIGIN: 0/<None>, DISC: 100]
 Mar 6 11:17:42.579 SRv6 BSID RW REQ: ID 2
 Mar 6 11:17:42.580 SRv6 BSID RW RES: ID 2 Status: success
 Mar 6 11:17:42.580 LSP PCRPT: 2, hops: fccc:cccl:2::
 Mar 6 11:17:42.580 CP PCRPT REMOVE: 1
 Mar 6 11:17:42.580 IM STATE CHANGE: UNKNOWN to UP, count: 0

```

Router# **show pce segment-routing traffic-eng policy private**

PCE's policy database:

-----

```

PCC Address: 192.168.2.1
Color: 200, Endpoint: 192::2
Name: srte_c_200_ep_192::2
Self Pointer: 0x26cd890
Active C-path Pointer: 0x26cda00
Candidate-paths:
 Symbolic-name: cfg_bar_discr_100 (Active)
 PLSP-ID: 2
 NB-API Notification pending flag: 0
 Self Pointer: 0x26cda00
 Tunnel Pointer: 0x26cd580
 Policy Pointer: 0x26cd890
Cached distribute LS element:
 Sense: TRUE

```

```

Refcount: 1
Is on queue: FALSE
Node identifiers:
 Protocol: 9 router ID: 192::1
Policy identifiers:
 Color: 200 Endpoint: 192::2
 Flags: 0x80 (endpoint-v6)
Candidate path identifiers:
 Originator: 0.0.0.0 protocol: 3 ASN: 0
 Flags: 0x00
 Discriminator: 100
Candidate path attributes:
 Name: bar
 Policy name: srte_c_200_ep_192::2
 State:
 Priority: 0 flags: 0x5A00 (active, evaluated, valid-sid-list, delegated)
 Preference: 100
SRv6 BSID:
 Flags: 0x8000 (alloc)
 BSID: fccc:cccl:1:e018::
 Specified BSID: ::
 Endpoint:
 Endpoint function: 71 flags: 0x00 algorithm: 0
 Structure:
 Locator block length: 32 locator node length: 16 function length: 16 arguments
length: 0
Constraints:
 Bitfield: 0x0020 flags: 0xC000 (dataplane-v6, protected) MT-ID: 0
 Algorithm: 0
 Bandwidth: 0 kbps
 Metric constraints[0]:
 Type: 0 flags: 0x80 (optimization)
 Margin: 0 bound: 0
 Metric constraints[1]:
 Type: 4 flags: 0x10 (bound)
 Margin: 0 bound: 10
Segment lists:
 Segment list[0]:
 Flags: 0xB800 (dataplane-v6, computed, verified, first-seg-resolved) MT-ID:
0 algorithm: 0 weight: 0
 Metric[0]:
 Type: 0 flags: 0x10 (value)
 Margin: 0 bound: 0 value: 10
 Segments:
 Segment[0]:
 Bitfield: 0x0003 type: 9 flags: 0x8000 (sid-present)
 SID: fccc:cccl:2::
 Descriptor:
 Algorithm: 0
 Local address: 192::2 [0] remote address: :: [0]
 Endpoint:
 Endpoint function: 48 flags: 0x00 algorithm: 0
 Structure:
 Locator block length: 32 locator node length: 16 function length: 0
arguments length: 80

Router# show bgp link-state link-state | i \\[SP\\]

*>i [SP] [SR] [I0x0] [N[c100] [b0.0.0.0] [q192.168.0.1] [te192::1]] [C[po0x3] [f0x80] [e192::2] [cl0xc8] [as0] [oa0.0.0.0] [di100]]/792
*>
[SP] [SR] [I0x0] [N[c100] [b0.0.0.0] [q192.168.0.3] [te192::1]] [C[po0x3] [f0x80] [e192::2] [cl0xc8] [as0] [oa0.0.0.0] [di100]]/792

```

```

Router# show bgp link-state link-state
[SP] [SR] [I0x0] [N[c100] [b0.0.0.0] [q192.168.0.1] [te192::1]] [C[p0x3] [f0x80] [e192::2] [c10xc8] [as0] [oa0.0.0.0] [di100]]/792
BGP routing table entry for
[SP] [SR] [I0x0] [N[c100] [b0.0.0.0] [q192.168.0.1] [te192::1]] [C[p0x3] [f0x80] [e192::2] [c10xc8] [as0] [oa0.0.0.0] [di100]]/792
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 128 128
Last Modified: Mar 6 11:17:43.000 for 00:04:09
Last Delayed at: ---
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local
 192.168.2.1 (metric 20) from 192.168.2.1 (192.168.0.1)
 Origin IGP, localpref 100, valid, internal, best, group-best
 Received Path ID 0, Local Path ID 1, version 128
 Link-state:
 SRTE-CP-State: Priority: 0 Flags: 0x5a00 Preference: 100
 SR-Policy-CP-name: bar
 SR-Policy-name: srte_c_200_ep_192::2
 SRTE-CP-Constraints: Flags: 0xc000 Mtid: 0 Algorithm: 0
 SRTE-CP-Constraints-Metric: Type: 0 Flags: 0x80 Margin: 0
 Bound: 0
 SRTE-CP-Constraints-Metric: Type: 4 Flags: 0x10 Margin: 0
 Bound: 10
 SRTE-Segment-List: Flags: 0xb800 Mtid: 0 Algorithm: 0
 Weight: 1
 SRTE-Segment: Segment-Type: 9 Flags: 0x8000 SID: fccc:cccl:2:: Algorithm:
0
 Local-node: 192::2
 SRv6-Endpoint-Fn: 48 Flags: 0x0 Algo: 0
 SRv6-SID-Struct: LBL: 32 LNL: 16 FL: 0 AL: 80
 SRTE-Segment-List-Metric: Type: 0 Flags: 0x10 Margin: 0
 Bound: 0 Value: 10
 SRTE-CP-SRV6-BSID: Flags: 0x8000 BSID: fccc:cccl:1:e018::
 Specified BSID: ::
 SRv6-Endpoint-Fn: 71 Flags: 0x0 Algo: 0
 SRv6-SID-Struct: LBL: 32 LNL: 16 FL: 16 AL: 0


```

## SR-TE Policy Path Protection

Table 46: Feature History Table

| Feature Name                 | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-TE Policy Path Protection | Release 7.4.2       | <p>You can now configure pre-programmed SR-TE policy Working and Protect candidate paths, and provide fast failure detection through SR Policy Liveness Monitoring probes. If there is a liveness failure on the Working candidate path, the headend triggers a switchover to the Protect candidate path.</p> <p>With this release, you can operate IP-centric (with ECMP and TI-LFA) and TDM-centric (with circuits and path protection) services over a common SR network. This eliminates the need for multiple parallel networks and reduces capital expenditures (CapEx) and operating expenditures (OpEx).</p> <p>For this feature, the following commands/keywords are added:</p> <ul style="list-style-type: none"> <li>• <b>policy path-protection</b></li> <li>• <b>policy candidate-paths preference lock duration</b></li> <li>• <b>backup</b> keyword is added to the <b>performance-measurement liveness-detection</b> command.</li> </ul> |

To provide SR policy path protection, headend router and liveness monitoring functions are introduced. The functions are explained with the 1:1 (*one-to-one*) path protection with SR policy liveness monitoring use case for TDM-centric networks. Pointers:






---

**Note** Path protection and local TI-LFA FRR are mutually exclusive functions.

---

- An SR-TE policy is enabled on the headend router. The headend router 1 sends traffic to endpoint router 7. The Working candidate path **Blue** spans routers 1-2-3-7, and the Protect candidate path **Green** spans routers 1-4-5-6-7.
- The headend Router maintains an independent liveness session on each candidate path using loopback measurement mode. After verifying liveness, it pre-programs Working and Protect paths in forwarding.
- The paths are manually configured in explicit segment lists using MPLS labels to ensure that unprotected adjacency SIDs are utilized.
- The headend router sends traffic over the Working candidate path, and detects any liveness failure. When there is a failure, it sends direct switchover notifications to the FIB, and triggers a switchover to the protected path.
- In 1:1 (**one-to-one**) path protection, when the Working candidate path fails, the Protect candidate path sends traffic.




---

**Note** SR-TE policy path protection and SR-TE path invalidation drop inter-working is not supported.

---

### Liveness Monitoring

- SR PM Liveness probes are performed over Working and Protect candidate paths.
- TWAMP Light (RFC 5357) is used for performance measurement and liveness monitoring.
- Separate PM liveness monitoring sessions are created for working and protect candidate-paths.
- Independent PM sessions are created at both endpoints of the SR Policy.
- Loopback measurement-mode (timestamps t1/t4) is used for liveness monitoring. Probe packets are not punted on the responder node. Round-trip delay is computed as (t4 – t1).
- From headend router 1, PM probe query packets are sent with forward and reverse (7->3->2->1) direction paths of the SR Policy's candidate-path in the header of the probe packet. Similarly, PM probe query packets are sent along the Protect path.
- For liveness monitoring:
  - Liveness is declared UP as soon as one probe packet is received back on all segment-lists of the candidate-path.
  - Liveness failure is detected when last N (user-configured value) consecutive probe packets are lost on any segment-list.
  - Fault in the forward and reverse direction of the segment-list (co-routed path) triggers liveness failure notification to SRTE and FIB. FIB triggers protection switchover upon PM notification (running on high priority thread).



## Configuration

- In this example, an SR-TE policy **foo** is created on the headend router and path-protection is enabled for the policy.

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng policy foo
RP/0/RSP0/CPU0:ios(config-sr-te-policy)# color 10 end-point ipv4 192.168.0.3
RP/0/RSP0/CPU0:ios(config-sr-te-policy)# path-protection
RP/0/RSP0/CPU0:ios(config-sr-te-path-pref-protection)#exit
```

- Under **candidate-paths**, the Protect and Working paths are specified through explicit segment lists.
- The Protect path's preference is 50, and it is lower than the Working path preference of 100. The forward (1->4->5->6->7) and reverse (7->6->5->4->1) Protect paths, and the forward (1->2->3->7) and reverse (7->3->2->1) Working paths are enabled as explicit segment lists.
- When the Working path is invalid, the Protect path becomes active. After the Working path has recovered, the Protect path remains active until the default lock duration (of 300 seconds) expires. You can configure a different lock duration using the **lock duration** command.

The duration range is 0 (disabled) to 3000 seconds. If the lock duration is 0 (disabled), then the Working path becomes active as soon as it recovers. If the duration is not specified, the Protect path remains active.

```
RP/0/RSP0/CPU0:ios(config-sr-te-policy)# candidate-paths
RP/0/RSP0/CPU0:ios(config-sr-te-policy-path)#preference 50
RP/0/RSP0/CPU0:ios(config-sr-te-policy-path-pref)#lock duration 30
RP/0/RSP0/CPU0:ios(config-sr-te-policy-path-pref)# explicit segment-list sl-protect-fwd
```

Type **Exit** three times to go to the SR-TE policy configuration mode.

```
RP/0/RSP0/CPU0:ios(config-sr-te-policy)#candidate-paths
RP/0/RSP0/CPU0:ios(config-sr-te-policy-path)#preference 100
RP/0/RSP0/CPU0:ios(config-sr-te-policy-path-pref)# explicit segment-list sl-working-fwd
RP/0/RSP0/CPU0:ios(config-sr-te-pp-info)# commit
```

## Working and Protect Segment Lists Configuration

Configure explicit segment lists for the candidate paths.



**Note** Segment lists must use only unprotected (dynamic or manual) Adjacency SID and BSIDs (as non-first-SID).

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# segment-list sl-working-fwd
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 1 mpls label 24000
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 2 mpls label 24004
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# exit
RP/0/RSP0/CPU0:ios(config-sr-te)# segment-list sl-working-bck
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 1 mpls label 24002
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 2 mpls label 24006
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# exit
RP/0/RSP0/CPU0:ios(config-sr-te)# segment-list sl-protect-fwd
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 1 mpls label 24000
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 2 mpls label 30201
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# exit
RP/0/RSP0/CPU0:ios(config-sr-te)# segment-list sl-protect-bck
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 1 mpls label 24002
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# index 2 mpls label 30201
RP/0/RSP0/CPU0:ios(config-sr-te-sl)# commit
```

## Performance Measurement Configuration For SR-TE Policy

- Enable SR-TE policy specific performance measurement configurations.
- Create a liveness profile for the Working and Protect paths.

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng policy foo
RP/0/RSP0/CPU0:ios(config-sr-te-policy)#performance-measurement
RP/0/RSP0/CPU0:ios(config-sr-te-policy-perf-meas)#liveness-detection
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# liveness-profile backup name
profile-PROTECT
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# liveness-profile name profile-WORKING
```

- The default Invalidation action is Down and it triggers path protection switching. The other action is None, which is enabled here.

```
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)#invalidation-action none
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)#commit
```

## Performance Measurement Global Profile Configuration

- Create a Working candidate path liveness profile.

```
RP/0/RSP0/CPU0:ios(config)#performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)#liveness-profile sr-policy name profile-WORKING
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# probe
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# measurement-mode loopback
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# tx-interval 30000
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# commit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# liveness-detection multiplier 4
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# commit
```

Type **Exit** to access the Performance Measurement config mode.

- Create a Protect candidate path liveness profile.

```
RP/0/RSP0/CPU0:ios(config-perf-meas)# liveness-profile sr-policy name profile-PROTECT
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# probe
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# measurement-mode loopback
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# tx-interval 100000
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# commit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# liveness-detection multiplier 3
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# commit
```

## Verification

Use the **show segment-routing traffic-eng policy candidate-path** command to display Working and Protect candidate-path details.

```
RP/0/RSP0/CPU0:ios# show segment-routing traffic-eng policy candidate-path name foo

SR-TE policy database

Color: 10, End-point: 192.168.0.3s
Name: srte_c_10_ep_192.168.0.3
Status:
 Admin: up Operational: Up for 00:11:55 (since Dec 15 07:02:08.709)
```

```

Candidate-paths:
 Preference: 100 (configuration) (active)
 Name: foo
 Requested BSID: dynamic
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Explicit: segment-list sl-working-fwd (active)
 Weight: 1, Metric Type: TE
 24000
 24004
 Protection Information:
 Role: WORKING
 Path Lock: Timed
 Lock Duration: 300(s)
 Preference: 50 (configuration) (active)
 Name: foo
 Requested BSID: dynamic
 Protection Type: protected-preferred
 Maximum SID Depth: 10
 Explicit: segment-list sl-protect-fwd (active)
 Weight: 1, Metric Type: TE
 24000
 30201
 Protection Information:
 Role: PROTECT
 Path Lock: Timed
 Lock Duration: 30(s)
..

```

## SR Policy Path Computation for IPv6

Table 47: Feature History Table

| Feature Name                        | Release Information | Feature Description                                                                                                                                                           |
|-------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR Policy Path Computation for IPv6 | Release 24.1.1      | We have introduced the capability for SR Policy to support segment lists with IPv6 addresses, which can be either dynamically computed or explicitly set at the SRTE headend. |

You can now use this feature when you want SR Policy to support segment lists with IPv6 addressed.



**Note** It uses the exiting configurations outlined in *Chapter: Configure SR-TE Policies* in *Segment Routing Configuration Guide for Cisco NCS 560 Series Routers*, with the supported features detailed in the Usage Guidelines section that follows.



## CHAPTER 9

# Segment Routing Tree Segment Identifier

Tree Segment Identifier (Tree-SID) is an SDN controller-based approach to build label switched multicast (LSM) Trees for efficient delivery of multicast traffic in an SR domain and without the need for multicast protocol running in the network. With Tree SID, trees are centrally computed and controlled by a path computation element (SR-PCE).

A Replication segment (as specified in IETF draft "[SR Replication segment for Multi-point Service Delivery](#)") is a type of segment which allows a node (Replication node) to replicate packets to a set of other nodes (Downstream nodes) in a Segment Routing Domain.

A Replication segment includes the following:

- Replication SID: The Segment Identifier of a Replication segment. This is an SR-MPLS label (Tree SID label).
- Downstream nodes: Set of nodes in Segment Routing domain to which a packet is replicated by the Replication segment.

A Point-to-Multipoint (P2MP) tree is formed by stitching Replication segments on the Root node, intermediate Replication nodes, and Leaf nodes. This is referred to as an SR P2MP Policy (as specified in IETF draft "[Segment Routing Point-to-Multipoint Policy](#)").

An SR P2MP policy works on existing MPLS data-plane and supports TE capabilities and single/multi routing domains. At each node of the tree, the forwarding state is represented by the same Replication segment (using a global Tree-SID specified from the SRLB range of labels).

An SR P2MP policy request contains the following:

- Policy name
- SID for the P2MP Tree (Tree-SID)
- Address of the root node
- Addresses of the leaf nodes
- Optimization objectives (TE, IGP metric)
- Constraints (affinity)

The SR-PCE is responsible for the following:

1. Learning the network topology - *to be added*

2. Learning the Root and Leaves of a Tree - *describe dynamic and static Tree SIDs (16-17) - Tree SID Policy Types and Behaviors*
3. Computing the Tree
4. Allocating MPLS label for the Tree
5. Signaling Tree forwarding state to the routers
6. Re-optimizing Tree

### Tree SID Policy Types and Behaviors

- Static P2MP Policies—can be configured in the following ways:
  - Tree SID parameters provided via Cisco Crosswork Optimization Engine (COE) UI
  - COE passes the policy configuration to the SR-PCE via REST API (no Tree-SID CLI at PCE). This method allows for SR-PCE High Availability (HA).



#### Note

Refer to the *Traffic Engineering in Crosswork Optimization Engine* chapter in the [Cisco Crosswork Optimization Engine](#) documentation.

- Tree SID parameters configured via Tree-SID CLI at the SR-PCE



#### Caution


With this method, SR-PCE HA is not supported. For this reason, this configuration method is not recommended.

- Dynamic P2MP Policies—can be configured in the following ways:
  - A BGP mVPN is configured in the network (PE nodes) – service configuration via CLI or Cisco NSO
  - As a result, BGP control plane is used for PE auto-discovery and customer multicast signaling.
  - Tree SID parameters are provided by mVPN PEs via PCEP to the PCE. This method allows for SR-PCE High Availability (HA).

### Tree SID Workflow Overview


This sections shows a basic workflow using a static Tree SID policy:

1. User creates a static Tree-SID policy, either via Crosswork Optimization Engine (preferred), or via CLI at the SR-PCE (not recommended).




521090

2. SR-PCE computes the P2MP Tree.




521091

3. SR-PCE instantiates the Tree-SID state at each node in the tree.



521092

4. The Root node encapsulates the multicast traffic, replicates it, and forwards it to the Transit nodes.
5. The Transit nodes replicate the multicast traffic and forward it to the Leaf nodes.
6. The Leaf nodes decapsulate the multicast traffic and forward it to the multicast receivers.




- [Bud Node Support](#), on page 367
- [Configure Static Segment Routing Tree-SID via CLI at SR-PCE](#), on page 368
- [Running Config](#), on page 370
- [Multicast VPN: Dynamic Tree-SID MVPN \(with TI-LFA\)](#), on page 372
- [Multicast VPN IPv6: Dynamic Tree-SID Multicast VPN IPv6](#), on page 387
- [Multicast: Cisco Nonstop Forwarding for Tree-SID](#), on page 395

## Bud Node Support

In a multicast distribution tree, a Bud node is a node that acts as a leaf (egress) node as well as a mid-point (transit) node toward the downstream sub-tree.

In the below multicast distribution tree topology with Root node {A} and Leaf nodes set {B, C, D}, node D is a Bud node. Similarly, if node E is later added to the Leaf set, it would also become a Bud node.





The tree computation algorithm on SR-PCE has been enhanced to detect a Bud node based on knowledge of the Leaf set, and to handle Leaf/Transit node transitions to Bud node. The role of the Bud node is also explicitly signaled in PCEP.

## Configure Static Segment Routing Tree-SID via CLI at SR-PCE



**Caution** With this configuration method, SR-PCE HA is not supported. For this reason, this configuration method is not recommended.

To configure static Segment Routing Tree-SID for Point-to-Multipoint (P2MP) SR policies, complete the following configurations:

1. Configure Path Computation Element Protocol (PCEP) Path Computation Client (PCC) on all nodes involved in the Tree-SID path (root, mid-point, leaf)
2. Configure Affinity Maps on the SR-PCE
3. Configure P2MP SR Policy on SR-PCE
4. Configure Multicast on the Root and Leaf Nodes

### Configure PCEP PCC on All Nodes in Tree-SID Path

Configure all nodes involved in the Tree-SID path (root, mid-point, leaf) as PCEP PCC. For detailed PCEP PCC configuration information, see [Configure the Head-End Router as PCEP PCC, on page 282](#).

### Configure Affinity Maps on the SR-PCE

Use the **affinity bit-map** *COLOR bit-position* command in PCE SR-TE sub-mode to define affinity maps. The bit-position range is from 0 to 255.

```
Router# configure
Router(config)# pce
Router(config-pce)# segment-routing traffic-eng
Router(config-pce-sr-te)# affinity bit-map RED 23
Router(config-pce-sr-te)# affinity bit-map BLUE 24
Router(config-pce-sr-te)# affinity bit-map CROSS 25
Router(config-pce-sr-te)#
```

### Configure P2MP SR Policy on SR-PCE

Configure the end-point name and addresses, Tree-SID label, and constraints for the P2MP policy.

Use the **endpoint-set** *NAME* command in SR-PCE P2MP sub-mode to enter the name of the end-point set and to define the set of end-point addresses.

```
Router(config-pce-sr-te)# p2mp
Router(config-pce-sr-te-p2mp)# endpoint-set BAR
Router(config-pce-p2mp-ep-set)# ipv4 10.1.1.2
Router(config-pce-p2mp-ep-set)# ipv4 10.1.1.3
Router(config-pce-p2mp-ep-set)# ipv4 10.1.1.4
Router(config-pce-p2mp-ep-set)# exit
Router(config-pce-sr-te-p2mp)#
```

Use the **policy** *policy* command to configure the P2MP policy name and enter P2MP Policy sub-mode. Configure the source address, endpoint-set color, Tree-SID label, affinity constraints, and metric type.

```
Router(config-pce-sr-te-p2mp)# policy FOO
Router(config-pce-p2mp-policy)# source ipv4 10.1.1.6
Router(config-pce-p2mp-policy)# color 10 endpoint-set BAR
Router(config-pce-p2mp-policy)# treesid mpls 15200
Router(config-pce-p2mp-policy)# candidate-paths
Router(config-pce-p2mp-policy-path)# constraints
Router(config-pce-p2mp-path-const)# affinity
Router(config-pce-p2mp-path-affinity)# exclude BLUE
Router(config-pce-p2mp-path-affinity)# exit
Router(config-pce-p2mp-path-const)# exit
Router(config-pce-p2mp-policy-path)# preference 100
Router(config-pce-p2mp-policy-path-preference)# dynamic
Router(config-pce-p2mp-path-info)# metric type te
Router(config-pce-p2mp-path-info)# root
Router(config)#
```

### Configure Multicast on the Root and Leaf Nodes

On the root node of the SR P2MP segment, use the **router pim** command to enter Protocol Independent Multicast (PIM) configuration mode to statically steer multicast flows into an SR P2MP policy.



**Note** Enter this configuration only on an SR P2MP segment. Multicast traffic cannot be steered into a P2P policy.

```
Router(config)# router pim
Router(config-pim)# vrf name
Router(config-pim-name)# address-family ipv4
Router(config-pim-name-ipv4)# sr-p2mp-policy FOO
Router(config-pim-name-ipv4-srp2mp)# static-group 235.1.1.5 10.1.1.6
Router(config-pim-name-ipv4-srp2mp)# root
Router(config)#
```

On the root and leaf nodes of the SR P2MP tree, use the **mdt static segment-routing** command to configure the multicast distribution tree (MDT) core as Tree-SID from the multicast VRF configuration submode.

```
Router(config)# multicast-routing
Router(config-mcast)# vrf TEST
Router(config-mcast-TEST)# address-family ipv4
Router(config-mcast-TEST-ipv4)# mdt static segment-routing
```

On the leaf nodes of an SR P2MP segment, use the **static sr-policy p2mp-policy** command to configure the static SR P2MP Policy from the multicast VRF configuration submode to statically decapsulate multicast flows.

```
Router(config)# multicast-routing
Router(config-mcast)# vrf TEST
Router(config-mcast-TEST)# address-family ipv4
Router(config-mcast-TEST-ipv4)# static sr-policy FOO
```

## Running Config

The following example shows how to configure the end point addresses and P2MP SR policy with affinity constraints on SR-PCE.

```
pce
segment-routing
traffic-eng
affinity bit-map
RED 23
BLUE 24
CROSS 25
!
p2mp
endpoint-set BAR
ipv4 10.1.1.2
ipv4 10.1.1.3
ipv4 10.1.1.4
!
policy FOO
source ipv4 10.1.1.6
color 10 endpoint-set BAR
treesid mpls 15200
candidate-paths
preference 100
dynamic
metric
type te
```



# Multicast VPN: Dynamic Tree-SID MVPN (with TI-LFA)

Table 48: Feature History Table

| Feature Name                                       | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multicast VPN: Dynamic Tree-SID MVPN (with TI-LFA) | Release 7.3.1       | <p>With this feature, you can use SR and MVPN for optimally transporting IP VPN multicast traffic over the SP network, using SR-PCE as a controller.</p> <p>With SR's minimal source router configuration requirement, its ability to implement policies with specific optimization objectives and constraints, protect against network failures using TI-LFA FRR mechanism, and use SR-PCE to dynamically generate optimal multicast trees (including when topology changes occur in the multicast tree), the SR-enabled SP network can transport IP multicast traffic efficiently.</p> |

## Prerequisites for Multicast VPN: Tree-SID MVPN With TI-LFA

- The underlay OSPF/IS-IS network is configured, and OSPF/IS-IS adjacency is formed between routers, across the network.
- BGP is configured for the network, and BGP adjacency is formed between routers. BGP MVPN configuration information is provided in this feature document.
- To understand the benefits, know-how, and configuration of SR and SR-TE policies, see About Segment Routing and Configure SR-TE Policies.


## Information About Multicast VPN: Tree-SID MVPN With TI-LFA

Typically, a customer's IP VPN is spread across VPN sites. IP VPN customer traffic is sent from one site to another over a VPN Service Provider (SP) network.

When IP multicast traffic within a (BGP/MPLS) IP VPN is transported over an SP network (say, from **VPN1-Site-A** to **VPN1-Site-B**, as shown in the image), the SP network requires protocols and procedures to optimally transport multicast traffic from a multicast sender in Site-A to multicast receivers in Site-B.

This use case explains how to enable SR multicast for an SP network, and efficiently transport IP VPN multicast traffic (sent from **VPN1-Site-A** and) received at PE router A, through to PE routers D and E, towards receivers in sites **VPN1-Site-B** and **VPN1-Site-C**.

Figure 20: IP VPN Multicast Traffic Flow Over An SP Network



To enable the *Multicast VPN: Tree-SID MVPN With TI-LFA* feature, the following protocols and software applications are used.

**OSPF/IS-IS** - The underlay network is created with OSPF/IS-IS routing protocol, and reachability is established across the network. See *Configure Segment Routing for IS-IS Protocol* or *Configure Segment Routing for OSPF Protocol* chapter for details.

**BGP Multicast VPN (MVPN)** – The PE routers (A, D, and E) are IP VPN end-points for IP multicast traffic arriving at the SP network (at PE router A) and exiting the SP network (at PE routers D and E). So, BGP MVPN is enabled on the PE routers. NSO is used to configure BGP MVPN on the PE routers.

**BGP Auto-Discovery (AD)** - To enable distributed VPN end-point discovery and C-multicast flow mapping and signalling, BGP AD function is configured on the PE routers. A BGP Auto-Discovery route contains multicast router (loopback IP address) and tree identity (segment ID) information. It carries the information in the Provider Multicast Service Interface (PMSI) Tunnel Attribute (PTA).

C-multicast states are signaled using BGP.

**SR** - To transport IP multicast traffic between the VPN end-points (PE routers A, D, and E), Provider (or P-) tunnels are used. In a P-tunnel, the PE devices are the tunnel end-points. P-tunnels can be generated using different technologies (RSVP-TE, P2MP LSPs, PIM trees, mLDP P2MP LSPs, and mLDP MP2MP LSPs). In this use case, Segment Routing (SR) is used for its benefits that were noted earlier.

With SR and SR-PCE, a Tree-SID Point-to-Multipoint (P2MP) segment is used to create P-Tunnels for MVPN. You can specify SR policy optimization objectives (such as *metrics*) and constraints (such as *affinity*) in an SR policy and send it to the SR-PCE controller, so that it can dynamically create SR multicast trees for traffic flow.

**SR-PCE** - This is a controller which, based on the provided SR policy information, computes optimal paths for a multicast tree, and deploys the tree forwarding state on the multicast routers. When a topology change occurs, SR-PCE automatically computes a new, optimal multicast tree, and deploys the new tree forwarding state on the multicast routers.

**TI-LFA** - In SR-TE, Topology-Independent Loop-Free Alternate (TI-LFA) fast reroute (FRR) function is used to reduce link and node failure reaction time. When the primary next-hop (router link) fails, a pre-computed alternate next hop is used to send traffic. TI-LFA FRR is used when transporting IP VPN multicast traffic.

### Overview of Multicast VPN: Tree-SID MVPN With TI-LFA


The following sections provide an overview of Tree-SID MVPN and TI-LFA. The topology remains the same, with PE routers A, D, and E acting as VPN end-points for carrying IP VPN multicast traffic.

#### Tree-SID MVPN Overview


1. For SR, A is designated as the SR head-end router, and D and E are designated as the SR end-points.

For multicast traffic, A is the root of the SR multicast tree, and D and E are leaf routers of the tree. B and C are the other multicast routers. The objective is to send the IP multicast traffic arriving at A to D and E, as needed


**Figure 21: Multicast Tree**



2. A discovers leaf routers' information through BGP MVPN.
3. Path Computation Element Protocol (PCEP) is used for the SR multicast policy communication between A and the SR-PCE server, and communication between PE routers and the SR-PCE server.
4. When the head-end router SR policy is created on A, and PCEP configurations are enabled on the SR-PCE server and all multicast routers, SR-PCE receives the SR policy and leaf router identity information from A.
5. Based on the policy information it receives, including TE objectives and constraints, SR-PCE builds multicast distribution trees in the underlay for efficient VPN traffic delivery.



6. SR-PCE assigns an SID for the SR multicast tree policy, and deploys the multicast tree forwarding state on the multicast routers.
7. When IP multicast traffic is sent from VPN1-SiteA to PE router A, it steers it into the SR policy, and sends it towards D and E, which forward it to multicast traffic receivers in the sites VPN1-SiteB and VPN1-SiteC.
8. When a leaf/multicast router is added or removed, PE router A updates the SR multicast policy and sends it to SR-PCE. SR-PCE computes new multicast routes, and deploys the multicast tree forwarding state information on the multicast routers.




### TI-LFA FRR Overview


High-level TI-LFA FRR function is depicted in these steps:

1. Tree-SID FRR state information.
  - The link from A to B is protected.
  - SID 16002 is the node SID of B.
  - A programs a backup path to B, through C.
2. IP multicast traffic arrives at A which steers the flow onto the tree.
3. A encapsulates and replicates to B, but the link to B is down.






4. A sends the traffic on the backup path, to C.
5. C sends the traffic to B where normal traffic processing resumes.



### SR Multicast Tree Types


This is an overview of the types of SR multicast trees you can configure, depending on your requirement. You can create a full mesh, on-demand, or optimal multicast tree for IP VPN multicast flow in the SP network.

Figure 22: Full Mesh Multicast Tree



1. A assigns Tree-ID 10 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 10) towards SR-PCE.
2. A announces BGP AD Inclusive PMSI (I-PMSI) route with the PTA (A, 10). Inclusive PMSI - Traffic that is multicast by a PE router on an I-PMSI is received by all other PEs in the MVPN. I-PMSIs are generated by Inclusive P-tunnels.
3. A discovers VPN endpoints D and E from their BGP AD Type I-PMSI route messages.
4. A invokes an Add SR multicast leaf router request (for D and E) to SR-PCE.
5. SR-PCE computes and generates the multicast tree forwarding state information on all the routers that are part of the tree.

Figure 23: On-Demand SR Multicast Tree




1. A assigns Tree-ID 20 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 20) towards SR-PCE.
2. A announces BGP AD Selective PMSI (or S-PMSI) route with PTA (A, 20). A sets the leaf-info-required to discover endpoint interest set.

*Selective PMSI* - Traffic multicast by a PE on an S-PMSI is received by some PEs in the MVPN. S-PMSIs are generated by Selective P-tunnels.

3. E has a receiver behind it, and announces a BGP-AD leaf route towards A. A discovers service endpoint E for the on-demand tree.
4. A invokes an Add SR multicast leaf router request (for E) to SR-PCE.
5. SR-PCE computes and generates the multicast tree information for all the routers that are part of the tree.

Figure 24: Optimal Multicast Tree



1. A decides to optimize a flow and assigns Tree-ID 30 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 30) towards SR-PCE.
2. A announces BGP AD I-PMSI route with PTA (A,30). A sets the leaf-info-required to discover endpoint interest set.
3. D has a receiver behind it, and announces a BGP-AD leaf route towards A. A discovers service endpoint D for optimized flow.
4. A invokes an Add SR multicast leaf router request (for D) to SR-PCE.
5. SR-PCE computes and generates the multicast tree information for all the routers that are part of the tree.

## Configurations

**Head End Router Configuration (Router A)** - The following configuration is specific to the head end router.

### Configure TE Constraints and Optimization Parameters

```
Router# configure terminal
Router(config)# segment-routing traffic-engineering
```

An affinity bit-map is created so that it can be applied to a link or interface.

```
Router(config-sr-te)# affinity-map name 10 bit-position 24
Router(config-sr-te)# commit
```

An affinity (or relationship) is created between the SR policy path and the link color so that SR-TE computes a path that includes or excludes links, as specified. The head-end router automatically follows the actions defined in the ODN template (for color 10) upon the arrival of VPN routes with a BGP color extended community that matches color 10.

```
Router(config)# segment-routing traffic-engineering
Router(config-sr-te)# on-demand color 10 dynamic
Router(config-sr-te-color-dyn)# affinity include-all name red
```

```
Router(config-sr-te-color-dyn)# affinity include-any name blue
Router(config-sr-te-color-dyn)# affinity exclude-any name green
Router(config-sr-te-color-dyn)# metric type te
Router(config-sr-te-color-dyn)# commit
```

The SR policy configuration on the head-end router A will be sent to the SR-PCE server, after a connection is established between A and SR-PCE.

## Multicast Router Configuration

### Configure PCEP Client on Multicast Routers

Associate each multicast router as a client of the SR-PCE server. The **pce address ipv4** command specifies the SR-PCE server's IP address.

```
Router# configure terminal
Router(config)# segment-routing traffic-engineering
Router(config-sr-te)# pcc pce address ipv4 3.3.3.3
Router(config-pcc-pce)# commit
```

### SR PCE Server Configuration

#### Configure Label Range for Multicast Trees

Configure the label range to be used for transporting IP multicast traffic in SP network.

```
Router(config)# pce segment-routing traffic-eng p2mp label-range min 30000 max 60000
Router(config)# commit
```

#### Configure FRR

The following configurations enable FRR for all SR multicast (P2MP) trees, including dynamic and static implementations.

The **lfa** keyword enables LFA FRR on the PCE server.

```
Router(config)# pce segment-routing traffic-eng p2mp fast-reroute lfa
Router(config)# commit
```

Alternatively, you can configure FRR for each individual tree using the following configuration. The **lfa** keyword under a specific multicast policy (**tree1** in this example) enables LFA FRR function for the specified SR multicast P2MP tree.

For dynamic trees, L-flag in LSP Attributes PCEP object controls FRR on a tree.

```
Router(config)# pce
Router(config-pce)# address ipv4 192.168.0.5
Router(config-pce)# segment-routing traffic-eng p2mp policy tree1 fast-reroute lfa
Router(config-pce)# commit
```

You can create FRR node sets using the **frr-node-set from ipv4 address** and **frr-node-set to ipv4 address** commands to specify the *from* and *to* paths on a multicast router that requires FRR protection. In this configuration, the PCE server is configured to manage the FRR function for traffic from 192.168.0.3 sent towards 192.168.0.4 and 192.168.0.5.

```
Router(config)# pce
Router(config-pce)# address ipv4 192.168.0.5
Router(config-pce)# segment-routing traffic-eng
Router(config-pce-sr-te)# p2mp
Router(config-pce-sr-te-p2mp)# frr-node-set from ipv4 192.168.0.3
Router(config-pce-sr-te-p2mp)# frr-node-set to ipv4 192.168.0.4
Router(config-pce-sr-te-p2mp)# frr-node-set to ipv4 192.168.0.5
Router(config-pce-sr-te-p2mp)# commit
```

### Disable ECMP load splitting

To disable ECMP load splitting of different trees on the SR-PCE server, configure the **multipath-disable** command.

```
Router(config)# pce segment-routing traffic-eng p2mp multipath-disable
Router(config)# commit
```

### Multicast Routing Configuration On PE Routers

The following MVPN configurations are required for VPN end-points, the 3 PE routers.

#### Configure Default MDT SR P2MP MVPN Profile

In this configuration, an MDT profile of the type *default* is created, and the SR multicast policy with color 10 will be used to send IP multicast traffic, as per the constraints and optimizations of the policy, through the multicast tree.

You can also specify the FRR LFA function with the **mdt default segment-routing mpls fast-reroute lfa** command.

```
Router(config)# multicast-routing vrf cust1
Router(config-mcast-cust1)# address-family ipv4
Router(config-mcast-cust1-ipv4)# mdt default segment-routing mpls color 10
Router(config-mcast-cust1-ipv4)# commit
```

#### Configure Partitioned MDT SR P2MP MVPN Profile

In this configuration, an MDT profile of the type *partitioned* is created, and the SR multicast policy with color 10 will be used to send IP multicast traffic, as per the constraints and optimizations of the policy, through the multicast tree.

You can also specify the FRR LFA function with the **mdt partitioned segment-routing mpls fast-reroute lfa** command.

```
Router(config)# multicast-routing vrf cust1
Router(config-mcast-cust1)# address-family ipv4
Router(config-mcast-cust1-ipv4)# mdt partitioned segment-routing mpls color 10
Router(config-mcast-cust1-ipv4)# commit
```

The following Data MVPN configuration is required at the Ingress PE (router A) where the multicast flows need to be steered onto the *data* MDT for SR multicast traffic flow.

**Note** - Data MDT can be configured for *Default* and *Partitioned* profiles.

#### Configure Data MDT for SR P2MP MVPN

In this configuration, an MDT profile of the type *data* is created, and the SR multicast policy with color 10 will be used to send IP multicast traffic, as per the constraints and optimizations of the policy, through the multicast tree.

- You can enable the FRR LFA function with the **mdt data segment-routing mpls fast-reroute lfa** command. This enables LFA FRR for SR multicast trees created for all data MDT profiles.
- As an alternative to the color keyword, you can specify a route policy in the **route-policy** command, and define the route policy separately (as mentioned in the next configuration).
- The **threshold** command specifies the threshold above which a multicast flow is switched onto the data MDT. The **immediate-switch** keyword enables an immediate switch of a multicast flow to the data MDT, without waiting for threshold limit to be crossed.
- The **customer-route-acl** keyword specifies an ACL to enable specific multicast flows to be put on to the data MDT.

- **color** and **fast-reroute lfa** keywords are mutually exclusive with the **route-policy** configuration. The objective is to apply constraints (through **color**) or FRR (through LFA protection) to either all data MDTs, or apply them selectively per data MDT, using the **set on-demand-color** and **set fast-reroute lfa** options in the route policy (configured in the **mdt data** configuration).

```
Router(config)# multicast-routing vrf cust1
Router(config-mcast-cust1)# address-family ipv4
Router(config-mcast-cust1-ipv4)# mdt data segment-routing mpls 2 color 10
Router(config-mcast-cust1-ipv4)# commit
```

### Route Policy Example

The route policy designates multicast flow-to-SR multicast policy mapping, with different colors.

- With this configuration, IP multicast flows for the 232.0.0.1 multicast group are steered into the SR multicast policy created with the on-demand color 10, while flows for 232.0.0.2 are steered into the policy created with color 20.
- The *data* MDT SR multicast tree created for the 232.0.0.2 multicast group is enabled with FRR LFA protection.
- Route policies can also be used to match other parameters, such as source address.

```
Router(config)# route-policy TSID-DATA
Router(config-rpl)# if destination in (232.0.0.1) then
Router(config-rpl-if)# set on-demand-color 10
Router(config-rpl-if)# pass
Router(config-rpl-if)# elseif destination in (232.0.0.2) then
Router(config-rpl-elseif)# set on-demand-color 20
Router(config-rpl-elseif)# set fast-reroute lfa
Router(config-rpl-elseif)# pass
Router(config-rpl-elseif)# endif
Router(config-rpl)# end-policy
Router(config)# commit
```

### Configure MVPN BGP Auto-Discovery for SR P2MP

The following configuration is required on all PE routers, and is mandatory for *default* MDT, *partitioned* MDT, and *data* MDT.

Configure the BGP Auto-Discovery function for transporting IP multicast traffic.

```
Router(config)# multicast-routing vrf cust1
Router(config-mcast-cust1)# address-family ipv4
Router(config-mcast-cust1-ipv4)# bgp auto-discovery segment-routing
Router(config-mcast-cust1-ipv4-bgp-ad)# commit
```

### Verification

**View MVPN Context Information** - You can view MVPN VRF context information with these commands.

#### View Default MDT Configuration

This command displays SR multicast tree information, including the MDT details (of *default* type, etc), and customer VRF information (route target, route distinguisher, etc).

```
Router# show mvpn vrf vpn1 context

MVPN context information for VRF vpn1 (0x9541cf0)

RD: 1:10 (Valid, IID 0x1), VPN-ID: 0:0
Import Route-targets : 2
```

```

RT:192.168.0.4:0, BGP-AD
RT:192.168.0.4:17, BGP-AD
BGP Auto-Discovery Enabled (I-PMSI added)
SR P2MP Core-tree data:
 MDT Name: TRmdtvpn1, Handle: 0x4150, idb: 0x956fc30
 MTU: 1376, MaxAggr: 255, SW_Int: 30, AN_Int: 60
 RPF-ID: 3, C:0, O:1, D:0, CP:0
 Static Type : - / -
Def MDT ID: 524289 (0x93993f0), added: 1, HLI: 0x80001, Cfg: 1/0
 Part MDT ID: 0 (0x0), added: 0, HLI: 0x00000, Cfg: 0/0
 Ctrl Trees : 0/0/0, Ctrl ID: 0 (0x0), Ctrl HLI: 0x00000

```

### View Partitioned MDT Configuration

This command displays SR multicast tree information, including the MDT details (of *partitioned* type, etc), and customer VRF information (route target, route distinguisher, etc).

```
Router# show mvpn vrf vpn1 context
```

```

MVPN context information for VRF vpn1 (0x9541cf0)

RD: 1:10 (Valid, IID 0x1), VPN-ID: 0:0
Import Route-targets : 2
 RT:192.168.0.4:0, BGP-AD
 RT:192.168.0.4:17, BGP-AD
BGP Auto-Discovery Enabled (I-PMSI added) , MS-PMSI sent
SR P2MP Core-tree data:
 MDT Name: TRmdtvpn1, Handle: 0x4210, idb: 0x956fc30
 MTU: 1376, MaxAggr: 255, SW_Int: 30, AN_Int: 60
 RPF-ID: 1, C:0, O:1, D:0, CP:0
 Static Type : - / -
 Def MDT ID: 0 (0x0), added: 0, HLI: 0x00000, Cfg: 0/0
Part MDT ID: 524292 (0x9399318), added: 1, HLI: 0x80004, Cfg: 1/0
 Ctrl Trees : 0/0/0, Ctrl ID: 0 (0x0), Ctrl HLI: 0x00000

```

### View Partitioned MDT Ingress PE Configuration

This command displays SR multicast tree information on the PE router that receives the multicast traffic on the SP network. The information includes PE router details, MDT details, Tree-SID details, and the specified customer VRF information.

```
Router# show mvpn vrf vpn1 pe
```

```

MVPN Provider Edge Router information

VRF : vpn1

PE Address : 192.168.0.3 (0x9570240)
 RD: 0:0:0 (null), RIB_HLI 0, RPF-ID 13, Remote RPF-ID 0, State: 0, S-PMSI: 2
 PPMP_LABEL: 0, MS_PMSI_HLI: 0x00000, Bidir_PMSI_HLI: 0x00000, MLDP-added: [RD 0, ID 0,
 Bidir ID 0, Remote Bidir ID 0], Counts(SHR/SRC/DM/DEF-MD): 0, 0, 0, 0, Bidir: GRE RP Count
 0, MPLS RP Count ORSVP-TE added: [Leg 0, Ctrl Leg 0, Part tail 0 Def Tail 0, IR added:
 [Def Leg 0, Ctrl Leg 0, Part Leg 0, Part tail 0, Part IR Tail Label 0
Tree-SID Added: [Def/Part Leaf 1, Def Egress 0, Part Egress 0, Ctrl Leaf 0]
 bgp_i_pmsi: 1,0/0 , bgp_ms_pmsi/Leaf-ad: 1/1, bgp_bidir_pmsi: 0, remote_bgp_bidir_pmsi:
 0, PMSIs: I 0x9570378, 0x0, MS 0x94e29d0, Bidir Local: 0x0, Remote: 0x0, BSR/Leaf-ad 0x0/0,
 Autorp-disc/Leaf-ad 0x0/0, Autorp-ann/Leaf-ad 0x0/0
 IIDs: I/6: 0x1/0x0, B/R: 0x0/0x0, MS: 0x1, B/A/A: 0x0/0x0/0x0

 Bidir RPF-ID: 14, Remote Bidir RPF-ID: 0
 I-PMSI: Unknown/None (0x9570378)
 I-PMSI rem: (0x0)
MS-PMSI: Tree-SID [524290, 192.168.0.3] (0x94e29d0)
 Bidir-PMSI: (0x0)
 Remote Bidir-PMSI: (0x0)

```

```

BSR-PMSI: (0x0)
A-Disc-PMSI: (0x0)
A-Ann-PMSI: (0x0)
RIB Dependency List: 0x0
Bidir RIB Dependency List: 0x0
Sources: 0, RPs: 0, Bidir RPs: 0

```

### View Partitioned MDT Egress PE Configuration

This command displays SR multicast tree information on the MVPN egress PE router that sends multicast traffic from the SP network towards multicast receivers in the destination sites. The information includes PE router, Tree-SID, MDT, and the specified customer VRF details.

```
Router# show mvpn vrf vpn1 pe
```

MVPN Provider Edge Router information

```

PE Address : 192.168.0.4 (0x9fa38f8)
RD: 1:10 (valid), RIB HLI 0, RPF-ID 15, Remote RPF-ID 0, State: 1, S-PMSI: 2
PPMP_LABEL: 0, MS_PMSI_HLI: 0x00000, Bidir_PMSI_HLI: 0x00000, MLDP-added: [RD 0, ID 0,
Bidir ID 0, Remote Bidir ID 0], Counts(SHR/SRC/DM/DEF-MD): 1, 1, 0, 0, Bidir: GRE RP Count
0, MPLS RP Count 0RSVP-TE added: [Leg 0, Ctrl Leg 0, Part tail 0 Def Tail 0, IR added:
[Def Leg 0, Ctrl Leg 0, Part Leg 0, Part tail 0, Part IR Tail Label 0
Tree-SID Added: [Def/Part Leaf 0, Def Egress 0, Part Egress 1, Ctrl Leaf 0]
bgp_i_pmsi: 1,0/0 , bgp_ms_pmsi/Leaf-ad: 1/0, bgp_bidir_pmsi: 0, remote_bgp_bidir_pmsi:
0, PMSIs: I 0x9f77388, 0x0, MS 0x9fa2f98, Bidir Local: 0x0, Remote: 0x0, BSR/Leaf-ad 0x0/0,
Autorp-disc/Leaf-ad 0x0/0, Autorp-ann/Leaf-ad 0x0/0
IIDs: I/6: 0x1/0x0, B/R: 0x0/0x0, MS: 0x1, B/A/A: 0x0/0x0/0x0

Bidir RPF-ID: 16, Remote Bidir RPF-ID: 0
I-PMSI: Unknown/None (0x9f77388)
I-PMSI rem: (0x0)
MS-PMSI: Tree-SID [524292, 192.168.0.4] (0x9fa2f98)
Bidir-PMSI: (0x0)
Remote Bidir-PMSI: (0x0)
BSR-PMSI: (0x0)
A-Disc-PMSI: (0x0)
A-Ann-PMSI: (0x0)
RIB Dependency List: 0x9f81370
Bidir RIB Dependency List: 0x0
Sources: 1, RPs: 1, Bidir RPs: 0

```

### View Data MDT Information

The commands in this section displays SR multicast tree information for *data* MDTs. The information includes cache, router-local, and remote MDT information.

#### View Data MDT Cache Information

```
Router# show pim vrf vpn1 mdt cache
```

| Core Source | Cust (Source, Group)    | Core Data        | Expires |
|-------------|-------------------------|------------------|---------|
| 192.168.0.3 | (26.3.233.1, 232.0.0.1) | [tree-id 524292] | never   |
| 192.168.0.4 | (27.3.233.6, 232.0.0.1) | [tree-id 524290] | never   |

Leaf AD: 192.168.0.3

#### View Local MDTs Information

```
Router# show pim vrf vpn1 mdt sr-p2mp local
```

| Tree Identifier            | MDT Source  | Cache Count | DIP | Local Entry | VRF Using | Routes Cache | On-demand Color |
|----------------------------|-------------|-------------|-----|-------------|-----------|--------------|-----------------|
| [tree-id 524290 (0x80002)] | 192.168.0.4 | 1           | N   | Y           | 1         |              | 10              |

Tree-SID Leaf: 192.168.0.3



## View Remote MDTs Information

Router # **show pim vrf vpn1 mdt sr-p2mp remote**

| Tree Identifier            | MDT Source  | Cache Count | DIP | Local Entry | VRF Using | Routes Cache | On-demand Color |
|----------------------------|-------------|-------------|-----|-------------|-----------|--------------|-----------------|
| [tree-id 524290 (0x80002)] | 192.168.0.4 | 1           | N   | N           | 1         |              | 0               |

## View MRIB MPLS Forwarding Information

This command displays labels used for transporting IP multicast traffic, on a specified router.

Router# **show mrib mpls forwarding**

```
LSP information (XTC) :
 LSM-ID: 0x00000, Role: Head, Head LSM-ID: 0x80002
 Incoming Label : (18000)
 Transported Protocol : <unknown>
 Explicit Null : None
 IP lookup : disabled

 Outsegment Info #1 [H/Push, Recursive]:
 OutLabel: 18000, NH: 192.168.0.3, Sel IF: GigabitEthernet0/2/0/0

LSP information (XTC) :
 LSM-ID: 0x00000, Role: Tail, Peek
 RPF-ID: 0x00011, Assoc-TIDs: 0xe0000011/0x0, MDT: TRmdtvpn1
 Incoming Label : 18001
 Transported Protocol : <unknown>
 Explicit Null : None
 IP lookup : enabled

 Outsegment Info #1 [T/Pop]:
 No info.
```

## SR-PCE Show Commands

### View Tree Information On PCE Server

This command displays SR multicast tree information on the SR-PCE server.



**Note** A cleanup process that activates every 30 minutes will delete any inconsistent entries between the PCEs, which might result from network or config changes. Inconsistent entries are expected during that time frame.

Router# **show pce lsp p2mp**

```
Tree: sr_p2mp_root_192.168.0.1_tree_id_524290
Label: 18000 Operational: up Admin: up
Metric Type: TE
Transition count: 3
Uptime: 00:00:03 (since Fri Jan 24 14:57:51 PST 2020)
Source: 192.168.0.1
Destinations: 192.168.0.4
Nodes:
Node[0]: 192.168.0.2 (rtrM)
 Role: Transit
 Hops:
 Incoming: 18000 CC-ID: 4
 Outgoing: 18000 CC-ID: 4 (17.17.17.4) [rtrR]
Node[1]: 192.168.0.1 (rtrL1)
 Role: Ingress
 Hops:
```

```

Incoming: 18000 CC-ID: 5
Outgoing: 18000 CC-ID: 5 (12.12.12.2) [rtrM]
Node[2]: 192.168.0.4 (rtrR)
Role: Egress
Hops:
Incoming: 18000 CC-ID: 6

```

For dynamic SR multicast trees created for MVPN, the **show** command has filters to view root multicast router and Tree-ID information. When the root router is specified, all multicast trees from that root are displayed. When root and Tree-ID are specified, only the specified tree information is displayed.

```
Router# show pce lsp p2mp root ipv4 10.1.1.1 524289
```

```

Tree: sr_p2mp_root_10.1.1.1_tree_id_524289, Root: 10.1.1.1 ID: 524289
Label: 20000 Operational: up Admin: up
PCC: 10.1.1.1

```

```

Local LFA FRR: Disabled
Metric Type: TE
Transition count: 11
Uptime: 00:03:37 (since Mon May 11 12:53:33 PDT 2020)
Destinations: 10.1.1.3, 10.1.1.4, 10.1.1.5
Nodes:
Node[0]: 10.1.1.1 (root1)
Role: Ingress
Hops:
Incoming: 20000 CC-ID: 26
Outgoing: 20000 CC-ID: 26 (192.168.114.4) [mid-4]
Outgoing: 20000 CC-ID: 26 (192.168.112.2) [mid-2]
Node[1]: 10.1.1.4 (mid-4)
Role: Egress
Hops:
Incoming: 20000 CC-ID: 27
Node[2]: 10.1.1.2 (mid-2)
Role: Transit
Hops:
Incoming: 20000 CC-ID: 28
Outgoing: 20000 CC-ID: 28 (192.168.123.3) [leaf-3]
Outgoing: 20000 CC-ID: 28 (192.168.125.5) [leaf-5]
Node[3]: 10.1.1.3 (leaf-3)
Role: Egress
Hops:
Incoming: 20000 CC-ID: 29
Node[4]: 10.1.1.5 (leaf-5)
Role: Egress
Hops:
Incoming: 20000 CC-ID: 30

```

The following output shows that LFA FRR is enabled on the hop from rtrR to rtrM. Unlike typical multicast replication where the address displayed is the remote address on the link to a downstream router, the IP address 192.168.0.3 (displayed with an exclamation mark) is the router-ID of the downstream router rtrM. The output also displays the LFA FRR state for the multicast tree.

```
Router# show pce lsp p2mp
```

```

Tree: sr_p2mp_root_192.168.0.4_tree_id_524290
Label: 18000 Operational: up Admin: up
LFA FRR: Enabled
Metric Type: TE
Transition count: 1
Uptime: 3d19h (since Thu Feb 13 13:43:40 PST 2020)
Source: 192.168.0.4
Destinations: 192.168.0.1, 192.168.0.2
Nodes:

```

```

Node[0]: 192.168.0.3 (rtrM)
 Role: Transit
 Hops:
 Incoming: 18000 CC-ID: 1
 Outgoing: 18000 CC-ID: 1 (12.12.12.1) [rtrL1]
 Outgoing: 18000 CC-ID: 1 (15.15.15.2) [rtrL2]
Node[1]: 192.168.0.4 (rtrR)
 Role: Ingress
 Hops:
 Incoming: 18000 CC-ID: 2
 Outgoing: 18000 CC-ID: 2 (192.168.0.3!) [rtrM]
Node[2]: 192.168.0.1 (rtrL1)
 Role: Egress
 Hops:
 Incoming: 18000 CC-ID: 3
Node[3]: 192.168.0.2 (rtrL2)
 Role: Egress
 Hops:
 Incoming: 18000 CC-ID: 4

```

### Multicast Tree Information on Routers

Router# **show segment-routing traffic-eng p2mp policy**

```

SR-TE P2MP policy database:

! - Replications with Fast Re-route

Policy: sr_p2mp_root_192.168.0.1_tree_id_524290 LSM-ID: 0x2
 Role: Leaf
 Replication:
 Incoming label: 18001 CC-ID: 6

Policy: sr_p2mp_root_192.168.0.4_tree_id_524290 LSM-ID: 0x80002 (PCC-initiated)
 Color: 0
 LFA FRR: Disabled
 Role: Root
 Replication:
 Incoming label: 18000 CC-ID: 2
 Interface: None [192.168.0.3!] Outgoing label: 18000 CC-ID: 2
 Endpoints:
 192.168.0.1, 192.168.0.2

```

For SR multicast policies originated locally on the router (root router of a dynamic MVPN multicast policy) additional policy information is displayed. The information includes color, end points, and whether LFA FRR is requested by the local application. When the SR-PCE server enables LFA FRR on a specific hop, the outgoing information shows the address of the next router with an exclamation mark and None is displayed for the outgoing interface.

