



Cisco Crosswork Optimization Engine 7.0 User Guide

First Published: 2024-08-27

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



CONTENTS

CHAPTER 1

About Cisco Crosswork Optimization Engine 1

- Audience 1
- Overview of Cisco Crosswork Optimization Engine 1
- Crosswork Network Controller Solution and Crosswork Optimization Engine 2
- Bandwidth Feature Packs 3
- Crosswork Optimization Engine APIs 3

CHAPTER 2

Traffic Engineering in Cisco Crosswork Optimization Engine 5

- Supported SR-TE Policies and RSVP Tunnels 5
- What is Segment Routing? 6
- Segment Routing Path Computation Element (SR-PCE) 8
- SR-TE Policy PCC and PCE Configuration Sources 8
 - PCC-Initiated SR-TE Policy Example 9
- What is Resource Reservation Protocol (RSVP)? 9
- RSVP-TE Tunnel PCC and PCE Configuration Sources 11
 - PCC-Initiated RSVP-TE Tunnel Example 11
- Get a Quick View of Traffic Engineering Services 11
- View TE Event and Utilization History 13
- View Traffic Engineering Device Details 14
- Configure Traffic Engineering Settings 15
 - Configure TE Timeout Settings 15
 - Configure How Device Groups Are Displayed for Traffic Engineering 16
 - Configure TE Data Retention Settings 16
- Resolve SR-TE Policies and RSVP-TE Tunnels 17

CHAPTER 3	SR-MPLS and SRv6	19
	View SR-MPLS and SRv6 Policies on the Topology Map	19
	View SR-MPLS and SRv6 Policy Details	22
	Visualize IGP Path and Metrics	23
	Find Multiple Candidate Paths (MCPs)	24
	Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label	27
	Visualize Native SR Paths	30
	Visualize Native Path Device Prerequisites	31
	Configure TE Link Affinities	33
	Create Explicit SR-MPLS Policies	34
	Create Dynamic SR-MPLS Policies Based on Optimization Intent	35
	Modify SR-MPLS Policies	36
<hr/>		
CHAPTER 4	Resource Reservation Protocol (RSVP)	37
	View RSVP-TE Tunnels on the Topology Map	37
	View RSVP-TE Tunnel Details	39
	Create Explicit RSVP-TE Tunnels	42
	Create Dynamic RSVP-TE Tunnels Based on Optimization Intent	42
	Modify RSVP-TE Tunnels	43
<hr/>		
CHAPTER 5	Flexible Algorithm	45
	Configure Flexible Algorithm Affinities	45
	Visualize Flexible Algorithm Topologies	46
	View Flexible Algorithm Details	47
<hr/>		
CHAPTER 6	Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering	51
	Visualize Tree-SID Policies	51
	View a Point-to-Multipoint Tree on the Topology Map	52
	Create Static Tree-SID Policies	55
	Static Tree-SID Policy Configuration Example through Crosswork UI	57
	Modify a Tree-SID Policy	59
	Tree-SID Important Notes	59

PART I**Bandwidth Feature Packs 61**

CHAPTER 7**SR Circuit Style Manager (CSM) 63**

- Circuit Style SR-TE Important Notes 63
- Workflow for Setting Up CS SR-TE Policy Visualization 68
- Enable SR Circuit Style Manager 69
- Configure Circuit Style SR Policies 70
- Review Circuit Style SR-TE Policy Bandwidth Utilization 72
- View Circuit Style SR-TE Policies 73
- Trigger CSM to Recalculate a Circuit Style SR-TE Policy 78
- What Happens When Bandwidth Reservation Settings are Exceeded? 79
- How Does CSM Handle Path Failures? 83

CHAPTER 8**Local Congestion Mitigation (LCM) 87**

- Local Congestion Mitigation Overview 87
- LCM Congestion Evaluation Requirements 88
 - Enable Strict SID for LCM Usage 88
- LCM Congestion Mitigation Requirements 90
- LCM Important Notes 91
 - BGP-LS Speaker Placement for Multiple AS Networks with a Dedicated IGP Instance Between ASBRs 92
- LCM Calculation Workflow 93
- Workflow Example: Mitigate Congestion on Local Interfaces 95
- Configure LCM 100
- Add Individual Interface Thresholds 102
- Monitor LCM Operations 104
- Temporarily Exclude an Interface from LCM 106

CHAPTER 9**Bandwidth on Demand (BWoD) 109**

- BWoD Important Notes 109
- PCC-Initiated BWoD SR-TE Policies 110
- Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example 111
- Configure Bandwidth on Demand 117

Troubleshoot BWoD 117



CHAPTER 1

About Cisco Crosswork Optimization Engine

- [Audience, on page 1](#)
- [Overview of Cisco Crosswork Optimization Engine, on page 1](#)
- [Crosswork Network Controller Solution and Crosswork Optimization Engine, on page 2](#)
- [Bandwidth Feature Packs, on page 3](#)
- [Crosswork Optimization Engine APIs, on page 3](#)

Audience

This guide is for experienced network administrators who want to use Cisco Crosswork Optimization Engine in their network. This guide assumes that you are experienced and familiar with using the following technologies:

- Networking technologies and protocols (BGP-LS, IGP (OSPF and IS-IS), PCEP, model-driven telemetry, and so on)
- Traffic Engineering (TE) tunnels:
 - RSVP-TE tunnel provisioning
 - Segment Routing Traffic Engineering (SR-TE) policy provisioning
- Cisco Segment Routing Path Computation Element (SR-PCE)
- Point to Multi Point Tree (Tree-SID)
- Flexible Algorithm

Overview of Cisco Crosswork Optimization Engine

Cisco Crosswork Optimization Engine is part of the Cisco Crosswork Network Automation suite of products and provides the ability to preserve network intent with network monitoring, network visualization, and closed-loop automation. It also provides real-time network optimization allowing operators to effectively maximize network utilization and increase service velocity.

Crosswork Optimization Engine is offered as an individual application and as the Traffic Engineering component of Cisco Crosswork Network Controller (see [Crosswork Network Controller Solution and Crosswork Optimization Engine, on page 2](#)).

Crosswork Optimization Engine provides the following:

- A topology map that gives valuable network visualization of:
 - devices
 - links and link utilization
 - SR-TE policies
 - SR-MPLS and SRv6
 - Tree-SID
 - Flexible Algorithm
 - Circuit-Style Segment Routing (CS-SR) policies
 - RSVP-TE tunnels
- A UI that allows the network operator to:
 - provision SR policies and RSVP-TE tunnels and modify or remove them using an intuitive workflow
 - preview an SR policy or RSVP-TE tunnel before deploying it to the network
 - continuously track SR policy dynamic path computations to maintain SLA objectives (with correct licensing)
 - provision static Point to Multi Point (Tree-SID) policies
 - provision circuit Style SR-TE policies
- APIs that extend Crosswork Optimization Engine functions to other Crosswork applications and third-party applications.

This guide covers all the Crosswork Optimization Engine capabilities. However, you may not be able to access the features and functions due to licensing or the configuration of the role associated with your user account (see the [Crosswork Network Controller Administration Guide](#)). For licensing and ordering information, work with your Cisco Partner or Cisco account representative.

Crosswork Network Controller Solution and Crosswork Optimization Engine

Cisco Crosswork Network Controller is a turnkey network automation solution for deploying and operating IP transport networks. It delivers increased service agility, cost efficiency, and optimization for faster time-to-customer value and lower operating costs. The solution combines intent-based network automation to deliver critical capabilities for service orchestration and fulfillment, network optimization, service path computation, device deployment and management, anomaly detection, and automatic remediation. For more information, see [Cisco Crosswork Network Controller](#).

Crosswork Optimization Engine can be activated by itself or in conjunction with other Crosswork components (see the [Crosswork Network Controller Installation Guide](#)). Throughout this document, when using Crosswork Optimization Engine as part of the Crosswork Network Controller solution, some options are unavailable or

slightly different. For example, to navigate to the Traffic Engineering UI, instead of **Traffic Engineering > Traffic Engineering**, the navigation within the Crosswork Network Controller solution is **Services & Traffic Engineering > Traffic Engineering**.

Bandwidth Feature Packs

Crosswork Optimization Engine feature packs (available with certain licensing) are tools designed to address congestion and manage SR-TE policies, ensuring the maintenance and optimization of requirements based on bandwidth, latency, and so on. Users can specify their optimization intent, and the tools will then implement it while continually monitoring, tracking, and reacting to uphold the original intent. To learn more about these feature packs, see the following topics:

- [Local Congestion Mitigation \(LCM\)](#), on page 87
- [SR Circuit Style Manager \(CSM\)](#), on page 63
- [Bandwidth on Demand \(BWoD\)](#), on page 109

Crosswork Optimization Engine APIs

Advanced users can integrate other Crosswork applications and third-party applications with Crosswork Optimization Engine functions by using application programming interfaces (APIs), delivering new capabilities to their network operations.

For more information, see the [Cisco Crosswork Network Automation API Documentation on Cisco DevNet](#).



CHAPTER 2

Traffic Engineering in Cisco Crosswork Optimization Engine

Traffic engineering (TE) is a method of optimizing and steering traffic in a network to achieve an operational goal or provide custom services, such as using guaranteed bandwidth routes for prioritized traffic. One way TE can improve network performance is by forcing traffic to take predetermined routes and by effectively using available resources.

One of the biggest advantages of using Crosswork is the ability to visualize SR-TE policies and RSVP-TE tunnels on a topology map. By visually examining your network, the complexity of provisioning and managing these SR-TE policies is significantly reduced.

Existing SR-TE policies and RSVP-TE in brownfield deployments

Crosswork discovers existing policies and tunnels when devices are imported, but cannot manage them. Crosswork can only manage policies that were provisioned in Crosswork.

This section contains the following topics:

- [Supported SR-TE Policies and RSVP Tunnels, on page 5](#)
- [What is Segment Routing?, on page 6](#)
- [Segment Routing Path Computation Element \(SR-PCE\), on page 8](#)
- [SR-TE Policy PCC and PCE Configuration Sources, on page 8](#)
- [What is Resource Reservation Protocol \(RSVP\)?, on page 9](#)
- [RSVP-TE Tunnel PCC and PCE Configuration Sources, on page 11](#)
- [Get a Quick View of Traffic Engineering Services , on page 11](#)
- [View TE Event and Utilization History, on page 13](#)
- [View Traffic Engineering Device Details, on page 14](#)
- [Configure Traffic Engineering Settings, on page 15](#)
- [Resolve SR-TE Policies and RSVP-TE Tunnels, on page 17](#)

Supported SR-TE Policies and RSVP Tunnels

Crosswork Traffic Engineering supports the visualization and provisioning of most SR-TE policies and RSVP tunnels. In networks where there are preexisting policies that were not provisioned in Crosswork, they will be discovered, but cannot be managed.

Table 1: Supported TE Technologies

TE Technology	Crosswork Optimization Engine	
	Visualize	Provision
SR-MPLS	✓	✓
SRv6	✓	✗
RSVP	✓	✓
Flexible Algorithm	✓	✗ ¹
Tree-SID	✓	✓ ²
Circuit Style	✓	✓

¹ When provisioning SR-TE policies, you can use segment lists with SIDs that are part of a Flexible Algorithm.

² Only static Tree-SID policies are supported. Dynamic Tree-SID policies can be provisioned manually on the device or via an API.



Note Crosswork supports the use of Role-based Access Control (RBAC) to limit not only what functions a user can perform, but also on which devices they are allowed to perform those functions, see the "[Cisco Crosswork Network Controller Administration Guide](#)".

For a list of known limitations, important notes, and what networking technologies are supported, see the [Cisco Crosswork Optimization Engine Release Notes](#).

What is Segment Routing?

Segment routing for traffic engineering takes place through a tunnel 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. 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 an unsigned 32-bit integer. 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 tunnel.

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. It is manually configured from the segment routing global block (SRGB) range of labels and distributed by IS-IS (Intermediate System to Intermediate System) or OSPF (Open Shortest Path First). The prefix segment steers 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 node's loopback address 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. This label represents a specific adjacency, such as an egress interface, to a neighboring router. The adjacency SID is distributed by IS-IS or OSPF. The adjacency segment steers 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 multi-path (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.

Segment Routing Policies

Segment routing for traffic engineering uses a “policy” to steer traffic through the network. An SR 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, instructing the network routers to follow the specified path instead of the shortest path calculated by the IGP. If a packet is steered into an SR policy, the head-end pushes the SID list on the packet. The rest of the network executes the instructions embedded in the SID list.

Crosswork supports the visualization (and some provisioning) of the following SR-related policies:

- [SR-MPLS and SRv6](#) , on page 19
- [Flexible Algorithm](#), on page 45
- [Tree Segment Identifier \(Tree-SID\) Multicast Traffic Engineering](#), on page 51
- [SR Circuit Style Manager \(CSM\)](#)

There are two types of SR policies: dynamic and explicit.

Dynamic SR Policy

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 path computation engine (PCE).

Explicit SR Policy

When configuring an explicit policy, you specify an explicit path consisting of a list of prefixes or adjacency SIDs, each representing a node or link along the path.

Disjointness

Crosswork uses the disjoint policy to compute two lists of segments that steer traffic from two source nodes to two destination nodes along disjoint paths. 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. The following disjoint path computations are supported:

- **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 Share Risk Link Group (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, a list of segments is computed, encoding the shortest path from the first source to the first destination. When the second request is received with the same disjoint-group ID, the information received in both requests is used to compute two disjoint paths: one path from the first source to the first destination and another from the second source to the second destination. Both paths are computed at the same time. The shortest lists of segments are calculated to steer traffic on the computed paths.

**Note**

- Disjointness is supported for two policies with the same disjoint ID.
- Configuring affinity and disjointness at the same time is not supported.

Segment Routing Path Computation Element (SR-PCE)

Crosswork Network Controller uses a combination of telemetry and data collected from the Cisco Segment Routing Path Computation Element (SR-PCE) to analyze and compute optimal TE tunnels.

Cisco SR-PCE is provided by the Cisco IOS XR operating system running on either a physical device or a virtual router running within a virtual machine. SR-PCE provides stateful PCE functionality that helps control and reroute TE tunnels to optimize the network. PCE describes a set of procedures by which a Path Computation Client (PCC) can report and delegate control of headend tunnels sourced from the PCC to a PCE peer. The PCC and PCE establish a Path Computation Element Communication Protocol (PCEP) connection that SR-PCE uses to push updates to the network.

Crosswork discovers all devices in the IGP domain, including those that do not establish PCEP peering with SR-PCE. However, PCEP peering is required to deploy TE tunnels to these devices.

**Note**

Certain features may not function as expected if the SR-PCE version is not supported. To avoid any compatibility issues, refer to the [Crosswork Optimization Engine Release Notes](#) for SR-PCE version support and compatibility.

For SR-PCE and HA configuration, see the "Prepare Infrastructure for Device Management: Manage Providers" section in the [Cisco Crosswork Network Controller Administration Guide](#).

SR-TE Policy PCC and PCE Configuration Sources

SR-TE policies discovered and reported by Crosswork may have been configured from the following sources:

- Path Computation Client (PCC) initiated—Policies configured on a PCC (see [PCC-Initiated SR-TE Policy Example, on page 9](#)). This type displays as Unknown in the UI.
- Path Computation Element (PCE) initiated—Policies configured on a PCE or created dynamically by Crosswork. Examples of PCE Initiated policy types:
 - **Dynamic**
 - **Explicit**
 - **Bandwidth on Demand** (can be either PCC or PCE)
 - **Local Congestion Mitigation**



Note SR policies that are configured using the UI are the only types of SR-TE policies that you can modify or delete in Crosswork.

PCC-Initiated SR-TE Policy Example

The following example shows a configuration of an SR-TE policy at the headend router. The policy has a dynamic path with affinity constraints computed by the headend router. See SR configuration documentation for your specific device to view descriptions and supported configuration commands (for example, [Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers](#)).

```
segment-routing
traffic-eng
policy foo
  color 100 end-point ipv4 1.1.1.2
  candidate-paths
  preference 100
  dynamic
  metric
  type te
  !
  !
constraints
  affinity
  exclude-any
  name RED
  !
  !
  !
  !
```

What is Resource Reservation Protocol (RSVP)?

Resource Reservation Protocol (RSVP) is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations.

The RSVP-TE process contains the following functionalities:

- Endpoint control - is associated with establishing and managing TE tunnels at the headend and tail end.
- Link-management - manages link resources to do resource-aware routing of TE Label-Switched Path (LSP) and to program MPLS labels.
- Fast Reroute (FRR) - manages the LSPs that need protection and assigns backup tunnel information to these LSPs.

The interactions between TE and RSVP assume the existence of the endpoint control, link-management, and FRR functionality within TE.

RSVP-TE Explicit Routing (Strict, Loose)

RSVP-TE explicit routes are particular paths in the network topology that you can specify as abstract nodes in the Explicit Route Object (ERO). These nodes could be a sequence of IP prefixes or a sequence of autonomous systems. The explicit path can be administratively specified or automatically computed using an algorithm such as constrained shortest path first (CSPF).

The explicit path specified in the ERO could be a strict path or a loose one.

A strict path means that a network node and its preceding node in the ERO must be adjacent and directly connected.

A loose hop means that a network node specified in the ERO must be in the path but is not required to be directly connected to its preceding node. If a loose hop is encountered during ERO processing, the node that processes the loose hop can update the ERO with one or more nodes along the path from itself to the next node in the ERO. The advantage of a loose path is that the entire path does not need to be specified or known when creating the ERO. The disadvantage of a loose path is that it can result in forwarding loops during transients in the underlying routing protocol.



Note RSVP-TE tunnels cannot be configured with loose hops when provisioning within the UI.

RSVP FRR

When a router's link or neighboring device fails, the router often detects this failure by receiving an interface-down notification. When a router notices that an interface has gone down, it switches LSPs going out of that interface onto their respective backup tunnels (if any).

The FRR object is used in the PATH message and contains a flag that identifies the backup method to be used as facility-backup. The FRR object specifies setup and hold priorities, which are included in a set of attribute filters and bandwidth requirements to be used in the selection of the backup path.

The Record Route Object (RRO) reports in the RESV message the availability or use of local protection on an LSP and whether bandwidth and node protection are available for that LSP.

The signaling of the FRR requirements is initiated at the TE tunnel headend. Points of Local Repair (PLR) along the path act on the FRR requirements based on the backup tunnel availability at the PLR and signal the backup tunnel selection information to the headend. When an FRR event is triggered, the PLR sends PATH messages through the backup tunnel to the merge point (MP), where the backup tunnel rejoins the original LSP. The MP also sends RESV messages to the PLR using the RSVP-Hop object that is included by the PLR in its PATH message. This process prevents the original LSP from being torn down by the MP. Also, the PLR signals the tunnel headend with a PATH-ERROR message to indicate the failure along the LSP and that FRR

is in active use for that LSP. The headend uses this information to signal a new LSP for the TE tunnel and to tear down the existing failed path after the new LSP is set up through make-before-break techniques.

RSVP-TE Tunnel PCC and PCE Configuration Sources

RSVP-TE tunnels discovered and reported by Crosswork may have been configured from the following sources:

- Path Computation Client (PCC) initiated—RSVP-TE tunnels configured on a PCC (see [PCC-Initiated RSVP-TE Tunnel Example, on page 11](#)).
- Path Computation Element (PCE) or PCC initiated dynamically.

PCC-Initiated RSVP-TE Tunnel Example

The following is a sample device configuration for a PCC-initiated RSVP-TE tunnel. See the appropriate documentation to view descriptions and supported RSVP-TE tunnel configuration commands for your particular device (for example, *MPLS Command Reference for Cisco NCS 5500 Series, Cisco NCS 540 Series, and Cisco NCS 560 Series Routers*).

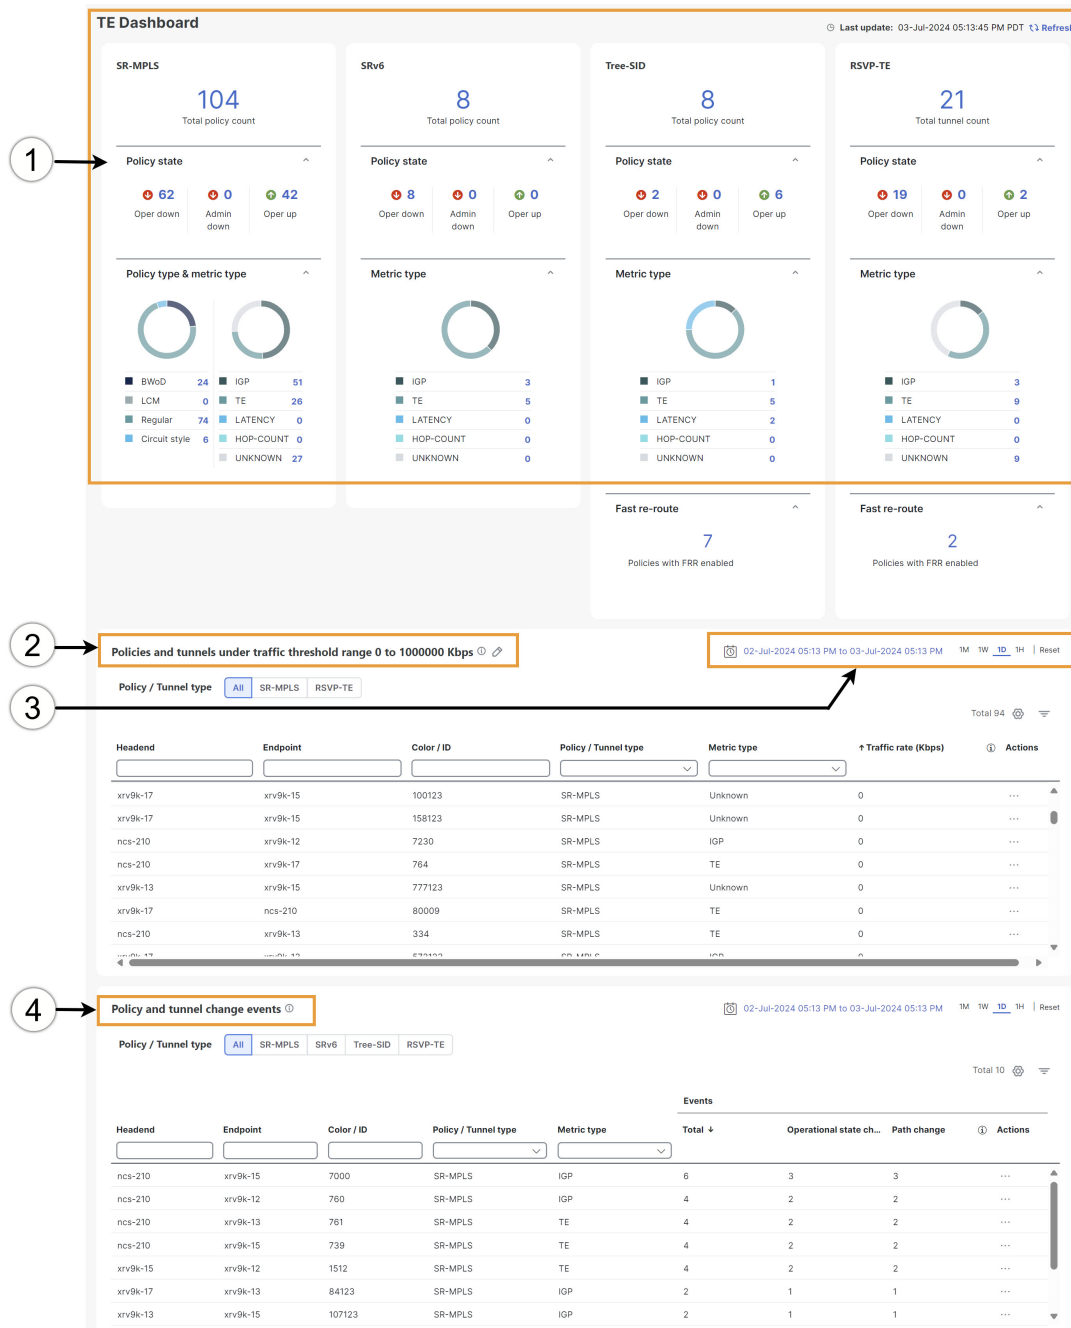
```
interface tunnel-te777
  ipv4 unnumbered Loopback0
  destination 192.168.0.8
  path-option 10 dynamic
  pce
  delegation
!
```

Get a Quick View of Traffic Engineering Services


The TE Dashboard provides a high-level summary of RSVP-TE tunnel, SR-MPLS, SRv6, and Tree-SID policy information.

To get to the TE Dashboard, choose **Traffic Engineering > TE Dashboard**.

Figure 1: Quick View of Traffic Engineering Services



Note If you are viewing the HTML version of this guide, click the images to view them in full-size.

Callout No.	Description
1	<p>Traffic Engineering Dashlet: Displays the total policy count and count of policies according to the policy state.</p> <p>It also displays the number of all TE policies and the number of policies or tunnels according to the metric types for all TE services.</p> <p>To drill down for more information, click on a value. The topology map and TE table appear, displaying only the filtered data you clicked on.</p>
2	<p>Policies and Tunnels Under Traffic Threshold:</p> <p>Displays RSVP-TE tunnels and SR-MPLS policies with traffic below the defined threshold in the selected time period. This information may be used to find and filter the unused policies or tunnels. Click  to update the LSP threshold range and change the units from Kbps to Mbps.</p> <p>Note Traffic utilization is not captured for SRv6 and Tree-SID policies.</p>
3	<p>Allows you to filter the data on the dashlet based on the time range you want to view (date, 1 month, 1 week, 1 day, and 1 hour).</p>
4	<p>Policy and Tunnel Change Events: Displays all the policies and tunnels that have had a path or state change event ordered by the event count, within the selected time range. This information helps identify the unstable policies and tunnels.</p> <p>Note The addition or deletion of leaf nodes for Tree-SID policies is captured as events.</p>



Note For a list of known limitations, see the [Cisco Crosswork Optimization Engine Release Notes](#)

View TE Event and Utilization History

The historical data captures the traffic rate and change events for a policy or tunnel. To view the historical data:



Note Traffic Rate is not captured for SRv6 and Tree-SID policies.

