



Cisco Crosswork Optimization Engine 4.1 User Guide

First Published: 2022-09-07

Last Modified: 2022-10-30

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	Overview of Cisco Crosswork Optimization Engine	1
	Audience	1
	Overview of Cisco Crosswork Optimization Engine	2
	Crosswork Optimization Engine APIs	3
	Crosswork Network Controller Solution and Crosswork Optimization Engine	3
	Segment Routing Path Computation Element (SR-PCE)	3
	About Segment Routing	4
	About Resource Reservation Protocol (RSVP)	7
CHAPTER 2	Set Up and Monitor Your Network View	9
	Get a Quick View in the Dashboard	9
	View Devices and Links on the Topology Map	10
	View Device and Link Details	12
	Use Device Groups to Filter Your Topology View	17
	Create and Modify Device Groups	21
	Enable Dynamic Device Grouping	22
	Customize Map Display Settings	23
	Customize the Display of Links and Devices	23
	Set Display Behavior of Device Groups for TE Tunnels	23
	Customize the Display of Traffic Engineering	23
	Save Topology Views for Easy Access	24
CHAPTER 3	Visualize Traffic Engineering Services	27
	Get a Quick View of Traffic Engineering Services	27
	View TE Event and Utilization History	29

Configure TE Data Dashboard Settings	31
View Traffic Engineering Device Details	32

CHAPTER 4**Visualize SR-MPLS and SRv6 Policies 33**

View SR-MPLS and SRv6 Policies on the Topology Map	33
View SR-MPLS and SRv6 Policy Details	35
Visualize SR-MPLS or SRv6 Policies Example	36
Find Multiple Candidate Paths (MCPs)	43
Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label	46
Visualizing Native SR Paths	48
Visualize Native Path Device Prerequisites	50

CHAPTER 5**Visualize Flexible Algorithms 53**

Configure Flexible Algorithm Affinities	53
Visualize Flexible Algorithm	54
Find Flexible Algorithms for Links and Devices	56

CHAPTER 6**Visualize Tree-SID Policies 59**

View a Point-to-Multipoint Tree on the Topology Map	60
Limitations for Tree-SID Policies	61
Tree SID Configuration Example	64
Static Tree-SID Policy Configuration Example	64
Dynamic Tree-SID Policy Configuration Example with VRF	65
Dynamic Tree-SID Policy Configuration Example without VRF	70

CHAPTER 7**Visualize RSVP-TE Tunnels 73**

View RSVP-TE Tunnels on the Topology Map	73
View RSVP-TE Tunnel Details	75
View Traffic Engineering Device Details	77

CHAPTER 8**Provision SR-MPLS Policies 79**

SR-TE Policy Configuration Sources	79
PCC-Initiated SR-TE Policy Example	80

Create Explicit SR-MPLS Policies 80

Configure Link Affinities 81

Create Dynamic SR-MPLS Policies Based on Optimization Intent 82

Modify SR-MPLS Policies 83

CHAPTER 9

Provision RSVP-TE Tunnels 85

RSVP-TE Tunnel Configuration Sources 85

 PCC-Initiated RSVP-TE Tunnel Example 85

Create Explicit RSVP-TE Tunnels 86

Configure Link Affinities 86

Create Dynamic RSVP-TE Tunnels Based on Optimization Intent 88

Modify RSVP-TE Tunnels 89

CHAPTER 10

Use Local Congestion Mitigation (LCM) to Mitigate Network Congestion Locally 91

Local Congestion Mitigation Overview 91

LCM Important Notes 92

 LCM Platform Requirements 93

 BGP-LS Speaker Placement for Multiple AS Networks with a Dedicated IGP Instance Between ASBRs 93

LCM Calculation Workflow 94

Workflow Example: Mitigate Congestion on Local Interfaces 96

Configure LCM 105

Add Individual Interface Thresholds 107

Monitor LCM Operations 109

CHAPTER 11

Use Bandwidth Optimization (BWOpt) to Optimize the Network 113

Bandwidth Optimization Overview 113

BWOpt Important Notes 113

Automated Network Congestion Mitigation Example 115

Configure Bandwidth Optimization 118

Add Individual Interface Thresholds 118

Troubleshoot Bandwidth Optimization 119

CHAPTER 12

Define and Maintain Intent-Based Bandwidth Requirements 121

BWoD Important Notes	121
Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example	122
PCC-Initiated BWoD SR-TE Policies	124
Configure Bandwidth on Demand	125
Troubleshoot BWoD	125



CHAPTER 1

Overview of Cisco Crosswork Optimization Engine

This is a post-installation document intended to cover the steps required to get up and running with Cisco Crosswork Optimization Engine and start using the user interface (UI). For administrative tasks including device and user management, see the *Cisco Crosswork Infrastructure and Applications Administration Guide*.

- [Audience, on page 1](#)
- [Overview of Cisco Crosswork Optimization Engine, on page 2](#)
- [Crosswork Optimization Engine APIs, on page 3](#)
- [Crosswork Network Controller Solution and Crosswork Optimization Engine, on page 3](#)
- [Segment Routing Path Computation Element \(SR-PCE\), on page 3](#)
- [About Segment Routing, on page 4](#)
- [About Resource Reservation Protocol \(RSVP\), on page 7](#)

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) on Topology Map
- Flexible Algorithms

Overview of Cisco Crosswork Optimization Engine

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 provides the following:

- A topology map that gives valuable real-time network visualization of the following:
 - devices
 - links and link utilization
 - provisioned SR-TE (SR-MPLS and SRv6) policies and RSVP-TE tunnels



Note For more information, see [View Devices and Links on the Topology Map, on page 10](#)

- A UI that allows the network operator to perform the following tasks:
 - Provision SR-MPLS policies and RSVP-TE tunnels and modify or remove them using an intuitive workflow
 - Preview an SR-MPLS policy or RSVP-TE tunnel before deploying it to the network
 - Continuously track SR-MPLS policy dynamic path computations to maintain SLA objectives (with correct licensing)
 - Visualize SR-TE policies and RSVP-TE tunnels that are created directly on the network devices providing a comprehensive view of the active network configuration
 - Visualize Flexible Algorithms in the network.
 - Visualize Point to Multi Point (Tree-SID) on Topology Map.
- APIs that extend Crosswork Optimization Engine functions to other Crosswork applications and third party applications.
- Crosswork Optimization Engine feature packs (available with correct licensing) provide congestion mitigation and closed loop bandwidth optimization. A user defines the optimization intent and the tools implement the intent, and continuously monitor, track, and react to maintain the original intent.

This guide covers the capabilities that are allowed by the Crosswork Optimization Engine. 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 Sales representative to review an option described in the "Cisco Crosswork Optimization Engine Ordering Guide".