For dynamic SR multicast trees created for MVPN, the **show** command has filters for displaying root multicast router and Tree-ID information. When the root router is specified, all multicast trees for that root are displayed. When root and Tree-ID are specified, only the specified tree information is displayed.

Router# **show segment-routing traffic-eng p2mp policy root ipv4 1.1\$**

```

SR-TE P2MP policy database:

! - Replications with Fast Re-route, * - Stale dynamic policies/endpoints

Policy: sr_p2mp_root_10.1.1.1_tree_id_524289 LSM-ID: 0x691
 Root: 10.1.1.1, ID: 524289
 Role: Transit
 Replication:
 Incoming label: 20000 CC-ID: 28

```

```

Interface: Bundle-Ether23 [192.168.123.3] Outgoing label: 20000 CC-ID: 28
Interface: Bundle-Ether25 [192.168.125.5] Outgoing label: 20000 CC-ID: 28

Policy: sr_p2mp_root_10.1.1.1_tree_id_524290 LSM-ID: 0x692
Root: 10.1.1.1, ID: 524290
Role: Transit
Replication:
 Incoming label: 19999 CC-ID: 28
 Interface: Bundle-Ether23 [192.168.123.3] Outgoing label: 19999 CC-ID: 28
 Interface: Bundle-Ether25 [192.168.125.5] Outgoing label: 19999 CC-ID: 28

```

## Multicast VPN IPv6: Dynamic Tree-SID Multicast VPN IPv6

**Table 49: Feature History Table**

| Feature Name                                             | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multicast VPN:<br>Dynamic Tree-SID<br>Multicast VPN IPv6 | Release<br>7.10.1   | This feature allows Dynamic Tree Segment Identifier (Tree-SID) deployment where IPv6 Multicast payload is used for optimally transporting IP VPN multicast traffic over the provider network, using SR-PCE as a controller. This implementation supports IPv6 only for the Dynamic Tree-SID. Currently, the Static Tree-SID supports IPV4 payloads only, not the IPv6 payloads. |


### Overview of Multicast VPN: Tree-SID Multicast VPN

Typically, a customer's IP VPN is spread across VPN sites. IP VPN customer traffic is sent from one site to another over a VPN Service Provider (SP) network.

When IP Multicast traffic within a (BGP/MPLS) IP VPN is transported over a provider network (say, from **VPN1-Site-A** to **VPN1-Site-B**, as shown in the image), the provider network requires protocols and procedures to optimally transport multicast traffic from a multicast sender in Site-A to multicast receivers in Site-B.

This use case explains how to enable SR multicast for a provider network, and efficiently transport IP VPN multicast traffic (sent from **VPN1-Site-A** and) received at PE router A, through to PE routers D and E, toward receivers in sites **VPN1-Site-B** and **VPN1-Site-C**.

Figure 25: IP VPN Multicast Traffic Flow Over A Provider Network



To enable the *Multicast VPN: Tree-SID multicast VPN* feature, the following protocols and software applications are used:

- **OSPF/IS-IS** - The underlay network is created with OSPF/IS-IS routing protocol, and reachability is established across the network. See *Configure Segment Routing for IS-IS Protocol* or *Configure Segment Routing for OSPF Protocol* chapter for details, within this Guide.
- **BGP Multicast VPN (multicast VPN)** – The PE routers (A, D, and E) are IP VPN endpoints for IP Multicast traffic arriving at the provider network (at PE router A) and exiting the provider network (at PE routers D and E). So, BGP multicast VPN is enabled on the PE routers. NSO is used to configure BGP multicast VPN on the PE routers. See, *Configure Segment Routing for BGP* chapter for details, within this guide
- **BGP Auto-Discovery (AD)** - To enable distributed VPN endpoint discovery and C-multicast flow mapping and signaling, BGP AD function is configured on the PE routers. A BGP Auto-Discovery route contains multicast router (loopback IP address) and tree identity (segment ID) information. It carries the

information in the Provider Multicast Service Interface (PMSI) Tunnel Attribute (PTA). See, *Configure Segment Routing for BGP* chapter for details, within this guide

- **C-multicast states** are signaled using BGP. See, *Configure Segment Routing for BGP* chapter for details, within this guide
- **SR** - To transport IP Multicast traffic between the VPN endpoints (PE routers A, D, and E), Provider (or P-) tunnels are used. In a P-tunnel, the PE devices are the tunnel endpoints. P-tunnels can be generated using different technologies (RSVP-TE, point-to-multipoint LSPs, PIM trees, mLDP point-to-multipoint LSPs, and mLDP MP2MP LSPs). In this use case, Segment Routing (SR) is used for its benefits that were noted earlier.
- With SR and SR-PCE, a Tree-SID point-to-multipoint (P2MP) segment is used to create P-Tunnels for multicast VPN. You can specify SR policy optimization objectives (such as *metrics*) and constraints (such as *affinity*) in an SR policy and send it to the SR-PCE controller, so that it can dynamically create SR multicast trees for traffic flow.
- **SR-PCE** - This is a controller which, based on the provided SR policy information, computes optimal paths for a multicast tree, and deploys the tree forwarding state on the multicast routers. When a topology change occurs, SR-PCE automatically computes a new, optimal multicast tree, and deploys the new tree forwarding state on the multicast routers.


### Tree-SID multicast VPN

The topology remains the same, with PE routers A, D, and E acting as VPN endpoints for carrying IP VPN multicast traffic.

1. For SR, A is designated as the SR headend router, and D and E are designated as the SR endpoints.


For multicast traffic, A is the root of the SR multicast tree, and D and E are leaf routers of the tree. B and C are the other multicast routers. The objective is to send the IP Multicast traffic arriving at A to D and E, as needed.

**Figure 26: Multicast Tree**




2. A discovers leaf routers' information through BGP multicast VPN.
3. Path Computation Element Protocol (PCEP) is used for the SR multicast policy communication between A and the SR-PCE server, and communication between PE routers and the SR-PCE server.

4. When the headend router SR policy is created on A, and PCEP configurations are enabled on the SR-PCE server and all multicast routers, SR-PCE receives the SR policy and leaf router identity information from A.
5. Based on the policy information it receives, including traffic engineering objectives and constraints, SR-PCE builds multicast distribution trees in the underlay for efficient VPN traffic delivery.



6. SR-PCE assigns an SID for the SR multicast tree policy, and deploys the multicast tree forwarding state on the multicast routers.
7. When IP Multicast traffic is sent from VPN1-SiteA to PE router A, it steers it into the SR policy, and sends it toward D and E, which forward it to multicast traffic receivers in the sites VPN1-SiteB and VPN1-SiteC.
8. When a leaf or multicast router is added or removed, PE router A updates the SR multicast policy and sends it to SR-PCE. SR-PCE computes new multicast routes, and deploys the multicast tree forwarding state information on the multicast routers.




### SR Multicast Tree Types

This is an overview of the types of SR multicast trees that you can configure, depending on your requirement. You can create the following tree types for IP VPN multicast flow in the provider network:

- **Full Mesh Multicast Tree**




Figure 27: Full Mesh Multicast Tree



1. A assigns Tree-ID 10 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 10) toward SR-PCE.
2. A announces BGP AD Inclusive PMSI (I-PMSI) route with the PTA (A, 10). Inclusive PMSI - Traffic that is multicast by a PE router on an I-PMSI is received by all other PEs in the multicast VPN. I-PMSIs are generated by Inclusive P-tunnels.
3. A discovers VPN endpoints D and E from their BGP AD Type I-PMSI route messages.
4. A invokes an Add SR multicast leaf router request (for D and E) to SR-PCE.
5. SR-PCE computes and generates the multicast tree forwarding state information on all the routers that are part of the tree.

#### • On-Demand SR Multicast Tree


Figure 28: On-Demand SR Multicast Tree



1. A assigns Tree-ID 20 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 20) toward SR-PCE.
2. A announces BGP AD Selective PMSI (or S-PMSI) route with PTA (A, 20). A sets the leaf-info-required to discover endpoint interest set.  
*Selective PMSI* - Traffic multicast by a PE on an S-PMSI is received by some PEs in the multicast VPN. S-PMSIs are generated by Selective P-tunnels.
3. E has a receiver behind it, and announces a BGP-AD leaf route toward A. A discovers service endpoint E for the on-demand tree.
4. A invokes an Add SR multicast leaf router request (for E) to SR-PCE.
5. SR-PCE computes and generates the multicast tree information for all the routers that are part of the tree.

#### • Optimal Multicast Tree

Figure 29: Optimal Multicast Tree



1. A decides to optimize a flow and assigns Tree-ID 30 and invokes a Create an SR multicast tree request by sending the multicast router and tree ID information (A, 30) toward SR-PCE.
2. A announces BGP AD I-PMSP route with PTA (A, 30). A sets the leaf-info-required to discover endpoint interest set.
3. D has a receiver behind it, and announces a BGP-AD leaf route toward A. A discovers service endpoint D for optimized flow.
4. A invokes an Add SR multicast leaf router request (for D) to SR-PCE.
5. SR-PCE computes and generates the multicast tree information for all the routers that are part of the tree.

# Multicast: Cisco Nonstop Forwarding for Tree-SID

Table 50: Feature History Table

| Feature Name                                     | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multicast: Cisco Nonstop Forwarding for Tree-SID | Release 7.10.1      | <p>Starting from this release, Multicast Nonstop Forwarding supports Tree-SID (Tree Segment Identifier). This ensures that traffic forwarding continues without interruptions whenever the active RSP fails over to the standby RSP.</p> <p>This feature prevents hardware or software failures on the control plane from disrupting the forwarding of existing packet flows through the router for Tree-SID. Thus, ensuring improved network availability, network stability, preventing routing flaps, and no loss of user sessions while the routing protocol information is being restored.</p> <p>The feature modifies the <a href="#">show mrib nsf private</a> command.</p> |



**Note** This section captures only the Cisco Nonstop Forwarding feature in relation with Tree-SID. For more information on the Cisco Nonstop Forwarding feature, see [Multicast Nonstop Forwarding](#).

Multicast now supports hitless Route Processor Fail Over (RPFO). During RPFO, the software deletes IP routes from the Static Tree-SID profile in the headend router. The Dynamic Tree-SID does not have this issue, because in this case, the BGP advertises the states that supports Nonstop Routing (NSR). To overcome this problem for static Tree-SID, there are checkpoints to check the feature in Protocol Independent Multicast (PIM). On switchover, the checkpoint reads to check if the feature is there or not and push Protocol Independent Multicast (PIM) to Cisco Nonstop Forwarding state.

## Verification Steps

The **show mrib nsf private** command is enhanced to display the XTC info as well.

```
Router#show mrib nsf private
Mon Jul 31 13:27:05.056 UTC
IP MRIB Non-Stop Forwarding Status:
Multicast routing state: Normal
 NSF Lifetime: 00:03:00
 Respawn Count: 6
 Last NSF On triggered: Tue Jul 25 13:20:49 2023, 6d00h
 Last NSF Off triggered: Tue Jul 25 13:22:49 2023, 6d00h
 Last NSF ICD Notification sent: Tue Jul 25 13:22:49 2023, 6d00h
 Last Remote NSF On triggered: Tue Jul 25 13:10:18 2023, 6d00h
 Last Remote NSF Off triggered: Tue Jul 25 13:10:27 2023, 6d00h
 Last Label TE NSF On triggered: Tue Jul 25 13:10:18 2023, 6d00h
 Last Label TE NSF Off triggered: Tue Jul 25 13:10:27 2023, 6d00h
 Last Label mLDP NSF On triggered: Tue Jul 25 13:10:18 2023, 6d00h
 Last Label mLDP NSF Off triggered: Tue Jul 25 13:10:27 2023, 6d00h
 Last Label PIM NSF On triggered: Tue Jul 25 13:20:49 2023, 6d00h
 Last Label PIM NSF Off triggered: Tue Jul 25 13:22:49 2023, 6d00h
 Last Label PIM6 NSF On triggered: Tue Jul 25 13:31:22 2023, 5d23h
```

Last Label PIM6 NSF Off triggered: Tue Jul 25 13:33:22 2023, 5d23h  
Last Label XTC NSF On triggered: Tue Jul 25 13:41:51 2023, 5d23h  
Last Label XTC NSF Off triggered: Tue Jul 25 13:41:52 2023, 5d23h

IP NSF :- Active: N, Assume N

MRIB connect timer: Inactive

NSF statistics:

Enabled Cnt - 4, Disabled Cnt - 4

Last Enabled: 6d00h, Last Disabled: 6d00h

Multicast COFO routing state: Normal

Current LMRIB clients: LDP RSVP\_TE PIM PIM6 XTC

LMRIB NSF clients: LDP RSVP\_TE PIM PIM6 XTC

Converged LMRIB clients: LDP RSVP\_TE PIM PIM6 XTC



## CHAPTER 10

# Configure Segment Routing Path Computation Element

The Segment Routing Path Computation Element (SR-PCE) provides stateful PCE functionality by extending the existing IOS-XR PCEP functionality with additional capabilities. SR-PCE is supported on the MPLS data plane and IPv4 control plane.

**Table 51: Feature History Table**

| Feature Name                         | Release Information | Feature Description                                                                                                        |
|--------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------|
| SR-PCE: Single PCE scale enhancement | Release 7.5.1       | With this feature, support for a single PCE is enhanced to 50000 nodes, 100000 LSPs, 500000 links, and 2000 PCEP sessions. |

- [About SR-PCE, on page 397](#)
- [Usage Guidelines and Limitations, on page 399](#)
- [Configure SR-PCE, on page 399](#)
- [PCE override of SR-TE policies, on page 403](#)
- [PCE-Initiated SR Policies, on page 412](#)
- [SR-PCE Flexible Algorithm Multi-Domain Path Computation, on page 413](#)
- [Exclude Network Resources during Path Computation over SR-TE Policies, on page 418](#)
- [Configure the Shortest Path for Disjoint Candidate Paths, on page 423](#)
- [Enable Strict Disjointness for SR-TE policies in the PCE, on page 425](#)
- [ACL Support for PCEP Connection, on page 427](#)
- [Inter-Domain Path Computation Using Redistributed SID, on page 427](#)
- [PCE Support for MPLS-TE LSPs, on page 430](#)
- [Configuring the North-Bound API on SR-PCE, on page 433](#)

## About SR-PCE

The path computation element protocol (PCEP) describes a set of procedures by which a path computation client (PCC) can report and delegate control of head-end label switched paths (LSPs) sourced from the PCC to a PCE peer. The PCE can request the PCC to update and modify parameters of LSPs it controls. The stateful

model also enables a PCC to allow the PCE to initiate computations allowing the PCE to perform network-wide orchestration.

SR-PCE learns topology information by way of IGP (OSPF or IS-IS) or through BGP Link-State (BGP-LS).

SR-PCE is capable of computing paths using the following methods:

- TE metric—SR-PCE uses the TE metric in its path calculations to optimize cumulative TE metric.
- IGP metric—SR-PCE uses the IGP metric in its path calculations to optimize reachability.
- LSP Disjointness—SR-PCE uses the path computation algorithms to compute a pair of disjoint LSPs. The disjoint paths can originate from the same head-end or different head-ends. Disjoint level refers to the type of resources that should not be shared by the two computed paths. SR-PCE supports the following disjoint path computations:
  - Link – Specifies that links are not shared on the computed paths.
  - Node – Specifies that nodes are not shared on the computed paths.
  - SRLG – Specifies that links with the same SRLG value are not shared on the computed paths.
  - SRLG-node – Specifies that SRLG and nodes are not shared on the computed paths.

When the first request is received with a given disjoint-group ID, the first LSP is computed, encoding the shortest path from the first source to the first destination. When the second LSP request is received with the same disjoint-group ID, information received in both requests is used to compute two disjoint paths: one path from the first source to the first destination, and another path from the second source to the second destination. Both paths are computed at the same time.

### TCP Authentication Option

TCP Message Digest 5 (MD5) authentication has been used for authenticating PCEP (TCP) sessions by using a clear text or encrypted password. This feature introduces support for TCP Authentication Option (TCP-AO), which replaces the TCP MD5 option.

TCP-AO uses Message Authentication Codes (MACs), which provides the following:

- Protection against replays for long-lived TCP connections
- More details on the security association with TCP connections than TCP MD5
- A larger set of MACs with minimal system and operational changes

TCP-AO is compatible with Master Key Tuple (MKT) configuration. TCP-AO also protects connections when using the same MKT across repeated instances of a connection. TCP-AO protects the connections by using traffic key that are derived from the MKT, and then coordinates changes between the endpoints.




---

**Note** TCP-AO and TCP MD5 are never permitted to be used simultaneously. TCP-AO supports IPv6, and is fully compatible with the proposed requirements for the replacement of TCP MD5.

---

# Usage Guidelines and Limitations

To ensure PCEP compatibility, we recommend that the Cisco IOS XR version on the SR-PCE be the same or later than the Cisco IOS XR version on the PCC or head-end.

## Configure SR-PCE

This task explains how to configure SR-PCE.

### Procedure

|               | Command or Action                                                                                                            | Purpose                                                                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# <b>configure</b>                                                | Enters mode.                                                                                                                                                       |
| <b>Step 2</b> | <b>pce</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>pce</b>                                                    | Enables PCE and enters PCE configuration mode.                                                                                                                     |
| <b>Step 3</b> | <b>address ipv4 address</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-pce)#<br><b>address ipv4 192.168.0.1</b>       | Configures a PCE IPv4 address.                                                                                                                                     |
| <b>Step 4</b> | <b>state-sync ipv4 address</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-pce)#<br><b>state-sync ipv4 192.168.0.3</b> | Configures the remote peer for state synchronization.                                                                                                              |
| <b>Step 5</b> | <b>tcp-buffer size size</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-pce)#<br><b>tcp-buffer size 1024000</b>        | Configures the transmit and receive TCP buffer size for each PCEP session, in bytes. The default buffer size is 256000. The valid range is from 204800 to 1024000. |
| <b>Step 6</b> | <b>password {clear   encrypted} password</b><br><b>Example:</b>                                                              | Enables TCP MD5 authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the                                    |



|                | Command or Action                                                                                                                                                                                                   | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                | <pre>RP/0/RP0/CPU0:router(config-pce)# password encrypted pwd1</pre>                                                                                                                                                | <p>configured password will be rejected. Specify if the password is encrypted or clear text.</p> <p><b>Note</b><br/>TCP-AO and TCP MD5 are never permitted to be used simultaneously.</p>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 7</b>  | <p><b>tcp-ao</b> <i>key-chain</i> [<b>include-tcp-options</b>] [<b>accept-ao-mismatch-connection</b>]</p> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-pce)# tcp-ao pce_tcp_ao include-tcp-options</pre> | <p>Enables TCP Authentication Option (TCP-AO) authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured key chain will be rejected.</p> <ul style="list-style-type: none"> <li>• <b>include-tcp-options</b>—Includes other TCP options in the header for MAC calculation.</li> <li>• <b>accept-ao-mismatch-connection</b>—Accepts connection even if there is a mismatch of AO options between peers.</li> </ul> <p><b>Note</b><br/>TCP-AO and TCP MD5 are never permitted to be used simultaneously.</p> |
| <b>Step 8</b>  | <p><b>segment-routing</b> {<b>strict-sid-only</b>   <b>te-latency</b>}</p> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-pce)# segment-routing strict-sid-only</pre>                                      | <p>Configures the segment routing algorithm to use strict SID or TE latency.</p> <p><b>Note</b><br/>This setting is global and applies to all LSPs that request a path from this controller.</p>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Step 9</b>  | <p><b>timers</b></p> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-pce)# timers</pre>                                                                                                                     | <p>Enters timer configuration mode.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Step 10</b> | <p><b>keepalive</b> <i>time</i></p> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-pce-timers)# keepalive 60</pre>                                                                                         | <p>Configures the timer value for locally generated keep-alive messages. The default time is 30 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Step 11</b> | <p><b>minimum-peer-keepalive</b> <i>time</i></p> <p><b>Example:</b></p> <pre>RP/0/RP0/CPU0:router(config-pce-timers)#</pre>                                                                                         | <p>Configures the minimum acceptable keep-alive timer that the remote peer may propose in the PCEP OPEN message during session establishment. The default time is 20 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                          |

|                | Command or Action                                                                                                              | Purpose                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
|                | <code>minimum-peer-keepalive 30</code>                                                                                         |                                                                          |
| <b>Step 12</b> | <b>reoptimization</b> <i>time</i><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-pce-timers)# reoptimization 600</pre> | Configures the re-optimization timer. The default timer is 1800 seconds. |
| <b>Step 13</b> | <b>exit</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-pce-timers)# exit</pre>                                     | Exits timer configuration mode and returns to PCE configuration mode.    |

## Configure the Disjoint Policy (Optional)

This task explains how to configure the SR-PCE to compute disjointness for a pair of LSPs signaled by PCCs that do not include the PCEP association group-ID object in their PCEP request. This can be beneficial for deployments where PCCs do not support this PCEP object or when the network operator prefers to manage the LSP disjoint configuration centrally.

### Procedure

|               | Command or Action                                                                                                                                                                                                                                | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>disjoint-path</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-pce)# disjoint-path</pre>                                                                                                                                            | Enters disjoint configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 2</b> | <b>group-id</b> <i>value</i> <b>type</b> { <b>link</b>   <b>node</b>   <b>srlg</b>   <b>srlg-node</b> } [ <b>sub-id</b> <i>value</i> ]<br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-pce-disjoint)# group-id 1 type node sub-id 1</pre> | Configures the disjoint group ID and defines the preferred level of disjointness (the type of resources that should not be shared by the two paths): <ul style="list-style-type: none"> <li>• <b>link</b>—Specifies that links are not shared on the computed paths.</li> <li>• <b>node</b>—Specifies that nodes are not shared on the computed paths.</li> <li>• <b>srlg</b>—Specifies that links with the same SRLG value are not shared on the computed paths.</li> </ul> |

|               | Command or Action                                                                                                                                                                                                                                                                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               |                                                                                                                                                                                                                                                                                                                                                      | <ul style="list-style-type: none"> <li>• <b>srlg-node</b>—Specifies that SRLG and nodes are not shared on the computed paths.</li> </ul> <p>If a pair of paths that meet the requested disjointness level cannot be found, then the paths will automatically fallback to a lower level:</p> <ul style="list-style-type: none"> <li>• If the requested disjointness level is SRLG or node, then link-disjoint paths will be computed.</li> <li>• If the requested disjointness level was link, or if the first fallback from SRLG or node disjointness failed, then the lists of segments encoding two shortest paths, without any disjointness constraint, will be computed.</li> </ul> |
| <b>Step 3</b> | <b>strict</b><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-pce-disjoint)#<br><b>strict</b>                                                                                                                                                                                                                                              | (Optional) Prevents the automatic fallback behavior of the preferred level of disjointness. If a pair of paths that meet the requested disjointness level cannot be found, the disjoint calculation terminates and no new path is provided. The existing path is not modified.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Step 4</b> | <b>lsp {1   2} pcc ipv4 address lsp-name</b><br><b>lsp_name [shortest-path]</b><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-pce-disjoint)#<br><b>lsp 1 pcc ipv4 192.168.0.1 lsp-name</b><br><b>rtrA_t1 shortest-path</b><br>RP/0/RP0/CPU0:router(config-pce-disjoint)#<br><b>lsp 2 pcc ipv4 192.168.0.5 lsp-name</b><br><b>rtrE_t2</b> | Adds LSPs to the disjoint group.<br><br>The <b>shortest-path</b> keyword forces one of the disjoint paths to follow the shortest path from the source to the destination. This option can only be applied to the the first LSP specified.                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Global Maximum-delay Constraint

This feature allows a PCE to compare the cumulative latency of a computed path against a global maximum-delay constraint value. If the latency of the computed path exceeds this global constraint, the path is not considered valid. This ensures that all latency-based paths computed by the PCE and signaled to the PCCs in the network do not exceed this maximum-delay constraint.

```
pce
constraints
 bounds
 cumulative
```

```

type
latency <1-4294967295> Bound metric value in microseconds

```

### Configuration

To configure a PCE for specifying maximum cumulative latency metric, you must complete the following configurations:

```

RP/0/RSP0/CPU0:ios(config)# pce
RP/0/RSP0/CPU0:ios(config-pce)# constraints
RP/0/RSP0/CPU0:ios(config-pce-constr)# bounds
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds)# cumulative
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)# type latency 1000000
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)#

```

### Verification

Verify using the **show** command:

```

RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)# show
Wed Oct 12 22:18:22.962 UTC
pce
constraints
 bounds
 cumulative
 type latency 1000000
 !
!
!
!

```

## PCE override of SR-TE policies

Segment routing for traffic engineering (SR-TE) uses a “policy” to steer traffic through the network. An SR-TE policy path is expressed as a list of segments that specifies the path, called a segment ID (SID) list. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the network to follow the specified path instead of following the shortest path calculated by the IGP. If a packet is steered into an SR-TE policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

An SR-TE policy is identified as an ordered list (head-end, color, end-point):

- Head-end – Where the SR-TE policy is instantiated
- Color – An unsigned non-zero 32-bit integer value that associates the SR Policy with an intent or objective (e.g., low latency).
- End-point – The destination of the SR-TE policy.

There are fewer PCEs in the network compared to the PCCs. However, PCE versions may be newer than the PCC version because it is easy to update the lesser PCEs than the higher number of PCCs. This ensures PCE supports more features with the latest Internet Assigned Numbers Authority (IANA) code points than PCCs.

IOS-XR PCE implementation is based on the PCEP standards and the implementations of the PCEP vary leading to interoperability challenges. To overcome these challenges, you can configure PCE with the override-rules (OVR). These OVRs have a filter part and an overriding part. The matching criteria filter

matches all the (Segment Routing) SR policies against these OVRs, suppose they match, then the overriding part applies to these SR policies attributes.

For example, suppose that a PCC is outdated and does not support reporting the latency metric type. In that case, PCC reports to the LSP with the metric-type that it supports and the PCE applies the OVRs that you configure to update the metric-type upon reception.

## Usage guidelines

Some usage guidelines that apply to the PCE override of SR-TE policies feature.

- Lower sequence numbers have higher precedence. LSP can be matched to one override sequence at a time. It will be matched to the highest preference sequence and further matching stops.
- Show commands such as “show pce lsp” display values received from PCC, so applied changes are visible only in values, which are based on path-computation.
- In a PCE High Availability case, where there are multiple PCEs for a particular PCC, configure this feature likewise on all PCEs, to allow seamless transition of delegation between the PCEs.
- PCE (North Bound) NB-API clients only see overridden attributes in their communication with PCEs.
- The PCE OVR feature is seamless for PCCs. This means that PCCs will not know if the tunnel’s attributes are overridden. PCEs uses overridden attributes internally and sends only the original attributes to the PCCs.

## Limitations

- Supports only Segment Routing (SR) setup type.
- Supports only ten override rules.
- Only one color or color range per override rule is supported.
- Does not support SRv6 policies.
- You cannot configure the filtration option "all" under peer and LSP with another filtration option under the same config submode. For example, you cannot provide all and regex together under the LSP matching criteria.
- Only IPv4 Access Control Lists (ACL)s are supported for matching PCEP peers.
  - If the ACL is empty or there are no entries within the ACL, then it is implicit permit.
  - If the ACL is not configured but the name is specified, then it is implicit permit.
  - In all other cases, if no entries are matched in the ACL, then it is implicit deny.
  - The behavior aligns with ACL behavior in other areas on an XR router.

## Prerequisites

Before you configure the PCE override of SR-TE policies, you can use the **try-regex** utility to test the regexes. This utility is an optional step and does not alter the system.

This example shows how to verify if the

**override-rules sequence** *sequence-number* [ **matching-criteria** **lsp** { **all** | **name** *lsp-name-regex* } ] works:

```
Router# pce try-regex ^cp_c_[0-9]+$ cp_c_5000
Regex verification utility
Regex: ^cp_c_[0-9]+$
Test string: cp_c_5000
Result: Matched
```

## PCE override of PCC-initiated SR-TE policies

The Path Computation Element (PCE) is a key component of Software-Defined Networking (SDN) for Traffic Engineering (TE), which can dynamically compute and recompute paths for Segment Routing Traffic Engineering (SR-TE) policies. These computations are guided by specified path optimization metrics and constraints. By leveraging the PCE override feature, the PCE can modify certain attributes or constraints of these policies, enabling support for a broader range of use cases and enhanced functionality.

**Table 52: Feature History Table**

| Feature Name                           | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PCE override of PCC-initiated policies | Release 7.7.1       | <p>You can now override the LSP attributes in the PCEs based on the matching criteria of the path computation element protocol (PCEP) peer and LSP attributes. This feature improves the path computation of the PCC-initiated SR-TE policies.</p> <p>You can configure the PCE override rule for one, many, or all PCCs.</p> <p>This feature introduces the following commands:</p> <ul style="list-style-type: none"> <li>• <a href="#">override-rules (PCE)</a></li> <li>• <a href="#">pce try-regex</a></li> <li>• <a href="#">show pce override-rules detail</a></li> </ul> |

Matching of SR-TE policies is done before PCInitiate is sent to PCC to allow modification of data included in messages sent to PCC before LSP is created.

The LSP matching is not re-evaluated again after receiving the next PCEP report message. Based on these matching criteria, you can override the values of the LSP attributes.

## Configure PCE override of PCC-initiated SR-TE policies

This configuration defines an **override rule** for a PCE on a network router, which allows the PCE to modify certain attributes of Label Switched Paths (LSPs) that meet specific criteria for PCC-initiated SR-TE policies.

## Procedure

- 
- Step 1**      **configure**
- Example:**  
 Router# **configure**  
 Enters the global configuration mode.
- Step 2**      **pce**
- Example:**  
 Router(config)# **pce**  
 Enters the PCE configuration mode.
- Step 3**      **override-rules**
- Example:**  
 Router(config-pce)# **override-rules**  
 Enters the override rules submode.
- Step 4**      **sequence value**
- Example:**  
 Router(config-pce-ovr-rule)# **sequence 100**  
 Creates a rule with the sequence number 100.
- Step 5**      **matching-criteria**
- Example:**  
 Router(config-pce-ovr-rule-seq)# **matching-criteria**  
 Specifies a matching criteria. For example, **peer all**, **lsp all**, and so on.
- Step 6**      **peer all**
- Example:**  
 Router(config-pce-ovr-rule-crit)# **peer**  
 Router(config-pce-ovr-rule-peer)# **all**  
 Router(config-pce-ovr-rule-peer)# **exit**  
 Specifies that this rule applies to all peers.
- Step 7**      **lsp all range**
- Example:**  
 Router(config-pce-ovr-rule-crit)# **lsp all**  
 Router(config-pce-ovr-rule-crit)# **exit**  
 Specifies all LSPs in the PCE tunnel database.
- Step 8**      **colors range**

**Example:**

```
Router(config-pce-ovr-rule-crit)# colors 0-50
Router(config-pce-ovr-rule-crit)# exit
```

Defines that the colors in the 0 to 50 range to match this rule.

**Step 9**      **override metric type** *metric-type***Example:**

```
Router(config-pce-ovr-rule-seq)# override metric type igp
```

Specifies that the IGP (Interior Gateway Protocol) metric for matching LSPs is overridden by the PCE.

**Step 10**      **override constraints bandwidth** *value***Example:**

```
Router(config-pce-ovr-rule-seq)# override constraints bandwidth 1000
```

Specifies that the bandwidth constraint for matching LSPs is overridden and set to 1000.

**Step 11**      **commit****Example:**

```
Router(config-pce-ovr-rule-seq)# commit
```

Saves and apply these configuration changes.

**Running configuration**

These are the running configurations for the PCE override for the PCC-initiated SR-TE policies.

- Match all LSPs from all peers and modify metric type to latency:

```
pce
 override-rules
 sequence 1
 matching-criteria
 peer
 all
 lsp
 all

 override
 metric
 type latency
 !
```

- Match LSP name using regex and peer based on provided ACL and change metric type to IGP:

```
pce
 override-rules
 sequence 1
 matching-criteria
 peer
 access-list ipv4 PCC1
 !
 lsp
 name ^cfg_test2_.*$
```



```

!
override
metric
type igp
!

```

- Match policy using specified Segment Routing (SR) policy color from all peers and change bandwidth to 1000 kbps:

```

pce
 override-rules
 sequence 1
 matching-criteria
 peer
 all
 !
 lsp
 colors 10
 !
 override
 constraints
 bandwidth 1000
 !
 !

```

## Verification

Show command displays operational values, which are applied during the override. For example, a bandwidth value is converted into IEEE float format. The value that is displayed in the show command may not be equal to the value configured, because of the limited precision of IEEE format used in PCEP.

Verify using the show command **show pce override-rules [sequence <sequence>]**

This is the show command output with two override rules that are configured and each matched one LSP are:

```

Router# show pce override-rules details

PCE's Override Rule database:

Sequence number: 1
Matching criteria:
 Peer:
 IPv4 ACL name: PCC1
 LSP
 Regex: ^cfg_test1_.*$
Override:
 Metric type: Latency
 Constraints:
Matching LSPs:
 Peer: 192.168.0.1, Tunnel name: cfg_test1_xxx_discr_5

Sequence number: 2
Matching criteria:
 Peer:
 IPv4 ACL name: PCC1
 LSP
 Regex: ^cfg_test2_.*$
Override:
 Metric type: IGP
 Constraints:
Matching LSPs:
 Peer: 192.168.0.1, Tunnel name: cfg_test2_xxx_discr_5

```

## PCE override of PCE-initiated SR-TE policies

The PCE override feature can be applied to PCE-initiated SR-TE policies and serves several purposes, including controlling PCEP communication with PCCs to ensure interoperability. This feature is particularly useful in scenarios where specific adjustments are needed to address compatibility issues, allowing seamless integration across diverse network devices and vendors.

*Table 53: Feature History Table*

| Feature Name                                                         | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interoperability between PCEP peers for PCE-initiated SR-TE policies | Release 25.2.1      | <p>This feature allows interoperability between PCEP peers for PCE-initiated SR-TE policies by blocking encoding PCEP objects and sub-objects thus ensuring seamless integration across diverse network devices and vendors.</p> <p>This feature applies only to the PCE-initiated SR-TE policies.</p> <p>This feature introduces these changes:</p> <p><b>CLI:</b></p> <p>New keywords, <b>pcep</b>, <b>advertise</b>, <b>block-object</b>, and <b>sub-type</b> were introduced in the <b>override-rules (PCE)</b> command.</p> |

### Configure PCE override of PCE-initiated SR-TE policies

This configuration enables the PCE to adjust PCEP communication with specific PCEP peers (PCCs) and LSPs, effectively addressing interoperability challenges and ensuring seamless integration across diverse network devices and vendors.

#### Procedure

- 
- |               |                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>Router# <b>configure</b><br>Enters the global configuration mode.                    |
| <b>Step 2</b> | <b>pce</b><br><b>Example:</b><br>Router(config)# <b>pce</b><br>Enters the PCE configuration mode.                           |
| <b>Step 3</b> | <b>override-rules</b><br><b>Example:</b><br>Router(config-pce)# <b>override-rules</b><br>Enters the override rules submode. |

**Step 4** *sequence value***Example:**

```
Router(config-pce-ovr-rule) # sequence 100
```

Creates a rule with the sequence number 100.

**Step 5** *matching-criteria***Example:**

```
Router(config-pce-ovr-rule-seq) # matching-criteria
```

Specifies a matching criteria. For example, **peer all**, **lsp all**, and so on.

**Step 6** *peer access-list ipv4acl-name***Example:**

```
Router(config-pce-ovr-rule-crit) # peer
Router(config-pce-ovr-rule-peer) # access-list ipv4 PCC1
Router(config-pce-ovr-rule-peer) # exit
```

Specifies that this rule applies to IPv4 peers.

**Step 7** *lsp name regex***Example:**

```
Router(config-pce-ovr-rule-crit) # lsp name ^cfg_test2_.*$
Router(config-pce-ovr-rule-crit) # exit
```

Defines that LSP name (specified as a regular expression) match this rule.

**Step 8** *override pcep advertise value sub-type value***Example:**

```
Router(config-pce-ovr-rule-seq) # override pcep advertise
```

Overrides the attributes that the PCE advertises to its PCC peers via the PCEP.

**Step 9** *block-object value sub-type value***Example:**

```
Router(config-pce-ovr-rule-seq) # block-object 40 sub-type 6
```

Blocks encoding of PCEP ASSOCIATION object with object class 40 and ASSOCIATION type 6 ("SR Policy Association") as described in the [Path Computation Element Protocol \(PCEP\) Numbers](#).

**Note**

Currently, **block-object 40** and **sub-type 6** are the only supported object and sub-object types.

**Step 10** *commit***Example:**

```
Router(config-pce-ovr-rule-seq) # commit
```

Saves and apply these configuration changes.

## Running configuration

This is the running configuration for the PCE override for the PCE-initiated SR-TE policies.

Match LSP name using regex and peer based on provided ACL and block the PCEP ASSOCIATION object of association type 6 [SR Policy Association](#) in PCInitiate message, which is sent from PCE to PCC to create LSP:

```
pce
 override-rules
 sequence 1
 matching-criteria
 peer
 access-list ipv4 PCC1
 !
 lsp
 name ^cfg_test2_.*$
 !
 override
 pcep
 advertise
 block-object 40
 sub-type 6
 !
 !
 !
 !
 !
```

## Verification

This show command output demonstrates an override rule configured to block a PCEP sub-object. This resolves interoperability issues with PCC implementations that do not support the PCEP object.




---

**Note** Interoperability with fully capable PCCs does not need this configuration.

---

```
Router# show pce override-rules details
```

```
PCE's Override Rule database:
```

```

```

```
Sequence number: 1
```

```
Matching criteria:
```

```
Peer:
```

```
 IPv4 ACL name: PCC1
```

```
LSP
```

```
 Regex: ^cfg_test1_.*$
```

```
Override:
```

```
 PCEP advertise:
```

```
 Objects blocked:
```

```
 Type: 40
```

```
 Sub-Type: 6
```

```
Constraints:
```




---

**Note** Unlike the PCC-initiated LSPs, the PCE-initiated LSPs are not tracked in `Matching LSPs:` in the show output as matching occurs only after the LSP is created.

---

## PCE-Initiated SR Policies

Use cases based on centralized optimization, such as congestion mitigation solutions, rely on the ability of the PCE to signal and instantiate SR-TE policies in the network. We refer to this as PCE-initiated SR-TE policies.

PCE-initiated SR-TE policies can be triggered via Crossworks Network Controller (recommended approach) or via CLI at the PCE.

For more information on configuring SR-TE policies, see the [SR-TE Policy Overview, on page 225](#).

The PCE deploys the SR-TE policy using PCC-PCE communication protocol (PCEP).

1. PCE sends a PCInitiate message to the PCC.
2. If the PCInitiate message is valid, the PCC sends a PCRpt message; otherwise, it sends PCErr message.
3. If the PCInitiate message is accepted, the PCE updates the SR-TE policy by sending PCUpd message.

You can achieve high-availability by configuring multiple PCEs with SR-TE policies. If the head-end (PCC) loses connectivity with one PCE, another PCE can assume control of the SR-TE policy.

### Configuration Example: PCE-Initiated SR Policy with Explicit SID List

To configure a PCE-initiated SR-TE policy, you must complete the following configurations:

1. Enter PCE configuration mode.
2. Create the segment list.



**Note** When configuring an explicit path using IP addresses of intermediate links, the SR-TE process prefers the protected Adj-SID of the link, if one is available.

3. Create the policy.

```
/* Enter PCE configuration mode and create the SR-TE segment lists */
Router# configure
Router(config)# pce

/* Create the SR-TE segment lists */
Router(config-pce)# segment-routing
Router(config-pce-sr)# traffic-eng
Router(config-pce-sr-te)# segment-list name addr2a
Router(config-pce-sr-te-sl)# index 10 address ipv4 10.1.1.2
Router(config-pce-sr-te-sl)# index 20 address ipv4 10.2.3.2
Router(config-pce-sr-te-sl)# index 30 address ipv4 10.1.1.4
Router(config-pce-sr-te-sl)# exit

/* Create the SR-TE policy */
Router(config-pce-sr-te)# peer ipv4 10.1.1.1
Router(config-pce-sr-te)# policy P1
Router(config-pce-sr-te-policy)# color 2 end-point ipv4 2.2.2.2
Router(config-pce-sr-te-policy)# candidate-paths
Router(config-pce-sr-te-policy-path)# preference 50
```

```

Router(config-pce-sr-te-policy-path-preference)# explicit segment-list addr2a
Router(config-pce-sr-te-pp-info)# commit
Router(config-pce-sr-te-pp-info)# end
Router(config)#

```

### Running Config

```

pce
 segment-routing
 traffic-eng
 segment-list name addr2a
 index 10 address ipv4 10.1.1.2
 index 20 address ipv4 10.2.3.2
 index 30 address ipv4 10.1.1.4
 !
 peer ipv4 10.1.1.1
 policy P1
 color 2 end-point ipv4 2.2.2.2
 candidate-paths
 preference 50
 explicit segment-list addr2a
 !
 !

```

## SR-PCE Flexible Algorithm Multi-Domain Path Computation

Flexible Algorithm provides a traffic engineered path automatically computed by the IGP to any destination reachable by the IGP. With the SR-PCE Flexible Algorithm Multi-Domain Path Computation feature, SR-PCE can use Flexible Algorithms to compute multi-domain paths. See the [Enabling Segment Routing Flexible Algorithm, on page 573](#) chapter for information about Segment Routing Flexible Algorithm.

The SR-PCE Flexible Algorithm Multi-Domain Path Computation feature incorporates the following functionality:

- BGP-LS has been augmented to allow selected nodes to advertise the Flexible Algorithm definition (FAD) to the SR-PCE
- PCEP has been augmented (vendor-specific object) to allow a PCC to indicate SR policy constraint based on the Flexible Algorithm instance number
- SR-PCE algorithms have been augmented to compute paths based on a Flexible Algorithm constraint

The SR-PCE Flexible Algorithm multi-domain path computation requires the following:

- The same Flexible Algorithm instance ID is used across domains.
- The metric for those Flexible Algorithm instances must be the same across domains.
- The affinity constraints for those Flexible Algorithm instances may be different across domains.
- Multiple Flexible Algorithms can exist in a domain.

For example, considering a multi-domain topology (Domain 1 and Domain 2), the following scenarios meet the requirements listed above:

| Scenario   | Domain 1                                                                   | Domain 2                                                                   |
|------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------|
| Scenario 1 | Flexible Algorithm 128, metric delay                                       | Flexible Algorithm 128, metric delay                                       |
| Scenario 2 | Flexible Algorithm 128, metric delay                                       | Flexible Algorithm 128, metric delay, exclude affinity blue                |
| Scenario 3 | Flexible Algorithm 128, metric delay, exclude affinity yellow              | Flexible Algorithm 128, metric delay, exclude affinity blue                |
| Scenario 4 | Flexible Algorithm 128, metric delay<br>Flexible Algorithm 129, metric IGP | Flexible Algorithm 128, metric delay<br>Flexible Algorithm 129, metric IGP |




**Note** The use of a Flexible Algorithm constraint in a multi-domain SR topology does not preclude the use of an SR policy that are optimized for a particular metric type. For example, a policy can request a PCE for a Multi Domain policy based on metric delay. SR-PCE computes the path and encodes it with regular prefix SIDs and Adj-SIDs as required. Alternatively, a policy can request to have a constraint for a Flexible Algorithm instance X, which is defined in multiple domains and it minimizes based on metric delay. In this case, the SR-PCE computes the multi-domain path and encodes it using only Flexible Algorithm prefix SIDs. This case benefits from the optimized label stack size that Flexible Algorithm provides (1 label per domain).

## Example: SR-PCE Flexible Algorithm Multi-Domain Path Computation Use Case

The following use case depicts a multi-domain topology with two IS-IS processes, each with a Flexible Algorithm instance of 128 that minimizes metric delay. A multi-domain SR policy programmed at Node 1 leverages a Flexible Algorithm 128 path computed by the SR-PCE toward Node 8.

**Figure 30: Multi-Domain Topology**



### Configuration on Node 8

#### IS-IS and Flexible Algorithm Configuration

```
router isis 2
 is-type level-2-only
 net 49.0002.0000.0000.0008.00
 distribute link-state
 flex-algo 128
 metric-type delay
 advertise-definition

 address-family ipv4 unicast
 metric-style wide
 router-id 10.1.1.8
 segment-routing mpls
 !
 interface Loopback0
 passive
 address-family ipv4 unicast
 prefix-sid absolute 16008
 prefix-sid algorithm 128 absolute 16808
 !
```

### Configuration on Node 4 (ABR/ASBR)

#### IS-IS and Flexible Algorithm Configuration

```
router isis 1
 is-type level-2-only
 net 49.0001.0000.0000.0004.00
 distribute link-state instance-id 100
 flex-algo 128
 metric-type delay
 advertise-definition

 address-family ipv4 unicast
 metric-style wide
 router-id 10.1.1.4
 segment-routing mpls
 !
 interface Loopback0
 passive
 address-family ipv4 unicast
 prefix-sid absolute 16004
 prefix-sid algorithm 128 absolute 16804
 !

router isis 2
 is-type level-2-only
 net 49.0002.0000.0000.0004.00
 distribute link-state instance-id 200
 flex-algo 128
 metric-type delay
 advertise-definition

 address-family ipv4 unicast
 metric-style wide
 router-id 10.1.1.4
 segment-routing mpls
 !
 interface Loopback0
 passive
 address-family ipv4 unicast
 prefix-sid absolute 16004
 prefix-sid algorithm 128 absolute 16804
```



!

**BGP-LS Configuration**

```

router bgp 65000
 bgp router-id 10.1.1.4
 address-family link-state link-state
 !
 neighbor-group AS65000-LS-group
 remote-as 65000
 update-source Loopback0
 address-family link-state link-state
 !
 !
 neighbor 10.1.1.10
 use neighbor-group AS65000-LS-group
 description *** To SR-PCE ***
 !
 !
 !

```

**Configuration on Node 1****IS-IS and Flexible Algorithm Configuration**

```

router isis 1
 is-type level-2-only
 net 49.0001.0000.0000.0001.00
 distribute link-state
 flex-algo 128
 metric-type delay
 advertise-definition

 address-family ipv4 unicast
 metric-style wide
 router-id 10.1.1.1
 segment-routing mpls
 !
 interface Loopback0
 passive
 address-family ipv4 unicast
 prefix-sid absolute 16001
 prefix-sid algorithm 128 absolute 16801
 !

```

**SR Policy Configuration**

```

segment-routing
 traffic-eng
 policy FOO
 color 100 end-point ipv4 10.1.1.8
 candidate-paths
 preference 100
 dynamic
 pcep
 !
 !
 constraints
 segments
 sid-algorithm 128
 !
 !
 !

```

```

!
!
!
!

```

### PCC Configuration

```

segment-routing
traffic-eng
pcc
source-address ipv4 10.1.1.1
pce address ipv4 10.1.1.10
precedence 10
!
report-all
!
!
!

```

### Configuration on PCE

```

pce
address ipv4 10.1.1.10
rest
!
!
router bgp 65000
bgp router-id 10.1.1.10
address-family link-state link-state
!
neighbor-group AS65000-LS-group
remote-as 65000
update-source Loopback0
address-family link-state link-state
!
!
neighbor 10.1.1.4
use neighbor-group AS65000-LS-group
description *** To Node-4 ***
!
!
neighbor 10.1.1.5
use neighbor-group AS65000-LS-group
description *** To Node-5 ***
!
!
!

```

# Exclude Network Resources during Path Computation over SR-TE Policies

Table 54: Feature History Table

| Feature Name                                                          | Release Information | Feature Description |
|-----------------------------------------------------------------------|---------------------|---------------------|
| Exclude Network Resources during Path Computation over SR-TE Policies | Release 24.1.1      |                     |

| Feature Name | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              |                     | <p>You can reduce the risk of attacks due to less secure network resources and IP addresses, improve network performance affected by overloaded or congested paths, and ensure higher levels of network stability and availability that could be impacted by resources experiencing issues and requiring maintenance. These improvements are possible because you can now exclude specific network resources using their IP addresses during path computation for traffic over SR-TE policies.</p> <p>Previously, you could not exclude network resources during path computation.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>resources</b> option is introduced in the <b>segment-routing traffic-eng policy</b> and <b>segment-routing traffic-eng on-demand color</b> commands.</li> <li>The <b>resource-list</b> keyword is introduced in the <b>segment-routing traffic-eng</b> command.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>Cisco-IOS-XR-infra-xtc-oper.yang</li> <li>Cisco-IOS-XR-infra-xtc-agent-oper.yang</li> <li>Cisco-IOS-XR-infra-xtc-agent-cfg.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

You can configure the SR-TE policies as either single or disjoint candidate paths that do not include certain IP addresses or subnets when traffic is routed through the network.

The key benefits of the feature are:

- **Security** - Excluding certain IP addresses or network resources can help protect sensitive resources from unauthorized access or attacks.
- **Performance** - By excluding certain resources, traffic can be directed away from congested parts of the network.
- **Isolation** - Excluding resources can help isolate different parts of the network, preventing issues in one part of the network from affecting others.

## Configuration to exclude network resources during Path Computation

To exclude network resources during path computation for traffic, you must complete these high-level tasks in order:

1. Configure the network resources or IP addresses that you want to exclude from the network list.
2. Associate the excluded network resources to candidate paths for SR-TE or ODN SR-TE policies.

### Configure the network resources or IP addresses to exclude from the network list

Perform the following task to configure a list of IPv4 addresses that you want to exclude from the network resource list:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#resource-list node_resc_list
Router(config-sr-te-rl)#index 1 ipv4 10.10.10.1
Router(config-sr-te-rl)#index 2 ipv4 10.10.10.8
```

#### Running Configuration

```
!
segment-routing
 traffic-eng
 resource-list node_resc_list
 index 1 ipv4 10.10.10.1
 index 2 ipv4 10.10.10.8
 !
!
```

### Associate the excluded network resources to candidate paths for SR-TE policies

Perform the following task to associate the excluded IPv4 addresses to one or more candidate paths for SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy dynamic pcep_policy
Router(config-sr-te-policy)#candidate-paths
Router(config-sr-te-policy-path)#preference 100
Router(config-sr-te-policy-path-pref)#constraints resources exclude resource-list
node_resc_list
```

#### Running Configuration

```
!
segment-routing
 traffic-eng
 policy dynamic pcep_policy
 candidate-paths
 constraints resources exclude resource-list
 node_resc_list
```

## Verification

```
Router#show segment-routing traffic-eng policy endpoint ipv4 100.2.1.1 color 8001
```

```

Wed Aug 30 22:21:50.014 UTC
SR-TE policy database

Color: 8001, End-point: 100.2.1.1
Name: srte_c_8001_ep_100.2.1.1
Status:
Admin: up Operational: up for 00:00:50 (since Aug 30 22:20:59.341)
Candidate-paths:
Preference: 100 (configuration) (active)
Name: dynamic_pcep_policy
Requested BSID: 8001
PCC info:
Symbolic name: cfg_dynamic_pcep_policy_discr_100
PLSP-ID: 14042
Constraints:
Protection Type: protected-preferred
Maximum SID Depth: 12
Exclude Resources: node_resc_list
10.10.10.1
10.10.10.8
Dynamic (pce 100.7.1.1) (valid)
Metric Type: TE, Path Accumulated Metric: 440
SID[0]: 41111 [Adjacency-SID, 101.1.5.1 - 101.1.5.2]
SID[1]: 21600 [Prefix-SID, 100.6.1.1]
SID[2]: 41600 [Adjacency-SID, 101.2.6.2 - 101.2.6.1]
Attributes:
Binding SID: 8001
Forward Class: Not Configured
Steering labeled-services disabled: yes
Steering BGP disabled: no
IPv6 caps enable: yes
Invalidation drop enabled: no
Max Install Standby Candidate Paths: 0

```

### Associate the excluded network resources to candidate paths for ODN SR-TE policies

Perform the following task to associate the excluded IPv4 addresses for ODN SR-TE policies:

```
Router(config)#segment-routing
Router(config-sr)#traffic-eng
Router(config-sr-te)#on-demand color 7001
Routerconfig-sr-te-color)#constraints resources exclude resource-list node resc_list
```

## Running Configuration

```
!
segment-routing
traffic-eng
on-demand color 7001
```

```

constraints
resources
 exclude resource-list node_resc_list
!
!
!
!
!

```

## Verification

Router#**show segment-routing traffic-eng policy endpoint ipv4 100.2.1.1 color 7001**

Wed Aug 30 22:53:01.079 UTC

SR-TE policy database

-----

Color: 7001, End-point: 100.2.1.1

Name: srte\_c\_7001\_ep\_100.2.1.1

Status:

Admin: up Operational: up for 00:31:56 (since Aug 30 22:21:04.869)

Candidate-paths:

Preference: 200 (BGP ODN) (inactive) (shutdown)

Requested BSID: dynamic

Constraints:

Protection Type: protected-preferred

Maximum SID Depth: 12

**Exclude Resources: node\_resc\_list**

**10.10.10.1**

**10.10.10.8**

Dynamic (inactive)

Metric Type: IGP, Path Accumulated Metric: 0

Preference: 100 (BGP ODN) (active)

Requested BSID: dynamic

PCC info:

Symbolic name: bgp\_c\_7001\_ep\_100.2.1.1\_discr\_100

PLSP-ID: 14044

Constraints:

Protection Type: protected-preferred

Maximum SID Depth: 12

**Exclude Resources: node\_resc\_list**

**10.10.10.1**

**10.10.10.8**

Dynamic (pce 100.7.1.1) (valid)

Metric Type: IGP, Path Accumulated Metric: 440

SID[0]: 41111 [Adjacency-SID, 101.1.5.1 - 101.1.5.2]

SID[1]: 21600 [Prefix-SID, 100.6.1.1]

SID[2]: 41600 [Adjacency-SID, 101.2.6.2 - 101.2.6.1]

Attributes:

Binding SID: 51679

Forward Class: Not Configured

Steering labeled-services disabled: yes

Steering BGP disabled: no

IPv6 caps enable: yes

Invalidation drop enabled: no

Max Install Standby Candidate Paths: 0

# Configure the Shortest Path for Disjoint Candidate Paths

Table 55: Feature History Table

| Feature Name                                             | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the Shortest Path for Disjoint Candidate Paths | Release 24.1.1      | <p>You can now configure the available disjoint paths for Label Switched Paths (LSPs) to prefer the shortest path between two points in the network. This configuration ensures that traffic is routed along the most efficient route in the network.</p> <p>Previously, you could not configure the shortest path preference for disjoint LSPs.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>shortest-path</b> keyword is introduced in the <b>policy candidate-paths constraints disjoint-path</b> command.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• Cisco-IOS-XR-infra-xtc-oper.yang</li> <li>• Cisco-IOS-XR-infra-xtc-agent-oper.yang</li> <li>• Cisco-IOS-XR-infra-xtc-agent-cfg.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

The configuration enables the available disjoint paths to indicate preference for the shortest path between two points in the network. This way the traffic is contained within specific network planes as per the affinity constraints while still provisioning Label Switched Paths (LSPs) that are diverse from each other to meet the path disjointness requirement.

In earlier releases, if path disjointness was configured for LSPs, the Path Computation Element (PCE) prioritized finding disjoint paths for the LSPs over adhering to the affinity constraints.