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering** .


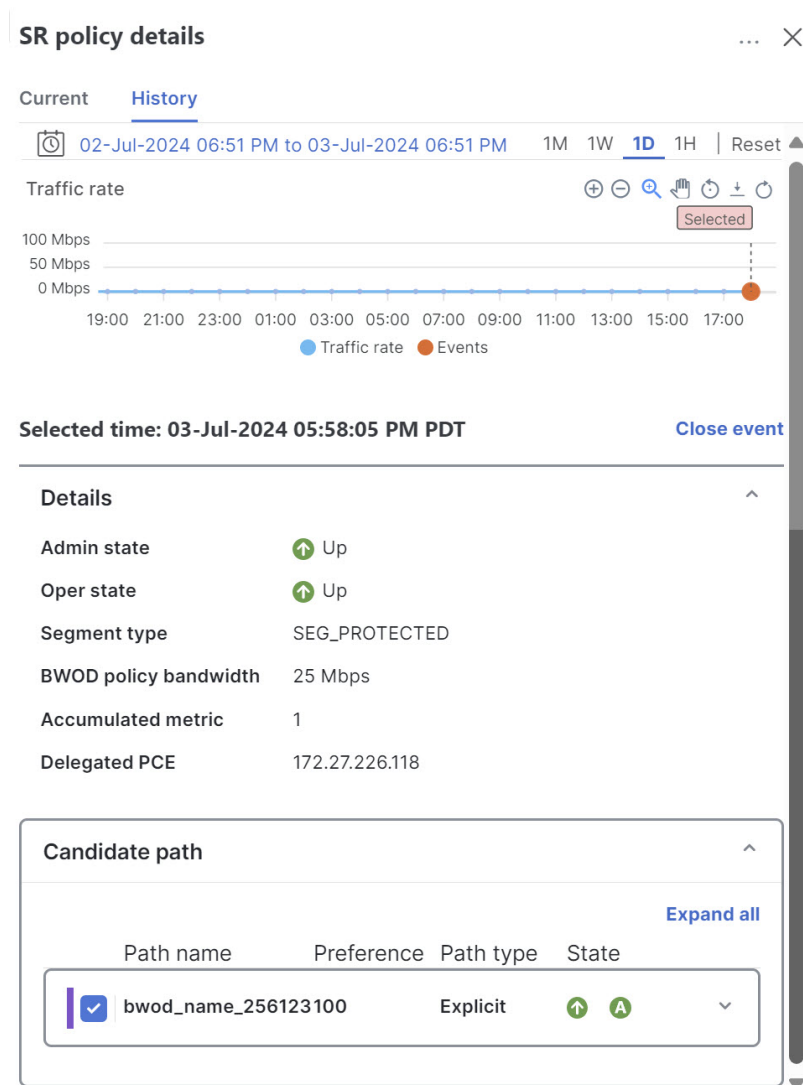
Step 2 From the **Actions** column of the Traffic Engineering table, click  > **View Details > History tab** tab for a policy or tunnel. The tab displays associated historical data for that device. Click on the event to see the path or state change event information.

Figure 2: TE Event and Utilization History



View Traffic Engineering Device Details

To view Traffic Engineering Device details (SR-MPLS, SRv6, RSVP-TE, and Flexible Algorithm information), do the following:

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering**.
- Step 2** From the Traffic Engineering topology map, click on a device.
- Step 3** From the **Traffic Engineering** tab, click on the policy type you are interested in. Each tab displays associated data for that device. From the browser, you can copy the URL and share with others.

The following example shows the Tree-SID information details for the selected device.

Note If you are viewing the HTML version of this guide, click on the image to view it in full-size.

Figure 3: Traffic Engineering Device Details

Device details

Details Links [Traffic engineering](#)

General SR-MPLS SRv6 [Tree-SID](#) RSVP-TE Flex Algo

Selected 0 / Total 5

	Root name	Root IP	Name	Tree ID	Label	Type	Programmin...	Fast reroute	PCE address	Admin status	Oper status	Actions
<input type="checkbox"/>												
<input type="checkbox"/>	xrv9k-13	192.168.0.3	DAY_O_TREE...	-	35	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-17	192.168.0.7	MY_FIRST_T...	-	15200	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	R4_TREE_SID	-	22	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	netflix	-	15202	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	ncs-210	192.168.0.6	prime	-	15203	Static	None	Enable	172.27.226.118			...

Configure Traffic Engineering Settings

Configure TE Timeout Settings

To configure timeout settings for the provisioning and retrieval of data for SR-TE policies, RSVP-TE tunnels, Bandwidth on Demand and IGP paths, select **Administration > Settings > System settings** tab > **Traffic engineering > General settings**. Enter the timeout duration options. For more information, click .



Note Timeouts change the response time of action if SR-PCE is slow in responding. You can modify the settings for a large-scale topology or to address slow SR-PCE response due to latency or load.

Figure 4: Traffic Engineering Timeout Settings

Configure How Device Groups Are Displayed for Traffic Engineering

You can configure what is shown on the topology map when a device group is selected, and a device in the selected SR policy, service, or RSVP-TE tunnel does not belong in the group. To set the behavior, choose **Administration > Settings > User settings** tab > **Switch device group** and select one of the behavior options.

By default, the user is asked to choose the device group view each time.

Configure TE Data Retention Settings

To see a historical view of LSP utilization (Historical tab), you must enable LSP utilization collection and specify how long data should be retained. To do this, click **Administration > System settings > Data retention > Network performance** and check the **LSP utilization** check box. Optionally, you can edit the default data retention periods.



Note If the retention period is reduced, all data older than the new retention period is lost. For example, if the daily retention interval is set to 31 days, then reduced to 7 days, then all data older than 7 days will be deleted.

Resolve SR-TE Policies and RSVP-TE Tunnels

Orphaned TE policies are any PCE initiated SR-TE policies (SRv6, SR-MPLS, and Tree-SID) or RSVP-TE tunnels that were created within Crosswork and *after* the last cluster data synchronization. After a switchover in a High Availability setup, Crosswork automatically checks for any orphaned TE policies. Orphaned policies/tunnels may also happen after a backup/restore operation. You can view policy details but not modify them since they were not included in the last data synchronization. Crosswork will display an alarm when it finds orphan TE policies (**Alerts > Alarms and Events**).

Crosswork provides APIs to help clear these orphans. To get a list of orphan SR-TE policies or RSVP-TE tunnels, use `cisco-crosswork-optimization-engine-sr-policy-operations:sr-datalist-oper` or `cisco-crosswork-optimization-engine-rsvp-te-tunnel-operations:rsvp-te-datalist-oper` where `is-orphan=True` and default action is GET. To make the orphans manageable again, use a SAVE action for the corresponding URL per policy type. For more information, see [API documentation on Devnet](#) (**Crosswork Optimization Engine APIs > release-id Release APIs**).



CHAPTER 3

SR-MPLS and SRv6

This section describes the SR-MPLS and SRv6 policy features that Crosswork supports. For a list of known limitations and important notes, see the [Cisco Crosswork Optimization Engine Release Notes](#).

- [View SR-MPLS and SRv6 Policies on the Topology Map, on page 19](#)
- [View SR-MPLS and SRv6 Policy Details, on page 22](#)
- [Visualize IGP Path and Metrics, on page 23](#)
- [Find Multiple Candidate Paths \(MCPs\), on page 24](#)
- [Visualize Underlying Paths Associated with a Defined Binding-Segment ID \(B-SID\) Label, on page 27](#)
- [Visualize Native SR Paths, on page 30](#)
- [Configure TE Link Affinities, on page 33](#)
- [Create Explicit SR-MPLS Policies, on page 34](#)
- [Create Dynamic SR-MPLS Policies Based on Optimization Intent, on page 35](#)
- [Modify SR-MPLS Policies, on page 36](#)

View SR-MPLS and SRv6 Policies on the Topology Map

To get to the Traffic Engineering topology map, choose **Traffic Engineering > Traffic Engineering**.

From the Traffic engineering table, click the checkbox of each SR-MPLS or SRv6 policy you want to view on the map. You can select up to 10 policies that will appear as separate colored links.

Figure 5: Traffic Engineering UI : SR-MPLS and SRv6 Policies



Callout No.	Description
1	<p>Click the appropriate check box to enable the following options:</p> <ul style="list-style-type: none"> • Show: IGP path—Displays the IGP path for the selected SR-TE policy. • Show: Participating only—Displays only links that belong to selected SR-TE policy. All other links and devices disappear.
2	<p>A device with an orange (🔴) outline indicates there is a node SID associated with that device or a device in the cluster.</p>
3	<p>When SR-TE policies are selected in the SR-MPLS or SRv6 tables, they show as colored directional lines on the map indicating source and destination.</p> <p>An adjacency segment ID (SID) is shown as an orange circle on a link along the path (🔴).</p>
4	<p>SR-MPLS and SRv6 Policy Origin and Destination: If both A and Z are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The A+ denotes that there is more than one SR-TE policy that originates from a node. The Z+ denotes that the node is a destination for more than one SR policy.</p>
5	<p>The content of this window depends on what has been selected or filtered. In this example, the SR-MPLS tab is selected, and the SR Policy table is displayed.</p>
6	<p>Click on either the SR-MPLS or SRv6 tabs to view the respective list of SR-TE policies.</p>
7	<p>Exports <i>all</i> data into a CSV file. You cannot export selected or filtered data.</p>
8	<p>The Mini Dashboard provides a summary of the operational SR-MPLS or SRv6 policy status. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the SR Policy and SRv6 Policy tables. In addition to the policy status, the SR-MPLS Mini Dashboard table displays the number of PCC and PCE initiated tunnels that are <i>currently</i> listed in the SR Policy table.</p>
9	<p>This option allows you to choose how the group filter (when in use) should be applied on the table data. For example, if Headend only was selected, then it would only display policies where the headend device of the policy is in the selected group. This filter allows you to see specific configurations and is useful when you have a large network.</p> <p>Filter options:</p> <ul style="list-style-type: none"> • Headend or Endpoint—Show policies with either the headend or endpoint device in the selected group. • Headend and Endpoint—Show policies if both the headend and endpoint are in the group. • Headend only—Show policies if the headend device of the policy is in the selected group. • Endpoint only—Show policies if the endpoint device of the policy is in the selected group.

View SR-MPLS and SRv6 Policy Details

View SR-MPLS or SRv6 TE policy level details as well segment lists and any path computation constraints configured on a per-candidate path basis.

Step 1 From the **Actions** column, click **⋮** > **View details** for one of the SR-MPLS or SRv6 policies.

Figure 6: View SR Policy Details

The screenshot displays the Traffic Engineering interface. On the left, a map shows the San Francisco Bay Area with various locations marked. A path is highlighted in purple, connecting San Francisco to San Jose. On the right, a table lists SR-MPLS policies. The policy 'xrv9k-22' is selected, and a 'View details' button is visible in the Actions column.

SR-MPLS	SRv6	Tree-SID	RSVP-TE
35	7	7	0
Total	Circuit style	BWd	LCM
		0	0
		Admin down	Oper up
		18	17
		Oper up	Oper down


Headend	Endpoint	Color	Admin status	Oper status	Actions
<input type="checkbox"/>	xrv9k-25	xrv9k-27	11111	●	● ⋮
<input type="checkbox"/>	xrv9k-24	xrv9k-26	1234	●	● ⋮
<input type="checkbox"/>	xrv9k-22	xrv9k-24	2151	●	● ⋮
<input checked="" type="checkbox"/>	xrv9k-22	xrv9k-24	666666	●	● ⋮
<input type="checkbox"/>	xrv9k-22	xrv9k-26	9090	●	● ⋮
<input type="checkbox"/>	xrv9k-22	xrv9k-26	2023	●	● ⋮
<input type="checkbox"/>	xrv9k-24	xrv9k-23	1982	●	● ⋮
<input type="checkbox"/>	xrv9k-25	xrv9k-24	1982	●	● ⋮
<input type="checkbox"/>	xrv9k-22	xrv9k-23	1992	●	● ⋮
<input type="checkbox"/>	xrv9k-23	xrv9k-24	1992	●	● ⋮
<input type="checkbox"/>	xrv9k-24	xrv9k-23	6915	●	● ⋮
<input type="checkbox"/>	xrv9k-23	xrv9k-25	6915	●	● ⋮
<input type="checkbox"/>	xrv9k-23	xrv9k-25	6925	●	● ⋮
<input type="checkbox"/>	xrv9k-24	xrv9k-23	6929	●	● ⋮
<input type="checkbox"/>	xrv9k-23	xrv9k-25	6929	●	● ⋮


Step 2 View SR-MPLS or SRv6 policy details. From the browser, you can copy the URL and share with others.

Figure 7: SR Policy Details - Headend, Endpoint, and Summary

SR policy details





Current History

Headend  xrv9k-22 | Source IP: 192.168.0.22
TE RID: 192.168.0.22 | IPv6 RID: 2001:192:168::22
PCC IP: 192.168.0.22

Endpoint  xrv9k-24 | Dest IP: 192.168.0.24
TE RID: 192.168.0.24 | IPv6 RID: 2001:192:168::24

color 2151



Summary

Admin state	 Up
Oper state	 Up
Binding SID	24007
Policy type	Regular
Profile ID	-
Description	-
Traffic rate	0 Mbps
Unused	True 
Delay	241 
Accumulated metric	1
Delegated PCE	172.27.226.126
Non-delegated PCEs	-
PCE computed time	15-Feb-2024 09:25:02 AM PDT
Last update	25-Feb-2024 11:38:37 PM PDT

[See less](#)

Candidate path

[Expand all](#)

Path name	Preference	Path type	State
<input checked="" type="checkbox"/> 2151-bwod	100	Unknown	 

Note The Delay value is calculated for all policies every 10 minutes. Hover your mouse over the "i" icon (next to the Delay value) to view the last time the value was updated.

Visualize IGP Path and Metrics

View the physical path and metrics between the endpoints of the selected SR-MPLS policies.

Step 1 From the **SR Policy** table, check the check box next to the SR-TE (SR-MPLS and SRv6) policies you are interested in.

Find Multiple Candidate Paths (MCPs)

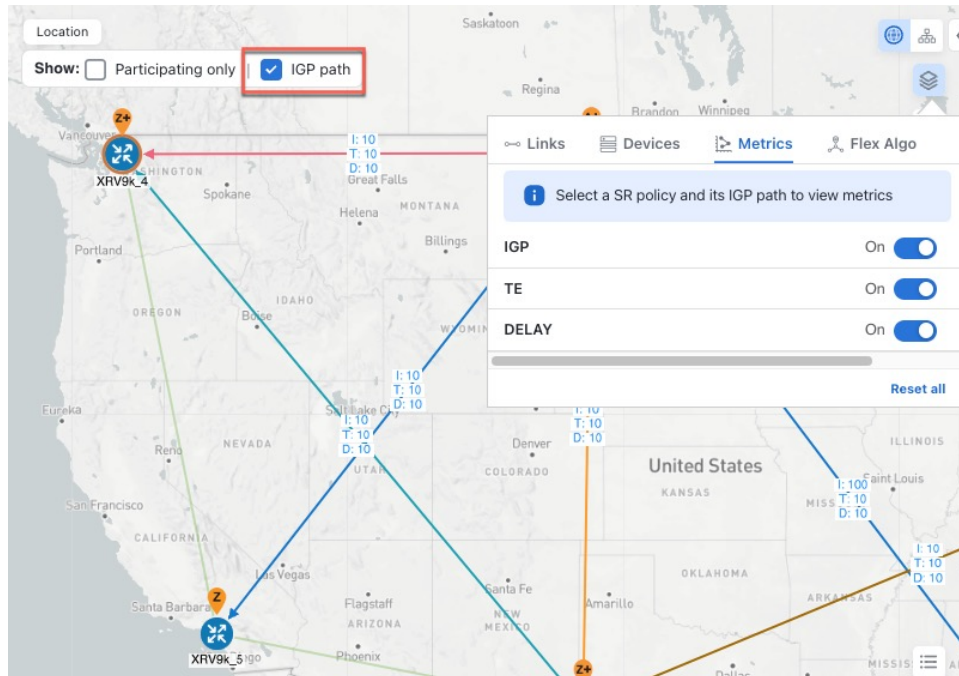
Step 2 Check the **Show IGP Path** check box. The IGP paths for the selected SR-MPLS policies are displayed as straight lines instead of the segment hops. In a dual-stack topology, the **Participating only** checkbox must also be checked to view metrics on participating links.

Step 3 Click  > **Metrics** tab.

Step 4 Toggle applicable metrics to **ON**.

Note You must check the **Show IGP Path** check box to view metrics.

Figure 8: View Physical Path and Metrics



Find Multiple Candidate Paths (MCPs)

Visualizing MCPs gives you insight into which paths might be a better alternative to the currently active ones. If you determine to do so, you can then manually configure the device and change which path becomes active.

Important Notes

- Only PCC-initialized SR-TE policies with MCPs are supported.
- Crosswork does not distinguish dynamic paths from explicit paths. The Policy Type field value displays as 'Unknown'.
- You can view active explicit paths but not inactive candidate explicit paths in the UI.

Before you begin

A policy must be configured with MCPs on devices before they can be visualized on the Traffic Engineering topology map. This configuration can be done manually or within the Crosswork Network Controller.

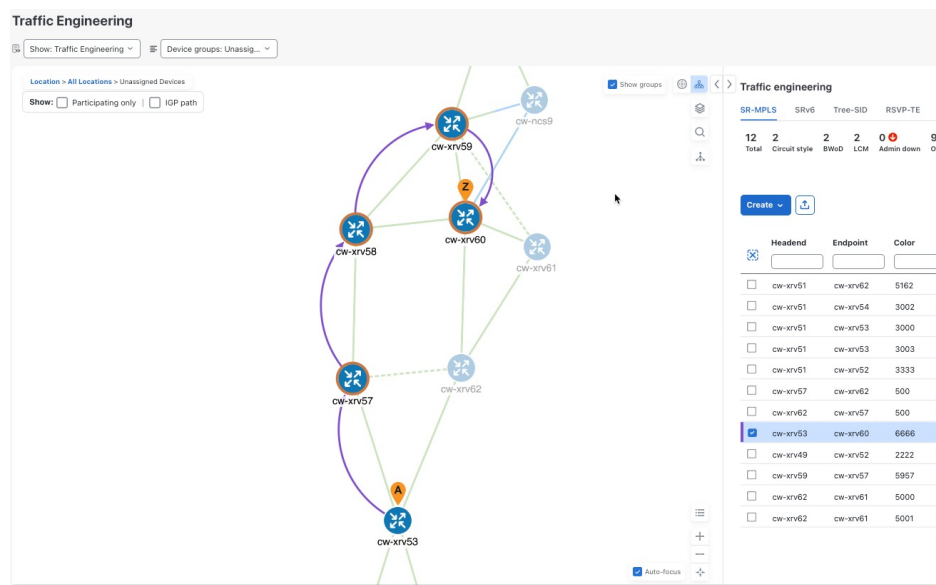
Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** or **SRv6** tab.

Step 2 Navigate to the active SR-TE policy that has MCPs configured and view it on the topology map.

- Check the check box next to the SR-TE policy that has MCPs configured.
- View the SR-TE policy that is highlighted on the topology map.

In this example, you see that the active path is going from **cw-xrv53 > cw-xrv57 > cw-xrv58 > cw-xrv59 > cw-xrv60**.

Figure 9: SR-TE policy on the Topology Map




Step 3 View the list of candidate paths.


- From the SR-MPLS or SRv6 Policy table **Actions** column, click ***** > View details**. A list of candidate paths appear along with policy details in the **SR policy details** window. The green A under the State column indicates the active path.

Figure 10: Candidate Path in SR Policy Details

SR policy details ...


Current **History**


Headend  cw-xrv53 | Source IP: 3.3.3.53
TE RID: 3.3.3.5 | IPv6 RID: bb:bb:bb:3:3:
PCC IP: 3.3.3.E3

Endpoint  cw-xrv60 | Dest IP: 3.3.3.60
TE RID: 3.3.3.5

color 6666

Summary ^

Admin state  Up

Oper state  Up


Binding SID 24035

Policy type Regular

Profile ID -

Description -




Traffic rate 0 Mbps

Unused True 

[See more](#) v

Candidate path ^

[Expand all](#)

Path name	Preference	Path type	State
<input checked="" type="checkbox"/> cfg_mcp-53-60_discr_25	25	Unknown	  v
<input checked="" type="checkbox"/> cfg_mcp-53-60_discr_20	20	Unknown	 v

Step 4 You can expand individual paths or click **Expand all** to view details of each path.

Step 5 Visualize the candidate path on the topology map.


- a) Check the check box next to any candidate path.

Note You will not be able to select or view explicit candidate paths.

- b) From the **Candidate path** area, hover your mouse over the candidate path name. The candidate path is highlighted on the topology map.

In this example, you see that the alternate path goes directly from **cw-xrv53** > **cw-xrv60**.

Figure 11: Candidate Path on the Topology Map



Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label

Crosswork Network Controller allows you to visualize the underlying path of a B-SID hop that you have manually configured on a device or configured using Crosswork Network Controller. In this example, we have assigned **15700** as a B-SID label on an SR-MPLS policy hop.


To view the B-SID underlying path for an SR-MPLS or SRv6 policy, do the following:

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS**.
- Step 2** From the SR Policy table, check the check box next to the policy that has a hop assigned with a B-SID label. Hover your mouse over any part of the SR-MPLS row to see the B-SID name. The B-SID path is highlighted in **orange** on the topology map.

In this example, you see that the B-SID path is going from **cw-xrv51** to **cw-xrv52**.

Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label

Figure 12: B-SID Label



Step 3 From the SR policy details window, click > **View details**.

Figure 13: View Details

	Head...	Endp...	Color	Admin ...	Oper s...	Actions
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	
<input type="checkbox"/>	CW-XI...	CW-XI...	3333			
<input type="checkbox"/>	CW-XF...	CW-XF...	500			
<input type="checkbox"/>	CW-XF...	CW-XF...	500			
<input type="checkbox"/>	CW-XF...	CW-XF...	6666			
<input checked="" type="checkbox"/>	CW-XF...	CW-XF...	2222			
<input type="checkbox"/>	CW-XF...	CW-XF...	5957			
<input type="checkbox"/>	CW-XF...	CW-XF...	5000			

View details

Edit / Delete

Step 4 Expand the active path and click the B-Sid Label ID to see the underlying path.

Figure 14: B-Sid Label ID

SR policy details

Current History

Candidate path

Collapse all

Path name	Preference	Path type	State
<input checked="" type="checkbox"/> cfg_test-bsid-policy_discr_1...100	Unknown	↑	A ^

Se...	Segm...	L...	Algo	IP	N...	Inter...	Sl...
0	N...	1...	0	3.3.3...	c...		R...
1	N...	1...	0	3.3.3...	c...		R...
2	B-Sid	15700		3.3.3...	c...		

Path name: cfg_test-bsid-policy_discr_100

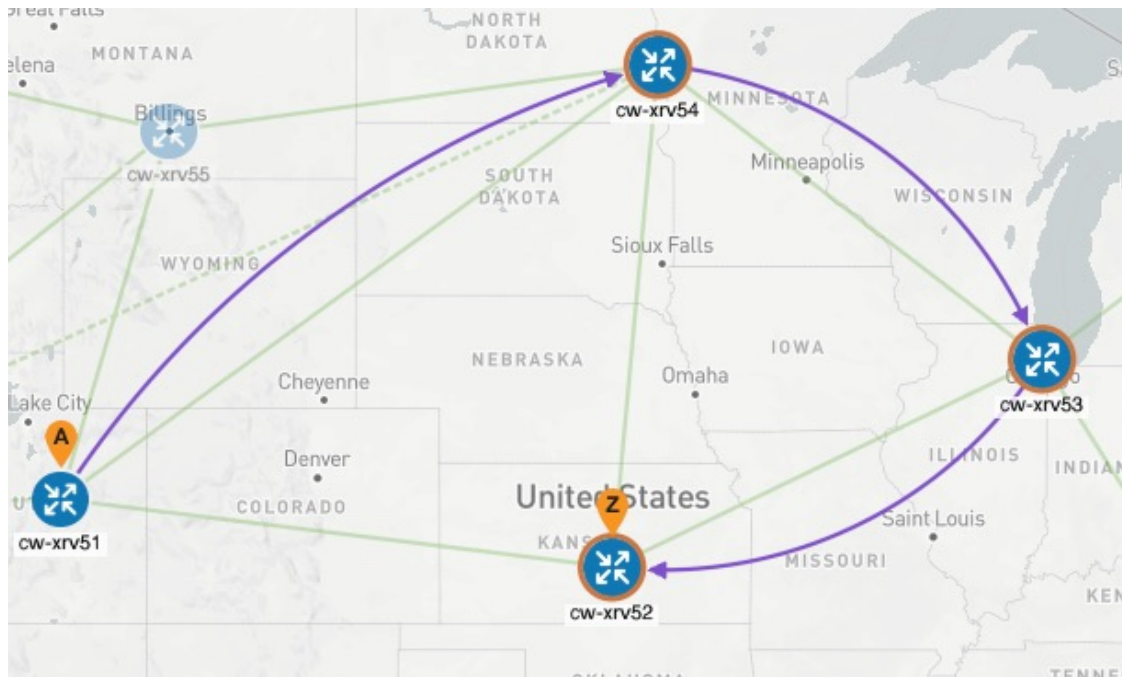
Oper state: ↑ Up | A Active

Metric type: TE

Bandwidth: -

In this example, the underlying path actually goes from **cw-xrv51** > **cw-xrv54** > **cw-xrv53** > **cw-xrv52**.

Figure 15: B-SID Path



Visualize Native SR Paths

Visualizing the native path will help you in OAM (Operations, Administration, and Maintenance) activities to monitor label-switched paths (LSPs) and quickly isolate forwarding problems to assist with fault detection and troubleshooting in the network. Since this feature uses multipaths, all ECMP paths are shown between the source and destination. You can visualize only native SR IGP paths.