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](#).

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](#).

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**.

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.

About Segment Routing

Segment routing is a method of forwarding packets on the network that are based on the source routing paradigm. The source selects a path and encodes it in the packet header as an ordered list of segments. Segments are an identifier for any type of instruction. For example, topology segments identify the next hop toward a destination. The segment ID (SID) consisting of an unsigned 32-bit integer identifies each segment.

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

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 that is 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.

The following diagram shows a basic network with the Node SID and the Adjacency SID for each of the devices and connections between the devices noted.



Segment Routing Policies

An SR policy path is expressed as a list of segments that specifies the path (SID list). 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 that is required by the next segment.

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 headend 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 headend does not have enough information about the topology, the headend might delegate the computation to a path computation engine (PCE). If a path isn't found, then the policy becomes operationally down (operation status down) and packets will not be routed based on the policy.

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. 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 headend. The rest of the network executes the instructions embedded in the SID list.



Note For PCC-initiated policies, if the explicit path is configured in the form of IP addresses, the policy goes operational status down if one of the hops goes down. If it is configured as a list of labels, then the policy goes operational status down only if it is the first hop that goes down. The remaining hops are not resolved by the PCC and so it will not take the policy operational status down if they fail.

Segment Routing over MPLS (SR-MPLS)

Segment Routing can be applied on an MPLS data plane. In an SR-MPLS enabled network, an MPLS label represents an instruction. The source nodes programs the path to a destination in the packet header as a stack of labels. For more information, see [IETF RFC 8660 Segment Routing with the MPLS Data Plane](#).

Segment Routing over IPv6 (SRv6)

Segment Routing over IPv6 (SRv6) extends Segment Routing support with an IPv6 data plane. SRv6 introduces the Network Programming framework that enables a network operator or an application to specify a packet processing program by encoding a sequence of instructions in the IPv6 packet header. Each instruction is implemented on one or several nodes in the network and identified by an SRv6 Segment Identifier (SID) in the packet. For more information, see [IETF RFC 8986 SRv6 Network Programming](#).

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

For more information, see <https://www.segment-routing.net/>.

SRv6 Limitations

- Cisco IOS XR 7.3.2 only supports SRv6 visualization with IS-IS IGP.
- Traffic collection on SRv6 policies is not currently supported.
- OSPFv3 IGP (PCE-initiated) SRv6 policies are not supported.
- SRv6 is not supported on Bandwidth Optimization, Bandwidth on Demand, or Local Congestion Mitigation feature packs.
- IPv4 and IPv6 topologies must be congruent. Different link metrics for IPv4 and IPv6 are not supported.
- Visualization of PCC-initiated dynamic path SRv6 policies only. PCE-initiated and explicit path are not supported.

Segment Routing for Traffic Engineering

SR-TE takes place through a policy between a source and destination pair. SR-TE uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment.

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

Disjointness

Crosswork can use a disjoint policy to compute two unique paths that steer traffic from the same source and destination avoiding common specified resources (links or nodes). This results in no single point of failure in steering traffic through the network. 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.



Note

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

Tree-SID Policies

Tree Segment Identifier (Tree-SID) is modern controller driven multicast technology based on Segment Routing. It is a tree-building solution that uses a Segment Routing Path Computation Element (SR-PCE) using path computation element protocol (PCEP) to calculate point-to-multipoint (P2MP) trees using SR policies. Tree-SID uses a single MPLS label to build a multicast replication tree in an SR network. The advantage of having a controller is that any sort of constraints can be applied to calculate the tree.

See [View a Point-to-Multipoint Tree on the Topology Map, on page 60](#)

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.

See [Visualize Flexible Algorithm](#), on page 54

Related Links

[Provision SR-MPLS Policies](#), on page 79

[Configure Link Affinities](#), on page 81

About Resource Reservation Protocol (RSVP)

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

The RSVP-TE process contains the following functionalities:

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

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

RSVP-TE Explicit Routing (Strict, Loose)

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

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

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

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



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

RSVP FRR

When a router's link or neighboring device fails, the router often detects this failure by receiving an interface-down notification. When a router notices that an interface has gone down, it switches LSPs going out 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.



CHAPTER 2

Set Up and Monitor Your Network View


Familiarize yourself with the UI and set up your network view before managing SR policies and RSVP-TE tunnels. This section contains the following topics:

- [Get a Quick View in the Dashboard, on page 9](#)
- [View Devices and Links on the Topology Map, on page 10](#)
- [Use Device Groups to Filter Your Topology View, on page 17](#)
- [Customize Map Display Settings, on page 23](#)
- [Save Topology Views for Easy Access, on page 24](#)





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 of a series of dashlets, and each dashlet represents different types of data belonging to the same category.

Figure 1: Crosswork Home page



522573


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 what Cisco Crosswork applications are installed.
2	Dashlets: Information varies depending on what Cisco Crosswork applications are installed. <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.
3	Settings icons: <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 the network topology map, from the main menu choose **Topology**.




For more information, see [View Device and Link Details](#), on page 12.

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 33 and View RSVP-TE Tunnels on the Topology Map, on page 73.</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>
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 cursor 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>

Callout No.	Description
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>
6	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.
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 the Common EMS Services application and configure host information for Syslog and SNMP traps on the devices you want to view alarms for. For more information, see the <i>Cisco Crosswork Infrastructure and Applications Installation Guide</i> and the <i>Cisco Crosswork Infrastructure and Applications Administration Guide</i>. The Alarm Status feature is available for select licensing packages.</p>
8	The content of this window changes depending on what Show is set to for the Topology Map and if you have selected to view more information on a 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 and Link Details

This example shows how you can view device and link details (including Link Aggregation Group (LAG) details, see Step 6) using the topology map.

-
- 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 Traffic Engineering interface. On the left, a map shows a network topology with nodes and links. A tooltip for a device provides the following details:

- Reachability State: Reachable
- Host Name: xrv9k-VM12_771-151
- Node IP: 40.40.40.18
- Type: CISCO-XRV9000

On the right, the SR Policy table is shown with the following data:

Headend	Endpoint	Color	Admin ...	Oper S...	Actions	
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	277	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	766	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	600	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	10	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	10	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	10	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	366	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	266	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM...	2777	●	●	...
<input type="checkbox"/>	xrv9k-VM...	xrv9k-VM8	299	●	●	...