## Configure the shortest path for disjoint candidate paths

Perform the following task to indicate the disjoint path preference for the shortest path in the network:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy dynamic_pcep_policy_disjoint
```



```

Router(config-sr-te-policy)#candidate-paths
Router(config-sr-te-policy-path)#preference 100
Router(config-sr-te-policy-path-pref)#constraints disjoint-path group-id 1 type link
shortest-path

```

### Running Configuration

```

!
segment-routing
 traffic-eng
 policy dynamic_pcep_policy
 candidate-paths
 preference 100
 constraints
 disjoint-path group-id 1 type link shortest-path
!
!
!
!
!
!

```

### Verification

```
Router#show pce lsp name cfg_dynamic_pcep_policy_disjoint_discr_100 detail
```

Output received:

Wed Aug 30 18:33:57.807 UTC

PCE's tunnel database:

-----

PCC 100.1.1.1:

Tunnel Name: cfg\_dynamic\_pcep\_policy\_disjoint\_discr\_100

Color: 30115

Interface Name: srte\_c\_30115\_ep\_100.2.1.1

LSPs:

LSP[0]:

source 100.1.1.1, destination 100.2.1.1, tunnel ID 9031, LSP ID 24

State: Admin up, Operation active

Setup type: Segment Routing

Binding SID: 30115

Maximum SID Depth: 12

Preference: 100

Bandwidth: requested 0 kbps, applied 0 kbps

Protection type: protected-preferred

Prefix-SID algorithm: 0 (set: 0)

PCEP information:

PLSP-ID 0x36e9, flags: D:1 S:0 R:0 A:1 O:2 C:0

LSP Role: Exclude LSP

State-sync PCE: None

PCC: 100.1.1.1

LSP is subdelegated to: None

Reported path:

Metric type: TE, Accumulated Metric 30

SID[0]: Node, Label 21200, Address 100.2.1.1

Computed path: (Local PCE)

Computed Time: Wed Aug 30 18:31:12 UTC 2023 (00:02:45 ago)

Metric type: TE, Accumulated Metric 30

SID[0]: Node, Label 21200, Address 100.2.1.1

Reverse path: (Local PCE)

None

Computed Time: Wed Aug 30 18:31:12 UTC 2023 (00:02:45 ago)

Recorded path:

None

```

Disjoint Group Information:
 Type Link-Disjoint, Group 1 (SP)
SR Policy Association Group:
 Color: 30115, Endpoint: 100.2.1.1
 Policy Name: srte_c_30115_ep_100.2.1.1
 Preference: 100
 CP Name: dynamic_pcep_policy_disjoint

```

## Enable Strict Disjointness for SR-TE policies in the PCE

Table 56: Feature History Table

| Feature Name                                             | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable Strict Disjointness for SR-TE Policies in the PCE | Release 24.1.1      | <p>You can now enforce strict disjoint constraints for Label Switched Paths (LSPs) and minimize the risk of a single point of failure that affects multiple LSPs. With this feature, if disjoint paths cannot be found for LSPs, the Path Computation Element (PCE) does not return any path.</p> <p>Previously, enforcing strict disjoint constraints for disjoint paths for LSPs was not possible.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>fallback disable</b> keyword is introduced in the <a href="#">segment-routing traffic-eng policy</a> and <a href="#">segment-routing traffic-eng on-demand color</a> commands.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-oper.yang</a></li> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-agent-oper.yang</a></li> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-agent-cfg.yang</a></li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

When you enable fallback disable, the Path Computation Element (PCE) does not return any path for LSPs that are not disjoint if there are no disjoint candidate paths. In other words, when the PCE does not find multiple LSPs that satisfy the disjointness requirement, it returns no path. This indicates that there might be a networking issue that prevents the establishment of the desired redundancy, and you can investigate and address the underlying problem.

Without the option, the PCE can relax disjointness by applying an objective function or use a local policy when no objective function is requested.

Enforcing strict disjointness ensures that LSPs do not share any common links or nodes along their paths, minimizing the risk of a single point of failure affecting multiple LSPs. The feature is critical to maintain network reliability and ensures that traffic is properly routed in the event of a network failure.

## Configure Strict Disjointness in the PCE

### Enable strict disjointness for ODN SR-TE policies

Perform the following steps to enable strict disjointness for ODN SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#on-demand color 4
Router(config-sr-te-color)#dynamic
Router(config-sr-te-color-dyn)#disjoint-path group-id 1 type node fallback disable
Router(config-sr-te-color-dyn)#commit
```

### Running Configuration

```
segment-routing
 traffic-eng
 on-demand color 4
 dynamic
 disjoint-path group-id 1 type node fallback disable
 !
 !
 !
 !
```

### Enable strict disjointness for SR-TE policies

Perform the following steps to enable strict disjointness for SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy foo
Router(config-sr-te-policy)#color 1 end-point ipv4 10.10.10.1
Router(config-sr-te-policy)#candidate-paths preference 100
Router(config-sr-te-policy-path-pref)#constraints disjoint-path group-id 1 type node fallback
disable
Router(config-sr-te-policy-path-pref)#commit
```

### Running Configuration

```
segment-routing
 traffic-eng
 policy foo
 color 1 end-point ipv4 10.10.10.1
 candidate-paths
 preference 100
 dynamic
 pcep
 !
 metric
 type latency
 !
 constraints
 disjoint-path group-id 1 type node fallback disable
 !
```

```

!
!
!
!
!

```

### Verification

```

Router#sh pce association
Wed Mar 13 14:38:27.173 PDT

PCE's association database:

Association: Type Node-Disjoint, Group 1, Strict
Associated LSPs:
 LSP[0]:
 PCC 10.10.10.1, tunnel name cfg_foo_discr_100, PLSP ID 4, tunnel ID 8, LSP ID 2,
 Configured on PCC
 LSP[1]:
 PCC 10.10.10.1, tunnel name bgp_c_4_ep_10.10.10.2_discr_100, PLSP ID 1, tunnel ID 5,
 LSP ID 3, Configured on PCC
Status: Satisfied

```

## ACL Support for PCEP Connection

PCE protocol (PCEP) (RFC5440) is a client-server model running over TCP/IP, where the server (PCE) opens a port and the clients (PCC) initiate connections. After the peers establish a TCP connection, they create a PCE session on top of it.

The ACL Support for PCEP Connection feature provides a way to protect a PCE server using an Access Control List (ACL) to restrict IPv4 PCC peers at the time the TCP connection is created based on the source address of a client. When a client initiates the TCP connection, the ACL is referenced, and the client source address is compared. The ACL can either permit or deny the address and the TCP connection will proceed or not.

Refer to the Implementing Access Lists and Prefix Lists chapter in the *IP Addresses and Services Configuration Guide* for detailed ACL configuration information.

To apply an ACL to the PCE, use the **pce peer-filter ipv4 access-list *acl\_name*** command.

The following example shows how to configure an ACL and apply it to the PCE:

```

pce
 address ipv4 10.1.1.5
 peer-filter ipv4 access-list sample-peer-filter
!
ipv4 access-list sample-peer-filter
 10 permit ipv4 host 10.1.1.6 any
 20 permit ipv4 host 10.1.1.7 any
 30 deny ipv4 any any
!

```

## Inter-Domain Path Computation Using Redistributed SID


A Path Computation Element (PCE) computes SR-TE paths based on SR topology database that stores connectivity, state, and TE attributes of SR network nodes and links. BGP Labeled Unicast (BGP-LU) provides

MPLS transport across IGP boundaries by advertising loopbacks and label binding of impact edge and border routers across IGP boundaries.

This feature adds new functionality to the SR-PCE that enables it to compute a path for remote non-SR end-point device distributed by BGP-LU.

The remote end-point device in the BGP-LU domain is unknown to the SR-PCE. For the SR-PCE to know about the end-point device, the gateway ABR/ASBR learns the end-point prefix via BGP-LU. The prefix is then redistributed to SR-PCE topology database from the gateway ABR/ASBR. SR-PCE then can compute the best path from the head-end device to the selected gateway router.


The following topology shows an SR domain and a BGP-LU domain, with a gateway ABR/ASBR between the two domains.



1. The gateway ABR/ASBR is configured with BGP/IGP helper to learn the remote prefix through BGP-LU and redistribute the remote prefix to the IGP helper, then to SR-PCE.
2. The SR-PCE selects the best gateway node to BGP-LU domain and computes the path to reach the remote prefix through the gateway node.
3. The head-end device in the SR domain requests a path to the remote destination and signals the SR profile interworking with the BGP-LU domain.

The BGP-LU prefix advertisement to SR-PCE Traffic Engineer Database (TED) is done by creating an IGP helper on the ABR/ASBR to redistribute BGP-LU prefix information to IGP. IGP then sends the prefix information to the SR-PCE via BGP-LS.

If there are multiple ABR/ASBRs advertising the same remote BGP-LU prefix, the SR-PCE selects the best gateway node to the BGP-LU domain using the accumulative metric from the head-end device to the gateway and the advertised metric from the gateway to the destination.



## Example: Inter-Domain Path Computation Using Redistributed SID

The following examples show the configurations for the IGP helper, BGP-LU, and proxy BGP-SR:

### Configuration on the End-Point Device

Configure the end-point device to allocate a label for the BGP-LU prefix on the end-point device:

```
router bgp 3107
 bgp router-id 1.0.0.8
 address-family ipv4 unicast
 network 1.0.0.8/32 route-policy bgplu-com
 allocate-label all

route-policy bgplu-com
 set community (65002:999)
end-policy
```

### Configuration on the Gateway ABR/ASBR

1. Configure the remote prefix set and create the route policy for the BGP-LU domain:

```
prefix-set bgplu
 1.0.0.7/32,
 1.0.0.8/32,
 1.0.0.101/32,
 1.0.0.102/32
end-set
!

route-policy bgp2isis
 if destination in bgplu then
 pass
 else
 drop
 endif
```

```
end-policy
!
end
```

2. Configure the helper IGP instance on the Loopback interface:

```
router isis 101
 is-type level-2-only
 net 49.0001.0000.1010.1010.00
 distribute link-state instance-id 9999
 nsf cisco
 nsf lifetime 120
 address-family ipv4 unicast
 metric-style wide
 maximum-paths 64
 router-id Loopback10
 redistribute bgp 3107 metric 200 route-policy bgp2isis
 segment-routing mpls sr-prefer
!
interface Loopback10 >>> this loopback is for gateway SR-TE node-id
 passive
 address-family ipv4 unicast
 prefix-sid index 2001 explicit-null
```

3. Configure the gateway proxy BGP-SR and SR Mapping Server to allocate SR labels:

```
router bgp 3107
 address-family ipv4 unicast
 segment-routing prefix-sid-map
 allocate-label all

segment-routing
 global-block 16000 23999
 mapping-server
 prefix-sid-map
 address-family ipv4
 1.0.0.7/32 2007
 1.0.0.8/32 2008
 1.0.0.101/32 2101
 1.0.0.102/32 2102
```

## PCE Support for MPLS-TE LSPs

This feature allows Cisco's SR-PCE to act as a Path Computation Element (PCE) for MPLS Traffic Engineering Label Switched Paths (MPLS-TE LSPs).



**Note** For more information about MPLS-TE, refer to the "Implementing MPLS Traffic Engineering" chapter in the *MPLS Configuration Guide*.

The supported functionality is summarized below:

- PCE type: Active Stateful PCE
- MPLS-TE LSP initiation methods:
  - PCE Initiated—An active stateful PCE initiates an LSP and maintains the responsibility of updating the LSP.

- PCC Initiated—A PCC initiates the LSP and may delegate the control later to the Active stateful PCE.
- MPLS-TE LSP metric—Metric optimized by the path computation algorithm:
  - IGP metric
  - TE metric
  - Latency metric
- MPLS-TE LSP constraints—TE LSP attributes to be taken into account by the PCE during path computation:
  - Resource Affinities
  - Path Disjointness
- MPLS-TE LSP parameters:
  - Setup priority—The priority of the TE LSP with respect to taking resources
  - Hold priority—The priority of the TE LSP with respect to holding resources
  - FRR L flag—The "Local Protection Desired" bit. Can be set from an application instantiating an MPLS-TE LSP via SR-PCE. SR-PCE passes this flag to the PCC, and the PCC will enable FRR for that LSP.
  - Signaled Bandwidth—This value can be set from an application instantiating an MPLS-TE LSP via SR-PCE. SR-PCE passes this value to the PCC.
  - Binding SID—A segment identifier (SID) that a headend binds to an MPLS-TE LSP. When the headend receives a packet with active segment (top MPLS label) matching the BSID of a local MPLS-TE LSP, the headend steers the packet into the associated MPLS-TE LSP.

Cisco Crosswork Optimization Engine is an application that leverages the SR-PCE in order to visualize and instantiate MPLS-TE LSPs. For more information, refer to the [Visualize SR Policies and RSVP-TE Tunnels](#) chapter in the [Cisco Crosswork Optimization Engine 1.2.1 User Guide](#).



---

**Note** No extra configuration is required to enable MPLS-TE support at SR-PCE.

---

#### Example: Configuring a PCEP Session (Stateful Mode) on MPLS-TE PCC

The following example shows the configuration for an MPLS-TE PCC to establish a PCEP session with a PCE (IPv4 address 10.1.1.100).



---

**Note** MPLS-TE PCC must operate in the stateful PCEP mode when connecting to SR-PCE.

---

The **instantiation** keyword enables the PCC to support MPLS-TE LSP instantiation by PCE (PCE-initiated).

The **report** keyword enables the PCC to report all the MPLS-TE LSPs configured on that node.






---

**Note** PCE-initiated LSPs are automatically reported to all configured PCEs.

---

The **autoroute-announce** keyword enables autoroute-announce globally for all PCE-initiated LSPs on the PCC.

The **redundancy pcc-centric** keywords enable PCC-centric high-availability model for PCE-initiated LSPs. The PCC-centric model changes the default PCC delegation behavior to the following:

- After LSP creation, LSP is automatically delegated to the PCE that computed it.
- If this PCE is disconnected, then the LSP is redelegated to another PCE.
- If the original PCE is reconnected, then the delegation fallback timer is started. When the timer expires, the LSP is redelegated back to the original PCE, even if it has worse preference than the current PCE.

```
mpls traffic-eng
pce
 peer ipv4 10.1.1.100
 !
 stateful-client
 instantiation
 report
 autoroute-announce
 redundancy pcc-centric
 !
!
end
```

#### Example: Configuring Multiple PCEP Sessions from a PCC Acting as MPLS-TE and SR-TE Headend Toward a Common PCE

The following example shows the configuration for a PCC (IPv4 addresses 10.1.1.1 and 10.1.1.2) to establish two PCEP sessions with a common PCE (IPv4 address 10.1.1.100). One session is configured under MPLS-TE, and the other under SR-TE.




---

**Note** The two PCEP sessions must use a different source address on the PCC when connecting to the same PCE.

---

For more information regarding PCEP configuration at SR-TE PCC, see the *Configure the Head-End Router as PCEP PCC* topic.

```
mpls traffic-eng
pce
 peer source ipv4 10.1.1.1
 peer ipv4 10.1.1.100
 !
!
end

segment-routing
traffic-eng
```

```
pcc
 source-address ipv4 10.1.1.2
 pce address ipv4 10.1.1.100
 !
!
!
end
```

## Configuring the North-Bound API on SR-PCE

*Table 57: Feature History Table*

*Table 58: Feature History Table*

| Feature Name                                  | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-PCE: Stateful North-Bound API for Tree-SID | Release 7.5.1       | <p>The SR-PCE provides a north-bound HTTP-based API to allow communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds stateful north-bound APIs to support real-time monitoring of Tree-SID states on the SR-PCE using a subscription model.</p> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p> |

Table 59: Feature History Table

| Feature Name                                                                                            | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR-PCE: North-Bound API for SRv6 and Flexible Algorithm in Cisco Optimization Engine (COE) v3.0 release | Release 7.3.2       | <p>The SR-PCE provides a north-bound HTTP-based API to allow communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds support for the following:</p> <ul style="list-style-type: none"> <li>• Reporting of Flexible Algorithm participation and definitions</li> <li>• SRv6 topology information (nodes, links, Node uSIDs and Adj uSIDs)</li> <li>• SRv6 uSID list and uB6 SIDs allocated for a policy</li> </ul> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p> |

The SR-PCE provides a north-bound HTTP-based API to allow communication between SR-PCE and external clients and applications.

Over this API, an external application can leverage the SR-PCE for topology discovery, SR policy discovery, and SR policy instantiation.

The Cisco Crosswork Optimization Engine is an application that leverages the SR-PCE. For more information, refer to the [Cisco Crosswork Optimization Engine User Guides](#).

Use the following commands under PCE configuration mode to configure the API to allow communication between SR-PCE and external clients or applications.

| Command                                                                                    | Description                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>api authentication</b> { <b>basic</b>   <b>digest</b> }                                 | Specify the type of authentication: <ul style="list-style-type: none"> <li>• <b>basic</b> – Use HTTP Basic authentication (plaintext)</li> <li>• <b>digest</b> – Use HTTP Digest authentication (MD5)</li> </ul> |
| <b>api username</b> <i>password</i> { <b>clear</b>   <b>encrypted</b> }<br><i>password</i> | Add credentials when connecting to API.                                                                                                                                                                          |

| Command                         | Description                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>api sibling ipv4 address</b> | <p>Opens a synchronization channel to another PCE in the same high availability (HA) pair.</p> <p><b>Note</b><br/>For more information regarding SR-PCE HA pairs, refer to the <a href="#">Multiple Cisco SR-PCE HA Pairs</a> chapter of the <a href="#">Cisco Crosswork Optimization Engine 1.2.1 User Guide</a>.</p> |

### Example: Configuring API on SR-PCE

```
pce
 address ipv4 10.1.1.100
 api
 user admin
 password encrypted 1304131F0202
 !
 authentication digest
 sibling ipv4 10.1.1.200
 !
end
```

The following example shows the current active connections:

```
RP/0/0/CPU0:pce1# show tcp brief | i 8080
Thu Aug 6 00:40:15.408 PDT
0xe9806fb8 0x60000000 0 0 :::8080 :::0 LISTEN
0xe94023b8 0x60000000 0 0 10.1.1.100:50487 10.1.1.200:8080 ESTAB
0xeb20bb40 0x60000000 0 0 10.1.1.100:8080 10.1.1.200:44401 ESTAB
0xe98031a0 0x60000000 0 0 0.0.0.0:8080 0.0.0.0:0 LISTEN
```

The first and fourth entries show the API server listening for IPv4 and IPv6 connections.

The second and third entries show the established sibling connection between PCE1 (10.1.1.100) and PCE2 (10.1.1.200).





## CHAPTER 11

# Configure Performance Measurement

Network performance metrics is a critical measure for traffic engineering (TE) in service provider networks. Network performance metrics include the following:

- Packet loss
- Delay
- Delay variation
- Bandwidth utilization

These network performance metrics provide network operators information about the performance characteristics of their networks for performance evaluation and help to ensure compliance with service level agreements. The service-level agreements (SLAs) of service providers depend on the ability to measure and monitor these network performance metrics. Network operators can use Segment Routing Performance Measurement (SR-PM) feature to monitor the network metrics for links and end-to-end TE label switched paths (LSPs).

The following table explains the functionalities supported by performance measurement feature for measuring delay for links or SR policies.

**Table 60: Performance Measurement Functionalities**

| Functionality                    | Details                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Profiles                         | You can configure different default profiles for different types of delay measurements. Use the "interfaces" delay profile type for link-delay measurement. The "sr-policy" delay profile type is used for SR policy delay measurements. Delay profile allows you to schedule probe and configure metric advertisement parameters for delay measurement. |
| Protocols                        |                                                                                                                                                                                                                                                                                                                                                          |
| Probe and burst scheduling       | Schedule probes and configure metric advertisement parameters for delay measurement.                                                                                                                                                                                                                                                                     |
| Metric advertisements            | Advertise measured metrics periodically using configured thresholds. Also supports accelerated advertisements using configured thresholds.                                                                                                                                                                                                               |
| Measurement history and counters | Maintain packet delay and loss measurement history, session counters, and packet advertisement counters.                                                                                                                                                                                                                                                 |

- [Liveness Monitoring](#), on page 438
- [Delay Measurement](#), on page 463
- [Fallback delay advertisement](#), on page 504
- [Two-Way Active Measurement Protocol Light Source Address Filtering](#), on page 507
- [Synthetic Loss Measurement](#), on page 511
- [Delay and synthetic loss measurement for GRE tunnel interfaces](#), on page 517
- [Delay Measurement Using Software Timestamp](#), on page 520

## Liveness Monitoring

Liveness refers to the ability of the network to confirm that a specific path, segment, or a node is operational and capable of forwarding packets. Liveness checks are essential for maintaining network availability and reliability. See *Configure PTP in System Management Configuration Guide* for more information on configuring PTP.

### Benefits

- **Fault Detection:** You can quickly identify if a device is down, which allows for immediate response and troubleshooting.
- **Load Balancing:** You can identify if the devices in a network are live, so work can be distributed more evenly across the network, preventing overloading of specific components and improving overall performance.
- **System Health:** You can provide an ongoing snapshot of a system's health, helping to identify potential issues before they become significant problems.
- **Maintenance Planning:** Liveness information can also help with maintenance planning, as system administrators can understand which components are live or down and plan maintenance and downtime accordingly without significant disruption to services.
- **Security:** Regular liveness checks can also play a role in maintaining network security. Administrators can take proactive steps to mitigate the damage and prevent future incidents by identifying unusual activity that might indicate a security breach or attack.

You can determine liveness for SR Policy and IP Endpoint.

## IP Endpoint Liveness Monitoring

*Table 61: Feature History Table*

| Feature Name                                          | Release Information | Feature Description                                                                            |
|-------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------|
| Liveness Monitoring for IP Endpoint over SRv6 Network | Release 24.2.11     | This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7. |

| Feature Name                                          | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Liveness Monitoring for IP Endpoint over SRv6 Network | Release 24.2.1      | <p>Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR:</p> <ul style="list-style-type: none"> <li>• N540-ACC-SYS</li> <li>• N540X-ACC-SYS</li> <li>• N540-24Z8Q2C-SYS</li> </ul> <p>In Segment Routing over an IPv6 network (SRv6), you can keep track of the operational status of both the forward and reverse paths of a particular node or IP endpoint. You can use this information for troubleshooting, network maintenance, and optimizing network performance.</p> <p>Additionally, you can use flow labels to verify the liveness of each subsequent hop path toward the IP endpoint of that path. So that, when network traffic is distributed across multiple available paths towards an IP endpoint, liveness detection tracks the operational status of each of these paths towards the IP endpoint.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• The <b>reverse-path</b> and <b>segment-list name</b> keywords are introduced in the <b>segment-routing traffic-eng explicit</b> command.</li> <li>• The <b>source-address ipv6</b> is introduced in the <b>performance-measurement endpoint</b> command.</li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>• Cisco-IOS-XR-um-performance-measurement-cfg</li> <li>• Cisco-IOS-XR-perf-meas-oper.yang</li> </ul> <p>(see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</p> |
| IP Endpoint Delay Measurement and Liveness Monitoring | Release 7.4.1       | <p>This feature measures the end-to-end delay and monitors liveness of a specified IP endpoint node, including VRF-aware (awareness of multiple customers belonging to different VRFs).</p> <p>This feature is supported on IPv4, IPv6, and MPLS data planes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

The Segment Routing Performance Measurement (SR-PM) for IP endpoint liveness is a type of node liveness that involves testing whether an IP endpoint or a device identified by an IP address is available to send and receive data.

IP endpoint liveness is verified by sending a request to the IP address of the endpoint and waiting for a response. The probe could be an ICMP echo request (Ping), a TCP packet, a UDP packet, or any other type of packet that the endpoint would respond to.

- If a response is received, the endpoint is considered *live*.



- If no response is received within a certain time frame, the endpoint is considered *down* or *unreachable*.

IP endpoint dynamically measures the liveness towards a specified IP endpoint. IP endpoints can be located in a default or nondefault VRFs. IP endpoint is any device in the network a device identified by an IP address.

Liveness of an IP endpoint is verified by sending a request to the IP address of the endpoint and waiting for a response, which is referred to as a probe.

The endpoint of a probe is defined by an IP address, which can be either IPv4 or IPv6. This IP address can be any address that the sender can reach, such as a local interface or a remote node or host, either within an operator's network or accessible via a VRF.

The endpoint of a probe can be any IP address reachable by the sender. For example, a local interface or a remote node or host located within an operator's network or reachable through a VRF.

The IP address of the endpoint can be reached through an IP path, MPLS, LSP SRV6, or IP tunnel (GRE).

- When the endpoint is reachable using an MPLS LSP (for example, SR, LDP, RSVP-TE, SR Policy), the forwarding stage imposes the corresponding MPLS transport labels.
- When the endpoint is reachable via a GRE tunnel, the forwarding stage imposes the corresponding GRE header.
- When the endpoint is reachable via a VRF in an MPLS network, the forwarding stage imposes the corresponding MPLS service labels. In the forward path, the sender node uses the configured VRF for the endpoint address. In the return path, the reflector node derives the VRF based on which incoming VRF label the probe packet is received with.
- When the endpoint is reachable using SRv6, the forwarding stage imposes the SRv6 encapsulation.

You can configure the following parameters in the **performance-measurement** command:

- **Endpoint:** The endpoint of a probe is defined by an IP address, which can be either IPv4 or IPv6. This IP address can be any address that the sender can reach, such as a local interface or a remote node or host, either within an operator's network or accessible via a VRF.

The endpoint of a probe can be any IP address reachable by the sender. For example, a local interface or a remote node or host located within an operator's network or reachable through a VRF.

Use the **performance-measurement endpoint** command to configure a probe endpoint source and destination addresses on a sender node.

- **VRF:** You can define the endpoint point IP address belonging to a specific VRF. Use the **performance-measurement endpoint {ipv4 | ipv6} ip\_addr [vrf WORD]** command to configure an endpoint to define the VRF. Endpoint segment list configuration is not supported under nondefault VRF.
  - VRF-awareness allows operators to deploy probes in the following scenarios:
    - Managed Customer Equipment (CE) scenarios:
      - PE to CE probes
      - CE to CE probes
    - Unmanaged Customer Equipment (CE) scenarios:
      - PE to PE probes
      - PE to PE (source from PE-CE interface) probes

- **Source address:** You can define the source of the endpoint using the endpoint specific source address and the global source address.

Global source address configuration is applied to all the endpoints when the endpoint specific source address configuration isn't specified. endpoint specific configuration overrides all the global source address configuration for those specific endpoints for which source addresses are configured.

For Micro-SID configuration for IPv4 endpoint sessions, if IPv6 global source address is configured, then it applies the configured global IPv6 source address for the IPv6 header in the SRv6 packet. If IPv6 global address is not configured, then It does not form a valid SRv6 packet.

You can use the **source-address** keyword under the **performance-measurement** command to define the global source address or use the keyword under **performance-measurement endpoint** to define endpoint specific source address.

- **Reverse Path:** To detect the liveness of the reverse of the segment, you can configure the reverse path using the **reverse-path** command.

The default reverse path configured under the endpoint submode is only used for sessions with segment list. The endpoint session without a segment list does not support reverse path configuration and will not use this reverse path.

The **reverse-path** under the **performance-measurement endpoint** is used as the default reverse path if there are no reverse paths configured under the segment list.

Use the **reverse-path** under the **performance-measurement endpoint segment-routing traffic-eng explicit segment-list name** to configure the reverse path under segment list.

The reverse type must be the same as the forward path. Using different types for forward and reverse paths is not supported. For example, uSID forward path and uSID reverse path; MPLS forward path and MPLS reverse path.

User-configured segment-list can also represent the reverse path (reflector to sender) when probe is configured in liveness detection mode. Up to 128 segment-lists can be configured under a probe. An additional PM session is created for each segment-list. Segment-lists are configured under **segment-routing traffic-eng segment-list** submode. See [SR-TE Policy with Explicit Path](#) for details about configuring segment lists.

- **Flow Label:** The flow label field in the IPv6 header is used to carry information that helps distribute traffic across multiple network paths. The flow label is a 20-bit field in the IPv6 header designed to carry information about the flow of packets, which routers can use to identify and differentiate between different traffic flows. Flow label sweeping uses a flow label to distribute the traffic load across multiple paths to the endpoint.

Use the **flow-label** keyword to configure flow label.

### Usage Guidelines and Limitations

- Liveness session without segment list for an endpoint in a non-default VRF is not supported.
- Liveness on SRv6 is supported on Cisco NCS 5508 starting from Cisco IOS XR Release 24.3.1.
- IPv6 Endpoint Liveness in Default VRF over SRv6 is not supported.
- IP Endpoint in VRF without segment list, over SRv6 underlay (dynamic uDT6 or uDT4) is not supported
- SR Performance Measurement endpoint session over BVI interface is not supported.

## IP Endpoint Liveness Detection in an SR MPLS Network

IP endpoint liveness detection leverages the loopback measurement-mode. The following workflow describes the sequence of events.

1. The sender creates and transmits the PM probe packets.

The IP destination address (DA) on the probe packets is set to the loopback value of the sender itself.

The transmit timestamp (T1) is added to the payload.

The probe packet is encapsulated with the label corresponding to the endpoint.

2. The network delivers the PM probe packets following the LSP toward the endpoint.

3. The end-point receives the PM probe packets.

Packets are forwarded back to the sender based on the forwarding entry associated with the IP DA of the PM probe packet. If an LSP exists, the probe packet is encapsulated with the label of the sender.

4. The sender node receives the PM probe packets.


The received timestamp (T4) stored.

If the sender node doesn't receive the specified number of probe packets (based on the configured multiplier), the sender node declares the PM session as down.

The following figure illustrates a liveness detection probe toward an IP endpoint learned by the IGP. The network interconnecting the sender and reflector provides MPLS connectivity with Segment Routing.

The liveness detection multiplier is set to 5 to specify the number of consecutive missed probe packets before the PM session is declared as down.

**Figure 31: IP Endpoint Liveness Detection**



### Configuration Example

```
RouterA(config)# performance-measurement
RouterA(config-perf-meas)# endpoint ipv4 10.1.1.5
RouterA(config-pm-ep)# source-address ipv4 10.1.1.1
```

```

RouterA(config-pm-ep) # liveness-detection
RouterA(config-pm-ep-ld) # exit
RouterA(config-pm-ep) # exit
RouterA(config-perf-meas) # liveness-profile endpoint default
RouterA(config-pm-ld-ep) # liveness-detection
RouterA(config-pm-ld-ep-ld) # multiplier 5
RouterA(config-pm-ld-ep-ld) # exit
RouterA(config-pm-ld-ep) # probe
RouterA(config-pm-ld-ep-probe) # measurement-mode loopback

```

### Running Configuration

```

performance-measurement
 endpoint ipv4 10.1.1.5
 source-address ipv4 10.1.1.1
 liveness-detection
 !
 !
 liveness-profile endpoint default
 liveness-detection
 multiplier 5
 !
 probe
 measurement-mode loopback
 !
 !
 !
end

```

### Verification

```
RouterA# show performance-measurement endpoint ipv4 10.1.1.5
```

```

0/RSP0/CPU0

```

```

Endpoint name: IPv4-10.1.1.5-vrf-default
Source address : 10.1.1.1
VRF name : default
Liveness Detection : Enabled
Profile Keys:
 Profile name : default
 Profile type : Endpoint Liveness Detection

Segment-list : None
Session State: Down
Missed count: 0

```


## IP Endpoint Liveness in an SRv6 Network

IP endpoint liveness detection leverages the loopback measurement-mode. The following workflow describes the sequence of events.

1. The sender creates and transmits the PM probe packets and sets the IP destination address (DA) on the probe packets to the loopback value of the sender itself. The sender adds the transmit timestamp (T1) to the payload and encapsulates the probe packet with the SRv6 labels corresponding to the endpoint.
2. The network delivers the PM probe packets by following the IPv6 or segment lists towards the endpoint.
3. The endpoint receives the PM probe packets.

- The sender node receives the PM probe packets and stores the received timestamp. If the sender node doesn't receive the specified number of probe packets based on the configured multiplier, it declares the PM session down.

Figure 32: IP Endpoint Liveness In an SRv6 Network



**Note** Liveness is not supported for non-default VRF.

### Configuration Example

```
Router(config)#performance-measurement
Router(config-perf-meas)#source-address ipv6 2020:1::1
Router(config-perf-meas)#endpoint ipv6 FCBB:0::5::
Router(config-pm-ep)#exit
Router(config-perf-meas)#liveness-profile endpoint default
Router(config-pm-ld-ep)#probe
Router(config-pm-ld-ep-probe)#exit
Router(config-pm-ld-ep)#liveness-detection
Router(config-pm-ld-ep-ld)#multiplier 3
Router(config-pm-ld-ep-ld)#
```

The following example shows how to configure liveness with segment list and reverse path.

```
Router(config-sr)#traffic-eng
Router(config-sr-te)#segment-lists
Router(config-sr-te-segment-lists)#srv6
Router(config-sr-te-sl-global-srv6)#sid-format usid-f3216
Router(config-sr-te-sl-global-srv6)#exit
Router(config-sr-te-sl-global)#segment-list test
Router(config-sr-te-sl)#srv6
Router(config-sr-te-sl-srv6)#index 10 sid ff::2
Router(config-sr-te-sl-srv6)#index 20 sid ff::3
```

The following example shows how to configure liveness reverse path under segment list and under endpoint:

```
Router(config)#performance-measurement
Router(config-perf-meas)#endpoint ipv6 ff::2

/* Configure reverse path under segment list name */
Router(config-pm-ep)#segment-routing traffic-eng explicit segment-list name fwd-path
```

```

Router(config-pm-ep-sl)#reverse-path segment-list name rev-path
Router(config-pm-ep-sl)#exit

/* Configure reverse path under performance measurement endpoint */
Router(config-pm-ep)# segment-routing traffic-eng explicit reverse-path segment-list name
rev-path-name

```

The following example shows how to configure liveness with flow label:

```

Router(config-perf-meas)#liveness-profile endpoint default
Router(config-pm-ld-ep)#probe
Router(config-pm-ld-ep-probe)#flow-label from 1000 to 20000 increment 16
Router(config-pm-ld-ep-probe)#liveness-detection
Router(config-pm-ld-ep-ld)#multiplier 3

```

The following example shows how to configure liveness with flow label sweeping:

```

Router#configure
Router(config)#performance-measurement
Router(config-perf-meas)#liveness-profile name profile-sweeping
Router(config-pm-ld-profile)# flow-label from 1000 to 20000 increment 16
Router(config-pm-ld-profile)#commit

```

## Verification

```

Router# show performance-measurement endpoint detail
Endpoint name: IPv6-FCBB:0::5::-vrf-default
 Source address : 2020:1::1
 VRF name : default
 Liveness Detection : Enabled
 Profile Keys:
 Profile name : default
 Profile type : Endpoint Liveness Detection
 Segment-list : None
 Liveness Detection session:
 Session ID : 4109
 Flow-label : 1000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.214
 Missed count: 0

 Liveness Detection session:
 Session ID : 4110
 Flow-label : 2000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.214
 Missed count: 0

 Segment-list : test-dm-two-carrier-sl2
 FCBB:0::5:2:e004::/64
 Format: f3216
 FCBB:0::5:3:e000::/64
 Format: f3216
 FCBB:0::5:2:e004::/64
 Format: f3216
 FCBB:0::5:2:e000::/64
 Format: f3216
 FCBB:0::5:1:e000::/64
 Format: f3216
 FCBB:0::5:1:e004::/64
 Format: f3216
 FCBB:0::5:4:e000::/64
 Format: f3216
 FCBB:0::5:4::/48

```

```

Format: f3216
Liveness Detection session:
 Session ID : 4111
 Flow-label : 1000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.217
 Missed count: 0

```

```

Liveness Detection session:
 Session ID : 4112
 Flow-label : 2000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.217
 Missed count: 0

```

## SR Policy Liveness Monitoring

**Table 62: Feature History Table**

| Feature Name                                                      | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR Policy Liveness Monitoring on Segment Routing over IPv6 (SRv6) | Release 7.11.1      | <p>In segment routing over IPv6 (SRv6), you can now verify end-to-end traffic forwarding over an SR policy candidate path by periodically sending probe messages. Performance monitoring on an SRv6 network enables you to track and monitor traffic flows at a granular level.</p> <p>Earlier releases supported SR policy liveness monitoring over an SR policy candidate path on MPLS.</p> |

| Feature Name                              | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR Performance Measurement Named Profiles | Release 7.3.1       | <p>You can use this feature to create specific performance measurement delay and liveness profiles, and associate it with an SR policy.</p> <p>This way, a delay or liveness profile can be associated with a policy for which the performance measurement probes are enabled, and performance measurement is precise, and enhanced.</p> <p>The <b>performance-measurement delay-profile sr-policy</b> command was updated with the <b>name profile</b> keyword-argument combination.</p> <p>The <b>performance-measurement liveness-profile sr-policy</b> command was updated with the <b>name profile</b> keyword-argument combination.</p> <p>The <b>performance-measurement delay-measurement</b> command was updated with <b>delay-profile name profile</b>.</p> <p>The <b>performance-measurement liveness-detection</b> command was updated with <b>liveness-profile name profile</b>.</p> |
| SR Policy Liveness Monitoring             | Release 7.3.1       | <p>This feature allows you to verify end-to-end traffic forwarding over an SR Policy candidate path by periodically sending performance monitoring (PM) packets.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

SR Policy liveness monitoring allows you to verify end-to-end traffic forwarding over an SR Policy candidate path by periodically sending performance monitoring (PM) packets. The head-end router sends PM packets to the SR policy's endpoint router, which sends them back to the head-end without any control-plane dependency on the endpoint router.

The following are benefits to using SR-PM liveness monitoring:

- Allows both liveness monitoring and delay measurement using a single-set of PM packets as opposed to running separate monitoring sessions for each purpose. This improves the overall scale by reducing the number of PM sessions required.
- Eliminates network and device complexity by reducing the number of monitoring protocols on the network (for example, no need for Bidirectional Failure Detection [BFD]). It also simplifies the network and device operations by not requiring any signaling to bootstrap the performance monitoring session.




- Improves interoperability with third-party nodes because signaling protocols aren't required. In addition, it leverages the commonly supported TWAMP protocol for packet encoding.
- Improves liveness detection time because PM packets aren't punted on remote nodes
- Provides a common solution that applies to data-planes besides MPLS, including IPv4, IPv6, and SRv6.

### How it works?

The workflow associated with liveness detection over SR policy is described in the following sequence.

Consider an SR policy programmed at head-end node router 1 towards end-point node router 5. This SR policy is enabled for liveness detection using the loopback measurement-mode.



- **A:** The head-end node creates and transmits the PM probe packets.

The IP destination address (DA) on the probe packets is set to the loopback value of the head-end node itself.

A transmit (Tx) timestamp is added to the payload.

Optionally, the head-end node may also insert extra encapsulation (labels) to enforce the reverse path at the endpoint node.

Finally, the packet is injected into the data-plane using the same encapsulation (label stack) of that of the SR policy being monitored.

- **B:** The network delivers the PM probe packets as it would user traffic over the SR policy.
- **C:** The end-point node receives the PM probe packets.

Packets are switched back based on the forwarding entry associated with the IP DA of the packet. This would typically translate to the end-point node pushing the prefix SID label associated with the head-end node.


If the head-end node inserted label(s) for the reverse path, then the packets are switched back at the end-point node based on the forwarding entry associated with the top-most reverse path label.

- **D:** Headend node receives the PM probe packets.

A received (Rx) timestamp stored.

If the head-end node receives the PM probe packets, the head-end node assume that the SR policy active candidate path is up and working.

If the head-end node doesn't receive the specified number of consecutive probe packets (based on configured multiplier), the head-end node assumes the candidate path is down and a configured action is triggered.



### Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- SR-PM liveness-detection over SR Policy is supported on manually configured SR Policies and On-Demand SR Policies (ODN).
- SR-PM liveness-detection over SR Policy is not supported on PCE-initiated SR Policies.
- SR-PM liveness-detection and delay-measurement aren't supported together
- When liveness-profile isn't configured, SR Policies use the default values for the liveness-detection profile parameters.
- The head-end router doesn't load-balance the liveness probes across bundle member links.

## Configure SR Policy Liveness Monitoring in an MPLS Network

Configuring SR Policy liveness monitoring involves the following steps:

- Configuring a performance measurement liveness profile to customize generic probe parameters
- Enabling liveness monitoring under SR Policy by associating a liveness profile, and customizing SR policy-specific probe parameters

Liveness monitoring parameters are configured under **performance-measurement liveness-profile** sub-mode. The following parameters are configurable:

- liveness-profile sr-policy {default | name name}**

Parameters defined under the **sr-policy default** liveness-profile apply to any SR policy with liveness monitoring enabled and that does not reference a non-default (named) liveness-profile.

- probe:** Configure the probe parameters.

- **measurement-mode:** Liveness detection must use loopback mode (see [Measurement Modes, on page 464](#)).
- **burst interval:** Interval for sending probe packet. The default value is 3000 milliseconds and the range is from 30 to 15000 milliseconds.
- **tos dscp value:** The default value is 48 and the range is from 0 to 63. You can modify the DSCP value of the probe packets, and use this value to prioritize the probe packets from headend to tailend.
- **sweep destination ipv4 127.x.x.x range range:** Configure SR Policy ECMP IP-hashing mode. Specify the number of IP addresses to sweep. The range is from 0 (default, no sweeping) to 128. The option is applicable to IPv4 packets.



**Note** The destination IPv4 headend address 127.x.x.x – 127.y.y.y is used in the Probe messages to take advantages of 3-tuple IP hashing (source-address, destination-address, and local router ID) for ECMP paths of SR-MPLS Policy. The destination IPv4 address must be 127/8 range (loopback), otherwise it will be rejected.



**Note** One PM session is always created for the actual endpoint address of the SR Policy.

- **liveness-detection:** Configure the liveness-detection parameters:
- **multiplier:** Number of consecutive missed probe packets before the PM session is declared as down. The range is from 2 to 10, and the default is 3.



**Note** The detection-interval is equal to (burst-interval \* multiplier).

### Enabling Liveness Monitoring under SR Policy

Enable liveness monitoring under SR Policy, associate a liveness-profile, and configure SR Policy-specific probe parameters under the **segment-routing traffic-eng policy performance-measurement** sub-mode. The following parameters are configurable:

- **liveness-detection:** Enables end-to-end SR Policy Liveness Detection for all segment-lists of the active and standby candidate-path that are in the forwarding table.
- **liveness-profile name name:** Specifies the profile name for named profiles.
- **invalidation-action {down | none}:**
  - **Down (default):** When the PM liveness session goes down, the candidate path is immediately operationally brought down.
  - **None:** When the PM liveness session goes down, no action is taken. If logging is enabled, the failure is logged but the SR Policy operational state isn't modified.

- **logging session-state-change**: Enables Syslog messages when the session state changes.
- **reverse-path label** {*BSID-value* | *NODE-SID-value*}: Specifies the MPLS label to be used for the reverse path for the reply. If you configured liveness detection with ECMP hashing, you must specify the reverse path. The default reverse path uses IP Reply.
  - *BSID-value*: The Binding SID (BSID) label for the reverse SR Policy. (This is practical for manual SR policies with a manual BSID.)
  - *NODE-SID-value*: The absolute SID label of the (local) Sender Node to be used for the reverse path for the reply.

## Configuration Examples

### Configure a Default SR-Policy PM Liveness-Profile

The following example shows a default sr-policy liveness-profile:

```
RP/0/RSP0/CPU0:ios(config)# performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)# liveness-profile sr-policy default
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# probe
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# measurement-mode loopback
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# burst-interval 1500
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# tos dscp 52
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# liveness-detection
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-ld)# multiplier 5
```

#### Running Configuration:

```
performance-measurement
 liveness-profile sr-policy default
 liveness-detection
 multiplier 5
 !
 probe
 tos dscp 52
 measurement-mode loopback
 burst-interval 1500
 !
 !
end
```

### Configure a Named (Non-Default) SR-Policy PM Liveness-Profile

The following example shows a named sr-policy liveness-profile:

```
RP/0/RSP0/CPU0:ios(config)# performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)# liveness-profile name sample-profile
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# probe
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# measurement-mode loopback
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# burst-interval 1500
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# tos dscp 52
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# liveness-detection
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-ld)# multiplier 5
```

#### Running Configuration:

```

performance-measurement
 liveness-profile sr-policy name sample-profile
 liveness-detection
 multiplier 5
 !
 probe
 tos dscp 52
 measurement-mode loopback
 burst-interval 1500
 !
!
end

```

### Configure a SR-Policy PM Liveness-Profile with Sweep Parameters

The following example shows a named liveness-profile with sweep parameters:

```

RP/0/RSP0/CPU0:ios(config)# performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)# liveness-profile name sample-profile
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# probe
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# measurement-mode loopback
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# burst-interval 1500
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# tos dscp 52
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# sweep
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe-sweep)# destination ipv4 127.0.0.1 range 25
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe-sweep)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-probe)# exit
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy)# liveness-detection
RP/0/RSP0/CPU0:ios(config-pm-ld-srpolicy-ld)# multiplier 5

```

### Running Configuration

```

performance-measurement
 liveness-profile sr-policy name sample-profile
 liveness-detection
 multiplier 5
 !
 probe
 tos dscp 52
 sweep
 destination ipv4 127.0.0.1 range 25
 !
 measurement-mode loopback
 burst-interval 1500
 !
!
end

```

### Enable Liveness Monitoring Under SR Policy

The following example shows how to enable liveness monitoring under SR Policy, associate a liveness-profile, and configure the invalidation action:

```

RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# policy FOO
RP/0/RSP0/CPU0:ios(config-sr-te-policy)# performance-measurement
RP/0/RSP0/CPU0:ios(config-sr-te-policy-perf-meas)# liveness-detection
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# liveness-profile name sample-profile
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# invalidation-action none

```

### Running Config

```

segment-routing
 traffic-eng

```

```

policy FOO
performance-measurement
liveness-detection
liveness-profile name sample-profile
invalidation-action none
!
!
!
!
end

```

### Enable Liveness Monitoring under SR Policy with Optional Parameters

The following example shows how to enable liveness monitoring under SR Policy, associate a liveness-profile, and configure reverse path label and session logging:

```

RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# policy BAA
RP/0/RSP0/CPU0:ios(config-sr-te-policy)# performance-measurement
RP/0/RSP0/CPU0:ios(config-sr-te-policy-perf-meas)# liveness-detection
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# liveness-profile name sample-profile
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# invalidation-action down
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# logging session-state-change
RP/0/RSP0/CPU0:ios(config-sr-te-policy-live-detect)# exit
RP/0/RSP0/CPU0:ios(config-sr-te-policy-perf-meas)# reverse-path label 16001

```

### Running Config

```

segment-routing
traffic-eng
policy BAA
performance-measurement
liveness-detection
logging
session-state-change
!
liveness-profile name sample-profile
invalidation-action down
!
reverse-path
label 16001
!
!
!
!
end

```

## Configure Segment Lists to Activate Candidate Paths in SRv6 for PM Liveness

Table 63: Feature History Table

| Feature Name                                                                | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Segment Lists to Activate Candidate Paths in SRv6 for PM Liveness | Release 7.11.1      | <p>You can now enable a candidate path to be up by configuring the minimum number of active segment lists associated with the candidate path. The head-end router determines that a candidate path is up based on the minimum number of active segment lists configured.</p> <p>In earlier releases, the head-end router identified a candidate path as up only when all the segment lists associated with the path were active.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>validation-cp minimum-active segment-lists</b> option is introduced in the <b>performance-measurement liveness-detection</b> command.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>Cisco-IOS-XR-infra-xtc-agent-cfg.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

The state of the segment lists in a candidate path determines whether a candidate path is up or down. You can now configure the minimum number of active segment lists associated with a candidate path. The head-end router identifies a candidate path as up when one or more segment lists are active.



**Note** If the configured minimum number of active segment lists is greater than the number of available segment lists in a candidate path, the head-end router determines the candidate path as up only when all the segment lists are active.

In earlier releases, the router identified a candidate path as up only when all the segment lists associated with the path were active.

## Configuration Example

### Configure the minimum number of segment lists in SRv6

Perform this task to activate three segment lists to have the PM liveness session up:

```
Router(config)#segment-routing
Router(config-sr)#traffic-eng
Router(config-sr-te)#policy po-103
Router(config-sr-te-policy)#performance-measurement
Router(config-sr-te-policy-perf-meas)#liveness-detection
Router(config-sr-te-policy-live-detect)#validation-cp minimum-active segment-lists 3
```

### Show Running Configuration

```
segment-routing
 traffic-eng
 policy po-103
 performance-measurement
 liveness-detection
 validation-cp minimum-active segment-lists 3
 !
!
!
!
!
```

### Verification

The following example shows three active segment-lists to have the PM liveness session up:

```
Router#show performance-measurement sr-policy liveness color 103 detail verbose private
Mon Oct 30 15:10:51.863 EDT
```

---

```
0/1/CPU0
```

---

```
SR Policy name: srte_c_103_ep_3::1
 Color : 103
 SRv6 Encap Source Address : 1::1
 Endpoint : 3::1
 Handle : 0x00000000
 Policy to be deleted : False
 Number of candidate-paths : 1

Candidate-Path:
 Instance : 5
 Preference : 300
 Protocol-origin : Configured
 Discriminator : 300
 Profile Keys:
 Profile name : default
 Profile type : SR Policy Liveness Detection
 Candidate path to be deleted: False
 Source address : 1::1
 Local label : Not set
 Fast notification for session down: Disabled
 No fast notifications have been sent
 Number of segment-lists : 3
 Liveness Detection: Enabled
 Minumum SL Up Required: 1
 Session State: Up
 Last State Change Timestamp: Oct 30 2023 15:10:16.322
```



Missed count: 0

**Segment-List : sl-1041**  
 fccc:cc00:1:fe10:: (Local Adjacency SID)  
 fccc:cc00:2:fe41::/64  
 Format: f3216  
 Segment List ID: 0  
 Reverse path segment-List: Not configured  
 Segment-list to be deleted: False  
 Number of atomic paths : 1  
**Liveness Detection: Enabled**  
**Session State: Up**  
 Last State Change Timestamp: Oct 30 2023 15:10:16.322  
 Missed count: 0

Atomic path:  
 Flow Label : 0  
 Session ID : 4198  
 Trace ID : 738913600  
 Atomic path to be deleted: False  
 NPU Offloaded session : False  
 Timestamping Enabled : True  
 Liveness Detection: Enabled  
 Session State: Up  
 Last State Change Timestamp: Oct 30 2023 15:10:16.322  
 Missed count: 0  
 Responder IP : 1::1  
 Number of Hops : 3

**Segment-List : sl-1042**  
 fccc:cc00:1:fe10:: (Local Adjacency SID)  
 fccc:cc00:2:fe42::/64  
 Format: f3216  
 Segment List ID: 0  
 Reverse path segment-List: Not configured  
 Segment-list to be deleted: False  
 Number of atomic paths : 1  
**Liveness Detection: Enabled**  
**Session State: Up**  
 Last State Change Timestamp: Oct 30 2023 15:10:16.322  
 Missed count: 0

Atomic path:  
 Flow Label : 0  
 Session ID : 4199  
 Trace ID : 954039677  
 Atomic path to be deleted: False  
 NPU Offloaded session : False  
 Timestamping Enabled : True  
 Liveness Detection: Enabled  
 Session State: Up  
 Last State Change Timestamp: Oct 30 2023 15:10:16.322  
 Missed count: 0  
 Responder IP : 1::1  
 Number of Hops : 3

**Segment-List : sl-1043**  
 fccc:cc00:1:fe10:: (Local Adjacency SID)  
 fccc:cc00:2:fe43::/64  
 Format: f3216  
 Segment List ID: 0  
 Reverse path segment-List: Not configured  
 Segment-list to be deleted: False  
 Number of atomic paths : 1

```

Liveness Detection: Enabled
Session State: Up
Last State Change Timestamp: Oct 30 2023 15:10:16.322
Missed count: 0

Atomic path:
Flow Label : 0
Session ID : 4200
Trace ID : 1119107116
Atomic path to be deleted: False
NPU Offloaded session : False
Timestamping Enabled : True
Liveness Detection: Enabled
Session State: Up
Last State Change Timestamp: Oct 30 2023 15:10:16.322
Missed count: 0
Responder IP : 1::1
Number of Hops : 3

```

0/RSP0/CPU0

## Configure Flow Labels in SRv6 Header for PM Liveness

**Table 64: Feature History Table**

| Feature Name                                         | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Flow Labels in SRv6 Header for PM Liveness | Release 7.11.1      | <p>You can now monitor the activeness of multiple paths for a given segment list using flow labels in the SRv6 header.</p> <p>In earlier releases, the SRv6 header didn't include flow labels.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>flow-label</b> keyword is introduced in the <b>performance-measurement liveness-profile</b> command.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>Cisco-IOS-XR-performance-meas-cfg.yang</li> <li>Cisco-IOS-XR-perf-meas-oper.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

To monitor the activeness of multiple paths for a given a segment list, you can configure the SRv6 header to include flow labels as the packet travels in the network. When there are multiple paths, different traffic flows

may use different paths. A flow label is a flow identifier and you can use different flow labels to monitor different ECMP paths. It's only used for IPv6 probe packets. Flow labels are 20-bit fields in the SRv6 header.

### Configure flow labels in the SRv6 header

Perform the following task in the global configuration mode to configure flow labels in the SRv6 header:

```
Router#configure
Router(config)#performance-measurement
Router(config-perf-meas)#liveness-profile name name1
Router(config-pm-ld-profile)#probe flow-label from 0 to 1000000 increment 10
```

### Running Configuration

```
performance-measurement
 liveness-profile name name1
 probe
 flow-label from 0 to 1000000 increment 10
 !
!
!
```

### Verification

The following example shows an SR-policy configured with flow labels:

```
Router#show performance-measurement sr-policy liveness color 1001 detail verbose private
```

```
Mon Oct 30 15:25:55.241 EDT
```

---

```
0/1/CPU0
```

---

```
SR Policy name: srte_c_1001_ep_3::1
Color : 1001
SRv6 Encap Source Address : 1::1
Endpoint : 3::1
Handle : 0x00000000
Policy to be deleted : False
Number of candidate-paths : 1

Candidate-Path:
 Instance : 3
 Preference : 300
 Protocol-origin : Configured
 Discriminator : 300
 Profile Keys:
 Profile name : profile-scale
 Profile type : Generic Liveness Detection
 Candidate path to be deleted: False
 Source address : 1::1
 Local label : Not set
 Fast notification for session down: Disabled
 No fast notifications have been sent
 Number of segment-lists : 2
 Liveness Detection: Enabled
 Minumum SL Up Required: 2
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0

Segment-List : sl-1041
```

```

fccc:cc00:1:fe10:: (Local Adjacency SID)
fccc:cc00:2:fe41::/64
 Format: f3216
Segment List ID: 0
Reverse path segment-List: Not configured
Segment-list to be deleted: False
Number of atomic paths : 2
Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0

Atomic path:
 Flow Label : 0
 Session ID : 4178
 Trace ID : 280178832
 Atomic path to be deleted: False
 NPU Offloaded session : False
 Timestamping Enabled : True
 Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0
 Responder IP : 1::1
 Number of Hops : 3

Atomic path:
 Flow Label : 10
 Session ID : 4179
 Trace ID : 1866227171
 Atomic path to be deleted: False
 NPU Offloaded session : False
 Timestamping Enabled : True
 Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0
 Responder IP : 1::1
 Number of Hops : 3

Segment-List : sl-scale
fccc:cc00:1:fe10:: (Local Adjacency SID)
fccc:cc00:2:fed1::/64
 Format: f3216
Segment List ID: 0
Reverse path segment-List: Not configured
Segment-list to be deleted: False
Number of atomic paths : 2
Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0

Atomic path:
 Flow Label : 0
 Session ID : 4180
 Trace ID : 2609815826
 Atomic path to be deleted: False
 NPU Offloaded session : False
 Timestamping Enabled : True
 Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0

```

```

Responder IP : 1::1
Number of Hops : 3

Atomic path:
 Flow Label : 10
 Session ID : 4181
 Trace ID : 170501506
 Atomic path to be deleted: False
 NPU Offloaded session : False
 Timestamping Enabled : True
 Liveness Detection: Enabled
 Session State: Up
 Last State Change Timestamp: Oct 26 2023 15:31:43.478
 Missed count: 0
 Responder IP : 1::1
 Number of Hops : 3

```

0/RSP0/CPU0

## Hardware offload of SR policy liveness monitoring

**Table 65: Feature History Table**

| Feature Name                                 | Release        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware offload of SRv6 liveness monitoring | Release 25.2.1 | <p>You can now offload SRv6 liveness monitoring for performance measurement to the router's hardware, which is the Network Processing Unit (NPU). This hardware-based approach improves efficiency and scalability, helping you meet delay-sensitive Service Level Agreements (SLAs). Previously, this monitoring was handled in software.</p> <p>Using hardware to offload performance monitoring tasks improves efficiency and reduce the load on the main processor.</p> <p>You can enable by using the existing <b>npu-offload</b> under the <b>performance-measurement liveness-profile name</b><i>liveness profile</i> command.</p> |
| Hardware Offload of MPLS Liveness Monitoring | Release 7.10.1 | <p>You can now offload MPLS liveness monitoring for performance measurement to the router's hardware, which is the Network Processing Unit (NPU). This hardware-based approach improves efficiency and scalability, helping you meet delay-sensitive Service Level Agreements (SLAs). Previously, this monitoring was handled in software.</p> <p>The feature introduces a new keyword <b>npu-offload</b> under the <b>performance-measurement liveness-profile name</b><i>liveness profile</i> command.</p>                                                                                                                              |

Performance Measurement (PM) hardware offload feature allows the offload of PM liveness monitoring session to the Network Processing Unit (NPU) on the platform, which considerably improves scale and reduces the overall network convergence detection time.

This improvement is done by sending rapid failure detection probes (messages) and detecting policy or path failures quickly and help routing protocols in recalculating the routing table.

This feature is required in order to quickly react on delay-bound Service Level Agreement (SLAs), for example 5G low-latency, where SRTE policy can quickly re-optimize once the SLA is violated.

Advantages of the PM Hardware Offloading feature are as listed:

- Probes are sent every 3.3 milliseconds
- Complete liveness of the endpoint is now reduced to 10ms from 50ms when the operator configures the multiplier to be 3 ( $10\text{ms} = 3.3\text{ms} * 3$ ).
- Currently, the hardware offload supports only liveness monitoring .



**Note** The hardware offload does not support delay and loss measurement yet.

## Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- The NPU offload generates PM probes (Packets Per Second) with a maximum limit. PPS is directly proportional with the number of sessions and transmit interval. If the PPS exceeds the limit supported by the offload engine, the stretch algorithm activates. This algorithm doubles the transmit interval until the PPS is within the supported limit. Use 'show performance-measurement pps' command to verify the maximum probes per second (PPS) supported by offload engine and total pps currently in use.

## Configuration Example



**Note** The **hw-module profile offload 4** command is a prerequisite for LC CPU sessions to work.

Once you use the **hw-module profile offload 4** command, the Bidirectional Forwarding Detection for IPv6 (BFDv6) at the router will not work even if the Performance Measurement sessions are hosted only on LC CPU and not offloaded to the offload processor.

Use the 4th option **4 PM-HW-Offload and Bsync** in **hw-module profile offload** command to configure the Hardware Profile Offloading.