Before you begin

Confirm that device requirements are met. See [Visualize Native Path Device Prerequisites](#), on page 31.

To create a path query, do the following:

- Step 1** From the main menu, choose **Traffic Engineering > Path Query**. The Path Query dashboard appears.
- Step 2** Click **New query**.
- Step 3** Enter the device information in the required fields to find available Native SR IGP Paths and click **Get paths**.

Note Path queries may take a moment to complete. When the Running Query ID pop-up appears, you can also select **View past queries** to return to the Path Query Dashboard. If you already had path queries in the list, you can view existing details as the new query continues to run in the background, which is indicated by the blue Running icon in the Query State column. When the new query state turns green, and is completed, it can be viewed.

Figure 16: New Path Query

New path query ✕

Select from the fields below to find available native SR IGP paths

Select service

Headend * ✕

Endpoint * ✕

Step 4

Click **View results** when it becomes available on the Running Query ID pop-up. The Path Details window appears with corresponding available paths details while the topology map displays the available Native SR IGP paths on the left.

Figure 17: Path Details

The screenshot shows a map of North America with a path highlighted between two devices (cw-xrv54 and cw-xrv60). The Path Details window is open, showing the following information:

Path details ✕

Select service

Headend * ✕

Endpoint * ✕

Available paths Last update: 1.3 Refresh

Path	Status
Path 0	Found
Path 1	
Path 2	

Hop details What is ret code?char?

Hop index:0 | Hop origin IP:3.3.3.54 | Hop destination IP:11.1.4.0 | MRU:1500 | Labels: [16060] | ret code:0 | multipaths:0

Hop index:1 | Hop origin IP:11.1.4.0 | Hop destination IP:3.3.3.60 | MRU:1500 | Labels: [16060] | ret code:8 | return char:L | multipaths:1

Hop index:2 | Hop origin IP:3.3.3.60 | Hop destination IP:12.1.8.0 | MRU:1500 | Labels: [implicit-null|16060] | ret code:8 | return char:L | multipaths:3

Hop index:3 | Hop origin IP:12.1.8.0 | Hop destination IP:12.1.8.0 | MRU:1500 | Labels: [implicit-null] | ret code:8 | return char:L | multipaths:1

Hop index:4 | Hop origin IP:12.1.8.0 | MRU:0 | ret code:3 | return char:| | multipaths:0

Visualize Native Path Device Prerequisites

Confirm the following device software and configurations are met prior to visualizing native paths.

1. Devices should be running Cisco IOS XR 7.3.2 or higher. Run `show version` command to verify it.
2. Devices should have GRPC enabled.
 - a. Run `show grpc` to confirm GRPC configuration. You should see something similar to this:

```
tpa
vrf default
address-family ipv4
default-route mgmt
!
```

```

address-family ipv6
default-route mgmt
!
!
!

or

linux networking
vrf default
address-family ipv4
default-route software-forwarding
!
address-family ipv6
default-route software-forwarding
!
!
!
```



-
- Note**
- `address-family` is only required in an IPv4 topology.
 - To enable GRPC with a secure connection, you must upload security certificates to connect to the device.
-

3. Devices should have GNMI capability enabled and configured.
 - a. From **Device Management > Network Devices**, click the IP address for the device you are interested in.
 - b. Confirm that GNMI is listed under **Connectivity details**.



-
- Note** Based on the type of devices, the following device encoding type are available:

- JSON
 - BYTES
 - PROTO
 - ASCII
 - JSON IETF
-

4. Devices should have the CDG router static address. Static route should be added from the device to the southbound CDG IP address. For example:

```
RP/0/RP0/CPU0:xrvr-7.3.2#config
```

```
RP/0/RP0/CPU0:xrvr-7.3.2(config)#router static
```

```
RP/0/RP0/CPU0:xrvr-7.3.2(config-static)#address-family ipv4 unicast <CDG Southbound
interface IP: eg. 172.24.97.110> <Device Gateway eg: 172.29.105.1>
```

```
RP/0/RP0/CPU0:xrvr-7.3.2(config-static)#commit
```

Configure TE Link Affinities

If you have any affinities you wish to account for when provisioning an SR policy, Tree-SID, or RSVP-TE tunnel, then you can optionally define affinity mapping on the Cisco Crosswork UI for consistency with affinity names in device configurations. Cisco Crosswork will only send bit information to SR-PCE during provisioning. If an affinity mapping is not defined in the UI, then the affinity name is displayed as "UNKNOWN". If you want to configure affinity mappings in Cisco Crosswork for visualization purposes, you should collect affinities on the device, then define affinity mapping in the Cisco Crosswork UI with the same name and bits that are used on the device.

The affinity configuration on interfaces simply turns on some bits. It is a 32-bit value, with each bit position (0–31) representing a link attribute. Affinity mappings can be colors representing a certain type of service profile (for example, low delay, high bandwidth, and so on). This makes it easier to refer to link attributes.

See SR, Tree-SID, or RSVP-TE configuration documentation for your specific device to view descriptions and supported configuration commands (for example, [Segment Routing Configuration Guide for Cisco ASR 9000 Series Router](#))

The following example shows the affinity configuration (`affinity-map`) on a device:

```
RP/0/RP0/CPU0:c12#sh running-config segment-routing traffic-eng affinity-map
Wed Jul 27 12:14:50.027 PDT
segment-routing
 traffic-eng
  affinity-map
   name red bit-position 1
   name blue bit-position 5
   name green bit-position 4
  !
 !
 !
```

Step 1 From the main menu, choose **Administration > Settings > System settings tab > Traffic engineering > Affinity > TE link affinities**. You can also define affinities while creating an SR-TE policy, Tree-SID, or RSVP-TE tunnel by clicking **Manage mapping** under the **Constraints > Affinity** field.

Step 2 To add a new affinity mapping, click **+ Create**.

Step 3 Enter the name and the bit it will be assigned. For example (using the above configuration):

Figure 18: Mapping Affinities

Name ⓘ	Bit position (0-31) ⓘ	Actions
<input type="text"/>	<input type="text"/>	
red	1	<button>Edit</button> <button>Delete</button>
blue	5	<button>Edit</button> <button>Delete</button>
green	4	<button>Edit</button> <button>Delete</button>

Step 4 Click **Save** to save the mapping. To create another mapping, you must click **+ Create** and save the entry.


Note You should remove the TE tunnel before removing the affinity to avoid orphan TE tunnels. If you have removed an affinity associated with a TE tunnel, the affinity is shown as "UNKNOWN" in the **SR policy / RSVP-TE tunnel details** window.


Create Explicit SR-MPLS Policies

This task creates SR-MPLS policies using an explicit (fixed) path consisting of a list of prefix or adjacency Segment IDs (SID list), each representing a node or link along on the path.



Tip If you plan to use affinities, collect affinity information from your devices, and then map them in Cisco Crosswork before creating an explicit SR-MPLS policy. For more information, see [Configure TE Link Affinities, on page 33](#).

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.
- Step 2** Click **+ Create**.
- Step 3** Under **Policy details**, enter or select the required SR-MPLS policy values. Hover the mouse pointer over the  to view a description of the field.
- Tip** If you have set up device groups, you can select the device group from the **Device Groups** drop-down list. Then navigate and zoom in on the topology map to click the device for headend or endpoint selection.
- Step 4** Under **Policy path**, click **Explicit path** and enter a path name.
- Step 5** Add segments that will be part of the SR-MPLS policy path.
- Step 6** Click **Preview** and confirm that the policy you created matched your intent.
- Step 7** If you want to commit the policy path, click **Provision** to activate the policy on the network or exit to abort the configuration process.
- Step 8** Validate the SR-MPLS policy creation:
- Confirm that the new SR-MPLS policy appears in the **Traffic engineering** table. You can also click the check box next to the policy to see it highlighted in the map.


Note The newly provisioned SR-TE policy may take some time, depending on the network size and performance, to appear in the table. The **Traffic engineering** table is refreshed every 30 seconds.
 - View and confirm the new SR-MPLS policy details. From the **Traffic engineering** table, click the  and select **View details**.
- Note** On a setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see [Configure TE Timeout Settings, on page 15](#).
-

Create Dynamic SR-MPLS Policies Based on Optimization Intent

This task creates an SR-MPLS policy with a dynamic path. SR-PCE computes a path for the policy based on metrics and path constraints (affinities or disjointness) defined by the user. A user can select from three available metrics to minimize in-path computation: IGP, TE, or latency. The SR-PCE will automatically re-optimize the path as necessary based on topology changes. If a link or interface fails, the network will find an alternate path that meets all the criteria specified in the policy and raise an alarm. If no path is found, an alarm is raised, and the packets are dropped.



Tip For visualization purposes, you can optionally collect affinity information from your devices and then map them in Cisco Crosswork before creating a dynamic SR-MPLS policy. For more information, see [Configure TE Link Affinities, on page 33](#) or [Configure Flexible Algorithm Affinities, on page 45](#).

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.
- Step 2** Click **+ Create**.
- Step 3** Under **Policy details**, enter or select the required SR-MPLS policy values. Hover the mouse pointer over  to view a description of each field.
- Tip** If you have set up device groups, you can select the device group from the **Device Groups** drop-down menu. Then navigate and zoom in on the topology map to click the device for headend or endpoint selection.
- Step 4** Under **Policy path**, click **Dynamic path** and enter a path name.
- Step 5** Under **Optimization objective**, select the metric you want to minimize.
- Step 6** Define any applicable constraints and disjointness.
- Note**
- Affinity constraints and disjointness cannot be configured on the same SR-MPLS policy. Also, there cannot be more than two SR-MPLS policies in the same disjoint group or subgroup. The configuration will not be allowed during Preview.
 - If there are existing SR-MPLS policies belonging to a disjoint group that you define here, all SR-MPLS policies that belong to that same disjoint group are shown during Preview.
- Step 7** Under **Segments**, select whether or not protected segments should be used when available.
- Step 8** If applicable, enter a SID constraint in the **SID Algorithm** field. Cisco Crosswork will try to find a path with this SID. If a path with the SID constraint cannot be found, the provisioned policy will remain operationally down until the conditions are met.
- Note**
- Flexible Algorithm: The values correspond to the Flexible Algorithm that are defined on the device and the 128-255 range is enforced by Cisco IOS XR.
 - Algorithm 0: This is a Shortest Path First (SPF) algorithm based on link metric. This shortest path algorithm is computed by the Interior gateway protocol (IGP).
 - Algorithm 1: This is a Strict Shortest Path First (SSPF) algorithm based on link metric. The algorithm 1 is identical to algorithm 0 but requires that all nodes along the path honor the SPF routing decision. Local policy does not alter the forwarding decision. For example, a packet is not forwarded through locally engineered path.

Step 9 Click **Preview**. The path is highlighted on the map.

Step 10 If you want to commit the policy path, click **Provision**.

Step 11 Validate the SR-MPLS policy creation:

- a. Confirm that the new SR-MPLS policy appears in the SR Policy table. You can also click the check box next to the policy to see it highlighted in the map.

Note The newly provisioned SR-MPLS policy may take some time, depending on the network size and performance, to appear in the **Traffic engineering** table. The table is refreshed every 30 seconds.

- b. View and confirm the new SR-MPLS policy details. From the **Traffic engineering** table, click and select **View details**.

Note On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see [Configure TE Timeout Settings, on page 15](#).

Modify SR-MPLS Policies

To view, modify, or delete an SR-MPLS policy, do the following:

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.

Step 2 From the **Traffic engineering** table, locate the SR-MPLS policy you are interested in and click .

Step 3 Choose **View details** or **Edit/Delete**.

- Note**
- You can only modify or delete SR-MPLS policies that have been created with the UI.
 - After updating the SR-MPLS policy details, you can preview the changes on the map before saving it.
-



CHAPTER 4

Resource Reservation Protocol (RSVP)

This section describes the RSVP-TE tunnel features that Crosswork Optimization Engine supports. For a list of known limitations and important notes, see the [Cisco Crosswork Optimization Engine Release Notes](#).

- [View RSVP-TE Tunnels on the Topology Map](#), on page 37
- [View RSVP-TE Tunnel Details](#), on page 39
- [Create Explicit RSVP-TE Tunnels](#), on page 42
- [Create Dynamic RSVP-TE Tunnels Based on Optimization Intent](#), on page 42
- [Modify RSVP-TE Tunnels](#), on page 43



View RSVP-TE Tunnels on the Topology Map

To get to the Traffic Engineering topology map for RSVP-TE visualization, choose **Traffic Engineering** > **Traffic Engineering** > **RSVP-TE** tab.

Figure 19: Traffic Engineering UI - RSVP-TE Tunnels

Tunnel ID	Headend	Endpoint	Admin Sta...	Oper Status	Actions	
<input checked="" type="checkbox"/>	3345	xrv9k-15	xrv9k-17	+	+	...
<input checked="" type="checkbox"/>	6000	xrv9k-16	xrv9k-14	+	-	...
<input checked="" type="checkbox"/>	10111	xrv9k-14	xrv9k-16	+	-	...
<input checked="" type="checkbox"/>	10112	xrv9k-15	xrv9k-17	+	-	...
<input type="checkbox"/>	10122	xrv9k-13	xrv9k-15	+	-	...
<input type="checkbox"/>	113	xrv9k-14	xrv9k-13	+	+	...
<input checked="" type="checkbox"/>	105	xrv9k-14	xrv9k-13	+	+	...
<input checked="" type="checkbox"/>	117	xrv9k-14	xrv9k-13	+	+	...
<input checked="" type="checkbox"/>	115	xrv9k-14	xrv9k-13	+	+	...
<input checked="" type="checkbox"/>	101	xrv9k-14	xrv9k-13	+	+	...

476165

Callout No.	Description
1	Click Show Participating Only to display links belonging to the selected RSVP-TE tunnels. All other links and devices disappear.
2	<p>A device with a solid orange outline () indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered.</p> <p>Note RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.</p>
3	<p>When RSVP-TE tunnels are selected in the RSVP-TE Tunnel table, they show as colored directional lines on the map indicating source and destination.</p> <ul style="list-style-type: none"> Record Route Object (RRO) paths are shown as straight lines. Explicit Route Object (ERO) paths are shown as curved lines. <p>Note If both RRO and ERO paths are available, the RRO path is displayed by default.</p> <ul style="list-style-type: none"> An adjacency segment ID (SID) is shown as a green dot on a link along the path () <p>If both A and Z are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The A+ denotes that there is more than one RSVP-TE tunnel that originates from a node. The Z+ denotes that the node is a destination for more than one RSVP-TE tunnel.</p>
4	<p>SR-MPLS and SRv6 Policy Origin and Destination: If both A and Z are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The A+ denotes that there is more than one SR-TE policy that originates from a node. The Z+ denotes that the node is a destination for more than one SR policy.</p>
5	<p>The content of this window depends on what has been selected or filtered. In this example, the RSVP-TE tab is selected and the RSVP-TE Tunnels table is displayed. Depending on what is selected on the topology map, or whether you are in the process of viewing and managing RSVP-TE tunnels, you can do the following:</p> <ul style="list-style-type: none"> Create Dynamic RSVP-TE Tunnels Based on Optimization Intent, on page 42 Create Explicit RSVP-TE Tunnels, on page 42 Modify RSVP-TE Tunnels, on page 43 View RSVP-TE Tunnel Details, on page 39
6	Click the RSVP-TE tab.
7	The Mini Dashboard provides a summary of the operational RSVP-TE tunnel status and the number of PCC and PCE initiated tunnels that are <i>currently</i> listed in the RSVP-TE tables. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the RSVP-TE table.

Callout No.	Description
8	<p>This option allows you to choose how the group filter (when in use) should be applied on the table data. For example, if Headend only was selected, then it would only display policies where the headend device of the policy is in the selected group. This filter allows you to see specific configurations and is useful when you have a large network.</p> <p>Filter options:</p> <ul style="list-style-type: none"> • Headend or Endpoint—Show policies with either the headend or endpoint device in the selected group. • Headend and Endpoint—Show policies if both the headend and endpoint are in the group. • Headend only—Show policies if the headend device of the policy is in the selected group. • Endpoint only—Show policies if the endpoint device of the policy is in the selected group.
9	Exports <i>all</i> data into a CSV file. You cannot export selected or filtered data.

View RSVP-TE Tunnel Details

View RSVP-TE tunnel details such as binding label, delegated PCE, metric type, ERO/RRO, delay, and so on.


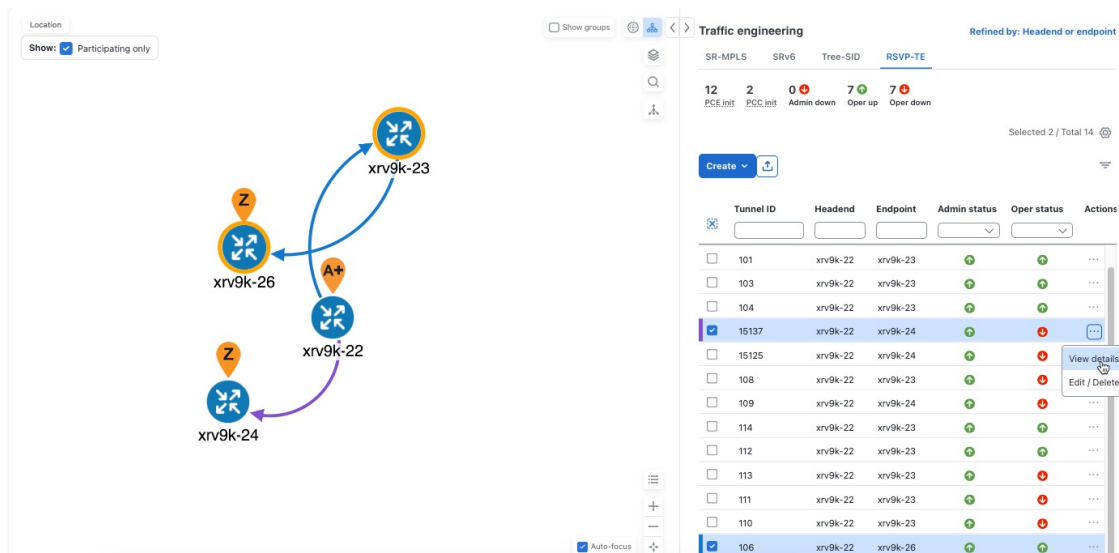
Step 1 From the **Actions** column, click  > **View details** for one of the RSVP-TE tunnels.

Figure 20: RSVP-TE > View details



The screenshot shows the 'Traffic engineering' section of the Cisco Crosswork Optimization Engine 7.0 interface. On the left, a network diagram illustrates a topology with four nodes: xrv9k-23 (top), xrv9k-26 (middle-left), xrv9k-22 (middle-right), and xrv9k-24 (bottom). Arrows indicate connections between these nodes. On the right, a table lists RSVP-TE tunnels. The table has columns for Tunnel ID, Headend, Endpoint, Admin status, Oper status, and Actions. Tunnel 15137 is selected, and the 'View details' option is highlighted in the Actions column. The table also shows other tunnels with their respective headend and endpoint devices and operational statuses.

Step 2 View RSVP-TE tunnel details. From the browser, you can copy the URL and share with others.

- Note**
- For end-to-end delays on RSVP-TE tunnels, inter-domain RSVP-TE tunnels must all be explicit (every interface along that path is specified as an adjacency hop).
 - If applicable, the Delay value is calculated for all policies every 10 minutes. Click the "i" icon (next to the Delay value) to view the last time the value was updated.

Figure 21: RSVP-TE Tunnel Details

The screenshot displays the 'RSVP-TE tunnel details' page. On the left, a network diagram shows two nodes, xrv9k-22 and xrv9k-23, connected by a purple line. Node xrv9k-23 is labeled with 'Z' and node xrv9k-22 with 'A'. The details panel on the right is divided into 'Current' and 'History' tabs. The 'Current' tab shows the following information:

- Headend: xrv9k-22 (192.168.0.22)
- Endpoint: xrv9k-23 (192.168.0.23)
- Tunnel ID: 104

The 'Summary' section lists various attributes:


- Description: -
- Path name: Path33442
- LSP ID: 4
- Path type: Unknown
- Admin state: Up
- Oper state: Up
- Traffic rate: 0 Mbps
- Unused: True
- Delay: 23 (with an 'i' icon next to it)
- Signaled bandwidth: 0 Mbps
- Setup / Hold priority: 7 / 7
- Metric type: TE
- Fast re-route (FRR): Disable
- Binding label: 24040
- Accumulated metric: 10
- Disjoint group: ID: 221

A tooltip is displayed over the 'Delay' value, showing 'Last updated' at '26-Mar-2024 03:41:56 PM PDT'. The interface also includes a 'Show groups' checkbox, a search icon, and an 'Auto-focus' checkbox at the bottom.

Figure 22: RSVP-TE Tunnel Details (close-up)

Create Explicit RSVP-TE Tunnels

This task creates RSVP-TE tunnels using an explicit (fixed) path consisting of a list of prefix or adjacency Segment IDs (SID list), each representing a node or link along the path.

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > RSVP-TE** tab.
- Step 2** Click **+ Create**.
- Step 3** Under **Tunnel details**, enter the required RSVP-TE tunnel values. Hover the mouse pointer over  to view a description of each field.
- Tip** If you have set up device groups, you can select the device group from the **Device groups: Location** drop-down menu. Then, navigate and zoom in on the topology map to click the device for headend or endpoint selection.
- Step 4** Under **Tunnel path**, click **Explicit path** and enter a path name.
- Step 5** Add segments that will be part of the RSVP-TE path.
- Step 6** Click **Preview**. The path is highlighted on the map.
- Step 7** If you want to commit the tunnel path, click **Provision**.
- Step 8** Validate the RSVP-TE tunnel creation:
- a. Confirm that the new RSVP-TE tunnel appears in the RSVP-TE Tunnels table. You can also click the check box next to the policy to see it highlighted in the map.

Note The newly provisioned RSVP-TE tunnel may take some time, depending on the network size and performance, to appear in the **Traffic engineering** table. The **Traffic engineering** table is refreshed every 30 seconds.
 - b. View and confirm the new RSVP-TE tunnel details. From the **Traffic engineering** table, click *** (in the same row as the RSVP-TE tunnel) and select **View details**.
- Note** A timeout may occur during policy deployment on a scaled setup with high node, policy, or interface counts. Please contact a Cisco representative to fine-tune the timers involved.
-

Create Dynamic RSVP-TE Tunnels Based on Optimization Intent

This task creates an RSVP-TE tunnel with a dynamic path. SR-PCE computes a tunnel path based on metrics and path constraints (affinity or disjointness) you defined. You can select from three available metrics to minimize in-path computation: IGP, TE, or delay. SR-PCE will also automatically re-optimize the path as necessary based on topology changes.



-
- Tip** If you plan to use affinities, collect affinity information from your devices and map them in Cisco Crosswork before creating a dynamic RSVP-TE tunnel. For more information, see [Configure TE Link Affinities, on page 33](#).
-

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > RSVP-TE** tab.
- Step 2** Click **+ Create**.
- Step 3** Under **Tunnel details**, enter the required RSVP-TE tunnel values. Hover the mouse pointer over ⓘ to view a description of each field.
- Tip** If you have set up device groups, you can select the device group from the **Device groups: Location** drop-down menu. Then, navigate and zoom in on the topology map to click the device for headend or endpoint selection.
- Step 4** Under **Tunnel path**, click **Dynamic path** and enter the Path Name.
- Step 5** Under **Optimization objective**, select the metric you want to minimize.
- Step 6** Define any applicable constraints and disjointness.
- Note** Affinity constraints and disjointness cannot be configured on the same RSVP-TE tunnel. Also, there can be up to two RSVP-TE tunnels in the same disjoint group or subgroup. If there are existing RSVP-TE tunnels belonging to a disjoint group that you define here, all RSVP-TE tunnels belonging to that same disjoint group are shown during Preview.
- Step 7** Click **Preview**. The path is highlighted on the map.
- Step 8** If you want to commit the tunnel path, click **Provision**.
- Step 9** Validate the RSVP-TE tunnel creation:
- Confirm that the new RSVP-TE tunnel appears in the RSVP-TE Tunnels table. You can also click the check box next to the policy to see it highlighted in the map.

Note The newly provisioned RSVP-TE tunnel may take some time, depending on the network size and performance, to appear in the **Traffic engineering** table. The **Traffic engineering** table is refreshed every 30 seconds.
 - View and confirm the new RSVP-TE tunnel details. From the **Traffic engineering** table, click and select **View details**.
- Note** A timeout may occur during policy deployment on a scaled setup with high node, policy, or interface counts. Please contact a Cisco representative to fine-tune the timers involved.
-

Modify RSVP-TE Tunnels

To view, modify, or delete an RSVP-TE tunnel, do the following:

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > RSVP-TE** tab.
- Step 2** Locate the RSVP-TE tunnel you are interested in and click .
- Step 3** Choose **View details** or **Edit/Delete**.

- Note**
- You can only modify or delete RSVP-TE tunnels that have been created with the UI or API.
 - After updating the RSVP-TE tunnel details, you can preview the changes on the map before saving it.
-



CHAPTER 5

Flexible Algorithm

Flexible Algorithm allow operators to customize and compute the IGP shortest path according to their needs and constraints (specific metrics and link properties). Many possible constraints can be used to compute a path over a network. For example, Flexible Algorithm can confine the path to a particular plane for networks with multiple logical planes. Since the meaning of the algorithm is not defined by any standard but is defined by the user, it is called a Flexible Algorithm.

Crosswork enables you to filter the IGP topology based on the Flexible Algorithm and visualize the network subset, which can provide a specific set of transport characteristics. Visualizing Flexible Algorithm topologies is an important tool to help you deploy, maintain, and verify that the configured Flexible Algorithm intent is realized in your network. For example, you may use Flexible Algorithm to improve service availability and define disjoint logical topologies to increase resiliency to network failures. Crosswork allows you to visualize both Flexible Algorithm topologies simultaneously and verify they have no common nodes or links. Or, if they do, help you determine the common network elements so that you can update Flexible Algorithm configurations.