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 shows the Device Details panel for the device xrv9k-VM12_771-151. The details are as follows:

- Host Name:** xrv9k-VM12_771-151
- Reachability:** ● Reachable
- IP Address:** 40.40.40.18
- Civic Address:** Seattle, Washington, United States, North America, 94539
- Geo Location:** Latitude 36.780000, Longitude -91.500000
- Device Type:** ● Router
- Device Group:** All Locations > Unassigned Devices
- Product Type:** CISCO-XRV9000
- Connect To Device:** ● Telnet IPv4, ● SSH IPv4
- Last Update:** 10-May-2022 09:55:56 AM GMT+5:30

Routing Details:

- TE Router ID:** 192.168.4.15
- ISIS System ID:** 0000.0000.0037 Level-2
- ASN:** 65000
- PCEP Session:** PCE - 172.23.209.75, Source Address - 192.168.5.15
 - Stateful:** true
 - Source Address:** 192.168.5.15
 - Capability Instantiate:** true
 - Capability SR:** false
 - Capability Update:** true

Note

If the Alarm Status feature is enabled, you will also see Alarm information here. To view the Alarm Status, you must install the Common EMS Services application 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 Infrastructure and Applications Installation Guide* and the *Cisco Crosswork Infrastructure and Applications 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




Figure 4: Multiple IGP: ISIS Processes



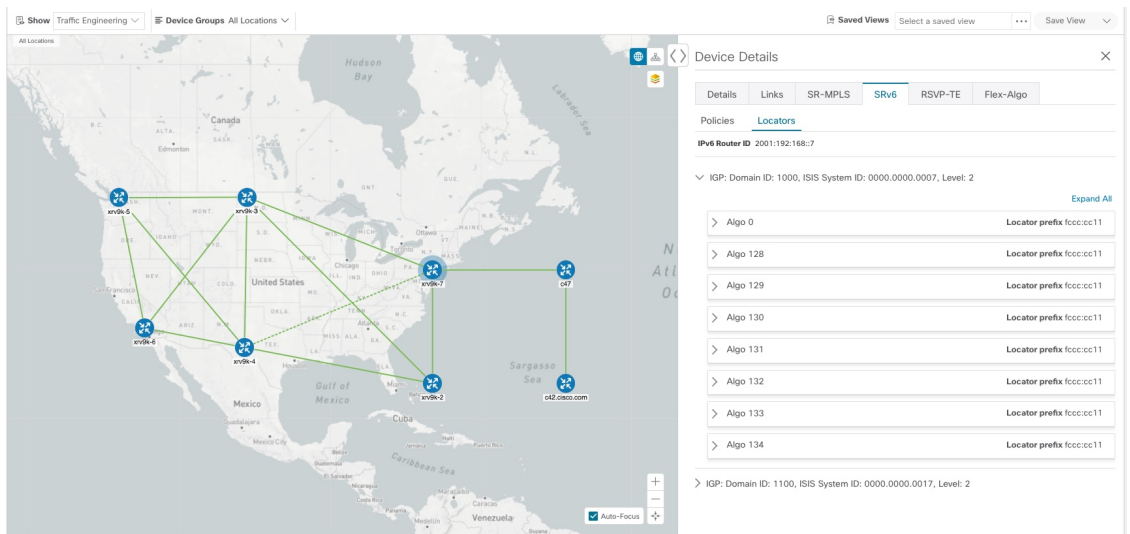


Figure 5: Multiple IGP: OSPF and ISIS Processes



- 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.

Links on Device PE-A

Total 4

State	Link Type	A Side	Utilization	Z Side
		Interface		Interface
	L3 ISIS IPv4	HundredGigE0/0/0/1	0% (0Bps/100Gbps)	
	L3 ISIS IPv6	HundredGigE0/0/0/0	0% (0Bps/100Gbps)	
	L3 ISIS IPv4	HundredGigE0/0/0/0	0% (0Bps/100Gbps)	
	L3 ISIS IPv6	HundredGigE0/0/0/1	0% (0Bps/100Gbps)	

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

The utilization shown on ipv4 and ipv6 links represents the aggregate traffic on the interface or sub-interface, not specific to each address family. The utilization shown on sub-interface links represents the bandwidth utilization on the main interface of the sub-interface's traffic.

Step 6 Collapse the side panel and close the **Device Details** window.

Step 7 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.

Links

Total 5

State	Link Type	A Side	Z Side
		Interface	Interface
	L3 ISIS IPv6	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/1
	L2 LLDP	GigabitEthernet0/0/0/6	GigabitEthernet0/0/0/6
	L3 ISIS IPv4	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/1
	L2 LLDP	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/1
	L2 LAG	Bundle-Ether2	Bundle-Ether2

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:

Links

Total 2

State	Link Type	A Si...	Z Si...	A Si...	Z Si...
Up	L2 LAG	Bundl...	Bundl...	0% (...)	0% (...)
Up	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...
Up	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)

Use Device Groups to Filter Your Topology View

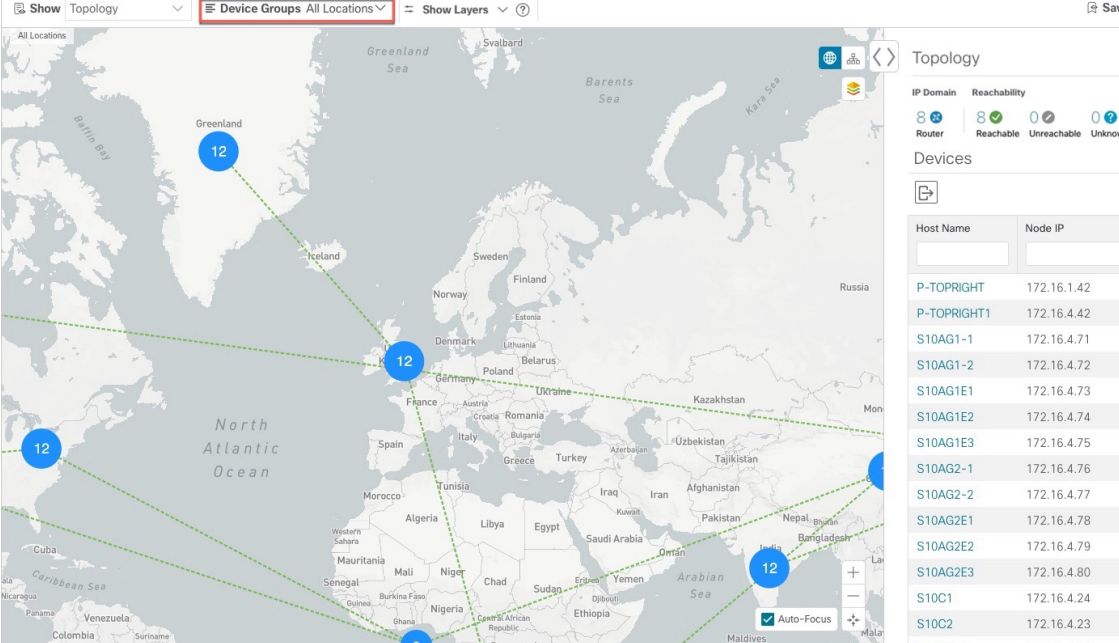
To help you identify, find, and group devices for a variety of purposes, you can create Device Groups. The Device Group window (**Device Management > Groups**) displays all devices and the device groups to which they belong. By default, all devices initially appear in the **Unassigned Devices** group.

