



Cisco Crosswork Optimization Engine 5.0 User Guide

First Published: 2023-04-23

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	Setup and Monitor Your Network	5
	Get a Quick View in the Dashboard	5
	View Devices and Links on the Topology Map	6
	View Device Details	9
	View Link Details	13
	Link States and Discovery Methods	17
	Protocols Used for Topology Services	17
	Customize Map Display Settings	18
	Customize the Display of Links and Devices	18
	Use Device Groups to Filter Your Topology View	18
	Create and Modify Device Groups	22
	Enable Dynamic Device Grouping	23
	Save Topology Views for Easy Access	23
CHAPTER 3	Traffic Engineering in Crosswork Optimization Engine	25
	Segment Routing Path Computation Element (SR-PCE)	26
	What is Segment Routing?	26
	SR-TE Policy PCC and PCE Configuration Sources	28

PCC-Initiated SR-TE Policy Example	28
What is Resource Reservation Protocol (RSVP)?	29
RSVP-TE Tunnel PCC and PCE Configuration Sources	30
PCC-Initiated RSVP-TE Tunnel Example	30
Get a Quick View of Traffic Engineering Services	30
View TE Event and Utilization History	32
View Traffic Engineering Device Details	33
Configure Traffic Engineering Settings	34
Configure TE Timeout Settings	34
Configure How Device Groups Are Displayed for Traffic Engineering	34
Configure TE Dashboard Settings	35

CHAPTER 4**SR-MPLS and SRv6 37**

View SR-MPLS and SRv6 Policies on the Topology Map	37
View SR-MPLS and SRv6 Policy Details	39
Visualize SR-MPLS or SRv6 Policies Example	40
Find Multiple Candidate Paths (MCPs)	46
Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label	49
Visualize Native SR Paths	51
Visualize Native Path Device Prerequisites	53
Configure TE Link Affinities	54
Create Explicit SR-MPLS Policies	55
Create Dynamic SR-MPLS Policies Based on Optimization Intent	56
Modify SR-MPLS Policies	58

CHAPTER 5**Resource Reservation Protocol (RSVP) 59**

View RSVP-TE Tunnels on the Topology Map	59
View RSVP-TE Tunnel Details	61
Create Explicit RSVP-TE Tunnels	62
Create Dynamic RSVP-TE Tunnels Based on Optimization Intent	63
Modify RSVP-TE Tunnels	64

CHAPTER 6**Flexible Algorithms 65**

Configure Flexible Algorithm Affinities	65
---	----

Visualize Flexible Algorithms 66
 Find Flexible Algorithms for Links and Devices 68

CHAPTER 7

Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering 71
 Visualize Tree-SID Policies 71
 View a Point-to-Multipoint Tree on the Topology Map 72
 Create Static Tree-SID Policies 73
 Static Tree-SID Policy Configuration Example through Crosswork UI 74
 Modify a Tree-SID Policy 76
 Tree-SID Important Notes 76

PART I

Bandwidth Feature Packs 79

CHAPTER 8

SR Circuit Style Manager (CSM) 81
 Circuit Style SR-TE Important Notes 82
 Workflow for Setting Up CS SR-TE Policy Visualization 86
 Enable SR Circuit Style Manager 87
 Configure Circuit Style SR Policies 88
 Review Circuit Style SR-TE Policy Bandwidth Utilization 90
 View Circuit Style SR-TE Policies 91
 Trigger CSM to Recalculate a Circuit Style SR-TE Policy 96
 What Happens When Bandwidth Reservation Settings are Exceeded? 96
 How Does CSM Handle Path Failures? 99

CHAPTER 9

Local Congestion Mitigation (LCM) 103
 Local Congestion Mitigation Overview 103
 LCM Important Notes 104
 LCM Platform Requirements 105
 BGP-LS Speaker Placement for Multiple AS Networks with a Dedicated IGP Instance Between ASBRs 105
 LCM Calculation Workflow 106
 Workflow Example: Mitigate Congestion on Local Interfaces 108
 Configure LCM 117
 Add Individual Interface Thresholds 119

Monitor LCM Operations 121

CHAPTER 10

Bandwidth on Demand (BWoD) 125

BWoD Important Notes 125

PCC-Initiated BWoD SR-TE Policies 126

Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example 128

Configure Bandwidth on Demand 129

Troubleshoot BWoD 130



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 Algorithms

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 proactive 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 is also a 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 real-time network visualization of the following:
 - devices
 - links and link utilization
 - SR-TE policies
 - SR-MPLS and SRv6
 - Tree-SID
 - Flexible Algorithms
 - Circuit-Style Segment Routing (CS-SR) policies
 - RSVP-TE tunnels
- A UI that allows the network operator to perform the following tasks:
 - 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.
- 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, either due to licensing or the configuration of the role that is associated with your user account, you may not be able to access the features and functions. 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 that delivers increased service agility, cost efficiency, and optimization for faster time-to-customer value and lower operating cost. 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, and anomaly detection and automatic remediation. For more information, see [Cisco Crosswork Network Controller](#).

Cisco Crosswork Optimization Engine is offered as an individual application and is also a component of Cisco Crosswork Network Controller where its functionality is integrated into Cisco Crosswork Network Controller's UI.

Throughout this document, when using the Crosswork Optimization Engine as part of the Crosswork Network Controller solution, some options are not available or are 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 that tackle congestion mitigation and the management of SR-TE policies to find and maintain intent based bandwidth requirements. Users can define the optimization intent and the tools implement the intent, while continuously monitoring, tracking, and reacting to maintain the original intent. To learn more about these feature packs, see the following topics:

- [Local Congestion Mitigation \(LCM\)](#), on page 103
- [SR Circuit Style Manager \(CSM\)](#), on page 81
- [Bandwidth on Demand \(BWoD\)](#), on page 125

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 into their network operations.

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



CHAPTER 2

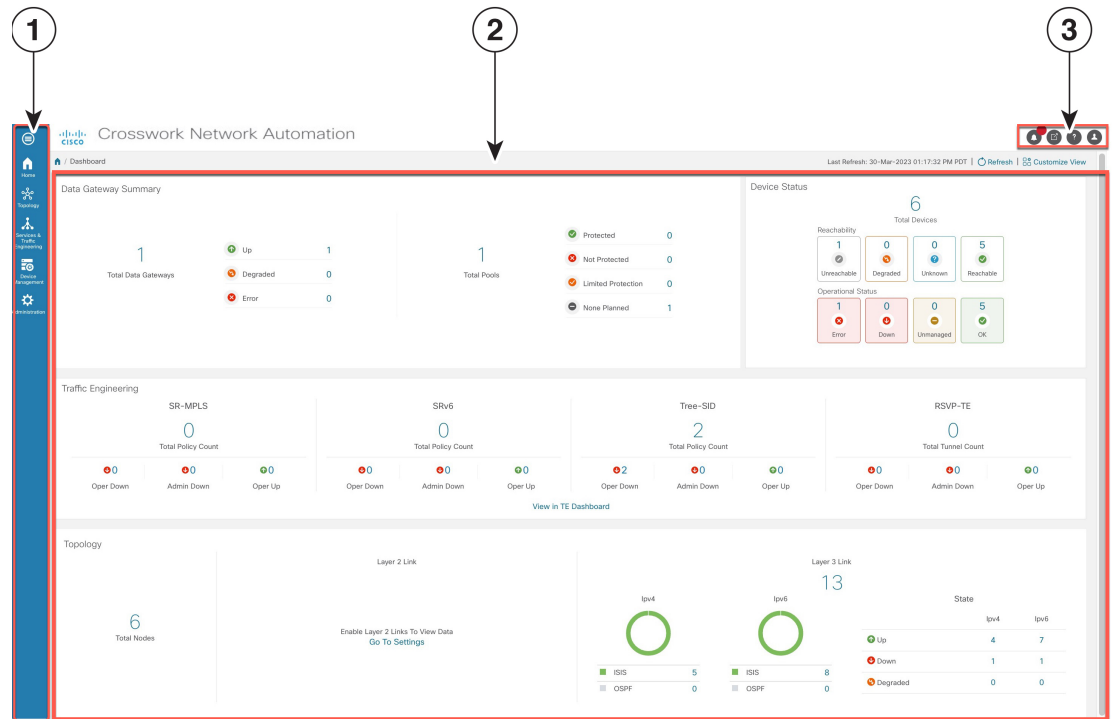
Setup and Monitor Your Network

- [Get a Quick View in the Dashboard, on page 5](#)
- [View Devices and Links on the Topology Map, on page 6](#)
- [Customize Map Display Settings, on page 18](#)
- [Save Topology Views for Easy Access, on page 23](#)





Get a Quick View in the Dashboard

The Home page displays a customizable collection of dashlets which provide an at-a-glance operational summary of the network being managed, including reachability and operational status of devices. The Dashboard is made up of a series of dashlets, and each dashlet represents different types of data belonging to the same category.

Figure 1: Crosswork Home page



523215

Callout No.	Description
1	Main Menu: The main menu allows you to navigate to installed Cisco Crosswork applications and device management and administrative tasks. Menu options may look slightly different depending on which Cisco Crosswork applications are installed.
2	<p>Dashlets: Information varies depending on which Cisco Crosswork applications are installed.</p> <ul style="list-style-type: none"> • To drill down for more information within a dashlet, click on a value. A window appears displaying only the filtered data you clicked on. • To add or change the layout of dashlets, click Customize View. Move the dashlets to your desired layout and click Save. • You can duplicate or remove TE dashlets using the trash or pencil icons within the top-right corner of the dashlet.
3	<p>Settings icons:</p> <ul style="list-style-type: none">  The Alerts icon notifies you of any current error conditions related to the system operations which require attention, and provides a link to detailed information about those conditions.  The Events icon notifies you of new events related to system operation, and also provides access to the history of all system events.  The About icon displays the current version of the Cisco Crosswork product.  The User Account icon lets you view your username, change your password, and log out.

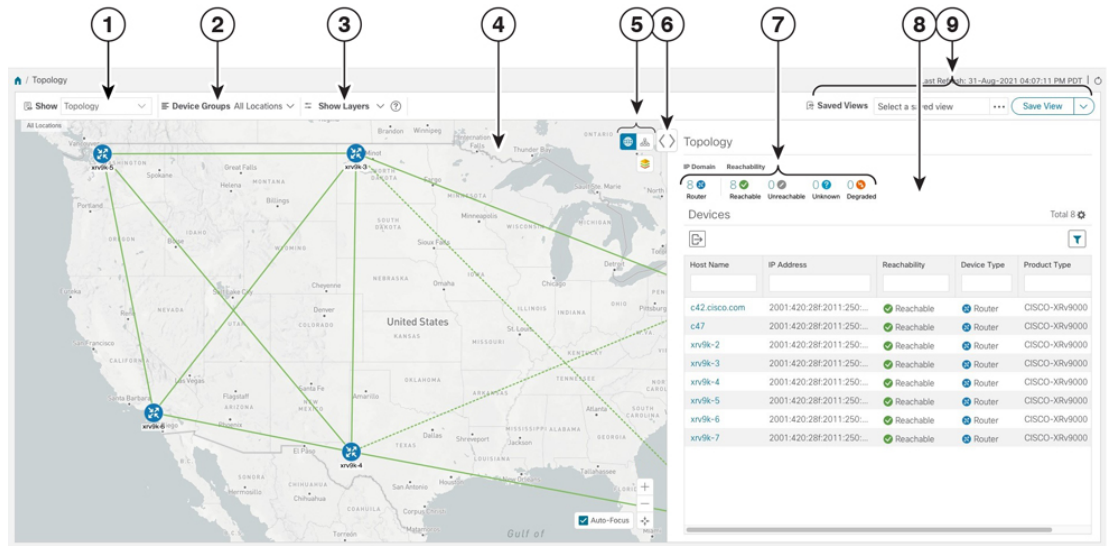
View Devices and Links on the Topology Map

To view your devices on the topology map, they must be onboarded. For more information refer to the [Cisco Crosswork Network Controller Administration Guide](#). To view the network topology map, from the main menu choose **Topology**.









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

Figure 2: Cisco Crosswork UI and Topology Map



522060

Callout No.	Description
1	<p>Topology Map View: From the Show drop-down list, click the option that displays the data that you would like to see on the map.</p> <p>If Topology is selected, devices and links in the network are displayed.</p> <p>If Traffic Engineering is selected, TE tunnel information is displayed. For more information on the Traffic Engineering topology map, see View SR-MPLS and SRv6 Policies on the Topology Map, on page 37 and View RSVP-TE Tunnels on the Topology Map, on page 59.</p>
2	<p>Device Groups: From the drop-down list, click the group of devices you want displayed on the map. All other device groups will be hidden.</p>
3	<p>Show Hide: From the drop-down list, click the network layers you want displayed on the map. All devices and links that belong to the selected layers are then displayed. By default, all layers are displayed.</p>

Callout No.	Description
4	<p>Topology Map: The network topology can be displayed on a logical map or a geographical map, where the devices and links are shown in their geographic context. From the map, you can drill down to get detailed information about devices and links.</p> <p>Devices:</p> <ul style="list-style-type: none"> • To view a device configuration summary, hover the mouse pointer over the device icon. A pop up window displaying the host name, state, node ID, and device type appears. • To view device details, click on the device icon. • If devices are in close physical proximity, the geographical map shows them as a cluster. <p>The number in a blue circle () indicates the number of devices in the cluster. Displaying devices in this manner helps prevent overlap and clutter on the map.</p> <p>Links:</p> <ul style="list-style-type: none"> • A solid line indicates a <i>single link</i> between two devices. If there is more than one link between two devices, or between a device and a cluster of devices, the line is shown dashed instead. A dashed line indicates an <i>aggregated</i> link that represents more than one link, or the use of multiple protocols (for example, IPv4 and IPv6) on the same physical link. • A and Z indicates headend and endpoint, respectively. • To view link information details, click on the link. <p>Note Although aggregated, dual stack links show as one single line.</p>
5	<p>: The logical map shows devices and their links, positioned according to an automatic layout algorithm, ignoring their geographical location. You can change the layout algorithm.</p> <p>: The geographical map shows single devices, device clusters, links, and tunnels, superimposed on a map of the world. Each device location on the map reflects the device's GPS coordinates (longitude and latitude) as defined in the device inventory.</p> <p>: The Display Preferences window allows you to change display settings for devices, links, utilization, Flexible Algorithms, and TE tunnel metrics.</p> <p>: The global search allows you to search the topology using device names, location or the device civic location.</p> <p>: Export to KML allows you to export the geo localised objects using a KML format. KML is an XML base file format used to display information in a geographic context, such as Google Earth.</p>
6	<p>Expand/Collapse/Hide Side Panel: Expand or collapse the contents of the side panel. Close the side panel to get a larger view of the topology map.</p>

Callout No.	Description
7	<p>The Mini Dashboard provides a summary of the IP Domain and device reachability status. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the Devices table.</p> <p>Note If the Alarm Status feature is enabled, you will also see Alarm information here. To view the Alarm Status, you must install Element Management Functions (part of the Crosswork Network Controller Essentials package) and configure host information for Syslog and SNMP traps on the devices you want to view alarms for. For more information, see the Cisco Crosswork Network Controller Installation Guide and the Cisco Crosswork Network Controller Administration Guide. The Alarm Status feature is available for select licensing packages.</p>
8	The content of this window changes depending on what applications you have installed, what Show is set to for the Topology Map and if you have selected to view more information on the device, link, SR-MPLS policy, SRv6 policy, or RSVP-TE tunnel..
9	Saved Custom Map Views: Lets you create a named custom view using the settings and layout for your current map, settings of the tables saved in the saved views, or display a custom view you have created previously. It also saves any filters applied to the Devices and Traffic Engineering tables.

View Device Details

This example shows how you can view device using the topology map.



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

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

Step 2 To quickly view the host name, reachability state, IP address and type of device, hover the mouse over the device icon.

The screenshot displays the Cisco Crosswork Optimization Engine 5.0 interface. On the left, a map shows a network topology with several devices connected. A tooltip is visible over one device, showing details: Reachability State (Reachable), Host Name (PCC4_80), Node IP (10.195.165.80), and Type (Cisco IOS XRv 9000 Router). On the right, the 'Traffic Engineering' section is active, showing a table of SR Policies. The table has columns for Headend, Endpoint, Color, Admin St..., Oper Status, and Actions. The table contains several rows of data, including PCC1_77, PCC2_78, PCC3_79, PCC5_81, and PCC5_81.

Step 3 To view more device details, click on the device icon.

a) The following examples show the Device details from the Topology map.

The screenshot displays a network map of the United States with several routers (PCC1_77, PCC1_80, PCC1_79, PCC1_76, PCC1_81, PCC1_82, PCC2_78) connected by green lines. The right-hand pane shows the 'Device Details' for PCC1_77, including Summary and Routing information.

Device Details

Summary

- Host Name: PCC1_77
- Reachability: ✔ Reachable
- IP Address: 10.195.165.77
- Civic Address: Chicago, Illinois, United States, North America, 7045
- Geo Location: Latitude 42.190000, Longitude -73.800000
- Device Type: 📡 Router
- Device Group: Location > All Locations > Unassigned Devices
- Product Type: Cisco IOS XRv 9000 Router
- Connect To Device: 🔑 SSH IPv4
- Last Update: 09-Apr-2023 04:22:21 PM GMT+5:30

Routing

- OSPF Router ID: 100.100.100.1 Area: 0.0.0.0 (0)
- TE Router ID: 100.100.100.1
- ASN: 1

Note If the Alarm Status feature is enabled, you will also see Alarm information here. To view the Alarm Status, you must install Element Management Functions (part of the Crosswork Network Controller Essentials package) and configure host information for Syslog and SNMP traps on the devices you want to view alarms for. For more information, see the [Cisco Crosswork Network Controller Installation Guide](#) and the [Cisco Crosswork Network Controller Administration Guide](#). The Alarm Status feature is available for select licensing packages.

In a multiple IGP setup, you can also view all the IGP, IS-IS, and OSPF processes in the Routing details. See the following examples:

Figure 3: Multiple IGP: OSPF Processes

The screenshot displays a network map of the United States with several routers (PCC1_77, PCC1_80, PCC1_79, PCC1_76, PCC1_81, PCC1_82, PCC2_78) connected by green lines. The right-hand pane shows the 'Device Details' for PCC1_77, including Summary and Routing information. The OSPF Router ID field is highlighted with a red box.

Device Details

Summary

- Host Name: PCC1_77
- Reachability: ✔ Reachable
- IP Address: 10.195.165.77
- Civic Address: Chicago, Illinois, United States, North America, 7045
- Geo Location: Latitude 42.190000, Longitude -73.800000
- Device Type: 📡 Router
- Device Group: Location > All Locations > Unassigned Devices
- Product Type: Cisco IOS XRv 9000 Router
- Connect To Device: 🔑 SSH IPv4
- Last Update: 09-Apr-2023 04:22:21 PM GMT+5:30

Routing

- OSPF Router ID: 100.100.100.1 Area: 0.0.0.0 (0)
- TE Router ID: 100.100.100.1
- ASN: 1

Figure 4: Multiple IGP: ISIS Processes

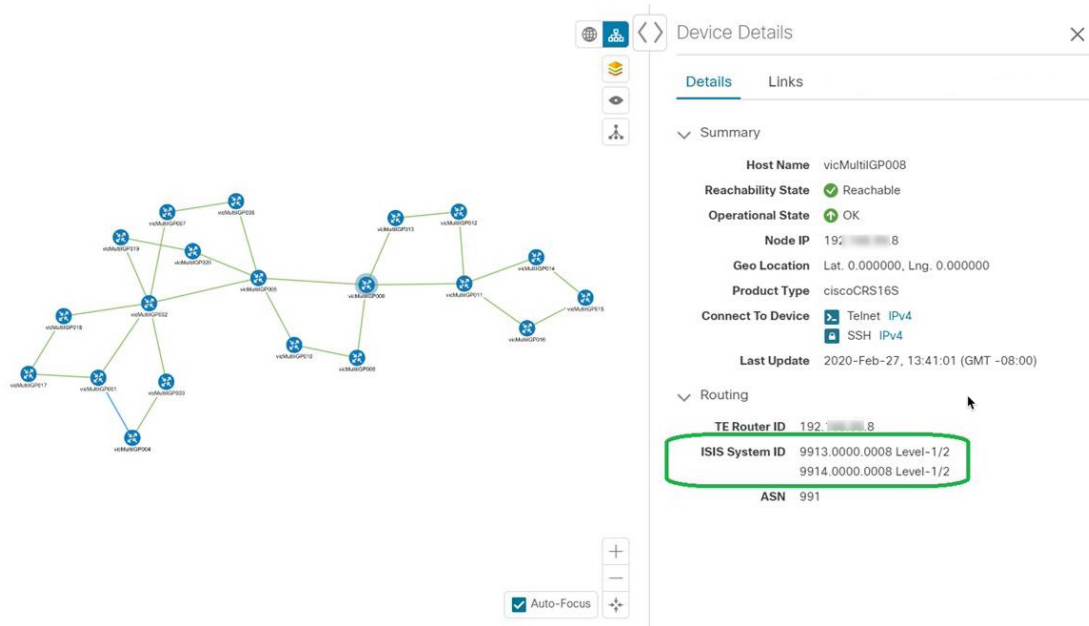
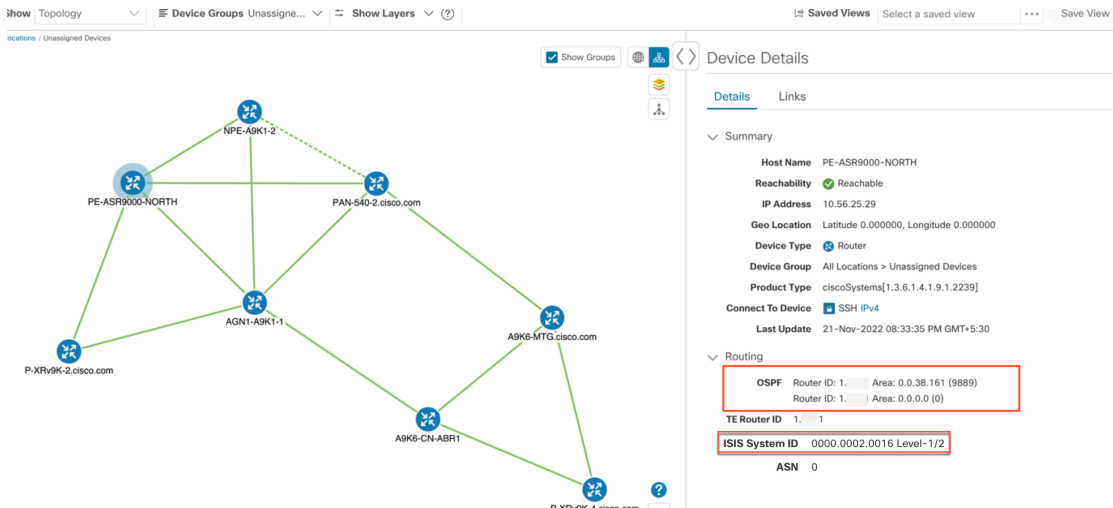



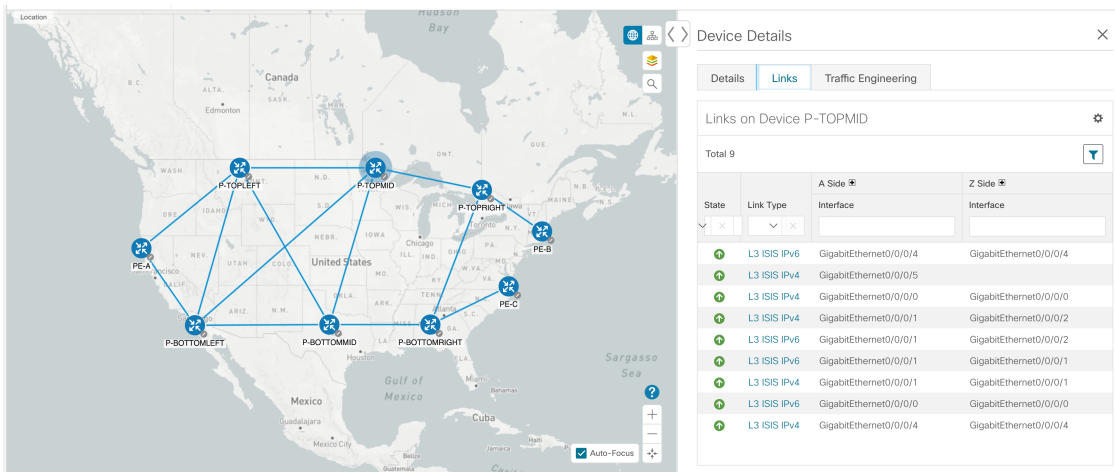
Figure 5: Multiple IGP: OSPF and ISIS Processes



- b) The following example shows additional Traffic Engineering Device details (SR-MPLS, SRv6, RSVP-TE, and Flexible Algorithm tabs) from the Traffic Engineering map. In this particular example, SRv6 Locators are listed for two domains.