This section contains the following topics:

- [Configure Flexible Algorithm Affinities, on page 45](#)
- [Visualize Flexible Algorithm Topologies, on page 46](#)
- [View Flexible Algorithm Details , on page 47](#)

Configure Flexible Algorithm Affinities

Flexible Algorithm affinity names that are defined on devices are not collected by Crosswork. You can optionally define affinity mapping on the Cisco Crosswork UI for consistency with Flexible Algorithm affinity names in device configurations. Cisco Crosswork will only send bit information to SR-PCE during provisioning. If an affinity mapping is not defined in the UI, then the affinity name is displayed as "UNKNOWN". If you want to configure affinity mappings in Cisco Crosswork for visualization purposes, you should collect affinities on the device, then define affinity mapping in the Cisco Crosswork UI with the same name and bits that are used on the device.

See SR configuration documentation for your specific device to view descriptions and supported configuration commands (for example, [Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers](#))

The following example shows the Flexible Algorithm affinity configuration (`affinity-map`) on a device:

```
router isis CORE
 is-type level-2-only
 net 49.0001.0000.0000.0002.00
 log adjacency changes
```

```

affinity-map b33 bit-position 33
affinity-map red bit-position 1
affinity-map blue bit-position 5
flex-algo 128
  priority 228
  advertise-definition
  affinity exclude-any blue indigo violet black
!
```

For visualization purposes, you must map the affinity names to the bits using the following procedure:

-
- Step 1** From the main menu, select **Administration > Settings > Traffic engineering > Affinity > Flex-Algo affinities** tab.
 - Step 2** To add a new Flexible Algorithm affinity mapping, click **+ Create** .
 - Step 3** Enter the name and the bit it will be assigned.
 - Step 4** Click **Save** to save the mapping. To view all Flexible Algorithm affinities for a link, see [View Flexible Algorithm Details](#) , on page 47.
-

Visualize Flexible Algorithm Topologies

Crosswork allows you to visualize Flexible Algorithm nodes and links on the topology map that have been manually configured or dynamically provisioned using the UI in your network.




Note To apply a Flexible Algorithm constraint when dynamically provisioning an SR-MPLS policy, see [Create Dynamic SR-MPLS Policies Based on Optimization Intent](#), on page 35.

Before you begin

You must understand and configure Flexible Algorithms in your network. See the SR Flexible Algorithm configuration documentation for your specific device to view descriptions and supported configuration commands (for example, [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#)).



Note You cannot visualize Flexible Algorithms if a Flexible Algorithm ID is the same across different domains.

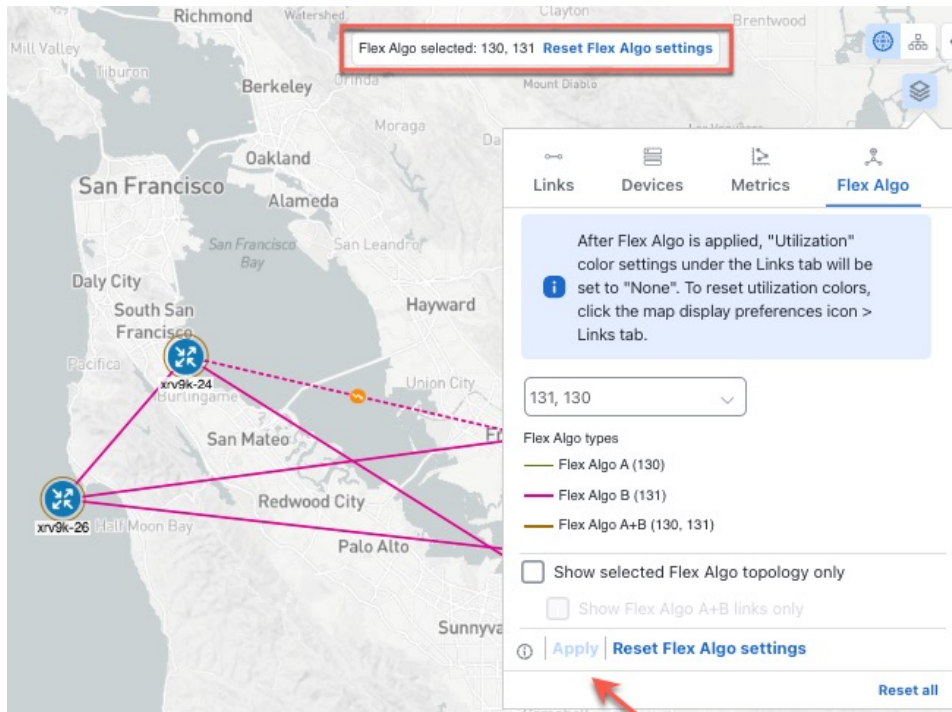
-
- Step 1** From the main menu, select **Traffic Engineering > Traffic Engineering**.
 - Step 2** From the topology map, click .
 - Step 3** Click the **Flex Algo** tab.
 - Step 4** From the drop-down list, select up to two Flexible Algorithm IDs.
 - Step 5** View the Flexible Algorithm Types and confirm that the selection is correct. Also, note the color assignments for each Flexible Algorithm.
 - Step 6** (Optional) Check the **Show selected Flex Algo topology only** check box to isolate the Flexible Algorithms on the topology map. When this option is enabled, SR policy selection is disabled.

- a) Check the **Show Flex Algo A+B links only** to show those links and nodes participating in both Flexible Algorithms.

Step 7

Click **Apply**. You must click **Apply** for any additional changes to Flexible Algorithm selections to be reflected on the topology map.

Figure 23: Flexible Algorithm on Map



Note If a selected Flexible Algorithm is defined with criteria but there are no link and node combinations that match it (for example, a defined affinity to include all nodes or links with the color blue), then the topology map will be blank. If a selected Flexible Algorithm is not configured on a node or link, the default blue link or node color appears.

Step 8

(Option) Click **Save View** to save the topology view and Flexible Algorithm selections.

View Flexible Algorithm Details

To view device or link Flex Algorithm details, do the following:

Step 1

From the main menu, choose **Traffic Engineering > Traffic Engineering**.

Step 2

To view device Flexible Algorithm details:

- From the topology map, click on a device.
- In the **Device details** window, navigate to the **Traffic engineering > Flex Algo** tab. For example:

Figure 24: Flex Algo Device Details


The screenshot shows a 'Device details' window with a close button (X) in the top right. The window has three main tabs: 'Details', 'Links', and 'Traffic engineering'. Under 'Traffic engineering', there are sub-tabs: 'General', 'SR-MPLS', 'SRv6', 'Tree-SID', 'RSVP-TE', and 'Flex Algo'. The 'Flex Algo' sub-tab is active, showing a list of Flexible Algorithms. At the top, it displays 'IGP: Domain ID: 1000, ISIS system ID: 0000.0000.0004, Level: 2' and a 'Collapse all' button. The list includes 'Algo 128', 'Algo 129', 'Algo 130', 'Algo 131', and 'Algo 132'. 'Algo 130' is expanded to show the following details:

Participating	Yes
Elected definition	Metric type:LATENCY
	Exclude-any affinity:
	Include-any affinity:
	Include-all Affinity:
Advertised	Yes
	Priority:128
	Definition equal to local:No

Step 3 To view whether a link is part of a Flexible Algorithm topology:

- From the topology map, click on a link. If a list of links appear, click on a link type.
- From the **Link details** window, click the **Traffic engineering** tab. If the link is a member, then the **FA topologies** row displays what Flexible Algorithm each source and destination device belong to.

Figure 25: Flex Algo Link Details

> Link details 

Summary **Traffic engineering**

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A side	Z side
Node	xrv9k-24	xrv9k-26
IF name	GigabitEthernet0/0/0/4	GigabitEthernet0/0/0/1
FA affinities		
FA TE metric		
FA delay metric		
FA topologies	128, 129, 130, 131, 132, ...	128, 129, 130, 131, 132, ...

Note

- Application-Specific Link Attribute (ASLA) is supported on PCC and core routers that are Cisco IOS XR 7.4.1 or later versions.
- Crosswork Network Controller only supports strict ASLA handling for Flexible Algorithm topologies.
- For Flexible Algorithms defined with Traffic Engineering (TE) or Delay metric types, only nodes advertising OSPF or IS-IS ASLA TE and ASLA Delay link metrics will be included in the corresponding Flexible Algorithm topology.



CHAPTER 6

Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering

Tree-SID is a method of implementing tree-like multicast flows over a segmented routing network. Using Tree-SID, an SDN controller (a device running SR-PCE using PCEP) calculates the tree. Each node (device) in the tree has a specific role in routing the multicast data through the tree. These roles include Ingress for the root or headend node, Transit or Bud for midpoint nodes that are not leaf nodes, and Egress for destination leaf nodes. The tree itself is assigned a single SID label, representing all of the tree segments and devices. The SDN controller is highly flexible, as it understands the segmentation and can construct routing paths using any constraints that network architects can specify.

The most interesting use case for constraint-based Tree-SID use is where routers are configured to deliver two P2MP streams with the same content over different paths. Here, the multicast flow is forwarded twice through the network, each copy following a unique path. The two copies never use the same node or link to reach the destination, reducing packet loss due to network failures on any one of the paths.

For detailed information on Tree-SID, see the Segment Routing Tree-SID configuration documentation for your specific device (for example, [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#)).

This section contains the following topics.

- [Visualize Tree-SID Policies, on page 51](#)
- [View a Point-to-Multipoint Tree on the Topology Map, on page 52](#)
- [Create Static Tree-SID Policies, on page 55](#)
- [Modify a Tree-SID Policy, on page 59](#)

Visualize Tree-SID Policies

Crosswork UI provides the ability to view details of the Tree-SID root, transit, leaf nodes, and bud nodes in the UI and allows you to easily confirm that Tree-SID is implemented correctly in your network (see [View a Point-to-Multipoint Tree on the Topology Map, on page 52](#)).

The Tree-SID policy has the following nodes:

- Root node—Encapsulates the multicast traffic, replicates it, and forwards it to the transit nodes.
- Transit node—Acts as a leaf (egress) node and a mid-point (transit) node toward the downstream sub-tree.
- Leaf node—Decapsulates the multicast traffic and forwards it to the multicast receivers.
- Bud Node—Has a separate leaf node path and is displayed separately in the topology map.

You can visualize the following Tree-SID policies:

- **Static:** A Static Tree-SID policy is configured via SR-PCE, directly using SR-PCE CLI or from the Crosswork UI. You can refer to the Tree-SID configuration documentation for your specific device for more information and examples of the supported configuration commands. (for example, [Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers](#))
- **Dynamic:** A Dynamic Tree-SID policy is not explicitly configured; it is configured as part of an L3VPN/mVPN service.



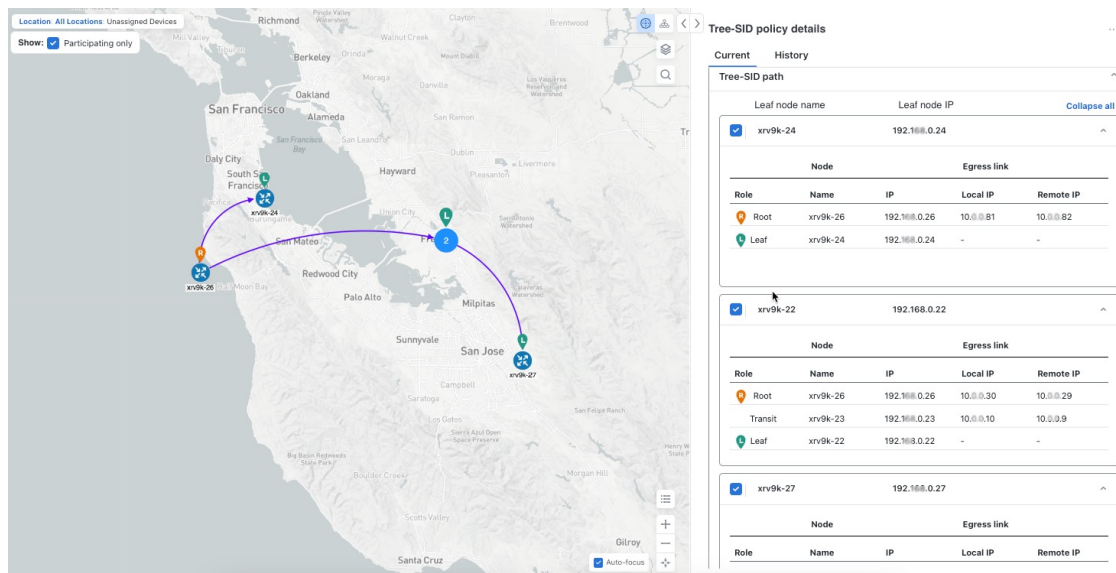
Note Static and Dynamic Tree-SID policies support fast reroute (FRR).

View a Point-to-Multipoint Tree on the Topology Map

Crosswork allows you to visualize Tree-SID policies configured in your network.

The following example shows a representation of a Tree-SID policy in the topology map. The root node (R) and leaf nodes (L) are marked, and the arrows denote the path through the transit nodes from the root to the leaf nodes.

Figure 26: Create a new Tree-SID Policy (static)



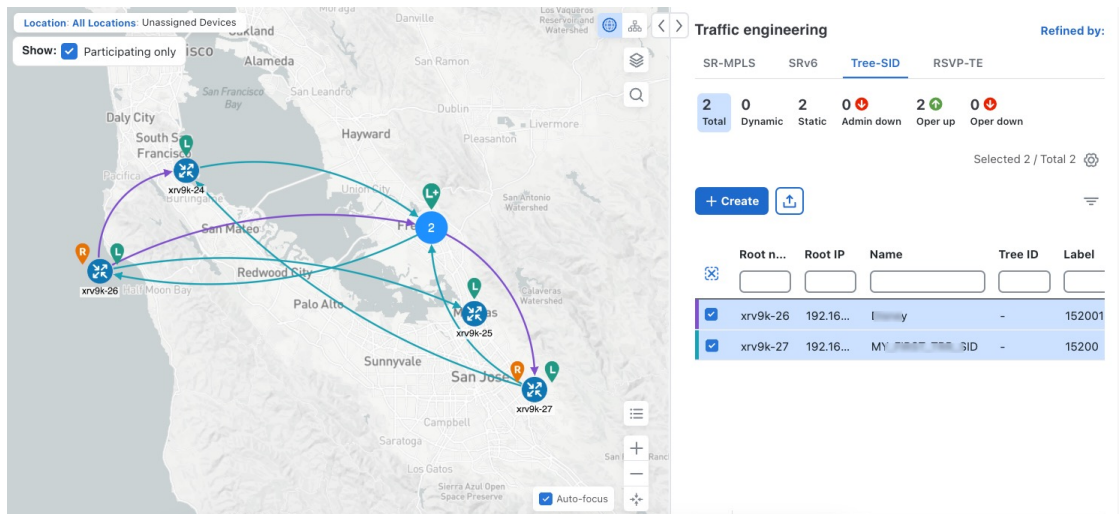
Before you begin

To visualize a multicast tree in the topology map, Tree-SID policies must be configured in your network. For more information, see the SR Tree-SID configuration documentation for your specific device (for example, [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#)).

Step 1 From the main menu, select **Traffic Engineering > Traffic Engineering > Tree-SID** tab.

Step 2 Select the Tree-SID policies you want to view on the topology map. You can view a maximum of two policies on the topology map at the same time.

Figure 27: Tree-SID Policies (static) on the Topology Map



Note Any change in end-points is captured as an event in the historical data tab. For information on Tree-SID Historical Data, see [View TE Event and Utilization History, on page 13](#)

Step 3 To view Tree-SID details, from the **Actions** column, click > **View details** for one of the Tree-SID policies. You will see a summary and Tree-SID path information.


Example:

- Note**
- A (Compute) label, next to the SR-PCE field, details the SR-PCE used to create the policies.
 - If a source node is unavailable, a warning icon and message appear next to the Oper Status field (hover your mouse over the warning icon), detailing where the connection issue resides.


Figure 28: Tree-SID Details Summary

Tree-SID policy details ... X




Current History

Root  xrv9k-26 | Root IP: 192. .0.26
TE RID: 192. .26 | IPv6 RID: 2001:192: .:26

Name Disney



Tree ID - 



Summary ^

Admin state	 Up
Oper status	 Up
Label	152001
Type	Static 
Programming state	None
Metric type	TE
Constraints	Exclude-Any: - Include-Any: - Include-All: -
FRR protected	Disable
Node count	Leaf: 3 Bud: 0 Transit: 1
Path compute elements (SR-PCEs)	172.27.226.126(Compute)
Last updated	05-Mar-2024 04:39:49 AM PDT

[See less ^](#)

Figure 29: Tree-SID Path Details

Leaf node name		Leaf node IP		Collapse all	
<input checked="" type="checkbox"/>	xrv9k-24	192.168.0.24			
Node		Egress link			
Role	Name	IP	Local IP	Remote IP	
 Root	xrv9k-26	192.168.0.26	10.0.0.81	10.0.0.82	
 Leaf	xrv9k-24	192.168.0.24	-	-	

<input checked="" type="checkbox"/>	xrv9k-22	192.168.0.22			
Node		Egress link			
Role	Name	IP	Local IP	Remote IP	
 Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29	
	Transit	xrv9k-23	192.168.0.23	10.0.0.10	10.0.0.9
 Leaf	xrv9k-22	192.168.0.22	-	-	

Create Static Tree-SID Policies

This task will explain how to create a static Tree-SID policy, each representing a leaf or a root node.



Tip If you plan to use affinities, collect affinity information from your devices and map them in Cisco Crosswork before creating a static Tree-SID policy. For more information, see [Configure TE Link Affinities, on page 33](#).

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > Tree-SID** tab and click **Create**.

Step 2 Enter or select the required Tree-SID policy values. Hover the mouse pointer over  to view a description of the field.

Note You can only add PCC nodes with a PCEP session to PCE as root nodes.

Figure 30: Create Static Tree-SID Policy

> Tree-SID policy (static)

Name *
tree-n9k

Tree-SID label * ⓘ
18

Root * ⓘ
Selected - cw-ncs9 [3.3.3.9] ⓘ Edit
cw-ncs9 [3.3.3.9]

Leaf (s) *
Selected - cw-xrv60 [3.3.3.60] ⓘ Edit
cw-xrv60 [3.3.3.60]

+ Add another

Optimization objective *
Interior gateway protocol (IGP) metric

LFA FRR ⓘ
 Enable Disable

Constraints

Affinity
Select Select or create mapping

+ Add another

Provision Cancel

Step 3 To commit the policy, click **Provision**.

Step 4 Validate the Tree-SID policy creation:

- a. Confirm that the new Tree-SID policy appears in the **Traffic engineering** table. You can also click the check box next to the policy to see it highlighted in the map.

Note The newly provisioned Tree-SID policy may take some time to appear in the **Traffic engineering** table, depending on the network size and performance. The **Traffic engineering** table is refreshed every 30 seconds.

Figure 31: Newly Added Tree-SID Policy on Topology Map

Root n...	Root IP	Name	Tree ID	Label	Admin st...	Oper sta...	Actions
cw-xrv...	3.3.3.50	MY_frd_TRE...	-	16	+	+	...
cw-ncs9	3.3.3.9	tree-n9k	-	18	+	+	...

- b. View and confirm the new Tree-SID policy details. From the Actions column, click and select **View details**.

Figure 32: Tree-SID Policy Details

Tree-SID policy details

Current History

Root cw-ncs9 (Root IP: 3.3.3.9)
TE RID: 3.3.3.9 (IPv4 RID: -)

Name tree-n9k
Tree ID -

Summary

Admin state + Up
Oper status + Up
Label 18
Type Static ?
Programming state None
Metric type IGP
Constraints Exclude-Any -
Include-Any -
[See more](#)

Tree-SID path

Leaf node name	Leaf node IP	Express link
cw-xrv0	3.3.3.60	-

Role	Name	IP	Local IP	Remote IP
Root	cw-ncs9	3.3.3.9	12.1.6.9	12.1.6.60
Leaf	cw-xrv0	3.3.3.60	-	-

Static Tree-SID Policy Configuration Example through Crosswork UI

The output below shows the static Tree-SID policy, configured from Crosswork UI, on the compute SR-PCE.

```
RP/0/RP0/CPU0:cw-xrv56#sh pce lsp p2mp
```

```
Tree: 50-52-54, Root: 3.3.3.50
PCC: 3.3.3.50
Label: 505254
Operational: up Admin: up Compute: Yes
Local LFA FRR: Disabled
Metric Type: IGP
Transition count: 1
Uptime: 00:01:45 (since Thu Apr 27 10:54:49 PDT 2023)
Destinations: 3.3.3.52, 3.3.3.54
```

```

Nodes:
Node[0]: 3.3.3.50 (cw-xrv50)
  Delegation: PCC
  PLSP-ID: 205
  Role: Ingress
  Hops:
    Incoming: 505254 CC-ID: 1
    Outgoing: 505254 CC-ID: 1 (11.1.28.54) [cw-xrv54]
    Outgoing: 505254 CC-ID: 1 (11.1.1.51) [cw-xrv51]
Node[1]: 3.3.3.54 (cw-xrv54)
  Delegation: PCC
  PLSP-ID: 148
  Role: Egress
  Hops:
    Incoming: 505254 CC-ID: 2
Node[2]: 3.3.3.51 (cw-xrv51)
  Delegation: PCC
  PLSP-ID: 187
  Role: Transit
  Hops:
    Incoming: 505254 CC-ID: 3
    Outgoing: 505254 CC-ID: 3 (11.1.2.52) [cw-xrv52]
Node[3]: 3.3.3.52 (cw-xrv52)
  Delegation: PCC
  PLSP-ID: 247
  Role: Egress
  Hops:
    Incoming: 505254 CC-ID: 4

```

The output below shows the same static Tree-SID policy on the High Availability (HA) peer SR-PCE.

```

RP/0/RP0/CPU0:cw-xrv63#sh pce lsp p2mp

Tree: 50-52-54, Root: 3.3.3.50
PCC: 3.3.3.50
Label: 505254
Operational: standby Admin: up Compute: No
Local LFA FRR: Disabled
Metric Type: IGP
Transition count: 0
Destinations: 3.3.3.52, 3.3.3.54
Nodes:
Node[0]: 3.3.3.54 (cw-xrv54)
  Delegation: PCE (3.3.3.56)
  PLSP-ID: 148
  Role: Egress
  Hops:
    Incoming: 505254 CC-ID: 2
Node[1]: 3.3.3.52 (cw-xrv52)
  Delegation: PCE (3.3.3.56)
  PLSP-ID: 247
  Role: Egress
  Hops:
    Incoming: 505254 CC-ID: 4
    Outgoing: 505254 CC-ID: 3 (11.1.2.52)
Node[2]: 3.3.3.51 (cw-xrv51)
  Delegation: PCE (3.3.3.56)
  PLSP-ID: 187
  Role: Transit
  Hops:
    Incoming: 505254 CC-ID: 3
Node[3]: 3.3.3.50 (cw-xrv50)
  Delegation: PCE (3.3.3.56)
  PLSP-ID: 205

```



```

Role: Ingress
Hops:
  Incoming: 505254 CC-ID: 1
  Outgoing: 505254 CC-ID: 1 (11.1.1.28.54)
  Outgoing: 505254 CC-ID: 1 (11.1.1.1.51)

```


Modify a Tree-SID Policy

To modify a Tree-SID policy, do the following:



Note You cannot modify the name, label and root of a Tree-SID policy.

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > Tree-SID** tab.

Step 2 Locate the Tree-SID policy you are interested in and click .

Step 3 Choose **Edit/Delete**.

- Note**
- You can only modify or delete a static Tree-SID policy that is created using the Crosswork UI or API as opposed to one created using SR-PCE CLI.
 - After updating the Tree-SID policy details, you can preview the changes on the map before saving it.

Tree-SID Important Notes

Limitation

- Tree-SID policies are only supported on devices running Cisco IOS XR software.
- PCE high-availability (HA) is supported for static Tree-SID policies configured via the Cisco Crosswork UI, but is not supported if configured directly on the SR-PCE CLI.
- Tree-SID policy details based on SRv6 are not supported.
- If a single instance of SR-PCE is used, and the SR-PCE restarts, all static Tree-SID policies that were configured from the Crosswork UI are deleted.
- IPV4 unnumbered interfaces are not supported.

Visualization of Tree-SID Paths with Missing Nodes

Missing Tree-SID nodes can cause the following issues:

There may be instances where a node on a Tree-SID policy path is not available in the Crosswork topology information. This could happen if the node is not added to the Crosswork device inventory. This affects the display of the Tree-SID policy path on the topology map, causing one or more root-to-leaf paths to appear broken. However, the path details in the right panel will still show the full path.

The screenshot displays a network management interface with a central diagram of a Tree-SID path. The path starts at a root node (xrv9k-VM3-...) and branches into two paths. One path goes through a transit node (xrv9k-VM8) to a bud node (xrv9k-VM7-151). The other path goes through another bud node (xrv9k-VM7-151) to a leaf node (xrv9k-VM8). The interface includes a 'Show' dropdown set to 'Traffic Engineering', 'Device Groups' set to 'All Locations', and a 'Show Groups' button. A 'Saved Views' section is visible at the top right.