```
Router(config)#hw-module profile offload ?
1 BFDv6 and Bsync
2 BFDv6 and Route download
3 Route download and Bsync
4 PM-HW-Offload and Bsync
```

Reload the router to apply the hardware offload profile. After reloading, the Performance Measurement application loads on the offload engine.



**Note** For SRv6 Liveness Monitoring - Hardware Offloading configuration, SRv6 uSID must be enabled. Use the following configuration to enable SRv6 uSID must be enabled

```
Router(config)# hw-module profile segment-routing srv6 mode micro-segment format f3216
```

The following example allows you to enable Performance Measurement Liveness Hardware (NPU) offload in the SR environment.



**Note** The configuration also applies to the default liveness profile and the configuration applies to the router.

```
Router(config)#performance-measurement
Router(config-perf-meas)#liveness-profile name hwo_profile
Router(config-pm-ld-profile)#npu-offload
Router(config-pm-ld-profile-npu-offload)#commit
```

### Running Configuration

The running configuration for this feature is as shown for both SRv6 and MPLS:

```
performance-measurement
 liveness-profile name hwo_profile
 npu-offload
 !
 !
 !
```

### Verification

Use the show command to verify the running configuration as shown for both MPLS and SRv6 Liveness Monitoring - Hardware Offloading :

#### • Example for MPLS

```
Router# show performance-measurement sessions detail

Transport type : SR Policy
Measurement type : Liveness Detection
Policy name : srte_c_90005_ep_10.2.2.2
Color : 90005
Endpoint : 10.2.2.2
Instance : 7
preference : 20
Protocol-origin : Configured
Discriminator : 20
Segment-list : route_12_2
Atomic path:
Hops : 10.2.2.2
Session ID : 45
Trace ID : 3111803555
NPU Offloaded session : True
NPU number : 0
NPU session state : Session created
Retry count : 0
```

```

Last NPU notification:
Session state : Up
Timestamp : Feb 28 2023 16:28:09.411
Timestamping Enabled : True
Liveness Detection: Enabled
Session State: Up
Last State Change Timestamp: Feb 28 2023 16:28:09.411
Missed count: 0

```

#### • Example for SRv6

```

Transport type : Endpoint
Measurement type : Liveness Detection
Endpoint name : IPv6-fccc:cccl:4::-vrf-default
endpoint : fccc:cccl:4::
source : 192::2
vrf : default
Segment-list : test-s11
Liveness Detection: Enabled
 Session ID : 3
 Profile Keys:
 Profile name : default
 Profile type : Endpoint Liveness Detection
 Session State: Down
 Missed count: 3
NPU Offloaded session : True
NPU number : 0
NPU session state : Session created
 Retry count : 0

```

## Delay Measurement

Delay measurement is a mechanism used to measure the latency or delay experienced by data packets when they traverse a network.

The PM for delay measurement uses the IP/UDP packet format defined in for probes. Two-Way Active Measurement Protocol (TWAMP) adds two-way or round-trip measurement capabilities. TWAMP employs time stamps applied at the echo destination (reflector) to enable greater accuracy. In the case of TWAMP Light, the Session-Reflector doesn't necessarily know about the session state. The Session-Reflector simply copies the Sequence Number of the received packet to the Sequence Number field of the reflected packet. The controller receives the reflected test packets and collects two-way metrics. This architecture allows for collection of two-way metrics.

### Benefits

- **Network Troubleshooting:** You can quickly and easily identify areas in your network with high delay and resolve network problems using delay measurement.
- **Network Planning and Optimization:** You can easily understand the performance of your network under various conditions and design a network that can handle expected traffic loads.
- **Quality of Service (QoS):** You can ensure quality of service standards are being met by continuously monitoring the delay in your network.



### Supported Delay Measurement Methods

You can measure delay using the following methods:

- [Link Delay Measurement, on page 467](#) Use to monitor delay experienced by data packets in a single link or path between two nodes in a network.
- [IP endpoint delay measurement](#): Use to monitor the amount of time it takes for a data packet to travel from a source device to a specific IP endpoint within a network.
- [SR Policy End-to-End Delay Measurement , on page 497](#): Use to to monitor the end-to-end delay experienced by the traffic sent over an SR policy.

## Measurement Modes

The following table compares the different hardware and timing requirements for the measurement modes that are supported in SR PM.

**Table 66: Measurement Mode Requirements**


| Measurement Mode | Sender:<br>PTP-Capable HW and HW<br>Timestamping | Reflector:<br>PTP-Capable HW and HW<br>Timestamping | PTP Clock Synchronization<br>between Sender and Reflector |
|------------------|--------------------------------------------------|-----------------------------------------------------|-----------------------------------------------------------|
| One-way          | Required                                         | Required                                            | Required                                                  |
| Two-way          | Required                                         | Required                                            | Not Required                                              |
| Loopback         | Required                                         | Not Required                                        | Not Required                                              |

### One-Way Measurement Mode

One-way measurement mode provides the most precise form of one-way delay measurement. PTP-capable hardware and hardware timestamping are required on both Sender and Reflector, with PTP Clock Synchronization between Sender and Reflector.

Delay measurement in one-way mode is calculated as  $(T2 - T1)$ .

Figure 33: One-Way




The PM query and response for one-way delay measurement can be described in the following steps:

1. The local-end router sends PM query packets periodically to the remote side once the egress line card on the router applies timestamps on packets.
2. The ingress line card on the remote-end router applies time-stamps on packets as soon as they are received.
3. The remote-end router sends the PM packets containing time-stamps back to the local-end router.
4. One-way delay is measured using the time-stamp values in the PM packet.

#### Far-end delay metrics in one-way measurement mode

Table 67: Feature History Table

| Feature Name                                      | Release Information | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Far-end delay metrics in one-way measurement mode | Release 24.4.1      | <p>Segment Routing Performance Monitoring (SR PM) now enables network operators to compute both far-end (T4 – T3) and near-end (T2 – T1) delay metrics, offering a comprehensive view of end-to-end delay across the data path. Measuring far-end delay, from the responder to the querier node, enhances visibility and allows operators to precisely monitor and assess network performance.</p> <p>Previously, you could measure the near-end delay metrics for a given data path.</p> |



The far-end delay metric measures the round-trip time a packet takes to travel from the source to the destination across the Segment Routing network. It calculates the time it takes for the packet to reach the far end of a segment and return to the source.

Starting from Cisco IOS XR Release 24.4.1, the PM session automatically calculates both near-end ( $T2 - T1$ ) and far-end ( $T4 - T3$ ) metrics. The far-end delay is the delay from the responder to the querier node. This metric is automatically included in the CLI and telemetry outputs.

The far-end metrics provide network operators a complete view of delay across the entire data path, and not just the near-end. The feature enables network operators to understand the full path latency and ensures accurate network monitoring.




**Note** If the measurement mode is not set to one-way or is set to loopback, SR PM skips the far-end delay metric computation.

### Two-Way Measurement Mode

Two-way measurement mode provides two-way measurements. PTP-capable hardware and hardware timestamping are required on both Sender and Reflector, but PTP clock synchronization between Sender and Reflector is not required.

Delay measurement in two-way mode is calculated as  $((T4 - T1) - (T3 - T2))/2$

**Figure 34: Two-Way**



The PM query and response for two-way delay measurement can be described in the following steps:

1. The local-end router sends PM query packets periodically to the remote side once the egress line card on the router applies timestamps on packets.
2. Ingress line card on the remote-end router applies time-stamps on packets as soon as they are received.
3. The remote-end router sends the PM packets containing time-stamps back to the local-end router. The remote-end router time-stamps the packet just before sending it for two-way measurement.
4. The local-end router time-stamps the packet as soon as the packet is received for two-way measurement.
5. Delay is measured using the time-stamp values in the PM packet.


### Loopback Measurement Mode

Loopback measurement mode provides two-way and one-way measurements. PTP-capable hardware and hardware timestamping are required on the Sender, but are not required on the Reflector.

Delay measurements in Loopback mode are calculated as follows:

- Round-Trip Delay =  $(T4 - T1)$
- One-Way Delay = Round-Trip Delay/2

**Figure 35: Loopback**



The PM query and response for Loopback delay measurement can be described in the following steps:

1. The local-end router sends PM probe packets periodically on the SR Policy.
2. The probe packets are loopback on the endpoint node (not punted), with no timestamping on endpoint node.
3. Round-trip Delay =  $T4 - T1$ .

## Link Delay Measurement



**Note** From Cisco IOS XR Release 7.6.1 onwards, routers support the following features:

- Link Delay Measurement
- Named Profiles
- Static Delay Value on an Interface

Table 68: Feature History Table

| Feature Name                                        | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Delay Measurement with IPv6 Link Local Address | Release 7.3.1       | The performance measurement for link delay determines the source and destination IP addresses used in the OAM packet based on the IP address of the interface, where the delay measurement operation is enabled. This feature enables using the IPv6 link-local address as the OAM packet source IP address, when no IPv4 or IPv6 address is configured in the interface. |

The PM for link delay uses the IP/UDP packet format defined in RFC 5357 (TWAMP-Light) for probes. Two-Way Active Measurement Protocol (TWAMP) adds two-way or round-trip measurement capabilities. TWAMP employs time stamps applied at the echo destination (reflector) to enable greater accuracy. In the case of TWAMP Light, the Session-Reflector doesn't necessarily know about the session state. The Session-Reflector simply copies the Sequence Number of the received packet to the Sequence Number field of the reflected packet. The controller receives the reflected test packets and collects two-way metrics. This architecture allows for collection of two-way metrics.

### Usage Guidelines and Restrictions for PM for Link Delay

The following restrictions and guidelines apply for the PM for link delay feature for different links.

- For broadcast links, only point-to-point (P2P) links are supported. P2P configuration on IGP is required for flooding the value.
- For link bundles, the hashing function may select a member link for forwarding but the reply may come from the remote line card on a different member link of the bundle.
- For one-way delay measurement, clocks should be synchronized on two end-point nodes of the link using PTP.
- Link delay measurement is supported on IPv4 unnumbered interfaces. An IPv4 unnumbered interface is identified by a node ID (a loopback address) and the local SNMP index assigned to the interface. Note that the reply messages could be received on any interface, since the packets are routed at the responder based on the loopback address used to identify the link.

### Configuration Example: PM for Link Delay

This example shows how to configure performance-measurement functionalities for link delay as a global default profile. The default values for the different parameters in the PM for link delay is given as follows:

- **probe measurement mode:** The default measurement mode for probe is two-way delay measurement. If you are configuring one-way delay measurement, hardware clocks must be synchronized between the local-end and remote-end routers using precision time protocol (PTP). See [Measurement Modes, on page 464](#) for more information.
- **protocol:** Interface delay measurement using RFC 5357 with IP/UDP encaps (TWAMP-Light).

- **burst interval:** Interval for sending probe packet. The default value is 3000 milliseconds and the range is from 30 to 15000 milliseconds.
- **computation interval:** Interval for metric computation. Default is 30 seconds; range is 1 to 3600 seconds.
- **periodic advertisement:** Periodic advertisement is enabled by default.
- **periodic-advertisement interval:** The default value is 120 seconds and the interval range is from 30 to 3600 seconds.
- **periodic-advertisement threshold:** Checks the minimum-delay metric change for threshold crossing for periodic advertisement. The default value is 10 percent and the range is from 0 to 100 percent.
- **periodic-advertisement minimum change:** The default value is 1000 microseconds (usec) and the range is from 0 to 100000 microseconds.
- **accelerated advertisement:** Accelerated advertisement is disabled by default.
- **accelerated-advertisement threshold:** Checks the minimum-delay metric change for threshold crossing for accelerated advertisement. The default value is 20 percent and the range is from 0 to 100 percent.
- **accelerated-advertisement minimum change:** The default value is 500 microseconds and the range is from 0 to 100000 microseconds.

```
RP/0/0/CPU0:router(config)# performance-measurement delay-profile interfaces
RP/0/0/CPU0:router(config-pm-dm-intf)# probe
RP/0/0/CPU0:router(config-pm-dm-intf-probe)# measurement-mode one-way
RP/0/0/CPU0:router(config-pm-dm-intf-probe)# burst-interval 60
RP/0/0/CPU0:router(config-pm-dm-intf-probe)# computation-interval 60
RP/0/0/CPU0:router(config-pm-dm-intf-probe)# exit

RP/0/0/CPU0:router(config-pm-dm-intf)# advertisement periodic
RP/0/0/CPU0:router(config-pm-dm-intf-adv-per)# interval 120
RP/0/0/CPU0:router(config-pm-dm-intf-adv-per)# threshold 20
RP/0/0/CPU0:router(config-pm-dm-intf-adv-per)# minimum-change 1000
RP/0/0/CPU0:router(config-pm-dm-intf-adv-per)# exit

RP/0/0/CPU0:router(config-pm-dm-intf)# advertisement accelerated
RP/0/0/CPU0:router(config-pm-dm-intf-adv-acc)# threshold 30
RP/0/0/CPU0:router(config-pm-dm-intf-adv-acc)# minimum-change 1000
RP/0/0/CPU0:router(config-pm-dm-intf-adv-per)# exit
```

### Configure the UDP Destination Port

Configuring the UDP port for TWAMP-Light protocol is optional. By default, PM uses port 862 as the TWAMP-reserved UDP destination port for delay.

The UDP port is configured for each PM measurement probe type (delay, loss, protocol, authentication mode, etc.) on querier and responder nodes. If you configure a different UDP port, the UDP port for each PM measurement probe type must match on the querier and the responder nodes.



**Note** The same UDP destination port is used for delay measurement for links and SR Policy.

This example shows how to configure the UDP destination port for delay.

```
Router(config)# performance-measurement
```

```
Router(config-perf-meas) # protocol twamp-light
Router(config-pm-protocol) # measurement delay unauthenticated
Router(config-pm-PROTO-mode) # querier-dst-port 12000
```

### Enable PM for Link Delay Over an Interface

This example shows how to enable PM for link delay over an interface.

```
RP/0/0/CPU0:router(config) # performance-measurement
RP/0/0/CPU0:router(config-perf-meas) # interface TenGigE0/0/0/0
RP/0/0/CPU0:router(config-pm-intf) # next-hop ipv4 10.10.10.2 // Optional IPv4 or IPv6
next-hop address
RP/0/0/CPU0:router(config-pm-intf) # delay-measurement
RP/0/0/CPU0:router(config-pm-intf-dm) # exit
```

The source and destination IP addresses used in the OAM packet are determined by the IP address present on the interface where the delay-measurement operation is enabled and the setting of the optional **next-hop** address.

When the **next-hop** address is not specified, the following rules apply to determine the source and destination IP addresses used in the OAM packet:

- If an IPv4 address is configured under the interface, then:
  - OAM packet source IP address = Interface's IPv4 address
  - OAM packet destination IP address = 127.0.0.0
- Else, if an IPv6 global address is configured under the interface, then:
  - OAM packet source IP address = Interface's IPv6 global address
  - OAM packet destination IP address = 0::ff:127.0.0.0
- Else, if an IPv6 link-local address is assigned to the interface, then:
  - OAM packet source IP address = Interface's IPv6 link-local address
  - OAM packet destination IP address = 0::ff:127.0.0.0

When the **next-hop** {**ipv4** | **ipv6**} address is configured, the following rules apply to determine the source and destination IP addresses used in the OAM packet:

- If a next-hop IPv4 address is configured, then:
  - OAM packet source IP address = Interface's IPv4 address
  - OAM packet destination IP address = Configured next-hop IPv4 address



#### Note

If there is no IPv4 address configured under the interface, then the delay-measurement probe does not send OAM packets.

- If a next-hop IPv6 address is configured, then:
  - OAM packet source IP address = Interface's IPv6 global address

- OAM packet destination IP address = Configured next-hop IPv6 address



**Note** If there is no IPv6 global address configured under the interface, then the delay-measurement probe does not send OAM packets.

This example shows how to enable PM for link delay over an interface with IPv4 address configured:

```
interface TenGigE0/0/0/0
 ipv4 address 10.10.10.1 255.255.255.0

performance-measurement
 interface TenGigE0/0/0/0
 delay-measurement
```

This example shows how to enable PM for link delay over an interface IPv6 address configured:

```
interface TenGigE0/0/0/0
 ipv6 address 10:10:10::1/64

performance-measurement
 interface TenGigE0/0/0/0
 delay-measurement
```

This example shows how to enable PM for link delay over an interface with a specified next-hop IPv4 address:

```
interface TenGigE0/0/0/0
 ipv4 address 10.10.10.1 255.255.255.0

performance-measurement
 interface TenGigE0/0/0/0
 next-hop ipv4 10.10.10.2
 delay-measurement
```

This example shows how to enable PM for link delay over an interface with a specified next-hop IPv6 address:

```
interface TenGigE0/0/0/0
 ipv6 address 10:10:10::1/64

performance-measurement
 interface TenGigE0/0/0/0
 next-hop ipv6 10:10:10::2
 delay-measurement
```

This example shows how to enable PM for link delay over an interface with only IPv6 link-local address:

```
interface TenGigE0/0/0/0
 ipv6 enable

performance-measurement
 interface TenGigE0/0/0/0
 delay-measurement
```



## Verification

```
RP/0/0/CPU0:router# show performance-measurement profile interface
Thu Dec 12 14:13:16.029 PST
```

```

0/0/CPU0

```

```
Interface Delay-Measurement:
```

```
Profile configuration:
```

```
Measurement Type : Two-Way
Probe computation interval : 30 (effective: 30) seconds
Type of services : Traffic Class: 6, DSCP: 48
Burst interval : 3000 (effective: 3000) mSec
Burst count : 10 packets
Encap mode : UDP
Payload Type : TWAMP-light
Destination sweeping mode : Disabled
Periodic advertisement : Enabled
Interval : 120 (effective: 120) sec
Threshold : 10%
Minimum-Change : 500 uSec
Advertisement accelerated : Disabled
Threshold crossing check : Minimum-delay
```

```
RP/0/0/CPU0:router# show performance-measurement summary detail location 0/2/CPU0
```

```
Thu Dec 12 14:09:59.162 PST
```

```

0/2/CPU0

```

```
Total interfaces : 1
Total SR Policies : 0
Total RSVP-TE tunnels : 0
Total Maximum PPS : 2000 pkts/sec
Total Interfaces PPS : 0 pkts/sec
Maximum Allowed Multi-hop PPS : 2000 pkts/sec
Multi Hop Requested PPS : 0 pkts/sec (0% of max allowed)
Dampened Multi Hop Requested PPS : 0% of max allowed
Inuse Burst Interval Adjustment Factor : 100% of configuration
```

```
Interface Delay-Measurement:
```

```
Total active sessions : 1
```

```
Counters:
```

```
Packets:
```

```
Total sent : 26
Total received : 26
```

```
Errors:
```

```
TX:
```

```
Reason interface down : 0
Reason no MPLS caps : 0
Reason no IP address : 0
Reason other : 0
```

```
RX:
```

```
Reason negative delay : 0
Reason delay threshold exceeded : 0
Reason missing TX timestamp : 0
Reason missing RX timestamp : 0
Reason probe full : 0
Reason probe not started : 0
Reason control code error : 0
Reason control code notif : 0
```

```

Probes:
 Total started : 3
 Total completed : 2
 Total incomplete : 0
 Total advertisements : 0

SR Policy Delay-Measurement:
 Total active sessions : 0
Counters:
 Packets:
 Total sent : 0
 Total received : 0
 Errors:
 TX:
 Reason interface down : 0
 Reason no MPLS caps : 0
 Reason no IP address : 0
 Reason other : 0
 RX:
 Reason negative delay : 0
 Reason delay threshold exceeded : 0
 Reason missing TX timestamp : 0
 Reason missing RX timestamp : 0
 Reason probe full : 0
 Reason probe not started : 0
 Reason control code error : 0
 Reason control code notif : 0
 Probes:
 Total started : 0
 Total completed : 0
 Total incomplete : 0
 Total advertisements : 0

RSVP-TE Delay-Measurement:
 Total active sessions : 0
Counters:
 Packets:
 Total sent : 0
 Total received : 0
 Errors:
 TX:
 Reason interface down : 0
 Reason no MPLS caps : 0
 Reason no IP address : 0
 Reason other : 0
 RX:
 Reason negative delay : 0
 Reason delay threshold exceeded : 0
 Reason missing TX timestamp : 0
 Reason missing RX timestamp : 0
 Reason probe full : 0
 Reason probe not started : 0
 Reason control code error : 0
 Reason control code notif : 0
 Probes:
 Total started : 0
 Total completed : 0
 Total incomplete : 0
 Total advertisements : 0

Global Delay Counters:
 Total packets sent : 26
 Total query packets received : 26

```

```

Total invalid session id : 0
Total missing session : 0

RP/0/0/CPU0:router# show performance-measurement interfaces detail
Thu Dec 12 14:16:09.692 PST

0/0/CPU0

0/2/CPU0

Interface Name: GigabitEthernet0/2/0/0 (ifh: 0x1004060)
Delay-Measurement : Enabled
Loss-Measurement : Disabled
Configured IPv4 Address : 10.10.10.2
Configured IPv6 Address : 10:10:10::2
Link Local IPv6 Address : fe80::3a:6fff:fec9:cd6b
Configured Next-hop Address : Unknown
Local MAC Address : 023a.6fc9.cd6b
Next-hop MAC Address : 0291.e460.6707
Primary VLAN Tag : None
Secondary VLAN Tag : None
State : Up

Delay Measurement session:
 Session ID : 1

Last advertisement:
 Advertised at: Dec 12 2019 14:10:43.138 (326.782 seconds ago)
 Advertised reason: First advertisement
 Advertised delays (uSec): avg: 839, min: 587, max: 8209, variance: 297

Next advertisement:
 Threshold check scheduled in 1 more probe (roughly every 120 seconds)
 Aggregated delays (uSec): avg: 751, min: 589, max: 905, variance: 112
 Rolling average (uSec): 756

Current Probe:
 Started at Dec 12 2019 14:15:43.154 (26.766 seconds ago)
 Packets 9, received: 9
 Measured delays (uSec): avg: 795, min: 631, max: 1199, variance: 164
 Next probe scheduled at Dec 12 2019 14:16:13.132 (in 3.212 seconds)
 Next burst packet will be sent in 0.212 seconds
 Burst packet sent every 3.0 seconds
 Probe samples:
 Packet Rx Timestamp Measured Delay (nsec)
 Dec 12 2019 14:15:43.156 689223
 Dec 12 2019 14:15:46.156 876561
 Dec 12 2019 14:15:49.156 913548
 Dec 12 2019 14:15:52.157 1199620
 Dec 12 2019 14:15:55.156 794008
 Dec 12 2019 14:15:58.156 631437
 Dec 12 2019 14:16:01.157 656440
 Dec 12 2019 14:16:04.157 658267
 Dec 12 2019 14:16:07.157 736880

```

You can also use the following commands for verifying the PM for link delay on the local-end router.

| Command                                                                                                       | Description                                                      |
|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>show performance-measurement history probe interfaces</b> [ <i>interface</i> ]                             | Displays the PM link-delay probe history for interfaces.         |
| <b>show performance-measurement history aggregated interfaces</b> [ <i>interface</i> ]                        | Displays the PM link-delay aggregated history for interfaces.    |
| <b>show performance-measurement history advertisement interfaces</b> [ <i>interface</i> ]                     | Displays the PM link-delay advertisement history for interfaces. |
| <b>show performance-measurement counters</b> [ <i>interface interface</i> ] [ <i>location location-name</i> ] | Displays the PM link-delay session counters.                     |

You can also use the following commands for verifying the PM for link-delay configuration on the remote-end router.

| Command                                                                                                                 | Description                                                                  |
|-------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| <b>show performance-measurement responder summary</b> [ <i>location location-name</i> ]                                 | Displays the PM for link-delay summary on the remote-end router (responder). |
| <b>show performance-measurement responder interfaces</b> [ <i>interface</i> ]                                           | Displays PM for link-delay for interfaces on the remote-end router.          |
| <b>show performance-measurement responder counters</b> [ <i>interface interface</i> ] [ <i>location location-name</i> ] | Displays the PM link-delay session counters on the remote-end router.        |

### Configure a Static Delay Value on an Interface

You can configure an interface to advertise a static delay value, instead of the measured delay value. When you configure a static delay value, the advertisement is triggered immediately. The average, minimum, and maximum advertised values will use the static delay value, with a variance of 0.

Scheduled probes will continue, and measured delay metrics will be aggregated and stored in history buffer. However, advertisement threshold checks are suppressed so that there are no advertisements of the actual measured delay values. If the configured static delay value is removed, the next scheduled advertisement threshold check will update the advertised measured delay values.

The static delay value can be configured from 1 to 16777215 microseconds (16.7 seconds).

This example shows how to configure a static delay of 1000 microseconds:

```
RP/0/0/CPU0:router(config)# performance-measurement
RP/0/0/CPU0:router(config-perf-meas)# interface TenGigE0/0/0/0
RP/0/0/CPU0:router(config-pm-intf)# delay-measurement
RP/0/0/CPU0:router(config-pm-intf-dm)# advertise-delay 1000
```

### Running Configuration

```
performance-measurement
 interface GigabitEthernet0/0/0/0
 delay-measurement
 advertise-delay 1000
```

```
!
!
!
```

## Verification

```
RP/0/RSP0/CPU0:ios# show performance-measurement interfaces detail
```

```

0/0/CPU0

```

```
Interface Name: GigabitEthernet0/0/0/0 (ifh: 0x0)
 Delay-Measurement : Enabled
```

```
. . .
```

```
Last advertisement:
 Advertised at: Nov 29 2021 21:53:00.656 (7.940 seconds ago)
 Advertised reason: Advertise delay config
 Advertised delays (uSec): avg: 1000, min: 1000, max: 1000, variance: 0
```

```
. . .
```

## SR Performance Measurement Named Profiles

You can create a named performance measurement profile for delay or liveness.

### Delay Profile

This example shows how to create a named SR performance measurement delay profile.

```
Router(config)# performance-measurement delay-profile sr-policy profile2
Router(config-pm-dm-srpolicy)# probe
Router(config-pm-dm-srpolicy-probe)# burst-interval 60
Router(config-pm-dm-srpolicy-probe)# computation-interval 60
Router(config-pm-dm-srpolicy-probe)# protocol twamp-light
Router(config-pm-dm-srpolicy-probe)# tos dscp 63
```

```
Router(config-pm-dm-srpolicy)# advertisement
Router(config-pm-dm-srpolicy-adv)# periodic
Router(config-pm-dm-srpolicy-adv-per)# interval 60
Router(config-pm-dm-srpolicy-adv-per)# minimum-change 1000
Router(config-pm-dm-srpolicy-adv-per)# threshold 20
Router(config-pm-dm-srpolicy-adv-per)# commit
```

Apply the delay profile for an SR Policy.

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy TEST
Router(config-sr-te-policy)# color 4 end-point ipv4 10.10.10.10
Router(config-sr-te-policy)# performance-measurement
Router(config-sr-te-policy-perf-meas)# delay-measurement delay-profile name profile2
```

```
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list LIST1
Router(config-sr-te-pp-info)# weight 2
```

```
Router(config-sr-te-policy-path-pref)# explicit segment-list LIST2
Router(config-sr-te-pp-info)# weight 3
```

## Running Configuration

```
Router# show run segment-routing traffic-eng policy TEST
```

```
segment-routing
 traffic-eng
 policy TEST
 color 4 end-point ipv4 10.10.10.10
 candidate-paths
 preference 100
 explicit segment-list LIST1
 weight 2
 !
 explicit segment-list LIST2
 weight 3
 !
 !
 !
performance-measurement
 delay-measurement
 delay-profile name profile2
```

### Verification

```
Router# show performance-measurement profile named-profile delay sr-policy name profile2
```

```

0/RSP0/CPU0

SR Policy Delay Measurement Profile Name: profile2
Profile configuration:
 Measurement mode : One-way
 Protocol type : TWAMP-light
 Encap mode : UDP
 Type of service:
 PM-MPLS traffic class : 6
 TWAMP-light DSCP : 63
 Probe computation interval : 60 (effective: 60) seconds
 Burst interval : 60 (effective: 60) mSec
 Packets per computation interval : 1000
 Periodic advertisement : Enabled
 Interval : 60 (effective: 60) sec
 Threshold : 20%
 Minimum-change : 1000 uSec
 Advertisement accelerated : Disabled
 Advertisement logging:
 Delay exceeded : Disabled (default)
 Threshold crossing check : Maximum-delay
 Router alert : Disabled (default)
 Destination sweeping mode : Disabled
 Liveness detection parameters:
 Multiplier : 3
 Logging state change : Disabled
```

### On-Demand SR Policy

```
Router(config-sr-te)# on-demand color 20
Router(config-sr-te-color)# performance-measurement delay-measurement
Router(config-sr-te-color-delay-meas)# delay-profile name profile2
Router(config-sr-te-color-delay-meas)# commit
```

### Running Configuration

```
Router# show run segment-routing traffic-eng on-demand color 20
```

```
segment-routing
 traffic-eng
 on-demand color 20
```

```

performance-measurement
delay-measurement
delay-profile name profile2

```

### Liveness Profile

This example shows how to create a *named* SR performance measurement liveness profile.

```

Router(config)# performance-measurement liveness-profile sr-policy name profile3
Router(config-pm-ld-srpolicy)# probe
Router(config-pm-ld-srpolicy-probe)# burst-interval 60
Router(config-pm-ld-srpolicy-probe)# measurement-mode loopback
Router(config-pm-ld-srpolicy-probe)# tos dscp 10
Router(config-pm-ld-srpolicy-probe)# liveness-detection
Router(config-pm-ld-srpolicy-probe)# multiplier 5
Router(config-pm-ld-srpolicy-probe)# commit

```

### Apply the Liveness Profile for the SR Policy

This example shows how to enable PM for SR policy liveness for a specific policy.

For the same policy, you cannot enable delay-measurement (delay-profile) and liveness-detection (liveness-profile) at the same time. For example, if delay measurement is enabled, use the **no delay-measurement** command to disable it, and then enable the following command for enabling liveness detection.

```

Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy TRST2
Router(config-sr-te-policy)# color 40 end-point ipv4 20.20.20.20
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 50
Router(config-sr-te-policy-path-pref)# explicit segment-list LIST3
Router(config-sr-te-pp-info)# weight 2

Router(config-sr-te-policy-path-pref)# explicit segment-list LIST4
Router(config-sr-te-pp-info)# weight 3

Router(config-sr-te-policy)# performance-measurement
Router(config-sr-te-policy-perf-meas)# liveness-detection liveness-profile name profile3

```

### Running Configuration

```
Router# show run segment-routing traffic-eng policy TRST2
```

```

segment-routing
 traffic-eng
 policy TRST2
 color 40 end-point ipv4 20.20.20.20
 candidate-paths
 preference 50
 explicit segment-list LIST3
 weight 2
 !
 explicit segment-list LIST4
 weight 3
 !
!
!
performance-measurement
 liveness-detection
 liveness-profile name profile3
!

```

### Verification

```

Router# show performance-measurement profile named-profile delay

0/RSP0/CPU0

SR Policy Liveness Detection Profile Name: profile1
 Profile configuration:
 Measurement mode : Loopback
 Protocol type : TWAMP-light
 Type of service:
 TWAMP-light DSCP : 10
 Burst interval : 60 (effective: 60) mSec
 Destination sweeping mode : Disabled
 Liveness detection parameters:
 Multiplier : 3
 Logging state change : Disabled

SR Policy Liveness Detection Profile Name: profile3
 Profile configuration:
 Measurement mode : Loopback
 Protocol type : TWAMP-light
 Type of service:
 TWAMP-light DSCP : 10
 Burst interval : 60 (effective: 60) mSec
 Destination sweeping mode : Disabled
 Liveness detection parameters:
 Multiplier : 3
 Logging state change : Disabled

```

### On-Demand SR Policy

For the same policy, you cannot enable delay-measurement (delay-profile) and liveness-detection (liveness-profile) at the same time. For example, to disable delay measurement, use the **no delay-measurement** command, and then enable the following command for enabling liveness detection.

```

Router(config-sr-te)# on-demand color 30
Router(config-sr-te-color)# performance-measurement
Router(config-sr-te-color-pm)# liveness-detection liveness-profile name profile1
Router(config-sr-te-color-delay-meas)# commit

```

### Running Configuration

```

Router# show run segment-routing traffic-eng on-demand color 30

segment-routing
 traffic-eng
 on-demand color 30
 performance-measurement
 liveness-detection
 liveness-profile name profile1
!
```

### Verification

```

Router# show performance-measurement profile named-profile liveness sr-policy name profile1

0/RSP0/CPU0

SR Policy Liveness Detection Profile Name: profile1
 Profile configuration:
 Measurement mode : Loopback
 Protocol type : TWAMP-light
 Type of service:

```



```

TWAMP-light DSCP : 10
Burst interval : 60 (effective: 60) mSec
Destination sweeping mode : Disabled
Liveness detection parameters:
Multiplier : 3
Logging state change : Disabled

```

## Delay Normalization

**Table 69: Feature History Table**

| Feature Name                       | Release Information | Feature Description                                                           |
|------------------------------------|---------------------|-------------------------------------------------------------------------------|
| SR-TE Delay Normalization for OSPF | Release 7.3.1       | This feature extends the current Delay Normalization feature to support OSPF. |

Performance measurement (PM) measures various link characteristics like packet loss and delay. Such characteristics can be used by IS-IS as a metric for Flexible Algorithm computation. Low latency routing using dynamic delay measurement is one of the primary use cases for Flexible Algorithm technology.

Delay is measured in microseconds. If delay values are taken as measured and used as link metrics during the IS-IS topology computation, some valid ECMP paths might be unused because of the negligible difference in the link delay.

The Delay Normalization feature computes a normalized delay value and uses the normalized value instead. This value is advertised and used as a metric during the Flexible Algorithm computation.

The normalization is performed when the delay is received from the delay measurement component. When the next value is received, it is normalized and compared to the previous saved normalized value. If the values are different, then the LSP generation is triggered.


The following formula is used to calculate the normalized value:

- **Dm** – measured Delay
- **Int** – configured normalized Interval
- **Off** – configured normalized Offset (must be less than the normalized interval Int)
- **Dn** – normalized Delay
- **a** = Dm / Int (rounded down)
- **b** = a \* Int + Off

If the measured delay (Dm) is less than or equal to **b**, then the normalized delay (Dn) is equal to **b**. Otherwise, Dn is **b + Int**.

### Example

The following example shows a low-latency service. The intent is to avoid high-latency links (1-6, 5-2). Links 1-2 and 5-6 are both low-latency links. The measured latency is not equal, but the difference is insignificant.



We can normalize the measured latency before it is advertised and used by IS-IS. Consider a scenario with the following:

- Interval = 10
- Offset = 3


The measured delays will be normalized as follows:

- **Dm** = 29

$$a = 29 / 10 = 2 \text{ (2.9, rounded down to 2)}$$

$$b = 2 * 10 + 3 = 23$$

In this case, **Dm** (29) is greater than **b** (23); so **Dn** is equal to **b+I** (23 + 10) = 33




- **Dm** = 31


$$a = 31 / 10 = 3 \text{ (3.1, rounded down to 3)}$$

$$b = 3 * 10 + 3 = 33$$

In this case, **Dm** (31) is less than **b** (33); so **Dn** is **b** = 33



The link delay between 1-2 and 5-6 is normalized to 33.



### Configuration

Delay normalization is disabled by default. To enable and configure delay normalization, use the **delay normalize interval** *interval* [**offset** *offset*] command.

- *interval* – The value of the normalize interval in microseconds.
- *offset* – The value of the normalized offset in microseconds. This value must be smaller than the value of normalized interval.

### IS-IS Configuration

```
router isis 1
 interface GigEth 0/0/0/0
 delay normalize interval 10 offset 3
 address-family ipv4 unicast
 metric 77
```

### OSPF Configuration

```
router ospf 1
 area 0
 interface GigabitEthernet0/0/0/0
 delay normalize interval 10 offset 3
 !
 !
 !
```

## Link Anomaly Detection with IGP Penalty

Table 70: Feature History Table

| Feature Name                            | Release Information | Feature Description                                                                                                                                                                                                                       |
|-----------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Anomaly Detection with IGP Penalty | Release 7.4.1       | This feature allows you to define thresholds above the measured delay that is considered “anomalous” or unusual. When this threshold is exceeded, an anomaly (A) bit/flag is set along with link delay attribute that is sent to clients. |

Customers might experience performance degradation issues, such as increased latency or packet loss on a link. Degraded links might be difficult to troubleshoot and can affect applications, especially in cases where traffic is sent over multiple ECMP paths where one of those paths is degraded.

The Anomaly Detection feature allows you to define a delay anomaly threshold to identify unacceptable link delays. Nodes monitor link performance using link delay monitoring probes. The measured value is compared against the delay anomaly threshold values. When the upper bound threshold is exceeded, the link is declared “abnormal”, and performance measurement sets an anomaly bit (A-bit). When IGP receives the A-bit, IGP can automatically increase the IGP metric of the link by a user-defined amount to make this link undesirable or unusable. When the link recovers (lower bound threshold), PM resets the A-bit.

For information on configuring IGP penalty, see the following:

- [IS-IS Penalty for Link Delay Anomaly](#)
- [OSPF Penalty for Link Delay Anomaly](#)

### Usage Guidelines and Limitations

This feature is not active when narrow metrics are configured because the performance measurement advertisement requires the “wide” metric type length values.

### Configuration Example

The following example shows how to configure the upper and lower anomaly thresholds. The range for *upper\_bound* and *lower\_bound* is from 1 to 200,000 microseconds. The *lower\_bound* value must be less than the *upper\_bound* value.

```
RP/0/0/CPU0:router(config)# performance-measurement delay-profile interfaces default
RP/0/0/CPU0:router(config-pm-dm-intf)# advertisement
RP/0/0/CPU0:router(config-pm-dm-intf-adv)# anomaly-check upper-bound 5000 lower-bound 1000
RP/0/0/CPU0:router(config-pm-dm-intf-adv)# commit
```

### Running Configuration

```
performance-measurement
 delay-profile interfaces default
 advertisement
 anomaly-check
 upper-bound 5000 lower-bound 1000
 !
!
```

```

!
!
end

```

## Delay Measurement for IP Endpoint

*Table 71: Feature History Table*

| Feature Name                                        | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Delay Measurement for IP Endpoint over SRv6 Network | Release 24.2.11     | This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Delay Measurement for IP Endpoint over SRv6 Network | Release 24.2.1      | <p>Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR:</p> <ul style="list-style-type: none"> <li>• N540-ACC-SYS</li> <li>• N540X-ACC-SYS</li> <li>• N540-24Z8Q2C-SYS</li> </ul> <p>In Segment Routing over an IPv6 network (SRv6), you can measure packet delay from the source to a specific IP endpoint. You can use this information for troubleshooting, network maintenance, and optimizing network performance.</p> <p>Additionally, you can use flow labels to verify the delay of each subsequent hop path towards the IP endpoint of that path. So that, when network traffic is distributed across multiple available paths towards an IP endpoint, delay measurement tracks the delay of each of these paths towards the IP endpoint.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• The <b>source-address ipv6</b> keyword is introduced in the <b>performance-measurement endpoint</b> command.</li> <li>• The <b>segment-list name</b> keyword is introduced in the <b>segment-routing traffic-eng explicit</b> command.</li> <li>• The <b>flow-label</b> keyword is introduced in the <b>performance-measurement delay-profile name</b> command.</li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>• Cisco-IOS-XR-um-performance-measurement-cfg</li> <li>• Cisco-IOS-XR-perf-meas-oper.yang</li> </ul> <p>(See <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</p> |

| Feature Name                             | Release Information | Feature Description                                                                                                                                                                                                                                        |
|------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Endpoint Delay Measurement Monitoring | Release 7.4.1       | This feature measures the end-to-end delay and monitors liveness of a specified IP endpoint node, including VRF-aware (awareness of multiple customers belonging to different VRFs).<br><br>This feature is supported on IPv4, IPv6, and MPLS data planes. |

Delay for an IP endpoint is the amount of time it takes for a data packet to travel from a source device to a specific IP endpoint within a network.

To measure a delay for a packet, also called a probe, is sent from a source device to the target IP endpoint.

The time from when the packet leaves the source to when it arrives at the endpoint is measured and recorded as the delay.

You can measure one-way delay, Two-way delay, and Roundtrip delay or delay in loop-back mode. For more information on Delay measurement, see Link Delay Measurement and Measurement Modes.

### Collecting IP Endpoint Probe Statistics

- Statistics associated with the probe for delay metrics are available via Histogram and Streaming Telemetry.
- Model Driven Telemetry (MDT) is supported for the following data:
  - Summary, endpoint, session, and counter show command bags.
  - History buffers data
- Model Driven Telemetry (MDT) and Event Driven Telemetry (EDT) are supported for the following data:
  - Delay metrics computed in the last probe computation-interval (event: probe-completed)
  - Delay metrics computed in the last aggregation-interval; that is, end of the periodic advertisement-interval (event: advertisement-interval expired)
  - Delay metrics last notified (event: notification-triggered)
- The following xpaths for MDT/EDT is supported:
  - `Cisco-IOS-XR-perf-meas-oper:performance-measurement/nodes/node/endpoints/endpoint-delay/endpoint-last-probes`
  - `Cisco-IOS-XR-perf-meas-oper:performance-measurement/nodes/node/endpoints/endpoint-delay/endpoint-last-aggregations`
  - `Cisco-IOS-XR-perf-meas-oper:performance-measurement/nodes/node/endpoints/endpoint-delay/endpoint-last-advertisements`

### Guidelines and Limitations

You can specify a custom labeled path through one or more user-configured segment-lists. User-configured segment-list represents the forwarding path from sender to reflector when the probe is configured in delay-measurement mode.

- Examples of the custom segment-list include:
  - Probe in delay-measurement mode with a segment-list that includes Flex-Algo prefix SID of the endpoint
  - Probe in delay-measurement mode with a segment-list that includes a SID-list with labels to reach the endpoint or the sender (forward direction)
  - Probe in delay-measurement mode with a segment-list that includes BSID associated with SR policy to reach the end point.
- Endpoint segment list configuration is not supported under nondefault VRF.
- SR Performance Measurement endpoint session over BVI interface is not supported.

### Unsupported Features for IP Endpoint Over SRv6 Network

- SRv6 PM is not supported on routers that use channelized ports
- The IPv6 Endpoint in vrf (dynamic uDT6) is not supported on NCS 5508 because the timestamp is not preserved among different LCs. However, it is supported when using an SRv6 segment list with the uDT case. In this scenario, use a global source address instead of the source address within the VRF.
- The IPv6 or IPv4 Endpoint in VRF is not supported for dynamic uDT on NCS 5508.

However, it is supported when using an SRv6 segment list with the uDT case. In this scenario, use a global source address instead of the source address within the VRF.

For example:

```
performance-measurement
 endpoint ipv4 10.2.2.2 vrf srv6_underlay
 delay-measurement
```

- SRv6 PM IPv6 endpoint delay session over SRv6 does not work on the Route Processor (RP) of NCS 5500 line cards and NCS 5700 line cards in the NCS 5508 Modular Router.


## IP Endpoint Delay Measurement over MPLS Network Usecases

The following use-cases show different ways to deploy delay measurement and liveness detection for IP endpoints.

### Use-Case 1: Delay Measurement Probe Toward an IP Endpoint Reachable in the Global Routing Table

The following figure illustrates a delay measurement probe toward an IP endpoint reachable in the global routing table. The network interconnecting the sender and the reflector provides plain IP connectivity.

Figure 36: Delay Measurement Probe Toward an IP Endpoint Reachable in the Global Routing Table



522018

### Configuration

```
RouterA(config)# performance-measurement
RouterA(config-perf-meas)# endpoint ipv4 10.1.1.5
RouterA(config-pm-ep)# source-address ipv4 10.1.1.1
RouterA(config-pm-ep)# delay-measurement
RouterA(config-pm-ep-dm)# exit
RouterA(config-pm-ep)# exit
RouterA(config-perf-meas)# delay-profile endpoint default
RouterA(config-pm-dm-ep)# probe
RouterA(config-pm-dm-ep-probe)# measurement-mode one-way
```

### Running Configuration

```
performance-measurement
 endpoint ipv4 10.1.1.5
 source-address ipv4 10.1.1.1
 delay-measurement
 !
!
 delay-profile endpoint default
 probe
 measurement-mode one-way
 !
!
```

### Verification

```
RouterA# show performance-measurement endpoint ipv4 10.1.1.5
```



```
0/RSP0/CPU0
```

---

```
Endpoint name: IPv4-10.1.1.5-vrf-default
Source address : 10.1.1.1
VRF name : default
Delay-measurement : Enabled
Description : Not set
Profile Keys:
 Profile name : default
 Profile type : Endpoint Delay Measurement

Segment-list : None
Delay Measurement session:
 Session ID : 33554433
 Last advertisement:
 No advertisements have occurred


Next advertisement:
 Threshold check scheduled in 4 more probes (roughly every 120 seconds)
 No probes completed

Current computation:
 Started at: Jul 19 2021 16:28:06.723 (17.788 seconds ago)
 Packets 6, received: 0
 Measured delays (uSec): avg: 0, min: 0, max: 0, variance: 0
 Next probe scheduled at: Jul 19 2021 16:28:36.718 (in 12.207 seconds)
 Next burst packet will be sent in 0.207 seconds
 Burst packet sent every 3.0 seconds
```

### Use-Case 2: Delay Measurement Probe Toward an IP Endpoint Reachable in a User-Specified VRF

The following figure illustrates a delay measurement probe toward an IP endpoint reachable in a user-specified L3VPN's VRF routing table. The L3VPN ingress PE (Router A) acts as the sender. The reflector is located in a CE device behind the L3VPN egress PE (Router E). The network interconnecting the L3VPN PEs provides MPLS connectivity with Segment Routing.

Figure 37: Delay Measurement Probe Toward an IP Endpoint Reachable in a User-Specified VRF



### Configuration

```

RouterA(config)# performance-measurement
RouterA(config-perf-meas)# endpoint ipv4 10.10.10.100 vrf green
RouterA(config-pm-ep)# source-address ipv4 10.1.1.1
RouterA(config-pm-ep)# delay-measurement
RouterA(config-pm-ep-dm)# exit
RouterA(config-pm-ep)# exit
RouterA(config-perf-meas)# delay-profile endpoint default
RouterA(config-pm-dm-ep)# probe
RouterA(config-pm-dm-ep-probe)# measurement-mode one-way

```

### Running Configuration

```

performance-measurement
 endpoint ipv4 10.10.10.100 vrf green
 source-address ipv4 10.1.1.1
 delay-measurement
 !
 delay-profile endpoint default
 probe
 measurement-mode one-way

```

```
!
!
!
```

## Verification

```
RouterA# show performance-measurement endpoint vrf green
```

```
0/RSP0/CPU0
```

```
Endpoint name: IPv4-10.10.10.100-vrf-green
Source address : 10.1.1.1
VRF name : green
Delay-measurement : Enabled
Description : Not set
Profile Keys:
 Profile name : default
 Profile type : Endpoint Delay Measurement
```

```
Segment-list : None
```

```
Delay Measurement session:
```

```
 Session ID : 33554434
```

```
 Last advertisement:
```

```
 No advertisements have occurred
```

```
 Next advertisement:
```

```
 Advertisement not scheduled as the probe is not running
```

```
Current computation:
```

```
 Not running: Unable to resolve (non-existing) vrf
```

## Use Case 3: Delay Measurement Probe Toward an IP Endpoint Using Custom Labeled Paths

The following figure illustrates a delay measurement probe toward an IP endpoint learned by the IGP. The network interconnecting the sender and reflector provides MPLS connectivity with Segment Routing.


The IP endpoint is advertised with multiple SR algorithms (Algo 0 and Flex Algo 128). The probe is configured with two custom-labeled paths in order to monitor the LSP for each algorithm separately.



**Note** Starting from Cisco IOS XR Release 24.1.1, use the **segment-routing traffic-eng explicit segment-list name** command to configure the segment list.

Earlier, the **segment-routing traffic-eng segment-list name** command was used to configure the segment list.

Figure 38: Delay Measurement Probe Toward an IP Endpoint Using Custom Labeled Paths



**Note** The probe response messages are not shown in the above figure.

### Configuration

```
RouterA(config)# segment-routing
RouterA(config-sr)# traffic-eng
RouterA(config-sr-te)# segment-list name SIDLIST1-Algo0
RouterA(config-sr-te-sl)# index 10 mpls label 16005
RouterA(config-sr-te-sl)# exit

RouterA(config-sr-te)# segment-list name SIDLIST2-FlexAlgo128
RouterA(config-sr-te-sl)# index 10 mpls label 16085
RouterA(config-sr-te-sl)# exit
RouterA(config-sr-te)# exit
RouterA(config-sr)# exit

RouterA(config)# performance-measurement
RouterA(config-perf-meas)# endpoint ipv4 10.1.1.5
RouterA(config-pm-ep)# source-address ipv4 10.1.1.1
RouterA(config-pm-ep)# segment-list name SIDLIST1-Algo0
RouterA(config-pm-ep-sl)# exit
RouterA(config-pm-ep)# segment-list name SIDLIST2-FlexAlgo128
RouterA(config-pm-ep-sl)# exit
RouterA(config-pm-ep)# delay-measurement
RouterA(config-pm-ep-dm)# exit
RouterA(config-pm-ep)# exit
RouterA(config-perf-meas)# delay-profile endpoint default
RouterA(config-pm-dm-ep)# probe
RouterA(config-pm-dm-ep-probe)# measurement-mode one-way
```

### Running Configuration

```
segment-routing
traffic-eng
segment-list SIDLIST1-Algo0
index 10 mpls label 16005
!
segment-list SIDLIST2-FlexAlgo128
index 10 mpls label 16085
!
!
```

```

!
!
performance-measurement
endpoint ipv4 10.1.1.5
 segment-list name SIDLIST1-Algo0
!
 segment-list name SIDLIST2-FlexAlgo128
!
 source-address ipv4 10.1.1.1
 delay-measurement
!
!
delay-profile endpoint default
probe
 measurement-mode one-way
!
!
!

```

## Verification

RouterA# **show performance-measurement endpoint ipv4 10.1.1.5**

---

0/RSP0/CPU0

---

Endpoint name: **IPv4-10.1.1.5-vrf-default**

```

Source address : 10.1.1.1
VRF name : default
Delay-measurement : Enabled
Description : Not set
Profile Keys:
 Profile name : default
 Profile type : Endpoint Delay Measurement

```

Segment-list : None

Delay Measurement session:

Session ID : 33554433

Last advertisement:

No advertisements have occurred

Next advertisement:

Threshold check scheduled in 4 more probes (roughly every 120 seconds)

No probes completed

Current computation:

Started at: Jul 19 2021 16:31:53.827 (15.844 seconds ago)

Packets 6, received: 0

Measured delays (uSec): avg: 0, min: 0, max: 0, variance: 0

Next probe scheduled at: Jul 19 2021 16:32:22.957 (in 13.286 seconds)

Next burst packet will be sent in 1.286 seconds

Burst packet sent every 3.0 seconds

Segment-list : **SIDLIST1-Algo0**

Delay Measurement session:

Session ID : 33554435

Last advertisement:

No advertisements have occurred

Next advertisement:

Threshold check scheduled in 4 more probes (roughly every 120 seconds)

No probes completed

```

Current computation:
Started at: Jul 19 2021 16:31:53.827 (15.844 seconds ago)
Packets 4, received: 0
Measured delays (uSec): avg: 0, min: 0, max: 0, variance: 0
Next probe scheduled at: Jul 19 2021 16:32:22.957 (in 13.286 seconds)
Next burst packet will be sent in 2.940 seconds
Burst packet sent every 3.0 seconds

Segment-list : SIDLIST2-FlexAlgo128
Delay Measurement session:
Session ID : 33554436
Last advertisement:
 No advertisements have occurred

Next advertisement:
 Threshold check scheduled in 4 more probes (roughly every 120 seconds)
 No probes completed

Current computation:
Started at: Jul 19 2021 16:31:53.827 (15.844 seconds ago)
Packets 4, received: 0
Measured delays (uSec): avg: 0, min: 0, max: 0, variance: 0
Next probe scheduled at: Jul 19 2021 16:32:22.957 (in 13.286 seconds)
Next burst packet will be sent in 2.940 seconds
Burst packet sent every 3.0 seconds

```

#### Use-Case 4: Liveness Detection Probe Toward an IP Endpoint


IP endpoint liveness detection leverages the loopback measurement-mode. The following workflow describes the sequence of events.

1. The sender creates and transmits the PM probe packets.  
 The IP destination address (DA) on the probe packets is set to the loopback value of the sender itself.  
 The transmit timestamp (T1) is added to the payload.  
 The probe packet is encapsulated with the label corresponding to the endpoint.
2. The network delivers the PM probe packets following the LSP toward the endpoint.
3. The end-point receives the PM probe packets.  
 Packets are forwarded back to the sender based on the forwarding entry associated with the IP DA of the PM probe packet. If an LSP exists, the probe packet is encapsulated with the label of the sender.
4. The sender node receives the PM probe packets.  
 The received timestamp (T4) stored.  
 If the sender node doesn't receive the specified number of probe packets (based on the configured multiplier), the sender node declares the PM session as down.

The following figure illustrates a liveness detection probe toward an IP endpoint learned by the IGP. The network interconnecting the sender and reflector provides MPLS connectivity with Segment Routing.

The liveness detection multiplier is set to 5 to specify the number of consecutive missed probe packets before the PM session is declared as down.

Figure 39: IP Endpoint Liveness Detection



### Configuration

```

RouterA(config)# performance-measurement
RouterA(config-perf-meas)# endpoint ipv4 10.1.1.5
RouterA(config-pm-ep)# source-address ipv4 10.1.1.1
RouterA(config-pm-ep)# liveness-detection
RouterA(config-pm-ep-ld)# exit
RouterA(config-pm-ep)# exit
RouterA(config-perf-meas)# liveness-profile endpoint default
RouterA(config-pm-ld-ep)# liveness-detection
RouterA(config-pm-ld-ep-ld)# multiplier 5
RouterA(config-pm-ld-ep-ld)# exit
RouterA(config-pm-ld-ep)# probe
RouterA(config-pm-ld-ep-probe)# measurement-mode loopback

```

### Running Configuration

```

performance-measurement
 endpoint ipv4 10.1.1.5
 source-address ipv4 10.1.1.1
 liveness-detection
 !
 !
 liveness-profile endpoint default
 liveness-detection
 multiplier 5
 !
 probe
 measurement-mode loopback
 !
 !
 !
end

```

### Verification

```

RouterA# show performance-measurement endpoint ipv4 10.1.1.5

```

```

0/RSP0/CPU0

```

```


Endpoint name: IPv4-10.1.1.5-vrf-default
Source address : 10.1.1.1
VRF name : default
Liveness Detection : Enabled
Profile Keys:
 Profile name : default
 Profile type : Endpoint Liveness Detection

Segment-list : None
Session State: Down
Missed count: 0

```

## IP Endpoint Delay Measurement over SRv6 Network Usecase

The following figure illustrates a delay measurement probe toward an IP endpoint over the SRv6 network. The network interconnecting the sender and the reflector provides plain IP connectivity.