Step 4 To view links on the device, click the **Links** tab and expand the right panel to see all the link details.



Step 5 To view interface utilization, expand **A side** or **Z side**.

The utilization shown on IPv4 and IPv6 links represents the aggregate traffic on the interface, not specific to each address family. Sub-interfaces will not show a utilization since they do not have a bandwidth like a physical interface. Traffic measurements will still be collected and displayed.

Interface	Utilization
GigabitEthernet0/0/0/5	0% (0Bps/1Gbps)
GigabitEthernet0/0/0/6	0% (0Bps/1Gbps)
GigabitEthernet0/0/0/2	0% (0Bps/1Gbps)
GigabitEthernet0/0/0/3.1	
GigabitEthernet0/0/0/2	0% (0Bps/1Gbps)
GigabitEthernet0/0/0/5	0% (0Bps/1Gbps)
GigabitEthernet0/0/0/1	0% (0Bps/1Gbps)

View Link Details



This example describes how you can view the following information:

- Link details (name, state, type, and endpoint interface information)
- Members of an aggregate link
- IPv4 unnumbered interfaces information (when available) is now displayed as either an index or a combination of the TE Router ID and the index in device, link, and topology details.
- Link Aggregation Group (LAG) details


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



Step 2 View link details.

Click a link on the topology map.

Link Details  

Summary

Name GigabitEthernet0/0/0/2-GigabitEthernet0/0/0/2
State  Up
Link Type L3 ISIS IPv4
ISIS Level 2
Last Update 13-Apr-2023 09:19:44 PM PDT

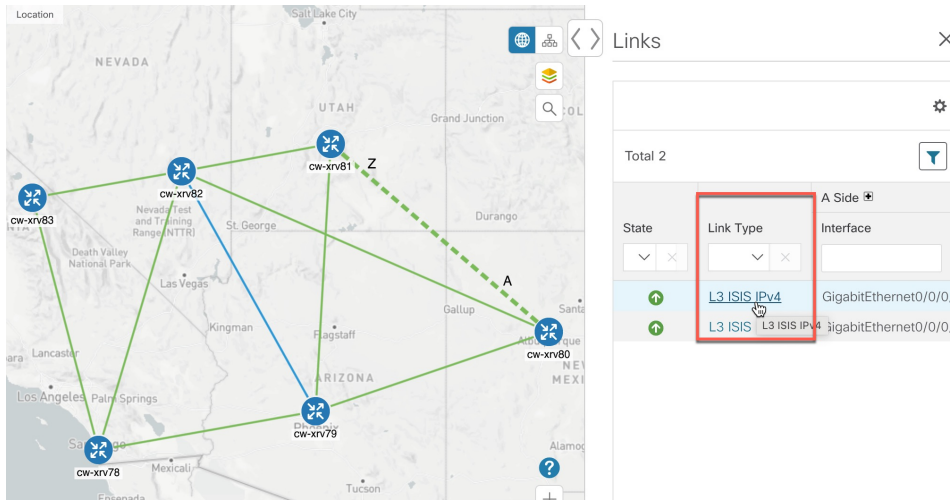
	A Side	Z Side
Node	cw-xrv78	cw-xrv83
TE Router ID	3.3.3.78	3.3.3.83
IPv6 Router ID	bb:bb:bb:3:3::78	bb:bb:bb:3:3::83
IF Name	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/2
IF Description	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/2
IF Alias	*** connect to xrv83 ***	*** connect to xrv78 ***
Type	ETHERNETCSMACD	ETHERNETCSMACD
Local IF ID	3.3.3.78 (10)	3.3.3.83 (8)
Utilization	 0% (0Bps/1Gbps)	 0% (0Bps/1Gbps)
Packet Drops	0%	0%
IGP Metric	10	10
Delay Metric	10	10
TE Metric	10	10
Admin Groups	5	

Step 3 View aggregate link details.

Click on a dashed line. A dashed line indicates an aggregated link that represents more than one link.

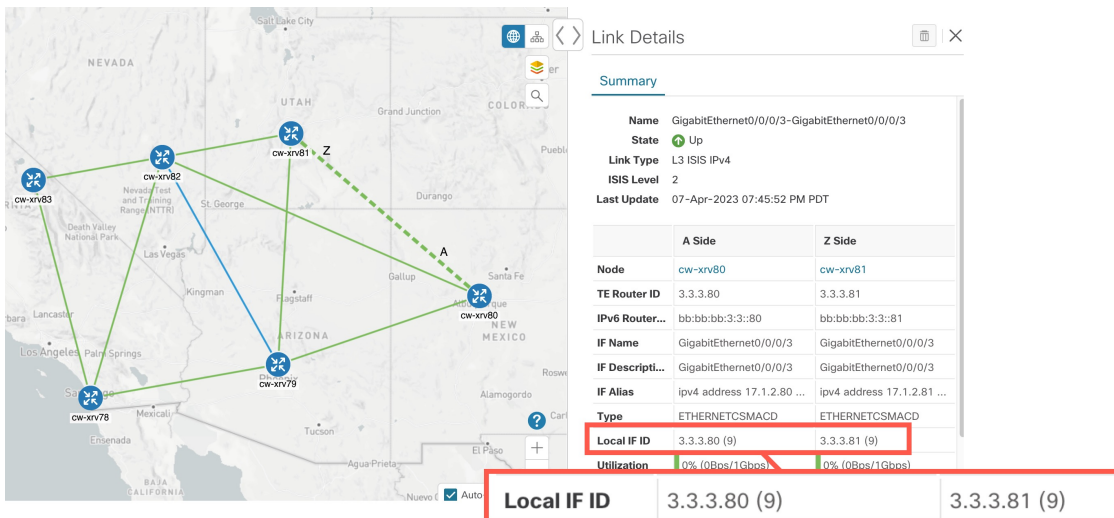
Note Dual stack links (although aggregate) are shown as one single line.

a) Under the **Link Type** column, click the link entry to see the link's details.



Step 4 View IPv4 unnumbered interface information (if available).

IPv4 unnumbered interfaces information is displayed as either an index or a combination of the TE Router ID and the index.



Step 5 View LAG details

To view different bundle members and member details in a Link Aggregation Group (LAG), confirm that LAG discovery is enabled (**Administration > Settings > System Settings** tab > **Discovery > LAG** checkbox):

Note It takes a few minutes for LAG collection to complete after LAG discovery is enabled.

a) Click on a LAG link. For example:

View Link Details

Links

Total 2

State	Link Type	A Si...	Z Si...	A Si...	Z Si...
🟢	L2 LAG	Bundl...	Bundl...	0% (...)	0% (...)
🟢	L2 CDP	Gigabi...	Gigabi...	0% (...)	0% (...)

b) Click the **Members** tab. In this example, only one link is displayed.

Link Details

Summary **Members**

Total 1

State	Link Type	A Si...	Z Si...	A Si...	Z Si...
🟢	L2 LAG MEM...	Gigabi...	Gigabi...	0% (...)	0% (...)

c) Click the LAG member link.

Link Details

Summary

Name GigabitEthernet0/0/0/3-GigabitEthernet0/0/0/3
State 🟢 Up
Link Type L2 LAG MEMBER
Last Update 25-Mar-2021 05:29:32 AM GMT+2

	A Side	Z Side
Node	P-BOTTOMRIGHT-L2	P-BOTTOMLEFT-L2
TE Router ID	101.101.101.4	101.101.101.3
IF Name	GigabitEthernet0/0/0/3	GigabitEthernet0/0/0/3
IF Description	GigabitEthernet0/0/0/3	GigabitEthernet0/0/0/3
Type	ETHERNETCSMACD	ETHERNETCSMACD
Utilization	0% (0Bps/1Gbps)	0% (0Bps/1Gbps)

Link States and Discovery Methods

Table 1: Link Types, Discovery and States

Link Type	Discovery	Link State
L3 link (ISIS and OSPF)	via SR-PCE	SR-PCE set it to UP or DOWN based on the link operational state
L2 link (CDP, LLDP, LAG)	via SNMP MIB: CDP, LLDP and LAG	<p>The link state is based on the two link endpoints operational states (via IF MIB).</p> <ul style="list-style-type: none"> • Link state is UP when initially discovered. • When one of the endpoint interfaces is operationally down, then the link state is set to DOWN. • When both endpoint interfaces are operationally up, then the link state is set to UP.

Table 2: Link State Definitions

Link State	Description
UP	Link is present in SR-PCE topology in both directions.
DEGRADED	Link is reported in SR-PCE topology in only one direction.
DOWN	Link is reported down in both directions.

Protocols Used for Topology Services

The following table lists the protocols and methods used for obtaining information.

Protocol/Method	Provides	Use Cases
IGP/ BGP-LS (via SR-PCE)	Real time topology (Nodes, links, link metrics, and so on.)	L3 topology visualization
PCEP(via SR-PCE)	Real time LSP status and CRUD of SR-PCE initiated LSPs	<ul style="list-style-type: none"> • SR/SRv6, RSVP-TE LSP visualization • SR-PCE initiated LSP Create/Update/Delete


Protocol/Method	Provides	Use Cases
SNMP(SNMPv2-MIB, IP-MIB, IF-MIB, LLDP-MIB, (CISCO CDB-MIB)(via CDG)	System info, Interface Table (interface and SR-TE/RSVP-TE traffic Utilization) IP Address Table, L2 adjacency information	Device management and details and Crosswork Optimization Engine model building: <ul style="list-style-type: none"> • L2/L3 Topology • Interface name, Admin/Oper Status • Interface and SR policy and RSVP-TE tunnel utilization Crosswork Optimization Engine model simulation: <ul style="list-style-type: none"> • IGP/LSP path simulation • Bandwidth use cases (for feature packs)
CLI (via CDG) -'show mpls	TE Router ID and so on.	To match the DLM node with the same TE Router ID that is learned from the SR-PCE

Customize Map Display Settings

You can configure visual settings on the topology map based on your needs and preferences. You can do the following:

- [Customize the Display of Links and Devices, on page 18](#)
- [Configure How Device Groups Are Displayed for Traffic Engineering , on page 34](#)

Customize the Display of Links and Devices

To set device and link map display preferences, choose **Topology** and click  on the topology map.

- Click **Links** to show aggregated links and how links should be colored so that you can easily see their state and utilization status. By default, aggregated links will be differentiated from single links on the map and links will be colored based on link utilization thresholds. Administrators can change the utilization thresholds and their corresponding colors.
- Click **Devices** to show the device state and how the devices should be labeled. By default, the device state is shown on the map and the host name is used to label devices.

Use Device Groups to Filter Your Topology View

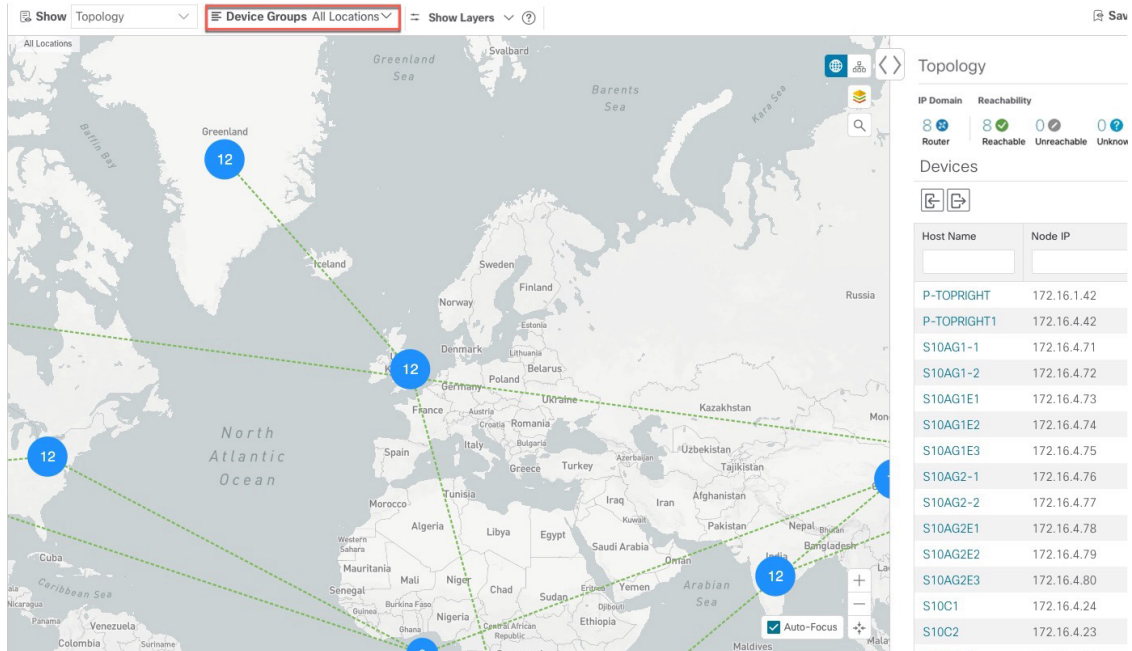
To help you identify, find, and group devices for a variety of purposes, you can create device groups. The device group window (**Device Management > Groups**) displays all devices and the device groups to which they belong. By default, all devices initially appear in the **Unassigned Devices** group.

To demonstrate the grouping and filtering functions, we have built an environment with devices distributed globally. You can sub-group the devices based on regions. For this example, we have a sub-group called **US West**.

Step 1 View devices on the geographical map:

a) From the main menu, choose **Topology**.

Note Devices without a geo-location appear in the **Devices** table only. To display these devices on the map, provide their geographical coordinates in the **Geo Location** column.



The screenshot shows the Cisco Crosswork Optimization Engine 5.0 interface. The main map displays a world map with several blue circular markers labeled '12' representing network devices. Green dashed lines connect these devices, representing network links. The 'Device Groups' dropdown menu is set to 'All Locations'. The 'Topology' panel on the right shows a table of devices with their host names and node IPs.

Host Name	Node IP
P-TOPRIGHT	172.16.1.42
P-TOPRIGHT1	172.16.4.42
S10AG1-1	172.16.4.71
S10AG1-2	172.16.4.72
S10AG1E1	172.16.4.73
S10AG1E2	172.16.4.74
S10AG1E3	172.16.4.75
S10AG2-1	172.16.4.76
S10AG2-2	172.16.4.77
S10AG2E1	172.16.4.78
S10AG2E2	172.16.4.79
S10AG2E3	172.16.4.80
S10C1	172.16.4.24
S10C2	172.16.4.23

b) From the **Device Group** drop-down list, select a group (US West). Only the devices in that group and related links are displayed on the geographical map. The Devices table has also been filtered to list only those devices in the group.


Use Device Groups to Filter Your Topology View

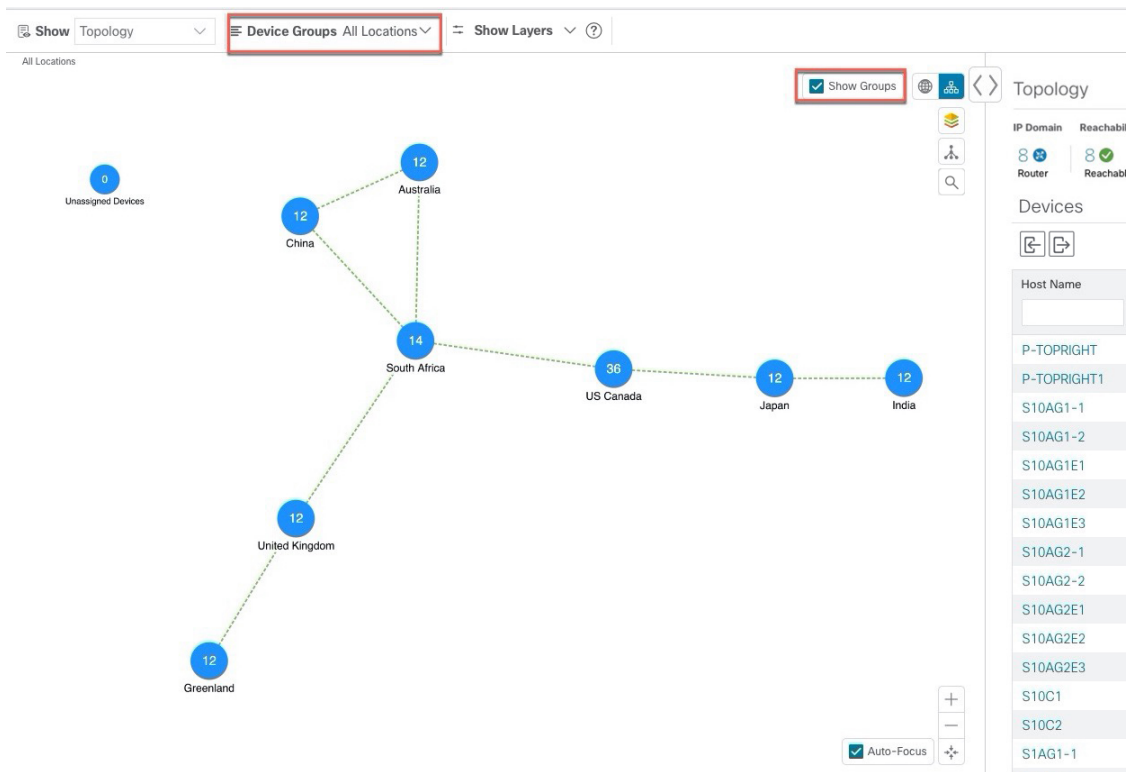
The screenshot shows the 'Topology' view in the Cisco Crosswork Optimization Engine 5.0. The 'Device Groups' dropdown is set to 'US West'. The logical map displays a network topology of devices (S7AG1E2, S7AG1E1, S7AG1E3, S7AG1-1, S7AG1-2, S7C1, S7C2, S7AG2E2) connected by lines. The right-hand panel shows a table of devices with their Host Name and Node IP. A red box highlights the following devices:

Host Name	Node IP
S7AG1-1	172.16.4.38
S7AG1-2	172.16.4.37
S7AG1E1	172.16.4.34
S7AG1E2	172.16.4.35
S7AG1E3	172.16.4.36
S7AG2-1	172.16.4.81
S7AG2-2	172.16.4.82
S7AG2E1	172.16.4.83
S7AG2E2	172.16.4.84
S7AG2E3	172.16.4.85
S7C1	172.16.4.46
S7C2	172.16.4.47

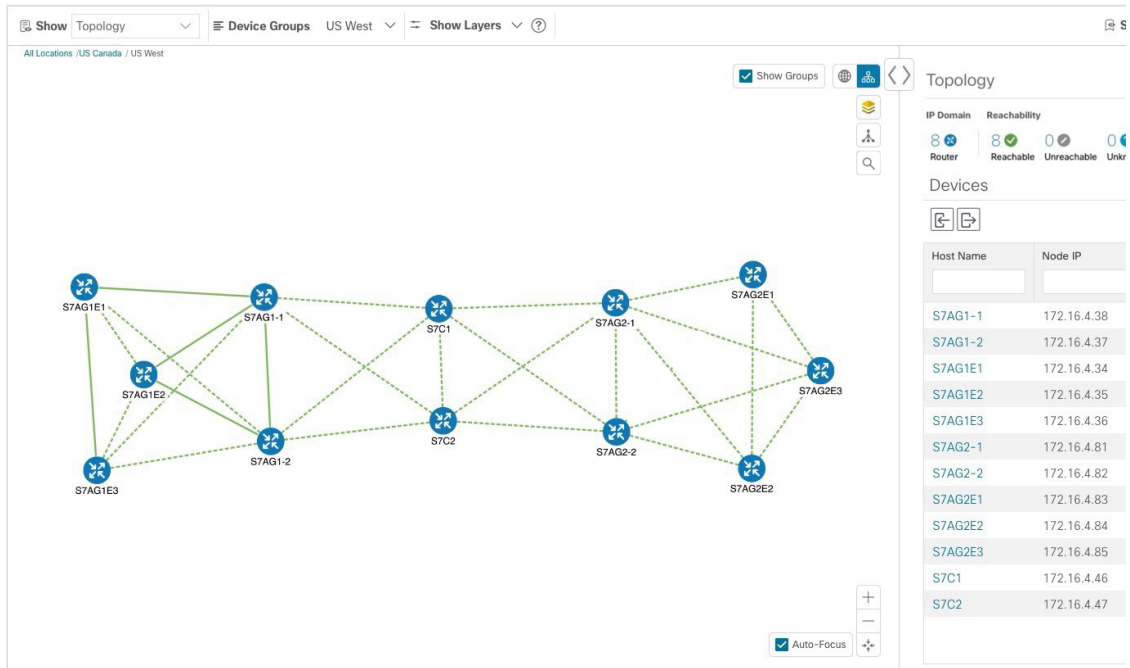
Step 2

View devices on the logical map:

- From the main menu, choose **Topology**.
- Click .
- From the **Device Group** drop-down list, select **All Locations** and check **Show Groups** if it is not already checked. You can see all device groups in this view. Device groups can be seen in this way only within the logical map.



d) From the **Device Group** drop-down list, select a group (US West). Devices that belong to this group are shown in the topology map and the **Devices** table.



e) Filter devices in the Device table by entering the partial host name or IP address in the text box (for example, **S7C** is entered in the **Host Name** text box for the current configuration). The Device table displays only devices that match

the filtering criteria. However, filtering the Device table does not filter the devices visually on the topology map. To visually filter devices on the geographical or logical maps is to use device groups.