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 4.1 interface. The main view is a geographical map of the world with network devices and links. The 'Device Groups' dropdown is set to 'All Locations'. The 'Devices' table on the right lists various host names and their corresponding node IP addresses.

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.

The screenshot shows the Cisco Crosswork Optimization Engine 4.1 interface. The 'Device Groups' dropdown is set to 'US West'. The logical map displays a network topology of devices in the US West region, including California and Nevada. The devices are connected in a mesh topology. The table on the right lists the devices and their IP addresses:


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.

Use Device Groups to Filter Your Topology View



- 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 network management interface. On the left, a topology map shows a mesh of devices labeled S7AG1E1 through S7AG2E3. On the right, a 'Devices' table is shown with the following data:

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


Device groups and assignment of devices to the groups can be done either manually (as described in this section) or automatically (as described in the next section).

- 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 gets added under **All Locations**.
- Step 3** To add a device to a group, from the right-pane, under **Unassigned Devices**, select a device and then from the **Move to Group** drop-down, select the appropriate group.
- Step 4** To edit, delete, or add a sub-group under an existing group, from the Device Groups tree, click  next to a group.

The screenshot shows the 'Device Management / Groups' interface. On the left, a tree view displays the following structure:

- Location
 - All Locations
 - User Defined
 - AAA (2)
 - DD
 - Edit Group Properties
 - Add a Sub-Group
 - ppp
 - Add Devices
 - Add Dynamic Rules
 - Delete Group
 - SSS

On the right, a large empty area contains a clipboard icon and the following text:

"Location" is an umbrella parent group that does not contain devices.
 Devices can be added to its sub-groups. Click the  button to create a new device group.
 Click on any group in the Device Groups pane to see all the devices in that group and move them if necessary.

Step 5 Choose to add, delete, or edit (rename or move) a 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.

Note Devices can belong to only one device group.

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.



Note 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 Dynamic Grouping Rule**.

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.

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 23](#)
- [Set Display Behavior of Device Groups for TE Tunnels , on page 23](#)

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.

Set Display Behavior of Device Groups for TE Tunnels

You can configure what is shown on the topology map when a device group is selected and a device in the selected 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.

Customize the Display of Traffic Engineering

To set Traffic Engineering display preferences, choose **Traffic Engineering > Traffic Engineering** 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.



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

- 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.
- Click **Metrics** to show IGP, TE, and delay (latency) metrics when viewing IGP paths. By default, these metrics are not enabled.



Note Metrics cannot be shown when the IGP path goes over an aggregate link. If you try to view an IPv6 network that has both IPv4 and IPv6 links you need to check the **Show Participating Only** checkbox to see IPv6 metrics.

- Click **Flex Algo** to show the Flex Algorithm paths. For more information see [Visualize Flexible Algorithms](#), on page 53.

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



Note 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 only the information you want and until the layout meets your needs.
- Step 2** When you have the view the way you want it, click **Save View**.

The screenshot shows the Cisco Crosswork Optimization Engine 4.1 interface. On the left, a geographical map displays a network topology with nodes labeled xrv9k-5, xrv9k-3, xrv9k-7, xrv9k-6, xrv9k-4, and srpce1. On the right, the 'Traffic Engineering' panel is visible, showing a table of SR Policies. The 'Save View' button in the top right corner of the interface is highlighted with a red box.

SR-MPLS	SRv6	RSVP-TE
15	15	0
PCE Init	PCC Init	Admin Down
		Oper Up
		Oper Down

SR POLICY						Selected 0 / Total 30
Hea...	End...	C...	Ad...	Op...	Actions	
<input type="checkbox"/>						
<input type="checkbox"/>	xrv9k-5	xrv9k-7	123...	↑	↑	...
<input type="checkbox"/>	xrv9k-5	xrv9k-7	222	↑	↑	...
<input type="checkbox"/>	xrv9k-5	xrv9k-7	333	↑	↑	...
<input type="checkbox"/>	xrv9k-6	xrv9k-7	607...	↑	↑	...
<input type="checkbox"/>	xrv9k-5	xrv9k-7	6521	↑	↑	...

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

Visualize Traffic Engineering Services

From the Traffic Engineering topology map, you can visualize the following TE services within your network:

- [Visualize SR-MPLS and SRv6 Policies](#)
- [Visualize Flexible Algorithms](#)
- [Visualize RSVP-TE Tunnels](#)
- [View a Point-to-Multipoint Tree on the Topology Map, on page 60](#)

The ability to visualize these services and use the Crosswork UI simplifies the process of monitoring and managing TE policies and tunnels.

This section applies to all TE services and describes how to:

- [Get a Quick View of Traffic Engineering Services , on page 27](#)
- [View TE Event and Utilization History, on page 29](#)
- [Configure TE Data Dashboard Settings, on page 31](#)
- [View Traffic Engineering Device Details, on page 32](#)

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**.

Get a Quick View of Traffic Engineering Services

1 → SR-MPLS dashlet showing Total Policy Count (4), Policy State (Oper Down: 0, Admin Down: 0, Oper Up: 4), and Policy Type & Metric Type breakdown.

2 → Filter: Policies and Tunnels Under Traffic Threshold 250 Kbps

3 → Table of Policy/Tunnel Type with columns: Headend, Endpoint, Color / ID, Policy / Tunnel Type, Metric Type, Traffic Rate (Kbps).


Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	Traffic Rate (Kbps)
PE-B	PE-A	70	SR-MPLS	IGP	0
PE-B	PE-C	1010	RSVP-TE	TE	0
PE-C	PE-B	1234	RSVP-TE	TE	0
PE-A	PE-B	1234	SR-MPLS	TE	0
PE-A	BOTTOM-LEFT	401	SR-MPLS	LATENCY	0
PE-A	PE-B	70	SR-MPLS	TE	0
PE-A	PE-B	123	RSVP-TE	TE	0
PE-A	PE-C	400	RSVP-TE	TE	0
PE-A	PE-C	417	RSVP-TE	TE	0
PE-A	PE-B	418	RSVP-TE	TE	0

4 → Policy and Tunnel Change Events table with columns: Headend, Endpoint, Color / ID, Policy / Tunnel Type, Metric Type, Events Total, Operational State Change, Path Change.

Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	Events Total	Operational State Cha...	Path Change
PE-A	BOTTOM-LEFT	401	SR-MPLS	LATENCY	2	1	1
PE-A	PE-B	70	SR-MPLS	TE	2	1	1
PE-A	PE-B	1234	SR-MPLS	TE	2	1	1
PE-B	PE-A	70	SR-MPLS	IGP	2	1	1
PE-A	PE-B	418	RSVP-TE	TE	2	1	1
PE-A	PE-C	400	RSVP-TE	TE	2	1	1

522714

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, and 1 day).</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>



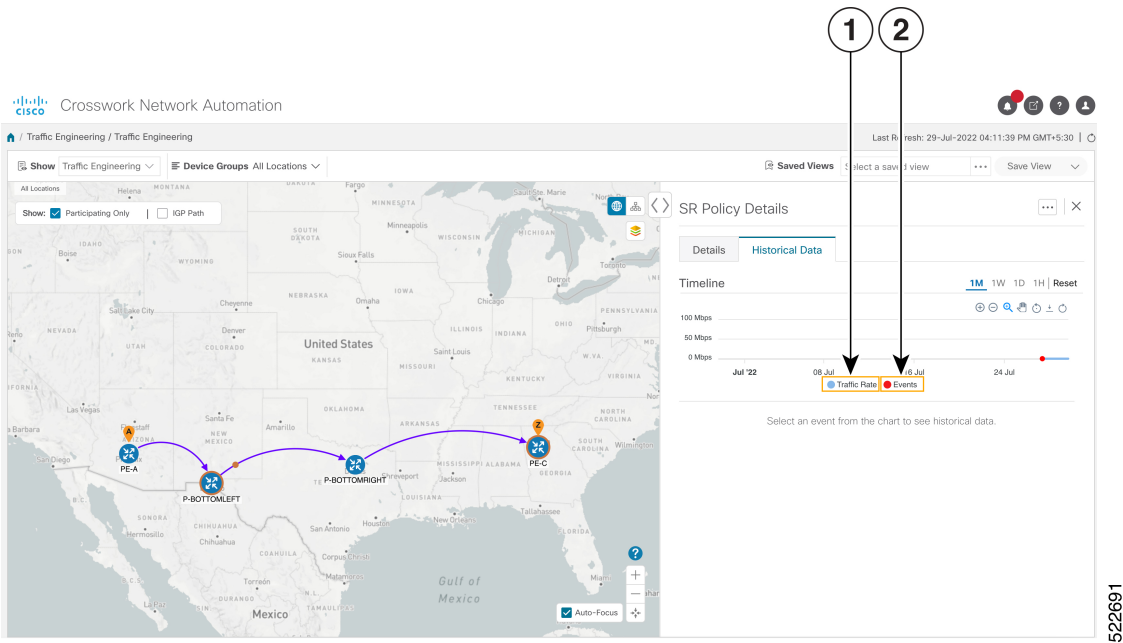
Note For a list of known limitations, see the [Cisco Crosswork Optimization Engine Release Notes](#).

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.

View TE Event and Utilization History




522691

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 at that point in time as shown in the following image: The path of the policy is displayed in the left pane.



Configure TE Data 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 > Traffic Engineering > TE Dashboard** .

The screenshot displays the 'TE Data Dashboard Settings' configuration page. On the left is a sidebar with a tree view containing categories like Servers, Maintenance Mode, Providers, Notifications, Topology, Traffic Engineering, and Affinity. The 'TE Dashboard' option is selected. The main area contains five settings, each with a callout box:

- 1**: LSP Metrics Collection (toggle switch set to On)
- 2**: LSP State Change Collection (toggle switch set to On)
- 3**: LSP Path Change Collection (toggle switch set to On)
- 4**: Retention Interval (input field with value 2, range 1 to 30 days)
- 5**: Traffic rate threshold for filtering underutilized LSPs (input field with value 1000, unit Kbps)

At the bottom of the main area are buttons for 'Save', 'Reset', and 'Reset to Default'. A vertical ID '522713' is on the right side.

Callout No.	Description
1	LSP Metric Collection: Turn on this field to capture the metric data in the TE Dashboard.
2	LSP State Change Collection: Turn on this field to capture the state change details in the TE Dashboard.
3	LSP Path Change Collection: Turn on this field to capture the path change details in the TE Dashboard.
	<p>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.</p> <p>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.</p>

Callout No.	Description
4	The LSPs for which the traffic has not exceeded the threshold value specified in this field are displayed under the Underutilized LSP dashlet in the TE Dashboard. The threshold value can also be configured on the dashlet.

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 **Device Details** page, click on the traffic engineering tab 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.

The screenshot displays the 'Device Details' panel for device xrv9k-16. The 'SR-MPLS' tab is selected, showing the following information:

- IGP: Domain ID: 1000, ISIS System ID: 0000.0000.0005, Level: 2
- SRGB 16000 - 23999
- SRLB 105000 - 105999
- MSD 10

Prefixes	Label	Algo
192.168.0.5	18115	0



CHAPTER 4

Visualize SR-MPLS and SRv6 Policies

Crosswork Optimization Engine allows you to visualize SR-MPLS and SRv6 policies in your network. The SR-PCE discovers policies and displays them in the Traffic Engineering topology map.

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

This section contains the following topics:

- [View SR-MPLS and SRv6 Policies on the Topology Map, on page 33](#)
- [View SR-MPLS and SRv6 Policy Details, on page 35](#)
- [Visualize SR-MPLS or SRv6 Policies Example, on page 36](#)
- [Find Multiple Candidate Paths \(MCPs\), on page 43](#)
- [Visualize Underlying Paths Associated with a Defined Binding-Segment ID \(B-SID\) Label, on page 46](#)
- [Visualizing Native SR Paths, on page 48](#)