523655

```

RP/0/RSP0/CPU0:ios#configure
RP/0/RSP0/CPU0:ios(config)#performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)#endpoint ipv6 FCBB:0:1::
RP/0/RSP0/CPU0:ios(config-pm-ep)#delay-measurement
RP/0/RSP0/CPU0:ios(config-pm-ep-dm)#delay-profile name test
RP/0/RSP0/CPU0:ios(config-pm-ep-dm)#exit
RP/0/RSP0/CPU0:ios(config-pm-ep)#exit
RP/0/RSP0/CPU0:ios(config-perf-meas)#liveness-profile name test
RP/0/RSP0/CPU0:ios(config-pm-ld-profile)#probe
RP/0/RSP0/CPU0:ios(config-pm-ld-probe)#flow-label explicit 100 200 300
RP/0/RSP0/CPU0:ios(config-pm-ld-probe)#

```

The following example shows how to use flow label for delay profile for a default endpoint:

```

RP/0/RSP0/CPU0:ios#configure
RP/0/RSP0/CPU0:ios(config)#performance-measurement
RP/0/RSP0/CPU0:ios(config-perf-meas)#delay-profile endpoint default
RP/0/RSP0/CPU0:ios(config-pm-dm-ep)#probe
RP/0/RSP0/CPU0:ios(config-pm-dm-ep-probe)#flow-label explicit 100 200 300

```

### Show Running Configuration

```

performance-measurement
endpoint ipv6 FCBB:0:1::
delay-measurement

```



```

 delay-profile name test
 !
!
liveness-profile name test
probe
 flow-label explicit 100 200 300
!
!
!
!

```

## Verification

The show output displays the delay information for the endpoint.

```

Router# show performance-measurement endpoint detail
Endpoint name: IPv6-FCBB:0:1::-vrf-default
Source address : 192::2
VRF name : default
Liveness Detection : Enabled
Profile Keys:
 Profile name : default
 Profile type : Endpoint Liveness Detection
Segment-list : None
Liveness Detection session:
 Session ID : 4109
 Flow-label : 1000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.214
 Missed count: 0

Liveness Detection session:
 Session ID : 4110
 Flow-label : 2000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.214
 Missed count: 0

Segment-list : test-dm-two-carrier-sl2
FCBB:0:1:2:e004::/64
 Format: f3216
FCBB:0:1:3:e000::/64
 Format: f3216
FCBB:0:1:2:e004::/64
 Format: f3216
FCBB:0:1:2:e000::/64
 Format: f3216
FCBB:0:1:1:e000::/64
 Format: f3216
FCBB:0:1:1:e004::/64
 Format: f3216
FCBB:0:1:4:e000::/64
 Format: f3216
FCBB:0:1:4::/48
 Format: f3216
Liveness Detection session:
 Session ID : 4111
 Flow-label : 1000
 Session State: Up
 Last State Change Timestamp: Jan 23 2024 16:06:01.217
 Missed count: 0

Liveness Detection session:
 Session ID : 4112
 Flow-label : 2000

```

```
Session State: Up
Last State Change Timestamp: Jan 23 2024 16:06:01.217
Missed count: 0
```

## SR Policy End-to-End Delay Measurement

The PM for SR Policy uses the IP/UDP packet format defined in RFC 5357 (TWAMP-Light) for probes. Two-Way Active Measurement Protocol (TWAMP) adds two-way or round-trip measurement capabilities. TWAMP employs time stamps applied at the echo destination (reflector) to enable greater accuracy. In the case of TWAMP Light, the Session-Reflector doesn't necessarily know about the session state. The Session-Reflector simply copies the Sequence Number of the received packet to the Sequence Number field of the reflected packet. The controller receives the reflected test packets and collects two-way metrics. This architecture allows for collection of two-way metrics.

The extended TE link delay metric (minimum-delay value) can be used to compute paths for SR policies as an optimization metric or as an accumulated delay bound.

There is a need to monitor the end-to-end delay experienced by the traffic sent over an SR policy to ensure that the delay does not exceed the requested "upper-bound" and violate SLAs. You can verify the end-to-end delay values before activating the candidate-path or the segment lists of the SR policy in forwarding table, or to deactivate the active candidate-path or the segment lists of the SR policy in forwarding table.



---

**Note** The end-to-end delay value of an SR policy will be different than the path computation result (for example, the sum of TE link delay metrics) due to several factors, such as queuing delay within the routers.

---

### Usage Guidelines and Limitations for PM for SR Policy Delay

The following usage guidelines and limitations apply:

- SR-PM delay measurement over SR Policy is supported on manually configured SR Policies and On-Demand SR Policies (ODN).
- SR-PM delay measurement over SR Policy is not supported on PCE-initiated SR Policies.
- Hardware clocks must be synchronized between the querier and the responder nodes of the link using PTP for one-way delay measurement.

### Configuring Performance Measurement Parameters

This example shows how to configure performance-measurement parameters for SR policy delay as a global default profile. The default values for the different parameters in the PM for SR policy delay is given as follows:

- **probe**: The default mode for probe is one-way delay measurement. Two-way delay and loopback modes are supported. See [Measurement Modes, on page 464](#) for more information.
- **burst interval**: Interval for sending probe packet. The default value is 3000 milliseconds and the range is from 30 to 15000 milliseconds.
- **computation interval**: Interval for metric computation. Default is 30 seconds; range is 1 to 3600 seconds.
- **protocol**:

- **twamp-light**: SR Policy delay measurement using RFC 5357 with IP/UDP encap. This is the default protocol.
- **tos**: Type of Service
  - **dscp value**: The default value is 48 and the range is from 0 to 63.
  - **traffic-class value**: The default value is 6 and the range is from 0 to 7.
- **advertisement threshold-check**: minimum-delay/maximum-delay - The default value of periodic advertisement threshold-check is maximum-delay.
- **periodic advertisement**: Periodic advertisement is enabled by default.
- **periodic-advertisement interval**: The default value is 120 seconds and the interval range is from 30 to 3600 seconds.
- **periodic-advertisement threshold**: Checks the minimum-delay metric change for threshold crossing for periodic advertisement. The default value is 10 percent and the range is from 0 to 100 percent.
- **periodic-advertisement minimum-change**: The default value is 500 microseconds (usec) and the range is from 0 to 100000 microseconds.
- **accelerated advertisement**: Accelerated advertisement is disabled by default.
- **accelerated-advertisement threshold**: Checks the minimum-delay metric change for threshold crossing for accelerated advertisement. The default value is 20 percent and the range is from 0 to 100 percent.
- **accelerated-advertisement minimum**: The default value is 500 microseconds and the range is from 1 to 100000 microseconds.

```

Router(config)# performance-measurement delay-profile sr-policy
Router(config-pm-dm-srpolicy)# probe
Router(config-pm-dm-srpolicy-probe)# burst-interval 60
Router(config-pm-dm-srpolicy-probe)# computation-interval 60
Router(config-pm-dm-srpolicy-probe)# protocol twamp-light
Router(config-pm-dm-srpolicy-probe)# tos dscp 63
Router(config-pm-dm-srpolicy-probe)# exit

Router(config-pm-dm-srpolicy)# advertisement
Router(config-pm-dm-srpolicy-adv)# periodic
Router(config-pm-dm-srpolicy-adv-per)# interval 60
Router(config-pm-dm-srpolicy-adv-per)# minimum-change 1000
Router(config-pm-dm-srpolicy-adv-per)# threshold 20
Router(config-pm-dm-srpolicy-adv-per)# exit

Router(config-pm-dm-srpolicy-adv)# accelerated
Router(config-pm-dm-srpolicy-adv-acc)# minimum-change 1000
Router(config-pm-dm-srpolicy-adv-acc)# threshold 10
Router(config-pm-dm-srpolicy-adv-acc)# exit

Router(config-pm-dm-srpolicy-adv)# threshold-check minimum-delay
Router(config-pm-dm-srpolicy-adv)# exit
Router(config-pm-dm-srpolicy)#

```

### Configure the UDP Destination Port

Configuring the UDP port for TWAMP-Light protocol is optional. By default, PM uses port 862 as the TWAMP-reserved UDP destination port for delay.

The UDP port is configured for each PM measurement probe type (delay, loss, protocol, authentication mode, etc.) on querier and responder nodes. If you configure a different UDP port, the UDP port for each PM measurement probe type must match on the querier and the responder nodes.



**Note** The same UDP destination port is used for delay measurement for links and SR Policy.

This example shows how to configure the UDP destination port for delay.

```
Router(config)# performance-measurement
Router(config-perf-meas)# protocol twamp-light
Router(config-pm-protocol)# measurement delay unauthenticated
Router(config-pm-proto-mode)# querier-dst-port 12000
```

### Enable Performance Measurement for SR Policy

This example shows how to enable PM for SR policy delay for a specific policy.

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy foo
Router(config-sr-te-policy)# performance-measurement
Router(config-sr-te-policy-perf-meas)# delay-measurement
```

### SR Policy Probe IP/UDP ECMP Hashing Configuration

This example shows how to configure SR Policy ECMP IP-hashing mode.

- The destination IPv4 address 127.x.x.x – 127.y.y.y is used in the Probe messages to take advantages of 3-tuple IP hashing (source-address, destination-address, and local router ID) for ECMP paths of SR-MPLS Policy.



**Note** The destination IPv4 address must be 127/8 range (loopback), otherwise it will be rejected.

- One PM session is always created for the actual endpoint address of the SR Policy.
- You can specify the number of IP addresses to sweep. The range is from 0 (default, no sweeping) to 128.
- Platforms may have a limitation for large label stack size to not check IP address for hashing.

```
Router(config)# performance-measurement delay-profile sr-policy
Router(config-pm-dm-srpolicy)# probe
Router(config-pm-dm-srpolicy-probe)# sweep
Router(config-pm-dm-srpolicy-probe-sweep)# destination ipv4 127.0.0.1 range 28
```

**Verification**

```
Router# show performance-measurement sr-policy
Mon Jan 20 18:48:41.002 PST
```

```

0/0/CPU0

```

| Policy Name              | LSP ID | Tx/Rx | Avg/Min/Max/Variance  |
|--------------------------|--------|-------|-----------------------|
| srte_c_10_ep_192.168.0.4 | 2      | 6/6   | 27012/26906/27203/106 |

```
Router# show performance-measurement sr-policy name srte_c_10_ep_192.168.0.4 detail verbose
Mon Jan 20 18:44:22.400 PST
```

```

0/0/CPU0

```

```
SR Policy name: srte_c_10_ep_192.168.0.4
```

```
Color : 10
Endpoint : 192.168.0.4
Number of candidate-paths : 1
```

## Candidate-Path:

```
Instance : 2
Preference : 100
Protocol-origin : Configured
Discriminator : 100
Source address : 192.168.0.2
Reverse path label : Not configured
Number of segment-lists : 1
```

## Last advertisement:

```
No advertisements have occurred
```

## Next advertisement:

```
Check scheduled at the end of the current probe (roughly every 30 seconds)
Aggregated delays (uSec): avg: 45218, min: 26512, max: 82600, variance: 18706
Rolling average (uSec): 45218
```

## Last probe:

```
Packets 9, received: 9
Measured delays (uSec): avg: 45218, min: 26512, max: 82600, variance: 18706
```

## Current Probe:

```
Started at Jan 20 2020 18:44:19.170 (3.453 seconds ago)
Packets 3, received: 3 Measured delays (uSec): avg: 26588, min: 26558, max: 26630, variance: 30
```

```
Next probe scheduled at Jan 20 2020 18:44:34.166 (in 11.543 seconds)
Next burst packet will be sent in 1.543 seconds
```

```
Burst packet sent every 5.0 seconds
```

```
Liveness Detection: Disabled
```

```
Segment-List : R4
16004
```

```
Number of atomic paths : 3
```

## Last advertisement:

```
No advertisements have occurred
```

## Next advertisement:

```
Aggregated delays (uSec): avg: 45218, min: 26512, max: 82600, variance: 18706
Rolling average (uSec): 45218
```

## Last probe:

```
Packets 9, received: 9
Measured delays (uSec): avg: 45218, min: 26512, max: 82600, variance: 18706
```

## Current probe:

```
Packets 3, received: 3
Measured delays (uSec): avg: 26588, min: 26558, max: 26630, variance: 30
```

Liveness Detection: Disabled

Atomic path:

Hops : 127.0.0.0

Session ID : 33554434

Last advertisement:

No advertisements have occurred

Next advertisement:

Aggregated delays (uSec): avg: 45407, min: 26629, max: 82600, variance: 18778

Rolling average (uSec): 45407

Last Probe:

Packets 3, received: 3

Measured delays (uSec): avg: 45407, min: 26629, max: 82600, variance: 18778

Current Probe:

Packets 1, received: 1

Measured delays (uSec): avg: 26630, min: 26630, max: 26630, variance: 0

Probe samples:

| Packet Rx Timestamp      | Measured Delay (nsec) |
|--------------------------|-----------------------|
| Jan 20 2020 18:44:19.198 | 26630730              |

Liveness Detection: Disabled

Atomic path:

Hops : 127.0.0.1

Session ID : 33554435

Last advertisement:

No advertisements have occurred

Next advertisement:

Aggregated delays (uSec): avg: 45128, min: 26521, max: 81961, variance: 18607

Rolling average (uSec): 45128

Last Probe:

Packets 3, received: 3

Measured delays (uSec): avg: 45128, min: 26521, max: 81961, variance: 18607

Current Probe:

Packets 1, received: 1

Measured delays (uSec): avg: 26576, min: 26576, max: 26576, variance: 0

Probe samples:

| Packet Rx Timestamp      | Measured Delay (nsec) |
|--------------------------|-----------------------|
| Jan 20 2020 18:44:19.198 | 26576938              |

Liveness Detection: Disabled

Atomic path:

Hops : 192.168.0.4

Session ID : 33554433

Last advertisement:

No advertisements have occurred

Next advertisement:

Aggregated delays (uSec): avg: 45119, min: 26512, max: 81956, variance: 18607

Rolling average (uSec): 45119

Last Probe:

Packets 3, received: 3

Measured delays (uSec): avg: 45119, min: 26512, max: 81956, variance: 18607

Current Probe:

Packets 1, received: 1

Measured delays (uSec): avg: 26558, min: 26558, max: 26558, variance: 0

Probe samples:

| Packet Rx Timestamp      | Measured Delay (nsec) |
|--------------------------|-----------------------|
| Jan 20 2020 18:44:19.198 | 26558375              |

Liveness Detection: Disabled

Router# **show performance-measurement history probe sr-policy**

Mon Jan 20 18:46:55.445 PST

-----  
0/0/CPU0

## SR Policy End-to-End Delay Measurement

-----

SR Policy name: srte\_c\_10\_ep\_192.168.0.4

Color : 10  
Endpoint : 192.168.0.4

Candidate-Path:

Preference : 100  
Protocol-origin : Configured  
Discriminator : 100

Delay-Measurement history (uSec):

| Probe  | Start | Timestamp    | Pkt (TX/RX) | Average | Min   | Max   |
|--------|-------|--------------|-------------|---------|-------|-------|
| Jan 20 | 2020  | 18:46:34.174 | 9/9         | 26880   | 26684 | 27070 |
| Jan 20 | 2020  | 18:46:19.174 | 9/9         | 26899   | 26822 | 27004 |
| Jan 20 | 2020  | 18:46:04.173 | 9/9         | 26813   | 26571 | 27164 |
| Jan 20 | 2020  | 18:45:49.172 | 9/9         | 26985   | 26713 | 27293 |
| Jan 20 | 2020  | 18:45:34.172 | 9/9         | 26744   | 26557 | 27005 |
| Jan 20 | 2020  | 18:45:19.171 | 9/9         | 26740   | 26435 | 27093 |
| Jan 20 | 2020  | 18:45:04.171 | 9/9         | 27115   | 26938 | 27591 |
| Jan 20 | 2020  | 18:44:49.171 | 9/9         | 26878   | 26539 | 27143 |
| Jan 20 | 2020  | 18:44:34.171 | 9/9         | 26824   | 26562 | 27265 |
| Jan 20 | 2020  | 18:44:19.170 | 9/9         | 26944   | 26558 | 27422 |
| Jan 20 | 2020  | 18:44:06.543 | 9/9         | 45218   | 26512 | 82600 |

Segment-List : R4

16004

Delay-Measurement history (uSec):

| Probe  | Start | Timestamp    | Pkt (TX/RX) | Average | Min   | Max   |
|--------|-------|--------------|-------------|---------|-------|-------|
| Jan 20 | 2020  | 18:46:34.174 | 9/9         | 26880   | 26684 | 27070 |
| Jan 20 | 2020  | 18:46:19.174 | 9/9         | 26899   | 26822 | 27004 |
| Jan 20 | 2020  | 18:46:04.173 | 9/9         | 26813   | 26571 | 27164 |
| Jan 20 | 2020  | 18:45:49.172 | 9/9         | 26985   | 26713 | 27293 |
| Jan 20 | 2020  | 18:45:34.172 | 9/9         | 26744   | 26557 | 27005 |
| Jan 20 | 2020  | 18:45:19.171 | 9/9         | 26740   | 26435 | 27093 |
| Jan 20 | 2020  | 18:45:04.171 | 9/9         | 27115   | 26938 | 27591 |
| Jan 20 | 2020  | 18:44:49.171 | 9/9         | 26878   | 26539 | 27143 |
| Jan 20 | 2020  | 18:44:34.171 | 9/9         | 26824   | 26562 | 27265 |
| Jan 20 | 2020  | 18:44:19.170 | 9/9         | 26944   | 26558 | 27422 |
| Jan 20 | 2020  | 18:44:06.543 | 9/9         | 45218   | 26512 | 82600 |

Atomic path:

Hops : 127.0.0.0

Delay-Measurement history (uSec):

| Probe  | Start | Timestamp    | Pkt (TX/RX) | Average | Min   | Max   |
|--------|-------|--------------|-------------|---------|-------|-------|
| Jan 20 | 2020  | 18:46:34.174 | 3/3         | 26927   | 26747 | 27070 |
| Jan 20 | 2020  | 18:46:19.174 | 3/3         | 26982   | 26970 | 27004 |
| Jan 20 | 2020  | 18:46:04.173 | 3/3         | 26895   | 26647 | 27164 |
| Jan 20 | 2020  | 18:45:49.172 | 3/3         | 27054   | 26764 | 27293 |
| Jan 20 | 2020  | 18:45:34.172 | 3/3         | 26801   | 26694 | 27005 |
| Jan 20 | 2020  | 18:45:19.171 | 3/3         | 26807   | 26524 | 27093 |
| Jan 20 | 2020  | 18:45:04.171 | 3/3         | 27226   | 26938 | 27591 |
| Jan 20 | 2020  | 18:44:49.171 | 3/3         | 26976   | 26644 | 27143 |
| Jan 20 | 2020  | 18:44:34.171 | 3/3         | 26880   | 26679 | 27265 |
| Jan 20 | 2020  | 18:44:19.170 | 3/3         | 26994   | 26630 | 27422 |
| Jan 20 | 2020  | 18:44:06.543 | 3/3         | 45407   | 26629 | 82600 |

Atomic path:

Hops : 127.0.0.1

Delay-Measurement history (uSec):

| Probe  | Start | Timestamp    | Pkt (TX/RX) | Average | Min   | Max   |
|--------|-------|--------------|-------------|---------|-------|-------|
| Jan 20 | 2020  | 18:46:34.174 | 3/3         | 26865   | 26705 | 26988 |
| Jan 20 | 2020  | 18:46:19.174 | 3/3         | 26846   | 26822 | 26881 |
| Jan 20 | 2020  | 18:46:04.173 | 3/3         | 26787   | 26581 | 26939 |
| Jan 20 | 2020  | 18:45:49.172 | 3/3         | 26954   | 26728 | 27180 |

|                          |     |       |       |       |
|--------------------------|-----|-------|-------|-------|
| Jan 20 2020 18:45:34.172 | 3/3 | 26724 | 26577 | 26957 |
| Jan 20 2020 18:45:19.171 | 3/3 | 26705 | 26452 | 27032 |
| Jan 20 2020 18:45:04.171 | 3/3 | 27043 | 26972 | 27124 |
| Jan 20 2020 18:44:49.171 | 3/3 | 26848 | 26550 | 27062 |
| Jan 20 2020 18:44:34.171 | 3/3 | 26800 | 26562 | 27204 |
| Jan 20 2020 18:44:19.170 | 3/3 | 26927 | 26576 | 27327 |
| Jan 20 2020 18:44:06.543 | 3/3 | 45128 | 26521 | 81961 |

Atomic path:

Hops : 192.168.0.4

Delay-Measurement history (uSec):

| Probe Start Timestamp    | Pkt (TX/RX) | Average | Min   | Max   |
|--------------------------|-------------|---------|-------|-------|
| Jan 20 2020 18:46:34.174 | 3/3         | 26848   | 26684 | 26967 |
| Jan 20 2020 18:46:19.174 | 3/3         | 26871   | 26833 | 26913 |
| Jan 20 2020 18:46:04.173 | 3/3         | 26759   | 26571 | 26876 |
| Jan 20 2020 18:45:49.172 | 3/3         | 26947   | 26713 | 27163 |
| Jan 20 2020 18:45:34.172 | 3/3         | 26708   | 26557 | 26939 |
| Jan 20 2020 18:45:19.171 | 3/3         | 26708   | 26435 | 27075 |
| Jan 20 2020 18:45:04.171 | 3/3         | 27078   | 27016 | 27138 |
| Jan 20 2020 18:44:49.171 | 3/3         | 26812   | 26539 | 27043 |
| Jan 20 2020 18:44:34.171 | 3/3         | 26793   | 26582 | 27181 |
| Jan 20 2020 18:44:19.170 | 3/3         | 26911   | 26558 | 27308 |
| Jan 20 2020 18:44:06.543 | 3/3         | 45119   | 26512 | 81956 |

Router# **show performance-measurement counters sr-policy name srte\_c\_10\_ep\_192.168.0.4**  
 Mon Jan 20 18:47:55.499 PST

-----  
0/0/CPU0  
-----

SR Policy name: srte\_c\_10\_ep\_192.168.0.4

Candidate-Path:

|                 |              |
|-----------------|--------------|
| Instance        | : 2          |
| Preference      | : 100        |
| Protocol-origin | : Configured |
| Discriminator   | : 100        |

Packets:

|                |       |
|----------------|-------|
| Total sent     | : 141 |
| Total received | : 141 |

Errors:

|                       |     |
|-----------------------|-----|
| Total sent errors     | : 0 |
| Total received errors | : 0 |

Probes:

|                      |      |
|----------------------|------|
| Total started        | : 16 |
| Total completed      | : 15 |
| Total incomplete     | : 0  |
| Total advertisements | : 2  |

Segment-List : R4

16004

Packets:

|                |       |
|----------------|-------|
| Total sent     | : 141 |
| Total received | : 141 |

Errors:

|                       |     |
|-----------------------|-----|
| Total sent errors     | : 0 |
| Total received errors | : 0 |

Probes:

|                      |      |
|----------------------|------|
| Total started        | : 16 |
| Total completed      | : 15 |
| Total incomplete     | : 0  |
| Total advertisements | : 2  |



# Fallback delay advertisement

Table 72: Feature History Table

| Feature Name                 | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fallback delay advertisement | Release 24.4.1      | <p>You can now advertise fallback delay value, retaining delay information in performance metrics even when delay metrics for interfaces are temporarily unavailable due to hardware, synchronization, or network connectivity issues. The feature ensures optimal routing decisions by maintaining network stability and consistent performance, even when real-time metrics are temporarily unavailable.</p> <p>Previously, the performance metrics did not include delay metrics when they were temporarily inaccessible, resulting in visibility gaps in the network and less effective routing.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"><li>• <a href="#">Cisco IOS XR performance measurement of yang</a></li></ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

## Fallback delay advertisement for interfaces

The feature enables network operators to handle situations where GRE tunnel delay interface delay metric is not available due to hardware, synchronization, or connectivity issues. Instead of removing the interface delay metrics from the network performance metrics, which might cause network visibility loss and potential suboptimal routing decisions, network operators can configure and advertise the fallback delay value. The network can make routing decisions based on the fallback delay value instead of assuming that the link is entirely unusable.

The feature provides a backup plan for hardware failures or other issues that prevent accurate delay measurement for GRE tunnel interfaces, ensuring operational continuity by avoiding the removal of delay metrics.

### Fallback delay mechanisms for seamless routing

With this feature, the network makes informed routing decisions by advertising either a maximum metric value with an anomaly flag or a user-configured fallback value.

- **Maximum Metric with an Anomaly Flag:** When you cannot measure the delay data for an interface due to hardware issues, connectivity problems, by default the performance measurement system advertises the highest possible metric value for that interface. The default fallback value advertised is 16777215 micro seconds. Also, it tags this metric with an anomaly flag, alerting the network that there is an issue with real-time delay data.
- **User-Configured Fallback Value:** When you configure a fallback value, you can advertise this fallback delay value instead of the maximum metric. The option is useful in cases where a predictable, user-defined delay value can help maintain stable routing, avoiding drastic changes in network behavior due to temporary loss of delay data.

### Benefits of fallback delay advertisement for interfaces

- **Resilience:** Ensures that the network continues to make optimal routing decisions when delay metrics are not available due to hardware failures or transient session errors.
- **Predictability:** By configuring a fallback value, network administrators can ensure that the network behaves predictably in the absence of real-time delay metrics.
- **Continuity:** Prevents the removal of delay metrics, maintaining continuous visibility and routing capability.

## Configure interface fallback delay value

Perform the following steps to configure a specific fallback delay value for interfaces:

### Procedure

**Step 1** Enable SR-PM.

**Example:**

```
Router#config
Router(config)#performance-measurement
```

**Step 2** Set the advertise-delay fallback value for the interface to ensure that the router advertises this value when the computed delay metric is unavailable.

In the following example, the advertised interface fallback delay value is **1000**.

**Example:**

```
Router(config-perf-meas)#interface GigabitEthernet 0/2/0/0
Router(config-pm-intf)#delay-measurement
Router(config-pm-intf-dm)#advertise-delay fallback 1000
Router(config-pm-intf-dm)#commit
```

**Step 3** Run the **show running-config** command to verify the running configuration.

**Example:**

```

!
performance-measurement
interface GigabitEthernet0/2/0/0
 delay-measurement
 advertise-delay fallback 1000
!
!
!

```

**Step 4** Run the **show performance-measurement sessions** command to verify whether the advertised metric is the default maximum metric value with an anomaly flag or the configured fallback delay value.

In the following show output, when the performance measurement session begins, by default the interface advertises the maximum metric value and the delay A flag is set.

**Example:**

```

Router#show performance-measurement sessions
Transport type : Interface
Measurement type : Delay Measurement
Interface name : GigabitEthernet0/2/0/0
Nexthop : Unknown
Delay Measurement session:
 Session ID : 4097
 Timestamp source : Hardware (local)
 Timestamp format : PTP
 Profile Keys:
 Profile name : default
 Profile type : Interface Delay Measurement
 Last advertisement:
 Advertised at: May 28 2024 13:53:41.117 (51.85 seconds ago)
 Advertised reason: Interface initial advertisement
 Advertised delays (uSec): avg: 16777215, min: 16777215, max: 16777215, variance: 0
 Packets Sent: 0, received: 0
 Min-Max A flag set: True
 Anomaly-Loss A flag set: False

 Next advertisement:
 Check scheduled at the end of the current computation (roughly every 60 seconds)
 Aggregated delays (uSec): avg: 1287, min: 623, max: 21048, variance: 628
 Packets Sent: 50, received: 50
 Rolling average (uSec): 991

 Current computation:
 Started at: May 28 2024 13:54:30.581 (1.621 seconds ago)
 Packets Sent: 2, received: 2
 Measured delays (uSec): avg: 704, min: 683, max: 724, variance: 21
 Next probe scheduled at: May 28 2024 13:54:40.570 (in 8.368 seconds)
 Next packet will be sent in 0.368 seconds
 Packet sent every 1.0 seconds
 Responder IP : 10.10.10.2
 Number of Hops : 1

```

# Two-Way Active Measurement Protocol Light Source Address Filtering

Table 73: Feature History Table

| Feature Name                                                       | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Two-Way Active Measurement Protocol Light Source Address Filtering | Release 7.11.1      | <p>You can now restrict unauthorized users from sending packets to the network and prevent compromising the network security and reliability. For a destination UDP port, you can configure the list of IP addresses that can send Two-Way Active Measurement Protocol (TWAMP)-light packets to responder or querier nodes.</p> <p>In earlier releases, the responder or querier node accepted TWAMP-light packets from all IP addresses.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>querier</b> and <b>responder</b> keywords are introduced in the <b>performance-measurement protocol twamp-light measurement delay</b> command.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>Cisco-IOS-XR-performance-meas-cfg.yang</li> <li>Cisco-IOS-XR-perf-meas-oper.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

Earlier, the responder node scanned all IP addresses of a querier on the destination UDP port. In other words, the responder node accepted packets from any IP address. See [Measurement Modes](#) for more information about querier and responder nodes.



**Note**

The responder is also called the reflector, and the querier is also called the sender.

With this configuration, you can specify the source IP addresses on both the responder and querier nodes. The responder or querier nodes accept packets only from the IP addresses configured in the TWAMP-light protocol, and reject the packets from an IP address that isn't included in the configured list.

All the configured addresses are available for use on all interfaces in the Local Packet Transport Services (LPTS). The configured address filter applies to both default and nondefault VRFs. The TWAMP delay measurement sessions use the configured addresses.

## Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- When you configure the prefix entries on the responder or querier nodes, the PM adds the responder or querier node source IP address to the LPTS. For each prefix, a new LPTS entry is added or created.
- For TWAMP liveness sessions, the PM automatically adds the source IP addresses to the LPTS for if you have configured the prefix entries on the responder or querier nodes.
- As the maximum number of LPTS hardware entries are limited, ensure that enough LPTS entries are allocated for the IP addresses on a line card. You can scale the LPTS configuration to maximum LPTS entries for the PM flow-type. For more details on configuring the LPTS entries for PM flow-type, refer to the *IP Addresses and Services Configuration Guide*.



**Note** The PM UDP port accepts all incoming IPv4 or IPv6 packets when there are no IPv4 or IPv6 prefix entries configured.

## Configure IP address on querier and responder nodes

- The length of the IPv4 and IPv6 prefixes must be less than 32 and 128 respectively.
- The length or mask of the source IP address must be:
  - For IPv4: 0–31
  - For IPv6: 0–127

### Configure the IP address on a responder

Perform this task to configure the IP address of a querier on a responder node for delay measurement.

```
Router#configure
Router(config)#performance-measurement
Router(config-perf-meas)#protocol twamp-light
Router(config-pm-protocol)#measurement delay
Router(config-pm-proto-meas)#responder
Router(config-pm-proto-responder)#allow-querier
Router(config-pm-allowed-querier)#address ipv4 10.10.10.1
```

### Running Configuration

```
performance-measurement
protocol twamp-light
```

```

measurement delay
 responder
 allow-querier
 address ipv4 10.10.10.1
!
!
!
!
End

```

### Verification

The following example shows output from the IP address of a querier, which is configured on a responder node for delay measurement.

```

Router#show performance-measurement allowed-querier summary
Wed Oct 11 10:41:43.268 UTC

```

```

0/RP0/CPU0

Allowed-querier IPv4 prefix : 1
 10.10.10.1/32
Allowed-querier IPv6 prefix : 0

RX UDP port status:
 TWAMP-Light Default Unauthenticated responder port : 862
 Opened IPv4 port : 862
 IPv4 Port Update Time : Oct 11 2023 10:37:48.118
 Opened IPv6 port : 862
 IPv6 Port Update Time : Oct 11 2023 10:37:47.778

```

### Configure the source IP address on a querier

Perform this task to configure the IP address of a responder on a querier node for delay measurement.

```

Router#configure
Router(config)#performance-measurement
Router(config-perf-meas)#protocol twamp-light
Router(config-pm-protocol)#measurement delay
Router(config-pm-proto-meas)#querier
Router(config-pm-proto-querier)#allow-responder
Router(config-pm-allowed-responder)#address ipv4 10.10.10.1

```

### Running Configuration

```

performance-measurement
 protocol twamp-light
 measurement delay
 querier
 allow-responder
 address ipv4 10.10.10.1
!
!
!
!
End

```

## Verification

The following example shows output from the IP address of a responder, which is configured on a querier node for delay measurement.

```
Router#show performance-measurement allowed-responder summary
Wed Mar 29 19:38:06.381 UTC
```

---

```
0/RP1/CPU0
```

---

```
Allowed-responder IPv4 prefix : 2
 10.10.10.1/32 [Auto]
 3.3.3.3/32
Allowed-responder IPv6 prefix : 1
 fc00:0:1::1/128 [Auto] [Pending Add]
Querier CPU UDP port status:
 TWAMP-Light Default Unauthenticated querier port : N/A
 Opened IPv4 port : 27643
 IPv4 Port Update Time : Mar 29 2023 18:43:49.080
 Opened IPv6 port : 28274
 IPv6 Port Update Time : Mar 24 2023 20:58:46.150
```

# Synthetic Loss Measurement

Table 74: Feature History Table

| Feature Name               | Release        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Synthetic Loss Measurement | Release 24.1.1 | <p>You can now proactively monitor and address potential network issues before they impact users by measuring key parameters everywhere, packet loss, and jitter. Using this information, you can plan network capacity optimally and ensure quality of service. Such proactive action is possible because this feature reports synthetic Two-Way Active Measurement Protocol (TWAMP) test packets deployed in delay-profile or delay measurement sessions.</p> <p>It also enables you to set the upper and lower limits and notifies when the synthetic packet loss metric is out of the set limit.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• The optional <b>anomaly-loss</b> and <b>anomaly-check</b> keywords are introduced in the <a href="#">performance-measurement delay-profile</a> command.</li> <li>• <a href="#">show performance-measurement history</a></li> </ul> <p><b>YANG Data Model</b></p> <ul style="list-style-type: none"> <li>• New XPaths for<br/><code>Cisco-IOS-XR-um-performance-measurement-cfg</code></li> <li>• New operations for<br/><code>Cisco-IOS-XR-perf-meas-oper.yang</code><br/>(see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> |

Synthetic packets are data packets that are artificially generated for the purpose of testing or managing a network, as opposed to being generated by normal network data traffic. These packets are unrelated to regular user activities or data. Instead, they are intentionally designed to serve monitoring, testing, and diagnostic purposes.

Synthetic TWAMP test packets deployed in delay profile are used to simulate network traffic, test network performance, and evaluate network infrastructure. They are particularly useful for assessing the quality of service (QoS) of a network, identifying network vulnerabilities, and troubleshooting network issues. These packets enable network administrators to identify potential problems before they impact the overall network performance and reliability. *For more information on TWAMP, refer to the IP Service Level Agreements chapter in the System Monitoring Configuration Guide for Cisco NCS 560 Series Routers.*



Overall, synthetic packets play a critical role in network management, providing network engineers with a powerful tool for optimizing network performance, identifying and resolving issues.



**Note** This is an inbuilt feature of delay measurement. To get the synthetic packet loss information for delay-measurement sessions, you must only configure the delay sessions. No additional configuration is required for Synthetic Loss Measurement. However, an optional **anomaly-loss** command is introduced.

Synthetic packets simulate various types of traffic that can be sent with the normal traffic to measure loss or latency, throughput or packet loss. This is often done in network performance monitoring and testing, where the goal is to understand how the network or a specific network device performs under typical conditions.

This feature enables you to measure and monitor the Synthetic packets lost and help make informed decisions on the following:

- **Network Performance Monitoring:** Synthetic packets can be used to simulate different types of network traffic in order to measure and test the performance of network devices or the network itself.
- **Troubleshooting:** They can be used to diagnose problems in a network by sending them through different parts of the system to see where they may fail or be slowed down.
- **Security:** They can also be used in security applications, such as penetration testing, where synthetic packets are sent to a system to find vulnerabilities.

## Configure Synthetic Loss measurement

The following example enables Synthetic Loss Measurement for an SR Policy.



**Note** A delay-measurement session is good enough to get the synthetic packet loss automatically, without any extra configuration. The configure shown here is for anomaly loss, which is optional.

```
Router(config)#performance-measurement
Router(config-perf-meas)#delay-profile sr-policy default
Router(config-pm-dm-srpolicy)#advertisement
Router(config-pm-dm-srpolicy-adv)#anomaly-loss
Router(config-pm-dm-srpolicy-adv-anom-loss)#upper-bound 30 lower-bound 20
Router(config-pm-dm-srpolicy-adv-anom-loss)#commit
```



**Note** Similarly, you can configure Synthetic Loss Measurement for **endpoint default**, **interface default** or **name** (named profile).

### Running Configuration

The running configuration for this feature is as shown:

```
performance-measurement
 delay-profile sr-policy default
 advertisement
```

```

 anomaly-loss
 upper-bound 30 lower-bound 20
 !
 !
!
!
!

```

## Verification

Use the show commands to verify the running configuration as shown:

**Router# show performance-measurement interfaces delay detail**

Interface Name: GigabitEthernet0/2/0/0 (ifh: 0x1000020)

```

...
 Last advertisement:
 ...
 Packets 40, received: 40
 ...
 Next advertisement:
 ...
 Packets 10, received: 10
 ...
 Current computation:
 ...
 Packets 4, received: 4

```

**Router# show performance-measurement history probe-computation interfaces**

Interface Name: GigabitEthernet0/2/0/0 (ifh: 0x1000020)

Delay-Measurement history (uSec):

| Probe       | Start        | Timestamp | Pkt(TX/RX) | Average | Min | Max |
|-------------|--------------|-----------|------------|---------|-----|-----|
| Aug 01 2023 | 08:04:15.230 |           | 10/10      | 704     | 651 | 779 |

## Endpoint

Use these show commands to verify the Endpoint configuration. The example include endpoint delay details and advertisement history. You can also verify the aggregation history.

**Router# show performance-measurement endpoint delay detail**

Endpoint name: IPv4-192.168.0.4-vrf-default

```

...
 Last advertisement:
 ...
 Packets 40, received: 40
 ...
 Next advertisement:
 ...
 Packets 30, received: 30
 ...
 Current computation:
 ...
 Packets 6, received: 6
 ...

```

**Router# show performance-measurement history advertisement endpoint**

Endpoint name: IPv4-192.168.0.4-vrf-default

```

...
 Delay-Measurement history (uSec):

```

| Advertisement Timestamp  | Pkt (TX/RX) | Average | Min  | Max   | Reason  |
|--------------------------|-------------|---------|------|-------|---------|
| Aug 01 2023 08:31:18.835 | 40/40       | 3948    | 3127 | 15503 | PER-MAX |

## RSVP-TE

Use these show commands to verify the RSVP-TE configuration. The example include rsvp-te delay details and the aggregation history. You can also verify the advertisement history.

**Router# show performance-measurement rsvp-te detail**

Tunnel name: tunnel-te1 (ifh: 0xd0)

```
...
 Last advertisement:
 ...
 Packets 40, received: 40
 ...
 Next advertisement:
 ...
 Packets 10, received: 10
 ...
 Last computation:
 ...
 Packets 10, received: 10
 ...
 Current computation:
 ...
 Packets 6, received: 6
 ...
...
```

**Router# show performance-measurement history aggregation rsvp-te**

```
...
Delay-Measurement history (uSec):
 Aggregation Timestamp Pkt (TX/RX) Average Min Max
 Aug 01 2023 08:37:23.702 40/40 3372 3172 4109
...
```

## SR-Policy

Use these show commands to verify the SR-Policy configuration. The example include delay details and advertisement history. You can also verify the aggregation history.

**Router# show performance-measurement sr-policy delay detail**

SR Policy name: srte\_c\_10\_ep\_192.168.0.4

```
...
 Last advertisement:
 ...
 Packets 88, received: 88
 ...
 Next advertisement:
 ...
 Packets 132, received: 132
 ...
 Last computation:
 ...
 Packets 4, received: 4
 ...
 Current computation:
 ...
 Packets 4, received: 4
 ...
 Segment-List : R4
 ...
 Last advertisement:
```

```

...
Packets 88, received: 88
...
Next advertisement:
...
Packets 132, received: 132
...
Last computation:
Packets 4, received: 4
...
Current computation:
Packets 4, received: 4
...

```

**Router# show performance-measurement history advertisement sr-policy**

```

...
Delay-Measurement history (uSec):
 Advertisement Timestamp Pkt (TX/RX) Average Min Max Reason
Aug 01 2023 10:05:14.072 24/24 3408 3408 3408 ACCEL-MAX

```

### Packet loss advertisement

Use this show command to verify multiple configurations with a single show command:

**Router# show performance-measurement history advertisement sr-policy/endpoint**

```

...
Delay-Measurement history (uSec):
 Advertisement Timestamp Pkt (TX/RX) Average Min Max Reason
Sep 14 2023 11:35:38.180 10/9 3595 3411 3696 NEW-SESSION

 Sep 14 2023 11:34:38.178 10/0 0 0 0 SESSION-ERROR

 Sep 14 2023 11:34:08.177 10/7 3733 3733 3733 ANOM-PKT-LOSS

 Sep 14 2023 11:34:02.177 8/7 3733 3733 3733 PKT-LOSS
 Sep 14 2023 11:33:38.176 10/10 3823 3617 4627 FIRST
...

```

**Router# show performance-measurement sessions**

```

...
Last advertisement:
 Advertised at: Sep 14 2023 08:47:16.540 (145.777 seconds ago)
 Advertised reason: Periodic timer, max delay threshold crossed
 Advertised delays (uSec): avg: 5373, min: 3992, max: 26212, variance: 0
 Packets 30, received: 30
 Min-Max A flag set: False
 Loss A flag set: False
...
...
Last advertisement:
 Advertised at: Sep 14 2023 08:52:39.603 (12.522 seconds ago)
 Advertised reason: PM session anomaly due to packet loss
 Advertised delays (uSec): avg: 5373, min: 3992, max: 26212, variance: 0
 Packets 30, received: 30
 Min-Max A flag set: False
 Loss A flag set: True
...
Last advertisement:
 Advertised at: Sep 15 2023 11:42:41.594 (47.660 seconds ago)
 Advertised reason: Performance measurement session error
...

```



**Note** No metrics or A bit is shown when there is a session error.

**Router# show logging**

```
RP/0/0/CPU0:Sep 14 11:33:38.176 PDT: perf_meas[1001]:
%ROUTING-PERF_MEAS-5-PM_DELAY_EXCEEDS_THRESHOLD : Reason: First advertisement, PM delay
metric avg: 3823, min: 3617, max: 4627, var: 138, min-max anomaly flag: Not set, pkt-loss
anomaly flag: Not set, exceeded threshold on SR Policy N:srte_c_10_ep_192.168.0.4, L:2
RP/0/0/CPU0:Sep 14 11:34:08.177 PDT: perf_meas[1001]:
%ROUTING-PERF_MEAS-5-PM_DELAY_EXCEEDS_THRESHOLD : Reason: PM session anomaly due to packet
loss, PM delay metric avg: 0, min: 0, max: 0, var: 0, min-max anomaly flag: Not set,
pkt-loss anomaly flag: Set, exceeded threshold on SR Policy N:srte_c_10_ep_192.168.0.4, L:2

RP/0/0/CPU0:Sep 14 11:35:38.180 PDT: perf_meas[1001]:
%ROUTING-PERF_MEAS-5-PM_DELAY_EXCEEDS_THRESHOLD : Reason: PM session anomaly due to packet
loss, PM delay metric avg: 3595, min: 3411, max: 3696, var: 58, min-max anomaly flag: Not
set, pkt-loss anomaly flag: Not set, exceeded threshold on SR Policy
N:srte_c_10_ep_192.168.0.4, L:2
```

=====

In this scenario, the router was running normally and then advertised a packet loss anomaly, the flag is set; later the router resumed to normal, the flag is unset, and it advertised a packet loss anomaly again.

# Delay and synthetic loss measurement for GRE tunnel interfaces

Table 75: Feature History Table

| Feature Name                                                   | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Delay and synthetic loss measurement for GRE tunnel interfaces | Release 24.4.1      | <p>You can now measure the latency or delay experienced by data packets when they traverse a network, and also proactively monitor and address potential network issues before they impact users by measuring key parameters such as packet loss, and jitter for GRE tunnel interfaces.</p> <p>This feature enables you to report synthetic Two-Way Active Measurement Protocol (TWAMP) test packets that are deployed in delay-profile or delay measurement sessions, and enables delay measurement for GRE tunnel interfaces.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>performance-measurement interface</b> command supports the <b>tunnel-ip</b> keyword.</p> |

## Delay and synthetic loss measurement for GRE tunnel interfaces

With this feature, you can configure a GRE interface with delay measurement under PM. You can now run sessions over a GRE tunnel as an interface session, and the PM advertises these measurement metrics to the Interface Manager (IM).

The PM for delay measurement uses the IP/UDP packet format that is defined in RFC 5357 (TWAMP-Light) for probes.

Synthetic packets are data packets that are artificially generated for testing or managing a network, as opposed to being generated by normal network data traffic. These packets are unrelated to regular user activities or data. Instead, they are intentionally designed to serve monitoring, testing, and diagnostic purposes. With synthetic loss measurement, you can set the upper and lower limits and notify when the synthetic packet loss metric is out of the set limit.

For more information about delay and synthetic loss measurement, see the [Delay Measurement, on page 463](#) and [Synthetic Loss Measurement, on page 511](#) sections.

**Measurement modes for GRE tunnel interfaces**

GRE tunnel interfaces support these modes to measure delay and synthetic loss:

- One-way
- Two-way

Both one-way and two-way measurement modes support round-trip synthetic loss measurement.

## Usage guidelines and limitations for delay and synthetic loss measurement for GRE tunnels

### Configure delay and synthetic loss measurement for GRE tunnel interfaces

#### Procedure

**Step 1** Run the **interface tunnel-ip** command to configure the tunnel mode, source, and destination for the interface.

**Note**

You must configure the responder node to decapsulate the GRE tunnel and establish a proper return path to the querier node.

In the following example, the tunnel uses GRE in IPv4 mode, encapsulating or decapsulating IP packets for transport across the tunnel. The source IP address of the tunnel is set to 209.165.201.2, indicating the origin of the IP packets. The destination IP address is set to 209.165.201.3, indicating where the encapsulated packets are sent.

**Example:**

```
Router#config
Router(config)#interface tunnel-ip 23
Router(config-if)#ipv4 address 209.165.201.1/27
Router(config-if)#ipv6 address 2001:DB8::/32
Router(config-if)#tunnel mode gre ipv4 encap
Router(config-if)#tunnel mode gre ipv4 decap
Router(config-if)#tunnel source 209.165.201.2
Router(config-if)#tunnel destination 209.165.201.3
Router(config-pm-intf-dm)#commit
```

**Step 2** Enable delay measurement on the configured GRE tunnel interface.

**Example:**

```
Router(config)#performance-measurement
Router(config-perf-meas)#interface tunnel-ip 23
Router(config-pm-intf)#delay-measurement
Router(config-pm-intf-dm)#commit
```

**Step 3** Run the **show running-config** command to verify the running configuration.

**Example:**

```
!
interface tunnel-ip23
```

```

ipv4 address 209.165.201.1 255.255.255.224
ipv6 address 2001:db8::/32
tunnel mode gre ipv4 decap
tunnel source 209.165.201.2
tunnel destination 209.165.201.3
!
performance-measurement
 interface tunnel-ip23
 delay-measurement
 !
!
!

```

**Step 4** Run the **performance-measurement delay-profile** command to configure synthetic loss measurement for the GRE tunnel interface.

For more information about configuring synthetic loss measurement, see the [Configure Synthetic Loss measurement, on page 512](#) section.

**Step 5** Run the **show performance-measurement interfaces delay** command to verify the delay and synthetic loss measurement that is enabled for the GRE tunnel interface.

**Example:**

```

Router#show performance-measurement interfaces delay detail
Mon Nov 4 12:44:56.849 UTC

```

---

```

0/RP0/CPU0

```

---

```

Interface Name: tunnel-ip23 (ifh: 0x78000014)
 Delay-Measurement : Enabled
 Loss-Measurement : Enabled
 Path-Tracing : Disabled
 Configured IPv4 Address : 209.165.201.1
 Configured IPv6 Address : 2001:db8::
 Link Local IPv6 Address : fe80::657e:adff:feec:d876
 Configured Next-hop Address : Unknown
 Local MAC Address : 0000.0000.0000
 Next-hop MAC Address : 0000.0000.0000
 In-use Source Address : Unknown
 In-use Destination Address : Unknown
 Primary VLAN Tag : None
 Secondary VLAN Tag : None
 State : Up

Delay Measurement session:
 Session ID : 4097
 Timestamp source : Hardware (local)
 Timestamp format : PTP
 Profile Keys:
 Profile name : default
 Profile type : Interface Delay Measurement
 Last advertisement:
 Advertised at: Nov 04 2024 11:26:41.674 (4695.200 seconds ago)
 Advertised delays (uSec): avg: 0, min: 0, max: 0, variance: 0
 Packets sent: 40, received: 40
 Min-Max A flag set: False
 Anomaly-Loss A flag set: False

Next advertisement:
 Threshold check scheduled in 1 more probe (roughly every 120 seconds)
 Aggregated delays (uSec): avg: 2265, min: 1967, max: 4572, variance: 248

```



```
Packets Sent: 20, received: 20
Rolling average (uSec): 2219
```

```
Current computation:
Started at: Nov 04 2024 11:31:59.071 (14.416 seconds ago)
Packets Sent: 5, received: 5
Measured delays (uSec): avg: 2405, min: 2232, max: 2714, variance: 173
Next probe scheduled at: Nov 04 2024 11:32:29.053 (in 15.566 seconds)
Next packet will be sent in 0.566 seconds
Packet sent every 3.0 seconds
Responder IP : 192::1
Number of Hops : 1
```

## Delay Measurement Using Software Timestamp

**Table 76: Feature History Table**

| Feature Name                               | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Delay Measurement Using Software Timestamp | Release 24.3.1      | <p>You can now identify performance issues caused by the network, disk I/O, processing, or other factors using software timestamping on your router by measuring the delay and loss of each network path, even if the existing hardware lacks timestamp support.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>timestamp-format NTP</b> keyword is introduced in the <b>performance-measurement delay-profile</b> command.</li> </ul> |

### Timestamps for Delay Measurement

The Performance Measurement (PM) is a crucial mechanism that measures the delay and loss on network paths. It measures the delay and loss by sending and receiving timestamped probe packets at the ingress and egress nodes.

PM uses probe packets to carry specific Type-Length-Value (TLVs) to encode the measurement information for delay measurement. One of the TLVs used for delay measurement is the Timestamp TLV that contains the time when a node sends or receives the probe packet.

The Timestamp TLV can have different formats, such as Precision Time Protocol (PTP) or Network Time Protocol (NTP).

You need to insert the Timestamp TLV close to the physical interface where the packet is transmitted or received to accurately measure the delay.

You can measure delay and loss using hardware timestamp. If the hardware does not support the timestamp capability and the corresponding format and offset, software timestamp is an option that you need to consider.

Before Release 24.3.1, only hardware timestamping was supported.

### Software Timestamps for Delay Measurement

A software timestamp is a method that uses software for the insertion of Type-Length-Value (TLVs) into probe packets..

In instances where the hardware device does not support timestamping, we enable the router to set the timestamp in software, which helps insert the Type-Length-Value (TLVs) into the probe packets. The default timestamp format is PTP. You can configure to use the NTP format using the **timestamp-format NTP** command.

### Measurement Modes for Software Timestamp

The software timestamp supports these three modes:

- One-way measurement mode
- Two-way measurement mode
- Loopback measurement mode

For more information, see [Measurement Modes, on page 464](#).

### Supported Features for Software Timestamp

The software timestamp supports these performance measurement features:

- Link delay measurement
- SR Policy
- IP Endpoint with both MPLS and SRv6
- RSVP-TE LSPs

## Key Concepts

- **Delay Measuremnt** : Delay measuremnt is a mechanism used to measure the latency or delay experienced by data packets when they traverse a network. For more information, see [Delay Measurement](#).
- **NTP** : Network Time Protocol (NTP) provides time synchronization to all devices on a network. The primary NTP servers are synchronized to a reference clock, such as GPS receivers and telephone modem services.

For more information, see *System Management Configuration Guide for Cisco NCS 560 Series Routers*

- **PTP**: Precision Time Protocol (PTP) defines a method for distributing time across a network with its foundation based on the IEEE 1588-2008 standard.

For more information, see *System Management Configuration Guide for Cisco NCS 560 Series Routers*

## Benefits of Software Timestamp

Software Timestamp helps you

- provide a software-based timestamping fallback option when the hardware lacks support for the necessary timestamping capability, format, or offset.
- insert the timestamp to probe packets when a hardware device cannot support the required format or offset.
- enable the packet loss measurement automatically since the software can also insert the Counter TLV, which contains the sequence number of the probe packet.

## How Delay is Measured Using Software Timestamp?

Here is how you measure a delay using software timestamp:

1. **Event Initiation:** When a client sends a packet, the current time is recorded as a timestamp using PTP or NTP.
2. **Event Completion:** When the client receives an acknowledgment from the server, another timestamp is recorded.
3. **Delay Calculation:** The delay is calculated by subtracting the first event's timestamp from the second event's timestamp. This calculation provides the round-trip time (RTT) or one-way delay, depending on the specific context.
4. **Accuracy and Precision:** The accuracy of software timestamps relies on the resolution of the system clock and any potential delays introduced by the software layers handling the timestamping.

## Guidelines and Limitations for Software Timestamp

Here are some of the guidelines and limitations that you need to keep in mind before you configure software timestamp:

- For devices that support hardware timestamping, the hardware timestamp has priority over the software timestamp.
- One-way delay measurement is not supported when there is a mix of hardware and software timestamping.
- Delay Measurement using Software Timestamp is not supported on the Compatibility mode of NCS 5500 modular routers using NCS 5500 line cards and NCS 5700 line cards.

## Configure Software Timestamp for Delay Measurement

The default software timestamp format is PTP. You can enable the router to use the NTP format using the **timestamp-format NTP** command.

You can configure software timestamp for

- Link delay

- SR Policy
- IP Endpoint with both MPLS and SRv6, and
- RSVP-TE LSPs.

## Procedure

**Step 1** Use the **timestamp-format NTP** command to configure software timestamp with NTP format on a router.

- This example shows how to configure a software timestamp for IP Endpoint.

```
Router#configure
Router(config)# performance-measurement
Router(config-perf-meas)#delay-profile endpoint default
Router(config-pm-dm-ep)#probe
Router(config-pm-dm-ep-probe)#timestamp-format NTP
Router(config-pm-dm-ep-probe)#
```

- This example shows how to configure software timestamp for SR policy.

```
Router#configure
Router(config)# performance-measurement
Router(config-perf-meas)#delay-profile sr-policy default
Router(config-pm-dm-srpolity)#probe
Router(config-pm-dm-srpolity-probe)#timestamp-format NTP
```

- This example shows how to configure software timestamp for delay profile name.

```
Router#configure
Router(config)# performance-measurement
Router(config-perf-meas)#delay-profile name blue
Router(config-pm-dm-profile)#probe
Router(config-pm-dm-probe)#timestamp-format NTP
Router(config-pm-dm-probe)#
```

- This example shows how to configure software timestamp for interface delay profile.

```
Router#configure
Router(config)# performance-measurement
Router(config-perf-meas)#delay-profile interfaces default
Router(config-pm-dm-intf)#probe
Router(config-pm-dm-intf-probe)#timestamp-format NTP
Router(config-pm-dm-intf-probe)#
```

- This example shows how to configure software timestamp for RSVP-TE LSPs.