SR-PCE Address: 172.23.209.75

Tree-SID path

Leaf Node Name		Leaf Node IP		Collapse All	
<input checked="" type="checkbox"/>	xrv9k-VM11-771-151	192.168.4.14			
	Node	Egress Link			
Role	Name	IP	Local IP	Remote IP	
Root	xrv9k-VM3-...	192.168.4.3	10.0.2.25	10.0.2.26	
Bud	xrv9k-VM5-...	192.168.4.5	20.10.0.14	20.10.0.15	
Transit	xrv9k-VM8	192.168.4.9	20.10.0.17	20.10.0.16	
Bud	xrv9k-VM7-151	192.168.4.6	10.0.3.42	10.0.3.41	
<input checked="" type="checkbox"/>	xrv9k-VM7_3_0_732_cco	192.168.4.7			
	Node	Egress Link			
Role	Name	IP	Local IP	Remote IP	
Root	xrv9k-VM3-...	192.168.4.3	10.0.2.41	10.0.2.42	
Leaf	xrv9k-VM7_...	192.168.4.7	-	-	
<input checked="" type="checkbox"/>	Spitfire	192.168.4.11			
	Node	Egress Link			



PART I

Bandwidth Feature Packs

- [SR Circuit Style Manager \(CSM\)](#), on page 63
- [Local Congestion Mitigation \(LCM\)](#), on page 87
- [Bandwidth on Demand \(BWoD\)](#), on page 109



CHAPTER 7

SR Circuit Style Manager (CSM)

- [Circuit Style SR-TE Important Notes](#), on page 63
- [Workflow for Setting Up CS SR-TE Policy Visualization](#), on page 68
- [Enable SR Circuit Style Manager](#), on page 69
- [Configure Circuit Style SR Policies](#), on page 70
- [Review Circuit Style SR-TE Policy Bandwidth Utilization](#), on page 72
- [View Circuit Style SR-TE Policies](#), on page 73
- [Trigger CSM to Recalculate a Circuit Style SR-TE Policy](#), on page 78
- [What Happens When Bandwidth Reservation Settings are Exceeded?](#), on page 79
- [How Does CSM Handle Path Failures?](#), on page 83

Circuit Style SR-TE Important Notes

This topic outlines the scope of Crosswork's support for Circuit Style SR-TE policies, including requirements and constraints on the policy attribute values set in each Circuit Style SR-TE policy and the processing logic followed during path reversions.



Note Role-based Access Control (RBAC) and task permissions have been introduced in this release. To provision a Circuit Style SR-TE policy, you must have write-access to the head-end device based on Device Access Groups and assigned roles. Only Circuit Style SR-TE admin users can modify Circuit Style SR-TE configuration settings. For more information on RBAC and user roles, see the "[Cisco Crosswork Network Controller Administration Guide](#)".

Policy Attribute Constraints

You set policy attribute values when you create a Circuit Style SR-TE policy using either the device's command line interface or Cisco Crosswork Network Controller. You can also change them later.

The table below describes the requirements for each attribute and how changes affect them. It is important to understand that all the attributes described in the table below act as constraints. Each corresponds to elements of the configuration that Cisco Crosswork uses to govern how Circuit-Style path hops are computed. Each value is effectively a path computation or optimization constraint since it either specifies a required property of a path or excludes possible choices for that path.

Table 2: Circuit Style SR-TE Policy Attribute Values and Constraints

Attribute	Description
Policy Path Protection	The path protection constraint is required for both sides of a Circuit Style SR-TE policy.
Bandwidth Constraint	<p>The bandwidth constraint is required and must be the same on both sides of a Circuit Style SR-TE policy. Bandwidth changes can be made to existing policies with these effects:</p> <ul style="list-style-type: none"> • Once you configure the new bandwidth on both sides, Crosswork will evaluate the path. This <u>will not</u> result in a recomputed path. • If the new bandwidth is higher, Crosswork checks the existing path to ensure sufficient resources. If all currently delegated paths can accommodate the new bandwidth, Crosswork returns the same path with the new bandwidth value, indicating to the path computation client (PCC) that it was successful. If any current paths cannot accommodate the new bandwidth, it returns the old bandwidth value, indicating that it was unsuccessful. This evaluation will only be retried if the bandwidth is changed again. • If the bandwidth is lower, Crosswork returns the same path with the new bandwidth value to indicate to the PCC that it was successful. <p>When you view the policy details, the user interface shows both the requested and reserved bandwidth under each candidate path. These values can differ if the requested bandwidth is increased but there is insufficient available CS pool bandwidth along one or more paths.</p>
Candidate Paths and Roles	<p>The Working path is defined as the highest preference Candidate Path (CP).</p> <p>The Protect path is defined as the CP of the second highest preference.</p> <p>The Restore path is defined as the lowest preference CP. The headend must have <code>backup-ineligible</code> configured.</p> <p>CPs of the same role in each direction must have the same CP preference.</p>
Bi-Directional	<p>All paths must be configured as co-routed.</p> <p>Paths of the same role on both sides must have the same globally unique bi-directional association ID.</p>

Attribute	Description
Disjointness	<p>Working and Protect paths on the same PCC must be configured with a disjointness constraint using the same disjoint association ID and disjointness type.</p> <p>The disjointness association ID for a Working and Protect path pair in one direction must be unique when compared with the corresponding pair in the opposite direction.</p> <p>Only the Node and Link disjoint types are supported. The disjoint type used must be the same in both directions of the same policy.</p> <p>The Restore path must not have a disjointness constraint set.</p> <p>Crosswork follows strict fallback behavior for all Working and Protect path disjointness computations. This means that if node type disjointness is configured but no path is available, Crosswork makes no automatic attempt to compute a less restrictive link type disjoint path.</p>
Metric Type	<p>Only the TE, IGP, Hop count, and Latency metric types are supported. The metric type must match Working, Protect and Restore paths in both directions.</p>
Segment Constraints	<p>All Working, Protect, and Restore paths must have the following segment constraints:</p> <ul style="list-style-type: none"> • protection unprotected-only • adjacency-sid-only <p>To ensure persistency through link failures, configure static adjacency SIDs on all interfaces that might be used by Circuit Style policies.</p>
Unsupported Configurations	<p>The following configurations are not supported:</p> <ul style="list-style-type: none"> • Metric-bounds • SID-Algo constraints • Partial recovery is not supported 7.8.x. • State-sync configuration between PCEs of a high-availability pair. These are not required with Circuit Style SR-TE policies. Use of this feature may result in degraded performance. • Multiple Circuit Style SR-TE policies between the same nodes with the same color but different endpoint IPs.

Attribute	Description
Supported Policy Changes	<p>The following constraints may be modified for an operationally "up" Circuit Style SR-TE policy that has been previously delegated:</p> <ul style="list-style-type: none"> • Metric type • Disjoint type • MSD • Affinities <p>Once configuration changes are consistent across all CPs and both PCCs (for example: the new metric type is the same for all CPs and both sides), Crosswork will initiate a recompute, which can result in new Working, Protect, and Restore paths.</p> <p>During any transitory period in which configurations are not in sync between paths on the same PCC or between PCCs, no path updates are sent to the PCCs.</p>
Unsupported Policy Changes	<p>The following configuration changes to a previously delegated and operationally "up" Circuit Style SR-TE policy are not supported:</p> <ul style="list-style-type: none"> • CP preference • Disjoint Association ID • Bi-directional Association ID <p>To change these configurations for an existing policy, you must first shut down the policy on both sides, make the change (complying with restrictions as detailed above in terms of consistency), and then "no shut" the policy.</p>
Path Computation	<p>Crosswork computes paths for circuit style policies only after a complete bi-directional, path-protected set of candidate paths has been delegated, including Working and Protect paths on both sides. In cases where there is insufficient bandwidth, and a path cannot be found, SR Circuit Style Manager will continue to retry after 30 minutes until a solution is found or if Circuit Style SR-TE is disabled.</p> <p>Crosswork computes the Restore path only after the Working and Protect paths are down. The SR Circuit Style Manager feature pack configuration interface provides a configurable delay timer to control how long to wait after Restore paths are delegated from both sides before computing the path. This delay allows topology and SR policy state changes to fully propagate to Crosswork in cases where these changes triggered the Restore path delegation.</p> <p>Automatic re-optimization is not supported for any paths based on changes in topology, LSP state, or any periodic event. Path computation is supported for Intra/Inter-area/Level and Intra/Inter IGP Domain (same AS). Path computation Inter-AS is not supported.</p>

Attribute	Description
Reversion Behavior	Reversion behavior is controlled by the configuration of the WTR lock timer option under the Protect and Revert paths (it is not relevant for the Working path): <ul style="list-style-type: none"> • No lock configuration: Revert after a default 5-minute lock • Lock with no duration specified: No reversion • Lock duration <value>: Revert after the specified number of seconds

Reversion Logic

Path reversion depends on the initial state of the Working, Protect, and Restore paths and the events affecting each path. The scenarios in the following table provide examples of typical reversion behavior.

Table 3: Path Reversion Scenarios

Initial State	Events	Behavior
Working path is down, Protect path is up/active	Working path comes back up	<ol style="list-style-type: none"> 1. Working path recovers to up/standby state. 2. Each PCC moves the Working path to active after the WTR timer expires. 3. Protect path moves to up/standby.
Working path is down, Protect path is down, Restore path is up/active	Working path comes back up, then Protect path comes back up	<ol style="list-style-type: none"> 1. Working path recovers and goes to up/active state 2. Restore path is removed 3. Protect path recovers and goes to up/standby

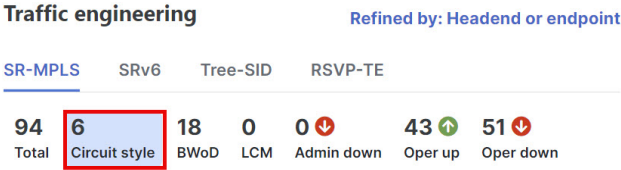
Initial State	Events	Behavior
Working path is down, Protect path is down, Restore path is up/active	Protect path comes back up, then Working path comes back up	<p>On side A: The Working path failure is local (the first Adj SID in the SegList is invalid):</p> <ol style="list-style-type: none"> 1. Protect path recovers and goes to up/active. 2. Restore path is removed. 3. Working path recovers and goes to up/standby. 4. Each PCC moves the Working path to active after the WTR timer expires, Protect path goes to up/standby. <p>On side Z: Working path failure is remote (first Adj SID in SegList is valid):</p> <ol style="list-style-type: none"> 1. Protect path recovers but is not brought up, Restore path remains up/active. 2. Working path recovers and goes up/active. 3. Restore path is removed. 4. Protect path goes to up/standby.

Workflow for Setting Up CS SR-TE Policy Visualization

The following tasks are necessary to start visualizing Circuit Style SR-TE policies in the topology map:

Table 4: Tasks to Complete to Start Visualizing Circuit Style SR-TE Policies

Step	Action
1. Enable the SR Circuit Style Manager (CSM) feature pack.	From the main menu, choose Services & Traffic Engineering > Traffic Engineering > Circuit Style SR-TE > Configuration . Follow the steps in Enable SR Circuit Style Manager, on page 69 .
2. Configure CS SR policies on the devices. Note If you do this step before enabling the Circuit Style SR-TE feature pack, then the CS SR policies will appear operationally down.	You can configure CS SR policies using one of the following methods: <ul style="list-style-type: none"> • Configure CS SR policies manually on the device using the CLI. For more information, see Configure Circuit Style SR Policies, on page 70. • If you are using Crosswork Optimization Engine within Crosswork Network Controller, you can configure CS SR policies using the UI. For more information, see the Cisco Crosswork Network Controller Solution Workflow Guide.

Step	Action
3. Verify that the CS SR policies appear in the SR Policy table.	<p>From the main menu, select Traffic Engineering > Traffic Engineering > SR-MPLS > Circuit style.</p>  <p>The SR Policy table now shows a filtered list containing only CS SR policies.</p>
4. Verify that the reserved bandwidth pool settings you defined in Step 1 are configured properly.	<p>Click on a CS SR node or policy and navigate to the Link Details > Traffic engineering page (see Review Circuit Style SR-TE Policy Bandwidth Utilization, on page 72). From the Circuit style section, view the reserved bandwidth pool size. You can also view current Circuit Style SR-TE bandwidth utilization and how much is still available for use.</p>

Enable SR Circuit Style Manager

To manage and visualize Circuit Style SR-TE policies on the topology map, you must first enable SR Circuit Style Manager (CSM) and set bandwidth reservation settings.

When enabled, CSM computes the best failover bidirectional paths with the requested bandwidth and other constraints defined in the Circuit Style SR policy configuration between two nodes.

Step 1 From the main menu, choose **Traffic Engineering > Circuit Style SR-TE > Configuration**.

Step 2 Toggle the **Enable** switch to **True**.

Step 3 Enter the required bandwidth pool size and threshold information. The following list describes additional field information. See also [What Happens When Bandwidth Reservation Settings are Exceeded?](#), on page 79.

Field	Description
Basic	
Link CS BW Pool Size	The percentage of each link's bandwidth reservable for Circuit Style SR-TE policies.
Link CS BW Min Threshold	The Link CS BW Pool utilization percentage beyond which a threshold crossing event notification will be generated.
Advanced	
Validation Interval	This is the interval that CSM policy will wait before the bandwidth that is reserved for an undelegated policy is returned to the Circuit Style SR-TE policy bandwidth Pool.
Timeout	The duration until which CSM will wait for the delegation request, to generate a notification.

Field	Description
Restore Delegation Delay	The duration until which CSM will pause before processing a restore path delegation.

- Step 4** Click **Commit Changes** to save the configuration. After enabling CSM, you must create Circuit Style SR policy configurations either manually on the device (see [Configure Circuit Style SR Policies, on page 70](#)) or through Cisco Crosswork Network Controller .

Configure Circuit Style SR Policies

A Circuit Style SR policy configuration must include the destination endpoint, the amount of requested bandwidth, and the bidirectional attribute (see [Circuit Style SR-TE Important Notes, on page 63](#) for additional requirements or notable constraints). The configuration should also include a Performance Measurement Liveness (PM) profile. A PM profile enables proper detection of candidate path liveness and effective path protection. PCCs do not validate past the first SID, so without PM, the path protection will not occur if the failure in the Circuit Style SR policy candidate path is not the first hop in the segment list. For more information, see [Configuring SR Policy Liveness Monitoring](#).

This section provides *guidance* on how to manually configure a Circuit Style SR policy and a Performance Measurement Liveness (PM) profile on a device.

- Step 1** If applicable, enable the hardware module on the device for PM configuration.

Example:

```
hw-module profile offload 4
reload location all
```

- Step 2** Configure the PM profile.

Example:

```
performance-measurement
  liveness-profile sr-policy name CS-active-path
  probe
    tx-interval 3300
  !
npu-offload enable    !! Required for hardware Offload only
!
!
  liveness-profile sr-policy name CS-protect-path
  probe
    tx-interval 3300
  !

npu-offload enable    !! Required for hardware Offload only
!
!
!
```

Step 3 Configure the Circuit Style SR policy with the PM profile. All configurations shown in the example are required in order for CSM to manage the Circuit Style SR-TE policy. Entries that are defined by the user are italicized. See [Circuit Style SR-TE Important Notes, on page 63](#) for additional requirements or notable constraints.

Example:

```

segment-routing
 traffic-eng
  policy cs1-cs4

  performance-measurement
    liveness-detection
      liveness-profile backup name CS-protect      !! Name must match liveness profile defined for
Protect path
      liveness-profile name CS-active            !! Name must match liveness profile defined for
Active path
    !
    !
    bandwidth 10000
    color 1000 end-point ipv4 192.168.20.4
    path-protection
    !
    candidate-paths
    preference 10
    dynamic
      pcep
      !
      metric
      type igp
      !
      !
    backup-ineligible
    !

    constraints
      segments
      protection unprotected-only
      adjacency-sid-only
      !
      !
    bidirectional
    co-routed
    association-id 1010
    !
    !
  preference 50
  dynamic
    pcep
    !
    metric
    type igp
    !
    !
  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 igp
  !
!
constraints
  segments
    protection unprotected-only
    adjacency-sid-only
  !
  disjoint-path group-id 3 type node
!
bidirectional
  co-routed
  association-id 1100
!
!
!
!
!
!

```

Review Circuit Style SR-TE Policy Bandwidth Utilization

You can verify that the reserved bandwidth pool settings (defined when enabling CSM, see [Enable SR Circuit Style Manager, on page 69](#)) are correctly configured. You can also view current Circuit Style SR-TE bandwidth utilization and how much is still available for use.



Note There are different ways to navigate to the **Link Details > Traffic Engineering** page from a participating Circuit Style SR-TE node or link. The following procedure assumes you have a Circuit Style SR-TE policy checked in the SR Policy table.

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** and click **Circuit style**. The SR Policy table lists all Circuit Style SR-TE policies.

Step 2 Check the check box next to the Circuit Style SR-TE policy you are interested in.



Step 3 From the topology map, click on a participating Circuit Style SR-TE policy node.

Step 4 From the **Device details** page, click **Links tab > Link_Type _entry > Traffic engineering tab > General**.

Under Circuit Style bandwidth pool, you can see the reserved bandwidth pool size, the amount of bandwidth currently being used, and what bandwidth (allocated to Circuit Style SR-TE policies) is still available.

This example shows the reserved bandwidth pool size as 800 Mbps for NCS-3 and NCS1. The configured settings were earlier defined as 80% for the bandwidth pool size. Since the interface is 1 Gbps, we can confirm that CSM has correctly allocated 80% of the bandwidth for Circuit Style SR-TE policies for these interfaces.

Figure 33: CS SR Policy Bandwidth Pool

Link details  

Summary **Traffic engineering**

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A Side	Z Side
Node	NCS-3	NCS1
IF Name	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/0
FA Affinities		
FA TE Metric		
FA Delay Metric		
FA Topologies	128, 129, 130, 131, 132	128, 129, 130, 131, 132...

Circuit style bandwidth pool

	A Side	Z Side
Pool Size	800 Mbps	800 Mbps
Used	4 Mbps	4 Mbps
Available	796 Mbps	796 Mbps

View Circuit Style SR-TE Policies




View Circuit Style SR-TE policy details such as the endpoints, bandwidth constraints, IGP metrics, and candidate (Working and Protect) paths.

Step 1 From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** and click **Circuit style**.

Figure 34: Select Circuit Style Tab

Traffic engineering Refined by: Headend or endpoint

SR-MPLS SRv6 Tree-SID RSVP-TE

94	6	18	0	0 	43 	51 
Total	Circuit style	BWoD	LCM	Admin down	Oper up	Oper down

The **SR Policy** table lists all Circuit Style SR-TE policies.

Step 2 From the **Actions** column, click  > **View Details** for one of the Circuit Style SR-TE policies.

Note You cannot edit or remove Circuit Style SR-TE policy configurations that have been created directly on the device.

Figure 35: View Circuit Style SR-TE Policy Details

The screenshot shows a network topology map on the left and a 'Traffic engineering' side panel on the right. The side panel displays a table of candidate paths with the following data:

SR-MPLS	SRv6	Tree-SID	RSVP-TE	Headend	Endp...	Color	Admin	Oper St...	Actions
90	6	2	0	0	31	59			
Total	Circuit style	BWoD	LCM	Admin down	Oper up	Oper down			
Selected 4 / Total 6 of 90									
<input checked="" type="checkbox"/>	xrv9k-16	xrv9k-15	11056						View details
<input type="checkbox"/>	xrv9k-15	xrv9k-16	11056						Edit / Delete
<input type="checkbox"/>	xrv9k-16	xrv9k-15	4294...						
<input checked="" type="checkbox"/>	xrv9k-15	xrv9k-16	4294...						
<input checked="" type="checkbox"/>	xrv9k-...	xrv9k-12	5600						
<input checked="" type="checkbox"/>	xrv9k-12	xrv9k-...	5600						

The **Circuit style policy details** window is displayed on the side panel. By default, the candidate path with an "active" state is displayed in the topology map. An active state is designated with a green "A" icon under **State**, indicating it is currently the operational active path. The map also has the **Bi-Dir path** checkbox checked by default, showing the bidirectional paths. The **Candidate path** list displays the candidate path with an active status (path that takes traffic) and other candidate paths.

Figure 36: CS-SR Policy Details Summary

The screenshot shows a network topology map on the left and a 'Circuit style policy details' side panel on the right. The side panel displays a summary of the candidate path configuration, including headend and endpoint information, performance metrics, and a table of candidate paths.

Current | History

Headend xrv9k-25 | TE RID: 192.168.6.6 | PCC IP: 192.168.6.6
Source IP: 192.168.6.6

Endpoint xrv9k-23 | TE RID: 192.168.6.6
Dest IP: 192.168.6.6

Color: 6905

Performance metrics

Traffic rate
0 Mbps avg.

Summary

Candidate path

Path name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.16...	100		Up
<input type="checkbox"/> cfg_srte_c_6905_ep_192.16...	50		Down

Note The Bandwidth Constraint value can differ from the bandwidth you requested if the value is increased and insufficient resources exist to satisfy demand on all Working and Protect candidate paths.

Step 3 View Candidate path configuration details.

- The **Circuit style policy details** window allows you to drill down to view more information about the candidate paths. You can also copy the URL and share this information with others.

The Working path (highest preference path) with an operational state (Oper state) "Up" will always have an active state indicating that it takes traffic (see [How Does CSM Handle Path Failures?](#), on page 83). If the Working path goes down, the Protect path is activated. In this example, the Protect path (with preference 50) is active and displayed on the topology map. Click **Expand all** to view more information about both paths.

Figure 37: Candidate Path on Topology Map

The screenshot displays a network topology map on the left and a 'Circuit style policy details' panel on the right. The map shows a route from Fremont to Mountain View with three candidate paths: a purple link (top preference), a blue link (second preference), and a pink link (third preference). The details panel for the selected path (ID: 50) shows the following configuration:

Path name	Pref	Role	State
cfg_srte_c_6905_ep_192.168.0.23...	100		Down
cfg_srte_c_6905_ep_192.168.0.23_desc_100			Down
Oper state			Down
Metric type			IGP
Bandwidth			Requested: 9,006 Mbps Reserved: 0 Mbps
Bi-Dir association ID			6906
Disjoint group			D: 463 Association source: 0.0.0.0 Type: Node-disjoint
PCE initiated			false
Affinity			Exclude-Any: - Include-Any: - Include-All: -
Segment type			Unprotected
SID algorithm			-
cfg_srte_c_6905_ep_192.168.0.23...	50		Active
cfg_srte_c_6905_ep_192.168.0.23_desc_50			Active
Oper state			Active
Metric type			IGP
Bandwidth			Requested: 9,006 Mbps Reserved: 0 Mbps
Bi-Dir association ID			6906

- Note**
- First preference paths are shown as purple links.
 - Second preference paths are shown as blue links.
 - Third preference paths are shown as pink links.

If the Circuit Style SR-TE policy configuration was done through Cisco Crosswork Network Controller, you have the option to view the Circuit Style SR-TE policy configuration. To see the configuration, click the link next to **Config ID**. For example:

Figure 38: View Candidate Path Details

Circuit style policy details ... | X

Current History

Path name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.168....	100		↑ A ^

Path name cfg_srte_c_6905_ep_192.168.0.25_disc

Oper state ↑ Up | A Active

Metric type IGP

Bandwidth Requested: 9.006 Mbps
Reserved: 0 Mbps

Bi-Dir association ID 5906

Config ID CS-CS-SR-WP-601-head-end-internal

Disjoint group ID: 567
Association source: 0.0.0.0
Type: Node-disjoint

PCE initiated false

Affinity Exclude-Any: -
Include-Any: -
Include-All: -

Segment type Unprotected

SID algorithm -

Here is a sample of a Circuit Style policy configuration. For more information, see [Configure Circuit Style SR Policies](#), on page 70.

Figure 39: Circuit Style Policy Configuration Example



- Step 4** To view the physical path and metrics between endpoints of the selected Circuit Style SR-TE policies, click  to turn applicable metrics on and check the **IGP path** checkbox.

Figure 40: IGP Metrics

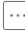
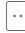


Trigger CSM to Recalculate a Circuit Style SR-TE Policy

Circuit Style SR-TE policies are static in nature, meaning once the paths are computed, they will not be automatically re-optimized based on topology or operational status changes that may affect their paths. You can manually trigger CSM to recalculate a CS-SR policy after the policy's operational status went from down to up or if bandwidth size and requirement changes have been configured.



Note You can only reoptimize a Working and Protect path. It will not work for a Restore path.

- Step 1** From the main menu, choose **Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** and click **Circuit style**. The **SR Policy** table lists all Circuit Style SR-TE policies.
- Step 2** From the **Actions** column, click  > **View Details** for the Circuit Style SR-TE policies you want CSM to recalculate a path for again.
- Step 3** From the top-right corner, click  > **Reoptimize**.

What Happens When Bandwidth Reservation Settings are Exceeded?

CSM discovers and updates the available and reservable bandwidth in the network. CSM maintains an accounting of all bandwidth reservations provided for CS SR policies to ensure that the total reserved bandwidth on all interfaces remains at or below the network-wide resource pool (bandwidth pool size).

This topic provides examples of how CSM handles policies that exceed the bandwidth pool size or bandwidth alarm threshold set on the CSM Configuration page.


Example: Bandwidth Utilization Surpasses Defined Threshold

- Link CS Bandwidth Pool Size: 10%
- Link CS Bandwidth Minimum Threshold: 10%

In this example, the bandwidth pool size for the 10 Gbps ethernet interfaces is 1Gbps and the alarm threshold is set at 100 Mbps (10% of pool size).

1. A Circuit Style SR-TE policy from node 5501-02 to node 5501-01 (r02 - r01) is created with a bandwidth of 100 Mbps.

Figure 41: CS-SR Policy 10 Mbps Up

Link details 

Summary [Traffic engineering](#)

[General](#) SR-MPLS SRv6 Tree-SID RSVP-TE


	A side	Z side
Node	xrv9k-15	xrv9k-16
IF name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA affinities		
FA TE metric		
FA delay metric		
FA topologies	128, 129, 130, 1...	128, 129, 130, 1...

Circuit style bandwidth pool

	A side	Z side
Pool size	100.00 Mbps	100.00 Mbps
Used	0 Mbps	0 Mbps
Available	100.00 Mbps	100.00 Mbps

- Later, the requested bandwidth configured for the policy is increased to 500 Mbps. CSM determines the additional bandwidth along the existing path is available and reserves it.

Figure 42: CS-SR Policy 500 Mbps Up

Link details 

Summary [Traffic engineering](#)

[General](#) SR-MPLS SRv6 Tree-SID RSVP-TE

	A side	Z side
Node	5501-02	5501-01
IF name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA affinities		
FA TE metric		
FA delay metric		
FA topologies	128, 129, 130, 1...	128, 129, 130, 1...