Note You can also double click on the device in the list to recenter the selected device on the geographical or logical maps.


The screenshot displays the 'Topology' view in the Cisco Crosswork Optimization Engine 5.0. On the left, a network topology map shows various devices (STAG1E1, STAG1E2, STAG1E3, STAG1-1, STAG1-2, STAG2-1, STAG2-2, STAG2E1, STAG2E2, STAG2E3) connected by lines. On the right, a 'Devices' table is visible, showing a list of devices with columns for Host Name, Node IP, Oper..., Reac..., and Product Type. The table is filtered to show two devices: S7C1 and S7C2. The S7C1 device is highlighted in red in the original image.

Host Name	Node IP	Oper...	Reac...	Product Type
S7C1	172.16.4.46	OK	Re...	ciscoCRS16S
S7C2	172.16.4.47	OK	Re...	ciscoCRS16S

Create and Modify Device Groups

You can create device groups and devices to the groups either manually (as described in this section) or automatically, as described in [Enable Dynamic Device Grouping, on page 23](#). A device can belong to only one device group.

Step 1 From the main menu choose **Device Management > Groups**.

Step 2 To add a new sub-group, click  next to **All Locations**. A new sub-group is added under **All Locations**.

Step 3 To add a device to a group, from the right-pane, under **Unassigned Devices**, select a device.

Step 4 From the **Move to Group** drop-down, select the appropriate group and click **Move**.

The screenshot shows the 'Device Management / Groups' page. On the left, the 'Device Groups' section is visible, showing a tree view with 'All Locations' and 'Unassigned' sub-groups. On the right, the 'Unassigned Devices' table is displayed, showing a list of devices with columns for Device IP, Host Name, and Product Type. The table is filtered to show several devices with IP addresses ranging from 1.1.1.16 to 1.1.1.45.

Device IP	Host Name	Product Type
<input type="checkbox"/>		
<input type="checkbox"/>	1.1.1.16	
<input type="checkbox"/>	1.1.1.93	
<input type="checkbox"/>	1.1.1.67	
<input type="checkbox"/>	1.1.1.36	
<input type="checkbox"/>	1.1.1.43	
<input type="checkbox"/>	1.1.1.70	
<input type="checkbox"/>	1.1.1.45	

Step 5 To delete from a group, click **Remove from Group**. If you delete a group, all devices that belong to that group are moved to the Unassigned Devices group. Also, deleting a group deletes all the sub-groups under it.

Step 6 Click **Save**.

Enable Dynamic Device Grouping


You can create a rule to dynamically create device groups and automatically add unassigned devices to these groups using a Regular Expression (regex) on the device hostname. Any newly added or discovered devices that match the rule will be placed in the appropriate group.

Dynamic rules do not apply to devices that already belong to groups. You must move them to Unassigned Devices if you want them to be considered by the rule.

Before you begin

While you can follow examples given in the Dynamic Groups dialog, it is helpful to be familiar with Regular Expressions.

Step 1 From the main menu choose **Device Management > Groups**.

Step 2 Click  next to **All Locations > Manage Location Dynamic Groups**.

Step 3 Click **Show more details and examples** to help you fill out the required Host Name and Group Name fields.


Step 4 If there are any existing devices in the Unassigned Devices group, click **Test Rule** to view a sampling of what type of group names will be created.

Step 5 Turn the **Enable Rule** toggle ON to enable the rule. After the rule is enabled, the system checks for unassigned devices every minute and will assign them to the appropriate group based on the rule.

Step 6 Click **Save**.

Step 7 Groups that are created this way initially appear under Unassigned Groups (created when a rule is enabled for the first time). Move newly created groups to the desired group hierarchy.

Step 8 To move newly created Unassigned groups to the correct group, do the following:

- a) Click  next to All Locations and click **Add a Sub-Group**.
 - b) Enter the New Group details and click **Create**.
 - c) Click on the unassigned devices from the left pane.
 - d) From the right pane, select the devices you want to move and click **Move to Group** to move to an appropriate group.
-

Save Topology Views for Easy Access

When you rearrange the devices and links on a map, your changes are not normally saved. To easily access a useful map layout, you can save it as a named custom view and quickly retrieve it, without having to rearrange the map each time. This is especially useful when managing large networks with many devices.

When you save a custom view, the following settings will be saved:

- Whether it is a geographical or logical map.
- Device positions in the logical map layout.

Save Topology Views for Easy Access

- Device and link display settings
- Any filters used in the Device and Traffic Engineering tables

All custom views can be seen by all users. However, only users with the admin role or users that created the custom view can modify the view.

- Step 1** Customize the current map view until it contains just the information you want and the layout meets your needs.
- Step 2** When you have the view the way you want it, click **Save View**.

- Step 3** Enter a unique name for the new custom view and click **Save**. You can later modify the view (click **Select a saved view**) and choose to edit the topology, rename, or delete the view.



CHAPTER 3

Traffic Engineering in 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 guaranteed bandwidth routes.

Crosswork Optimization Engine allows you to visualize Traffic Engineering SR policies and RSVP tunnels discovered in your network, whether they were configured manually on the network devices or through the Crosswork UI. The following table lists what Traffic Engineering SR policies and RSVP tunnels can be visualized and provisioned through the Crosswork UI:

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

Table 3:

TE Technology	Crosswork Optimization Engine		Crosswork Network Controller	
	Visualize	Provision	Visualize	Provision
SR-MPLS	✓	✓	✓	✓
SRv6	✓	✗	✓	✓
RSVP	✓	✓	✓	✓
Flexible Algorithms	✓	✗ 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.

- [Segment Routing Path Computation Element \(SR-PCE\)](#), on page 26
- [What is Segment Routing?](#), on page 26

- [SR-TE Policy PCC and PCE Configuration Sources](#), on page 28
- [What is Resource Reservation Protocol \(RSVP\)?](#), on page 29
- [RSVP-TE Tunnel PCC and PCE Configuration Sources](#), on page 30
- [Get a Quick View of Traffic Engineering Services](#), on page 30
- [View TE Event and Utilization History](#), on page 32
- [View Traffic Engineering Device Details](#), on page 33
- [Configure Traffic Engineering Settings](#), on page 34

Segment Routing Path Computation Element (SR-PCE)

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

Cisco SR-PCE (formerly Cisco XR Traffic Controller (XTC)) runs on the Cisco IOS XR operating system. 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 that are 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 that are part of the IGP domain including those that do not establish PCEP peering with SR-PCE. However, PCEP peering is required to deploy TE tunnels to the device.



Note Features may not work as expected if the SR-PCE version is not supported. It is important to 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](#).

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 a 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. The prefix SID is manually configured from the segment routing global block (SRGB) range of labels, and is distributed by IS-IS or OSPF. The prefix segment steers the traffic along the shortest path to its destination. A node SID is a special type of prefix SID that

identifies a specific node. It is configured under the loopback interface with the loopback address of the node as the prefix.

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

- An adjacency segment is identified by a label called an adjacency SID, which represents a specific adjacency, such as egress interface, to a neighboring router. The adjacency SID is distributed by IS-IS or OSPF. The adjacency segment steers the traffic to a specific adjacency.

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

By combining prefix (node) and adjacency segment IDs in an ordered list, any path within a network can be constructed. At each hop, the top segment is used to identify the next hop. Segments are stacked in order at the top of the packet header. When the top segment contains the identity of another node, the receiving node uses equal cost multipaths (ECMP) to move the packet to the next hop. When the identity is that of the receiving node, the node pops the top segment and performs the task required by the next segment.

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, and instructs the routers in the network to follow the specified path instead of the shortest path calculated by the IGP. If a packet is steered into an SR policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

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

- [SR-MPLS and SRv6](#) , on page 37
- [Flexible Algorithms](#), on page 65
- Circuit Style (when [SR Circuit Style Manager \(CSM\)](#), on page 81 is enabled)
- [Tree Segment Identifier \(Tree-SID\) Multicast Traffic Engineering](#), on page 71



Note Crosswork discovers existing SR policies when devices are imported, but cannot manage them. SR policies can be managed only if they were provisioned using the UI.

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 you configure an explicit policy, you specify an explicit path which consists of a list of prefix or adjacency SIDs, each representing a node or link along on 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, information received in both requests is used to compute two disjoint paths: one path from the first source to the first destination, and another path from the second source to the second destination. Both paths are computed at the same time. The shortest lists of segments is 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.

SR-TE Policy PCC and PCE Configuration Sources

SR-TE policies discovered and reported by Crosswork Optimization Engine 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 28](#)). This policy type displays as **Unknown** in the UI.