```
Router#configure
Router(config)# performance-measurement
Router(config-perf-meas)#delay-profile rsvp-te default
Router(config-pm-dm-rsvpte)#probe
Router(config-pm-dm-rsvpte-probe)#timestamp-format NTP
Router(config-pm-dm-rsvpte-probe)#
```

**Step 2** Execute the **show performance-measurement sessions** command to verify the timestamp format.

- This example shows how to verify the timestamp format for the IP endpoint:

```
Router# show performance-measurement endpoint detail
Endpoint name: IPv4-192.168.0.4-vrf-default
Source address : 192.168.0.2
```

```

Destination Port : 862
Source Port : 35923
VRF name : default
Delay-measurement : Enabled
Description : Not set
Profile Keys:
 Profile name : default
 Profile type : Endpoint Delay Measurement

Segment-list : None
Delay Measurement session:
 Session ID : 4275
Timestamp source : Software (local)
Timestamp format : NTP
Last advertisement:
 No advertisements have occurred

Next advertisement:
 Threshold check scheduled in 1 more probe (roughly every 30 seconds)
 No probes completed

Current computation:
 Started at: Apr 01 2024 17:31:41.452 (11.141 seconds ago)
 Packets Sent: 5, received: 5
 Measured delays (uSec): avg: 3538, min: 2836, max: 4379, variance: 702
 Next probe scheduled at: Apr 01 2024 17:32:08.459 (in 15.866 seconds)
 Next packet will be sent in 0.866 seconds
 Packet sent every 3.0 seconds
 Responder IP : 192.168.0.4
 Number of Hops : 2
 Probe samples:
 Packet Rx Timestamp Measured Delay (nsec)
 Apr 01 2024 17:31:41.457 3507559
 Apr 01 2024 17:31:41.480 2836021
 Apr 01 2024 17:31:44.481 3072131
 Apr 01 2024 17:31:47.483 3896705
 Apr 01 2024 17:31:50.484 4379536

```

```

.
.
.

```

- This example shows how to verify the timestamp format for the SR policy:

```

Router# show performance-measurement sessions

Transport type : SR Policy
Measurement type : Delay Measurement
Policy name : srte_c_10_ep_192.168.0.4
Color : 10
Endpoint : 192.168.0.4
Instance : 2
preference : 100
Protocol-origin : Configured
Discriminator : 100
Segment-list : R4
Atomic path:
 Hops : 192.168.0.4
 Session ID : 4145
 Trace ID : 2101013379
 NPU Offloaded session : False
 Timestamping Enabled : True
Timestamp source : Software (local)
Timestamp format : NTP

```

```

Last advertisement:
 No advertisements have occurred
Next advertisement:
 No probes completed
Last computation:
 None
Current computation:
 Packets Sent: 2, received: 2
 Measured delays (uSec): avg: 3531, min: 3424, max: 3639, variance: 107
Probe samples:
 Packet Rx Timestamp Measured Delay (nsec)
 Apr 01 2024 16:47:26.876 3424443
 Apr 01 2024 16:47:27.927 3639431
Liveness Detection: Disabled
Responder IP : 192.168.0.4
Number of Hops : 2

```

- This example shows how to verify the timestamp format for interfaces:

```

Router# show performance-measurement profile default interface delay
Interface Delay Measurement (default)
Profile configuration:
 Measurement mode : Two-way
 Timestamp format : NTP
 Protocol type: TWAMP-light
 Encap mode : UDP
 Type of service:
 PM-MPLS traffic class : 6
 TWAMP-light DSCP : 48
 Probe computation interval : 30 (effective: 30) seconds
 Burst interval : 3000 (effective: 3000) mSec
 TX interval : 3000000 (effective: 3000000) uSec
 Packets per computation interval : 10
 Periodic advertisement : Enabled
 Interval : 120 (effective: 120) sec
 Threshold : 10%
 Minimum-change : 500 uSec
 Advertisement accelerated : Disabled
 Anomaly delay advertisement: Disabled
 Anomaly loss advertisement: Disabled
 Advertisement logging:
 Delay exceeded : Disabled (default)
 Threshold crossing check : Minimum-delay
 Router alert : Disabled (default)
 Destination sweeping mode : Disabled
 Flow Label Count: 0

```

- This example shows how to verify the timestamp format for RSVP TE:

```

Router#show performance-measurement profile default rsvp-te delay
RSVP-TE Delay Measurement (default)
Profile configuration:
 Measurement mode : One-way
 Timestamp format : NTP
 Protocol type : TWAMP-light
 Encap mode : UDP
 Type of service:
 PM-MPLS traffic class : 6
 TWAMP-light DSCP : 48
 Probe computation interval : 30 (effective: 30) seconds
 Burst interval : 3000 (effective: 3000) mSec
 TX interval : 3000000 (effective: 3000000) uSec
 Packets per computation interval : 10
 Periodic advertisement : Enabled

```

```
Interval : 120 (effective: 120) sec
Threshold : 10%
Minimum-change : 500 uSec
Advertisement accelerated : Disabled
Anomaly delay advertisement : Disabled
Anomaly loss advertisement : Disabled
Advertisement logging:
 Delay exceeded : Disabled (default)
Threshold crossing check : Maximum-delay
Router alert : Disabled (default)
Destination sweeping mode : Disabled
Flow Label Count : 0
```

---



## CHAPTER 12

# Configure Topology-Independent Loop-Free Alternate (TI-LFA)

Table 77: Feature History Table

| Feature Name                                         | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| mLDP LSP protection with optimal TI-LFA backup paths | Release 24.4.1      | <p>Introduced in this release on: NCS 5500 fixed port routers; NCS 5700 fixed port routers; NCS 5500 modular routers (NCS 5500 line cards; NCS 5700 line cards [Mode: Compatibility; Native])</p> <p>You can now enhance the protection and convergence for Multicast Label Distribution Protocol Label Switched Paths (mLDP LSP) using Topology Independent Loop-Free Alternate (TI-LFA) backup paths. TI-LFA calculates an optimal, loop-free post-convergence path, and you can encode it as a list of segments.</p> <p>With this change, the TI-LFA supports zero-segment and single-segment backup paths, enabling native Point-to-Multipoint (P2MP) over SR code with a Tree Segment Identifier (Tree-SID) profile.</p> |

Topology-Independent Loop-Free Alternate (TI-LFA) uses segment routing to provide link, node, and Shared Risk Link Groups (SRLG) protection in topologies where other fast reroute techniques cannot provide protection.

- Classic Loop-Free Alternate (LFA) is topology dependent, and therefore cannot protect all destinations in all networks. A limitation of LFA is that, even if one or more LFAs exist, the optimal LFA may not always be provided.



- Remote LFA (RLFA) extends the coverage to 90-95% of the destinations, but it also does not always provide the most desired repair path. RLFA also adds more operational complexity by requiring a targeted LDP session to the RLFA to protect LDP traffic.

TI-LFA provides a solution to these limitations while maintaining the simplicity of the IPFRR solution.


The goal of TI-LFA is to reduce the packet loss that results while routers converge after a topology change due to a link or node failure. Rapid failure repair (< 50 msec) is achieved through the use of pre-calculated backup paths that are loop-free and safe to use until the distributed network convergence process is completed.

The optimal repair path is the path that the traffic will eventually follow after the IGP has converged. This is called the post-convergence path. This path is preferred for the following reasons:

- Optimal for capacity planning — During the capacity-planning phase of the network, the capacity of a link is provisioned while taking into consideration that such link will be used when other links fail.
- Simple to operate — There is no need to perform a case-by-case adjustments to select the best LFA among multiple candidate LFAs.
- Fewer traffic transitions — Since the repair path is equal to the post-convergence path, the traffic switches paths only once.

The following topology illustrates the optimal and automatic selection of the TI-LFA repair path.

**Figure 40: TI-LFA Repair Path**



Node 2 protects traffic to destination Node 5.

With classic LFA, traffic would be steered to Node 4 after a failure of the protected link. This path is not optimal, since traffic is routed over edge node Node 4 that is connected to lower capacity links.

TI-LFA calculates a post-convergence path and derives the segment list required to steer packets along the post-convergence path without looping back.

In this example, if the protected link fails, the shortest path from Node2 to Node5 would be:

Node2 → Node6 → Node7 → Node3 → Node5

Node7 is the PQ-node for destination Node5. TI-LFA encodes a single segment (prefix SID of Node7) in the header of the packets on the repair path.

### TI-LFA Protection Types


TI-LFA supports the following protection:

- Link protection — The link is excluded during the post-convergence backup path calculation.
- Node protection — The neighbor node is excluded during the post convergence backup path calculation.
- Shared Risk Link Groups (SRLG) protection — SRLG refer to situations in which links in a network share a common fiber (or a common physical attribute). These links have a shared risk: when one link fails, other links in the group might also fail. TI-LFA SRLG protection attempts to find the post-convergence backup path that excludes the SRLG of the protected link. All local links that share any SRLG with the protecting link are excluded.

When you enable link protection, you can also enable node protection, SRLG protection, or both, and specify a tiebreaker priority in case there are multiple LFAs.

The following example illustrates the link, node, and SRLG protection types. In this topology, Node2 applies different protection models to protect traffic to Node7.

**Figure 41: TI-LFA Protection Types**



- [Limitations, on page 530](#)
- [Usage guidelines and limitations for TI-LFA, on page 530](#)

- [Configuring TI-LFA for IS-IS, on page 531](#)
- [Configuring TI-LFA for OSPF, on page 533](#)
- [TI-LFA Node and SRLG Protection: Examples, on page 534](#)
- [Configuring Global Weighted SRLG Protection, on page 535](#)
- [SR-MPLS over GRE as TI-LFA Backup Path, on page 538](#)
- [Unlabeled IPv6 Traffic Protection, on page 549](#)

## Limitations

Only two backup labels are supported.

## Usage guidelines and limitations for TI-LFA

### Guidelines

- IGP directly programs a TI-LFA backup path requiring 3 or fewer labels, including the label of the protected destination prefix.

### Limitations

- The platform does not support programming of TI-LFA backup paths requiring more than 3 labels.

| TI-LFA Functionality                           | IS-IS <sup>1</sup>  | OSPFv2      |
|------------------------------------------------|---------------------|-------------|
| <b><i>Protected Traffic Types</i></b>          |                     |             |
| Protection for SR labeled traffic              | Supported           | Supported   |
| Protection of IPv4 unlabeled traffic           | Supported (IS-ISv4) | Supported   |
| Protection of IPv6 unlabeled traffic           | Supported (IS-ISv6) | N/A         |
| <b><i>Protection Types</i></b>                 |                     |             |
| Link Protection                                | Supported           | Supported   |
| Node Protection                                | Supported           | Supported   |
| Local SRLG Protection                          | Supported           | Supported   |
| Weighted Remote SRLG Protection                | Supported           | Supported   |
| Line Card Disjoint Protection                  | Supported           | Unsupported |
| <b><i>Interface Types</i></b>                  |                     |             |
| Ethernet Interfaces                            | Supported           | Supported   |
| Ethernet Bundle Interfaces                     | Supported           | Supported   |
| TI-LFA over GRE Tunnel as Protecting Interface | Supported           | Supported   |

| TI-LFA Functionality                                                                                         | IS-IS <sup>1</sup> | OSPFv2    |
|--------------------------------------------------------------------------------------------------------------|--------------------|-----------|
| <b>Additional Functionality</b>                                                                              |                    |           |
| Maximum number of labels that can be pushed on the backup path (including the label of the protected prefix) | 3                  | 3         |
| BFD-triggered                                                                                                | Supported          | Supported |
| BFDv6-triggered                                                                                              | Supported          | N/A       |
| Prefer backup path with lowest total metric                                                                  | Supported          | Supported |
| Prefer backup path from ECMP set                                                                             | Supported          | Supported |
| Prefer backup path from non-ECMP set                                                                         | Supported          | Supported |
| Load share prefixes across multiple backups paths                                                            | Supported          | Supported |
| Limit backup computation up to the prefix priority                                                           | Supported          | Supported |

<sup>1</sup> Unless specified, IS-IS support is IS-ISv4 and IS-ISv6

## Configuring TI-LFA for IS-IS

This task describes how to enable per-prefix Topology Independent Loop-Free Alternate (TI-LFA) computation to converge traffic flows around link, node, and SRLG failures.

### Before you begin

Ensure that the following topology requirements are met:

- Routers are configured with IS-IS.
- Segment routing for IS-IS is configured. See [Enabling Segment Routing for IS-IS Protocol, on page 147](#).

### Procedure

|               | Command or Action                                                                                              | Purpose                                                                                                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# configure                                         | Enters mode.                                                                                                                                                                                                                                                                       |
| <b>Step 2</b> | <b>router isis instance-id</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>router isis</b> <b>1</b> | Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.<br><br><b>Note</b><br>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command. |

|               | Command or Action                                                                                                                                                                                                                                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 3</b> | <b>interface</b> <i>type interface-path-id</i><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-isis)# interface GigabitEthernet0/0/0/1</pre>                                                                                                  | Enters interface configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Step 4</b> | <b>address-family ipv4 [unicast]</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast</pre>                                                                                                              | Specifies the IPv4 address family, and enters router address family configuration mode.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 5</b> | <b>fast-reroute per-prefix</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix</pre>                                                                                                                     | Enables per-prefix fast reroute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Step 6</b> | <b>fast-reroute per-prefix ti-lfa</b><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix ti-lfa</pre>                                                                                                       | Enables per-prefix TI-LFA fast reroute link protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 7</b> | <b>fast-reroute per-prefix tiebreaker</b><br><b>{node-protecting   srlg-disjoint} index</b><br><i>priority</i><br><b>Example:</b><br><pre>RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix tie-breaker srlg-disjoint index 100</pre> | <p>Enables TI-LFA node or SRLG protection and specifies the tiebreaker priority. Valid <i>priority</i> values are from 1 to 255. The lower the <i>priority</i> value, the higher the priority of the rule. Link protection always has a lower priority than node or SRLG protection.</p> <p><b>Note</b><br/>The same attribute cannot be configured more than once on an interface.</p> <p><b>Note</b><br/>For IS-IS, TI-LFA node protection and SRLG protection can be configured on the interface or the instance.</p> |

TI-LFA has been successfully configured for segment routing.

# Configuring TI-LFA for OSPF

This task describes how to enable per-prefix Topology Independent Loop-Free Alternate (TI-LFA) computation to converge traffic flows around link, node, and SRLG failures.



**Note** TI-LFA can be configured on the instance, area, or interface. When configured on the instance or area, all interfaces in the instance or area inherit the configuration.

## Before you begin

Ensure that the following topology requirements are met:

- Routers are configured with OSPF.
- Segment routing for OSPF is configured. See [Enabling Segment Routing for OSPF Protocol, on page 179](#).

## Procedure

|               | Command or Action                                                                                                                              | Purpose                                                                                                     |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# <b>configure</b>                                                                  | Enters mode.                                                                                                |
| <b>Step 2</b> | <b>router ospf process-name</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>router ospf</b><br><b>1</b>                             | Enables OSPF routing for the specified routing process, and places the router in router configuration mode. |
| <b>Step 3</b> | <b>area area-id</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-ospf)# <b>area</b><br><b>1</b>                                           | Enters area configuration mode.                                                                             |
| <b>Step 4</b> | <b>interface type interface-path-id</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-ospf-ar)#<br><b>interface GigabitEthernet0/0/0/1</b> | Enters interface configuration mode.                                                                        |
| <b>Step 5</b> | <b>fast-reroute per-prefix</b><br><b>Example:</b>                                                                                              | Enables per-prefix fast reroute.                                                                            |

|               | Command or Action                                                                                                                                                                                                                                                      | Purpose                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | RP/0/RP0/CPU0:router(config-ospf-ar-if) #<br><b>fast-reroute per-prefix</b>                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Step 6</b> | <b>fast-reroute per-prefix ti-lfa</b><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-ospf-ar-if) #<br><b>fast-reroute per-prefix ti-lfa</b>                                                                                                                 | Enables per-prefix TI-LFA fast reroute link protection.                                                                                                                                                                                                                                                                                                                          |
| <b>Step 7</b> | <b>fast-reroute per-prefix tiebreaker</b><br><b>{node-protecting   srlg-disjoint} index</b><br><b>priority</b><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-ospf-ar-if) #<br><b>fast-reroute per-prefix tie-breaker</b><br><b>srlg-disjoint index 100</b> | Enables TI-LFA node or SRLG protection and specifies the tiebreaker priority. Valid <i>priority</i> values are from 1 to 255. The higher the <i>priority</i> value, the higher the priority of the rule. Link protection always has a lower priority than node or SRLG protection.<br><br><b>Note</b><br>The same attribute cannot be configured more than once on an interface. |

TI-LFA has been successfully configured for segment routing.

## TI-LFA Node and SRLG Protection: Examples

The following examples show the configuration of the tiebreaker priority for TI-LFA node and SRLG protection, and the behavior of post-convergence backup-path. These examples use OSPF, but the same configuration and behavior applies to IS-IS.

### Example: Enable link-protecting and node-protecting TI-LFA

```
router ospf 1
 area 1
 interface GigabitEthernet0/0/2/1
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa
 fast-reroute per-prefix tiebreaker node-protecting index 100
```

Both link-protecting and node-protecting TI-LFA backup paths will be computed. If the priority associated with the node-protecting tiebreaker is higher than any other tiebreakers, then node-protecting post-convergence backup paths will be selected, if it is available.

### Example: Enable link-protecting and SRLG-protecting TI-LFA

```
router ospf 1
 area 1
 interface GigabitEthernet0/0/2/1
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa
```

```
fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Both link-protecting and SRLG-protecting TI-LFA backup paths will be computed. If the priority associated with the SRLG-protecting tiebreaker is higher than any other tiebreakers, then SRLG-protecting post-convergence backup paths will be selected, if it is available.

#### Example: Enable link-protecting, node-protecting and SRLG-protecting TI-LFA

```
router ospf 1
 area 1
 interface GigabitEthernet0/0/2/1
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa
 fast-reroute per-prefix tiebreaker node-protecting index 200
 fast-reroute per-prefix tiebreaker srlg-disjoint index 100
```

Link-protecting, node-protecting, and SRLG-protecting TI-LFA backup paths will be computed. If the priority associated with the node-protecting tiebreaker is highest from all tiebreakers, then node-protecting post-convergence backup paths will be selected, if it is available. If the node-protecting backup path is not available, SRLG-protecting post-convergence backup path will be used, if it is available.

## Configuring Global Weighted SRLG Protection

A shared risk link group (SRLG) is a set of links sharing a common resource and thus shares the same risk of failure. The existing loop-free alternate (LFA) implementations in interior gateway protocols (IGPs) support SRLG protection. However, the existing implementation considers only the directly connected links while computing the backup path. Hence, SRLG protection may fail if a link that is not directly connected but shares the same SRLG is included while computing the backup path. Global weighted SRLG protection feature provides better path selection for the SRLG by associating a weight with the SRLG value and using the weights of the SRLG values while computing the backup path.

To support global weighted SRLG protection, you need information about SRLGs on all links in the area topology. For IS-IS, you can flood SRLGs for remote links or manually configuring SRLGs on remote links. For OSPF, the SRLG values are advertised in the Extended Link LSA when SRLG is configured for that interface and Segment Routing is enabled.

The administrative weight (cost) of the SRLG can be configured using the **admin-weight** command. This command can be applied for all SRLG (global), or for a specific (named) SRLG. The default (global) admin-weight value is 1 for IS-IS, and 1000 for OSPF.

#### Configuration Examples: Global Weighted SRLG Protection for IS-IS

There are three types of configurations that are supported for the global weighted SRLG protection feature for IS-IS:

- Local SRLG with global weighted SRLG protection
- Remote SRLG flooding
- Remote SRLG static provisioning

This example shows how to configure the local SRLG with global weighted SRLG protection feature.



```

RP/0/RP0/CPU0:router(config)# srlg
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/1
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg)# exit
RP/0/RP0/CPU0:router(config-srlg)# name group1 value 100
RP/0/RP0/CPU0:router(config-srlg)# exit
RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix srlg-protection weighted-global
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix tiebreaker srlg-disjoint index
1
RP/0/RP0/CPU0:router(config-isis-af)# exit
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-isis-if)# point-to-point
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
RP/0/RP0/CPU0:router(config-isis-if-af)# exit
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# srlg
RP/0/RP0/CPU0:router(config-isis-srlg)# name group1
RP/0/RP0/CPU0:router(config-isis-srlg-name)# admin-weight 5000

```

This example shows how to configure the global weighted SRLG protection feature with remote SRLG flooding. The configuration includes local and remote router configuration. On the local router, the global weighted SRLG protection is enabled by using the **fast-reroute per-prefix srlg-protection weighted-global** command. In the remote router configuration, you can control the SRLG value flooding by using the **advertise application lfa link-attributes srlg** command. You should also globally configure SRLG on the remote router.

The local router configuration for global weighted SRLG protection with remote SRLG flooding is as follows:

```

RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix srlg-protection weighted-global
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix tiebreaker srlg-disjoint index
1
RP/0/RP0/CPU0:router(config-isis-af)# exit
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-isis-if)# point-to-point
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
RP/0/RP0/CPU0:router(config-isis-if-af)# exit
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# srlg
RP/0/RP0/CPU0:router(config-isis-srlg)# name group1
RP/0/RP0/CPU0:router(config-isis-srlg-name)# admin-weight 5000

```

The remote router configuration for global weighted SRLG protection with remote SRLG flooding is as follows:

```

RP/0/RP0/CPU0:router(config)# srlg
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/1

```

```
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# name group1 value 100
RP/0/RP0/CPU0:router(config-srlg)# exit
RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# advertise application lfa link-attributes srlg
```

This example shows configuring the global weighted SRLG protection feature with static provisioning of SRLG values for remote links. You should perform these configurations on the local router.

```
RP/0/RP0/CPU0:router(config)# srlg
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/1
RP/0/RP0/CPU0:router(config-srlg-if)# name group1
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# name group1 value 100
RP/0/RP0/CPU0:router(config-srlg)# exit
RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix srlg-protection weighted-global
RP/0/RP0/CPU0:router(config-isis-af)# fast-reroute per-prefix tiebreaker srlg-disjoint index
1
RP/0/RP0/CPU0:router(config-isis-af)# exit
RP/0/RP0/CPU0:router(config-isis)# interface TenGigE0/0/0/0
RP/0/RP0/CPU0:router(config-isis-if)# point-to-point
RP/0/RP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix
RP/0/RP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix ti-lfa
RP/0/RP0/CPU0:router(config-isis-if-af)# exit
RP/0/RP0/CPU0:router(config-isis-if)# exit
RP/0/RP0/CPU0:router(config-isis)# srlg
RP/0/RP0/CPU0:router(config-isis-srlg)# name group1
RP/0/RP0/CPU0:router(config-isis-srlg-name)# admin-weight 5000
RP/0/RP0/CPU0:router(config-isis-srlg-name)# static ipv4 address 10.0.4.1 next-hop ipv4
address 10.0.4.2
RP/0/RP0/CPU0:router(config-isis-srlg-name)# static ipv4 address 10.0.4.2 next-hop ipv4
address 10.0.4.1
```

### Configuration Examples: Global Weighted SRLG Protection for OSPF

There are two types of configurations that are supported for the global weighted SRLG protection feature for OSPF:

- Local SRLG with global weighted SRLG protection
- Remote SRLG static provisioning



**Note** There is no specific configuration to enable SRLG flooding in OSPF. The SRLG values are advertised in the Extended Link LSA if SRLG is configured for that interface and Segment Routing is enabled.

This example shows how to configure the local SRLG with global weighted SRLG protection feature.

```

Router(config)# srlg
Router(config-srlg)# interface hundredGigE 0/0/0/0
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# interface hundredGigE 0/0/0/1
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# name group1 value 100
Router(config-srlg)# exit
Router(config)# router ospf 1
Router(config-ospf)# fast-reroute per-prefix srlg-protection weighted-global
Router(config-ospf)# fast-reroute per-prefix tiebreaker srlg-disjoint index 1
Router(config-ospf-ar)# interface hundredGigE 0/0/0/0
Router(config-ospf-ar-if)# fast-reroute per-prefix
Router(config-ospf-ar-if)# fast-reroute per-prefix ti-lfa
Router(config-ospf-ar-if)# exit
Router(config-ospf-ar)# exit
Router(config-ospf)# srlg
Router(config-ospf-srlg)# admin-weight 3000
Router(config-ospf-srlg)# name group1
Router(config-ospf-srlg-name)# admin-weight 5000

```

This example shows configuring the global weighted SRLG protection feature with static provisioning of SRLG values for remote links. You should perform these configurations on the remote router.

```

Router(config)# srlg
Router(config-srlg)# interface hundredGigE 0/0/0/0
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# interface hundredGigE 0/0/0/1
Router(config-srlg-if)# name group1
Router(config-srlg-if)# exit
Router(config-srlg)# name group1 value 100
Router(config-srlg)# exit
Router(config)# router ospf 1
Router(config-ospf)# area 1
Router(config-ospf-ar)# interface hundredGigE 0/0/0/0
Router(config-ospf-ar-if)# exit
Router(config-ospf-ar)# exit

```

## SR-MPLS over GRE as TI-LFA Backup Path

This feature allows the router (as ABR) to program a Generic Routing Encapsulation (GRE) tunnel as an outgoing interface for TI-LFA backup paths computed by the IGP in a Segment Routing network. Single-segment TI-LFA scenario is supported. In this scenario, the router pushes one extra label when programming the backup path.




**Note** GRE is a tunneling protocol that provides a simple generic approach to transport packets of one protocol over another protocol by means of encapsulation. See the [Configuring GRE Tunnels](#) chapter in the *Interface and Hardware Component Configuration Guide*.

### Multi-Level Network Topology

The following example shows a multi-level network topology with interconnecting links between ABRs.




**Note** This could also be a multi-instance network topology.



Two links between ABR 2 and ABR 3 are required, one in each IS-IS level. These links provide protection in each direction and ensure that there is always an alternate path inside the IGP domain.




**Note** Alternatively, a single link with two logical sub-interfaces could be used between the ABRs.




TI-LFA performs the backup path calculation inside the domain (process, level, or area) of the failed link.

For example, if the link between nodes 2 and 5 failed, the link between ABR 2 and 3 would create a TI-LFA path in L1 IS-IS level. If the link between nodes 1 and 2 failed, the link between ABR 2 and 3 would create a TI-LFA path in L2 IS-IS level.


However, if the interconnecting link between ABRs are in the same Shared Risk Link Groups (SRLG) as other links inside the domain (for example, the link between Nodes 2 and 3 are in the same SRLG as link between Nodes 2 and 5), TI-LFA with local SRLG protection would not find an alternate path.




In cases where it is not feasible to provide interconnecting links between ABRs (for example, the ABR nodes might be in different locations with no connectivity options), TI-LFA will not be able to compute backup paths for all of the prefixes.



To address these issues, you can create a GRE tunnel in each domain, between the ABRs, which can be used as TI-LFA backup paths.




Now, if a link failure occurs in either IS-IS level (for example, between nodes 1 and 2 or between nodes 2 and 5), the path is protected by the GRE tunnel.




### Backup Path for Link Failure Between Nodes 2 and 5

Traffic from node 1 is rerouted over the GRE tunnel TI-LFA backup path between ABR nodes 2 and 3.



Traffic flowing in the opposite direction, from node 5 to node 1, is simply routed over nodes 6-3-4 to node 1.



## Limitations


The following behaviors and limitations apply to the router when a GRE tunnel is programmed as backup interface for TI-LFA:

- The MPLS label of a protected prefix must be the same in the primary and backup paths (SWAP scenario)

- Single-segment TI-LFA is supported. In this scenario, the router pushes one extra label when programming the backup path. The total label stack is 2, including the primary label and backup label.
- Double-segment (or more) TI-LFA is not supported. In this scenario, the router pushes two or more extra labels when programming the backup path.
- GRE tunnel as a primary or backup path for an SR policy with TI-LFA protection is not supported.

## Example: SR-MPLS over GRE as TI-LFA Backup Path

The examples in this section use the following network topology:



### Configurations Without Interconnecting ABR Links

The following sample configurations show OSPF configurations for nodes 2, 3 and 5. Nodes 2 and 3 are ABRs between Area 0 and Area 100. There is no connection between the ABRs.

#### Configuration on ABR 2 for Area 0 and Area 100

```
router ospf 100
 router-id 2.2.2.2
 segment-routing mpls
 segment-routing forwarding mpls
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa enable
 segment-routing sr-prefer
 area 0
 interface Loopback0
 prefix-sid index 2
 !
 !
 interface TenGigE0/0/1/10
 network point-to-point
 !
 !
 area 100
 interface TenGigE0/0/1/11
 network point-to-point
```

```
RP/0/RSP0/CPU0:ABR2# show ospf neighbor area-sorted
Fri Jul 19 09:43:59.328 UTC
```

```
Neighbors for OSPF 100
Area 0
```

| Neighbor ID | Pri | State | Dead Time  | Address  | Up Time | Interface  |
|-------------|-----|-------|------------|----------|---------|------------|
| 10.1.1.1    | 1   | FULL/ | - 00:00:35 | 10.1.2.1 | 1d20h   | Te0/0/1/10 |

```
Total neighbor count: 1
```

```
Area 100
```

| Neighbor ID | Pri | State | Dead Time  | Address  | Up Time | Interface  |
|-------------|-----|-------|------------|----------|---------|------------|
| 5.5.5.5     | 1   | FULL/ | - 00:00:33 | 10.2.5.5 | 1d20h   | Te0/0/1/11 |

```
Total neighbor count: 1
```

### Configuration on ABR 3 for Area 0 and Area 100

```
router ospf 100
router-id 3.3.3.3
segment-routing mpls
segment-routing forwarding mpls
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa enable
segment-routing sr-prefer
area 0
 interface Loopback0
 prefix-sid index 3
 !
 interface TenGigE0/0/0/9
 network point-to-point
 !
 !
area 100
 interface TenGigE0/0/0/3
 network point-to-point
 !
```

```
RP/0/RSP0/CPU0:ABR3# show ospf neighbor area-sorted
Fri Jul 19 09:33:35.816 UTC
```

```
Neighbors for OSPF 100
Area 0
```

| Neighbor ID | Pri | State | Dead Time  | Address  | Up Time | Interface |
|-------------|-----|-------|------------|----------|---------|-----------|
| 4.4.4.4     | 1   | FULL/ | - 00:00:36 | 10.3.4.4 | 2d17h   | Te0/0/0/9 |

```
Total neighbor count: 1
```

```
Area 100
```

| Neighbor ID | Pri | State | Dead Time  | Address  | Up Time | Interface |
|-------------|-----|-------|------------|----------|---------|-----------|
| 6.6.6.6     | 1   | FULL/ | - 00:00:36 | 10.3.6.6 | 2d19h   | Te0/0/0/3 |

```
Total neighbor count: 1
```

### Configuration on Node 5

```
segment-routing mpls
!
set-attributes
address-family ipv4
sr-label-preferred
```



## Example: SR-MPLS over GRE as TI-LFA Backup Path

```

!
connected-prefix-sid-map
 address-family ipv4
 5.5.5.5/32 index 5 range 1
!
interface TenGigabitEthernet0/0/26
 description ***Connected to ABR 2
 ip address 10.2.5.5 255.255.255.0
 ip ospf network point-to-point
 cdp enable
!
interface TenGigabitEthernet0/0/27
 description ***Connected to Node 6
 ip address 10.5.6.5 255.255.255.0
 ip ospf network point-to-point
 cdp enable

router ospf 100
 router-id 5.5.5.5
 segment-routing area 100 mpls
 segment-routing mpls
 fast-reroute per-prefix enable prefix-priority low
 fast-reroute per-prefix ti-lfa
 fast-reroute per-prefix ti-lfa area 100
 passive-interface default
 no passive-interface TenGigabitEthernet0/0/26
 no passive-interface TenGigabitEthernet0/0/27
 network 10.2.5.0 0.0.0.255 area 100
 network 10.5.5.0 0.0.0.255 area 100
 network 10.5.6.0 0.0.0.255 area 100
 network 5.5.5.5 0.0.0.0 area 100

```

```

RP/0/RSP0/CPU0:Node5# show ip ospf neighbor
Load for five secs: 4%/1%; one minute: 4%; five minutes: 4%
Time source is NTP, 09:50:51.417 UTC Fri Jul 19 2019

```

| Neighbor ID | Pri | State   | Dead Time | Address  | Interface                |
|-------------|-----|---------|-----------|----------|--------------------------|
| 6.6.6.6     | 0   | FULL/ - | 00:00:32  | 10.5.6.6 | TenGigabitEthernet0/0/27 |
| 2.2.2.2     | 0   | FULL/ - | 00:00:36  | 10.5.2.5 | TenGigabitEthernet0/0/26 |

## TI-LFA Fast Reroute Coverage on Node 5

The following output shows that this configuration provides only 52% TI-LFA fast reroute coverage on Node 5:

```

RP/0/RSP0/CPU0:Node5# show ip ospf fast-reroute prefix-summary
Load for five secs: 4%/1%; one minute: 4%; five minutes: 4%
Time source is NTP, 10:32:20.236 UTC Fri Jul 19 2019
 OSPF Router with ID (5.5.5.5) (Process ID 100)
 Base Topology (MTID 0)

Area 100:
Interface Protected Primary paths Protected paths Percent protected
 All High Low All High Low All High Low
Lo0 Yes 0 0 0 0 0 0 0% 0% 0%
Te0/0/27 Yes 7 4 3 1 1 0 14% 25% 0%
Te0/0/26 Yes 10 5 5 8 4 4 80% 80% 80%

Area total: 17 9 8 9 5 4 52% 55% 50%

Process total: 17 9 8 9 5 4 52% 55% 50%


```

### GRE Tunnel Configuration

The following examples show how to configure GRE tunnels between the ABRs in each area to provide TI-LFA backup paths for the Segment Routing network.


**GRE BLU** is configured in Area 0 using Loopback50 (on ABR2) and Loopback 60 (on ABR3). These loopbacks are advertised in Area 100:

Figure 42: GRE BLU



**GRE RED** is configured in Area 100 using Loopback55 (on ABR2) and Loopback 66 (on ABR3). These loopbacks are advertised in Area 0:

Figure 43: GRE RED



### Configuration on ABR 2

```
interface Loopback0
 ipv4 address 2.2.2.2 255.255.255.255
!
interface Loopback50
 description Lo for GRE BLU
 ipv4 address 50.0.0.50 255.255.255.0
!
interface Loopback55
```

```

description Lo for GRE RED
ipv4 address 55.55.55.55 255.255.255.255
!
interface tunnel-ip5060
description GRE virtual link for Area 0 BLU
ipv4 address 66.3.2.2 255.255.255.0
tunnel source Loopback50
tunnel destination 60.0.0.60
!
interface tunnel-ip5566
description GRE virtual link for Area 100 RED
ipv4 address 100.3.2.2 255.255.255.0
tunnel source Loopback55
tunnel destination 66.66.66.66

router ospf 100
router-id 2.2.2.2
segment-routing mpls
segment-routing forwarding mpls
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa enable
segment-routing sr-prefer
area 0
interface Loopback0
prefix-sid index 2
!
interface Loopback55
passive enable
!
interface tunnel-ip5060
cost 1000
!
interface TenGigE0/0/1/10
network point-to-point
!
!
area 100
interface Loopback50
passive enable
!
interface tunnel-ip5566
cost 1000
!
interface TenGigE0/0/1/11
network point-to-point

```



**Note** In the above configuration, GRE tunnel-ip5060 belongs to area 0, but its source and destination addresses are advertised in area 100. This ensures disjointness between the GRE tunnel and the links in area 0 that it protects. The same applies to GRE tunnel-ip5566 which belongs to area 100 and its source and destination addresses are advertised in area 0.

A high cost is applied to the GRE tunnel interfaces so that they are used only as a backup path.

### Configuration on ABR 3

```

interface Loopback0
ipv4 address 3.3.3.3 255.255.255.255
!
interface Loopback60
description Lo for GRE BLU

```

```

 ipv4 address 60.0.0.60 255.255.255.0
 !
interface Loopback66
 description Lo for GRE RED
 ipv4 address 66.66.66.66 255.255.255.255
 !
interface tunnel-ip5060
 description GRE virtual link for Area 0 BLU
 ipv4 address 66.3.2.3 255.255.255.0
 tunnel source Loopback60
 tunnel destination 50.0.0.50
 !
interface tunnel-ip5566
 description GRE virtual link for Area 100 RED
 ipv4 address 100.3.2.3 255.255.255.0
 tunnel source Loopback66
 tunnel destination 55.55.55.55

router ospf 100
 router-id 3.3.3.3
 segment-routing mpls
 segment-routing forwarding mpls
 fast-reroute per-prefix
 fast-reroute per-prefix ti-lfa enable
 segment-routing sr-prefer
 area 0
 interface Loopback0
 prefix-sid index 3
 !
 interface TenGigE0/0/0/9
 network point-to-point
 !
 interface Loopback66
 passive enable
 !
 interface tunnel-ip5060
 cost 1000
 !
 area 100
 interface TenGigE0/0/0/3
 network point-to-point
 !
 interface Loopback60
 passive enable
 !
 interface tunnel-ip5566
 cost 1000

```



**Note** In the above configuration, GRE tunnel-ip5060 belongs to area 0, but its source and destination addresses are advertised in area 100. This ensures disjointness between the GRE tunnel and the links in area 0 that it protects. The same applies to GRE tunnel-ip5566 which belongs to area 100 and its source and destination addresses are advertised in area 0.

A high cost is applied to the GRE tunnel interfaces so that they are used only as a backup path.

### TI-LFA Fast Reroute Coverage on Node 5 After GRE Tunnel Configuration

The following output shows that this configuration provides 100% TI-LFA fast reroute coverage on Node 5:

## Example: SR-MPLS over GRE as TI-LFA Backup Path

```

RP/0/RSP0/CPU0:Node5# show ip ospf fast-reroute prefix-summary
Load for five secs: 5%/1%; one minute: 4%; five minutes: 4%
Time source is NTP, 11:20:31.743 UTC Fri Jul 19 2019
 OSPF Router with ID (5.5.5.5) (Process ID 100)
 Base Topology (MTID 0)

Area 100:
Interface Protected Primary paths Protected paths Percent protected
 Yes All High Low All High Low All High Low
Lo0 Yes 0 0 0 0 0 0 0% 0% 0%
Te0/0/27 Yes 9 6 3 9 6 3 100% 100% 100%
Te0/0/26 Yes 11 6 5 11 6 5 100% 100% 100%


Area total: 20 12 8 20 12 8 100% 100% 100%

Process total: 20 12 8 20 12 8 100% 100% 100%


```

## Traffic Flow with GRE Tunnel as TI-LFA Backup


With a link failure between Node 1 and ABR 2, traffic flowing from Node 1 to Node 5 is simply routed through Nodes 4-3-6 to Node 5.




With GRE tunnel as TI-LFA backup, traffic flowing from Node 5 to Node 1 will be encapsulated at ABR2 and routing over the GRE tunnel.



With a link failure between Node 5 and ABR 2, traffic flowing from Node 5 to Node 1 is simply routed through Nodes 6-3-4 to Node 1.



With GRE tunnel as TI-LFA backup, traffic flowing from Node 1 to Node 5 will be encapsulated at ABR2 and routing over the GRE tunnel.



## Unlabeled IPv6 Traffic Protection

Table 78: Feature History Table

| Feature Name                                  | Release Information | Feature Description                                                                                                                                                                                                                  |
|-----------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv6 Unlabeled Traffic protection with TI-LFA | Release 7.3.1       | <p>TI-LFA provides protection for SR-labeled traffic (IPv4 and IPv6 prefixes associated with a prefix SID) and for other unlabeled IPv4 prefixes.</p> <p>This feature introduces support for protecting unlabeled IPv6 prefixes.</p> |

This feature introduces support for protecting unlabeled IPv6 prefixes. IS-IS can calculate and install TI-LFA backup paths for unlabeled IPv6 prefixes.

By default, all IPv6 prefixes without a prefix SID are eligible to have a dynamic local label assigned to them. This behavior provides the greatest degree of protection, but in some deployments, it may consume too many MPLS labels. This feature provides the ability to disable local label allocation or to restrict local label allocation to a subset of prefixes based on a prefix list or route policy.

Use the following commands in IS-IS IPv6 address family configuration mode to specify the local label allocation behavior:

- **segment-routing mpls unlabeled protection disable**—Disable local label allocation.
- **segment-routing mpls unlabeled protection prefix-list *sample\_prefix\_list***—Restricts local label allocation to the prefixes based on a prefix list.
- **segment-routing mpls unlabeled protection route-policy *sample\_rpl***—Restricts local label allocation to the prefixes based a route policy.

## Configuration

The following example shows how to disable local label allocation:

```
Router(config)# router isis 1
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing mpls unlabeled protection disable
Router(config-isis-af)#
```

The following example shows how to enable local label allocation for prefixes in a prefix list:

```
Router(config)# ipv6 prefix-list sample_prefix_list
Router(config-ipv6-pfx)# 10 permit 333::333:0:0/96 ge 112
Router(config-ipv6-pfx)# 20 permit 666::666:0:0/96 ge 112
Router(config-ipv6-pfx)# exit
Router(config)#

Router(config)# router isis 1
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing mpls unlabeled protection prefix-list
sample_prefix_list
Router(config-isis-af)# commit
```

The following example shows how to enable local label allocation for prefixes in a route policy:

```
Router(config)# prefix-set sample_prefix_set
Router(config-pfx)# 333::333:1:0/112
Router(config-pfx)# end-set
Router(config)# route-policy sample_rpl
Router(config-rpl)# if destination in sample_prefix_set then
Router(config-rpl-if)# pass
Router(config-rpl-if)# else drop endif
Router(config-rpl)# end-policy
Router(config)#

Router(config)# router isis 1
Router(config-isis)# address-family ipv6 unicast
Router(config-isis-af)# segment-routing mpls unlabeled protection route-policy sample_rpl
Router(config-isis-af)# commit
```

## Verification

In the following **show** command output, 24103 is the local label used to program the prefix with a TI-LFA backup.

```

Router# show isis ipv6 unicast route 333::333:1:0/112 detail
L2 333::333:1:0/112 [120/115] Label: 24103, low priority
 via fe80::28a:96ff:fe4:5403, TenGigE0/1/0/3/5, r6, SRGB Base: 16000, Weight: 0
 src R1.00-01, 1:1:1::1

Router# show isis ipv6 fast-reroute 333::333:1:0/112 detail
L2 333::333:1:0/112 [120/115] Label: 24103, low priority
 via fe80::28a:96ff:fe4:5403, TenGigE0/1/0/3/5, R6, SRGB Base: 16000, Weight: 0
 Backup path: TI-LFA (link), via fe80::2c1:64ff:fe60:39b9, TenGigE0/1/0/3/0 R4, SRGB
 Base: 16000, Weight: 0, Metric: 140
 P node: R3.00 [3:3:3::3], Label: 16333
 Prefix label: None
 Backup-src: R1.00
 P: No, TM: 140, LC: No, NP: No, D: No, SRLG: Yes
 src R1.00-a9, 1:1:1::1

Router# show route ipv6 333::333:1:0/112 detail
Routing entry for 333::333:1:0/112
 Known via "isis 1", distance 115, metric 120, type level-2
 Installed Jul 26 19:11:58.840 for 00:00:47
Routing Descriptor Blocks
 fe80::2c1:64ff:fe60:39b9, from 1:1:1::1, via TenGigE0/1/0/3/0, Backup (TI-LFA)
 Repair Node(s): 3:3:3::3
 Route metric is 140
 Label: 0x3fcd (16333)
 Tunnel ID: None
 Binding Label: None
 Extended communities count: 0
 Path id:65 Path ref count:1
 NHID:0x20008(Ref:20026)
 fe80::28a:96ff:fe4:5403, from 1:1:1::1, via TenGigE0/1/0/3/5, Protected
 Route metric is 120
 Label: None
 Tunnel ID: None
 Binding Label: None
 Extended communities count: 0
 Path id:1 Path ref count:0
 NHID:0x20009(Ref:20021)
 Backup path id:65
Route version is 0x1e (30)
Local Label: 0x5e27 (24103)
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_NON_RECURSIVE_LOW (8) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 2, Download Version 4200135
No advertising protos.

```

```
Router# show mpls forwarding labels 24103
```

| Local Label | Outgoing Label | Prefix or ID     | Outgoing Interface | Next Hop                 | Bytes Switched |
|-------------|----------------|------------------|--------------------|--------------------------|----------------|
| 24103       | Unlabelled     | 333::333:1:0/112 | Te0/1/0/3/5        | fe80::28a:96ff:fe4:5403  | \              |
|             |                |                  |                    | 0                        |                |
|             | 16333          | 333::333:1:0/112 | Te0/1/0/3/0        | fe80::2c1:64ff:fe60:39b9 | \              |
|             |                |                  |                    | 0                        | (!)            |







## CHAPTER 13

# Configure Segment Routing Microloop Avoidance

The Segment Routing Microloop Avoidance feature enables link-state routing protocols, such as IS-IS and OSPF, to prevent or avoid microloops during network convergence after a topology change.

- [About Segment Routing Microloop Avoidance, on page 553](#)
- [Usage Guidelines and Limitations, on page 556](#)
- [Configure Segment Routing Microloop Avoidance for IS-IS, on page 556](#)
- [Configure Segment Routing Microloop Avoidance for OSPF, on page 561](#)

## About Segment Routing Microloop Avoidance

IP hop-by-hop routing may induce microloops (uLoops) at any topology transition. Microloops are a day-one IP challenge. Microloops are brief packet loops that occur in the network following a topology change:

- Link down or up (remote or local)
- Metric increase or decrease (remote or local)
- OSPFv2 only — Single-node cost-out: This occurs when a Router LSA (Link State Advertisement) is received with all non-stub links set to the maximum metric (link cost of 65535). It indicates that the node is being taken out of the routing path by setting the links to an unreachable state.
- OSPFv2 only — Single-node cost-in: This occurs when a Router LSA is received that brings the node back into the routing path by changing at least one non-stub link from the maximum metric mode.

Microloops are caused by the non-simultaneous convergence of different nodes in the network. If a node converges and sends traffic to a neighbor node that has not converged yet, traffic may be looped between these two nodes, resulting in packet loss, jitter, and out-of-order packets.

Segment Routing can be used to resolve the microloop problem. A router with the Segment Routing Microloop Avoidance feature detects if microloops are possible for a destination on the post-convergence path following a topology change associated with a remote link event.


If a node determines that a microloop could occur on the new topology, the IGP computes a microloop-avoidant path by updating the forwarding table and temporarily (based on a RIB update delay timer) installing the SID-list imposition entries associated with the microloop-avoidant path for the destination. Traffic is steered to that destination loop-free.

After the RIB update delay timer expires, IGP updates the forwarding table and removes the microloop-avoidant SID list. Traffic now natively follows the post-convergence path.


SR microloop avoidance is a local behavior and therefore not all nodes need to implement it to get the benefits.

In the topology below, microloops can occur after the failure of the link between Node6 and Node7, or if Node6 costs out.


At steady state, Node1 sends traffic to node 6 (16006) via Node7. Node 7 is configured with TI-LFA to protect traffic to Node6.



TI-LFA on Node7 pre-computes a backup path for traffic to Node6 (prefix SID 16006) that will be activated if the link between Node7 and Node6 goes down. In this network, the backup path would steer traffic toward Node5 (prefix SID 16005) and then via link between Node5 and Node6 (adj-SID 24056). All nodes are notified of the topology change due to the link failure.




However, if nodes along the path do not converge at the same time, microloops can be introduced. For example, if Node2 converged before Node3, Node3 would send traffic back to Node2 as the shortest IGP path to Node6. The traffic between Node2 and Node3 creates a microloop.




With microloop avoidance configured on Node1, a post-convergence path is computed and possible microloops on the post-convergence path for any destination are detected.

If microloops are possible on the post-convergence path to Node6, a microloop-avoidant path is constructed to steer the traffic to Node6 loop-free over the microloop-avoidant path {16005, 24056, 16006}.

Node1 updates the forwarding table and installs the SID-list imposition entries for those destinations with possible microloops, such as Node6. All nodes converge and update their forwarding tables, using SID lists where needed.



After the RIB update delay timer expires, the microloop-avoidant path is replaced with regular forwarding paths; traffic now natively follows the post-convergence path.



## Usage Guidelines and Limitations

### Configure Segment Routing Microloop Avoidance for IS-IS

This task describes how to enable Segment Routing Microloop Avoidance and set the Routing Information Base (RIB) update delay value for IS-IS.

#### Before you begin

Ensure that the following topology requirements are met:

- Routers are configured with IS-IS.
- Segment routing for IS-IS is configured. See [Enabling Segment Routing for IS-IS Protocol](#), on page 147.

#### Procedure

|               | Command or Action                                                                                                 | Purpose                                                                                                                                                                                                                                                             |
|---------------|-------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# configure                                            | Enters mode.                                                                                                                                                                                                                                                        |
| <b>Step 2</b> | <b>router isis instance-id</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>router isis</b><br><b>1</b> | Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.<br><br>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command. |
| <b>Step 3</b> | <b>address-family ipv4 [ unicast ]</b><br><b>Example:</b>                                                         | Specifies the IPv4 address family and enters router address family configuration mode.                                                                                                                                                                              |

|               | Command or Action                                                                                                                                                                        | Purpose                                                                                                                                                                                                         |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | RP/0/RP0/CPU0:router(config-isis)#<br><b>address-family ipv4 unicast</b>                                                                                                                 |                                                                                                                                                                                                                 |
| <b>Step 4</b> | <b>microloop avoidance segment-routing</b><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-isis-af)#<br><b>microloop avoidance segment-routing</b>                             | Enables Segment Routing Microloop Avoidance.                                                                                                                                                                    |
| <b>Step 5</b> | <b>microloop avoidance rib-update-delay</b><br><i>delay-time</i><br><br><b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-isis-af)#<br><b>microloop avoidance rib-update-delay 3000</b> | Specifies the amount of time the node uses the microloop avoidance policy before updating its forwarding table. The <i>delay-time</i> is in milliseconds. The range is from 1-60000. The default value is 5000. |

## Microloop Avoidance for IS-IS with Per-Prefix Filtering

Table 79: Feature History Table

| Feature Name                                            | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Microloop Avoidance for IS-IS with Per-Prefix Filtering | Release 7.11.1      | <p>Currently, when SR Microloop Avoidance for IS-IS is enabled, it applies to all prefixes.</p> <p>This feature allows you to selectively allow or deny specific IPv4 or IPv6 prefixes or routes that may cause microloops, which allows for efficient use of hardware resources and ensures overall network stability.</p> <p>This feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>microloop avoidance segment-routing</b> command is modified with the new <b>route-policy name</b> option for IS-IS.</li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>This feature extends the native <code>Cisco-IOS-XR-um-router-isis-cfg.yang</code> model (see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> |

Per-prefix filtering is an enhancement to the existing IOS XR IS-IS SR microloop avoidance feature. Per-prefix filtering allows network administrators to specify a subset of IP prefixes (v4 and v6) to which micro loop avoidance mechanisms can be applied.



**Note** Per-prefix filtering is available only for SR microloop avoidance and is not supported for local microloop avoidance.

When SR microloop avoidance is enabled, it applies to all prefixes. However, it might be important to preserve hardware resources for certain prefixes. In such a cases, it is beneficial to use SR microloop avoidance per-prefix filtering to allow only those prefixes without such limitations to be subjected to SR microloop avoidance. Per-prefix filtering provides a level of granularity that allows you to apply microloop avoidance only for the prefixes that require it, and to avoid consumption of resources that might otherwise be exhausted.

SR Microloop avoidance per-prefix filtering is configured under the IPv4 or IPv6 address family (AF). It will only be used for filtering in that specific AF. Filtering is applied to prefixes from all algorithms (Algo 0, Flexible Algorithms 128 to 255).

- For SR MLPS – If a prefix has multiple Flexible Algorithm paths and the filtering configuration permits SR microloop avoidance for that prefix, then SR microloop avoidance is allowed for "all" Flexible Algorithm paths associated with that prefix. On the other hand, if a prefix has multiple Flexible Algorithm paths and the filtering configuration prohibits SR microloop avoidance for that prefix, then microloop avoidance is disabled for all Flexible Algorithm paths associated with that prefix.
- For SRv6 – Regardless of the association of the prefix to the algorithm, filtering is applied solely on a per-prefix basis.

SR microloop avoidance per-prefix filtering uses route policies to identify the prefixes subjected to microloop avoidance. When SR microloop avoidance per-prefix filtering is enabled, the prefixes are verified against the route policy as follows:

- If the route policy permits the prefixes (pass), SR microloop avoidance computes the explicit path for the prefixes.
- If the route policy prevents the prefixes from being considered for SR microloop avoidance (drop), it is treated as if there is no explicit path defined for that prefix. The network will rely on the standard routing mechanisms to determine the path for those prefixes after convergence.

### Usage Guidelines and Limitations

- SR microloop avoidance per-prefix filtering is supported only for IS-IS.
- A route policy must be defined before it can be attached to the SR microloop avoidance configuration.
- Inline modification of a route policy is not supported. Once a route policy is defined and attached to the SR microloop avoidance configuration, it cannot be modified or removed until the route policy is removed from the SR microloop avoidance configuration.
- The following match types are supported for route policies used for SR microloop avoidance per-prefix filtering:
  - Destination-based match ([prefix](#) or [prefix set](#))
  - [Tag-based match](#)

### Example: Configuration

1. Identify the prefixes to be filtered by defining a prefix set or applying a tags to prefixes.

- Prefix Set

```
prefix-set pset-sample-ipv4
 2.3.3.3/32,
 2.4.4.4/32
 2.5.5.5/32
end-set

prefix-set pset-sample-ipv6
 2001:0:0:1::/64,
 2001:0:0:2::/64,
 2001:0:0:2::/64,
 2001:0:0:2::/64
end-set
```

- Tag

```
router isis 1
 interface Loopback1
 address-family ipv4 unicast
 tag 7
 prefix-sid index 7
```

2. Create a route policy for the destinations (prefix set) or tagged prefixes.

- Destination

```
route-policy BAR
 if destination in pset-sample-ipv4 then
 pass
 else
 drop
 endif
end-policy

route-policy BAR2
 if destination in (2.3.3.3/32, 2.4.4.4/32) then
 pass
 else
 drop
 endif
end-policy

route-policy BAR3
 if destination in pset-sample-ipv6 then
 drop
 else
 pass
 endif
end-policy
```

- Tag

```
route-policy FOO
 if tag eq 7 then
 drop
 endif
pass
```



```

end-policy

route-policy FOO2
 if tag eq 7 then
 pass
 else
 drop
 endif
end-policy

```

3. Use the **microloop avoidance segment-routing route-policy name** command to attach the route policy to the SR Microloop Avoidance configuration.

```

router isis 1
 address-family ipv4 unicast
 microloop avoidance segment-routing route-policy FOO2
 !
 address-family ipv6 unicast
 microloop avoidance segment-routing route-policy BAR3

```

## Verify

Use the **show isis** command to verify that SR microloop avoidance is enabled under the AF and the route policy is applied for per-prefix filtering.

Router# **show isis**

```

IS-IS Router: 1
 System Id: 0000.0000.0001
 IS Levels: level-2-only
 Manual area address(es):
 49.0001
 Routing for area address(es):
 49.0001
 Multi-Instance Id: 0
 Job Id: 1013
 PID: 61171
 Respawn count: 1
 Started: Thu Feb 23 02:57:58 2023
 LSP MTU: 1400
 LSP Full: level-1: No, level-2: No
 Non-stop forwarding: Cisco Proprietary NSF Restart enabled
 Most recent startup mode: Cold Restart
 TE connection status: Up
 XTC connection status: Up
 Overload Bit: not configured
 Maximum Metric: not configured
 Topologies supported by IS-IS:
 IPv4 Unicast
 Rib connected
 Level-2
 Metric style (generate/accept): Wide/Wide
 Metric: 10
 Microloop avoidance: Enabled
 Configuration: Type: Segment Routing, RIB update delay: 60000 msec, Policy: FOO2
 No protocols redistributed
 Distance: 115
 Advertise Passive Interface Prefixes Only: No
 IPv6 Unicast

```

```

Rib connected
Level-2
Metric: 10
Microloop avoidance: Enabled
 Configuration: Type: Segment Routing, RIB update delay: 5000 msec, Policy: BAR3
No protocols redistributed
Distance: 115
Advertise Passive Interface Prefixes Only: No
SR-MPLS:
 SRLB allocated: 15000 - 15999
 SRGB allocated: 16000 - 23999
SRv6:
 Configured locators:
 USID_ALG0 (Active)
 USID_ALG128 (Active)
Interfaces supported by IS-IS 1:
 Loopback0 is running actively (active in configuration)
 GigabitEthernet0/2/0/0 is running actively (active in configuration)
 GigabitEthernet0/2/0/3 is running actively (active in configuration)
 GigabitEthernet0/2/0/4 is running actively (active in configuration)
 GigabitEthernet0/2/0/6 is running actively (active in configuration)
 GigabitEthernet0/2/0/7 is running actively (active in configuration)

```

## Configure Segment Routing Microloop Avoidance for OSPF

**Table 80: Feature History Table**

| Feature Name                                                                             | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Microloop Avoidance for OSPFv2<br>Single-Node Cost-in and<br>Single-Node Cost-out Events | Release 7.11.1      | <p>Microloops disrupt network connectivity and cause suboptimal routing decisions. This feature avoids microloops by implementing the Greedy walk algorithm, which is similar to TI-LFA computation.</p> <p>This feature extends the microloop avoidance support for additional scenarios in OSPFv2, such as cost-in and cost-out events.</p> <p>This feature introduces these changes:</p> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>Cisco-IOS-XR-ipv4-ospf-oper.yang</li> </ul> <p>(see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</p> |

This task describes how to enable Segment Routing Microloop Avoidance and set the Routing Information Base (RIB) update delay value for OSPF.