Circuit style bandwidth pool

	A side	Z side
Pool size	1000 Mbps	1000 Mbps
Used	500 Mbps	500 Mbps
Available	500 Mbps	500 Mbps

- Since the bandwidth utilization (500 Mbps) with the updated policy is above the configured pool utilization threshold (100 Mbps), an event is triggered.

Figure 43: Threshold Alerts

Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/21
Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/20
Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/2
Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/0
Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/1/0/1
Optima CSM App	 Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/0/0

Example: Bandwidth Pool Size and Utilization Exceeded

- Link CS Bandwidth Pool Size: 10%
- Link CS Bandwidth Minimum Threshold: 90%

In this example, the bandwidth pool size for the 10 Gbps ethernet interfaces is 1Gbps and the alarm threshold is set for 900 Mbps.

- An existing Circuit Style SR-TE policy from node 5501-02 to node 5501-01 (*r02 - r01*) uses a bandwidth of 500 Mbps.

2. Later, a new policy requiring a bandwidth of 750 Mbps with a path from node 5501-02 to node 5501-01 to 5501-2 (*r02 - r01 - r2*) is requested. The only paths available between these two nodes are the paths computed for the first CS policy.
 - CSM cannot compute a path for the new Circuit Style SR-TE policy *r02 - r01 - r2* and remains operationally down. CSM will try again every 30 minutes to find a path that meets the bandwidth requirements.

Figure 44: CS-SR Policy Exceeds Bandwidth Pool Size

Summary

- Admin State Up
- Oper State Down
- Binding SID 0
- Policy Type Circuit-Style
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True ⓘ
- [See more](#) ▾

Candidate Path

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_100	100		Up ⓘ ▾
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_50	50		Up ⓘ ▾

- Alerts are triggered.

Figure 45: Threshold Alerts

Source	Severity	Description
Optima CSM App	Warning	Unable to compute path for 10.255.255.1 <-> 10.255.255.2 color 2000 due to CsmUpdateStatus.NO_PATH
SR Policy [10.255.255.2#10.255.255.1]	Warning	Policy 'srte_c_2000_ep_10.255.255.1' has operational status as DOWN.
SR Policy [10.255.255.1#10.255.255.2]	Warning	Policy 'srte_c_2000_ep_10.255.255.2' has operational status as DOWN.

3. Later, the Circuit Style SR-TE policy *r02 - r01 - r2* is updated and only requires 10 Mbps. The following behaviors occur:
 - Since the total bandwidth required for the two policies (10 Mbps + 500 Mbps = 510 Mbps) now requires less than the bandwidth pool size (1Gbps), Circuit Style SR-TE policy *r02 - r01 - r2* receives a path computed by CSM and becomes operationally up.

Figure 46: Updated CS-SR Policy Operational

Summary			
Admin State	Up		
Oper State	Down		
Binding SID	0		
Policy Type	Circuit-Style		
Profile ID	-		
Description	-		
Traffic Rate	0 Mbps		
Unused	True		
			See more

Candidate Path			
Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_100	100		Up
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_50	50		Up

- Since the second Circuit Style SR-TE policy with the reduced bandwidth is now provided a path by CSM, alerts are cleared.

Figure 47: Cleared Alerts

Source	Severity	Description
SR Policy [10.10.255.1#10.255.1.2]	Clear	Policy 'srte_c_2000_ep_10.10.255.1.2' has operational status back to UP.
SR Policy [10.10.255.2#10.255.1.1]	Clear	Policy 'srte_c_2000_ep_10.10.255.1.1' has operational status back to UP.

How Does CSM Handle Path Failures?

Cisco Crosswork computes paths for Circuit Style SR-TE policies only after a complete bidirectional, path-protected set of candidate paths has been delegated. Three types of candidate paths are used during path failures:

- **Working**—This candidate path has the highest preference value.
- **Protect**—This candidate path has the second-highest preference value. If the Working path goes down, the Protect path (with the lower preference value) is activated. After the Working path recovers, the Protect path remains active until the default lock duration expires.
- **Restore**—This candidate path has the lowest preference value. Crosswork computes the Restore path only after the Working and Protect paths are down. You can control how long after Restore paths are delegated from both sides to wait before the path is computed (see [Enable SR Circuit Style Manager, on page 69](#)). This delay allows topology and policy state changes to fully propagate to Crosswork in cases where these changes triggered the Restore path delegation.

You can configure Performance Measurement (PM) to address path failures effectively and switch from the Working path to the Protect path. For more information, see [Configure Circuit Style SR Policies, on page 70](#).

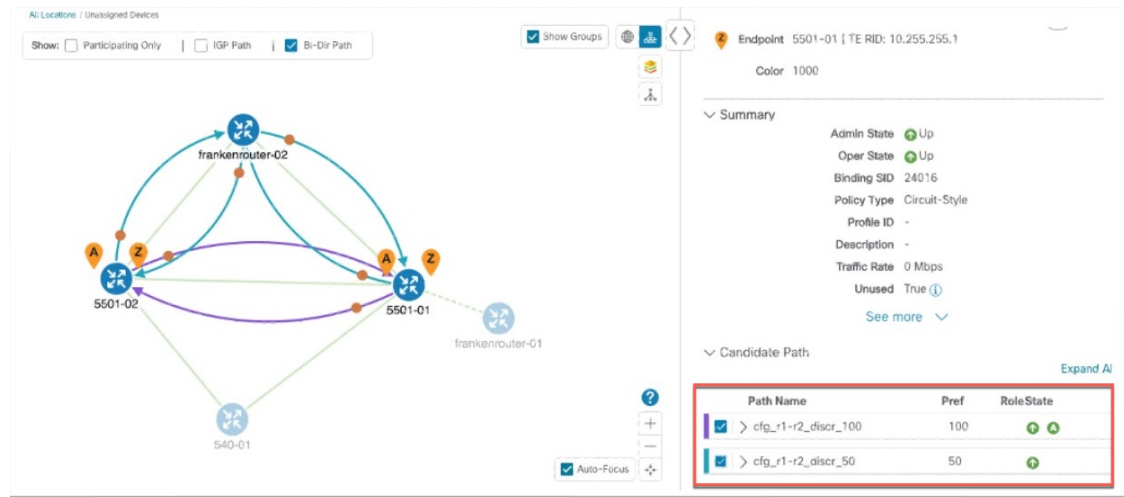
Examples



Note Illustrations are for demonstration purposes only and may not always reflect the exact UI or data described in the workflow content. If you are viewing the HTML version of this guide, click the images to view them in full size.

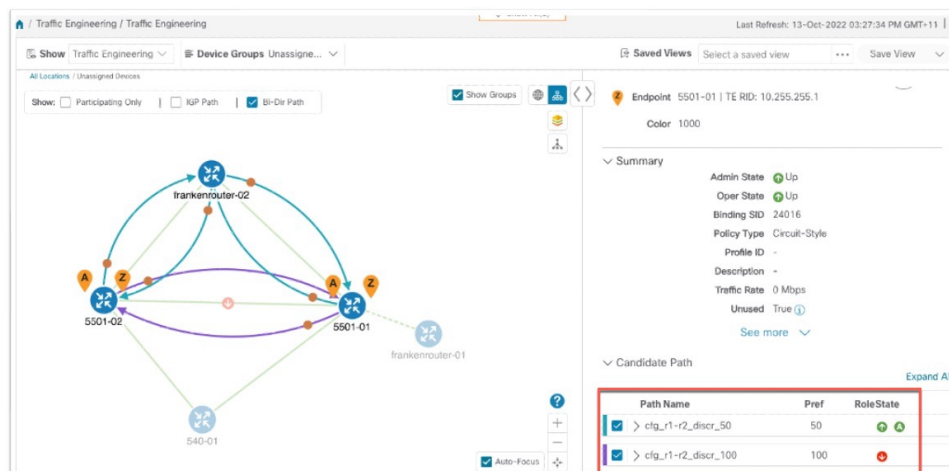
The following image shows that the Working and Protect paths of the Circuit Style SR-TE policy are operational. The *active* path is indicated by the "A" icon.

Figure 48: Initial Candidate Paths



When a Working path having an active status goes down, the Protect path immediately becomes "active." When the Working path recovers, the Protect path moves to up/standby, and the Working path (with preference 100 in the example) becomes active.

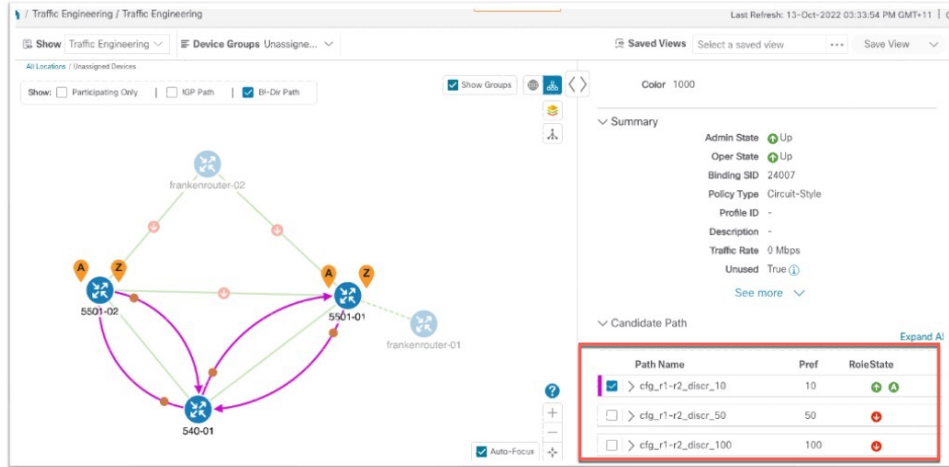
Figure 49: Protected Path Becomes Active



When both the Working and Protect paths go down, CSM calculates a Restore path, which becomes active. The Restore path only appears in this specific scenario. Note that the Restore path has the lowest preference

value of 10 in the example. If the Working or Protected paths become operational again, the Restore path will no longer be visible on the topology map and will be removed from the **Candidate path** list.

Figure 50: Restore Path





CHAPTER 8

Local Congestion Mitigation (LCM)

- [Local Congestion Mitigation Overview, on page 87](#)
- [LCM Congestion Evaluation Requirements, on page 88](#)
- [LCM Congestion Mitigation Requirements, on page 90](#)
- [LCM Important Notes, on page 91](#)
- [LCM Calculation Workflow, on page 93](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 95](#)
- [Configure LCM, on page 100](#)
- [Add Individual Interface Thresholds, on page 102](#)
- [Monitor LCM Operations, on page 104](#)
- [Temporarily Exclude an Interface from LCM, on page 106](#)

Local Congestion Mitigation Overview

Local Congestion Mitigation (LCM) searches for congestion on a configurable cadence (as opposed to a triggered event). It provides localized mitigation recommendations in surrounding interfaces (local interface-level optimization) within a domain. LCM computes the shortest paths for one or more tactical policies to divert minimal traffic on a congested interface to alternate paths with sufficient bandwidth. It attempts to keep as much of the traffic on the original IGP path. With LCM, you can do the following:

- Monitor congestion as defined by the interface thresholds you specify.
- Visually preview LCM recommendations on your network before you decide whether to commit the Tactical Traffic Engineering (TTE) SR policy deployment.
- Enable LCM to delete any down, failed, or uncommitted LCM TTE policies when there is an imminent risk of network failures based on LCM solution configurations. For more information, see the advanced configuration options (**Auto Repair Solution** and **Adjacency Hop Type**) in [Configure LCM, on page 100](#).

LCM enables a wider application of the solution in different network topologies, including those with multiple IGP areas, because of its simpler path computation and limitation to specific network elements. Focusing on the problem locally within a domain eliminates the need to simulate edge-to-edge traffic flows in the network through a full traffic matrix. It allows for better scaling of large networks. Also, LCM collects TTE SR policy and interface counters via SNMP and does not require SR-TM.



Note Take a look at the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 95](#) to see how to use LCM in your network.

LCM Congestion Evaluation Requirements

For LCM to properly evaluate congestion, LCM requires traffic statistics from interface and headend SR-TE policy traffic measurements.

To ensure LCM is receiving these traffic statistics,

- enable SNMP or gNMI on the devices whose traffic you want to monitor, including the headend device. For more information on how to configure these protocols, see the specific device platform configuration guide, for example, [Configuring SNMP Support](#)
- confirm that all the devices are [reachable](#) from the Crosswork Data Gateway, and
- enable strict SID labels on all devices within an LCM domain.

For device configuration and examples, see

Enable Strict SID for LCM Usage

All devices in the LCM domain must have strict SID enabled.

Step 1 Enable strict SID labels for all devices in the LCM domain.

Example:

Cisco IOS XR with ISIS:

```
router isis core
interface Loopback0
  address-family ipv4 unicast
    prefix-sid absolute 16003
    prefix-sid strict-spf absolute 16503
  !
  address-family ipv6 unicast
  !
!
```

Example:

Cisco IOS XR with OSPF:

```
router ospf 100
area 0
  mpls traffic-eng
  segment-routing mpls
  interface Loopback0
    passive enable
    prefix-sid absolute 16002
    prefix-sid strict-spf absolute 16502
  !
```

Example:

Cisco IOS XE:

```
segment-routing mpls
!
connected-prefix-sid-map
  address-family ipv4
    <ipv4-address> absolute 16010 range 1
  exit-address-family
  address-family ipv4 strict-spf
    <ipv4-address> absolute 16510 range 1
  exit-address-family
!
!
```

Step 2 Enable segment routing on the headend device and confirm that *all* devices

- are either using the same default segment routing global block (SRGB) range or a custom range you specify, and
- have the maximum SID depth explicitly configured *if* there are devices along the path that impose restrictions on the label stack depth.

Example:

```
segment-routing
global-block 16000 80000
traffic-eng
  maximum-sid-depth 8
```

Step 3 If there are existing SR policies, the headend device must be configured to use strict SPF SID labels.

Example:

For PCC-initiated or computed SR policies:

```
segment-routing
traffic-eng
  policy srte_c_8000_ep
    color 8000 end-point ipv4 <ipv4-address>
    candidate-paths
      preference 100
      dynamic
        metric
          type igp
    !
  !
  constraints
    segments
      sid-algorithm 1
```

Example:

For PCE computed or delegated SR policies:

```
policy srte_c_8001_ep_198.19.1.4
  color 8001 end-point ipv4 198.19.1.4
  candidate-paths
    preference 100
  dynamic
    pcep
  !
  metric
    type igp
```

This SR-PCE configuration returns paths with strict SID only. For example:

```
pce
segment-routing
strict-sid-only
```

LCM Congestion Mitigation Requirements

For LCM to correctly calculate and mitigate congestion, the headend device must support autoroute steering and Equal Cost Multi-Path (ECMP).

Autoroute Steering

The headend device must support PCE-initiated SR-TE policies with autoroute steering. However, LCM will not work if the headend is a Cisco NCS device and there is L2VPN traffic in the network.

Devices should be configured with `include ipv4 all` and `force-sr-include` to enable traffic steering into SR-TE policies with autoroute.

For example:

```
segment-routing
traffic-eng
pcc
  profile 10      !!The profile ID must match the value in the UI LCM Configuration > Basic
> Profile ID
  autoroute
    include ipv4 all
    force-sr-include
```

The `ID` parameter in this command identifies the PCC profile associated with the SR-TE policy that PCE has provisioned. The ID value can be any integer from 1 to 65535, but it must match the profile ID that PCE uses to instantiate the policy. If not, the policy will not be activated. For example, if PCE provisions a policy with profile ID 10, you must configure `segment-routing traffic-eng pcc profile 10 autoroute force-sr-include` on the headend router to enable autoroute announcement for that policy. See the specific device platform configuration guide, for example, [Segment Routing Configuration Guide, Cisco IOS XE 17 \(Cisco ASR 920 Series\)](#), [COE-PCE Initiated SR Policy with OSPF and IS-IS SR-TE Autoroute Announce](#).



Note The ID that is configured under the PCC profile, must match the Profile ID option set on the LCM Configuration page.

Equal Cost Multi-Path (ECMP)

The headend device must support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies. To verify that a device can support SR-TE policies using ECMP, check that the device has the following:

- Segment Routing is enabled and configured with a Segment Routing Global Block (SRGB) that matches the SRGB of the SR-TE policy headend and tailend routers. Use the `show segment-routing mpls state` command to verify the SRGB configuration on the device.
- BGP-LS is enabled and configured to advertise and receive link-state information from the SR-TE policy headend and tailend routers. Use the `show bgp link-state link-state` command to verify the BGP-LS

status and the `show bgp link-state link-state database` command to verify the link-state information on the device.

- ECMP is enabled and configured to load-balance traffic across multiple equal-cost paths based on flows. Use the `show ip route` command to verify the ECMP routes and the `show ip cef` command to verify the ECMP load-balancing algorithm on the device.

LCM Important Notes

Consider the following information when using LCM:

- You must have the Advanced RTM license package to use LCM.
- User roles must be granted with LCM task permissions for a domain in order to configure LCM and commit LCM recommendations. For more information on RBAC and user roles, see the "[Cisco Crosswork Network Controller Administration Guide](#)".
- Device Access Group (DAG) access is *not* supported by LCM. Users that have been granted with LCM task permissions in a domain are able to configure and commit LCM recommendations regardless of whether or not they have DAG access for any devices in that domain.
- LCM does not support LDP-labeled traffic. LDP-labeled traffic *must not* be steered into LCM autoroute TTE SR policies.
- The use of LCM is not recommended on networks with Tree SID policies. Initial calculations are skewed because full traffic measurements are unavailable.
- LCM supports domains with up to 2000 devices. A *domain* is an identifier that is assigned to an IGP process. Domains are learned from the network. The domain ID is taken from the PCC router configuration (`link-state instance-id`) that you use to advertise IGP with BGP-LS.
- LCM recommended solutions use the resources within a single domain only.
- LCM evaluates network utilization on a regular, configurable cadence of 1 minute or more. The cadence is typically set to be greater than or equal to the SNMP traffic polling interval but can be set lower to improve responsiveness. The default cadence is 10 minutes.
- The traffic statistics collection interval affects how quickly LCM can respond to topology changes and LSP deployments that affect interface and LSP traffic measurements. It can take up to twice the traffic statistics collection interval plus the LCM evaluation interval for LCM recommendations to fully reflect these changes. During this period, LCM recommendations may evolve as the traffic measurements are updated and eventually fully converge in Crosswork.
- LCM leverages ECMP across parallel TTE SR policies and assumes roughly equal splitting of traffic. The degree to which actual ECMP splitting adheres to this assumption depends on the presence of large elephant flows and the level of traffic aggregation.

You can configure LCM to detect excessive uneven ECMP splitting among parallel TTE SR policies and issue an event to notify. To mitigate the effects of uneven ECMP, the overprovisioning factor is used in LCM. For more information, see [Configure LCM](#).
- LCM assumes traffic in an *existing* SR-TE policy is ineligible for optimization and should not be steered into LCM TTE SR policies. To enforce this assumption, any existing non-LCM SR-TE policies should

not use regular Algo-0 prefix SIDs. Any combination of Algo-1 Strict, Flexible Algorithm, or adjacency SIDs is recommended to prevent this traffic from being steered into LCM TTE SR policies.

- When domain interfaces and links are removed (intentionally or unintentionally), the following occurs:
 - If all links in a domain go down (LINK_DOWN state), LCM configuration and the Domain UI card (see [Configure LCM, on page 100](#)) will remain available until the links are aged out (after 4 hours). This behavior is intentional as it gives you time to recover domain interfaces and links if this was done by mistake.
 - If you want to force domain removal before links age out, then you can remove links manually from the UI. The domain will remain in a "ready for deletion" status until the last link is removed.
- After an HA switchover, you can manually add missing interfaces that were previously monitored or update domain configuration options after the system is stable. Missing interfaces and other configuration options occur if they were added after the last cluster data synchronization.

BGP-LS Speaker Placement for Multiple AS Networks with a Dedicated IGP Instance Between ASBRs

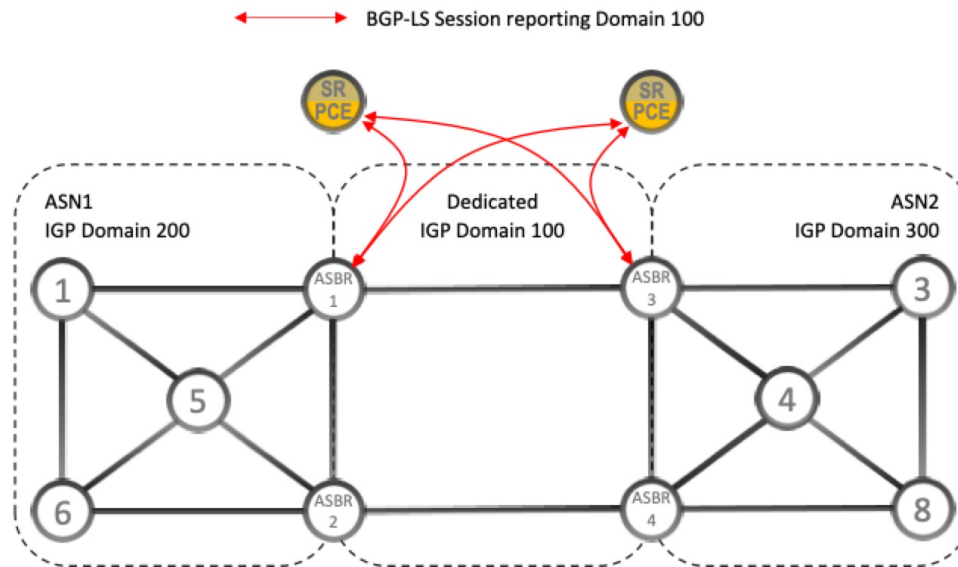
To support interdomain latency-optimized SR policy path computation by an SR-PCE (or other use cases where egress peer engineering (EPE) is not supported), a dedicated IGP instance may be configured between autonomous system border routers (ASBRs) in different ASNs. In these cases, it is important to identify which ASBRs report the topology via BGP-LS for proper topology discovery.

In the following example, at least one ASBR in each AS participating in the dedicated inter-AS IGP (Domain 100) must have BGP-LS enabled to report the IGP between each ASBR. Each ASBR must report the domain with the same BGP-LS identifier.



Note More than one ASBR per AS reporting the BGP-LS topology is also supported.

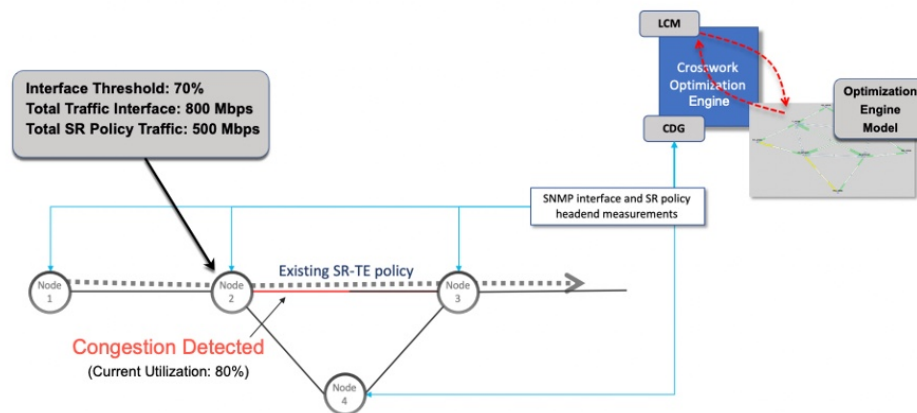
Figure 51: BGP-LS Session Reporting Domain 100



LCM Calculation Workflow

This example walks you from congestion detection to the calculations LCM performs before recommending tactical tunnel deployment. These calculations are done on a per-domain basis, allowing better scalability and faster calculation for larger networks.

Figure 52: LCM Configuration Workflow Example



- Step 1** LCM first analyzes the Optimization Engine Model (a realtime topology and traffic representation of the physical network) on a regular cadence.
- Step 2** In this example, LCM detects congestion after a congestion check interval when Node 2 utilization goes above the 70% utilization threshold.
- Step 3** LCM calculates how much traffic is eligible to divert.

LCM only diverts traffic that is not already routed on an existing SR policy or RSVP-TE tunnel (for example, unlabeled, IGP routed, or carried via FlexAlgo-0 SIDs). Traffic within an SR-TE policy will not be included in the LCM calculation and will continue to travel over the original programmed path.

Eligible traffic is computed by taking the interface traffic statistics that account for all traffic on the interface and subtracting the sum of traffic statistics for all SR-TE policies that flow over the interface.

Total interface traffic – SR policy traffic and RSVP-TE tunnels = Eligible traffic that can be optimized

This process must account for any ECMP splitting of SR policies to ensure the proper accounting of SR policy traffic. In this example, the total traffic on congested Node 2 is 800 Mbps, and the total traffic of all SR policies routed over Node 2 is 500 Mbps.

The total traffic that LCM can divert in this example is 300 Mbps: 800 Mbps – 500 Mbps = 300 Mbps

Step 4 LCM calculates the amount that must be sent over alternate paths by subtracting the threshold equivalent traffic from the total traffic on the interface. In this example, the amount to be diverted is 100Mbps:

800 Mbps – 700 Mbps (70% threshold) = 100 Mbps

LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see [Configure LCM, on page 100](#).