View SR-MPLS and SRv6 Policies on the Topology Map

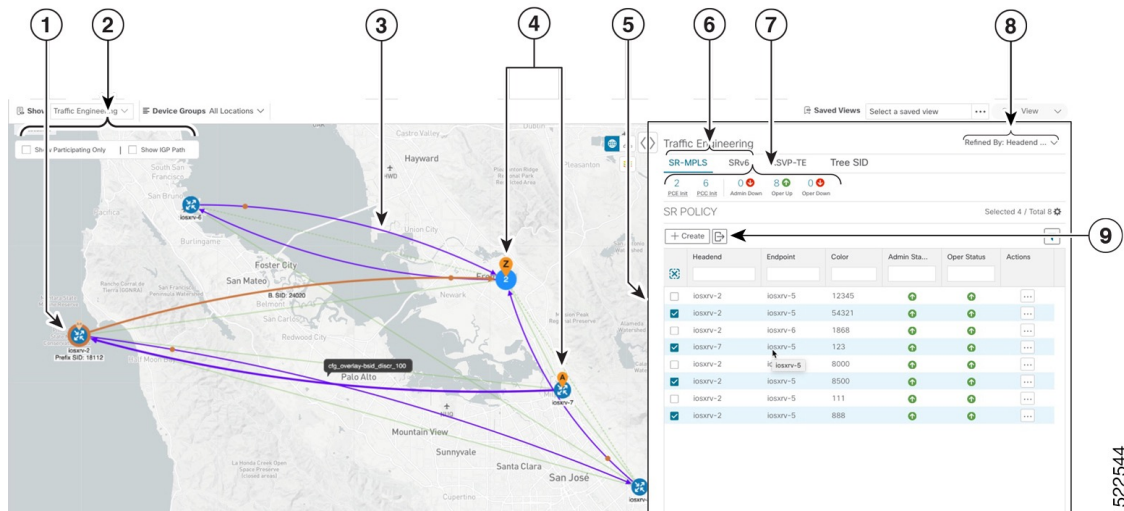
Crosswork Optimization Engine visualization provides the most value by giving you the ability to easily view and manage SR-MPLS and SRv6 policies. By visually examining your network, the complexity of provisioning and managing these SR-TE policies is significantly reduced.



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



Note Throughout this section, the navigation is documented as **Traffic Engineering > Traffic Engineering**. However, when using Crosswork Optimization Engine within the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering** and select either the **SR-MPLS** or **SRv6** tabs.

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. Depending on what is selected on the topology map, or whether you are in the process of viewing and managing SR-TE policies, you can do the following: <ul style="list-style-type: none"> • Visualize SR-MPLS or SRv6 Policies Example, on page 36 • Provision SR-MPLS Policies, on page 79 • View Device and Link Details, on page 12

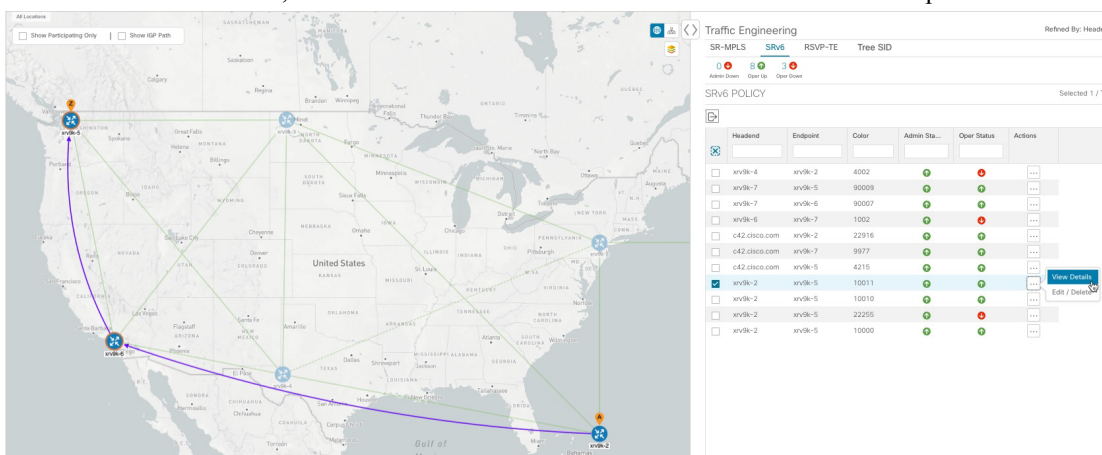
522544

Callout No.	Description
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.
8	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. Filter options: <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 policy details such as disjoint groups, metric type, candidate path, segment hop information, and so on.

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 of the United States shows a network topology with nodes and links. On the right, a table lists SRv6 policies. The table has columns for 'Headend', 'Endpoint', 'Color', 'Admin Sta...', 'Oper Status', and 'Actions'. One policy is selected, and the 'View Details' button in the 'Actions' column is highlighted.

Headend	Endpoint	Color	Admin Sta...	Oper Status	Actions
<input type="checkbox"/> svr6-4	svr6-2	4502		●	
<input type="checkbox"/> svr6-7	svr6-5	90009		●	
<input type="checkbox"/> svr6-7	svr6-5	90007		●	
<input type="checkbox"/> svr6-6	svr6-7	1002		●	
<input type="checkbox"/> c42.cisco.com	svr6-2	22516		●	
<input type="checkbox"/> c42.cisco.com	svr6-7	9977		●	
<input type="checkbox"/> c42.cisco.com	svr6-5	4215		●	
<input checked="" type="checkbox"/> svr6-2	svr6-5	10011		●	 View Details Edit / Delete
<input type="checkbox"/> svr6-2	svr6-5	10010		●	
<input type="checkbox"/> svr6-2	svr6-5	22255		●	
<input type="checkbox"/> svr6-2	svr6-5	10000		●	

Step 2 View SR-MPLS or SRv6 policy details.

522553

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 a network management interface with a 'Details' tab. The 'Summary' section lists various policy attributes:

- Admin State: Up
- Oper State: Up
- Binding SID: 24006
- Policy Type: Regular
- Profile ID: -
- Description: -
- Traffic Rate: 0 Mbps
- Unused: True
- Delay: 15 (with an information icon 'i')
 - Tooltip: Last Updated 16-May-2022 10:23:15 AM GMT+5:30
- 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

Below the summary is a 'Candidate Path' table:

Path Name	Preference	Path Type	State
cto-test-1-discr-100	100	Unknown	Up

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. The topology map displays SR-TE policies that are provisioned using the UI along with policies that are discovered from the network by SR-PCE. Then you can drill down to details and visualization of participating SR-TE policies.

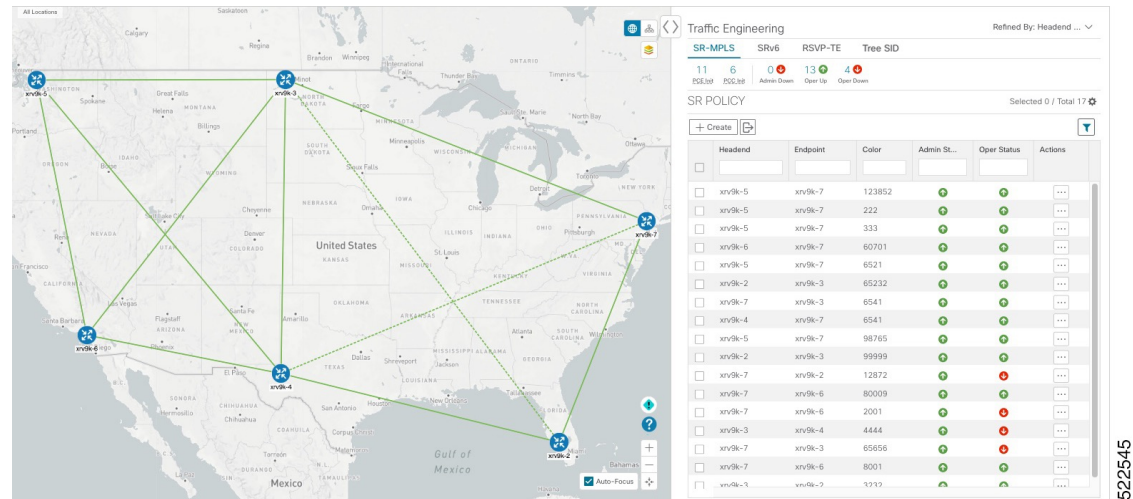
In this example, we assume that devices and SR-MPLS policies have been added and device groups have been created.