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
```


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 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. This information is used by the headend 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 Optimization Engine 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 30](#)).
- 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**.

The screenshot shows the TE Dashboard with the following data:

Service	Total Policy Count	Oper Down	Admin Down	Oper Up
SR-MPLS	4	1	0	3
SRv6	0	0	0	0
Tree-SID	3	0	0	3
RSVP-TE	1	1	0	0

Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	Traffic Rate (Kbps)
xrv9k-22	xrv9k-26	2023	SR-MPLS	TE	0
xrv9k-22	xrv9k-23	7777	SR-MPLS	Unknown	0
xrv9k-22	xrv9k-24	100	RSVP-TE	Unknown	0
xrv9k-25	xrv9k-23	15130	SR-MPLS	IGP	0


Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	Total	Operational State Change	Path Change	Actions
xrv9k-25	xrv9k-23	15130	SR-MPLS	IGP	2	1	1	
xrv9k-22	-	-	Tree-SID	IGP	1	1	0	

523195



Note If you are viewing the HTML version of this guide, click on 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 SR-MPLS, BWoD and LCM policies and the number of policies/tunnel 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 that you clicked on.</p>


Callout No.	Description
2	<p>Policies and Tunnels Under Traffic Threshold for Historic Data:</p> <p>Displays RSVP-TE tunnels and SR-MPLS policies that have 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 underutilized LSP threshold value.</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:

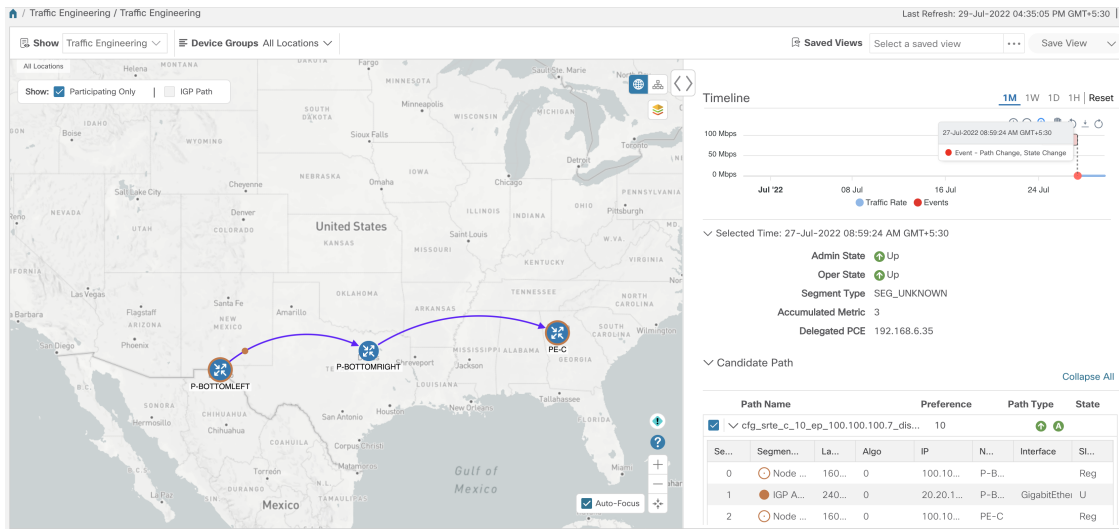
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering** .
- Step 2** From the **Actions** column of the Traffic Engineering table, click  > **View Details > Historical Data** tab for a policy or tunnel. The tab displays associated historical data for that device. The following example shows the traffic rate and event history for an SR-MPLS policy.



523228

Callout No.	Description
1	<p>Traffic Rate: Displays the traffic rate for the policies.</p> <p>Note Traffic Rate is not captured for SRv6 and Tree-SID policies.</p>
2	<p>Events:</p> <p>Displays the path or state change event.</p>

Step 3 Click the event, to view the state of the policy or tunnel as shown in the following image:
The policy path is displayed in the left pane.



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** page, click on the policies you are interested in. Each tab displays associated data for that device.

The following example shows SR-MPLS Prefix information which includes the MSD value for the device.

Device Details

Details Links **Traffic Engineering**

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

IGP: Domain ID: 1000, ISIS System ID: 0000.0000.0004, Level: 2

SR-MPLS

TE Router ID 192.168.0.24

SRGB 16000 - 23999

SRLB 15000 - 15999

MSD 10 ⓘ

Prefixes	Label	Algo
192.168.0.24	18114	0

SRv6

PCEP Sessions

PCE : 172.27.226.126, PCC/Source - 192.168.0.24

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 > Traffic Engineering > General Settings** tab. Enter the timeout duration options. For more information, click ⓘ.



Note Timeouts change the response time of each of the actions 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.

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 and select one of the behavior options.

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

Configure TE Dashboard Settings

To configure the TE Dashboard (and Historical Data) settings for the collection of policy and tunnel metrics, state changes, path changes, data retention interval, and the utilization threshold for underutilized LSPs, select **Administration > Settings > System Settings tab > Performance Monitoring & Analytics > Historical Data**.

The screenshot shows the 'System Settings' tab with the 'Historical Data' sub-tab selected. The 'TE Historical Data Settings' section includes the following controls:

- LSP Traffic Rate:** A toggle switch set to 'On'.
- LSP State Change:** A toggle switch set to 'On'.
- LSP Path Change:** A toggle switch set to 'On'.
- Retention Interval:** A text input field containing the value '2', with a help icon (question mark) to its right. Below the field, it says 'Range: 1 to 30 days'.

At the bottom of the settings area, there are three buttons: 'Save', 'Reset', and 'Reset to Default'.

Historical Data Settings	Description
LSP Traffic Rate	Turn on this field to capture the metric data in the TE Dashboard.
LSP State Change	Turn on this field to capture the state change details in the TE Dashboard.
LSP Path Change	Turn on this field to capture the path change details in the TE Dashboard.
Retention Interval	The interval for which the historical data is collected and retained before being deleted. The default retention interval is set to two days. Note If the Retention Interval is reduced, all data older than the new retention interval is lost. For example, if the retention interval is set to 30 days and later it is reduced to 7 days, all the data older than 7 days will be deleted.



CHAPTER 4

SR-MPLS and SRv6

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

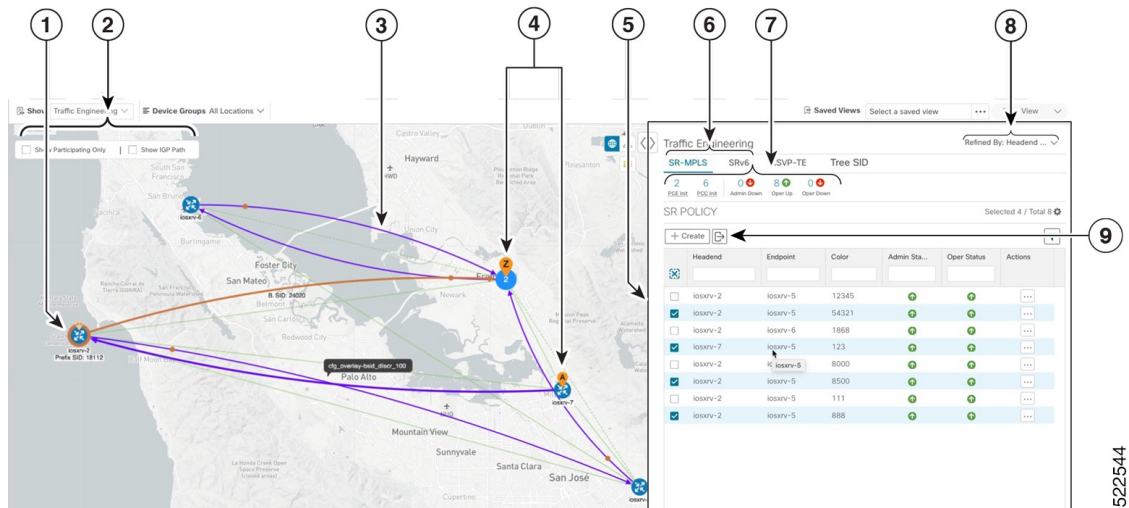
- [View SR-MPLS and SRv6 Policies on the Topology Map, on page 37](#)
- [View SR-MPLS and SRv6 Policy Details, on page 39](#)
- [Visualize SR-MPLS or SRv6 Policies Example, on page 40](#)
- [Find Multiple Candidate Paths \(MCPs\), on page 46](#)
- [Visualize Underlying Paths Associated with a Defined Binding-Segment ID \(B-SID\) Label, on page 49](#)
- [Visualize Native SR Paths, on page 51](#)
- [Configure TE Link Affinities, on page 54](#)
- [Create Explicit SR-MPLS Policies, on page 55](#)
- [Create Dynamic SR-MPLS Policies Based on Optimization Intent, on page 56](#)
- [Modify SR-MPLS Policies, on page 58](#)



View SR-MPLS and SRv6 Policies on the Topology Map

Crosswork Optimization Engine visualization provides the most value by giving you the ability to easily view and manage SR-MPLS and SRv6 policies. By visually examining your network, the complexity of provisioning and managing these SR-TE policies is significantly reduced.

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

Figure 6: Traffic Engineering UI : SR-MPLS and SRv6 Policies



Callout No.	Description
1	A device with an orange () outline indicates there is a node SID associated with that device or a device in the cluster.
2	Click the appropriate check box to enable the following options: <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.
3	When SR-TE policies are selected in the SR-MPLS or SRv6 tables, they show as purple directional lines on the map indicating source and destination. An adjacency segment ID (SID) is shown as an orange circle on a link along the path ().
4	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.
5	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.
6	Click on either the SR-MPLS or SRv6 tabs to view the respective list of SR-TE policies.
7	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.

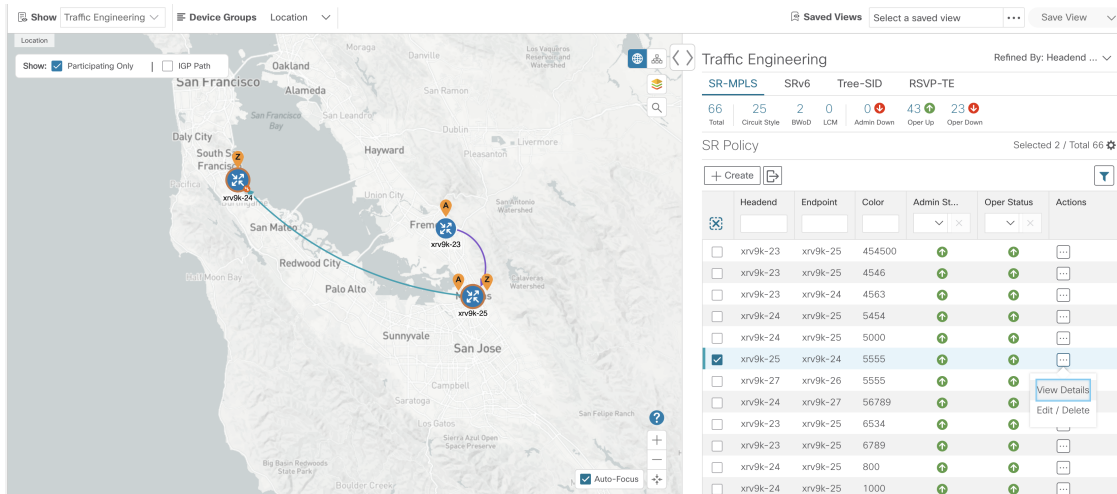
522544

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 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.


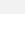
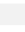



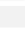
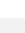
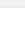

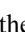

View SR-MPLS and SRv6 Policy Details

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.



The screenshot displays the Traffic Engineering interface. On the left, a map shows the San Francisco Bay Area with several nodes (xrv9k-24, xrv9k-25) and a path highlighted. On the right, a table lists SR-MPLS policies. The table has the following columns: Headend, Endpoint, Color, Admin St..., Oper Status, and Actions. The selected policy is xrv9k-25 (Headend) to xrv9k-24 (Endpoint) with a color of 5555. A tooltip over the Actions column for this policy shows 'View Details', 'Edit', and 'Delete' options.

Headend	Endpoint	Color	Admin St...	Oper Status	Actions
<input type="checkbox"/>	xrv9k-23	xrv9k-25	454500	●	● 
<input type="checkbox"/>	xrv9k-23	xrv9k-25	4546	●	● 
<input type="checkbox"/>	xrv9k-23	xrv9k-24	4563	●	● 
<input type="checkbox"/>	xrv9k-24	xrv9k-25	5454	●	● 
<input type="checkbox"/>	xrv9k-24	xrv9k-25	5000	●	● 
<input checked="" type="checkbox"/>	xrv9k-25	xrv9k-24	5555	●	●  View Details
<input type="checkbox"/>	xrv9k-27	xrv9k-26	5555	●	●  Edit / Delete
<input type="checkbox"/>	xrv9k-24	xrv9k-27	56789	●	● 
<input type="checkbox"/>	xrv9k-23	xrv9k-25	6534	●	● 
<input type="checkbox"/>	xrv9k-23	xrv9k-25	6789	●	● 
<input type="checkbox"/>	xrv9k-24	xrv9k-25	800	●	● 
<input type="checkbox"/>	xrv9k-24	xrv9k-25	1000	●	● 

Step 2 View SR-MPLS or SRv6 policy details.

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.

The screenshot displays the 'SR Policy Details' window. It features two tabs: 'Current' (selected) and 'History'. Under the 'Current' tab, the 'Endpoint' is 'xrv9k-24' with a destination IP of '192.168.0.24'. Below this, it shows 'TE RID: 192.168.0.24 | IPv6 RID: 2001:192:168::24' and 'Color 5555'. A 'Summary' section is expanded, listing various policy attributes:

- Admin State: Up (green arrow icon)
- Oper State: Up (green arrow icon)
- Binding SID: 24027
- Policy Type: Regular
- Profile ID: -
- Description: -
- Traffic Rate: 0 Mbps
- Unused: True (info icon)
- Delay: 532 (info icon)
- Bandwidth Constraint: 0 Mbps
- Accumulated Metric: 1
- Delegated PCE: 172.27.226.126
- Non-delegated PCEs: -
- PCE Computed Time: 06-Feb-2023 03:11:16 PM GMT+5:30
- Last Update: 07-Mar-2023 11:45:11 AM GMT+5:30

A tooltip for the 'Last Updated' field shows the timestamp '09-Mar-2023 02:22:54 PM GMT+5:30'. At the bottom of the summary, there is a 'See less' link with an upward arrow icon.

Visualize SR-MPLS or SRv6 Policies Example

This example walks you through several SR-TE (SR-MPLS and SRv6) policy visualization features that are available from the topology map.


In this example, we assume that devices and SR-MPLS policies have been added and device groups have been created.



Note Although this example uses SR-MPLS policies, the basic functionality of the maps for both SR-MPLS policies and SRv6 policies is the same.

Click images to zoom in for a closer look.

Figure 7: Topology Map Example



Step 1

Select SR-MPLS policies for visualization and isolate them on the map.

- From the main menu, choose **Traffic Engineering > Traffic Engineering**.
- From the **SR Policy** table, check the check box next to the SR-MPLS policies you are interested in.
- Check the check box next to **Show Participating Only** so that other links and devices that are not part of the selected SR-TE policies are hidden.

In the following example, the topology map displays the following:

- Four SR-MPLS policies are selected.
- SR-MPLS policies appear as purple links with arrows that indicate the path direction.
- The **PCC1_77** node is the destination for two of the selected policies. Both **PCC5_81** and **PCC3_79** are destinations for the selected policies. SR-MPLS policy origin and destination are marked with **A** and **Z**, respectively. The **A+** denotes that there is more than one policy that originates from a device. A **Z+** denotes that the device is a destination for more than one policy.

Visualize SR-MPLS or SRv6 Policies Example

	Headend	Endpoint	Color	Admin St...	Oper Status	Actions
<input type="checkbox"/>	PCC1_77	PCC2_78	112	🟢	🟢	⋮
<input type="checkbox"/>	PCC3_79	PCC5_81	44	🟢	🟢	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	879	🟢	🟢	⋮
<input type="checkbox"/>	PCC3_79	PCC5_81	99	🟢	🟢	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	50	🟢	🟢	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	51	🟢	🟢	⋮
<input checked="" type="checkbox"/>	PCC5_81	PCC1_77	520	🟢	🟢	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	10555	🟢	🔴	⋮
<input checked="" type="checkbox"/>	PCC1_77	PCC2_78	10555	🟢	🔴	⋮
<input checked="" type="checkbox"/>	PCC2_78	PCC1_77	10666	🟢	🔴	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	10666	🟢	🔴	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	10777	🟢	🔴	⋮

Step 2

Highlight and view more details for a particular SR-MPLS policy.


- a) From the **SR Policy** table, *hover* over a selected policy. The topology map displays the following details:

- The path is emphasized on the map. The path goes through **PCC5_81 > PCC3_79 > PCC1_77**.
- The prefix SID for PCC3_79 and PCC1_77 is displayed.

	Headend	Endpoint	Color	Admin St...	Oper Status	Actions
<input type="checkbox"/>	PCC1_77	PCC2_78	112	🟢	🟢	⋮
<input type="checkbox"/>	PCC3_79	PCC5_81	44	🟢	🟢	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	879	🟢	🟢	⋮
<input type="checkbox"/>	PCC3_79	PCC5_81	99	🟢	🟢	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	50	🟢	🟢	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	51	🟢	🟢	⋮
<input checked="" type="checkbox"/>	PCC5_81	PCC1_77	520	🟢	🟢	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	10555	🟢	🔴	⋮
<input checked="" type="checkbox"/>	PCC1_77	PCC2_78	10555	🟢	🔴	⋮
<input checked="" type="checkbox"/>	PCC2_78	PCC1_77	10666	🟢	🔴	⋮
<input type="checkbox"/>	PCC1_77	PCC2_78	10666	🟢	🔴	⋮
<input type="checkbox"/>	PCC2_78	PCC1_77	10777	🟢	🔴	⋮

Step 3

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

- a) Check the **Show IGP Path** check box. The IGP paths for the selected SR-MPLS policies are displayed, with straight lines, instead of the segment hops.
- b) Click .
- c) Click the **Metrics** tab.
- d) Toggle applicable metrics to **ON**.

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

Headend	Endpoint	Color	Admin St...	Oper Status	Actions
<input type="checkbox"/>	PCC1_77	PCC2_78	112	↑	↑
<input type="checkbox"/>	PCC3_79	PCC5_81	44	↑	↑
<input type="checkbox"/>	PCC1_77	PCC2_78	879	↑	↑
<input type="checkbox"/>	PCC3_79	PCC5_81	99	↑	↑
<input type="checkbox"/>	PCC2_78	PCC1_77	50	↑	↑
<input type="checkbox"/>	PCC1_77	PCC2_78	51	↑	↑
<input checked="" type="checkbox"/>	PCC5_81	PCC1_77	520	↑	↑
<input type="checkbox"/>	PCC2_78	PCC1_77	10555	↑	↓
<input checked="" type="checkbox"/>	PCC1_77	PCC2_78	10555	↑	↓
<input checked="" type="checkbox"/>	PCC2_78	PCC1_77	10666	↑	↓
<input type="checkbox"/>	PCC1_77	PCC2_78	10666	↑	↓
<input type="checkbox"/>	PCC2_78	PCC1_77	10777	↑	↓

Step 4

View SR-MPLS policy details such as disjoint groups, metric type, segment hop information, delay (calculated for all policies every 10 minutes), and so on.

- From the **Actions** column, click > **View Details** for one of the SR-MPLS policies. The **SR Policy Details** window is displayed in the side panel. Note that only the selected policy is displayed on the topology map.

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> > 1111_path	100	Unknown	↑

Step 5

Customize and save a logical view of the topology.

- Click to display the logical view of selected SR-MPLS policies.
- Arrange the nodes to your preference.
- To save the topology layout (*not SR-MPLS policy selection*), clear all selected SR-MPLS policies, and click **Save View**.

Example:

Figure 8: Logical Map (Save Without SR-MPLS Policies Selected)

The screenshot shows the Cisco Crosswork Optimization Engine 5.0 interface. On the left, a logical map displays several interconnected nodes labeled iosxrv-2, iosxrv-5, iosxrv-6, iosxrv-4, iosxrv-7, iosxrv-3, and xtc-iosxrv. On the right, the 'Traffic Engineering' configuration page is shown with the 'SR-MPLS' tab selected. The 'SR POLICY' table is visible, showing a list of policies with columns for Headend, Endpoint, Color, Admin Status, and Oper Status. A 'Save View' button is highlighted in the top right corner.

522547

Step 6 Close (X) the current view to return to the **SR Policy** table.

Step 7 To understand how device groups are displayed with the selection of SR-MPLS policies, uncheck any SR-MPLS policies that might be selected and check **Show Groups**.

The screenshot shows the Cisco Crosswork Optimization Engine 5.0 interface. On the left, a logical map displays several interconnected nodes labeled Australia, China, South Africa, US Canada, Japan, India, and Greenland. On the right, the 'Traffic Engineering' configuration page is shown with the 'SR-MPLS' tab selected. The 'SR POLICY' table is visible, showing a list of policies with columns for Headend, Endpoint, Color, Admin Status, and Oper Status. A 'Show Groups' button is highlighted in the top right corner.

522548

Step 8 Selecting a specific group from the **Device Groups** drop-down list, will only display that group in the map. In this example, **Australia** is selected and the associated SR-MPLS policy is selected and displayed.

The screenshot shows the Traffic Engineering interface. On the left is a network diagram with nodes labeled S1AG1E1, S1AG1E2, S1AG1E3, S1AG1-1, S1AG1-2, S1C1, S1AG2-1, S1AG2-2, S1AG2E1, S1AG2E2, and S1AG2E3. On the right is a table of SR Policies:

Headend	Endpoint	Color	Admin St...	Oper Stat...	Actions
<input type="checkbox"/>	S1AG1E1	SSAG1E2	63212	✔	✔
<input type="checkbox"/>	S1C2	S2C1	100	✔	✔
<input checked="" type="checkbox"/>	S1AG2E1	S1AG1E2	100	✔	✔
<input type="checkbox"/>	SSC1	S1C2	111	✔	✘
<input type="checkbox"/>	S7C1	S1C2	202	✔	✘
<input checked="" type="checkbox"/>	S1AG1-2	S1C2	4521	✔	✔
<input type="checkbox"/>	S10AG1-1	S1AG1-1	3256	✔	✘

522549

Step 9

If you select a policy where participating devices are not part of the selected group, then a dialog appears giving you an option to switch the group view. This is the default behavior. If this window does not appear, then the administrator has configured the display to automatically switch view or stay in the current view. For more information, see [Configure How Device Groups Are Displayed for Traffic Engineering](#), on page 34.

The screenshot shows the same Traffic Engineering interface as in Step 9, but with a dialog box overlaid. The dialog box contains the following text:

Some of the participating devices are not in the current device group.
Click "Switch Device Group" to automatically switch to the device group that will show all participating devices.

Buttons: **Switch Device Group** (highlighted), **Don't Switch**

Checkbox: Don't show this message again

Step 10

If you select **Switch Device Group**, then the group will change and you will see all participating devices for the SR-MPLS policies you have selected.

Find Multiple Candidate Paths (MCPs)

To go back to the previous group view, click **Back** (this link appears later in the yellow text area indicated in the following figure).

The screenshot shows the Cisco Crosswork Network Automation interface. On the left, there is a navigation sidebar with options like Home, Topology, Traffic Engineering, Device Management, and Administration. The main area displays a network topology with nodes representing different regions: United Kingdom, Australia, South Africa, US Canada, Japan, India, and China. A path is highlighted in purple. On the right, the 'Traffic Engineering' section is visible, showing a summary for SR-MPLS (15 PCE Init, 6 PCC Init, 0 Admin Down, 17 Oper Up, 4 Oper Down) and a table of SR Policies. The table has columns for Headend, Endpoint, Color, Admin St., Oper Stat., and Actions. Several policies are listed, including S1AG1E1, S8AG1-1, S3AG1-1, S3C1, S10C1, S10AG2E3, S2AG1-1, S10AG1-1, S2AG1E3, S10AG1-1, SSC2, S1C2, S2AG1-1, S1AG2E1, and SSC1.

522551

Step 11 You can also use the Mini Dashboard to drill down and focus on certain SR-TE policies.

To filter the SR Policy table to show only PCE-initiated policies, click the value for PCE Init from the SR-MPLS Mini Dashboard. Note that the **Filters Applied** text appears.

The screenshot shows the same Cisco Crosswork Network Automation interface, but with a different network topology. The nodes are labeled with device names like PE4-ASRtk.cisco.com, P3-NCS5501, PE1-ASRtk, PE3-ASRtk, PE2-ASRtk, PE1-XRtk, PE2-XRtk, PE3-XRtk, and PE4-ASRtk. On the right, the 'Traffic Engineering' section shows a summary for SR-MPLS (4 PCE Init, 1 PCC Init, 0 Admin Down, 5 Oper Up, 0 Oper Down). The SR Policy table is filtered to show only PCE-initiated policies. The table has columns for Headend, Endpoint, Color, Admin St., Oper Stat., and Actions. The filtered policies are: PE1-AS..., P3-NCS..., 345; PE4-AS..., PE7-XR..., 123; PE7-XR..., P4-NCS..., 234; and PE4-AS..., PE2-AS..., 2258. A 'Filters Applied (1)' dropdown is visible above the table.

522552

Step 12 To remove filter criteria, click **Filters Applied > Clear All Filters**. You can also select individual filters if more than one filter has been applied.

Find Multiple Candidate Paths (MCPs)

Visualizing MCPs gives you insight into which paths might be a better alternative to the currently active one. 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 Optimization Engine does not distinguish dynamic paths versus 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

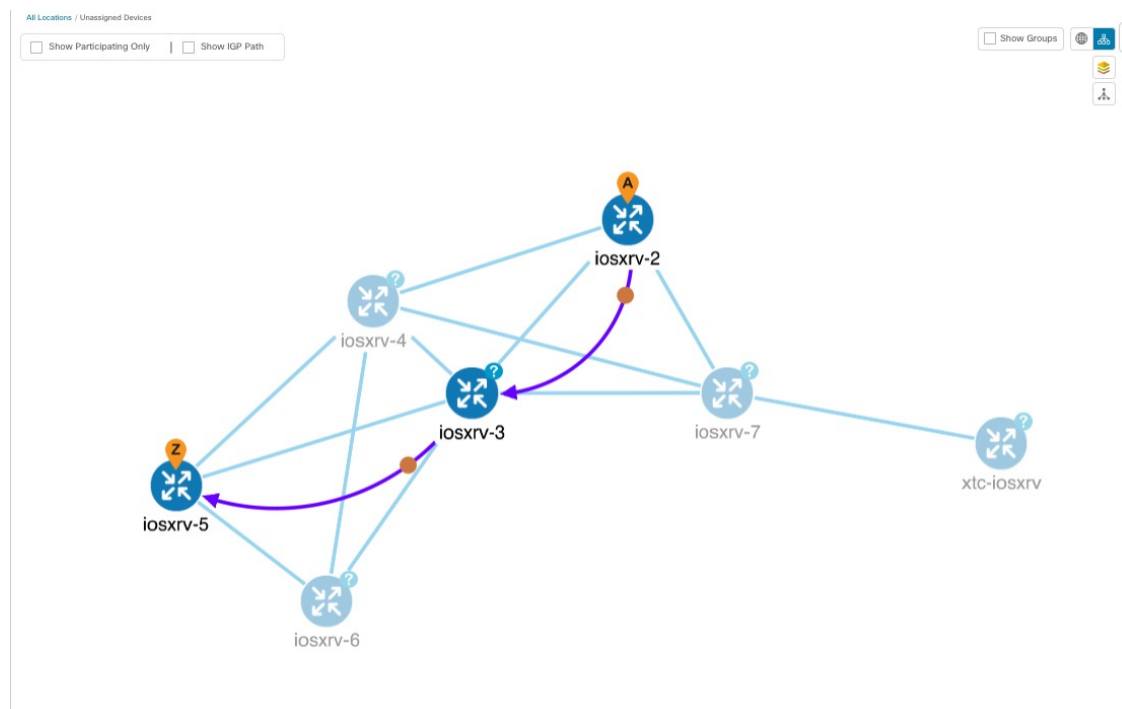
A policy must be configured with MCPs on devices before visualizing them on the Traffic Engineering topology map. This configuration can be done manually or within Crosswork Network Controller.

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.

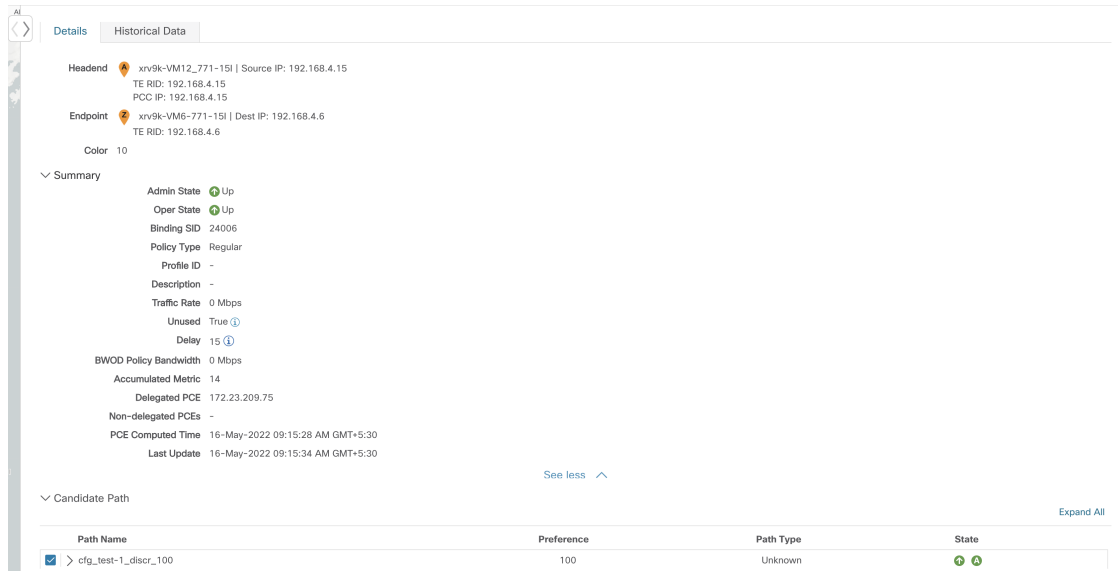
In this example, you see that the active path is going from **iosxrv-2 > iosxrv-3 > iosxrv-5**.



Step 3 View the list of candidate paths.

Find Multiple Candidate Paths (MCPs)

- a) From the SR-TE 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 in the state column indicates the active



The screenshot shows the 'SR Policy Details' window with the following information:

- Headend:** xrv9k-VM12_771-151 | Source IP: 192.168.4.15
TE RID: 192.168.4.15
PCC IP: 192.168.4.15
- Endpoint:** xrv9k-VM6-771-151 | Dest IP: 192.168.4.6
TE RID: 192.168.4.6
- Color:** 10
- Summary:**
 - Admin State: Up
 - Oper State: Up
 - Binding SID: 24006
 - Policy Type: Regular
 - Profile ID: -
 - Description: -
 - Traffic Rate: 0 Mbps
 - Unused: True
 - Delay: 15
 - BWOD Policy Bandwidth: 0 Mbps
 - Accumulated Metric: 14
 - Delegated PCE: 172.23.209.75
 - Non-delegated PCEs: -
 - PCE Computed Time: 16-May-2022 09:15:28 AM GMT+5:30
 - Last Update: 16-May-2022 09:15:34 AM GMT+5:30
- Candidate Path Table:**

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> > cfg_test-1_discr_100	100	Unknown	Up

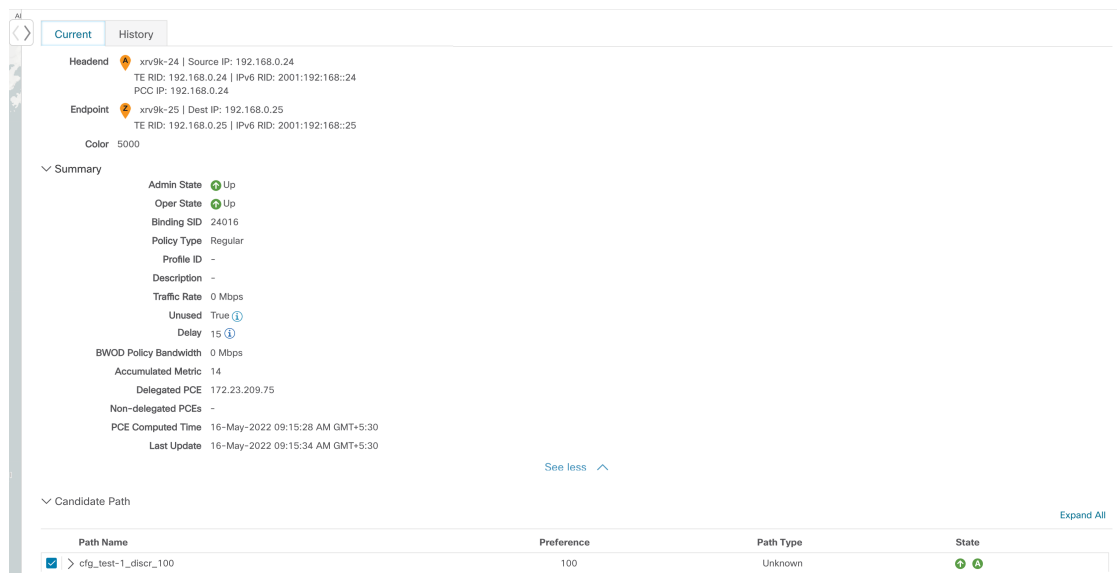
path.

- Step 4** You can expand individual paths or click **Expand All** to view details of each path. As you hover each segment, the segment is highlighted on the map.

- 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.



The screenshot shows the 'SR Policy Details' window with the following information:

- Headend:** xrv9k-24 | Source IP: 192.168.0.24
TE RID: 192.168.0.24 | IPv6 RID: 2001:192:168::24
PCC IP: 192.168.0.24
- Endpoint:** xrv9k-25 | Dest IP: 192.168.0.25
TE RID: 192.168.0.25 | IPv6 RID: 2001:192:168::25
- Color:** 5000
- Summary:**
 - Admin State: Up
 - Oper State: Up
 - Binding SID: 24016
 - Policy Type: Regular
 - Profile ID: -
 - Description: -
 - Traffic Rate: 0 Mbps
 - Unused: True
 - Delay: 15
 - BWOD Policy Bandwidth: 0 Mbps
 - Accumulated Metric: 14
 - Delegated PCE: 172.23.209.75
 - Non-delegated PCEs: -
 - PCE Computed Time: 16-May-2022 09:15:28 AM GMT+5:30
 - Last Update: 16-May-2022 09:15:34 AM GMT+5:30
- Candidate Path Table:**

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> > cfg_test-1_discr_100	100	Unknown	Up

- 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 **iosxrv-2** > **iosxrv-5**.

The screenshot displays the Cisco Crosswork Optimization Engine interface. On the left, a network topology map shows several nodes (iosxrv-2, iosxrv-3, iosxrv-4, iosxrv-5, iosxrv-6, iosxrv-7, and xtc-iosxrv) connected by lines. A path is highlighted in orange, labeled "Candidate Path". A tooltip for this path shows "cfg_test_mcp_diff_paths_discr_5000". On the right, the "SR Policy Details" panel is open, showing a table of candidate paths. The path "cfg_test_mcp_diff_paths_discr_5000" is selected and highlighted in red. Below the table, the details for this path are shown, including the path name, policy type, metric type, disjoint group, PCE initiated status, and affinity settings.

Path Name	Preference	Path Type
cfg_test_mcp_diff_paths_discr_10000	10000	Unknown
cfg_test_mcp_diff_paths_discr_5000	5000	Unknown

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

Crosswork Optimization Engine 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** or **SRv6** tab.
- Step 2** Check the check box next to the SR-MPLS policy that contains a hop assigned with a B-SID label and 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**.

Note Click image examples to zoom in for a closer look.

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

The screenshot shows the Traffic Engineering interface. On the left, a map of the United States displays network paths between various locations. A path is highlighted from a headend node (orange) to an endpoint node (blue) via several intermediate nodes. A specific path is labeled with a B-SID: 15700. On the right, the SR Policy table is visible, showing a list of policies with columns for Headend, Endpoint, Color, Admin State, Oper State, and Actions. The policy 'cw-xrv23' is selected, and its details are shown in a red box.

Headend	Endpoint	Color	Admin...	Oper ...	Actions	
<input type="checkbox"/>	cw-xrv51	cw-xrv52	3333	↑	↑	...
<input type="checkbox"/>	cw-asr23	cw-xrv62	4455	↑	↓	...
<input type="checkbox"/>	cw-xrv52	cw-xrv54	2700	↑	↑	...
<input type="checkbox"/>	cw-xrv61	cw-xrv58	2900	↑	↑	...
<input checked="" type="checkbox"/>	cw-asr23	cw-xrv52	2222	↑	↑	...
<input type="checkbox"/>	cw-xrv51	cw-xrv52	300	↑	↑	...
<input type="checkbox"/>	cw-xrv50	cw-xrv52	2022	↑	↓	...
<input type="checkbox"/>	cw-xrv62	cw-xrv53	2001	↑	↑	...
<input type="checkbox"/>	cw-xrv60	cw-xrv61	2001	↑	↑	...

Step 3 From the **Actions** column, click **⋮** > **View Details**.

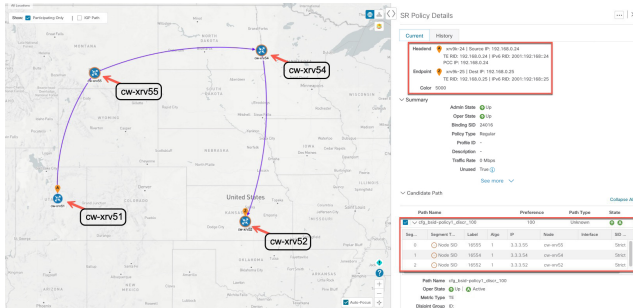
Step 4 From the **SR Policy Details** window, expand the active path name and click the **B-SID label**. In this example, the B-SID label is **15700**.