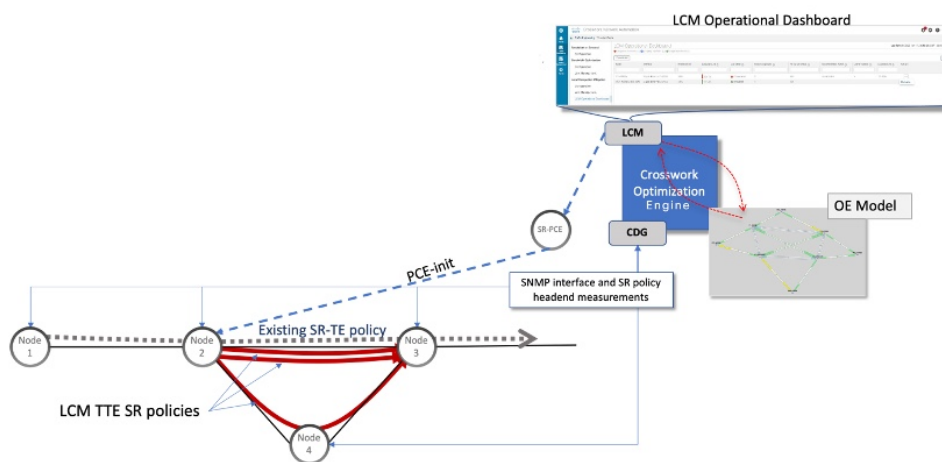
Step 5 LCM determines how many TTE SR policies are needed and their paths. The ratio of how much LCM-eligible traffic can stay on the shortest path to the amount that must be detoured will determine the number of TTE SR policies needed on the shortest versus alternate paths, respectively.

In this example, LCM must divert one-third of the total eligible traffic (100Mbps out of 300Mbps) away from the congested link. Assuming a perfect ECMP, LCM estimates three tactical SR-TE policies are required to create this traffic split: one tactical SR-TE policy will take the diversion path and two tactical SR-TE policies will take the original path. There is sufficient capacity in the path between Node 2 and Node 4. Therefore, LCM recommends three TTE SR policies (each expected to route approximately 100Mbps) to be deployed from Node 2 to Node 3 via SR-PCE:

- 2 TTE SR policies to take a direct path to Node 3 (200 Mbps)
- 1 TTE SR policy takes hop via Node 4 (100 Mbps)

These recommendations will be listed in the **LCM Operational Dashboard**.

Figure 53: LCM Recommendation Example



Step 6 Assuming you deploy these TTE SR policies, LCM continues to monitor the deployed TTE policies and will recommend modifications or deletions as needed in the LCM **Operational Dashboard**. LCM recommends deleting deployed TTE SR policies if the mitigated interface will remain uncongested after they are removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Workflow Example: Mitigate Congestion on Local Interfaces



Note If you are viewing the HTML version of this guide, click on the images to view them in full size.

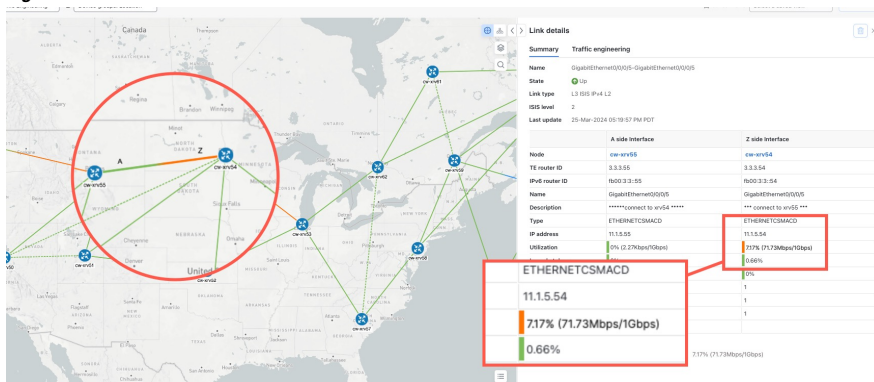
In this example, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE SR policies when utilization on a device's interface surpasses a defined utilization threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion. The example goes through the following steps:

1. View uncongested topology.
2. Set utilization thresholds for individual interfaces.
3. Enable and configure LCM in manual mode. Manual mode allows you to view recommended TTE policies prior and decide whether or not to deploy them.
4. After LCM detects congestion, view LCM recommendations on the Operational dashboard.
5. Preview the recommended LCM TTE policies to deploy visually on the topology map.
6. Commit and deploy all LCM TTE policy recommendations to mitigate the congestion.
7. Verify that the LCM TTE policies have been deployed.

Step 1 View initial topology and utilization prior to LCM configuration.

- a) In this example, note that the node cw-xrv54 has a utilization of 7.17%.

Figure 54: Initial Utilization



Step 2 Define any individual interface thresholds.

LCM allows you to configure a *global* utilization threshold that can be used for all interfaces. When traffic utilization surpasses the threshold, LCM will try to find bypass policies to remediate the congestion. You set the global utilization threshold on the **LCM Configuration** page. However, if you want to define different thresholds for individual interfaces, we recommend defining them on the **Customized interface threshold** page *before* enabling LCM.

- a) In this example, we will define an individual interface threshold. Go to the **Customized interface thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-Identifier > ... > Interface thresholds**). You can add interfaces individually or upload a CSV file with a list of nodes and interfaces with custom utilization thresholds. For more information, see [Add Individual Interface Thresholds, on page 102](#).





See the following example and note the defined threshold for cw-xrv54 with interface GigabitEthernet0/0/0/1 is 20%.



Note The utilization thresholds in this example are extremely low and best used for lab environments.

Figure 55: Customized Interface Thresholds

Customized interface thresholds

Interfaces to monitor: Selected interfaces - LCM monitors only the interfaces with custom thresholds.

+ Create   | Edit mode: off Total 0  

Node	Interface	Threshold (%)	Select for deletion 
<input type="text"/>	<input type="text"/>	<input type="text"/>	
<input checked="" type="checkbox"/> cw-xrv54	GigabitEthernet0/0/0/5	20	

Note By default, LCM monitors all interfaces. This includes any individual thresholds that are imported to this page. The rest of the interfaces will be monitored using the global **Utilization threshold** defined on the **LCM Configuration** page.

- b) After adding interfaces and defining thresholds, click **Save**.

Step 3 Enable LCM and configure the global utilization thresholds.

- a) From the main menu, choose **Traffic Engineering > Local Congestion Mitigation > Domain-Identifier** and click **Configuration**. Toggle the **Enable** switch to **True** and configure other LCM options. In this example, the global threshold is set at 80%, and the **Interfaces to monitor > All interfaces** option is selected.

Figure 56: LCM Configuration Page

Configuration

Basic Advanced

Enable True
False True

Color *
Range: 1 to 4294967294

Utilization threshold * %
Range: 0 to 100

Utilization hold margin * %
Range: 0 to Utilization threshold

Delete tactical SR policies when disabled True
False True

Profile ID *
Range: 0 to 65534

Congestion check interval * seconds
Range: 60 to 86400 seconds

Max LCM policies per set *
Range: 1 to 8

Interfaces to monitor Selected interfaces All interfaces

Description

- b) Click **Commit changes** to save your configuration. After committing the configuration changes, LCM will display *recommendations* on the **LCM Operational Dashboard** if congestion occurs on any monitored interfaces. Later, you will be able to preview the recommended TTE policies and decide whether or not to commit and deploy them onto your network.

Step 4 After some time, congestion occurs, surpassing the custom LCM threshold defined at 20% for node cw-xrv54 with interface GigabitEthernet0/0/0/5.

Figure 57: Observed Congestion

Link details

Summary **Traffic engineering**

Name GigabitEthernet0/0/0/5-GigabitEthernet0/0/0/5

State Up

Link type L3 ISIS IPv4 L2

ISIS level 2

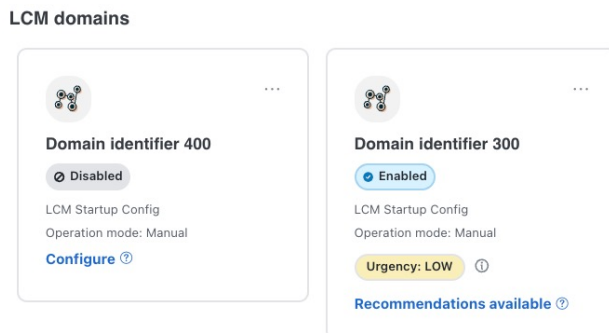
Last update 25-Mar-2024 05:19:57 PM PDT

	A side interface	Z side interface
Node	cw-xrv55	cw-xrv54
TE router ID	3.3.3.55	3.3.3.54
IPv6 router ID	fb00:3:3::55	fb00:3:3::54
Name	GigabitEthernet0/0/0/5	GigabitEthernet0/0/0/5
Description	*****connect to xrv54 *****	*** connect to xrv55 ***
Type	ETHERNETCSMACD	ETHERNETCSMACD
IP address	11.1.5.55	11.1.5.54
Utilization	0% (2.25Kbps/1Gbps)	28.5% (285Mbps/1Gbps)
In packet drops	0%	0.66%
In packet errors	0%	0%
IGP metric	1	1
Delay metric	1	1
TE metric	1	1
Admin groups		

Step 5 View TTE SR policy recommendations in the LCM Operational Dashboard.

- a) Navigate to **Traffic Engineering > Local Congestion Mitigation**. When congestion is detected, the domain displays the urgency type and recommendations that are available. Click the question mark icons to display more information about the urgency type and when the most recent recommendation was given.

Figure 58: Congested Detected and LCM Recommendations



- b) (Optional) View LCM events.

From the top-right corner of the Crosswork UI, click > **Events** tab to view LCM events. You can also monitor this window to view LCM events as they occur. You should see events for LCM recommendations, commit actions, and any exceptions.

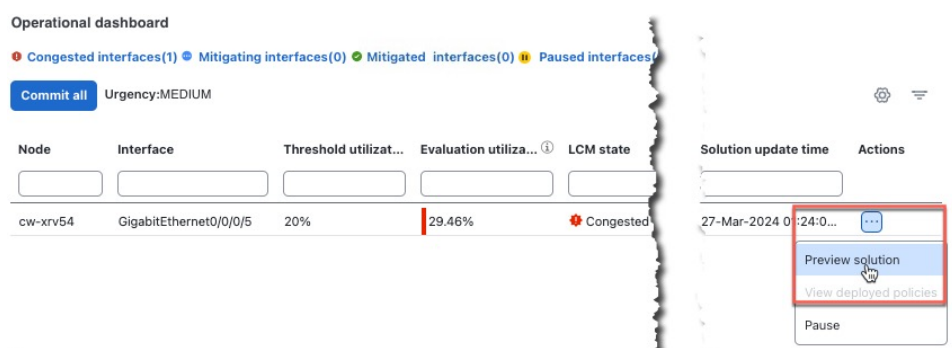
- c) Open the **Operational dashboard (Traffic Engineering > Local Congestion Mitigation > Domain-Identifier > ... > Operational dashboard)**.

The dashboard shows that cw-xrv54 utilization has surpassed 20% and is now at 29.46%. In the **Recommended action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended action - Create set**) to address the congestion on the interface. For more information, see [Monitor LCM Operations, on page 104](#).

Note If LCM cannot find a solution (**Recommended action - No solution**), it may be due to constraints enabled when configuring LCM ([Configure LCM, on page 100](#)).

- d) Before committing TTE policies, you can preview the deployment of each TTE policy solution set. Click in the **Actions** column and choose **Preview solution**.

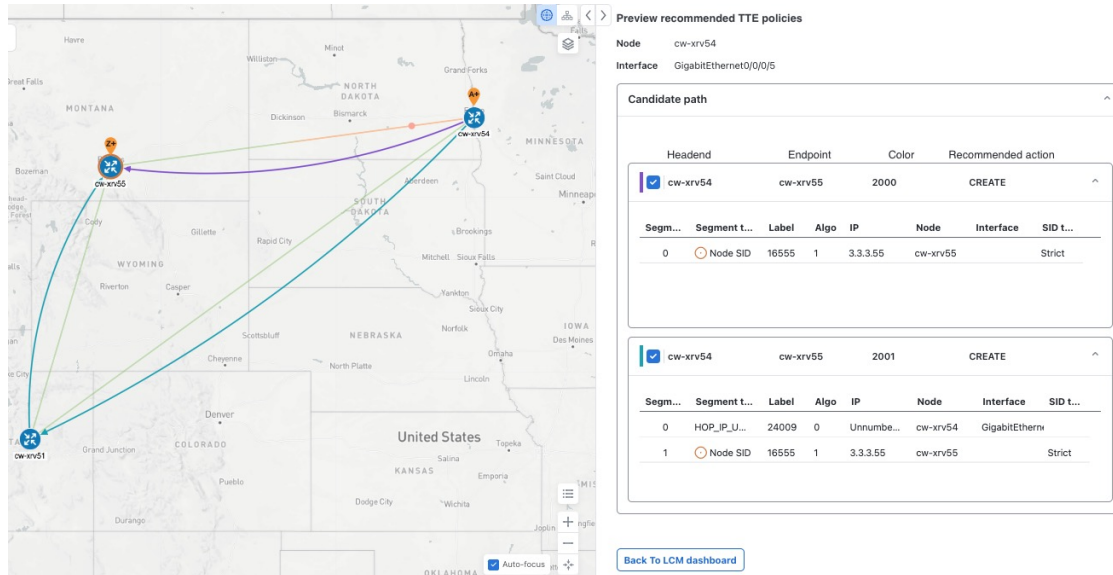
Figure 59: Preview Solution



The resulting window displays the node, interface, and the recommended action for each TTE policy. From the **Preview** window, you can select the individual TTE policies and view different aspects and information as you would normally on the topology map. You can expand each policy to view individual segments. After reviewing the potential implications on your network, you can decide whether or not to deploy the bypass policies that LCM recommends.

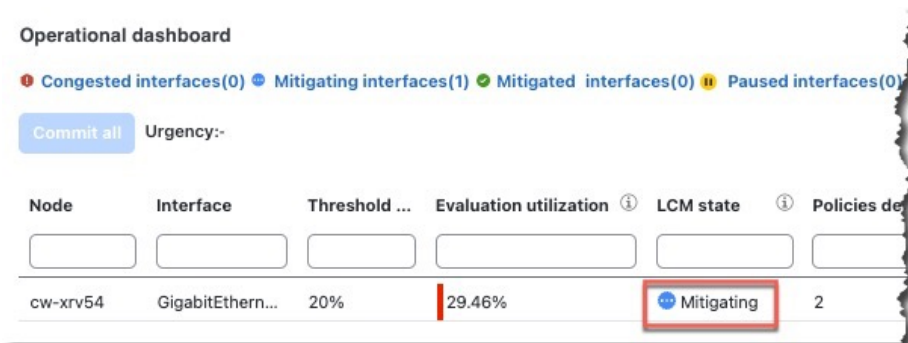
The following figure shows the recommended TTE policies for node cw-xrv54.

Figure 60: LCM TTE Deployment Preview



- e) After you are done viewing the recommended TTE policies on the map, go back to the **Operational dashboard** and click **Commit all**. The **LCM state** column changes to **Mitigating**.

Figure 61: Mitigating State



Note All LCM recommendations per domain must be committed to mitigate congestion and produce the expected utilization as shown in the **Operational dashboard**. The mitigating solution is based on *all* LCM recommendations being committed because of dependencies between solution sets.

Step 6 Validate TTE SR policy deployments.


- a) Click > **Events** tab. Note which LCM events are listed in the **Events** window.

Note Crosswork will report network events detected based on the policies and features you have enabled. For example, if a link drop causes an SR-TE policy to go down or if LCM detects congestion an event is displayed. These alerts are reported in the UI and, if desired, can be forwarded to third-party alerting/monitoring tools.

- b) Return to the **Operational dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

Note The LCM state change will take up to 2 times longer than the SNMP cadence.

- c) Confirm the TTE policy deployment by viewing the topology map.

Click  in the **Actions** column and choose **View deployed policies**. The deployed policies are displayed in focus within the topology map.

Step 7 Remove the TTE SR policies based on the LCM recommendation.

- After some time, the deployed TTE SR policies may no longer be needed. This occurs if the utilization continues to stay under the threshold without the LCM-initiated TTE tunnels. If this is the case, LCM generates new recommended actions to delete the TTE SR policy sets.
- Click **Commit all** to remove the previously deployed TTE SR policies.
- Confirm the removal by viewing the topology map and SR Policy table.

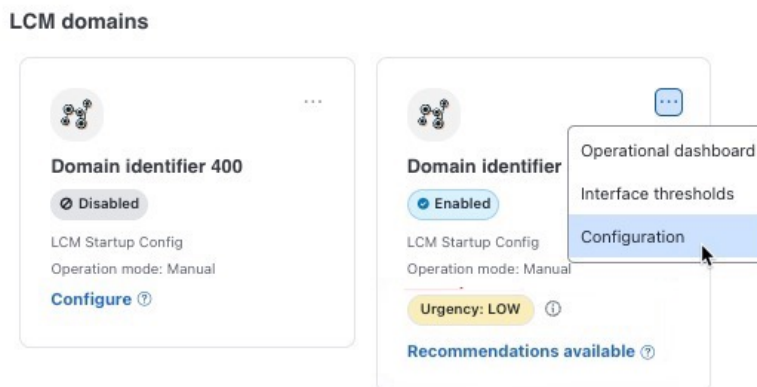
In this scenario, we observed how to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands, but at the same time, gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations and see how the potential deployment will take effect in your network before you deploy them. As traffic changes, LCM tracks the deployed TTE SR-TE policies and decides whether or not they are still needed. If not, LCM recommends deleting them.

Configure LCM


To enable and configure LCM:

Step 1 From the main menu, choose **Traffic Engineering > Local Congestion Mitigation > Domain-ID-card** and click *****> Configuration**.

Figure 62: LCM Configuration



Step 2 Toggle the **Enable** switch to **True**.

Step 3 Enter the required information. Hover the mouse pointer over  to view a description of each field.

Note If LCM is enabled, but cannot find a solution (**Recommended action - No solution**), it may be due to constraints enabled on this page.

The following list describes additional **Basic** field information not described in hover text:

- **Utilization threshold**—Set the utilization percent at which LCM will consider an interface to be congested. This value applies to all interfaces unless you specify thresholds to individual interfaces on the **Customized interface thresholds** page.
- **Profile ID**—This is a required configuration to enable traffic steering onto LCM policies. Autoroute (steers traffic into the tactical SR-TE policies LCM creates) is applied to SR-TE policies through the proper **Profile ID** option that is set here to align with the configuration on the PCC associating that Profile ID with autoroute feature.
- **Congestion check interval** (seconds)—This value determines the interval at which LCM will evaluate the network for congestion. Under a steady state, when there are no recommendation commits, it uses this interval to re-evaluate the network to determine if changes are required. For example, if the interval is set to 600 seconds (10 minutes), LCM will evaluate the network every 10 minutes for new congestion and determine whether a new recommendation or modifications to existing recommendations are needed. Examples of modifications can include removal or updates to individual policies that were previously recommended. This option is typically set to greater than or equal to the SNMP polling cadence but can be set as low as 60 sec to improve responsiveness within the bounds imposed by the traffic collection interval.
- **Interfaces to monitor**—By default, this is set to **Selected interfaces**, and you will need to add thresholds to individual interfaces by importing a CSV file on the **Customized interface thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-identifier > ... > Interface thresholds**). Only interfaces defined on the **Customized interface thresholds** page will be monitored. If set to **All interfaces**, LCM will monitor the interfaces with custom thresholds that are uploaded on the **Customized interface thresholds** page and the rest of the interfaces using the **Utilization threshold** value configured on this page.

The following list describes additional **Advanced** field information not described in hover text:

- **Advanced > Congestion check suspension interval** (seconds)—This interval determines the time to wait (after a **Commit all** is performed) before resuming congestion detection and mitigation. Since this interval should allow time for network model convergence, set the interval to no less than twice the SNMP collection cadence.
- **Advanced > Auto repair solution**—If set to **True**, LCM will automatically delete any down, failed, or uncommitted LCM TTE policies. This option is mainly to address a failure in a policy.

If this option is disabled, and the **Urgency** status of the recommendation shown in the LCM Operational Dashboard is **High**, then the recommended solution is a candidate for the **Auto repair solution**. This means that a network failure will most likely occur if the solution is not deployed.

- **Advanced > Adjacency hop type**—If set to **Protected**, LCM will create SR policies using protected adjacency SIDs. This allows for Topology-Independent Loop-Free Alternate (TI-LFA) to compute a path for any adjacency failures.
- Note** This option should only be set to **Protected** if all nodes in the same IGP area as LCM is operating are strict SPF SID capable.
- **Advanced > Optimization objective**—LCM calculates tactical SR policies based on the metric type chosen to minimize.
 - **Advanced > Deployment timeout**—Enter the maximum number of seconds allowed to confirm deployment of tactical SR policies.

- **Advanced > Over-provisioning factor (OPF)**—This option helps address unequal ECMP traffic distribution (elephant flows). This value determines the percentage of extra traffic that should be accounted for when computing a path for a by-pass policy. If LCM needs to divert x amount of traffic due to congestion, then it will search for a path that can support $x * (1 + OPF)$ traffic. For more information, see [LCM Calculation Workflow, on page 93](#). The default value is 0.
- **Advanced > Maximum segment hops**—When calculating bypass TTE policies, LCM uses the effective Maximum SID Depth (MSD) value (as entered here) for specified device tags. You can assign up to five device tags with specific MSD values.

Note A 0 value will not result in a solution. Setting a 0 value is equivalent to LCM monitoring and indicating when there is congestion in the network without providing a recommendation.

Crosswork learns from SR-PCE the MSD for each platform advertising the hardware limit in the IGP and BGP-LS. It represents the hardware limit that can be imposed exclusive of any service/transport/special labels. Therefore, you may want to use this new option to assign less than the advertised MSD value that LCM can use for bypass TTE policy calculation. To view the MSD value for a device, navigate to the **Traffic Engineering** topology map and click on the device. From the **Device details** page, click **SR-MPLS > Prefixes** tab > **Expand all**.

Note Prior to using this option, you must create device tag groups to which you want to assign certain MSD values. For information on creating tags and assigning them to devices, see the [Cisco Crosswork Network Controller Administration Guide](#)

Step 4 To save your configuration, click **Commit changes**. If congestion occurs on any monitored interfaces, LCM will display *recommendations* (LCM will *not* automatically commit or deploy new TTE policies) on the **LCM Operational dashboard**. You can then preview the recommended TTE policies and decide whether or not to commit and deploy them onto your network.

Add Individual Interface Thresholds

Networks have many different links (10G, 40G, 100G) that require different thresholds to be set. The **Customized interface thresholds** page allows you to manage and assign individual thresholds to nodes and interfaces.

Figure 63: Customized Interface Thresholds

Customized interface thresholds

1 → Interfaces to monitor: Selected interfaces - LCM monitors only the interfaces with custom thresholds.

2 →


3 → + Create

4 → Edit mode: off

5 →

Node ↑	Interface	Threshold (%)	Select for deletion
F1.cisco.com	GigabitEthernet0/0/0/2	70	
F3.cisco.com	GigabitEthernet0/0/0/1	25	

6 → Total 0

Callout No.	Description
1	Interfaces to monitor: Displays the option that is currently configured on the Configure LCM page.
2	Import CSV file: All interfaces currently in the table will be replaced with the data in the CSV file you import. Export CSV file: All interfaces are exported to a CSV file. You cannot filter data for export.
3	+ Create: Click this button to add new interface threshold rows.
4	Edit mode: When Edit mode is ON , you can edit multiple fields in one session, then click Save .
5	Filter: By default, this row is available for you to enter text in which to filter content.
6	Select for deletion: Click  to delete the row. When Edit mode is ON , you can check multiple rows to delete, then click Save .

To assign specific threshold values for individual interfaces when using LCM, do the following:

Step 1 From the main menu, choose **Traffic Engineering > Local Congestion Mitigation > Domain-identifier > ... > Interface thresholds** and click one of the following:

- **Import CSV file**—Edit a CSV file to include a list of interfaces and thresholds, then later import the file into LCM.
- **Add new interface**—Manually add individual interfaces and thresholds.

Step 2 If you import a CSV file:

- Click the **Download sample configuration file** link.
- Click **Cancel**.
- Open and edit the configuration file (LCMLinkManagementTemplate.csv) you just downloaded. Replace the sample text with your specific node, interface, and threshold information.
- Rename and save the file.
- Navigate back to the **Customized interface thresholds** page.
- Click **Import CSV file** and navigate to the CSV file you just edited.
- Click **Import**.

Step 3 If you manually add individual interfaces:

- Click the first empty row and enter the appropriate node, interface, and threshold values.

Figure 64: Add First Interface



- Click **+ Create** to add more interfaces.

Step 4 Confirm that the information appears correctly on the **Customized interface thresholds** page.

Note To update the table, you can either turn on Edit Mode or import a CSV file that replaces all current data in the table.

Monitor LCM Operations

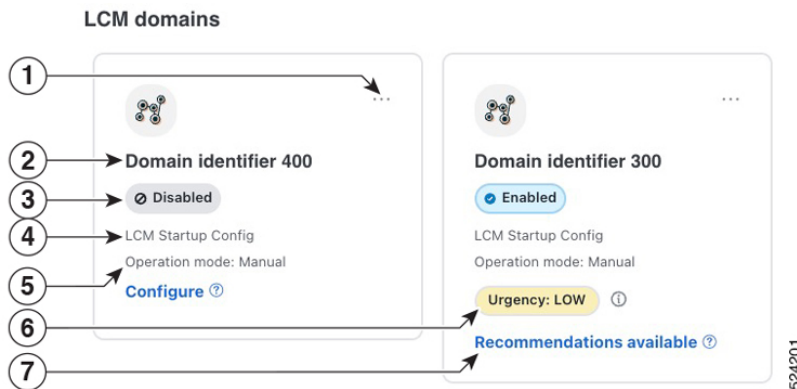


Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the [Workflow Example: Mitigate Congestion on Local Interfaces](#), on page 95 topic.

LCM Domains Dashboard

The LCM Domain Dashboard (**Traffic Engineering > Local Congestion Mitigation**) displays all the domains discovered by Crosswork. A *domain* is an identifier assigned to an IGP process.


Figure 65: LCM Domains Dashboard




Callout No.	Description
1	Main Menu: Allows you to navigate to the following pages: <ul style="list-style-type: none"> Operational Dashboard Add Individual Interface Thresholds Configure LCM
2	Domain identifier: The domain ID is taken from the router configuration (<code>link-state instance-id</code>) that you use to advertise IGP with BGP-LS.
3	LCM status: Indicates whether LCM has been enabled for the domain or can be deleted.