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

Click images to zoom in for a closer look.

Figure 7: Topology Map Example



522545

Step 1

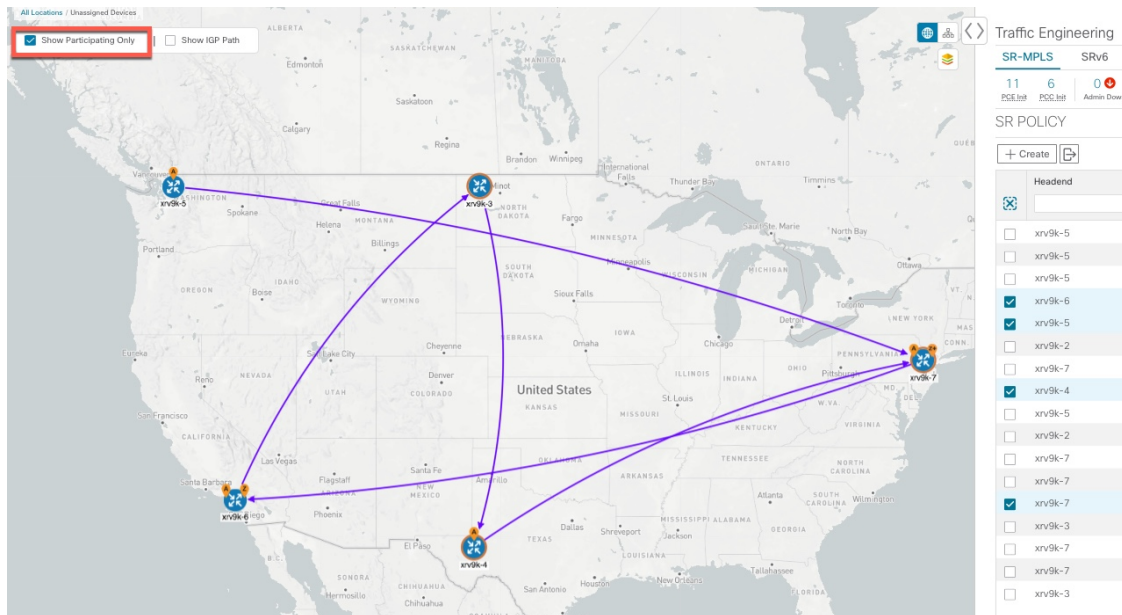
Select SR-MPLS policies for visualization and isolate them on the map.

- 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 **xrv9k-7** node is the destination for two of the selected policies. Both **xrv9k-3** and **xrv9k-2** 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.
- The orange outline (🔴) indicates that **xrv9k-3**, **xrv9k-7**, and **xrv9k-4** have node SIDs.

Visualize SR-MPLS or SRv6 Policies Example




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 **xrv9k-6 > xrv9k-3 > xrv9k-4 > xrv9k-7**.
- The prefix SID for xrv9k-3 and xrv9k-4 are displayed.
- The path name is displayed: **60701-expl-napo**



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.

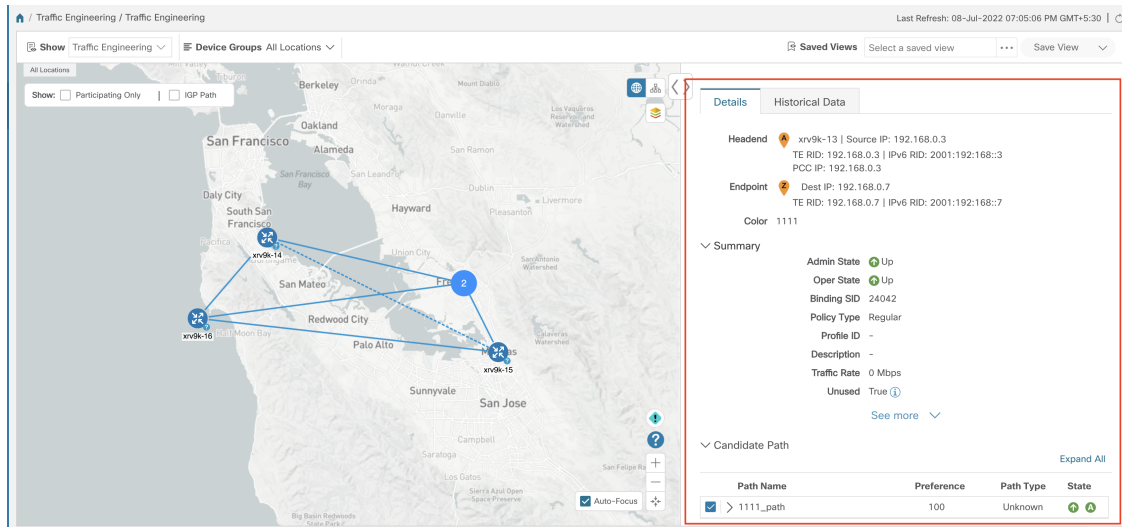
Note You must check the **Show IGP Path** check box in order to view metrics.



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.

a) 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.