The screenshot shows the SR Policy Details window. The 'Current' tab is selected. The Headend is 'xrv9k-24' and the Endpoint is 'xrv9k-25'. The Summary section shows the Admin State is 'Up', Oper State is 'Up', Binding SID is '24016', Policy Type is 'Regular', Profile ID is '-', Description is '-', Traffic Rate is '0 Mbps', and Unused is 'True'. The Candidate Path section is expanded, showing a table of paths. The path 'cfg_test-bsid-policy_discr_100' is selected, and its details are shown below the table.

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> √ cfg_test-bsid-policy_discr_100	100	Unknown	↑ A
S...	Segme...	L...	Algo
0	Nod...	16...	1
1	IGP ...	24...	0
2	B-Sid	15700	3.3.3.51

Path Name: cfg_test-bsid-policy_discr_100
 Oper State: ↑ Up | A Active
 Metric Type: TE

- Step 5** In the **SR Policy Details** window for the underlying path, expand the active path name to view more details. In this example, you see the underlying path actually goes from **cw-xrv51 > cw-xrv55 > cw-xrv54 > cw-xrv52**.



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 SR-MPLS policies.

Before you begin

Confirm that device requirements are met. See [Visualize Native Path Device Prerequisites, on page 53](#).

To create a path query, do the following:

- Step 1** From the main menu, choose **Traffic Engineering > Path Query**.
- Step 2** On the Query Path Dashboard, click **New Query**.
- Step 3** Under the New Path Query, select the required values and click **Get Paths**.
- Step 4** Click **View Result** to view the query result.
- Step 5** (Optional) On the result pop-up click, **View Past Result**. Check the query ID to view the available results.

Example:

In the below example, you can view the available paths : **Path 0**

Figure 9: Path Details

The screenshot shows a map of the United States with a path highlighted between PE-A (San Francisco) and PE-B (Boston). The Path Details panel on the right shows the following information:

Select from the fields below to find available Native SR IGP Paths * Required field

Select Service: Select Type | Select Instance

Headend * PE-A (100.100.100.5)

Endpoint * PE-B (100.100.100.6)

Get Paths

Available Paths

Path 0

Output: tunnel-te220
 Nexthop: 20.20.10.2
 Source: 100.100.100.5
 Destination: 127.0.0.0

Hop Details

Hop Index:0 | Hop Origin IP:100.100.100.5 | Hop Destination IP:20.20.10.2 | MRU:1500 | Labels: [24007/implicit-null] | ret code:0 | return char: | multipaths:0

Hop Index:1 | Hop Origin IP:20.20.10.2 | Hop Destination IP:20.20.10.14 | MRU:1500 | Labels: [24022/implicit-null] | ret code:8 | return char:L | multipaths:1

Hop Index:2 | Hop Origin IP:20.20.10.14 | Hop Destination IP:20.20.10.26 | MRU:1500 | Labels:[implicit-null/implicit-null] | ret code:8 | return char:L | multipaths:1

Hop Index:3 | Hop Origin IP:20.20.10.26 | MRU:0 | ret code:3 | return char:! | multipaths:0

Step 6 From the **Actions** column, click **View Details**.

If you have not provided the longitude and latitude information for your devices, the path is visualized in the logical view.

Step 7 From the available paths, click **Path 0** to expand and view the active path.

Example:

Figure 10: Path Details

The screenshot shows the expanded view of Path 0 in the Path Details panel. The information is as follows:

Path Details

Select from the fields below to find available Native SR IGP Paths * Required field

Select Service: Select Type | Select Instance

Headend * PE-A (100.100.100.5)

Endpoint * PE-B (100.100.100.6)

Get Paths

Available Paths

Path 0

Output: tunnel-te220
 Nexthop: 20.20.10.2
 Source: 100.100.100.5
 Destination: 127.0.0.0

Hop Details

Hop Index:0 | Hop Origin IP:100.100.100.5 | Hop Destination IP:20.20.10.2 | MRU:1500 | Labels: [24007/implicit-null] | ret code:0 | return char: | multipaths:0

Hop Index:1 | Hop Origin IP:20.20.10.2 | Hop Destination IP:20.20.10.14 | MRU:1500 | Labels: [24022/implicit-null] | ret code:8 | return char:L | multipaths:1

Hop Index:2 | Hop Origin IP:20.20.10.14 | Hop Destination IP:20.20.10.26 | MRU:1500 | Labels:[implicit-null/implicit-null] | ret code:8 | return char:L | multipaths:1

Hop Index:3 | Hop Origin IP:20.20.10.26 | 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**, click on a device and view device details (.
 - b. Confirm that GNMI capability and connectivity details are configured.

Connectivity Details

Protocol *	IP Address / Subnet Mask *	Port *	Timeout	Encoding Type
TELNET	172.29.105.236 / 24	23	30	
SNMP	172.29.105.236 / 24	161	30	
SSH	172.29.105.236 / 24	22	30	
GNMI	172.29.105.236 / 24	57400	30	JSON

+ Add Another

Capability*

YANG MDT
 TL1
 YANG CLI
 YANG EPNM
 SNMP
 GNMI



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 must define them on the Cisco Crosswork UI. Affinity names defined on devices are not collected by Cisco Crosswork. The affinity mapping is only used for visualization. For this reason, you should collect affinities on the device, then define affinity mapping in Cisco Crosswork UI with the same name and bits that are used on the device. Crosswork Optimization Engine 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".

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 > 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**.

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):

Example:

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

Step 4 Click **Save** to save the mapping.


Note You should remove the TE tunnel before removing the affinity to avoid orphan TE tunnels. If you have removed an affinity associated to 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 54](#).

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.
- Step 2** From the **SR Policies** table, click **+ Create**. If you are using Crosswork Network Controller, select either **PCE init** or **PCC init**.
- Step 3** Enter or select the required SR-MPLS policy values. Hover the mouse pointer over  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 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 **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:
- 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-TE policy may take some time, depending on the network size and performance, to appear in the **SR Policy** table. The **SR Policy** table is refreshed every 30 seconds.
 - b. View and confirm the new SR-MPLS policy details. From the **SR Policy** table, click  and select **View**.

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 34](#).
-

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. In the event of a link or interface failing, the network will find an alternate path that meets all the criteria specified in the policy and raise an alarm. The alarm is also raised in case no path is found, the packets are then dropped.



Tip If you plan to use affinities, 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 54](#) or [Configure Flexible Algorithm Affinities, on page 65](#).

Before you begin


If you have affinities configured on your devices, define an affinity mapping using the Cisco Crosswork UI prior to provisioning an SR-MPLS policy (see [Configure TE Link Affinities, on page 54](#)).

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.
- Step 2** From the **SR Policy** table, click + **Create**. If you are using Crosswork Optimization Engine within Crosswork Network Controller, select either **PCE Init** or **PCC Init**.
- 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 **SR Policy** table. The **SR Policy** table is refreshed every 30 seconds.

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

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 34](#).

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 SR Policy table, locate the SR-MPLS policy you are interested in and click .

Step 3 Choose **View** 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 5

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](#).


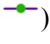
- [View RSVP-TE Tunnels on the Topology Map](#), on page 59
- [View RSVP-TE Tunnel Details](#), on page 61
- [Create Explicit RSVP-TE Tunnels](#), on page 62
- [Create Dynamic RSVP-TE Tunnels Based on Optimization Intent](#), on page 63
- [Modify RSVP-TE Tunnels](#), on page 64

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 11: Traffic Engineering UI - RSVP-TE Tunnels

Tunnel...	He...	En...	A...	O...	Actions	
<input type="checkbox"/>	33005	xrv9k...	xrv9k...	↑	↓	...
<input checked="" type="checkbox"/>	33004	xrv9k...	xrv9k...	↑	↑	...
<input type="checkbox"/>	33006	xrv9k...	xrv9k...	↑	↑	...
<input checked="" type="checkbox"/>	33007	xrv9k...	xrv9k...	↑	↑	...
<input type="checkbox"/>	33009	xrv9k...	xrv9k...	↑	↓	...
<input checked="" type="checkbox"/>	33010	xrv9k...	xrv9k...	↑	↓	...
<input checked="" type="checkbox"/>	33000	xrv9k...	xrv9k...	↑	↑	...
<input type="checkbox"/>	33000	xrv9k...	xrv9k...	↑	↑	...
<input type="checkbox"/>	33011	xrv9k...	xrv9k...	↑	↑	...
<input type="checkbox"/>	33001	xrv9k...	xrv9k...	↑	↑	...

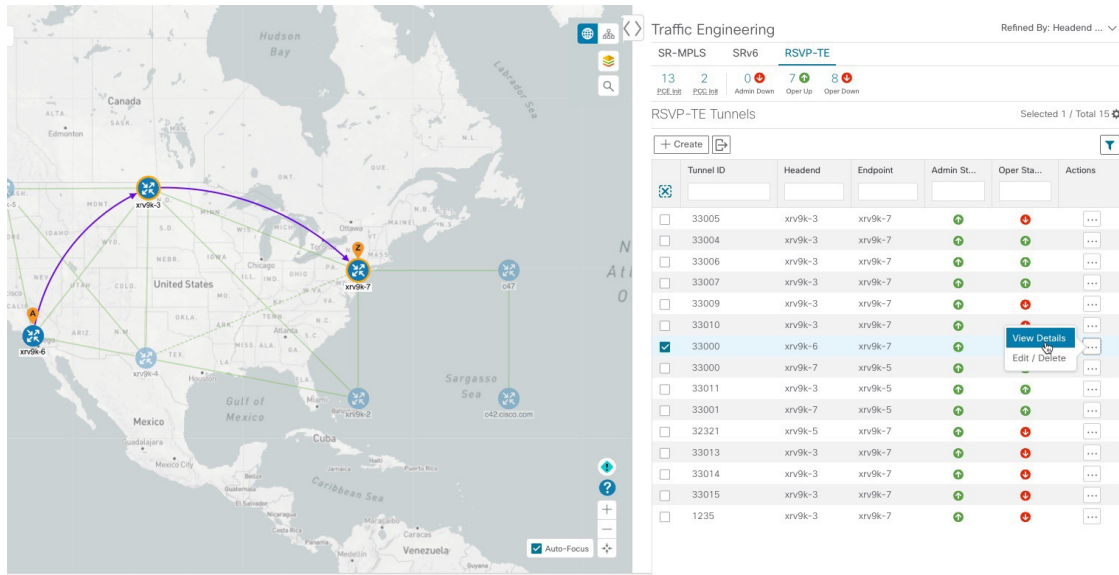
Callout No.	Description
1	Click Show Participating Only to display links that only belong to the selected RSVP-TE tunnels. All other links and devices disappear.
2	<p>A device with a solid orange outline () indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered.</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 purple 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>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 63 Create Explicit RSVP-TE Tunnels, on page 62 Modify RSVP-TE Tunnels, on page 64 View RSVP-TE Tunnel Details, on page 61
5	Exports <i>all</i> data into a CSV file. You cannot export selected or filtered data.
6	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
7	<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 endpoint device of the policy is in the selected group.

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.



Traffic Engineering

SR-MPLS SRv6 **RSVP-TE**

13 2 0 7 8
ESL: 13 ESL: 2 Admin Down Open Up Open Down

RSVP-TE Tunnels Selected 1 / Total 15

Tunnel ID	Headend	Endpoint	Admin St...	Oper Sta...	Actions
<input type="checkbox"/> 33005	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33004	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33006	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33007	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33009	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33010	xrv9k-3	xrv9k-7	●	●	...
<input checked="" type="checkbox"/> 33000	xrv9k-6	xrv9k-7	●	●	View Details Edit / Delete
<input type="checkbox"/> 33000	xrv9k-7	xrv9k-5	●	●	...
<input type="checkbox"/> 33011	xrv9k-3	xrv9k-5	●	●	...
<input type="checkbox"/> 33001	xrv9k-7	xrv9k-5	●	●	...
<input type="checkbox"/> 32321	xrv9k-5	xrv9k-7	●	●	...
<input type="checkbox"/> 33013	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33014	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 33015	xrv9k-3	xrv9k-7	●	●	...
<input type="checkbox"/> 1235	xrv9k-3	xrv9k-7	●	●	...

Step 2 View RSVP-TE tunnel details.

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).
- 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.

> RSVP-TE Tunnel Details
⋮ | ✕

Headend A xrv9k-6 (192.168.0.6)
Endpoint Z xrv9k-7 (192.168.0.7)
Tunnel ID 33000

▾ Summary

- Description** -
- Path Name** 60701-rsvp
- LSP ID** 6
- Path Type** Unknown
- Admin State** ↑ Up
- Oper State** ↑ Up
- Utilization** 0 Mbps
- Delay** 109 ↻
- Signaled Bandwidth** 0 Mbps
- Setup / Hold Priority** 7 / 7
- Metric Type** IGP
- Fast Re-route (FRR)** Disable
- Binding Label** 24012
- Accumulated Metric** 20
- Disjoint Group** ID:
Association Source: -
Type: -
- PCE Initiated** true
- Delegated PCE** 2001:420:28f:2011:250:56ff:fe85:a025
- Non-delegated PCEs** -
- Affinity** Exclude-Any: -
Include-Any: -
Include-All: -
- PCE Computed Time** 27-Oct-2021 12:33:03 PM PDT
- Last Update** 27-Oct-2021 12:39:58 PM PDT

Last Updated ✕
 27-Oct-2021 06:41:22 PM PDT


Explicit Route Object (ERO)

Hop	Node	IP	Interface Name	Type
0	xrv9k-3	10.0.0.29	GigabitEthernet0/0/0/4	Strict
1	xrv9k-7	10.0.0.42	GigabitEthernet0/0/0/1	Strict

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 on the path.

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering**.
- Step 2** From the right window, click **RSVP-TE**.
- Step 3** Under **RSVP-TE Tunnels**, click **+ Create**.
- Step 4** If you are using Crosswork Optimization Engine within Crosswork Network Controller, select either **PCE Initor** or **PCC Initor**.

- Step 5** 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** drop-down menu. Then navigate and zoom in on the topology map to click the device for headend or endpoint selection.
- Step 6** Under Policy Path, click **Explicit Path** and enter a path name.
- Step 7** Add segments that will be part of the RSVP-TE path.
- Step 8** Click **Preview**. The path is highlighted on the map.
- Step 9** If you want to commit the tunnel path, click **Provision**.
- Step 10** 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 **RSVP-TE Tunnels** table. The **RSVP-TE Tunnels** table is refreshed every 30 seconds.
 - b. View and confirm the new RSVP-TE tunnel details. From the **RSVP-TE** table, click *** (in the same row as the RSVP-TE tunnel), and select **View**.
- Note** On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. 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 path for the tunnel that is based on metrics and path constraints (affinity or disjointness) defined by you. 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 then map them in Cisco Crosswork before creating a dynamic RSVP-TE tunnel. For more information, see [Configure TE Link Affinities, on page 54](#).

- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering**.
- Step 2** From the right window, click **RSVP-TE**.
- Step 3** Under **RSVP-TE Tunnels**, click + **Create**.
- Step 4** 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** drop-down menu. Then navigate and zoom in on the topology map to click the device for headend or endpoint selection.

Step 5 Under **Tunnel Path**, click **Dynamic Path** and enter the Path Name.

Step 6 Under **Optimization Objective**, select the metric you want to minimize.

Step 7 Define any applicable constraints and disjointness.

Note Affinity constraints and disjointness cannot be configured on the same RSVP-TE tunnel. Also, there cannot be more than 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 that belong to that same disjoint group are shown during Preview.

Step 8 Click **Preview**. The path is highlighted on the map.

Step 9 If you want to commit the tunnel path, click **Provision**.

Step 10 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 **RSVP-TE Tunnels** table. The **RSVP-TE Tunnels** table is refreshed every 30 seconds.

b. View and confirm the new RSVP-TE tunnel details. From the **RSVP-TE** table, click and select **View**.

Note On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. 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**.

Step 2 From the **Traffic Engineering** window select the **RSVP-TE** tab.

Step 3 Locate the RSVP-TE tunnel you are interested in and click .

Step 4 Choose **View** 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 6

Flexible Algorithms

Flexible Algorithm allows operators to customize and compute the IGP shortest path according to their own 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 Flexible Algorithm and visualize the subset of the network that is capable of providing a specific set of transport characteristics. The ability to visualize Flexible Algorithm topologies provides an important tool to help you deploy, maintain, and verify that the configured Flexible Algorithm intent is realized in your network. For example, to improve service availability, you may use Flexible Algorithm to 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.



Note When using Crosswork Optimization Engine within the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering**.

This section contains the following topics:

- [Configure Flexible Algorithm Affinities, on page 65](#)
- [Visualize Flexible Algorithms, on page 66](#)
- [Find Flexible Algorithms for Links and Devices, on page 68](#)

Configure Flexible Algorithm Affinities

Flexible Algorithm affinity names that are defined on devices are not collected by Crosswork. The affinity mapping name is used for visualization and should be configured prior to visualizing Flexible Algorithms. For this reason, you should manually configure and collect Flexible Algorithm affinities on the device, then define the affinity mapping in the UI with the same name and bits that are used on the device. Crosswork only sends 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".

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 > System Settings > Traffic Engineering > Affinity > Flex-Algorithm 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. For example (using the above configuration):

Example:

Name	Bit Position (0-255)	Actions
b33	33	Edit Delete
red	1	Edit Delete
blue	5	Edit Delete

Step 4 Click **Save** to save the mapping. To view all Flexible Algorithm affinities for a link, see [Find Flexible Algorithms for Links and Devices](#), on page 68.

Visualize Flexible Algorithms

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 56.

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](#)).
- You should know the Flexible Algorithm IDs that are used in your network. To view Flexible Algorithm membership, see [Find Flexible Algorithms for Links and Devices](#), on page 68.



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 only those links and nodes that participate in both Flexible Algorithms.
- Step 7** Click **Apply**. You must click **Apply** for any additional changes to Flexible Algorithm selections to see the update on the topology map.

Example:



- Note**
- You cannot filter Flexible Algorithm IDs that are on multiple domains. Domain filtering is not supported based on Flexible Algorithms.
 - 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, then the default blue link or node color appears.

Step 8 (Option) Click **Save View** to save the topology view and Flexible Algorithm selections.

Find Flexible Algorithms for Links and Devices

If you want to know if a device or link is a member of a Flexible Algorithm, do the following:

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

Step 2 To view whether a device is part of a Flexible Algorithm:

- a) From the topology map, click on a device.
- b) In the **Device Details** window, click the **Flex-Algo** tab. If the device is part of a Flexible Algorithm then Algo ID and information appears. For example:

The screenshot shows the 'Device Details' window with the 'Flex-Algo' tab selected. The window title is 'Device Details' with a close button (X). Below the title bar are tabs for 'Alarms', 'SR-MPLS', 'SRv6', 'Tree-SID', 'RSVP-TE', and 'Flex-Algo'. Below the tabs, there is a dropdown menu showing 'IGP: Domain ID: 1001, ISIS System ID: 0000.0000.0005, Level: 2' and an 'Expand All' link. Two algorithm entries are visible, each with a dropdown arrow and a title: 'Algo 128' and 'Algo 129'. Each entry contains the following information:

- Participating** Yes
- Elected Definition** Metric Type: IGP
 - Exclude-Any Affinity:
 - Include-Any Affinity:
 - Include-All Affinity:
- Advertised** Yes
 - Priority: 228
 - Definition Equal to Local: No

Note If the device is not a member, then you will only see IGP domain and OSPF ID information.

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

- a) From the topology map, click a link.
- b) In the **Links** page, click one of the link types.
- c) By default, the **Summary** tab is displayed within the **Link Details** window. If the link is a member, then the **FA Topologies** row displays what Flexible Algorithm each source and destination device belong to. You can also view any affinities in the **FA Affinities** row.

Link Details



Summary	Alarms	SR-MPLS	SRv6	Tree-SID	RSVP-TE
<p>Name GigabitEthernet0/0/0/2-GigabitEthernet0/0/0/2</p> <p>State Up</p> <p>Link Type L3 ISIS IPV4</p> <p>ISIS Level 2</p> <p>Last Update 28-Jul-2022 03:41:47 PM PDT</p>					
	A Side	Z Side			
Node	xrv9k-PE6	xrv9k-5			
TE Router ID	192.168.0.6	192.168.0.5			
IPv6 Router ID	2001:192:168::6	2001:192:168::5			
IF Name	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/2			
IF Description	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/2			
Type	ETHERNETCSMACD	ETHERNETCSMACD			
IP Address	10.0.0.50	10.0.0.49			
Utilization	0% (0Bps/1Gbps)	0% (0Bps/1Gbps)			
IGP Metric	10	10			
Delay Metric	10	10			
TE Metric	10	10			
FA Affinities					
Admin Groups	2,4	2,4			
FA Topologies	128, 129, 130, 131, 132, 134	128, 129, 130, 131, 132			



CHAPTER 7

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, which represents all of the tree segments and devices in it. The SDN controller is highly flexible, as it understands the segmentation and can construct routing paths using any kind of 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 71](#)
- [View a Point-to-Multipoint Tree on the Topology Map, on page 72](#)
- [Create Static Tree-SID Policies, on page 73](#)
- [Modify a Tree-SID Policy, on page 76](#)

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 72](#)).

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 as well as a mid-point (transit) node toward the downstream sub-tree.
- Leaf node—Decapsulates the multicast traffic and forward 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, either directly using SR-PCE CLI or from the Crosswork UI. You can refer to the Tree-SID configuration documentation for your specific device to find out 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 a 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 Crosswork topology map. The root node (R) and leaf nodes (L) are clearly marked, and the arrows denote the path through the transit nodes from the root to the leaf nodes.

You can drill down on the nodes and the links to see more details about the Tree-SID policy and validate the configuration.

Tree-SID Policy Details																																									
Leaf Node Name		Leaf Node IP																																							
<table border="1"> <thead> <tr> <th>Role</th> <th>Name</th> <th>IP</th> <th>Local IP</th> <th>Remote IP</th> <th>Egress Link</th> </tr> </thead> <tbody> <tr> <td>Root</td> <td>xrv9k-26</td> <td>192.168.0.26</td> <td>10.0.0.30</td> <td>10.0.0.29</td> <td></td> </tr> <tr> <td>Transit</td> <td>xrv9k-23</td> <td>192.168.0.23</td> <td>10.0.0.41</td> <td>10.0.0.42</td> <td></td> </tr> <tr> <td>Leaf</td> <td>xrv9k-27</td> <td>192.168.0.27</td> <td>-</td> <td>-</td> <td></td> </tr> </tbody> </table>						Role	Name	IP	Local IP	Remote IP	Egress Link	Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29		Transit	xrv9k-23	192.168.0.23	10.0.0.41	10.0.0.42		Leaf	xrv9k-27	192.168.0.27	-	-													
Role	Name	IP	Local IP	Remote IP	Egress Link																																				
Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29																																					
Transit	xrv9k-23	192.168.0.23	10.0.0.41	10.0.0.42																																					
Leaf	xrv9k-27	192.168.0.27	-	-																																					
<table border="1"> <thead> <tr> <th colspan="2">Leaf Node Name</th> <th colspan="4">Leaf Node IP</th> </tr> </thead> <tbody> <tr> <td colspan="6">192.168.0.22</td> </tr> <tr> <th>Role</th> <th>Name</th> <th>IP</th> <th>Local IP</th> <th>Remote IP</th> <th>Egress Link</th> </tr> <tr> <td>Root</td> <td>xrv9k-26</td> <td>192.168.0.26</td> <td>10.0.0.81</td> <td>10.0.0.82</td> <td></td> </tr> <tr> <td>Bud</td> <td>xrv9k-24</td> <td>192.168.0.24</td> <td>10.0.0.14</td> <td>10.0.0.13</td> <td></td> </tr> <tr> <td>Leaf</td> <td>xrv9k-22</td> <td>192.168.0.22</td> <td>-</td> <td>-</td> <td></td> </tr> </tbody> </table>						Leaf Node Name		Leaf Node IP				192.168.0.22						Role	Name	IP	Local IP	Remote IP	Egress Link	Root	xrv9k-26	192.168.0.26	10.0.0.81	10.0.0.82		Bud	xrv9k-24	192.168.0.24	10.0.0.14	10.0.0.13		Leaf	xrv9k-22	192.168.0.22	-	-	
Leaf Node Name		Leaf Node IP																																							
192.168.0.22																																									
Role	Name	IP	Local IP	Remote IP	Egress Link																																				
Root	xrv9k-26	192.168.0.26	10.0.0.81	10.0.0.82																																					
Bud	xrv9k-24	192.168.0.24	10.0.0.14	10.0.0.13																																					
Leaf	xrv9k-22	192.168.0.22	-	-																																					
<table border="1"> <thead> <tr> <th colspan="2">Leaf Node Name</th> <th colspan="4">Leaf Node IP</th> </tr> </thead> <tbody> <tr> <td colspan="6">192.168.0.24</td> </tr> </tbody> </table>						Leaf Node Name		Leaf Node IP				192.168.0.24																													
Leaf Node Name		Leaf Node IP																																							
192.168.0.24																																									

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.

Note You can view a maximum of two policies on the topology map at the same time.

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 32](#)

Step 3 To view the Tree-SID Details, from the **Actions** column, click > **View Details** for one of the Tree-SID policies.

Step 4 You can view the Tree-SID details, and verify the path and node details to ensure that the Tree-SID is configured correctly.

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.

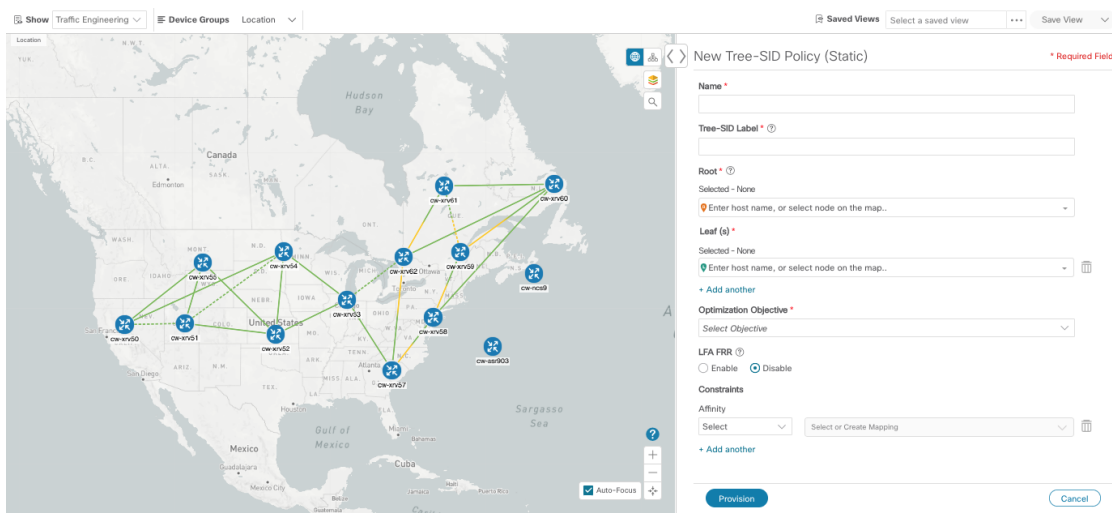


Tip If you plan to use affinities, collect affinity information from your devices and then map them in Cisco Crosswork before creating a static Tree-SID policy. For more information, see [Configure TE Link Affinities, on page 54](#).

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.



Step 3 Add leaf nodes that will be part of the Tree-SID policy.

Step 4 Under **Affinity**, define any applicable affinities. For more information on Affinities, see [Configure TE Link Affinities, on page 54](#)

Step 5 To commit the policy, click **Provision**.

Step 6 Validate the Tree-SID policy creation:

- a. Confirm that the new Tree-SID policy appears in the Tree-SID 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 Tree-SID table, depending on the network size and performance. The **Tree-SID** table is refreshed every 30 seconds.

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

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
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
  Outgoing: 505254 CC-ID: 3 (11.1.2.52)
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 (11.1.28.54)
  Outgoing: 505254 CC-ID: 1 (11.1.1.51)