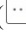


Callout No.	Description
4	LCM Configuration Description: The description is defined on the Configure LCM page. The default description is "LCM Startup Config".
5	Operation mode: Manual —This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
6	Urgency: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following: <ul style="list-style-type: none"> • Low—Indicates that LCM instantiated policies can be removed because they are no longer needed or that no changes are required. • Medium—Indicates new or modified recommendations. • High—Indicates network failures and recommendations should be deployed. This is a candidate that can be addressed automatically if the Auto Repair Solution advanced option was enabled. See Configure LCM, on page 100.
7	Configure: This link appears if LCM has not yet been configured. Click Configure to go to the Configure LCM page. Recommendations available: This link appears if LCM has detected congestion and has TTE policy recommendations. To view LCM recommendations, click the link to go to the LCM Operational dashboard . Delete: Indicates that the domain card can be removed from LCM monitoring.

LCM Operational Dashboard

The LCM Operational Dashboard (**Traffic Engineering > Local Congestion Mitigation > Domain-ID >  > Operational Dashboard**) shows congested interfaces as defined by the configured utilization threshold.

Each interface lists details such as current utilization, recommended action, status, expected utilization after committing recommendations, and so on. Hover the mouse pointer over  to view a description of the type of information each column provides.

From the Actions column, you can do the following:

- Preview TTE policies prior to deployment ( > **Preview Solution**)
- Verify deployment ( > **View Deployed Policies**)
- Pause or resume an interface ( > **Resume / Pause**)

To gain a better understanding of what information the LCM Operational Dashboard provides, see the following example:



Note If you are viewing the HTML version of this guide, click on the image to view it in full-size.

Figure 66: LCM Operational Dashboard


Operational Dashboard

● Congested interfaces (1) ● Mitigating interfaces (0) ● Mitigated interfaces (1) ● Paused interfaces (0) ● Resuming interfaces (0)

Commit All Urgency: MEDIUM

Node	Interface	Threshold Utiliz...	Evaluation Util...	LCM State	Policies Deplo...	Policy Set St...	Recommended ...	Commit St...	Expected Utiliz...	Solution Up...	Actions
L2-NCSS5...	GigabitEthern...	30%	25.85%	Mitigated	2	DEGRADED	Delete Set	None	12.92%	14-Nov-2023...	...
L5-0201-L...	FortyGigE0/0/...	8%	15.78%	Congested	0	-	Create Set	None	7.89%	14-Nov-2023...	...

In this example, the following information is conveyed:

- The first row is an interface that is currently in a Mitigated state. It shows that two policies have been deployed (**Policies Deployed - 2**) to mitigate a previous congestion. However, the current recommendation (**Recommended Action - Delete Set**) is to delete the policies since they are no longer needed (congestion should not occur even if the previously deployed policies are removed). Since the current recommendation has not been committed, the current Commit Status is None.
- The second row is an interface that is currently in a Congested state. LCM detects congestion and suggests to deploy policies to remediate the congestion (**Recommended Action - Create Set**). You can choose to preview the solution ( > **Preview Solution**).



Note If LCM cannot find a solution (**Recommended Action - No Solution**), it may be due to constraints enabled on the **LCM Configuration** page. For more information, see [Configure LCM, on page 100](#).

Recommendations are listed as part of a set, and if deployed, all changes are committed. You must click **Commit All**.

Temporarily Exclude an Interface from LCM



You can temporarily pause LCM from including an interface for mitigations. When an interface is paused, it will no longer be considered as part of a recommendation, and any existing solutions that the interface participates in will be removed. Pausing operations may be necessary in many use cases, such as the following:

- Where deployed solutions do not result in the intended resolution
- When there is uneven ECMP traffic
- When there are policies that do not carry traffic
- When an interface is continuously throttling between different solutions

LCM Function Pack may automatically pause an interface when certain anomalies are detected, for example, when there is:

- No LCM SR policy traffic
- Excessive LCM Policy Traffic Imbalance
- Excessive LCM Oscillations/removals per hour

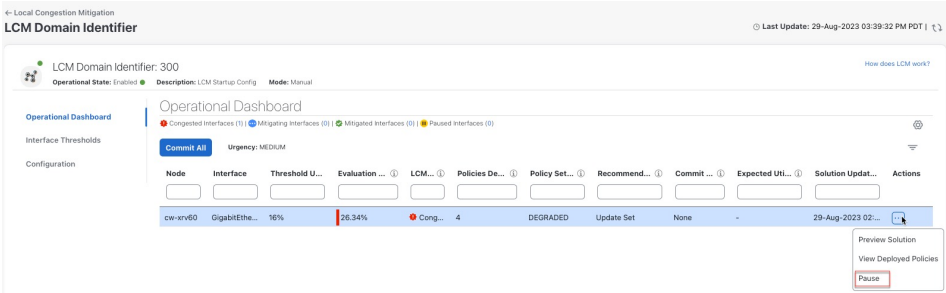
In these circumstances, the user may perform a corrective action, and manually Unpause the interface.

From the Actions column of the LCM Operational Dashboard, click  > **Pause** for the interface you would like to exclude from LCM calculations. To include the interface in LCM calculations again, click  > **Resume**.



Note Pausing multiple interfaces at the same time may result in requests timing out. However, each request will be queued and displayed on the dashboard.

Figure 67: Pause Interface



The screenshot shows the LCM Operational Dashboard for LCM Domain Identifier: 300. The dashboard displays a table of congested interfaces. The first row is highlighted, and a context menu is open over the 'Actions' column, showing the 'Pause' option.

Node	Interface	Threshold U...	Evaluation ...	LCM...	Policies De...	Policy Set...	Recommend...	Commit ...	Expected Util...	Solution Updat...	Actions
cw-vr60	GigabitEthe...	16%	26.34%	Cong...	4	DEGRADED	Update Set	None	-	29-Aug-2023 02...	<ul style="list-style-type: none"> Preview Solution View Deployed Policies Pause



CHAPTER 9

Bandwidth on Demand (BWoD)

Bandwidth on Demand (BWoD) provides a bandwidth-aware Path Computation Element (PCE) in conjunction with SR-PCE for segment routing policies (SR policies). BWoD policies can be PCC-initiated (PCE-delegated) or PCE-initiated. BWoD is designed to deliver soft bandwidth guarantee services over SR policies. BWoD monitors network conditions and re-optimizes BWoD paths to prevent total BWoD traffic on any interface from exceeding the configured threshold percent.

BWoD does not track total interface utilization, and therefore, interfaces can still be congested if the combined BWoD traffic and non-BWoD traffic exceed the interface capacity. In addition, BWoD does not enforce the total amount of traffic entering BWoD SR policy. BWoD policies may traverse Equal Cost Multi-Path (ECMP) and assume even traffic distribution over these paths. However, actual ECMP distribution can be uneven, especially with large flows.



Note Functionality described within this section is only available with certain licensing options.

This section contains the following topics:

- [BWoD Important Notes, on page 109](#)
- [PCC-Initiated BWoD SR-TE Policies, on page 110](#)
- [Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example, on page 111](#)
- [Configure Bandwidth on Demand, on page 117](#)
- [Troubleshoot BWoD, on page 117](#)

BWoD Important Notes

Consider the following information when using BWoD:

- You must have the Advanced RTM license package to use BWoD.
- Role-based Access Control (RBAC) and task permissions have been introduced in this release. To provision a BWoD policy, you must have write-access to the head-end device based on Device Access Groups and assigned roles. Only BWoD admin users can modify BWoD configuration settings. For more information on RBAC and user roles, see the [Cisco Crosswork Network Controller Administration Guide](#).
- If BWoD cannot find a path for a policy that guarantees its requested bandwidth, BWoD will attempt to find a *best effort* path if this option is enabled.

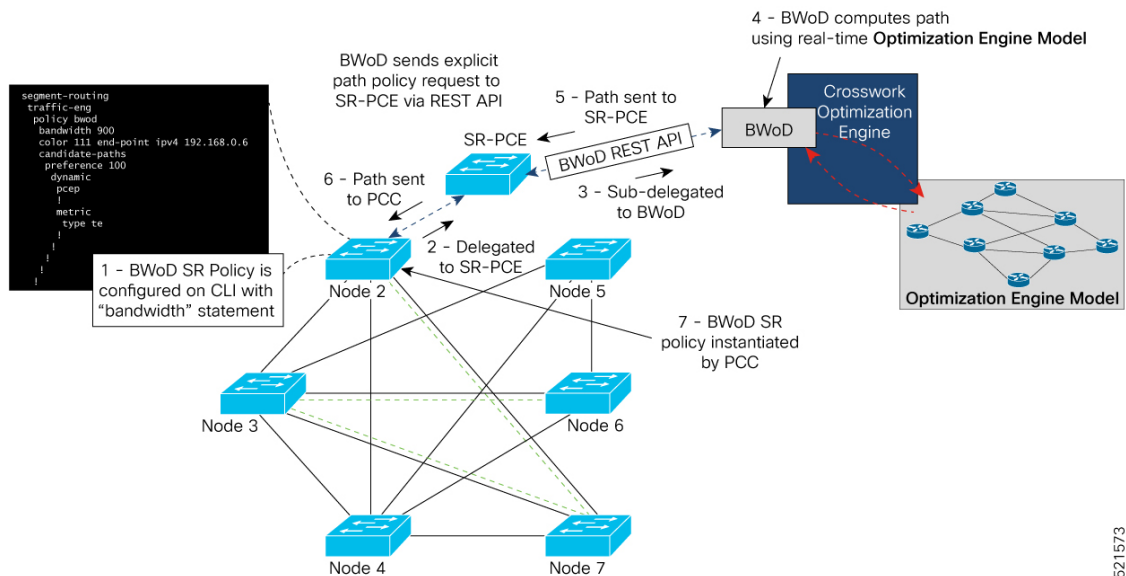
- BWoD will disable itself when an unexpected error is encountered to avoid network disruption.
- BWoD temporarily pauses operation whenever the Optimization Engine model is unavailable due to an Optimization Engine restart or a rebuild of the topology from Topology Services. Any requests to BWoD during this time are rejected. When the model becomes available and BWoD receives two traffic updates from the Optimization Engine, BWoD will resume normal operation.
- If the Policy Violation advanced field is set to **Strict**, then the SR Policy Traffic option should be set to **Max Measured Requested**.
- After a switchover in a High Availability setup, BWoD policies created after the last cluster data synchronization will not be manageable and are considered orphaned TE policies. Crosswork will display an alarm when it finds orphan TE policies (**Administration > Alarms**). You can use APIs to help clear these orphan policies so that they are manageable. For more information, see [API documentation on Devnet](#).

PCC-Initiated BWoD SR-TE Policies

When enabled, BWoD automatically connects to all SR-PCE providers configured in Crosswork. A persistent connection is made to the SR-PCE BWoD Rest API, which is registered as a PCE for bandwidth-constrained SR-TE policies.

The following figure shows the PCC-initiated workflow for BWoD:

Figure 68: PCC-Initiated BWoD SR-TE Policies



521573

Callout No.	Description
1	<p>A BWoD policy is configured on a PCC via the CLI. For example:</p> <pre> segment-routing traffic-eng policy bwod bandwidth 900 color 100 end-point ipv4 1.1.1.2 candidate-paths preference 100 dynamic pcep ! metric type te ! ! constraints affinity exclude-any name RED ! ! ! !</pre>
2	The bandwidth statement is added to a PCE-delegated SR policy to create a BWoD policy. Once committed, the PCC delegates the path compute to SR-PCE.
3, 4	SR-PCE then sub-delegates the policy to BWoD, which attempts to compute a path that meets the bandwidth constraint.
5, 6	If a bandwidth-compliant path is found, the segment list is returned to SR-PCE, which forwards it over PCEP to the PCC, and the PCC instantiates it. If BWoD is unable to compute a bw-compliant path for the policy or doing so will force an existing BWoD policy to not have a bw-compliant path, best effort paths may be computed by BWoD, which attempts to minimize violations. This occurrence will also trigger BWoD to issue an event to the COE events UI indicating which BWoD policies are now on best-effort paths.
7	A BWoD SR-TE policy is instantiated.

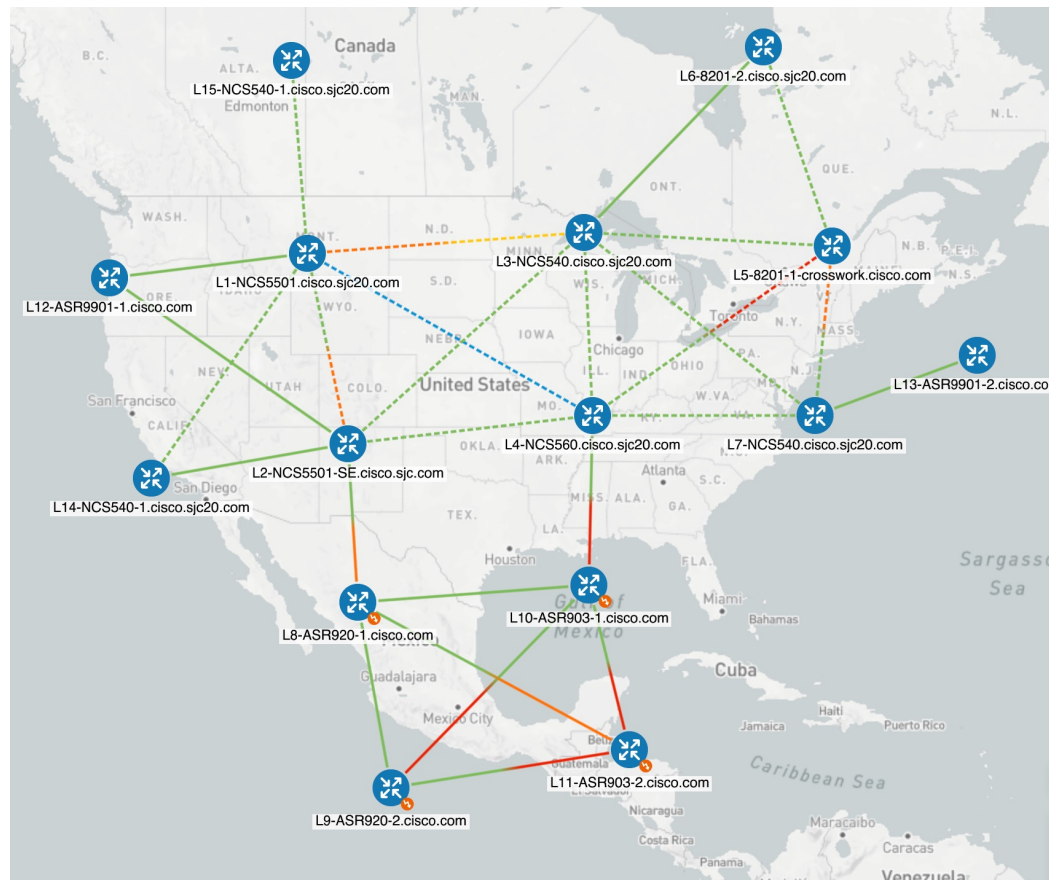
Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example

This example demonstrates

- how to enable and configure Bandwidth on Demand (BWoD)
- how to create BWoD policies
- how BWoD calculates paths, and
- how BWoD calculates new policies when the Policy violation option is set to Loose or Strict.

In particular, three BWoD policies will be created using the *same* headend (L1-NCS5501.cisco.sjc20.com) and endpoint (L5-8201-1-crosswork.cisco.com) with a bandwidth requirement of 700 and 1000 Mbps, while keeping the utilization at 80%. In this example, all interfaces have the capacity of **1 Gbps**.

Figure 69: Initial BWoD Topology Example



Step 1 Enable and configure BWoD.

- a) From the main menu, choose **Traffic Engineering > Bandwidth on Demand > Configuration**.
- b) Toggle the Enable switch to **True**, enter **80** in the **Link utilization** field, and confirm that **Advance > Policy violations** is set to **Loose**. To find descriptions of other options, simply hover the mouse over **?**.
- c) Click **Commit changes**.

Step 2 Create the first PCE-initiated BWoD SR-TE policy.

- a) From the main menu, choose **Traffic Engineering > SR-TE** tab and click **+Create**.
- b) Enter the required policy details. In this example, we are creating a policy with these values:
 - Headend: **L1-NCS5501.cisco.sjc20.com**
 - Endpoint: **L5-8201-1-crosswork.cisco.com**
 - Color: **70000**

Example:

Figure 70: Policy Details

Policy details

Headend * ⓘ
Selected - L1-NCS5501.cisco.sjc20.com [192.168.] ⓘ [Edit](#)
[2001:192:168::1]

L1-NCS5501.cisco.sjc20.com [192.168.] [2001:192: .. ▼

Endpoint *
Selected - L5-8201-1-crosswork.cisco.com [192.168.] ⓘ [Edit](#)
[2001:192:]

L5-8201-1-crosswork.cisco.... 192.168. ▼

Color * ⓘ
70000

- c) In the **Policy path** area, click **Bandwidth on demand**, and enter the required policy path details. In this example, we use these values:

- Path name: **bwod-70000**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **7000 Mbps**

Example:

Figure 71: Policy Path Details

Policy path

Explicit path Dynamic path Bandwidth on demand

Path name * ⓘ
bwod-70000

Optimization objective *
Interior gateway protocol (IGP) metric ▼

Bandwidth * ⓘ
700 Mbps

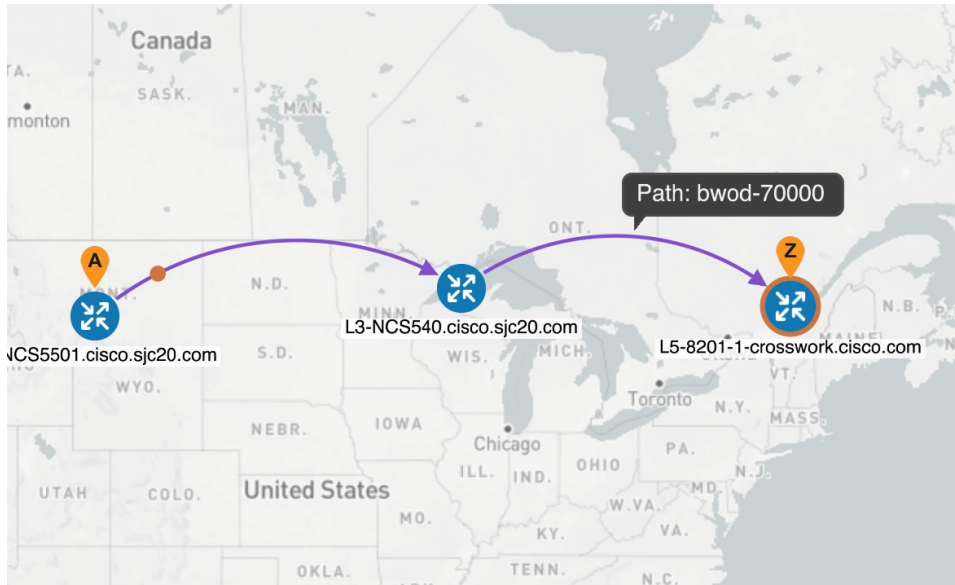
SID algorithm ⓘ

- d) Click **Preview**. BWoD only takes into account current interface utilization that has been reserved by another BWoD policy. Otherwise, BWoD only considers the capacity of the interface in its calculations. In this example, all interfaces

have the capacity of 1 Gbps. Since there are no existing BWoD policies, BWoD considers the capacity of all nodes and takes the shortest route.

Example:

Figure 72: First BWoD Policy (bwod-70000)



e) If you are satisfied with the proposed BWoD SR-TE policy deployment, click **Provision**.

Step 3 Verify that the new BWoD SR-TE policy has been created.

a) From the main menu, choose **Traffic Engineering > SR-TE**.

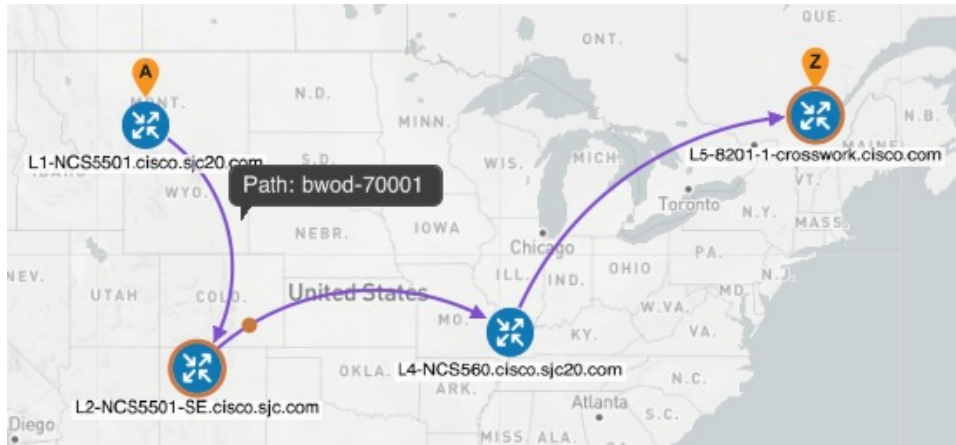
b) Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View**).

Step 4 Create a second BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70001**
- Path name: **bwod-70001**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **700 Mbps**

BWoD considers the existing BWoD policy (bwod-70000) and its bandwidth requirement into its interface capacity calculations. So, a new path is created for the bwod-70001 policy.

Figure 73: New bwod-70001 Policy

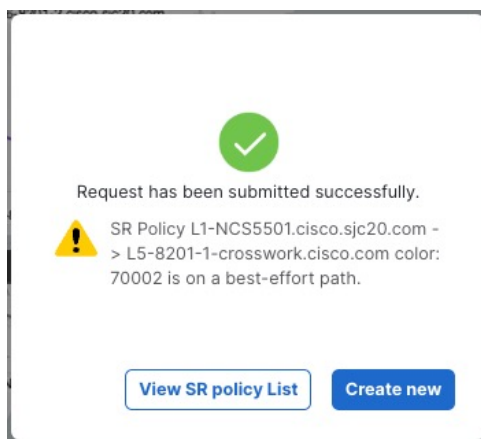


Step 5 Create a third BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70002**
- Path name: **bwod-70002**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **1000 Mbps**

Since BWoD takes into account all previous BWoD policy requirements and the BWoD policy violation option was set to **Loose**, BWoD creates a best effort path for the bwod-70002 policy. You will receive this message when you provision the new policy:

Figure 74: Best Effort Message



Note that existing paths for bwod-7000 and bwod-70001 are moved to accommodate the new bwod-70002 policy.

Example:

Figure 75: BWoD Policies with Loose Option



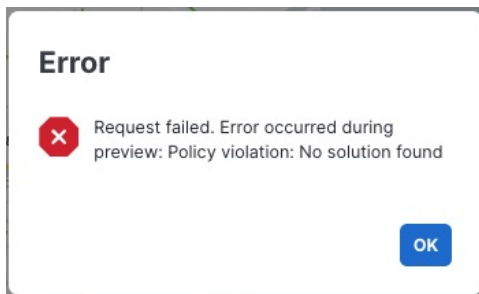
Step 6 Change the BWoD policy violation option to **Strict** (Traffic Engineering > Bandwidth on Demand > Configuration > Advanced).

Step 7 Create a fourth BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70003**
- Path name: **bwod-70003**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **1000 Mbps**

Since the BWoD policy violation option is set to **Strict**, BWoD is not be able to overwrite existing BWoD policies, and the request for additional 1000 Mbps policy results in a "No solution found" message.


Figure 76: No Solution Found



Configure Bandwidth on Demand

There are two main steps in using Bandwidth on Demand (BWoD):

1. Enable and configure BWoD options.
2. Create BWoD SR policies. As long as BWoD is enabled, you can create multiple BWoD SR policies.

-
- Step 1** From the main menu, choose **Traffic Engineering > Bandwidth on Demand > Configuration**.
- Step 2** Toggle the **Enable** switch to **True**.
- Step 3** Configure additional options. Hover the mouse pointer over  to view a description of each field.
- Step 4** Click **Commit changes** to save the configuration.
- Step 5** To create BWoD SR policies, choose **Traffic Engineering > Traffic Engineering**.
- Step 6** From the SR Policy table, choose **Create > PCE Init**.
- Step 7** In addition to entering the required SR policy details, click the **Bandwidth on demand** option and enter the required bandwidth.
- Step 8** If applicable, enter a Flexible Algorithm constraint in the **SID Algorithm** field. The values correspond to the Flexible Algorithm that are defined on the device and the 128-255 range is enforced by Cisco IOS XR. Cisco Crosswork will try to find a path with this SID. If a path with the SID constraint cannot be found, the provisioned policy will remain operationally down until the conditions are met.
- Step 9** Click **Preview** to view the proposed SR policy.
- Step 10** Click **Provision** to commit the SR policy.
-

Troubleshoot BWoD

The following are some of the most common error conditions for BWoD and some possible corrective actions that may fix the issue.

Table 5: Errors

Error Event Message	Possible Causes and Recommended Corrective Action
OptimaModelError	<p>The network model used by BWoD from the Optimization Engine is corrupt or is missing key data needed to properly support BWoD. Possible causes include network discovery issues or synchronization problems between the Optimization Engine and Topology Services. Try restarting the Optimization Engine pod to rebuild the model.</p> <p>This error can also occur if the time required to discover a policy and add it to the model after it has been deployed exceeds the Deployment Timeout option set for BWoD. The default is 30 seconds, sufficient for small to medium-sized networks. However, larger networks may require additional time.</p>
NATSTimedOutError	<p>The deployment of a bandwidth policy through SR-PCE exceeds the Deployment Timeout option set for BWoD. Increase the Deployment Timeout option to allow for additional time for deployments in larger networks.</p>
Traceback or other errors found in the log file	<p>Please contact your Cisco service representative.</p>

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)

© 2024 Cisco Systems, Inc. All rights reserved.