**Before you begin**

Ensure that the following topology requirements are met:

- Routers are configured with OSPF.
- Segment routing for OSPF is configured. See [Enabling Segment Routing for OSPF Protocol](#), on page 179.

**Procedure**

|               | Command or Action                                                                                                                                                | Purpose                                                                                                                                                                                                       |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# <b>configure</b>                                                                                    | Enters mode.                                                                                                                                                                                                  |
| <b>Step 2</b> | <b>router ospf process-name</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>router ospf 1</b>                                                         | Enables OSPF routing for the specified routing process, and places the router in router configuration mode.                                                                                                   |
| <b>Step 3</b> | <b>microloop avoidance segment-routing</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-ospf)# <b>microloop avoidance segment-routing</b>                   | Enables Segment Routing Microloop Avoidance.                                                                                                                                                                  |
| <b>Step 4</b> | <b>microloop avoidance rib-update-delay delay-time</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-ospf)# <b>microloop avoidance rib-update-delay 3000</b> | Specifies the amount of time the node uses the microloop avoidance path before updating its forwarding table. The <i>delay-time</i> is in milliseconds. The range is from 1-60000. The default value is 5000. |



## CHAPTER 14

# Configure Segment Routing Mapping Server

The mapping server is a key component of the interworking between LDP and segment routing. It enables SR-capable nodes to interwork with LDP nodes. The mapping server advertises Prefix-to-SID mappings in IGP on behalf of other non-SR-capable nodes.

- [Segment Routing Mapping Server, on page 563](#)
- [Segment Routing and LDP Interoperability, on page 565](#)
- [Configuring Mapping Server, on page 568](#)
- [Enable Mapping Advertisement, on page 570](#)
- [Enable Mapping Client, on page 572](#)

## Segment Routing Mapping Server

The mapping server functionality in Cisco IOS XR segment routing centrally assigns prefix-SIDs for some or all of the known prefixes. A router must be able to act as a mapping server, a mapping client, or both.

- A router that acts as a mapping server allows the user to configure SID mapping entries to specify the prefix-SIDs for some or all prefixes. This creates the local SID-mapping policy. The local SID-mapping policy contains non-overlapping SID-mapping entries. The mapping server advertises the local SID-mapping policy to the mapping clients.
- A router that acts as a mapping client receives and parses remotely received SIDs from the mapping server to create remote SID-mapping entries.
- A router that acts as a mapping server and mapping client uses the remotely learnt and locally configured mapping entries to construct the non-overlapping consistent active mapping policy. IGP instance uses the active mapping policy to calculate the prefix-SIDs of some or all prefixes.

The mapping server automatically manages the insertions and deletions of mapping entries to always yield an active mapping policy that contains non-overlapping consistent SID-mapping entries.

- Locally configured mapping entries must not overlap each other.
- The mapping server takes the locally configured mapping policy, as well as remotely learned mapping entries from a particular IGP instance, as input, and selects a single mapping entry among overlapping mapping entries according to the preference rules for that IGP instance. The result is an active mapping policy that consists of non-overlapping consistent mapping entries.
- At steady state, all routers, at least in the same area or level, must have identical active mapping policies.

## Usage Guidelines and Restrictions

*Table 81: Feature History Table*

| Feature Name                                              | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advertisement of SID-Mapping Entries Between IS-IS Levels | Release 7.3.1       | <p>The Segment Routing Mapping Server (SRMS) is a key component of the interworking between LDP and segment routing, enabling SR-capable nodes to interwork with LDP nodes.</p> <p>This release introduces support for SRMS SID-mapping entries to be advertised between IS-IS levels (for example, from Level 1 to Level 2-only and from Level 2 to Level 1), where previously, the mappings were advertised only within the same IS-IS level, but not between IS-IS levels. This feature simplifies and centralizes the deployment of SRMS by removing the requirement of having a mapping server for each IS-IS area.</p> |

- The position of the mapping server in the network is not important. However, since the mapping advertisements are distributed in IGP using the regular IGP advertisement mechanism, the mapping server needs an IGP adjacency to the network.
- The role of the mapping server is crucial. For redundancy purposes, you should configure multiple mapping servers in the networks.
- The mapping server functionality supports the advertisement of SID-mapping entries between IS-IS levels (for example, from L1 to L2-only and from L2 to L1). A mapping server is not required for each IS-IS area.

For example, mapping entries learned from IS-IS Type Level-1 (intra-area) routers can be used to calculate prefix-SIDs for prefixes learned or advertised by IS-IS Type Level-2-only (backbone) routers.

Use the **domain-wide** option to advertise the prefix-SID mappings between Level 1 and Level 2 IS-IS routers.

- The mapping server functionality does not support a scenario where SID-mapping entries learned through one IS-IS instance are used by another IS-IS instance to determine the prefix-SID of a prefix. For example, mapping entries learnt from remote routers by 'router isis 1' cannot be used to calculate prefix-SIDs for prefixes learnt, advertised, or downloaded to FIB by 'router isis 2'. A mapping server is required for each IS-IS instance.
- Segment Routing Mapping Server does not support Virtual Routing and Forwarding (VRF) currently.

# Segment Routing and LDP Interoperability


IGP provides mechanisms through which segment routing (SR) interoperate with label distribution protocol (LDP). The control plane of segment routing co-exists with LDP.

The Segment Routing Mapping Server (SRMS) functionality in SR is used to advertise SIDs for destinations, in the LDP part of the network, that do not support SR. SRMS maintains and advertises segment identifier (SID) mapping entries for such destinations. IGP propagates the SRMS mapping entries and interacts with SRMS to determine the SID value when programming the forwarding plane. IGP installs prefixes and corresponding labels, into routing information base (RIB), that are used to program the forwarding information base (FIB).

## Example: Segment Routing LDP Interoperability

Consider a network with a mix of segment routing (SR) and label distribution protocol (LDP). A continuous multiprotocol label switching (MPLS) LSP (Labeled Switched Path) can be established by facilitating interoperability. One or more nodes in the SR domain act as segment routing mapping server (SRMS). SRMS advertises SID mappings on behalf of non-SR capable nodes. Each SR-capable node learns about SID assigned to non-SR capable nodes without explicitly configuring individual nodes.

Consider a network as shown in the following figure. This network is a mix of both LDP and SR-capable nodes.




In this mixed network:

- Nodes PE1, P2, P3, and P4 are SR-capable
- Nodes P4, P5, P6, and PE7 are LDP-capable
- Nodes PE1, P2, P3, and P4 are configured with segment routing global block (SRGB) range of 16000 to 23999
- Nodes PE1, P2, P3, and P4 are configured with node segments of 16001, 16002, 16003, and 16004 respectively

A service flow must be established from PE1 to PE3 over a continuous MPLS tunnel. This requires SR and LDP to interoperate.

## LDP-to-SR Traffic Direction




The traffic flow in the LDP-to-SR direction involves the following:

1. PE7 learns a service route with service label 40001 and BGP nhop PE1.
2. PE7 has an LDP label binding (24010) from the nhop P6 for the FEC PE1. PE7 forwards the packet to P6.
3. P6 has an LDP label binding (24003) from its nhop P5 for the FEC PE1. P6 forwards the packet to P5.
4. P5 has an LDP label binding (24008) from its nhop P4 for the FEC PE1. P5 forwards the packet to P4.
5. P4 does not have an LDP binding from its nhop P3 for the FEC PE1. But P4 has an SR node segment to the IGP route PE1. P4 forwards the packet to P3 and swaps its local LDP label (24008) for FEC PE1 by the equivalent node segment 16001. This process is called label merging.
6. P3 pops 16001, assuming PE1 has advertised its node segment 16001 with the penultimate-pop flag set and forwards to PE1.
7. PE1 receives the packet and processes the service label.

The end-to-end MPLS LSP is established from an LDP LSP from PE7 to P4 and the related node segment from P4 to PE1.

## SR-to-LDP Traffic Direction



Suppose that the operator configures P2 as a Segment Routing Mapping Server (SRMS) and advertises the mappings (1.1.1.7, 16007 for PE7). Because PE7 is non-SR capable, the operator configures that mapping policy at the SRMS; the SRMS advertises the mapping on behalf of the non-SR capable nodes. Multiple SRMS servers can be provisioned in a network for redundancy. The mapping server advertisements are only understood by the SR-capable nodes. The SR-capable routers install the related node segments in the MPLS data plane in exactly the same manner as if node segments were advertised by the nodes themselves.

The traffic flow in the SR to LDP direction involves the following:

1. PE1 learns a service route with service label 50010 and BGP nhop PE7.
2. PE1 has an SR label binding (16007) learned from the SRMS (P2) for PE7.
3. PE1 installs the node segment 16007 following the IGP shortest-path with nhop P3.
4. P3 swaps 16007 for 16007 and forwards to P4.
5. The nhop for P4 for the IGP route PE7 is non-SR capable, since P5 does not advertise the SR capability. However, P4 has an LDP label binding from that nhop for the same FEC (for example, LDP label 24009). P4 would then swap 16007 for 24009 and forward to P5. We refer to this process as label merging.
6. P5 swaps this label with the LDP label received from P6 (for example, LDP label 24006) and forwards to P6.
7. P6 pops the LDP label and forwards to PE7.
8. PE7 receives the packet and processes the service label.

The end-to-end MPLS LSP is established from an SR node segment from PE1 to P4 and an LDP LSP from P4 to PE7.

Observe that the capabilities provided by the SRMS are only required in the SR-to-LDP direction.



# Configuring Mapping Server

Perform these tasks to configure the mapping server and to add prefix-SID mapping entries in the active local mapping policy.

## Procedure

|               | Command or Action                                                                                                                                                                                                                                                                                                | Purpose                                                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router# <b>configure</b>                                                                                                                                                                                                                                    | Enters mode.                                                                                                                                             |
| <b>Step 2</b> | <b>segment-routing</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)#<br><b>segment-routing</b>                                                                                                                                                                                                             | Enables segment routing.                                                                                                                                 |
| <b>Step 3</b> | <b>mapping-server</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-sr)#<br><b>mapping-server</b>                                                                                                                                                                                                            | Enables mapping server configuration mode.                                                                                                               |
| <b>Step 4</b> | <b>prefix-sid-map</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-sr-ms)#<br><b>prefix-sid-map</b>                                                                                                                                                                                                         | Enables prefix-SID mapping configuration mode.<br><br><b>Note</b><br>Two-way prefix SID can be enabled directly under IS-IS or through a mapping server. |
| <b>Step 5</b> | <b>address-family ipv4   ipv6</b><br><b>Example:</b><br>This example shows the address-family for ipv4:<br>RP/0/RP0/CPU0:router(config-sr-ms-map)#<br><b>address-family ipv4</b><br><br>This example shows the address-family for ipv6:<br>RP/0/RP0/CPU0:router(config-sr-ms-map)#<br><b>address-family ipv6</b> | Configures address-family for IS-IS.                                                                                                                     |
| <b>Step 6</b> | <i>ip-address/prefix-length first-SID-value</i><br><b>range range</b>                                                                                                                                                                                                                                            | Adds SID-mapping entries in the active local mapping policy. In the configured example:                                                                  |

|               | Command or Action                                                                                                                                                          | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | <b>Example:</b><br><br>RP/0/RP0/CPU0:router(config-sr-ms-map-af) #<br>10.1.1.1/32 10 range 200<br>RP/0/RP0/CPU0:router(config-sr-ms-map-af) #<br>20.1.0.0/16 400 range 300 | <ul style="list-style-type: none"> <li>Prefix 10.1.1.1/32 is assigned prefix-SID 10, prefix 10.1.1.2/32 is assigned prefix-SID 11,..., prefix 10.1.1.199/32 is assigned prefix-SID 200</li> <li>Prefix 20.1.0.0/16 is assigned prefix-SID 400, prefix 20.2.0.0/16 is assigned prefix-SID 401,..., and so on.</li> </ul>                                                                                                                                                                                  |
| <b>Step 7</b> | Use the <b>commit</b> or <b>end</b> command.                                                                                                                               | <b>commit</b> —Saves the configuration changes and remains within the configuration session.<br><br><b>end</b> —Prompts user to take one of these actions: <ul style="list-style-type: none"> <li><b>Yes</b> — Saves configuration changes and exits the configuration session.</li> <li><b>No</b> —Exits the configuration session without committing the configuration changes.</li> <li><b>Cancel</b> —Remains in the configuration session, without committing the configuration changes.</li> </ul> |

Verify information about the locally configured prefix-to-SID mappings.



**Note** Specify the address family for IS-IS.

```
RP/0/RP0/CPU0:router# show segment-routing mapping-server prefix-sid-map ipv4
Prefix SID Index Range Flags
20.1.1.0/24 400 300
10.1.1.1/32 10 200
```

Number of mapping entries: 2

```
RP/0/RP0/CPU0:router# show segment-routing mapping-server prefix-sid-map ipv4 detail
Prefix
20.1.1.0/24
 SID Index: 400
 Range: 300
 Last Prefix: 20.2.44.0/24
 Last SID Index: 699
 Flags:
10.1.1.1/32
 SID Index: 10
 Range: 200
 Last Prefix: 10.1.1.200/32
 Last SID Index: 209
 Flags:
```

Number of mapping entries: 2

**What to do next**

Enable the advertisement of the local SID-mapping policy in the IGP.

## Enable Mapping Advertisement

In addition to configuring the static mapping policy, you must enable the advertisement of the mappings in the IGP.

Perform these steps to enable the IGP to advertise the locally configured prefix-SID mapping.

## Configure Mapping Advertisement for IS-IS

**Procedure**

|               | Command or Action                                                                                                                                                                                                                                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>router isis</b> <i>instance-id</i><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config) # <b>router isis</b><br><b>1</b>                                                                                                                                                                                                                       | Enables IS-IS routing for the specified routing instance, and places the router in router configuration mode.<br><ul style="list-style-type: none"> <li>You can change the level of routing to be performed by a particular routing instance by using the <b>is-type</b> router configuration command.</li> </ul> |
| <b>Step 2</b> | <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } [ <b>unicast</b> ]<br><b>Example:</b><br>The following is an example for ipv4 address family:<br>RP/0/RP0/CPU0:router(config-isis) #<br><b>address-family ipv4 unicast</b>                                                                                                                  | Specifies the IPv4 or IPv6 address family, and enters router address family configuration mode.                                                                                                                                                                                                                   |
| <b>Step 3</b> | <b>segment-routing prefix-sid-map</b><br><b>advertise-local</b> [ <b>domain-wide</b> ]<br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-isis-af) #<br><b>segment-routing prefix-sid-map</b><br><b>advertise-local</b><br>RP/0/RP0/CPU0:router(config-isis-af) #<br><b>segment-routing prefix-sid-map</b><br><b>advertise-local domain-wide</b> | Configures IS-IS to advertise locally configured prefix-SID mappings. Use the <b>domain-wide</b> option to advertise the prefix-SID mappings between IS-IS Level 1 and Level 2 routers.                                                                                                                           |

|               | Command or Action                            | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 4</b> | Use the <b>commit</b> or <b>end</b> command. | <b>commit</b> —Saves the configuration changes and remains within the configuration session.<br><b>end</b> —Prompts user to take one of these actions: <ul style="list-style-type: none"> <li>• <b>Yes</b> — Saves configuration changes and exits the configuration session.</li> <li>• <b>No</b> —Exits the configuration session without committing the configuration changes.</li> <li>• <b>Cancel</b> —Remains in the configuration session, without committing the configuration changes.</li> </ul> |

Verify IS-IS prefix-SID mapping advertisement and TLV.

```
RP/0/RP0/CPU0:router# show isis database verbose
```

```
<...removed...>
```

```
SID Binding: 10.1.1.1/32 F:0 M:0 S:0 D:0 A:0 Weight:0 Range:200
SID: Start:10, Algorithm:0, R:0 N:0 P:0 E:0 V:0 L:0
SID Binding: 20.1.1.0/24 F:0 M:0 S:0 D:0 A:0 Weight:0 Range:300
SID: Start:400, Algorithm:0, R:0 N:0 P:0 E:0 V:0 L:0
```

## Configure Mapping Advertisement for OSPF

### Procedure

|               | Command or Action                                                                                                                                                    | Purpose                                                                                                      |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>router ospf process-name</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config)# <b>router ospf 1</b>                                                             | Enables OSPF routing for the specified routing instance, and places the router in router configuration mode. |
| <b>Step 2</b> | <b>segment-routing prefix-sid-map advertise-local</b><br><b>Example:</b><br>RP/0/RP0/CPU0:router(config-ospf)# <b>segment-routing prefix-sid-map advertise-local</b> | Configures OSPF to advertise locally configured prefix-SID mappings.                                         |
| <b>Step 3</b> | Use the <b>commit</b> or <b>end</b> command.                                                                                                                         | <b>commit</b> —Saves the configuration changes and remains within the configuration session.                 |

|  | Command or Action | Purpose                                                                                                                                                                                                                                                                                                                                                                                                    |
|--|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  |                   | <b>end</b> —Prompts user to take one of these actions: <ul style="list-style-type: none"> <li>• <b>Yes</b> — Saves configuration changes and exits the configuration session.</li> <li>• <b>No</b> —Exits the configuration session without committing the configuration changes.</li> <li>• <b>Cancel</b> —Remains in the configuration session, without committing the configuration changes.</li> </ul> |

Verify OSP prefix-SID mapping advertisement and TLV.

```
RP/0/RP0/CPU0:router# show ospf database opaque-area
```

```
<...removed...>
```

```
Extended Prefix Range TLV: Length: 24
```

```
AF : 0
Prefix : 10.1.1.1/32
Range Size: 200
Flags : 0x0
```

```
SID sub-TLV: Length: 8
```

```
Flags : 0x60
MTID : 0
Algo : 0
SID Index : 10
```

## Enable Mapping Client

By default, mapping client functionality is enabled.

You can disable the mapping client functionality by using the **segment-routing prefix-sid-map receive disable** command.

You can re-enable the mapping client functionality by using the **segment-routing prefix-sid-map receive** command.

The following example shows how to enable the mapping client for IS-IS:

```
RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# segment-routing prefix-sid-map receive
```

The following example shows how to enable the mapping client for OSPF:

```
RP/0/RP0/CPU0:router(config)# router ospf 1
RP/0/RP0/CPU0:router(config-ospf)# segment-routing prefix-sid-map receive disable
RP/0/RP0/CPU0:router(config-ospf)# commit
```



## CHAPTER 15

# Enabling Segment Routing Flexible Algorithm

Segment Routing Flexible Algorithm allows operators to customize IGP shortest path computation according to their own needs. An operator can assign custom SR prefix-SIDs to realize forwarding beyond link-cost-based SPF. As a result, Flexible Algorithm provides a traffic engineered path automatically computed by the IGP to any destination reachable by the IGP.

The SR architecture associates prefix-SIDs to an algorithm which defines how the path is computed. Flexible Algorithm allows for user-defined algorithms where the IGP computes paths based on a user-defined combination of metric type and constraint.

This document describes the IS-IS and OSPF extensions to support Segment Routing Flexible Algorithm on an MPLS data-plane.

- [Prerequisites for Flexible Algorithm, on page 573](#)
- [Building Blocks of Segment Routing Flexible Algorithm, on page 573](#)
- [Configuring Flexible Algorithm, on page 582](#)
- [Example: Configuring IS-IS Flexible Algorithm, on page 599](#)
- [Example: Configuring OSPF Flexible Algorithm, on page 599](#)
- [User-Defined Generic Metric Support for IS-IS Flex Algo, on page 600](#)
- [Example: Traffic Steering to Flexible Algorithm Paths, on page 605](#)
- [Delay Normalization, on page 608](#)

## Prerequisites for Flexible Algorithm

Segment routing must be enabled on the router before the Flexible Algorithm functionality is activated.

## Building Blocks of Segment Routing Flexible Algorithm

This section describes the building blocks that are required to support the SR Flexible Algorithm functionality in IS-IS and OSPF.

## Flexible Algorithm Definition

Many possible constraints may be used to compute a path over a network. Some networks are deployed with multiple planes. A simple form of constraint may be to use a particular plane. A more sophisticated form of constraint can include some extended metric, like delay, as described in [RFC7810]. Even more advanced

case could be to restrict the path and avoid links with certain affinities. Combinations of these are also possible. To provide a maximum flexibility, the mapping between the algorithm value and its meaning can be defined by the user. When all the routers in the domain have the common understanding what the particular algorithm value represents, the computation for such algorithm is consistent and the traffic is not subject to looping. Here, since the meaning of the algorithm is not defined by any standard, but is defined by the user, it is called a Flexible Algorithm.

## Flexible Algorithm Membership

An algorithm defines how the best path is computed by IGP. Routers advertise the support for the algorithm as a node capability. Prefix-SIDs are also advertised with an algorithm value and are tightly coupled with the algorithm itself.

An algorithm is a one octet value. Values from 128 to 255 are reserved for user defined values and are used for Flexible Algorithm representation.

## Flexible Algorithm Definition Advertisement

To guarantee the loop free forwarding for paths computed for a particular Flexible Algorithm, all routers in the network must share the same definition of the Flexible Algorithm. This is achieved by dedicated router(s) advertising the definition of each Flexible Algorithm. Such advertisement is associated with the priority to make sure that all routers will agree on a single and consistent definition for each Flexible Algorithm.

Definition of Flexible Algorithm includes:

- Metric type
- Affinity constraints
- Exclude SRLG constraint
- Minimum bandwidth constraint
- Maximum delay constraint

To enable the router to advertise the definition for the particular Flexible Algorithm, **advertise-definition** command is used. At least one router in the area, preferably two for redundancy, must advertise the Flexible Algorithm definition. Without the valid definition being advertised, the Flexible Algorithm will not be functional.

## Flexible Algorithm Link Attribute Advertisement

Various link attributes may be used during the Flexible Algorithm path calculation. For example, include or exclude rules based on link affinities can be part of the Flexible Algorithm definition, as defined in [RFC9350](#) (IGP Flexible Algorithm).

Link attribute advertisements used during Flexible Algorithm calculation must use the Application-Specific Link Attribute (ASLA) advertisements, as defined in [RFC8919](#) (IS-IS) and [RFC8920](#) (OSPF). In the case of IS-IS, if the L-Flag is set in the ASLA advertisement, then legacy advertisements (IS-IS Extended Reachability TLV) are used instead.

## Flexible Algorithm Prefix-SID Advertisement

To be able to forward traffic on a Flexible Algorithm specific path, all routers participating in the Flexible Algorithm will install a MPLS labeled path for the Flexible Algorithm specific SID that is advertised for the prefix. Only prefixes for which the Flexible Algorithm specific Prefix-SID is advertised is subject to Flexible Algorithm specific forwarding.

## Calculation of Flexible Algorithm Path

*Table 82: Feature History Table*

| Feature Name                                     | Release Information | Feature Description                                                                                    |
|--------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------|
| OSPF: Microloop Avoidance for Flexible Algorithm | Release 7.4.1       | This feature extends the current OSPF Flexible Algorithm functionality to support Microloop Avoidance. |

*Table 83: Feature History Table*

| Feature Name                                     | Release Information | Feature Description                                                                                    |
|--------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------|
| OSPF: Microloop Avoidance for Flexible Algorithm | Release 7.3.2       | This feature extends the current OSPF Flexible Algorithm functionality to support Microloop Avoidance. |
| OSPF: TI-LFA for Flexible Algorithm              | Release 7.3.1       | This feature extends the current OSPF Flexible Algorithm functionality to support TI-LFA.              |

A router may compute path for multiple Flexible Algorithms. A router must be configured to support particular Flexible Algorithm before it can compute any path for such Flexible Algorithm. A router must have a valid definition of the Flexible Algorithm before Flexible Algorithm is used.

The router uses the following rules to prune links from the topology during the Flexible Algorithm computation:

- All nodes that don't advertise support for Flexible Algorithm are pruned from the topology.
- Affinities:
  - Check if any exclude affinity rule is part of the Flexible Algorithm Definition. If such exclude rule exists, check if any color that is part of the exclude rule is also set on the link. If such a color is set, the link must be pruned from the computation.
  - Check if any include-any affinity rule is part of the Flexible Algorithm Definition. If such include-any rule exists, check if any color that is part of the include-any rule is also set on the link. If no such color is set, the link must be pruned from the computation.
  - Check if any include-all affinity rule is part of the Flexible Algorithm Definition. If such include-all rule exists, check if all colors that are part of the include-all rule are also set on the link. If all such colors are not set on the link, the link must be pruned from the computation.






---

**Note** See [Flexible Algorithm Affinity Constraint](#).

---

- If the Flexible Algorithm definition includes an "exclude SRLG" rule, then all links that are part of such SRLG are pruned from the topology.




---

**Note** See [Flexible Algorithm with Exclude SRLG Constraint, on page 588](#).

---

- If the minimum bandwidth constraint is configured, any link that does not meet the minimum link bandwidth threshold is pruned from the topology.
- If the maximum delay constraint is configured, any link that exceeds the maximum unidirectional link delay is pruned from the topology.
- Router uses the metric that is part of the Flexible Algorithm definition. If the metric isn't advertised for the particular link, the link is pruned from the topology.

### Configuring Microloop Avoidance for Flexible Algorithm

By default, Microloop Avoidance per Flexible Algorithm instance follows Microloop Avoidance configuration for algo-0. For information about configuring Microloop Avoidance, see [Configure Segment Routing Microloop Avoidance, on page 553](#).

You can disable Microloop Avoidance for Flexible Algorithm using the following commands:

```
router isis instance flex-algo algo microloop avoidance disable
router ospf process flex-algo algo microloop avoidance disable
```

### Configuring LFA / TI-LFA for Flexible Algorithm

By default, LFA/TI-LFA per Flexible Algorithm instance follows LFA/TI-LFA configuration for algo-0. For information about configuring TI-LFA, see [Configure Topology-Independent Loop-Free Alternate \(TI-LFA\), on page 527](#).

You can disable TI-LFA for Flexible Algorithm using the following commands:

```
router isis instance flex-algo algo fast-reroute disable
router ospf process flex-algo algo fast-reroute disable
```

## Flexible Algorithm Affinity Constraint


Table 84: Feature History Table

| Feature Name                               | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IS-IS: Flexible Algorithm Reverse Affinity | Release 7.9.1       | <p>This feature enhances the IS-IS Flexible Algorithm link admin group (affinity) constraint to include link colors on links in the reverse direction toward the calculating router.</p> <p>The ability to apply affinity constraints in the reverse direction provides additional control for IS-IS Flexible Algorithm path computation.</p> <p>This feature introduces the <b>reverse</b> keyword to the <b>router isis instance flex-algo algo affinity</b> command.</p> |

You can apply a color or name to links or interfaces by assigning affinity bit-maps to them. You can then specify an affinity (or relationship) between a Flexible Algorithm path and link colors in the forwarding direction. Flexible Algorithm computes a path that includes or excludes links that have specific colors, or combinations of colors.

- Affinity “blue” is assigned to interface on node A; exclude affinity "blue": Link A-B is pruned from path calculation

### FA 128: Metric IGP and Exclude Affinity “blue”




In Cisco IOS XR release 7.9.1, for IS-IS Flexible Algorithm, you can also specify a reverse affinity between a Flexible Algorithm path and link colors (in the direction toward the computing router). Flexible Algorithm computes a path that includes or excludes links in the reverse direction that have specific colors, or combinations of colors.

For example, on a point-to-point link between endpoints A and B and for the traffic flowing in the direction from A to B, the input errors can only be detected at node B. You may measure the rate of such input errors and set certain 'color' on a link locally on node B when the input error rate crosses a certain threshold.

- Affinity “blue” is assigned to interface on node B; exclude affinity "blue": Link A-B is not pruned from path calculation


### FA 128: Metric IGP and Exclude Affinity “blue”



With IS-IS Flexible Algorithm Reverse Affinity, when Flex-Algorithm calculation processes link A to B, it may look at the 'colors' of the link in the reverse direction (link B to A). This enables you to exclude this link from the Flex-Algorithm topology.

- Affinity “blue” is assigned to interface on node B; exclude reverse-affinity "blue": Link A-B is pruned from path calculation

### FA 128: Metric IGP and Exclude-Reverse Affinity “blue”



## Installation of Forwarding Entries for Flexible Algorithm Paths

Flexible Algorithm path to any prefix must be installed in the forwarding using the Prefix-SID that was advertised for such Flexible Algorithm. If the Prefix-SID for Flexible Algorithm is not known, such Flexible Algorithm path is not installed in forwarding for such prefix..

Only MPLS to MPLS entries are installed for a Flexible Algorithm path. No IP to IP or IP to MPLS entries are installed. These follow the native IGP paths computed based on the default algorithm and regular IGP metrics.

## Flexible Algorithm Prefix-SID Redistribution

Table 85: Feature History Table

| Feature Name                                                                | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Flexible Algorithm Prefix-SID Redistribution for External Route Propagation | Release 7.5.2       | <p>You can now propagate flexible algorithm prefix-SIDs and their algorithm-specific metric between different IGP domains, such as OSPF to IS-IS RIP to OSPF. With this functionality enabling interdomain traffic engineering, you can export flexible algorithm labels from the OSPF domain to other domains and import the labels from other domains into OSPF.</p> <p>The <b>show ospf route flex-algo</b> command has been modified to include additional attributes to indicate the external routes.</p> |

Prefix redistribution from IS-IS to another IS-IS instance or protocol was limited to SR algorithm 0 (regular SPF) prefix SIDs; SR algorithm 1 (Strict SPF) and SR algorithms 128-255 (Flexible Algorithm) prefix SIDs were not redistributed along with the prefix. The Segment Routing IS-IS Flexible Algorithm Prefix SID Redistribution feature allows redistribution of strict and flexible algorithms prefix SIDs from IS-IS to another IS-IS instance or protocols. This feature is enabled automatically when you configure redistribution of IS-IS Routes with strict or flexible algorithm SIDs.

### Configuration Example

The following example shows how to configure redistribute and flexible algorithm to enable external routes.

```
RP/0/RP0/CPU0:ios(config)#router ospf 1
RP/0/RP0/CPU0:ios(config-ospf)#segment-routing mpls
RP/0/RP0/CPU0:ios(config-ospf)#segment-routing forwarding mpls
RP/0/RP0/CPU0:ios(config-ospf)#redistribute isis 2 route-policy loopback-type
RP/0/RP0/CPU0:ios(config-ospf)#flex-algo 240
RP/0/RP0/CPU0:ios(config-ospf-flex-algo)#metric-type delay
RP/0/RP0/CPU0:ios(config-ospf-flex-algo)#prefix-metric
RP/0/RP0/CPU0:ios(config-ospf-flex-algo)#advertise-definition
```

### Verification

This following show output displays the route-type as 'Extern' for the external routes.

```
Router#show ospf routes flex-algo 240 route-type external detail
Route Table of ospf-1 with router ID 192.168.0.2 (VRF default)

Algorithm 240

Route entry for 192.168.4.3/32, Metric 220, SID 536, Label 16536
Priority : Medium
```

```

Route type : Extern Type 1
Last updated : Apr 25 14:30:12.718
Flags: Inuse

Prefix Contrib Algo 240 SID 536
From 192.168.0.4 Route-type 5
Total Metric : 220 Base metric 20 FAPM 20
Contrib Flags : Inuse, Reachable
SID Flags : PHP off, Index, Global, Valid

Path: 10.1.1.3, from 192.168.0.4, via GigabitEthernet0/2/0/2
Out Label : 16536
Weight : 0
Area : 0

Path: 10.1.2.3, from 192.168.0.4, via GigabitEthernet0/2/0/3
Out Label : 16536
Weight : 0
Area : 0

Path: 10.2.1.5, from 192.168.0.4, via GigabitEthernet0/2/0/4
Out Label : 16536
Weight : 0
Area : 0

Route entry for 192.168.4.5/32, Metric 120, SID 556, Label 16556
Priority : Medium

Route type : Extern Type 1
Last updated : Apr 25 14:30:12.724
Flags: Inuse

Prefix Contrib Algo 240 SID 556
From 192.168.0.3 Route-type 5
Total Metric : 120 Base metric 1 FAPM 20
Contrib Flags : Inuse, Reachable
SID Flags : PHP off, Index, Global, Valid

Path: 10.1.1.3, from 192.168.0.3, via GigabitEthernet0/2/0/2
Out Label : 16556
Weight : 0
Area : 0

Path: 10.1.2.3, from 192.168.0.3, via GigabitEthernet0/2/0/3
Out Label : 16556
Weight : 0
Area : 0

```

The following show output displays label information for flexible algorithm and its corresponding metric as added in RIB:

```

RP/0/RP0/CPU0:ios# show route 192.168.0.2/32 detail
Wed Apr 6 16:24:46.021 IST

Routing entry for 192.168.0.2/32
 Known via "ospf 1", distance 110, metric 2, labeled SR, type intra area
 Installed Apr 6 15:51:57.973 for 00:32:48
 Routing Descriptor Blocks
 10.10.10.2, from 192.168.0.2, via GigabitEthernet0/2/0/0, Protected
 Route metric is 2
 Label: 0x3 (3)
 Tunnel ID: None
 Binding Label: None

```

```

Extended communities count: 0
Path id:1 Path ref count:0
NHID:0x1(Ref:1)
Backup path id:65
OSPF area: 1
10.11.11.2, from 192.168.0.2, via GigabitEthernet0/2/0/1, Backup (Local-LFA)
Route metric is 6
Label: 0x3 (3)
Tunnel ID: None
Binding Label: None
Extended communities count: 0
Path id:65 Path ref count:1
NHID:0x2(Ref:1)
OSPF area:
Route version is 0x12 (18)
Local Label: 0x3ee6 (16102)
Local Label Algo Set (ID, Label, Metric): (1, 16202, 0),(128, 17282, 2)
IP Precedence: Not Set
QoS Group ID: Not Set
Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_NON_RECURSIVE_MEDIUM (7) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 1, Download Version 38
No advertising protos.

```

## Flexible Algorithm Prefix Metric

**Table 86: Feature History Table**

| Feature Name                                      | Release Information | Feature Description                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prefix Metric support for OSPF Flexible Algorithm | Release 7.5.1       | This feature extends the current OSPF Flexible Algorithm functionality to introduce a Flexible Algorithm-specific prefix-metric in the OSPF prefix advertisement. The prefix-metric provides a way to compute the best end-to-end Flexible Algorithm optimized paths across multiple areas or domains. |

A limitation of the existing Flexible Algorithm functionality in IS-IS and OSPF is the inability to compute the best path to a prefix in a remote area or remote IGP domain. Prefixes are advertised between IS-IS areas, OSPF processes, or between protocol domains, but the existing prefix metric does not reflect any of the constraints used for Flexible Algorithm path. Although the best Flexible Algorithm path can be computed to the inter-area or redistributed prefix inside the area, the path may not represent the overall best path through multiple areas or IGP domains.

The Flexible Algorithm Prefix Metric feature introduces a Flexible Algorithm-specific prefix-metric in the IS-IS and OSPF prefix advertisement. The prefix-metric provides a way to compute the best end-to-end Flexible Algorithm optimized paths across multiple areas or domains.



**Note** The Flexible Algorithm definition must be consistent between domains or areas. Refer to section 8 and section 9 in IETF draft <https://datatracker.ietf.org/doc/draft-ietf-lsr-flex-algo/>.

# Configuring Flexible Algorithm

Table 87: Feature History Table

| Feature Name                                          | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IS-IS Enhancements: max-metric and data plane updates | Release 7.8.1       | The new <b>anomaly</b> optional keyword is introduced to <b>affinity flex-algo</b> command. This keyword helps to advertise the flex-algo affinity when the performance measurement signals a link anomaly, such as an excessive delay on a link. You could use the anomaly option to exclude the link from flex-algo path computations.<br><a href="#">affinity flex-algo</a> |

Table 88: Feature History Table

| Feature Name                          | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TE Metric Support for IS-IS Flex Algo | Release 7.4.1       | Flexible Algorithm allows for user-defined algorithms where the IGP computes paths based on a user-defined combination of metric type (path optimization objective) and constraint.<br><br>This feature adds support for TE metric as a metric type for IS-IS Flexible Algorithm. This allows the TE metric, along with IGP and delay metrics, to be used when running shortest path computations. |

The following IS-IS and OSPF configuration sub-mode is used to configure Flexible Algorithm:

```
router isis instance flex-algo algo
router ospf process flex-algo algo
algo—value from 128 to 255
```

## Configuring Flexible Algorithm Definitions

The following commands are used to configure Flexible Algorithm definition under the flex-algo sub-mode:

```
• router isis instance flex-algo algo metric-type {delay | te}
 router ospf process flex-algo algo metric-type {delay | te-metric}
```



**Note** By default the IGP metric is used. If delay or TE metric is enabled, the advertised delay or TE metric on the link is used as a metric for Flexible Algorithm computation.



**Note** See [Flexible Algorithm Link Attribute Advertisement Behavior, on page 584](#) for TE metric behaviors.

```
• router isis instance flex-algo algo affinity [reverse] { include-any | include-all |
 exclude-any} name1, name2, ...

router ospf process flex-algo algo affinity { include-any | include-all | exclude-any}
 name1, name2, ...

name—name of the affinity map
```



**Note** See [Flexible Algorithm Affinity Constraint, on page 577](#) for information about affinity constraint behaviors.

```
• router isis instance flex-algo algo priority priority value

router ospf process flex-algo algo priority priority value

priority value—priority used during the Flexible Algorithm definition election.
```

The following command is used to include the Flexible Algorithm prefix metric in the advertised Flexible Algorithm definition in IS-IS and OSPF :

```
router isis instance flex-algo algo prefix-metric

router ospf process flex-algo algo prefix-metric
```

The following command is used to enable advertisement of the Flexible Algorithm definition in IS-IS:

```
router isis instance flex-algo algo advertise-definition
```

### Configuring Affinity

The following command is used for defining the affinity-map. Affinity-map associates the name with the particular bit positions in the Extended Admin Group bitmask.

```
router isis instance flex-algo algo affinity-map name bit-position bit number

router ospf process flex-algo algo affinity-map name bit-position bit number
```

- *name*—name of the affinity-map.
- *bit number*—bit position in the Extended Admin Group bitmask.



With the IOS XR Release 7.8.1, the new optional keyword **anomaly** is introduced to the **interface** submode of **affinity flex-algo**. This keyword option helps to advertise flex-algo affinity on PM anomaly. The following command is used to associate the affinity with an interface:

```
router isis instance interface type interface-path-id affinity flex-algo anomaly name 1,
name 2, ...

router ospf process area area interface type interface-path-id affinity flex-algo anomaly
name 1, name 2, ...
```

*name*—name of the affinity-map

You can configure both normal and anomaly values. For the following example, the **blue** affinity is advertised. However, if a metric is received with the anomaly flag set, it will change to **red**:

```
Router# configure
Router(config)# router isis 1
Router(config-isis)#flex-algo 128
Router(config-isis-flex-algo)# interface GigabitEthernet0/0/0/2
Router(config-isis-flex-algo)# affinity flex-algo blue
Router(config-isis-flex-algo)# affinity flex-algo anomaly red
```

### Configuring Prefix-SID Advertisement

The following command is used to advertise prefix-SID for default and strict-SPF algorithm:

```
router isis instance interface type interface-path-id address-family {ipv4 | ipv6} [unicast]
prefix-sid [strict-spf | algorithm algorithm-number] [index | absolute] sid value

router ospf process area area interface Loopback interface-instance prefix-sid [strict-spf
| algorithm algorithm-number] [index | absolute] sid value
```

- *algorithm-number*—Flexible Algorithm number
- *sid value*—SID value

## Flexible Algorithm Link Attribute Advertisement Behavior

Table 89: Feature History Table

| Feature Name                                                  | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advertisement of Link Attributes for IS-IS Flexible Algorithm | Release 7.4.1       | <p>Link attribute advertisements used during Flexible Algorithm path calculation must use the Application-Specific Link Attribute (ASLA) advertisements, as defined in IETF draft <a href="#">draft-ietf-lsr-flex-algo</a>.</p> <p>This feature introduces support for ASLA advertisements during IS-IS Flexible Algorithm path calculation.</p> |

The following tables explain the behaviors for advertising (transmitting) and processing (receiving) Flexible Algorithm link attributes.

**Table 90: OSPF**

| Link Attribute                        | Transmit                                                                                                                                                                                                                                                                                                                                                                                                                | Receive                                                                                                                                                                                                         |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Delay Metric                     | IOS XR OSPF Flex Algo implementation advertises the link delay metric value using the OSPF ASLA sub-TLV with the F-bit set.                                                                                                                                                                                                                                                                                             | IOS XR OSPF only uses the link delay metric advertised in the ASLA sub-TLV for Flex Algo.<br><br>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks.           |
| Link TE Metric                        | IOS XR OSPF Flex Algo implementation advertises the link TE metric value using the OSPF ASLA sub-TLV with the F-bit set.<br><br>The link TE metric values advertised are configured under SR-TE.                                                                                                                                                                                                                        | IOS XR OSPF only uses the TE metric advertised in the ASLA sub-TLV for Flex Algo.<br><br>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks.                   |
| Link Admin Group/Extended Admin Group | IOS XR OSPF Flex Algo implementation advertises the link admin group value using both link admin group (AG) and link extended admin group (EAG) encoding using the OSPF ASLA sub-TLV with the F-bit set.<br><br>The link admin group values advertised can be configured directly under the IGP and are therefore FA-specific. Otherwise, they will be derived from the link admin group values configured under SR-TE. | IOS XR OSPF only uses the AG/EAG (either one or both) advertised in the ASLA sub-TLV for Flex Algo.<br><br>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks. |

Table 91: IS-IS

| Link Attribute                        | Transmit                                                                                                                                                                                                                                                                                                                                                                                                                              | Receive                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Delay Metric                     | IOS XR IS-IS Flex Algo implementation advertises the link delay metric value using <b>both</b> the IS-IS Extended Reachability TLV and the IS-IS ASLA.                                                                                                                                                                                                                                                                                | <p>By default, IOS XR IS-IS Flex Algo implementation prefers the link delay metric value received in the IS-IS ASLA. Otherwise, it will use link delay metric value received in the IS-IS Extended Reachability TLV.</p> <p>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks.</p> <p>If the incoming ASLA includes the L-Flag, implementation derives the link delay metric value from the IS-IS Extended Reachability TLV.</p> <p>You can configure the IOS XR IS-IS Flex Algo implementation to strictly use the link delay metric value received in the IS-IS ASLA. See <a href="#">Strict IS-IS ASLA Link Attribute</a>, on page 587.</p> |
| Link TE Metric                        | <p>IOS XR IS-IS Flex Algo implementation advertises the link TE metric value using the IS-IS ASLA.</p> <p>The link TE metric values advertised can be configured directly under the IGP and are therefore FA-specific. Otherwise, they will be derived from the link TE metric values configured under SR-TE.</p> <p>See <a href="#">Flexible Algorithm-Specific TE Metric</a>, on page 587.</p>                                      | <p>IOS XR IS-IS Flex Algo implementation processes the link TE metric value received in the IS-IS ASLA.</p> <p>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks.</p> <p>If incoming ASLA includes the L-Flag, implementation derives the link TE metric value from the IS-IS Extended Reachability TLV.</p>                                                                                                                                                                                                                                                                                                                                   |
| Link Admin Group/Extended Admin Group | <p>IOS XR IS-IS Flex Algo implementation advertises the affinity value as both the link admin group (AG) TLV and the link extended admin group (EAG) TLV using the IS-IS ASLA when its value falls within the first 32 bits. Otherwise, the affinity value is advertised only as link EAG TLV using the IS-IS ASLA.</p> <p>The admin group values advertised are configured directly under the IGP and are therefore FA-specific.</p> | <p>IOS XR IS-IS Flex Algo implementation processes the affinity value received as either the link admin group TLV or link extended admin group TLV in the IS-IS ASLA.</p> <p>ASLA sub-TLV is supported with non-zero-length or with zero-length Application Identifier Bit Masks.</p> <p>If incoming ASLA includes the L-Flag, implementation derives the affinity value from the IS-IS Extended Reachability TLV.</p>                                                                                                                                                                                                                                                                           |

| Link Attribute | Transmit                                                                          | Receive                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link SRLG      | IOS XR IS-IS LFA implementation advertises the link SRLG value in the IS-IS ASLA. | IOS XR IS-IS LFA implementation processes the link SRLG value received in the IS-IS ASLA.<br><br>If incoming ASLA includes the L-Flag, implementation derives the link SRLG value from the IS-IS Extended Reachability TLV. |

## Strict IS-IS ASLA Link Attribute

Use the following command to configure the IOS XR IS-IS Flex Algo implementation to strictly use the link delay metric value received in the IS-IS ASLA:

```
router isis instance-id receive application flex-algo delay app-only
```

## Flexible Algorithm-Specific TE Metric

Use the following command to configure the Flexible Algorithm-specific TE metric value under IS-IS, where *metric\_value* is from 1 to 16777214:

- **router isis instance interface type interface-path-id address-family { ipv4 | ipv6 } [unicast] te-metric flex-algo metric\_value [level {1 | 2}]**

The following example shows how to configure the IS-IS Flexible Algorithm-specific TE metric value to 50:

```
Router(config)# router isis 1
Router(config-isis)# interface HundredGigE 0/0/0/2
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-if-af)# te-metric flex-algo 50
```

Use the following command to configure the Flexible Algorithm-specific TE metric value under OSPF, where *metric\_value* is from 1 to 2147483647:

- **router ospf process-name area area interface type interface-path-id te-metric flex-algo metric\_value**

The following example shows how to configure the OSPF Flexible Algorithm-specific TE metric value to 50:

```
Router(config)# router ospf 1
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface HundredGigE 0/0/0/2
Router(config-ospf-ar-if)# te-metric flex-algo 50
```

## Flexible Algorithm with Exclude SRLG Constraint

Table 92: Feature History Table

| Feature Name                                      | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Flexible Algorithm to Exclude SRLGs for OSPF      | Release 7.5.2       | You can now configure the flexible algorithm to exclude any link belonging to the Shared Risk Link Groups (SRLGs) from the path computation for OSPF. The ability to exclude the at-risk links ensures that the rest of the links in the network remain unaffected.                                                                                                                                                            |
| IS-IS Flexible Algorithm: Exclude-SRLG Constraint | Release 7.5.1       | <p>This feature allows the Flexible Algorithm definition to specify Shared Risk Link Groups (SRLGs) that the operator wants to exclude during the Flex-Algorithm path computation. The ability to exclude the at-risk links ensures that the rest of the links in the network remain unaffected.</p> <p>This allows the setup of disjoint paths between two or more Flex Algos by leveraging deployed SRLG configurations.</p> |

This feature allows the Flexible Algorithm definition to specify Shared Risk Link Groups (SRLGs) that the operator wants to exclude during the Flex-Algorithm path computation. A set of links that share a resource whose failure can affect all links in the set constitute a SRLG. An SRLG provides an indication of which links in the network might be at risk from the same failure.

This allows the setup of disjoint paths between two or more Flex Algos by leveraging deployed SRLG configurations. For example, multiple Flex Algos could be defined by excluding all SRLGs except one. Each FA will prune the links belonging to the excluded SRLGs from its topology on which it computes its paths.

This provides a new alternative to creating disjoint paths with FA, in addition to leveraging FA with link admin group (affinity) constraints.

The Flexible Algorithm definition (FAD) can advertise SRLGs that you want to exclude during the Flexible Algorithm path computation. The IS-IS Flexible Algorithm Exclude SRLG Sub-TLV (FAESRLG) is used to advertise the exclude rule that is used during the Flexible Algorithm path calculation, as specified in IETF draft <https://datatracker.ietf.org/doc/draft-ietf-lsr-flex-algo/>

The Flexible Algorithm path computation checks if an “exclude SRLG” rule is part of the FAD. If an “exclude SRLG” rule exists, it then checks if the link is part of an SRLG that is also part of the “exclude SRLG” rule. If the link is part of an excluded SRLG, the link is pruned from the path computation.

The figure below shows a topology configured with the following flex algos:

- Flex algo 128: metric IGP and exclude SRLG X constraint


- Flex algo 129: metric IGP and exclude SRLG Y constraint

The horizontal links between nodes 3 and 4 and between 2 and 5 are part of SRLG group X. The diagonal links between nodes 3 and 5 and between 2 and 4 are part of SRLG group Y. As a result, traffic from node 1 to node 6's FA 128 prefix SID (16806) avoids interfaces part of SRLG X. While traffic from node 1 to node 6's FA 129 prefix SID (16906) avoids interfaces part of SRLG Y.



**Note** See [Constraints, on page 267](#) section in the *Configure SR-TE Policies* chapter for information about configuring SR policies with Flex-Algo constraints.

**Figure 44: Flex Algo with Exclude SRLG Constraint**



## Configuration

Use the **router isis instance address-family ipv4 unicast advertise application flex-algo link-attributes srlg** command to enable the Flexible Algorithm ASLA-specific advertisement of SRLGs.

Use the **router isis instance flex-algo algo srlg exclude-any srlg-name . . . srlg-name** command to configure the SRLG constraint which is advertised in the Flexible Algorithm definition (FAD) if the FAD advertisement is enabled under the flex-algo sub-mode. You can specify up to 32 SRLG names.

The SRLG configuration (value and port mapping) is performed under the global SRLG sub-mode. Refer to [MPLS Traffic Engineering Shared Risk Link Groups](#) for more information.

## Example

The following example shows how to enable the Flexible Algorithm ASLA-specific advertisement of SRLGs and to exclude SRLG groups from Flexible Algorithm path computation:

```
RP/0/RP0/CPU0:router(config)# srlg
RP/0/RP0/CPU0:router(config-srlg)# interface HunGigE0/0/0/0
RP/0/RP0/CPU0:router(config-srlg-if)# name groupX
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/1
RP/0/RP0/CPU0:router(config-srlg-if)# name groupX
RP/0/RP0/CPU0:router(config-srlg-if)# exit

RP/0/RP0/CPU0:router(config-srlg)# interface HunGigE0/0/1/0
```

```

RP/0/RP0/CPU0:router(config-srlg-if)# name groupY
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/1/1
RP/0/RP0/CPU0:router(config-srlg-if)# name groupY
RP/0/RP0/CPU0:router(config-srlg-if)# exit

RP/0/RP0/CPU0:router(config-srlg)# name groupX value 100
RP/0/RP0/CPU0:router(config-srlg)# name groupY value 200
RP/0/RP0/CPU0:router(config-srlg)# exit

RP/0/RP0/CPU0:router(config)# router isis 1
RP/0/RP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-isis-af)# advertise application flex-algo link-attributes srlg
RP/0/RP0/CPU0:router(config-isis-af)# exit
RP/0/RP0/CPU0:router(config-isis)# flex-algo 128
RP/0/RP0/CPU0:router(config-isis-flex-algo)# advertise-definition
RP/0/RP0/CPU0:router(config-isis-flex-algo)# srlg exclude-any groupX
RP/0/RP0/CPU0:router(config-isis-flex-algo)# exit
RP/0/RP0/CPU0:router(config-isis)# flex-algo 129
RP/0/RP0/CPU0:router(config-isis-flex-algo)# advertise-definition
RP/0/RP0/CPU0:router(config-isis-flex-algo)# srlg exclude-any groupY
RP/0/RP0/CPU0:router(config-isis-flex-algo)# commit
RP/0/RP0/CPU0:router(config-isis-flex-algo)# exit
RP/0/RP0/CPU0:router(config-isis)#

```

The following example shows how to enable the Flexible Algorithm ASLA-specific advertisement of SRLGs and to exclude SRLG groups from Flexible Algorithm path computation for OSPF:

```

RP/0/RP0/CPU0:router(config)# srlg
RP/0/RP0/CPU0:router(config-srlg)# interface HunGigE0/0/0/0
RP/0/RP0/CPU0:router(config-srlg-if)# name groupX
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/0/1
RP/0/RP0/CPU0:router(config-srlg-if)# name groupX
RP/0/RP0/CPU0:router(config-srlg-if)# exit

RP/0/RP0/CPU0:router(config-srlg)# interface HunGigE0/0/1/0
RP/0/RP0/CPU0:router(config-srlg-if)# name groupY
RP/0/RP0/CPU0:router(config-srlg-if)# exit
RP/0/RP0/CPU0:router(config-srlg)# interface TenGigE0/0/1/1
RP/0/RP0/CPU0:router(config-srlg-if)# name groupY
RP/0/RP0/CPU0:router(config-srlg-if)# exit

RP/0/RP0/CPU0:router(config-srlg)# name groupX value 100
RP/0/RP0/CPU0:router(config-srlg)# name groupY value 200
RP/0/RP0/CPU0:router(config-srlg)# exit

RP/0/0/CPU0:r1(config)#router ospf 1
RP/0/0/CPU0:r1(config-ospf)#flex-algo 128
RP/0/0/CPU0:r1(config-ospf-flex-algo)#srlg exclude-any
RP/0/0/CPU0:r1(config-ospf-flex-algo-srlg-exclude-any)#groupX
RP/0/0/CPU0:r1(config-ospf-flex-algo-srlg-exclude-any)#groupY
RP/0/0/CPU0:r1(config-ospf-flex-algo-srlg-exclude-any)#commit

```

## Verification

The following example shows how to verify the number of SRLGs excluded for OSPF:

```

RP/0/RP0/CPU0:router# show ospf topology summary

```

```
Process ospf-1
Instance default
 Router ID : 192.168.0.1
 Number of Areas : 1
 Number of Algos : 1
 Max Path count : 16
 Route count : 10
 SR Global Block : 16000 - 23999

Area 0
 Number of Nodes : 6
 Algo 128
 FAD Advertising Router : 192.168.0.1
 FAD Area ID : 0
 Algo Type : 0
 Metric Type : 0
 Number of Exclude SRLGs : (2)
 [1]: 100 [2]: 200
 FAPM supported : No
```



## Flexible Algorithm with Exclude Minimum Bandwidth Constraint

Table 93: Feature History Table

| Feature Name                                                       | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IS-IS Flexible Algorithm with Exclude Minimum Bandwidth Constraint | Release 7.11.1      | <p>Traffic engineering in networks can be optimized by avoiding low-bandwidth links that may not be capable of handling high volumes of traffic.</p> <p>This feature allows you to use Flexible Algorithm to create topologies in your network that explicitly exclude high bandwidth traffic from utilizing links below a specified capacity. This constraint is achieved by introducing a new bandwidth-based metric type within the Flexible Algorithm framework. Links that do not satisfy the constraint are ignored when computing the associated Flexible Algorithm topology.</p> <p>This feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>router isis instance flex-algo algo</b> command is modified with the new <b>minimum-bandwidth value</b> option.</li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>This feature extends the native <code>Cisco-IOS-XR-clns-isis-cfg.yang</code> model (see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> |

This feature allows you to configure a minimum bandwidth value for computing a Flexible Algorithm path.

The IS-IS Flex-Algorithm Exclude Minimum Bandwidth sub-TLV (FAEMB) is a way to set a minimum bandwidth requirement for links in the Flex-Algorithm topology.

To determine if a link should be excluded based on this minimum bandwidth requirement, we compare the Minimum Bandwidth specified in the FAEMB sub-TLV with the Maximum Link Bandwidth advertised in the Area Supported by the Link Attribute (ASLA) sub-TLV.

If the Maximum Link Bandwidth is lower than the Minimum bandwidth specified, the link is excluded from the Flex-Algorithm topology. However, if the FAD includes the FAEMB sub-TLV but the Maximum Link Bandwidth is not advertised for the link, it should not be excluded based on the Minimum Bandwidth constraint.

Use the **router isis instance flex-algo algo minimum-bandwidth value** command to configure the minimum bandwidth value in kbps.