```

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**.

Step 2 From the **Traffic Engineering** window, click the **Tree-SID** tab.

Step 3 Locate the Tree-SID policy you are interested in and click .

Step 4 Choose **Edit**.

- 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 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 to happen:

A node on a Tree-SID policy path may not be available in the Crosswork topology information for various reasons. For example, the node is not added to the Crosswork device inventory. This affects how the Tree-SID policy path is displayed on the topology map. The unavailable node will cause one or more root-to-leaf paths to look broken on the map, as shown in the figure below. However, the path details in the right panel will still show the full path.

The screenshot shows the Cisco Crosswork Optimization Engine 5.0 interface. The left panel displays a topology map with a Tree-SID path highlighted in purple. The path starts at a root node (xrv9k-VM3), goes through a missing node (xrv9k-VM7), and ends at a leaf node (xrv9k-VM8). The right panel shows the path details for the selected SR-PCE Address (172.23.209.75). The path is broken at the missing node, and the details table shows the full path including the missing node.

Leaf Node Name		Leaf Node IP		Egress Link	
xrv9k-VM11-771-151		192.168.4.14			
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-...	192.168.4.6	10.0.3.42	10.0.3.41	
xrv9k-VM7_3_0_732_cco		192.168.4.7			
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	-	-	
Splitfire		192.168.4.11			
Role	Name	IP	Local IP	Remote IP	
Leaf	Splitfire	192.168.4.11	-	-	



PART I

Bandwidth Feature Packs

- [SR Circuit Style Manager \(CSM\)](#), on page 81
- [Local Congestion Mitigation \(LCM\)](#), on page 103
- [Bandwidth on Demand \(BWoD\)](#), on page 125



CHAPTER 8

SR Circuit Style Manager (CSM)

The SR Circuit Style Manager (CSM) feature pack provides a bandwidth-aware Path Computation Element (PCE) to compute Circuit Style SR-TE policy paths that you can visualize in your network. Circuit Style enables segment routed transport tailored for circuit-oriented services over a packet based network through the use of bi-directional, co-routed, path protected SR-TE policies. Circuit Style SR-TE policies are typically used for high priority services, such as crucial monetary transactions or important live video feed, which *require committed bandwidth with fast and fail-safe connections*. The CSM feature pack ensures dynamic Circuit Style SR-TE policies are provisioned along paths that meet strict bandwidth requirements while at the same time respecting any additional user configured constraints such as latency minimization and disjointness.

Centralized bandwidth accounting in the CSM feature pack allows the user to monitor resource reservation levels and quickly identify hot spots where available bandwidth in the circuit style bandwidth pool is low. The ability to visualize Circuit Style SR-TE policies in your network topology enables easy verification of Circuit Style SR-TE policy configurations, details, and path states. With a few clicks you can view Active and Protect paths, operational status, reserved bandwidth pool size and monitor path failover behavior for individual Circuit Style SR-TE policies.



Note

- Functionality described within this section is only available with certain licensing options.
- When using the Crosswork Network Controller solution, the UI navigation starts from **Traffic Engineering & Services** instead of **Traffic Engineering**.

This section contains the following topics:

- [Circuit Style SR-TE Important Notes](#), on page 82
- [Workflow for Setting Up CS SR-TE Policy Visualization](#), on page 86
- [Enable SR Circuit Style Manager](#), on page 87
- [Configure Circuit Style SR Policies](#), on page 88
- [Review Circuit Style SR-TE Policy Bandwidth Utilization](#), on page 90
- [View Circuit Style SR-TE Policies](#), on page 91
- [Trigger CSM to Recalculate a Circuit Style SR-TE Policy](#), on page 96
- [What Happens When Bandwidth Reservation Settings are Exceeded?](#), on page 96
- [How Does CSM Handle Path Failures?](#), on page 99

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.

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 of the attributes described in the table below act as constraints. Each of them 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 they either specify a required property of a path or exclude possible choices for that path.

Table 4: 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 of the current paths cannot accommodate the new bandwidth, it returns the old bandwidth value indicating that it was unsuccessful. This evaluation will not be retried unless 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>The user interface shows both the requested and reserved bandwidth under each candidate path when you view the policy details. These values can differ if the requested bandwidth is increased but there is insufficient available CS pool bandwidth along one or more of the paths.</p>

Attribute	Description
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>
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 and Latency metric types are supported. The metric type used must match across 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>

Attribute	Description
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.
Supported Policy Changes	<p>The following constraints may be changed 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 made in a consistent manner 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>

Attribute	Description
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.</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 after Restore paths are delegated from both sides to wait before the path is computed. 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) Not supported path computation Inter-AS</p>
Reversion Behavior	<p>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):</p> <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 5: 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 6: Tasks to Complete to Start Visualizing Circuit Style SR-TE Policies

Step	Action
1. Enable the SR Circuit Style Manager (CSM) feature pack.	<p>From the main menu, choose Traffic Engineering > Circuit Style SR-TE > Configuration.</p> <p>Follow the steps in Enable SR Circuit Style Manager, on page 87.</p>
<p>2. Configure CS SR policies on the devices.</p> <p>Note If you do this step before enabling the Circuit Style SR-TE feature pack, then the CS SR policies will appear operationally down.</p>	<p>You can configure CS SR policies using one of the following methods:</p> <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 88. • 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 (version 5.0 or later).

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 90). 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

In order 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 CSM is enabled, it 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 96.

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 88](#)) 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 82](#) 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-protected-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 82](#) for additional requirements or notable constraints.



Example:

```
segment-routing
traffic-eng
  policy cs1-cs4

  performance-measurement
    liveness-detection
      liveness-profile backup name CS-protected      !! 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
  !
```


In this example, the reserved bandwidth pool size is displayed 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 bandwidth to be used for Circuit Style SR-TE policies for these interfaces.

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 Topologies		
<input type="checkbox"/> Circuit Style Bandwidth Pool		
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**.

Traffic Engineering

SR-MPLS SRv6 Tree-SID

29	24	0	0	0
Total	Circuit Style	BWoD	LCM	Admin Dow.

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.

View Circuit Style SR-TE Policies

The screenshot shows the Traffic Engineering interface. On the left is a map of the United States with several Network Control Sites (NCS) marked: NCS-1 in San Diego, NCS-3 in Miami, and NCS-4 in Denver. A path is shown connecting NCS-1 to NCS-3. On the right, the 'SR Policy' table is displayed with the following data:

	Head...	End...	C...	Admi...	Oper ...	Actions
<input type="checkbox"/>	NCS-3	NCS1	124		↑ ↓	...
<input type="checkbox"/>	NCS1	NCS-3	124		↑ ↓	...
<input type="checkbox"/>	NCS-3	NCS1	134		↑ ↓	...
<input type="checkbox"/>	NCS1	NCS-3	134		↑ ↓	...
<input checked="" type="checkbox"/>	NCS-3	NCS1	173		↑ ↓	...
<input checked="" type="checkbox"/>	NCS1	NCS-3	173		↑ ↓	View Details
<input type="checkbox"/>	NCS-3	NCS1	183		↑ ↓	Edit / Delete
<input type="checkbox"/>	NCS1	NCS-3	183		↑ ↓	...
<input type="checkbox"/>	NCS-3	NCS1	194		↑ ↓	...

The **Circuit Style Policy Details** window is displayed in the side panel. By default, the Active path is displayed and shows the bidirectional paths (Bi-Dir Path checkbox is checked) on the topology map.

The screenshot shows the 'Circuit Style Policy Details' window. The 'Bi-Dir Path' checkbox is checked. The 'Summary' section displays the following details:

- Headend: NCS-3 | TE RID: 100.100.100.5 PCC IP: 100.100.100.5
- Endpoint: NCS1 | TE RID: 100.100.100.4
- Color: 173
- Admin State: Up
- Oper State: Up
- Binding SID: 24033
- Policy Type: Circuit-Style
- Profile ID: -
- Description: -
- Traffic Rate: 0 Mbps
- Unused: True

The 'Candidate Path' section shows the following paths:

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_disc...	100		↑ ↓
<input type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_disc...	50		↑

Here is a closer look at the type of Summary details available. The Candidate Path list displays the Active (path that currently takes traffic) and Protected paths.

Circuit Style Policy Details ⋮ ×

Current History

Headend NCS-3 | TE RID: 100.100.100.5 PCC IP: 100.100.100.5

----- ↕

Endpoint NCS1 | TE RID: 100.100.100.4

Color 173

▼ Summary

Admin State ↑ Up

Oper State ↑ Up

Binding SID 24033

Policy Type Circuit-Style

Profile ID -

Description -

Traffic Rate 0 Mbps

Unused True ?

Delay 10 ?

Bandwidth Constraint 0 Mbps

Accumulated Metric 10

Protection Status PROTECTED

Delegated PCE 172.20.100.240

Non-delegated PCEs -

PCE Computed Time -


Last Update 16-Feb-2023 09:18:08 AM PST


[See less](#) ^

▼ Candidate Path Expand All

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_dis...	100		↑ ⚠
<input type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_dis...	50		↑

Note The Bandwidth Constraint value can differ from the bandwidth you requested if the value was increased and insufficient resources existed to satisfy demand on all Active and Protected candidate paths.

Step 3 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.



Step 4 View Active and Protected path configuration details.

- a) Within the **CS-SR Policy Details** window, you can drill down to view more information about the Active and Protected paths. In the following example, the Active path has a preference value of 100 and the Protected path has preference value of 50. The operational (Oper State Up) candidate path with the highest preference will always be the Active path (see [How Does CSM Handle Path Failures?, on page 99](#)). Click **Expand All** to view more information about both the Active and Protected paths.

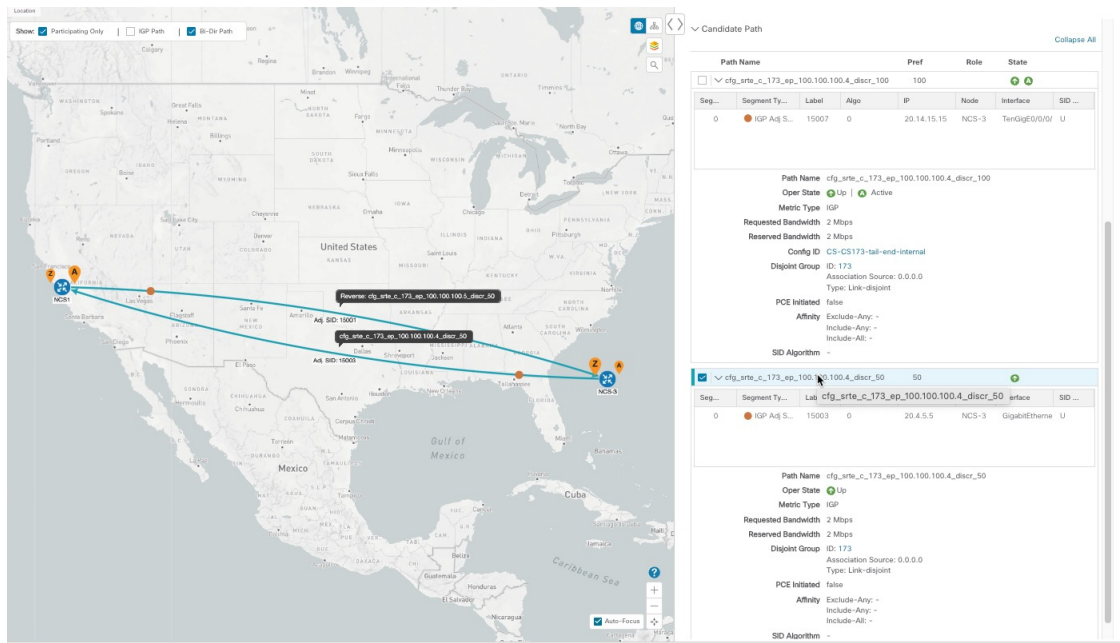
Candidate Path

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_...	100	↑ A	
<input type="checkbox"/> > cfg_srte_c_173_ep_100.100.100.4_...	50	↑	

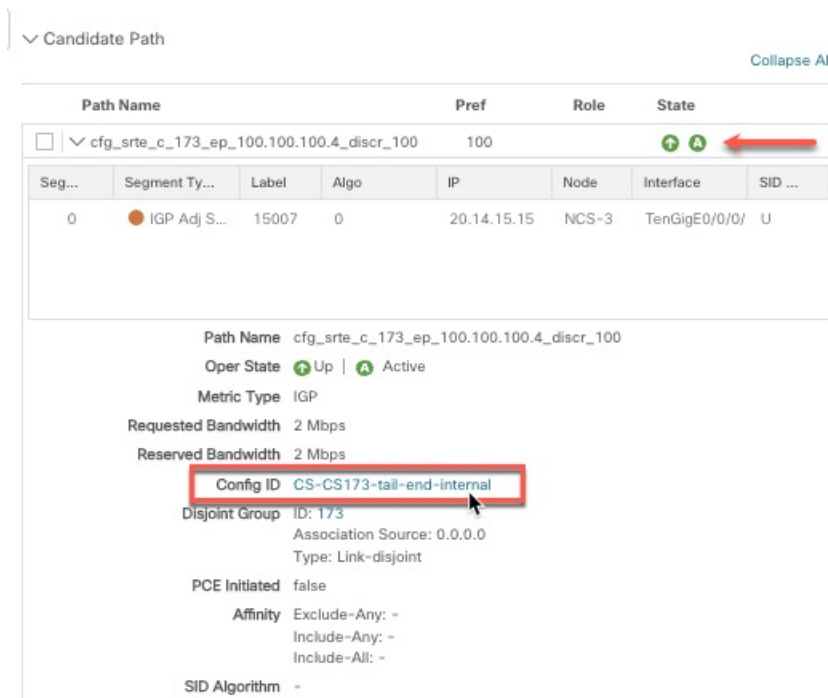
Expand All

- 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.

- b) In the following example, the Protected path is checked and displayed on the topology map. If you hover your mouse over the path name, forward and reverse paths are displayed on the topology map.





- c) Here is a closer view of an Active path's configuration details. Notice that it is designated with the "A" icon under State to clearly indicate it is currently the operational Active path. Also, 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**.



Here is a sample of a Circuit Style policy configuration. For more information, see [Configure Circuit Style SR Policies](#), on page 88.

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.

-
- 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 either the bandwidth pool size or bandwidth alarm threshold that were set in 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 for 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.

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 Affi...		
FA Top...		
∨ Circuit...		
Pool ...	1000 Mbps	1000 Mbps
Used	10 Mbps	10 Mbps
Avail...	990 Mbps	990 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..

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 Affi...		
FA Top...		
∨ Circuit...		
Pool ...	1000 Mbps	1000 Mbps
Used	500 Mbps	500 Mbps
Avail...	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.

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.

1. 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 2 nodes are the paths computed for the first CS policy.
 - The new Circuit Style SR-TE policy *r02 - r01 - r2* gets no path computed from CSM and therefore remains operationally down.

Endpoint 5501-01 | TE RID: 10.255.255.1

Color 2000

Summary

Admin State ↑ Up

Oper State ↓ Down

Binding SID 0

Policy Type Circuit-Style

Profile ID -

Description -

Traffic Rate 0 Mbps

Unused True i

[See more](#) v

Candidate Path


Path Name	Pref	RoleState
<input checked="" type="checkbox"/> > cfg_r1-r2-2_discr_100	100	↓ A
<input type="checkbox"/> > cfg_r1-r2-2_discr_50	50	↓

- Alerts are triggered.

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.

 Endpoint 5501-01 | TE RID: 10.255.255.1

Color 2000

Summary


Admin State  Up
 Oper State  Up
 Binding SID 24005
 Policy Type Circuit-Style
 Profile ID -
 Description -
 Traffic Rate 0 Mbps
 Unused True 

[See more](#) 

Candidate Path

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

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

Source	Severity	Description
SR Policy [10.255.255.1#10.255.255.1]	 Clear	Policy 'srte_c_2000_ep_10.255.255.2' has operational status back to UP.
SR Policy [10.255.255.2#10.255.255.1]	 Clear	Policy 'srte_c_2000_ep_10.255.255.1' 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. There are three types of candidate paths used during path failures:

- **Working**—This is the path with the highest preference candidate path.
- **Protected**—This path is defined as the second highest preference candidate path. If the Working path goes down, the Protected path (with the lower preference value) is activated. After the Working path recovers, the Protected path remains active until the default lock duration expires.
- **Restore**—This path is defined as the lowest preference candidate path. 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 87](#)). This delay allows topology and policy state changes to fully propagate to Crosswork, in cases where these changes triggered the Restore path delegation.

To address path failures effectively and switchover from working path to protect path, you can configure Performance Measurement (PM). For more information, see [Configure Circuit Style SR Policies, on page 88](#).

Examples



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

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

The screenshot displays a network topology with routers labeled frankenrouter-01, frankenrouter-02, 5501-01, 5501-02, and 540-01. The candidate path table is as follows:

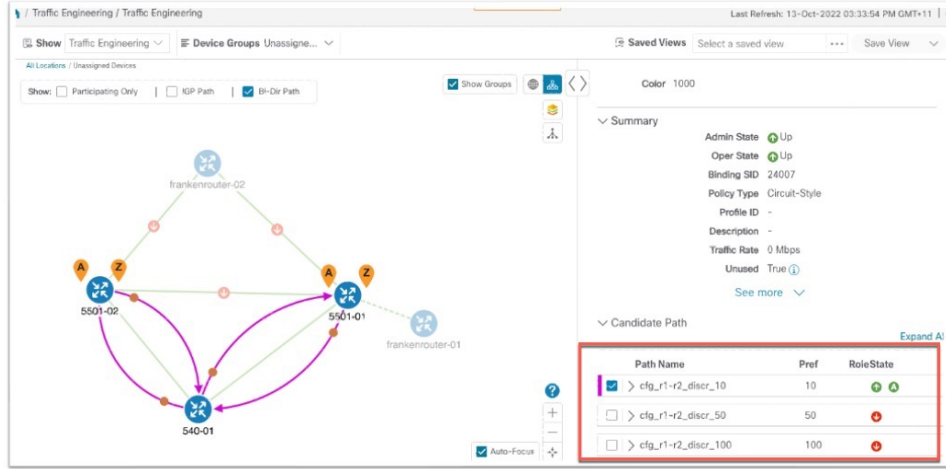
Path Name	Pref	Role State
> cfg_r1-r2_discr_100	100	Up (A)
> cfg_r1-r2_discr_50	50	Up (P)

When the Active path goes down, the Protected path immediately becomes "active". When the Active path goes back up, then the Protected path takes the role of "protected" again and the Active path (with preference 100) becomes active.

The screenshot displays the same network topology as the previous image. The candidate path table is as follows:

Path Name	Pref	Role State
> cfg_r1-r2_discr_50	50	Up (A)
> cfg_r1-r2_discr_100	100	Down (P)

In the case where both Working and Protected paths go down, CSM calculates a Restore path and it becomes the active path. Note that the Restore path has the lowest preference value of 10. The Restore path only appears in this particular case. If either the Working or Protected paths become operational again, the Restore path disappears from the topology map and from the Candidate Path list.





CHAPTER 9

Local Congestion Mitigation (LCM)

- [Local Congestion Mitigation Overview, on page 103](#)
- [LCM Important Notes, on page 104](#)
- [LCM Calculation Workflow, on page 106](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 108](#)
- [Configure LCM, on page 117](#)
- [Add Individual Interface Thresholds, on page 119](#)
- [Monitor LCM Operations, on page 121](#)

Local Congestion Mitigation Overview

Local Congestion Mitigation (LCM) searches for congestion on a configurable cadence (as opposed to a triggered event) and 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 the minimal amount of 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. If the user approves, LCM performs the mitigation through the deployment of Tactical Traffic Engineering (TTE) SR policies. LCM will not modify paths of existing deployments of SR policies to mitigate congestion. With LCM, you are able to 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 deploy changes in the network automatically to address congestion and 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 117](#).

LCM allows for a wider applicability of the solution in various network topologies such as that involving multiple IGP areas due to its simpler path computation and limitation to specific network elements. Focusing on the problem locally within a domain eliminates the need for simulating edge-to-edge traffic flows in the network through a full traffic matrix and allows for better scaling of large networks. Also, LCM performs the collection of TTE SR policy and interface counters via SNMP and does not require the use of SR-TM.



Note Take a look at the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 108](#) to see how to use LCM in your network.

LCM Important Notes

Consider the following information when using LCM:

- You must have the Advanced RTM license package to use LCM.
- 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 assigned to an IGP process. Domains are learned from the network. The domain ID is taken from 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 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 over-provisioning 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.
- LCM will not operate in portions of the network carrying Tree-SID LSPs.
- When domain interfaces and links are removed (intentionally or unintentionally), the following occurs:
 - As links go down (LINK_DOWN state), LCM configuration and the Domain UI card (see [Configure LCM, on page 117](#)) 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.

LCM Platform Requirements

The following is a non-exhaustive list of high-level requirements for proper LCM operation:

Congestion Evaluation:

- LCM requires traffic statistics from the following:
 - SNMP interface traffic measurements
 - SNMP headend SR-TE policy traffic measurements
- Strict SID labels should be configured for SR.

Congestion Mitigation:

- The headend device should support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies
- The headend device must support PCE-initiated SR-TE policies with autoroute steering

Devices should be configured with `force-sr-include` to enable traffic steering into SR-TE policies with autoroute. For example:

```
segment-routing traffic-eng pcc profile <id> autoroute include ipv4 all
segment-routing traffic-eng pcc profile <id> autoroute force-sr-include
```

where `<id>` is the user configured ID (any number as allowed per 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](#))

Contact your Cisco sales representative for an exhaustive list of platform requirements.

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.

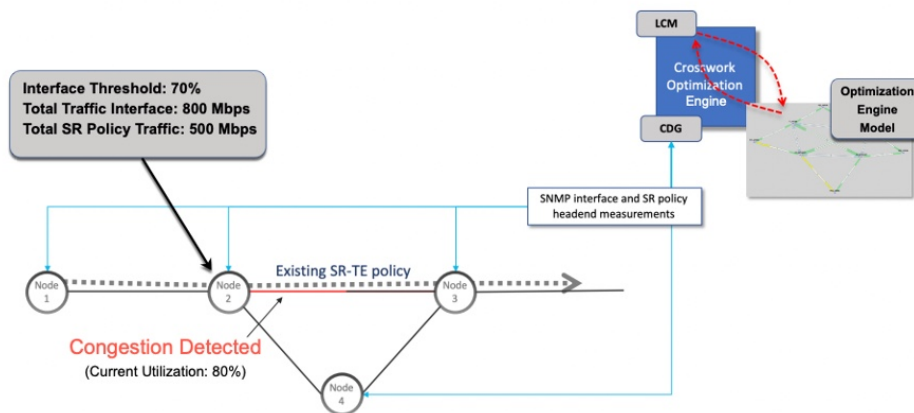
Figure 12: BGP-LS Session Reporting Domain 100



LCM Calculation Workflow

This example walks you from congestion detection to the calculations LCM performs prior to recommending tactical tunnel deployment. With the release of Crosswork Optimization Engine 3.0, these calculations are done on a per domain basis which allows better scalability and faster calculation for larger networks.

Figure 13: 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, after a congestion check interval, LCM detects congestion 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 (for example: unlabeled, IGP routed, or carried via FlexAlgo-0 SIDs). The traffic within an SR-TE policy will not be included in 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 = 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. 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

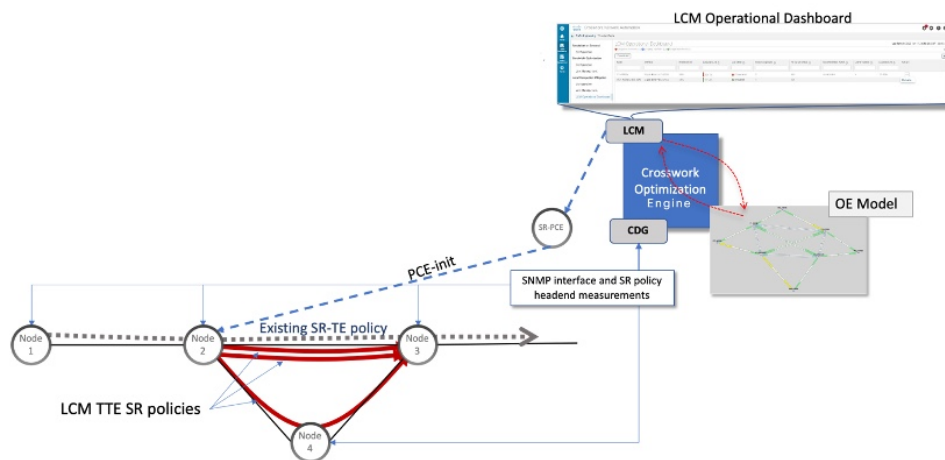
LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see [Configure LCM, on page 117](#).

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 that are needed on the shortest versus alternate paths, respectively.

In this example, LCM needs to divert one-third of the total eligible traffic (100Mbps out of 300Mbps) away from the congested link. Assuming a perfect ECMP, LCM estimates that 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**.



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**. TTE SR policy removal recommendations will occur if the mitigated interface would not be congested if these policies were 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.

This example demonstrates the following workflow:

1. View uncongested topology.
2. Set utilization thresholds for individual interfaces.
3. Enable and configure LCM.
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.

The following image shows the topology that will be used for this example.

Figure 14: Initial Topology



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

- a) Click on the link between F3.cisco.com and F5.cisco.com to view link details. Note that utilization on F3.cisco.com is 9%.

Figure 15: Initial Utilization

The screenshot shows the 'Link Details' window for the connection between F3.cisco.com and F5.cisco.com. The window is divided into a map on the left and a details panel on the right. The details panel includes a summary and a table of link attributes.

	A Side	Z Side
Name	GigabitEthernet0/0/0/1-GigabitEthernet0/0/0/0	
State	Up	
Link Type	L3 OSPF V2	
Last Update	15-Apr-2022 10:20:22 PM PDT	
Node	F3.cisco.com	F5.cisco.com
TE Router ID	192.168.100.3	192.168.100.5
IF Name	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/0
IF Description	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/0
Type	ETHERNETCSMACD	ETHERNETCSMACD
IP Address	100.100.1.17	100.100.1.18
Utilization	9% (159Mbps/1Gbps)	0% (0Bps/1Gbps)
IGP Metric	1	
Delay Metric	1	
TE Metric	1	
OSPF Router ID	192.168.100.3	
OSPF Area	0	
FA Affinities		
Admin Groups		

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 polices to remediate the congestion. You set the global utilization threshold in the **LCM Configuration** page. However, if you want to define different thresholds for individual interfaces, we recommend that you define them in the **Customized Interface Threshold** page *prior* to enabling LCM.

- a) In this example, we will define some individual interface thresholds. Go to the **Customized Interface Thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Interface Thresholds**). Add interfaces 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 119](#).

See the following example and note the defined thresholds for F3.cisco.com with interface GigabitEthernet0/0/0/1 (13%) and F5.cisco.com with interface GigabitEthernet0/0/0/1 (11%).

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

Figure 16: Customized Interface Thresholds

Customized Interface Thresholds

Interfaces to Monitor: All Interfaces - LCM monitors the interfaces with custom thresholds. All other interfaces are monitored using the Utilization Threshold defined in the Configuration page.

Total 5

| Edit Mode: OFF

Node	Interface	Threshold (%)	Select to Delete
<input type="text"/>	<input type="text"/>	<input type="text"/>	
F4.cisco.com	GigabitEthernet0/0/0/1	14	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/1	13	<input type="checkbox"/>
F5.cisco.com	GigabitEthernet0/0/0/1	11	<input type="checkbox"/>
F1.cisco.com	GigabitEthernet0/0/0/1	20	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/2	10	<input type="checkbox"/>

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 in the **LCM Configuration** page (see **Step 3**).

- 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-ID** 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. For more information on all the available options, see [Configure LCM, on page 117](#).

Figure 17: LCM Configuration Page

Configuration

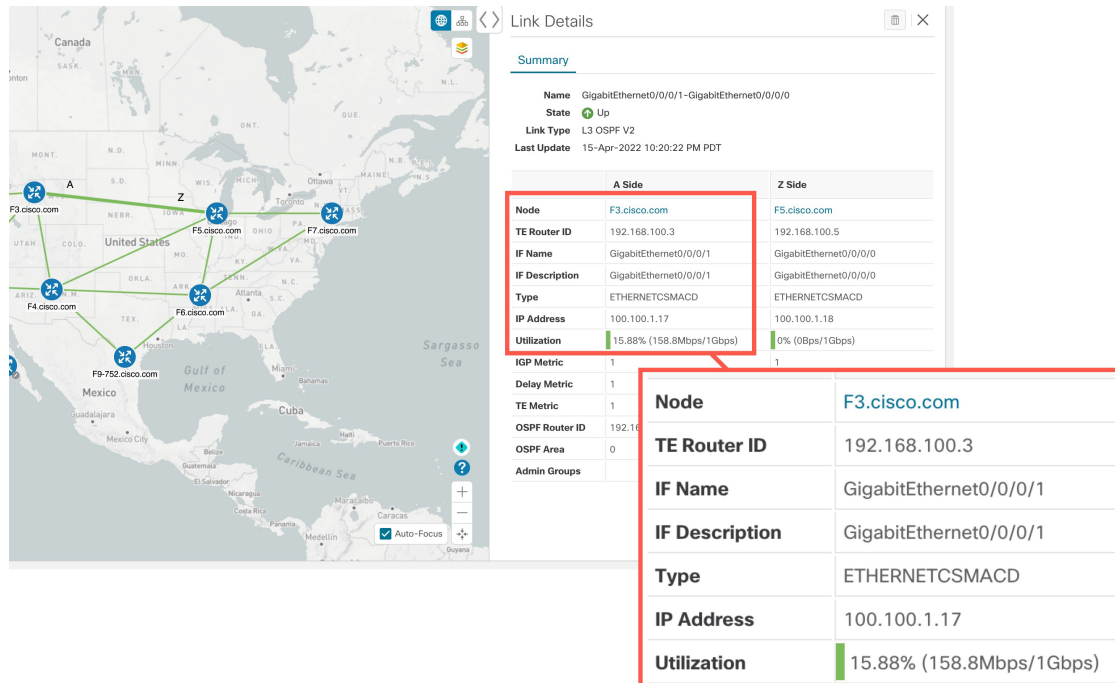
Basic Advanced

Enable ? False <input checked="" type="checkbox"/> True	Color ? 2000 <small>Range: 1 to 4294967295</small>	Utilization Threshold ? 80 <small>Range: 0 to 100</small>
Utilization Hold Margin ? 5 <small>Range: 0 to Utilization Threshold</small>	Delete Tactical SR Policies when Disabled ? False <input checked="" type="checkbox"/> True	Profile ID ? 1981 <small>Range: 0 to 65535</small>
Congestion Check Interval ? 300 seconds <small>Range: 60 to 86400 seconds</small>	Max LCM Policies per Set ? 8 <small>Range: 1 to 8</small>	Interfaces to Monitor ? <input type="radio"/> Selected Interfaces <input checked="" type="radio"/> All Interfaces
Description ? LCM Startup Config		

- 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. LCM will *not* commit or deploy new TTE policies automatically. 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 13% for node F3.cisco.com with interface GigabitEthernet0/0/0/1.

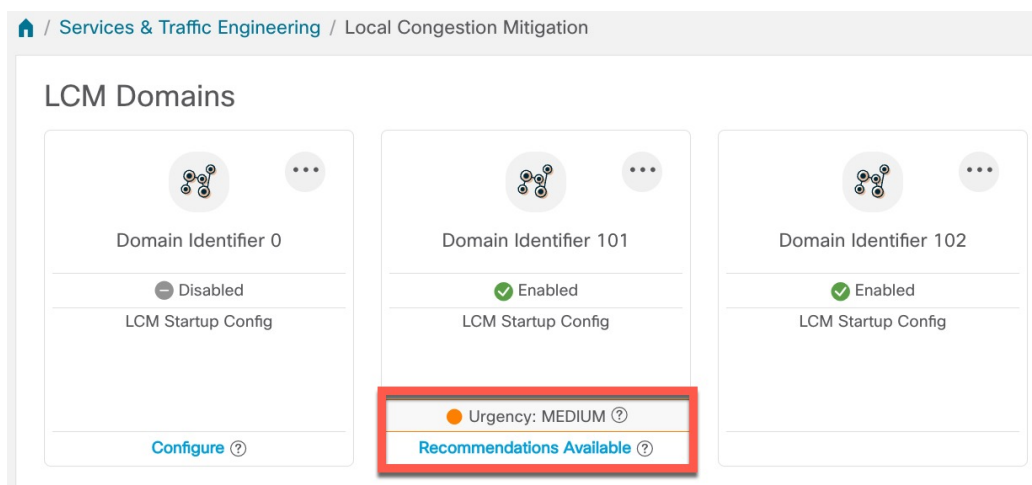
Figure 18: Observed Congestion




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 19: 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-ID > ... > Operational Dashboard)**.

The dashboard shows that F3.cisco.com utilization has surpassed 13% and is now at 16.05%. It also shows that F5.cisco.com utilization has also surpassed the 11% threshold and is now 19.26%. 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. The **Expected Utilization** column shows the expected utilization of each of the interface after the recommended action is committed. For more information, see [Monitor LCM Operations, on page 121](#).

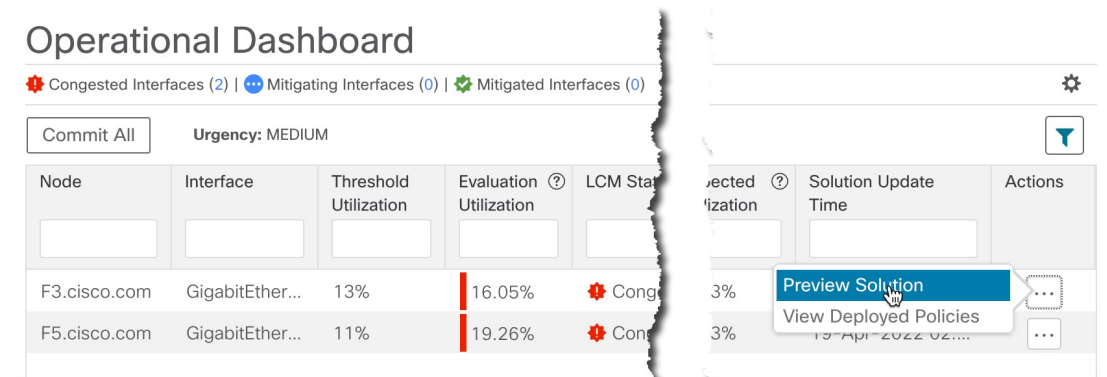
Figure 20: LCM Operational Dashboard

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F3.cisco.com	GigabitEther...	13%	16.05%	Congested	0	-	Create Set	None	8.03%	19-Apr-2022 02:...	...
F5.cisco.com	GigabitEther...	11%	19.26%	Congested	0	-	Create Set	None	9.63%	19-Apr-2022 02:...	...

Note If LCM cannot find a solution (**Recommended Action - No Solution**), it may be due to constraints enabled in this page.

- 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 21: Select 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 do in 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 F3.cisco.com and interface GigabitEthernet0/0/0/1. The top path shows the node SID (orange outline), headend and endpoint (A and Z) because the mouse pointer hovers over that segment.

Figure 22: LCM TTE Deployment Preview

Preview Recommended TTE Policies

Node F3.cisco.com
Interface GigabitEthernet0/0/0/1

Headend	Endpoint	Color	Recommended Action				
<input checked="" type="checkbox"/> F3.cisco.com	F5.cisco.com	2000	CREATE				
Se...	Segme...	L...	Algo	IP	N...	Interf...	St...
0	Nod...	16...	1	192.16...	F5...		Strict

Headend	Endpoint	Color	Recommended Action				
<input checked="" type="checkbox"/> F3.cisco.com	F5.cisco.com	2001	CREATE				
Se...	Segme...	L...	Algo	IP	N...	Interf...	St...
0	Nod...	16...	1	192.16...	F9...		Strict
1	IGP...	10...	0	100.10...	F9...	GigabitEthe	U
2	Nod...	16...	1	192.16...	F5...		Strict

[Back To LCM Dashboard](#)

- 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**.

Note All LCM recommendations per domain must be committed in order 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.

Figure 23: Mitigating LCM State

Operational Dashboard

🔴 Congested Interfaces (0) | 🟡 Mitigating Interfaces (2) | 🟢 Mitigated Interfaces (0)

[Commit All](#) | Urgency: LOW

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F5.cisco.com	GigabitEther...	11%	19.78%	🟡 Mitigating	2	OK	No Change	CONFIRMED	9.89%	19-Apr-2022 03:...	⋮
F3.cisco.com	GigabitEther...	13%	15.88%	🟡 Mitigating	2	OK	No Change	CONFIRMED	7.94%	19-Apr-2022 03:...	⋮

Step 6 Validate TTE SR policy deployments.

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

Note Crosswork Optimization Engine will report network events that are 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.