### Example: Configuring IS-IS Flexible Algorithm with Minimum Bandwidth Constraint

```
router isis 1
 address-family ipv4 unicast
 segment-routing mpls
 !
 address-family ipv6 unicast
 segment-routing srv6
 locator L1_A129
 !
 !
 !
 flex-algo 129
 advertise-definition
 minimum-bandwidth 10000000
 !
 interface Loopback0
 address-family ipv4 unicast
 prefix-sid index 100
 prefix-sid algorithm 129 index 300
 !
 address-family ipv6 unicast
 !
 !
 segment-routing
 srv6
 locators
 locator L1_A129
 micro-segment behavior unode psp-usd
 prefix cafe:0:2100::/48
 algorithm 129
 !
 !
 !
 !
```

## Flexible Algorithm with Exclude Maximum Delay Constraint

Table 94: Feature History Table

| Feature Name                                                   | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IS-IS Flexible Algorithm with Exclude Maximum Delay Constraint | Release 7.11.1      | <p>This feature enables you to configure topologies that exclude links that have delays over a specific threshold. This is especially critical for high-frequency trading applications, in satellite networks, or wherever there are fluctuations in link delays.</p> <p>This feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>The <b>router isis instance flex-algo algo</b> command is modified with the new <b>maximum-delay value</b> option.</li> </ul> <p><b>YANG Data Model:</b></p> <ul style="list-style-type: none"> <li>This feature extends the native Cisco-IOS-XR-clns-isis-cfg.yang model (see <a href="#">GitHub</a>, <a href="#">YANG Data Models Navigator</a>)</li> </ul> |

This feature allows you to configure a maximum delay value for computing a Flexible Algorithm path.

The Flexible Algorithm Exclude Minimum Delay (FAEMD) sub-TLV is used to specify the maximum delay requirement for links in a Flex-Algorithm topology. To ensure proper functioning, the FAEMD sub-TLV must appear only once in the FAD sub-TLV (Flexible Algorithm Definition). If it appears more than once, it should be ignored by the receiver. The maximum link delay advertised in the FAEMD sub-TLV is compared with the minimum unidirectional link delay advertised in the ASLA sub-TLV.

If the minimum unidirectional link delay is higher than the maximum link delay advertised in the FAEMD sub-TLV, the link must be excluded from the Flex-Algorithm topology.

However, if a link does not have the minimum unidirectional link delay advertised but the FAD contains the FAEMD sub-TLV, then based on the maximum delay constraint, that link should not be excluded from the topology.

Use the **router isis instance flex-algo algo maximum-delay delay** command to configure the maximum delay value in microseconds.

**Example: Configuring IS-IS Flexible Algorithm with Maximum Delay Constraint**

```

router isis 1
 address-family ipv4 unicast
 segment-routing mpls
 !
 address-family ipv6 unicast
 segment-routing srv6
 locator L1_A128
 !
 !
 !
flex-algo 128
 advertise-definition
 maximum-delay 300
 !
 interface Loopback0
 address-family ipv4 unicast
 prefix-sid index 100
 prefix-sid algorithm 128 index 200
 !
 address-family ipv6 unicast
 !
 !
segment-routing
 srv6
 locators
 locator L1_A128
 micro-segment behavior unode psp-usd
 prefix cafe:0:1100::/48
 algorithm 128
 !
 !
 !
 !

```

## Maximum Paths Per IS-IS Flexible Algorithm

*Table 95: Feature History Table*

| Feature Name                                | Release       | Description                                                                                                                                                                                       |
|---------------------------------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IS-IS: Maximum Paths Per Flexible Algorithm | Release 7.8.1 | This feature introduces a new subcommand under <a href="#">flex-algo</a> command. This feature allows for maximum number of Equal-Cost Multi-path (ECMP) to be set for individual Flex Algorithms |

A new subcommand under **flex-algo** is introduced.

The **flex-algo** command now includes the **address-family** <ipv4/ipv6> **unicast** subcommand, and the **maximum-paths** <maximum-paths> subcommand.



**Note** For information on IS-IS Algo0 Maximum Paths, refer to the "Implementing IS-IS" chapter in the *Routing Configuration Guide for Cisco NCS 560 Series Routers*.

The new subcommands allow for maximum number of Equal-Cost Multi-path (ECMP) to be set for individual algorithms. The value that is configured on a per-flex-algo per address-family basis overrides any value that is configured under the IS-IS global address-family submode.

### Usage Guidelines and Limitations

- The **maximum-paths** configuration is not part of the Flexible Algorithm Definitions (FAD). If the advertised definition is configured for the flexible algorithm, the **maximum-paths** will not be propagated by the IS-IS.
- The maximum-paths per algorithm takes precedence over maximum-paths per address-family.
- The maximum paths effective for each algorithm are as follows:
  - For Flex-Algo 128:
    - IPv4: 5
    - IPv6: 3
  - For Flex-Algos 129 through 255:
    - IPv4: 12
    - IPv6: 8

### Configuration Example - Max Path

This example shows how you can set the maximum paths per-Flex-Algo:

```
Router(config)# router isis 10
Router(config-isis)# flex-algo 128
Router(config-isis-flex-algo)# metric-type te
Router(config-isis-flex-algo)# address-family ipv4 unicast
Router(config-isis-flex-algo-af)# maximum-paths 5
Router(config-isis-flex-algo)# exit
Router(config-isis-flex-algo-af)# address-family ipv6 unicast
Router(config-isis-flex-algo-af)# maximum-paths 3
Router(config-isis-flex-algo-af)# exit
```

## Maximum Paths Per IS-IS Flexible Algorithm Per Prefix

Table 96: Feature History Table

| Feature Name                                          | Release        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Paths Per IS-IS Flexible Algorithm Per Prefix | Release 7.11.1 | <p>Previously, you could configure a maximum number of Equal-Cost Multi-path (ECMP) to be set for individual Flex Algorithms.</p> <p>This feature provides additional granularity to the IS-IS Maximum Paths Per-Algorithm feature by allowing you to specify a set of prefixes for Flexible Algorithm.</p> <p>Now you can achieve a balance between path diversity and computational and memory requirements by controlling the number of paths for each specific algorithm and destination prefix combination.</p> <p>This feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• <b>maximum-paths</b><br/><b>route-policy</b> <i>name</i></li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• This feature extends the native <code>Cisco-IOS-XR-clns-isis-cfg.yang</code> model</li> </ul> <p>See <a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a></p> |

Previously, you could set the maximum paths for a Flexible Algorithm per address-family.

With this feature, you can further refine the maximum paths configuration by associating it with specific prefixes for each Flexible Algorithm. The existing **maximum-paths** command is extended to include a **route-policy** qualifier to configure the maximum paths per algorithm per prefix-list.

When installing paths into the Routing Information Base (RIB) for Segment Routing with IPv6 (SRv6) or the Label Switched Database (LSD) for Segment Routing with MPLS (SR-MPLS), the system checks if a maximum paths value has been configured for the algorithm and the associated prefix. If such a configuration exists, it will be used instead of the existing address-family value to determine the number of paths to be installed.



**Note** Route policies that have the attribute **set maximum-paths number** are supported.



**Note** For information on maximum paths per prefix for IS-IS algo 0 (SPF), refer to the "Implementing IS-IS" chapter in the *Routing Configuration Guide for Cisco NCS 560 Series Routers*.

### Usage Guidelines and Limitations

- The **maximum-paths maximum-paths** and **maximum-paths route-policy name** configurations are mutually exclusive. You can configure either an unqualified number or a route-policy for any given IS-IS instance.
- The maximum paths per-algorithm per-prefix configuration takes precedence over maximum paths per-algorithm configuration. Likewise, the maximum paths per-algorithm configuration takes precedence over maximum paths per-address-family configuration. This hierarchy ensures that the most specific configuration is prioritized when determining the maximum paths for a given algorithm and prefix combination.

### Example

The following example shows how to configure the maximum paths for Flex Algo 128:

#### • Define a Prefix Set:

```
prefix-set isis-ipv4-L1
 10.1.0.101/32
end-set
```

#### • Create a Route Policy:

```
route-policy isis-mp-if-L1
 if destination in isis-ipv4-L1 then
 set maximum-paths 2
 endif
end-policy
```

#### • Apply Route Policy to Configure Maximum Paths Per-Algo Per-Prefix:

```
router isis 10
 flex-algo 128
 address-family ipv4 unicast
 maximum-paths route-policy isis-mp-if-L1
```

### Verification

Router# **show isis route flex-algo 128**

IS-IS 10 IPv4 Unicast routes Flex-Algo 128

Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)  
 df - level 1 default (closest attached router), su - summary null  
 C - connected, S - static, R - RIP, B - BGP, O - OSPF  
 E - EIGRP, A - access/subscriber, M - mobile, a - application  
 i - IS-IS (redistributed from another instance)

Maximum parallel path count: as defined in isis-mp-if-L1

```
L1 10.1.0.101/32 [121/115]
 via 15.15.15.2, GigabitEthernet0/0/0/5, hare, SRGB Base: 16000, Weight: 0
 via 16.16.16.2, GigabitEthernet0/0/0/6, hare, SRGB Base: 16000, Weight: 0
```

## Example: Configuring IS-IS Flexible Algorithm

```
router isis 1
 affinity-map red bit-position 65
 affinity-map blue bit-position 8
 affinity-map green bit-position 201

 flex-algo 128
 advertise-definition
 affinity exclude-any red
 affinity include-any blue
 !
 flex-algo 129
 affinity exclude-any green
 !
!
address-family ipv4 unicast
 segment-routing mpls
!
interface Loopback0
 address-family ipv4 unicast
 prefix-sid algorithm 128 index 100
 prefix-sid algorithm 129 index 101
!
!
interface GigabitEthernet0/0/0/0
 affinity flex-algo red
!
interface GigabitEthernet0/0/0/1
 affinity flex-algo blue red
!
interface GigabitEthernet0/0/0/2
 affinity flex-algo blue
!
```

## Example: Configuring OSPF Flexible Algorithm

```
router ospf 1
 flex-algo 130
 priority 200
 affinity exclude-any
 red
 blue
 !
 metric-type delay
!
 flex-algo 140
 affinity include-all
 green
 !
 affinity include-any
 red
```



```

!
!

interface Loopback0
 prefix-sid index 10
 prefix-sid strict-spf index 40
 prefix-sid algorithm 128 absolute 16128
 prefix-sid algorithm 129 index 129
 prefix-sid algorithm 200 index 20
 prefix-sid algorithm 210 index 30
!
!

interface GigabitEthernet0/0/0/0
 flex-algo affinity
 color red
 color blue
 !
!

affinity-map
 color red bit-position 10
 color blue bit-position 11
!

```

## User-Defined Generic Metric Support for IS-IS Flex Algo

*Table 97: Feature History Table*

| Feature Name                                            | Release Information | Feature Description                                                                            |
|---------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------|
| User-Defined Generic Metric Support for IS-IS Flex Algo | Release 24.2.11     | This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7. |

| Feature Name                                            | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User-Defined Generic Metric Support for IS-IS Flex Algo | Release 24.2.1      | <p>Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR:</p> <ul style="list-style-type: none"> <li>• N540-ACC-SYS</li> <li>• N540X-ACC-SYS</li> <li>• N540-24Z8Q2C-SYS</li> </ul> <p>This feature adds support for user-defined generic metric as a metric type for IS-IS Flexible Algorithm.</p> <p>You can now have more control over traffic flows using user-defined generic metrics. You can define a family of user-defined generic metrics that can advertise different types of administrative metrics such as jitter, reliability, and fiscal cost depending on the traffic class for Flexible Algorithms. You can selectively define and assign semantics of these metrics as per the network requirement.</p> <p>The feature introduces the following changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• The feature introduces the <b>generic-metric flex-algo</b> and <b>metric-type generic</b> commands.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• <code>Cisco-IOS-XR-um-router-isis-cfg.yang</code></li> </ul> |

### Control Traffic Flow with User-Defined Generic Metrics

With the addition of different traffic types, the need for alternate types of metrics has evolved. Flexible Algorithm already supports IGP, TE, and delay metrics. However a network operator might want to minimize their operational costs and might want a metric that reflects the actual fiscal costs of using a link. Other traffic may require low jitter, leading to an entirely different set of metrics. With Flexible Algorithm, all these different metrics could be used concurrently on the same network. These improvements are possible as you can now define a family of user-defined generic metrics that can advertise various types of administratively assigned metrics. These metrics are not predefined, which provides network administrators with the flexibility to assign

their own meanings and semantics to the metrics. This means you can create metrics tailored to specific operational goals or traffic requirements.


### Apply User-Defined Generic Metrics in Flexible Algorithm Definition

You can apply the user-defined generic metrics within the Flexible Algorithm Definition (FAD) similar to how you would use TE or delay metrics. For example, you can assign specific user-defined generic metrics to individual Flex Algos, allowing for a customized path selection criteria. This capability enables different paths to use different metrics, ensuring that each path is optimized for the specific type of traffic it handles.

### Split traffic with User-Defined Generic Metrics

Consider a scenario where for some networks, traffic engineering might require splitting east and west traffic for different Flexible Algorithms across multiple hardware data planes. In the following figure, the Flex Algo 128 (Orange) uses the upper primary path, while the Flex Algo 129 (Blue) uses the lower primary path. You can achieve this by defining specific user-defined generic metrics for the Flex Algo to determine the best path for each Flex Algo. The numbers 10 and 100 represent the user-defined generic metrics configured for the two different Flex Algos.

**Figure 45: Splitting traffic using user-defined generic metrics for different Flex Algos**



### Advantages of splitting traffic using user-defined generic metrics

This approach offers several advantages. You can use Equal-Cost Multi-Path (ECMP) routing within the backbone, allowing multiple links of the same color between core nodes. It also ensures efficient backup routes and prevents U-loops with Flex Algo specific Loop-Free Alternates (LFA) and Topology-Independent Loop-Free Alternates (TI-LFA).

In the above example, the vertical core links are used as backup paths for all Flex Algos, and you can activate them only if a core link or node fails. These links are part of all Flex Algos, and you can use user-defined generic metrics to configure and advertise such metrics for each individual Flex Algo to ensure smooth operation.

### Benefits of User-Defined Generic Metrics

The key benefits of user-defined generic metrics are:

- You can ensure precise control over routing decisions based on the assigned metrics, as the links that do not advertise the user-defined generic metric are excluded from the Flex Algo topology.
- You can use the metrics to reduce operational expenses by choosing cost-effective paths.
- You can customize metrics to fit unique operational goals and traffic requirements, providing tailored solutions for different scenarios.
- You can simultaneously apply multiple metrics across the same network, enhancing overall network performance and reliability.

## Usage Guidelines and Limitations for User-Defined Generic Metrics

The user-defined generic metric is disabled by default. In other words, the metric is not advertised unless it is configured.

## Configure User-Defined Generic Metrics

This section includes configuration for user-defined generic metrics for IS-IS Flex Algos.

### Procedure

- Step 1** Define the metric type and value for an interface using the **generic-metric flex-algo** command in the IS-IS interface address-family submode. In the following example, the user-defined generic metric type is 177 and the metric value as 100.

**Example:**

```
Router(config)#router isis 1
Router(config-isis)#interface GigabitEthernet 0/2/0/7
Router(config-isis-if)#address-family ipv4 unicast
Router(config-isis-if-af)#generic-metric flex-algo type 128 100
Router(config-isis-if-af)#generic-metric flex-algo type 177 100
Router(config-isis-if-af)#generic-metric flex-algo type 188 1000
```

- Step 2** Associate or advertise the configured metric type to a Flexible Algorithm Definition using the **metric-type generic** command. In the following example, the metric type advertised is 177.

```
Router(config)#router isis 1
Router(config-isis)#flex-algo 128
Router(config-isis-flex-algo)#priority 254
Router(config-isis-flex-algo)#metric-type generic 177
Router(config-isis-flex-algo)#advertise-definition
```

- Step 3** Verify the running configuration using the **show running-config** command.

**Example:**

```
router isis 1
flex-algo 128
priority 254
metric-type generic 177
advertise-definition
```

```

!
interface GigabitEthernet 0/2/0/7
 address-family ipv4 unicast
 generic-metric flex-algo type 128 100
 generic-metric flex-algo type 177 100
 generic-metric flex-algo type 188 1000
!
!
!

```

**Step 4** Verify the Flexible Algorithm Definition configuration using the **show isis flex-algo** command.

**Example:**

```

Router#show isis flex-algo 128
Thu Dec 7 03:10:56.452 PST
IS-IS 1 Flex-Algo Database
Flex-Algo 128:
Level-2:
Definition Priority: 254
Definition Source: plzen.00, (Local)
Definition Equal to Local: Yes
Definition Metric Type: User-defined: 177.
Definition Flex-Algo Prefix Metric: No
Exclude Any Affinity Bit Positions:
Include Any Affinity Bit Positions:
Include All Affinity Bit Positions:
Reverse Exclude Any Affinity Bit Positions:
Reverse Include Any Affinity Bit Positions:
Reverse Include All Affinity Bit Positions:
Exclude SRLGs:
Minimum Link Bandwidth: 0 kbits/s
Maximum Link Delay: 0 us
Disabled: No
Topologies supported:
IPv4 Unicast
Local Priority: 254
FRR Disabled: No
Microloop Avoidance Disabled: No
UCMP Disabled: No
Data Plane Segment Routing: Yes
Data Plane IP: No

```

**Step 5** Optionally, you can also verify the application-specific user-defined generic metric configured for an interface by using the **show isis interface** command.

**Example:**

```

Router#show isis interface GigabitEthernet 0/2/0/7
Thu Dec 7 03:13:34.140 PST
GigabitEthernet0/2/0/7 Enabled
Adjacency Formation: Enabled
Prefix Advertisement: Enabled
Bandwidth: 1000000
...
...
IPv4 Unicast Topology: Enabled
Adjacency Formation: Running
Prefix Advertisement: Running
Policy (L1/L2): -/-
Metric (L1/L2): 10/10
Metric fallback:
Bandwidth (L1/L2): Inactive/Inactive
Anomaly (L1/L2): Inactive/Inactive
Weight (L1/L2): 0/0

```

```

L1 Flex-algo Generic-metrics:
 Type: 128 100
 Type: 177 100
 Type: 188 1000
L2 Flex-algo Generic-metrics:
 Type: 128 100
 Type: 177 100
 Type: 188 1000

```

---

## Example: Traffic Steering to Flexible Algorithm Paths

### BGP Routes on PE – Color Based Steering

SR-TE On Demand Next-Hop (ODN) feature can be used to steer the BGP traffic towards the Flexible Algorithm paths.

The following example configuration shows how to setup BGP steering local policy, assuming two router: R1 (2.2.2.2) and R2 (4.4.4.4), in the topology.

#### Configuration on router R1:

```

vrf Test
address-family ipv4 unicast
 import route-target
 1:150
 !
 export route-policy SET_COLOR_RED_HI_BW
 export route-target
 1:150
 !
 !
 !
interface Loopback0
ipv4 address 2.2.2.2 255.255.255.255
!
interface Loopback150
vrf Test
ipv4 address 2.2.2.222 255.255.255.255
!
interface TenGigE0/1/0/3/0
description exr1 to cxr1
ipv4 address 10.0.20.2 255.255.255.0
!
extcommunity-set opaque color129-red-igp
 129
end-set
!
route-policy PASS
 pass
end-policy
!
route-policy SET_COLOR_RED_HI_BW
 set extcommunity color color129-red-igp
 pass
end-policy
!
router isis 1

```

```

is-type level-2-only
net 49.0001.0000.0000.0002.00
log adjacency changes
affinity-map RED bit-position 28
flex-algo 128
 priority 228
!
address-family ipv4 unicast
 metric-style wide
 advertise link attributes
 router-id 2.2.2.2
 segment-routing mpls
!
interface Loopback0
 address-family ipv4 unicast
 prefix-sid index 2
 prefix-sid algorithm 128 index 282
 !
!
interface TenGigE0/1/0/3/0
 point-to-point
 address-family ipv4 unicast
 !
!
router bgp 65000
 bgp router-id 2.2.2.2
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 retain route-target all
 !
 neighbor-group RR-services-group
 remote-as 65000
 update-source Loopback0
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
 !
 neighbor 4.4.4.4
 use neighbor-group RR-services-group
 !
vrf Test
 rd auto
 address-family ipv4 unicast
 redistribute connected
 !
segment-routing
traffic-eng
 logging
 policy status
 !
 segment-list sl-cxr1
 index 10 mpls label 16294
 !
 policy pol-foo
 color 129 end-point ipv4 4.4.4.4
 candidate-paths
 preference 100
 explicit segment-list sl-cxr1
 !
 !
 !

```

```

!
!
!

```

### Configuration on router R2:

```

vrf Test
address-family ipv4 unicast
 import route-target
 1:150
 !
 export route-policy SET_COLOR_RED_HI_BW
 export route-target
 1:150
 !
!
!
interface TenGigE0/1/0/1
description cxr1 to exr1
ipv4 address 10.0.20.1 255.255.255.0
!
extcommunity-set opaque color129-red-igp
 129
end-set
!
route-policy PASS
 pass
end-policy
!
route-policy SET_COLOR_RED_HI_BW
 set extcommunity color color129-red-igp
 pass
end-policy
!
router isis 1
is-type level-2-only
net 49.0001.0000.0000.0004.00
log adjacency changes
affinity-map RED bit-position 28
affinity-map BLUE bit-position 29
affinity-map GREEN bit-position 30
flex-algo 128
 priority 228
!
flex-algo 129
 priority 229
!
flex-algo 130
 priority 230
!
address-family ipv4 unicast
 metric-style wide
 advertise link attributes
 router-id 4.4.4.4
 segment-routing mpls
!
interface Loopback0
 address-family ipv4 unicast
 prefix-sid index 4
 prefix-sid algorithm 128 index 284
 prefix-sid algorithm 129 index 294
 prefix-sid algorithm 130 index 304
 !
!

```



```

interface GigabitEthernet0/0/0/0
 point-to-point
 address-family ipv4 unicast
 !
!
interface TenGigE0/1/0/1
 point-to-point
 address-family ipv4 unicast
 !
!
router bgp 65000
 bgp router-id 4.4.4.4
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
 neighbor-group RR-services-group
 remote-as 65000
 update-source Loopback0
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
 !
 neighbor 10.1.1.1
 use neighbor-group RR-services-group
 !
 neighbor 2.2.2.2
 use neighbor-group RR-services-group
 !
vrf Test
 rd auto
 address-family ipv4 unicast
 redistribute connected
 !
 neighbor 25.1.1.2
 remote-as 4
 address-family ipv4 unicast
 route-policy PASS in
 route-policy PASS out
 !
!
!
segment-routing
!
end

```

## Delay Normalization

**Table 98: Feature History Table**

| Feature Name                       | Release Information | Feature Description                                                           |
|------------------------------------|---------------------|-------------------------------------------------------------------------------|
| SR-TE Delay Normalization for OSPF | Release 7.3.1       | This feature extends the current Delay Normalization feature to support OSPF. |

Performance measurement (PM) measures various link characteristics like packet loss and delay. Such characteristics can be used by IS-IS as a metric for Flexible Algorithm computation. Low latency routing using dynamic delay measurement is one of the primary use cases for Flexible Algorithm technology.

Delay is measured in microseconds. If delay values are taken as measured and used as link metrics during the IS-IS topology computation, some valid ECMP paths might be unused because of the negligible difference in the link delay.

The Delay Normalization feature computes a normalized delay value and uses the normalized value instead. This value is advertised and used as a metric during the Flexible Algorithm computation.

The normalization is performed when the delay is received from the delay measurement component. When the next value is received, it is normalized and compared to the previous saved normalized value. If the values are different, then the LSP generation is triggered.


The following formula is used to calculate the normalized value:

- **Dm** – measured Delay
- **Int** – configured normalized Interval
- **Off** – configured normalized Offset (must be less than the normalized interval Int)
- **Dn** – normalized Delay
- **a** =  $Dm / Int$  (rounded down)
- **b** =  $a * Int + Off$

If the measured delay (Dm) is less than or equal to **b**, then the normalized delay (Dn) is equal to **b**. Otherwise, Dn is **b + Int**.

### Example

The following example shows a low-latency service. The intent is to avoid high-latency links (1-6, 5-2). Links 1-2 and 5-6 are both low-latency links. The measured latency is not equal, but the difference is insignificant.



We can normalize the measured latency before it is advertised and used by IS-IS. Consider a scenario with the following:


- Interval = 10
- Offset = 3

The measured delays will be normalized as follows:

- **Dm** = 29
- **a** =  $29 / 10 = 2$  (2.9, rounded down to 2)

$$b = 2 * 10 + 3 = 23$$

In this case, **Dm** (29) is greater than **b** (23); so **Dn** is equal to **b+I** (23 + 10) = 33




- $D_m = 31$


$$a = 31 / 10 = 3 \text{ (3.1, rounded down to 3)}$$

$$b = 3 * 10 + 3 = 33$$

In this case, **Dm** (31) is less than **b** (33); so **Dn** is **b** = 33



The link delay between 1-2 and 5-6 is normalized to 33.



### Configuration

Delay normalization is disabled by default. To enable and configure delay normalization, use the **delay normalize interval** *interval* [*offset* *offset*] command.

- *interval* – The value of the normalize interval in microseconds.

- *offset* – The value of the normalized offset in microseconds. This value must be smaller than the value of normalized interval.

### IS-IS Configuration

```
router isis 1
 interface GigEth 0/0/0/0
 delay normalize interval 10 offset 3
 address-family ipv4 unicast
 metric 77
```

### OSPF Configuration

```
router ospf 1
 area 0
 interface GigabitEthernet0/0/0/0
 delay normalize interval 10 offset 3
 !
 !
 !
```





## CHAPTER 16

# Using Segment Routing OAM

Segment Routing Operations, Administration, and Maintenance (OAM) helps service providers to monitor label-switched paths (LSPs) and quickly isolate forwarding problems to assist with fault detection and troubleshooting in the network. The Segment Routing OAM feature provides support for BGP prefix SIDs, IGP prefix SIDs, and Nil-FEC (forwarding equivalence classes) LSP Ping and Traceroute functionality.

- [MPLS Ping and Traceroute for BGP and IGP Prefix-SID, on page 613](#)
- [Examples: MPLS Ping, Traceroute, and Tree Trace for Prefix-SID, on page 614](#)
- [MPLS LSP Ping and Traceroute Nil FEC Target, on page 616](#)
- [Examples: LSP Ping and Traceroute for Nil\\_FEC Target, on page 616](#)
- [Segment Routing Ping and Traceroute, on page 618](#)
- [Segment Routing Ping and Traceroute for Flexible Algorithm, on page 623](#)
- [Segment Routing over IPv6 OAM, on page 624](#)
- [Data Plane Validation Support for SR-MPLS IPv6-based LSPs, on page 626](#)
- [MPLS OAM support for SR-TE Policies using MPLS IPv6-based LSPs, on page 629](#)

## MPLS Ping and Traceroute for BGP and IGP Prefix-SID

MPLS Ping and Traceroute operations for Prefix SID are supported for various BGP and IGP scenarios, for example:

- Within an IS-IS level or OSPF area
- Across IS-IS levels or OSPF areas
- Route redistribution from IS-IS to OSPF and from OSPF to IS-IS
- Anycast Prefix SID
- Combinations of BGP and LDP signaled LSPs

The MPLS LSP Ping feature is used to check the connectivity between ingress Label Switch Routers (LSRs) and egress LSRs along an LSP. MPLS LSP ping uses MPLS echo request and reply messages, similar to Internet Control Message Protocol (ICMP) echo request and reply messages, to validate an LSP. The destination IP address of the MPLS echo request packet is different from the address used to select the label stack. The destination IP address is defined as a 127.x.y.z/8 address and it prevents the IP packet from being IP switched to its destination, if the LSP is broken.


The MPLS LSP Traceroute feature is used to isolate the failure point of an LSP. It is used for hop-by-hop fault localization and path tracing. The MPLS LSP Traceroute feature relies on the expiration of the Time to Live (TTL) value of the packet that carries the echo request. When the MPLS echo request message hits a transit node, it checks the TTL value and if it is expired, the packet is passed to the control plane, else the message is forwarded. If the echo message is passed to the control plane, a reply message is generated based on the contents of the request message.

The MPLS LSP Tree Trace (traceroute multipath) operation is also supported for BGP and IGP Prefix SID. MPLS LSP Tree Trace provides the means to discover all possible equal-cost multipath (ECMP) routing paths of an LSP to reach a destination Prefix SID. It uses multipath data encoded in echo request packets to query for the load-balancing information that may allow the originator to exercise each ECMP. When the packet TTL expires at the responding node, the node returns the list of downstream paths, as well as the multipath information that can lead the operator to exercise each path in the MPLS echo reply. This operation is performed repeatedly for each hop of each path with increasing TTL values until all ECMP are discovered and validated.

MPLS echo request packets carry Target FEC Stack sub-TLVs. The Target FEC sub-TLVs are used by the responder for FEC validation. The BGP and IGP IPv4 prefix sub-TLV has been added to the Target FEC Stack sub-TLV. The IGP IPv4 prefix sub-TLV contains the prefix SID, the prefix length, and the protocol (IS-IS or OSPF). The BGP IPv4 prefix sub-TLV contains the prefix SID and the prefix length.

## Examples: MPLS Ping, Traceroute, and Tree Trace for Prefix-SID

These examples use the following topology:



### MPLS Ping for Prefix-SID

```
RP/0/RP0/CPU0:router-arizona# ping mpls ipv4 10.1.1.4/32
Thu Dec 17 01:01:42.301 PST
```

```
Sending 5, 100-byte MPLS Echos to 10.1.1.4,
 timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
```

**MPLS Traceroute for Prefix-SID**

```
RP/0/RP0/CPU0:router-arizona# traceroute mpls ipv4 10.1.1.4/32
Thu Dec 17 14:45:05.563 PST
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 12.12.12.1 MRU 4470 [Labels: 16004 Exp: 0]
L 1 12.12.12.2 MRU 4470 [Labels: 16004 Exp: 0] 3 ms
L 2 23.23.23.3 MRU 4470 [Labels: implicit-null Exp: 0] 3 ms
! 3 34.34.34.4 11 ms
```

**MPLS Tree Trace for Prefix-SID**

```
RP/0/RP0/CPU0:router-arizona# traceroute mpls multipath ipv4 10.1.1.4/32
Thu Dec 17 14:55:46.549 PST
```

Starting LSP Path Discovery for 10.1.1.4/32

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
LL!
Path 0 found,
 output interface TenGigE0/0/0/0 nexthop 12.12.12.2 source 12.12.12.1 destination 127.0.0.0
 L!
Path 1 found,
 output interface TenGigE0/0/0/0 nexthop 12.12.12.2 source 12.12.12.1 destination 127.0.0.2
 LL!
Path 2 found,
 output interface TenGigE0/0/0/1 nexthop 15.15.15.5 source 15.15.15.1 destination 127.0.0.1
 L!
Path 3 found,
 output interface TenGigE0/0/0/1 nexthop 15.15.15.5 source 15.15.15.1 destination 127.0.0.0

Paths (found/broken/unexplored) (4/0/0)
Echo Request (sent/fail) (10/0)
Echo Reply (received/timeout) (10/0)
Total Time Elapsed 53 ms
```



## MPLS LSP Ping and Traceroute Nil FEC Target

The Nil-FEC LSP ping and traceroute operations are extensions of regular MPLS ping and traceroute.

Nil-FEC LSP Ping/Traceroute functionality supports segment routing and MPLS Static. It also acts as an additional diagnostic tool for all other LSP types. This feature allows operators to provide the ability to freely test any label stack by allowing them to specify the following:

- label stack
- outgoing interface
- nexthop address

In the case of segment routing, each segment nodal label and adjacency label along the routing path is put into the label stack of an echo request message from the initiator Label Switch Router (LSR); MPLS data plane forwards this packet to the label stack target, and the label stack target sends the echo message back.

The following table shows the syntax for the ping and traceroute commands.

**Table 99: LSP Ping and Traceroute Nil FEC Commands**

| Command Syntax                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ping mpls nil-fec labels</b> {label[,label]} [ <b>output</b> { <b>interface</b> tx-interface} [ <b>nexthop</b> nexthop-ip-addr]]       |
| <b>traceroute mpls nil-fec labels</b> {label[,label]} [ <b>output</b> { <b>interface</b> tx-interface} [ <b>nexthop</b> nexthop-ip-addr]] |

## Examples: LSP Ping and Traceroute for Nil\_FEC Target

These examples use the following topology:

```

Node loopback IP address: 172.18.1.3 172.18.1.4 172.18.1.5 172.18.1.7
Node label: 16004 16005 16007
Nodes: Arizona ---- Utah ----- Wyoming ---- Texas

Interface: GigabitEthernet0/0/0/1 GigabitEthernet0/0/0/1
Interface IP address: 10.1.1.3 10.1.1.4

```

```
RP/0/RP0/CPU0:router-utah# show mpls forwarding
```

```

Tue Jul 5 13:44:31.999 EDT
Local Outgoing Prefix Outgoing Next Hop Bytes
Label Label or ID Interface Hop Switched

16004 Pop No ID Gi0/0/0/1 10.1.1.4 1392
 Pop No ID Gi0/0/0/2 10.1.2.2 0
16005 16005 No ID Gi0/0/0/0 10.1.1.4 0
 16005 No ID Gi0/0/0/1 10.1.2.2 0
16007 16007 No ID Gi0/0/0/0 10.1.1.4 4752
 16007 No ID Gi0/0/0/1 10.1.2.2 0
24000 Pop SR Adj (idx 0) Gi0/0/0/0 10.1.1.4 0
24001 Pop SR Adj (idx 2) Gi0/0/0/0 10.1.1.4 0

```

|       |     |                |           |             |   |
|-------|-----|----------------|-----------|-------------|---|
| 24002 | Pop | SR Adj (idx 0) | Gi0/0/0/1 | 10.1.2.2    | 0 |
| 24003 | Pop | SR Adj (idx 2) | Gi0/0/0/1 | 10.1.2.2    | 0 |
| 24004 | Pop | No ID          | tt10      | point2point | 0 |
| 24005 | Pop | No ID          | tt11      | point2point | 0 |
| 24006 | Pop | No ID          | tt12      | point2point | 0 |
| 24007 | Pop | No ID          | tt13      | point2point | 0 |
| 24008 | Pop | No ID          | tt30      | point2point | 0 |

## Ping Nil FEC Target

```
RP/0/RP0/CPU0:router-arizona# ping mpls nil-fec labels 16005,16007 output interface
GigabitEthernet 0/0/0/1 nexthop 10.1.1.4 repeat 1
```

```
Sending 1, 72-byte MPLS Echos with Nil FEC labels 16005,16007,
 timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'd' - see DDMAP for return code,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
!
```

```
Success rate is 100 percent (1/1), round-trip min/avg/max = 1/1/1 ms
Total Time Elapsed 0 ms
```

## Traceroute Nil FEC Target

```
RP/0/RP0/CPU0:router-arizona# traceroute mpls nil-fec labels 16005,16007 output interface
GigabitEthernet 0/0/0/1 nexthop 10.1.1.4
```

```
Tracing MPLS Label Switched Path with Nil FEC labels 16005,16007, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'd' - see DDMAP for return code,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
0 10.1.1.3 MRU 1500 [Labels: 16005/16007/explicit-null Exp: 0/0/0]
L 1 10.1.1.4 MRU 1500 [Labels: implicit-null/16007/explicit-null Exp: 0/0/0] 1 ms
L 2 10.1.1.5 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 1 ms
! 3 10.1.1.7 1 ms
```

# Segment Routing Ping and Traceroute

Table 100: Feature History Table

| Feature Name                                                          | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SR OAM for SR Policy (Policy Name / Binding SID / Custom label stack) | Release 7.3.1       | This feature extends SR OAM ping and traceroute function for an SR policy (or binding SID)-LSP end-point combination.<br><br>This addresses the limitations of the Nil-FEC LSP Ping and Traceroute function which cannot perform a ping operation to a segment list that is not associated with an installed SR policy. Also, it cannot validate egress device-specific SR policies. |

## Segment Routing Ping

The MPLS LSP ping feature is used to check the connectivity between ingress and egress of LSP. MPLS LSP ping uses MPLS echo request and reply messages, similar to Internet Control Message Protocol (ICMP) echo request and reply messages, to validate an LSP. Segment routing ping is an extension of the MPLS LSP ping to perform the connectivity verification on the segment routing control plane.



**Note** Segment routing ping can only be used when the originating device is running segment routing.

You can initiate the segment routing ping operation only when Segment Routing control plane is available at the originator, even if it is not preferred. This allows you to validate the SR path before directing traffic over the path. Segment Routing ping can use either generic FEC type or SR control-plane FEC type (SR-OSPF, SR-ISIS). In mixed networks, where some devices are running MPLS control plane (for example, LDP) or do not understand SR FEC, generic FEC type allows the device to successfully process and respond to the echo request. By default, generic FEC type is used in the target FEC stack of segment routing ping echo request. Generic FEC is not coupled to a particular control plane; it allows path verification when the advertising protocol is unknown or might change during the path of the echo request. If you need to specify the target FEC, you can select the FEC type as OSPF, IS-IS, or BGP. This ensures that only devices that are running segment routing control plane, and can therefore understand the segment routing IGP FEC, respond to the echo request.

### Configuration Examples

These examples show how to use segment routing ping to test the connectivity of a segment routing control plane. In the first example, FEC type is not specified. You can also specify the FEC type as shown in the other examples.

```
RP/0/RP0/CPU0:router# ping sr-mpls 10.1.1.2/32
```

```
Sending 5, 100-byte MPLS Echos to 10.1.1.2/32,
timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
```

'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,  
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,  
 'P' - no rx intf label prot, 'p' - premature termination of LSP,  
 'R' - transit router, 'I' - unknown upstream index,  
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!!!  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/5 ms

RP/0/RP0/CPU0:router# **ping sr-mpls 10.1.1.2/32 fec-type generic**

Sending 5, 100-byte MPLS Echos to 10.1.1.2/32,  
 timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,  
 'L' - labeled output interface, 'B' - unlabeled output interface,  
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,  
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,  
 'P' - no rx intf label prot, 'p' - premature termination of LSP,  
 'R' - transit router, 'I' - unknown upstream index,  
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!!!  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

RP/0/RP0/CPU0:router# **ping sr-mpls 10.1.1.2/32 fec-type igp ospf**

Sending 5, 100-byte MPLS Echos to 10.1.1.2/32,  
 timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,  
 'L' - labeled output interface, 'B' - unlabeled output interface,  
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,  
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,  
 'P' - no rx intf label prot, 'p' - premature termination of LSP,  
 'R' - transit router, 'I' - unknown upstream index,  
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!!!  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

RP/0/RP0/CPU0:router# **ping sr-mpls 10.1.1.2/32 fec-type igp isis**

Sending 5, 100-byte MPLS Echos to 10.1.1.2/32,  
 timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,  
 'L' - labeled output interface, 'B' - unlabeled output interface,  
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,  
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,  
 'P' - no rx intf label prot, 'p' - premature termination of LSP,  
 'R' - transit router, 'I' - unknown upstream index,  
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!!!

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

```
RP/0/RP0/CPU0:router# ping sr-mpls 10.1.1.2/32 fec-type bgp
```

```
Sending 5, 100-byte MPLS Echos to 10.1.1.2/32,
 timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

### Ping for SR Policy

You can perform the ping operation for an SR policy (or binding SID), and LSP end-point combination. Use the **ping** command's **policy name lsp-end-point** and **policy binding-sid lsp-end-point** options to perform this task. You can instantiate the policy through the CLI, Netconf, PCEP or BGP-TE process.

IPv6 policies are not supported for SR OAM function.




---

**Note** As a prerequisite, you must enable the MPLS OAM function.

---

```
Router(config)# mpls oam
Router(config)# commit
```

```
Router# ping sr-mpls policy name srte_c_4_ep_10.0.0.1 lsp-end-point 209.165.201.1
Router# ping sr-mpls policy binding-sid 1000 lsp-end-point 209.165.201.1
```

## Segment Routing Traceroute

The MPLS LSP traceroute is used to isolate the failure point of an LSP. It is used for hop-by-hop fault localization and path tracing. The MPLS LSP traceroute feature relies on the expiration of the Time to Live (TTL) value of the packet that carries the echo request. When the MPLS echo request message hits a transit node, it checks the TTL value and if it is expired, the packet is passed to the control plane, else the message is forwarded. If the echo message is passed to the control plane, a reply message is generated based on the contents of the request message. Segment routing traceroute feature extends the MPLS LSP traceroute functionality to segment routing networks.

Similar to segment routing ping, you can initiate the segment routing traceroute operation only when Segment Routing control plane is available at the originator, even if it is not preferred. Segment Routing traceroute can use either generic FEC type or SR control-plane FEC type (SR-OSPF, SR-ISIS). By default, generic FEC type is used in the target FEC stack of segment routing traceroute echo request. If you need to specify the target FEC, you can select the FEC type as OSPF, IS-IS, or BGP. This ensures that only devices that are running segment routing control plane, and can therefore understand the segment routing IGP FEC, respond to the echo request.

The existence of load balancing at routers in an MPLS network provides alternate paths for carrying MPLS traffic to a target router. The multipath segment routing traceroute feature provides a means to discover all possible paths of an LSP between the ingress and egress routers.

### Configuration Examples

These examples show how to use segment routing traceroute to trace the LSP for a specified IPv4 prefix SID address. In the first example, FEC type is not specified. You can also specify the FEC type as shown in the other examples.

```
RP/0/RP0/CPU0:router# traceroute sr-mpls 10.1.1.2/32
```

```
Tracing MPLS Label Switched Path to 10.1.1.2/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 10.12.12.1 MRU 1500 [Labels: implicit-null Exp: 0]
! 1 10.12.12.2 3 ms
```

```
RP/0/RP0/CPU0:router# traceroute sr-mpls 10.1.1.2/32 fec-type generic
```

```
Tracing MPLS Label Switched Path to 10.1.1.2/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 10.12.12.1 MRU 1500 [Labels: implicit-null Exp: 0]
! 1 10.12.12.2 2 ms
```

```
RP/0/RP0/CPU0:router# traceroute sr-mpls 10.1.1.2/32 fec-type igp ospf
```

```
Tracing MPLS Label Switched Path to 10.1.1.2/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 10.12.12.1 MRU 1500 [Labels: implicit-null Exp: 0]
! 1 10.12.12.2 2 ms
```

```
RP/0/RP0/CPU0:router# traceroute sr-mpls 10.1.1.2/32 fec-type igp isis
```

```
Tracing MPLS Label Switched Path to 10.1.1.2/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
 0 10.12.12.1 MRU 1500 [Labels: implicit-null Exp: 0]
! 1 10.12.12.2 2 ms
```

```
RP/0/RP0/CPU0:router#traceroute sr-mpls 10.1.1.2/32 fec-type bgp
```

```
Tracing MPLS Label Switched Path to 10.1.1.2/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
 0 10.12.12.1 MRU 1500 [Labels: implicit-null/implicit-null Exp: 0/0]
! 1 10.12.12.2 2 ms
```

This example shows how to use multipath traceroute to discover all the possible paths for a IPv4 prefix SID.

```
RP/0/RP0/CPU0:router# traceroute sr-mpls multipath 10.1.1.2/32
```

```
Starting LSP Path Discovery for 10.1.1.2/32
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
!
Path 0 found,
 output interface GigabitEthernet0/0/0/2 nexthop 10.13.13.2
source 10.13.13.1 destination 127.0.0.0
!
Path 1 found,
 output interface Bundle-Ether1 nexthop 10.12.12.2
source 10.12.12.1 destination 127.0.0.0
```

```
Paths (found/broken/unexplored) (2/0/0)
Echo Request (sent/fail) (2/0)
```

```
Echo Reply (received/timeout) (2/0)
Total Time Elapsed 14 ms
```

### Traceroute for SR Policy

You can perform the traceroute operation for an SR policy (or binding SID), and LSP end-point combination. Use the **traceroute** command's **policy name lsp-end-point** and **policy binding-sid lsp-end-point** options to perform this task. You can instantiate the policy through the CLI, Netconf, PCEP or BGP-TE process.

IPv6 policies are not supported for SR OAM function.



**Note** As a prerequisite, you must enable the MPLS OAM function.

```
Router(config)# mpls oam
Router(config)# commit
```

```
Router# traceroute sr-mpls policy name srte_c_4_ep_10.0.0.1 lsp-end-point 209.165.201.1
Router# traceroute sr-mpls policy binding-sid 1000 lsp-end-point 209.165.201.1
```

## Segment Routing Ping and Traceroute for Flexible Algorithm

Flexible Algorithm validation method is based on segment identifier (SID) label and label switched path (LSP) destination, instead of being based on IP address. The assigner is validated against the topology prefix information provided by SR-PCE database. If the assigner is valid, then the label given is also validated against the SR-PCE database. On the egress side, the destination label is contained in a new SR Label sub-TLV. This label is verified against a SID list provided by SR-PCE.



**Note** Observe the following guidelines and restrictions:

- All routers within an area must share the same Flexible Algorithm definition for a Flexible Algorithm to be valid.
- All routers within the domain must be configured with the same SRGB range of values.
- BGP-LS must be enabled.
- Only prefix SIDs and Flexible Algorithm SIDs are supported.
- Only single label stack is supported.

## Segment Routing Ping for Flexible Algorithm

```
Router# ping sr-mpls labels 16131 lsp-end-point 10.1.1.5
Fri Dec 13 19:26:29.517 IST
```

```
Sending 5, 100-byte MPLS Echos with SR Label FEC with lsp end point 10.1.1.5, SID Label(s)
[16131],
```

```
timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
```



```
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/6 ms

## Segment Routing Traceroute for Flexible Algorithm

```
Router# traceroute sr-mpls labels 16130 lsp-end-point 10.1.1.5
Fri Dec 13 19:26:59.368 IST
```

Tracing MPLS Label Switched Path to SR Label FEC with lsp end point 10.1.1.5, SID Label(s) [16130], timeout is 2 seconds

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 13.13.13.1 MRU 1500 [Labels: 16130 Exp: 0]
L 1 13.13.13.3 MRU 1500 [Labels: 16130 Exp: 0] 5 ms
L 2 16.16.16.4 MRU 1500 [Labels: implicit-null Exp: 0] 4 ms
! 3 18.18.18.5 4 ms
```

## Segment Routing over IPv6 OAM

Segment Routing over IPv6 data plane (SRv6) implementation adds a new type of routing extension header. Hence, the existing ICMPv6 mechanisms including ping and traceroute can be used in the SRv6 network. There is no change in the way ping and traceroute operations work for IPv6- or SRv6-capable nodes in an SRv6 network.

### Restrictions and Usage Guidelines

The following restriction applies for SRv6 OAM:

- Ping to an SRv6 SID is not supported.

### Examples: SRv6 OAM

The following example shows using ping in an SRv6 network.

```
RP/0/RP0/CPU0:Router# ping ipv6 2001::33:33:33:33
Mon Sep 17 20:04:10.068 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::33:33:33:33, timeout is 2 seconds:
```

```
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

The following example shows using traceroute in an SRv6 network.

```
RP/0/RP0/CPU0:Router# traceroute ipv6 2001::33:33:33:33 probe 1 timeout 0 srv6
Fri Sep 14 15:59:25.170 UTC
Type escape sequence to abort.
Tracing the route to 2001::33:33:33:33
 1 2001::22:22:22:22[IP tunnel: DA=cafe:0:0:a4:1::: SRH =(2001::33:33:33:33 ,SL=1)] 2
msec
 2 2001::2:2:2:2[IP tunnel: DA=cafe:0:0:a4:1::: SRH =(2001::33:33:33:33 ,SL=1)] 2 msec
 3 2001::44:44:44:44 2 msec
 4 2001::33:33:33:33 3 msec
```

The following example shows using traceroute in an SRv6 network without an SRH.

```
RP/0/RSP1/CPU0:Router# traceroute ipv6 2001::44:44:44:44 srv6
Wed Jan 16 14:35:27.511 UTC
Type escape sequence to abort.
Tracing the route to 2001::44:44:44:44
 1 2001::2:2:2:2 3 msec 2 msec 2 msec
 2 2001::44:44:44:44 3 msec 3 msec 3 msec
```

The following example shows using ping for a specified IP address in the VRF.

```
RP/0/RP0/CPU0:Router# ping 10.15.15.1 vrf red
Mon Sep 17 20:07:10.085 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.15.15.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

The following example shows using traceroute for a specified IP address in the VRF.

```
RP/0/RP0/CPU0:Router# traceroute 10.15.15.1 vrf red
Mon Sep 17 20:07:18.478 UTC

Type escape sequence to abort.
Tracing the route to 10.15.15.1
 1 10.15.15.1 3 msec 2 msec 2 msec
```

The following example shows using traceroute for CE1 (4.4.4.5) to CE2 (5.5.5.5) in the VRF:

```
RP/0/RP0/CPU0:Router# traceroute 5.5.5.5 vrf a
Wed Jan 16 15:08:46.264 UTC

Type escape sequence to abort.
Tracing the route to 5.5.5.5
 1 14.14.14.1 5 msec 1 msec 1 msec
 2 15.15.15.1 3 msec 2 msec 2 msec
 3 15.15.15.2 2 msec * 3 msec
```

# Data Plane Validation Support for SR-MPLS IPv6-based LSPs

*Table 101: Feature History Table*

| Feature Name                                                              | Release Information | Feature Description                                                                            |
|---------------------------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------|
| Data Plane Validation for SR-MPLS IPv6-based Controller Instantiated LSPs | Release 24.2.11     | This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7. |

| Feature Name                                                              | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Plane Validation for SR-MPLS IPv6-based Controller Instantiated LSPs | Release 24.2.1      | <p>Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR:</p> <ul style="list-style-type: none"> <li>• N540-ACC-SYS</li> <li>• N540X-ACC-SYS</li> <li>• N540-24Z8Q2C-SYS</li> </ul> <p>You can now verify the network configuration and paths and policies set up, without interrupting or potentially disrupting live network traffic, for SR-MPLS (Segment Routing over Multiprotocol Label Switching) IPv6-based Label Switched Paths (LSPs). With this feature, you can validate controller instantiated LSPs programmed directly into the forwarding hardware.</p> <p>Previously, SR data plane validation was possible over IPv4-based LSPs.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>dataplane-only</b> keyword is introduced in the <b>traceroute sr-mpls</b> and <b>ping sr-mpls</b> commands.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Cisco-IOS-XR-mpls-traceroute-act.yang</a></li> <li>• <a href="#">Cisco-IOS-XR-mpls-ping-act.yang</a></li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

With this configuration, you can validate the SR-MPLS (Segment Routing over Multiprotocol Label Switching) IPv6-based LSPs and policies without disrupting the live network traffic. You can also validate Service-layer Application Programming Interface (SL-API) initiated LSPs such as controller instantiated LSPs. For more information about SL-API, refer *Use Service Layer API to Bring your Controller on Cisco IOS XR Router* chapter in *Programmability Configuration Guide*.

In the earlier releases, you could perform SR-MPLS data plane validation over IPv4-based LSPs. For more information, refer [Segment Routing Traceroute, on page 620](#) and [Segment Routing Ping, on page 618](#) sections.

## Examples: SR-MPLS Data Plane Validation over IPv6-based LSPs

The following example shows how to use segment routing ping to validate SR-MPLS over IPv6-based LSPs:

```
Router#ping sr-mpls dataplane-only 2001:DB8::1/32
Tue Jan 16 15:05:19.120 EST

Sending 5, 100-byte MPLS Echos with Nil FEC to 2001:DB8::1/32,
timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
```

The following example shows how to use segment routing traceroute to validate SR-MPLS over IPv6-based LSPs:

```
Router#traceroute sr-mpls dataplane-only 2001:DB8::1/32
Tue Jan 16 15:08:54.681 EST

Tracing MPLS Label Switched Path with Nil FEC to 2001:DB8::1/32, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 11:11:11::1 MRU 1500 [Labels: 18004/explicit-null Exp: 0/0]
L 1 11:11:11::2 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 3 ms
! 2 15:15:15::4 3 ms
```

The following example shows how to trace the SR-MPLS LSPs with Nil-FEC that includes labels:

```
Router#traceroute sr-mpls nil-fec labels 18004 output interface GigabitEthernet 0/0/0/0
nexthop 10:10:10::2
Tue Jan 16 15:28:03.162 EST

Tracing MPLS Label Switched Path with Nil FEC with labels [18004], timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 10:10:10::1 MRU 1500 [Labels: 18004/explicit-null Exp: 0/0]
```

```
L 1 10:10:10::2 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 2 ms
! 2 15:15:15::4 2 ms
```

## MPLS OAM support for SR-TE Policies using MPLS IPv6-based LSPs

*Table 102: Feature History Table*

| Feature Name                                                   | Release Information | Feature Description                                                                            |
|----------------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------|
| MPLS OAM support for SR-TE Policies using MPLS IPv6-based LSPs | Release 24.2.11     | This feature now extends support on the Cisco NCS 540 Series routers running on Cisco IOS XR7. |

| Feature Name                                                   | Release Information | Feature Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MPLS OAM support for SR-TE Policies using MPLS IPv6-based LSPs | Release 24.2.1      | <p>Introduced in this release on the following Cisco NCS 540 router variants running on Cisco IOS XR:</p> <ul style="list-style-type: none"> <li>• N540-ACC-SYS</li> <li>• N540X-ACC-SYS</li> <li>• N540-24Z8Q2C-SYS</li> </ul> <p>You can now verify the network configuration and paths and SR-TE policies set up, without interrupting or potentially disrupting live network traffic, for SR-MPLS (Segment Routing over Multiprotocol Label Switching) IPv6-based Label Switched Paths (LSPs).</p> <p>Previously, MPLS OAM support was only for IPv4-based LSPs.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>traceroute sr-mpls</b> and <b>ping sr-mpls</b> commands are extended to support IPv6 nexthop addresses.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• Cisco-IOS-XR-mpls-traceroute-act.yang</li> <li>• Cisco-IOS-XR-mpls-ping-act.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p> |

With this feature, you can now verify the SR-MPLS (Segment Routing over Multiprotocol Label Switching) IPv6-based LSPs and Segment routing for traffic engineering (SR-TE) policies without disrupting the live network traffic.

In the earlier releases, SR-MPLS support was limited to IPv4-based LSPs. For more information, refer [Segment Routing Ping](#) and [Segment Routing Traceroute](#) sections.

## Usage Guidelines and Limitations

The following usage guidelines and limitations apply:

- Provide a valid LSP endpoint for non-Null-FEC ping and traceroute operations for SR-TE policies.
- When a path is fully used with traffic, MPLS OAM might incorrectly report packet drops and failures.

## Examples: MPLS OAM support for SR-TE Policies with IPv6-based LSPs

The following example shows how to use segment routing traceroute for SR-TE policies with IPv6-based LSPs:

```
Router#traceroute sr-mpls nil-fec policy name srte_c_40_ep_2001:DB8::1
Tue Feb 6 12:07:38.295 EST

Tracing MPLS Label Switched Path with Nil FEC for SR-TE Policy srte_c_40_ep_2001:DB8::1,
timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

 0 12:12:12::1 MRU 1500 [Labels: 26134/explicit-null Exp: 0/0]
L 1 12:12:12::3 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 16 ms
! 2 16:16:16::4 16 ms
```

The following example shows how to use segment routing ping for SR-TE policies with IPv6-based LSPs:

```
Router#ping sr-mpls nil-fec policy name srte_c_40_ep_2001:DB8::1
Tue Feb 6 12:08:28.277 EST

Sending 5, 100-byte MPLS Echos with Nil FEC for SR-TE Policy srte_c_40_ep_2001:DB8::1,
timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
```

For more information about configuring SR-TE policies, refer [SR-TE Policy Overview, on page 225](#) section.

The following example shows how to use segment routing traceroute with labels using IPv6 LSPs:

```
Router#traceroute sr-mpls labels 18004 lsp-end-point 2001:DB8::1
Tue Feb 6 12:10:41.928 EST

Tracing MPLS Label Switched Path to NIL FEC with lsp end point 2001:DB8::1, SID Label(s)
[18004], timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0
```



Type escape sequence to abort.

```

0 11:11:11::1 MRU 1500 [Labels: 18004/explicit-null Exp: 0/0]
L 1 11:11:11::2 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 7 ms
! 2 15:15:15::4 3 ms

```

The following example shows how to use segment routing ping with labels using IPv6 LSPs:

```

Router#ping sr-mpls labels 18004 lsp-end-point 2001:DB8::1
Tue Feb 6 12:11:05.349 EST

```

```

Sending 5, 100-byte MPLS Echos with NIL FEC with lsp end point 2001:DB8::1, SID Label(s)
[18004],
 timeout is 2 seconds, send interval is 0 msec:

```

```

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
 'L' - labeled output interface, 'B' - unlabeled output interface,
 'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
 'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
 'P' - no rx intf label prot, 'p' - premature termination of LSP,
 'R' - transit router, 'I' - unknown upstream index,
 'X' - unknown return code, 'x' - return code 0

```

Type escape sequence to abort.

```

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/4 ms

```