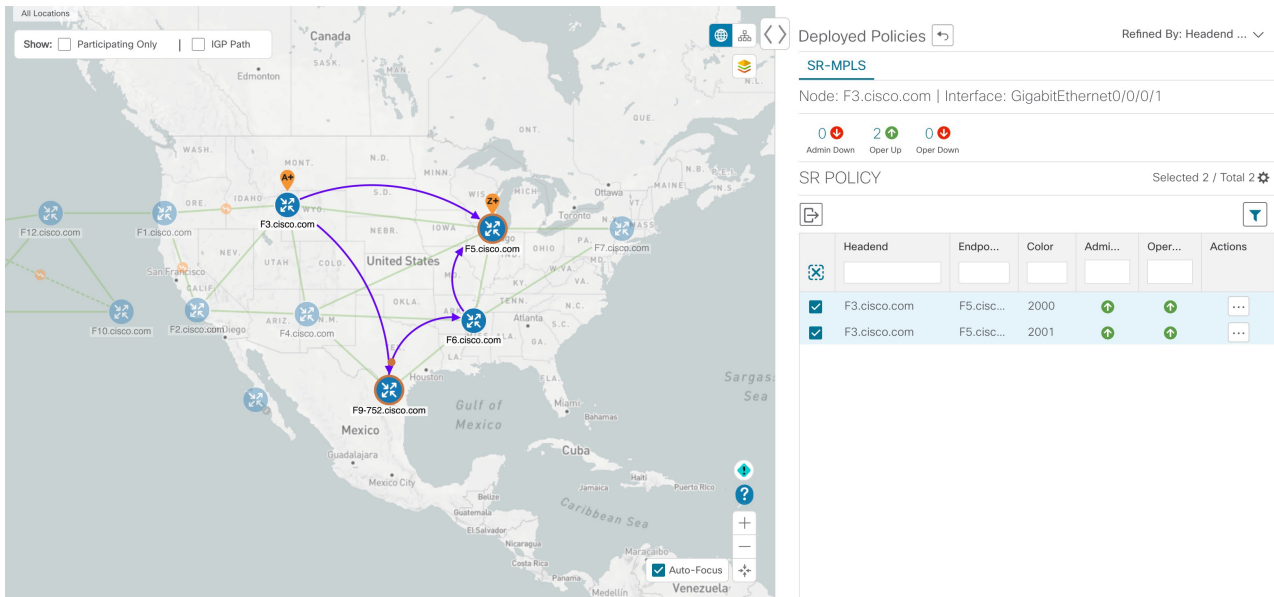


Click  in the **Actions** column and choose **View Deployed Policies**. The deployed policies are displayed in focus within the topology map. All other policies are dimmed.




Figure 24: View TTE Deployment Policies on Topology Map




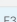
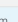

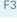
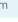
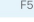
Deployed Policies  Refined By: Headend ... 

SR-MPLS

Node: F3.cisco.com | Interface: GigabitEthernet0/0/0/1

0  Admin Down 2  Oper Up 0  Oper Down

SR POLICY Selected 2 / Total 2 

	Headend	Endpo...	Color	Admi...	Oper...	Actions
<input checked="" type="checkbox"/>	F3.cisco.com	F5.cisc...	2000			
<input checked="" type="checkbox"/>	F3.cisco.com	F5.cisc...	2001			

- d) View SR policy details.

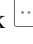
From the **Actions** column of one of the deployed policies click  and choose **View Details**. Note that the **Policy Type** is **Local Congestion Mitigation**.

Figure 25: SR Policy Details

SR Policy Details
⋮ | ✕

Details

Historical Data

Headend A F3.cisco.com | Source IP: 192.168.100.3
 TE RID: 192.168.100.3
 PCC IP: 192.168.100.3

Endpoint Z F5.cisco.com | Dest IP: 192.168.100.5
 TE RID: 192.168.100.5

Color 2000

∨ Summary

- Admin State ↑ Up
- Oper State ↑ Up
- Binding SID 1005011
- Policy Type Local Congestion Mitigation
- Profile ID 1981
- Description -
- Traffic Rate 39.28 Mbps
- Unused False
- Delay 1 ⓘ
- BWOD Policy Bandwidth 0 Mbps
- Accumulated Metric 0
- Delegated PCE 10.194.60.51
- Non-delegated PCEs -
- PCE Computed Time -
- Last Update 22-Apr-2022 01:31:10 PM PDT

[See less](#) ^

∨ Candidate Path [Collapse All](#)

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> ∨ lcm_to_F5_cisco_com_c_2000	100	Explicit	↑ A

Segment	Segment T...	Label	Algo	IP	Node	Interface	SID
0	⊙ Node SID	16505	1	192.168.1...	F5.cisco.com		Stric

Path Name lcm_to_F5_cisco_com_c_2000

Oper State ↑ Up | A Active

Metric Type UNKNOWN

Disjoint Group ID:
 Association Source: -
 Type: -

PCE Initiated true

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

Segment Type Unprotected

SID Algorithm -

Step 7 Remove the TTE SR policies upon LCM recommendation.

- After some time, the deployed TTE SR policies may no longer be needed. This occurs if the utilization will continue 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.

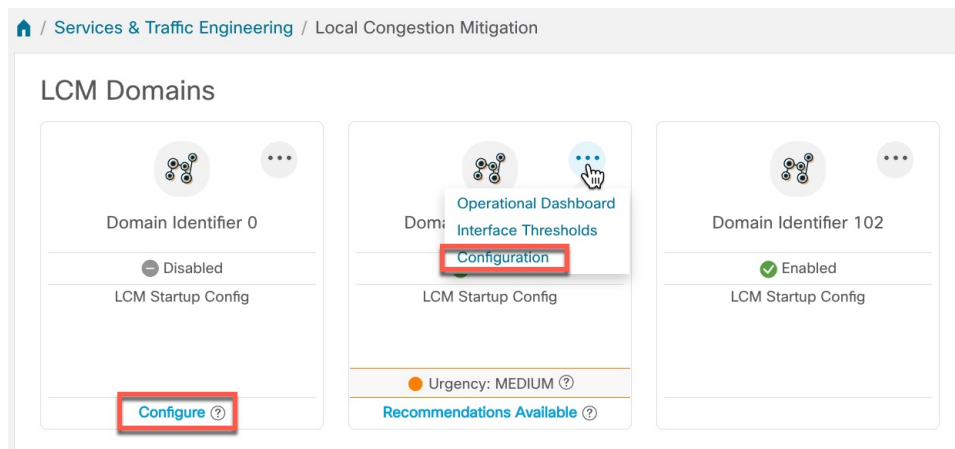
In this scenario we observed how to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands but at the same time gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations 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 one of the following:

- **Configuration**
- **Configure**



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 in this page.

The following list describes additional 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 in 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 to recommendations. 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 in the **Customized Interface Thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Customized Interface Thresholds**). Only interfaces defined in the **Customized Interface Thresholds** page will be monitored. If set to **All Interfaces**, LCM will monitor the interfaces with custom thresholds that are uploaded in the **Customized Interface Thresholds** page and the rest of the interfaces using the **Utilization Threshold** value configured on this page.
- **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 amount 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 how much extra traffic 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 106. 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, and click **SR-MPLS > Prefixes** tab > **Expand All**.

Note Prior to using this option, you must create device tag groups that you want to assign certain MSD values to. 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 26: Customized Interface Thresholds

Interface	Threshold (%)	Select to Delete
F4.cisco.com	14	<input type="checkbox"/>
F3.cisco.com	13	<input type="checkbox"/>
F5.cisco.com	11	<input type="checkbox"/>
F1.cisco.com	20	<input type="checkbox"/>
F3.cisco.com	10	<input type="checkbox"/>

Callout No.	Description
1	Interfaces to Monitor: Displays the option that is currently configured in the Configure LCM page.

Callout No.	Description
2	Import CSV File: All interfaces currently in the table will be replaced with the data in the CSV file you import.
3	Add: Click this icon to add new interface threshold rows.
4	Export CSV File: All interfaces are exported to a CSV file. You cannot filter data for export.
5	Edit Mode: When Edit Mode is ON , you can edit multiple fields in one session, then click Save .
6	Filter: By default, this row is available for you to enter text in which to filter content. To disable or enable the filtering feature, click  .
7	Select to Delete: 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-ID > ... > 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 27: Add First Interface



Node	Interface	Threshold (%)	Select for Dele... 
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>

- Click  to add more interfaces.

- Step 4** Confirm that the information appears correctly in 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. For more information, see Figure 15.

Monitor LCM Operations

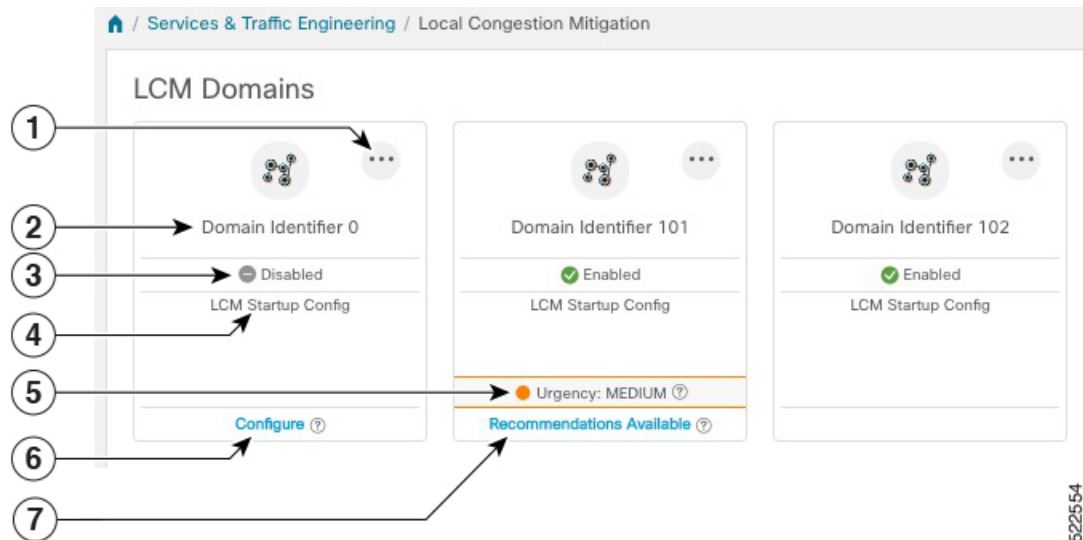


Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 108](#) topic.

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 28: 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 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 had been enabled for the domain.

Callout No.	Description
4	LCM Configuration Description: The description is defined in the Configure LCM page. The default description is "LCM Startup Config".
5	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 117.
6	Configure: This link appears if LCM has not yet been configured. Click Configure to go to the Configure LCM page.
7	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 .

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. For each interface, it lists details such as current utilization, recommended action, status, expected utilization after committing recommendations, and so on. You can also preview TTE policies prior to deployment (**... > Preview Solution**) or to verify deployment (**... > View Deployed Policies**) visually on a topology map. Hover the mouse pointer over  to view a description of what type of information each column provides. 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 29: LCM Operational Dashboard

Operational Dashboard

● Congested Interfaces (2) | ● Mitigating Interfaces (0) | ● Mitigated Interfaces (1)

Commit All | Urgency: MEDIUM

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F3.cisco.com	GigabitEthernet0/0/0/1	13%	4.59%	Mitigated	2	OK	Delete Set	None	8.14%	20-Apr-2022 09:07:44 PM PDT	...
F3.cisco.com	GigabitEthernet0/0/0/2	10%	13.47%	Congested	0	-	No Solution	None	-	25-Apr-2022 04:32:10 PM PDT	...
F5.cisco.com	GigabitEthernet0/0/0/1	11%	13.56%	Congested	0	-	Create Set	None	6.78%	20-Apr-2022 08:52:43 PM PDT	...

In this example, the following information is conveyed:

- f3.cisco.com with interface GogabitEthernet0/0/1—The current LCM state is Mitigated and 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.
- f3.cisco.com with interface GogabitEthernet0/0/2—LCM detects congestion however it cannot find bypass policies to remediate the congestion (**Recommended Action - No Solution**).



Note If LCM cannot find a solution (**No Solution**), it may be due to constraints enabled in the **LCM Configuration** page. For more information, see [Configure LCM, on page 117](#).

- f5.cisco.com with interface GogabitEthernet0/0/1—LCM detects congestion and has recommended to deploy policies to remediate the congestion (**Recommended Action - Create Set**).

Recommendations are listed as part of a set, and if deployed, all changes are committed. You must click **Commit All** if you want to remediate the congestion on F5.cisco.com with interface GogabitEthernet0/0/1.



CHAPTER 10

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 for the delivery of 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 when there are large flows.



Note

- Functionality described within this section is only available with certain licensing options.
- When using the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering**.

This section contains the following topics:

- [BWoD Important Notes, on page 125](#)
- [PCC-Initiated BWoD SR-TE Policies, on page 126](#)
- [Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example, on page 128](#)
- [Configure Bandwidth on Demand, on page 129](#)
- [Troubleshoot BWoD, on page 130](#)

BWoD Important Notes

Consider the following information when using BWoD:

- You must have the Advanced RTM license package to use BWoD.
- 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 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 2 traffic updates from the Optimization Engine, BWoD will resume normal operation.

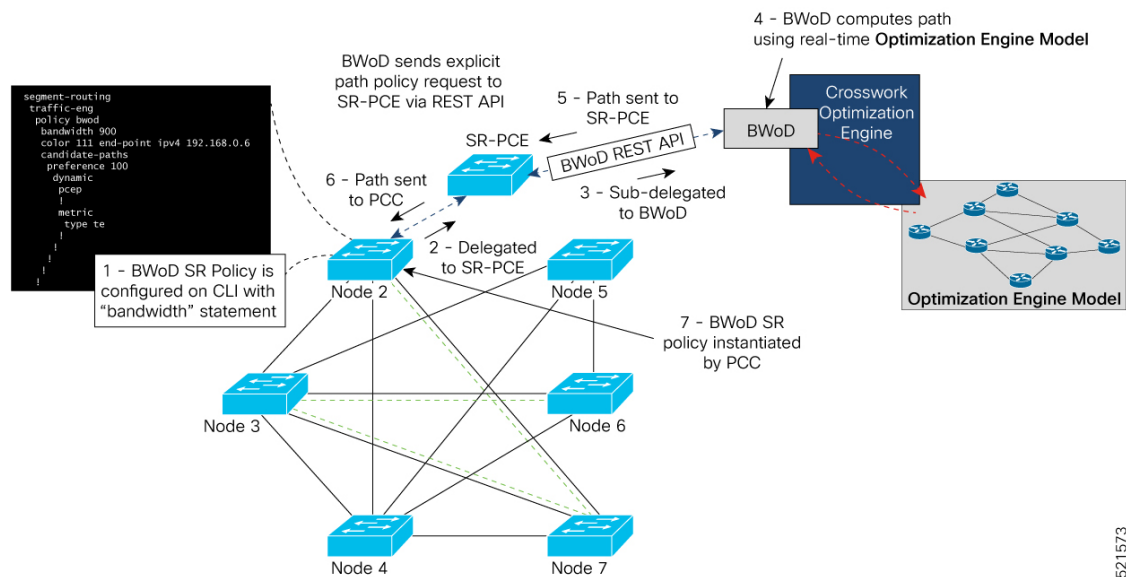
- If the Policy Violation advanced field is set to either **Strict Policy** or **Strict Network**, then the SR Policy Traffic option should be set to **Max Measured Requested**.

PCC-Initiated BWoD SR-TE Policies

When enabled, BWoD automatically connects to all SR-PCE providers configured in Crosswork Optimization Engine. The persistent connection is made to the SR-PCE BWoD Rest API, registering it as a PCE for bandwidth constrained SR-TE policies.

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

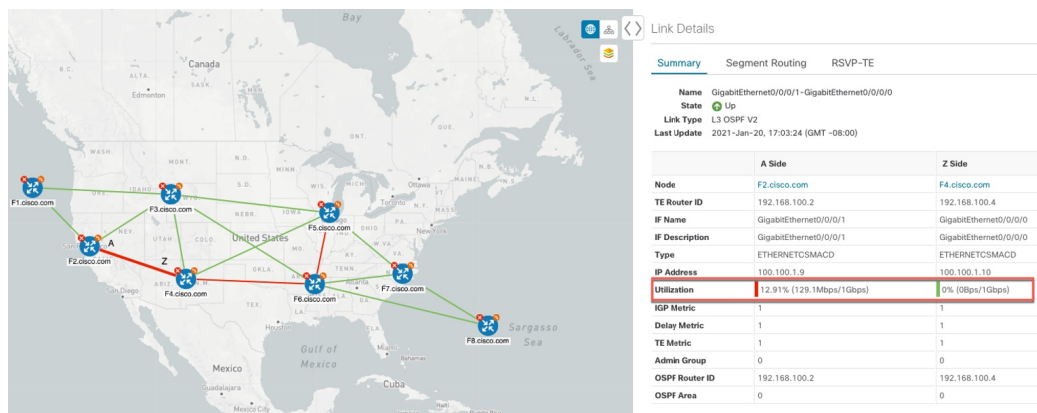
Figure 30: PCC-Initiated BWoD SR-TE Policies



521573


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

Figure 31: Initial BWoD Topology Example



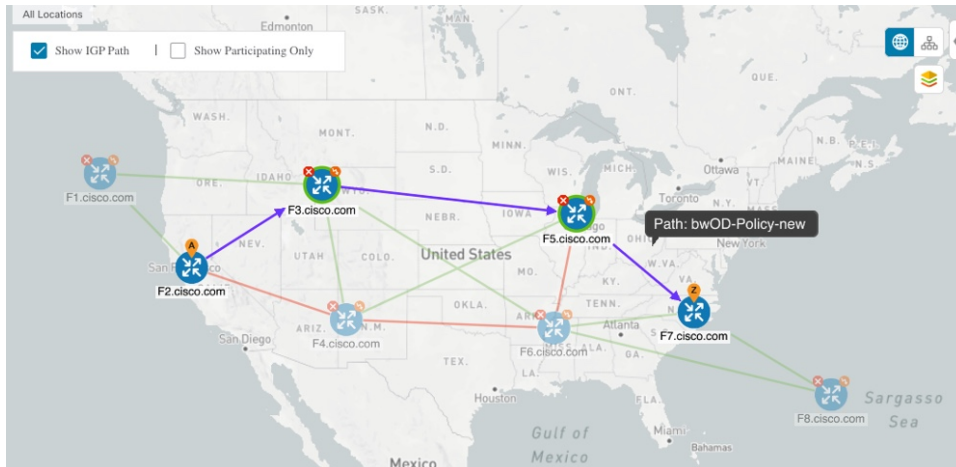
In this scenario we are using the above topology. The goal is to create a path from F2.cisco.com to F7.cisco.com that can accommodate 920 Mbps of traffic while keeping the utilization at 80%. The above example highlights the utilization on nodes F2.cisco.com and node F4.cisco.com to show that the link is being utilized and has a capacity of 1 Gbps. BWoD will initially try to find a single path that does not include this link since the addition of the requested bandwidth would exceed the utilization threshold.

Step 1 Enable and Configure BWoD.

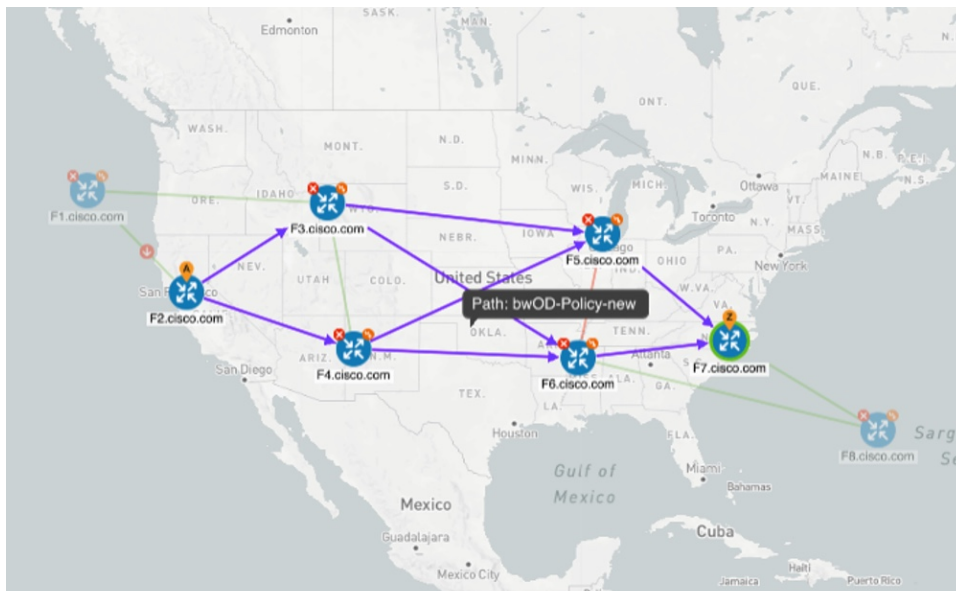
- From the main menu, choose **Traffic Engineering** > **Bandwidth on Demand** > **Configuration**.
- Toggle the Enable switch to **True** and enter **80** to set the utilization threshold percentage. To find descriptions of other options, simply hover the mouse over .
- Click **Commit Changes**.

Step 2 Create a PCE-initiated BWoD SR-TE Policy.

- From the main menu, choose **Traffic Engineering** > **SR-TE** tab and click **+Create**.
- Enter the required SR-TE policy details.
- In the **Policy Path** field, click **Bandwidth on Demand** and enter a unique name for the BWoD path. In this case, **bwOD-Policy-new**.
- From the **Optimization Objective** drop-down list, select **Traffic Engineering (TE) Metric**.
- In the **Bandwidth** field enter the requested bandwidth. In this case, we are requesting **920** Mbps.
- Click **Preview**.



In the below example, the BWoD SR policy uses the existing ECMP paths on the network without explicitly splitting the traffic to avoid congestion. The traffic may be distributed by ECMP, but BWoD does not influence that. It is only aware of it and takes it into consideration if it occurs on the path computed.



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

Step 3


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

- From the main menu, choose **Traffic Engineering** > **SR-TE**.
- Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View**). Note that the Policy Type is **Bandwidth on Demand**.

Configure Bandwidth on Demand

There are two parts to configure 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, navigate to **Traffic Engineering > Traffic Engineering**.
- Step 6** From the SR Policy table, click **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** Click **Preview** to view the proposed SR policy.
- Step 9** 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 7: 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 that is 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 which should suffice 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	Please contact your Cisco service representative.

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)

© 2020 Cisco Systems, Inc. All rights reserved.