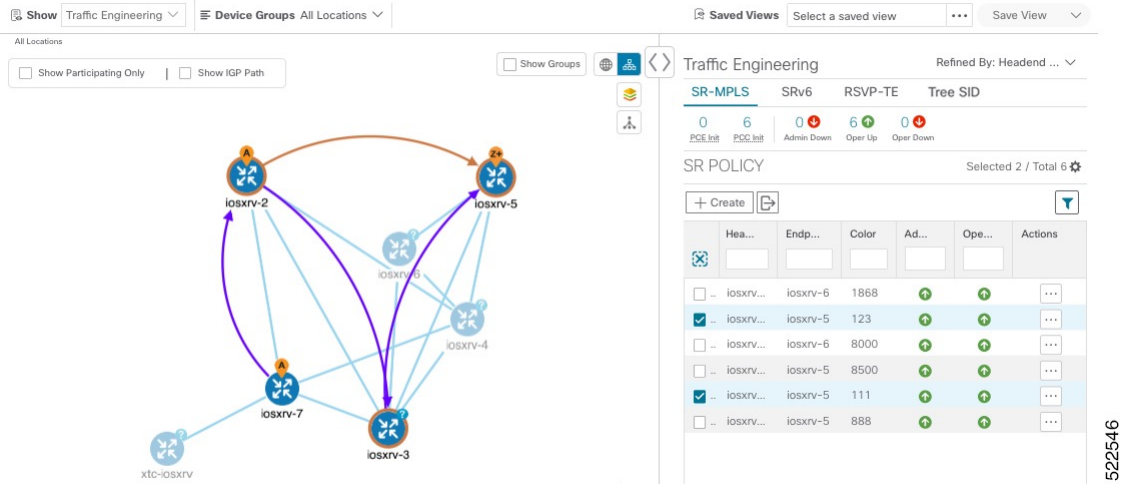
Step 5

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 (SR-MPLS Policies Selected)



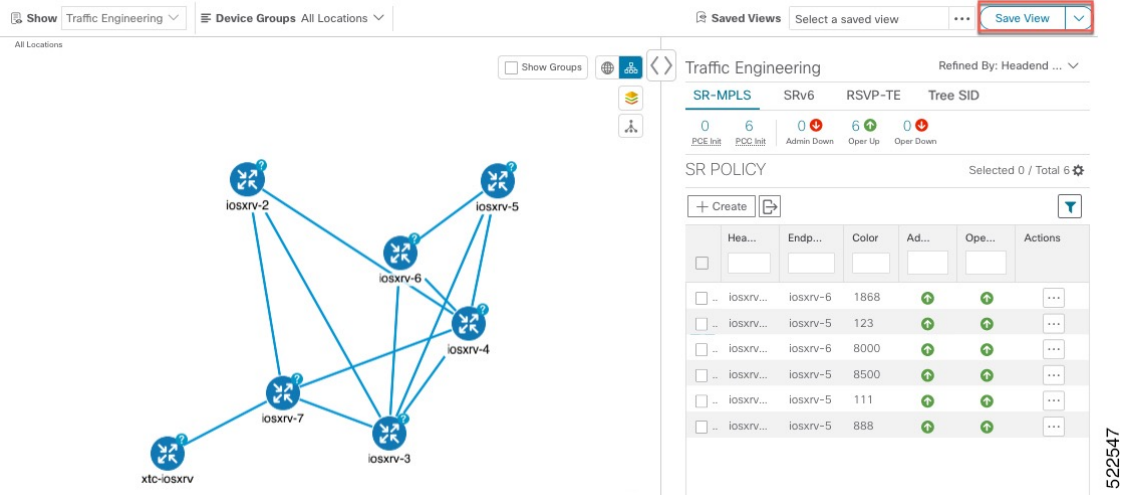
The screenshot shows the 'Traffic Engineering' interface. On the left, a network diagram displays several nodes (iosxrv-2, iosxrv-3, iosxrv-4, iosxrv-5, iosxrv-6, iosxrv-7, xtc-iosxrv) connected by lines. Some connections are highlighted in purple and orange. On the right, the control panel shows 'SR-MPLS' selected. Below the tabs, there are statistics for PCE Init, Admin Down, Oper Up, and Oper Down. The 'SR POLICY' section shows a table with 6 rows. Two rows are checked: iosxrv-2 to iosxrv-5 (123) and iosxrv-3 to iosxrv-5 (111). The 'Save View' button in the top right of the control panel is highlighted with a red box.

Hea...	Endp...	Color	Ad...	Ope...	Actions
<input type="checkbox"/>	iosxrv-2	iosxrv-6	1868	+	+
<input checked="" type="checkbox"/>	iosxrv-2	iosxrv-5	123	+	+
<input type="checkbox"/>	iosxrv-3	iosxrv-6	8000	+	+
<input type="checkbox"/>	iosxrv-3	iosxrv-5	8500	+	+
<input checked="" type="checkbox"/>	iosxrv-3	iosxrv-5	111	+	+
<input type="checkbox"/>	iosxrv-7	iosxrv-5	888	+	+

522546

Example:

Figure 9: Logical Map (Save Without SR-MPLS Policies Selected)



The screenshot shows the same network diagram as Figure 8, but with no SR-MPLS policies selected. The control panel shows 'SR-MPLS' selected, but the table below it has all checkboxes unchecked. The 'Save View' button in the top right of the control panel is highlighted with a red box.

Hea...	Endp...	Color	Ad...	Ope...	Actions
<input type="checkbox"/>	iosxrv-2	iosxrv-6	1868	+	+
<input type="checkbox"/>	iosxrv-2	iosxrv-5	123	+	+
<input type="checkbox"/>	iosxrv-3	iosxrv-6	8000	+	+
<input type="checkbox"/>	iosxrv-3	iosxrv-5	8500	+	+
<input type="checkbox"/>	iosxrv-7	iosxrv-5	111	+	+
<input type="checkbox"/>	iosxrv-7	iosxrv-5	888	+	+

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 Traffic Engineering interface. On the left, a network map displays various geographical locations: United Kingdom, Greenland, US Canada, Japan, India, South Africa, Australia, and China. The 'Device Groups' dropdown is currently set to 'All Locations'. On the right, the 'SR POLICY' table is visible, listing various policies with their headend, endpoint, color, and status.

Headend	Endpoint	Color	Admin St...	Oper Sta...	Actions	
<input type="checkbox"/>	S1AG1E1	SSAG1E2	63212	🟢	🟢	...
<input type="checkbox"/>	S8AG1-1	S6C1	5522	🟢	🟢	...
<input type="checkbox"/>	S5AG1-1	S3AG1-1	102	🟢	🟢	...
<input type="checkbox"/>	S5C1	S7C1	22332	🟢	🟢	...
<input type="checkbox"/>	S10C1	S3C1	5123	🟢	🟢	...
<input type="checkbox"/>	S10AG2E3	S6AG2E1	3215	🟢	🟢	...
<input type="checkbox"/>	S2AG1-1	S2AG2-2	106	🟢	🟢	...
<input type="checkbox"/>	S10AG1-1	S10AG2E2	434	🟢	🟢	...
<input type="checkbox"/>	S2AG1E3	S2C1	6325	🟢	🟢	...
<input type="checkbox"/>	S10AG1-1	S10AG1E1	6325	🟢	🟢	...
<input type="checkbox"/>	S5C2	P-TOPRIGHT	100	🟢	🟢	...
<input type="checkbox"/>	S1C2	S2C1	100	🟢	🟢	...
<input type="checkbox"/>	S2AG1-1	S2AG2E3	100	🟢	🟢	...
<input type="checkbox"/>	S1AG2E1	S1AG1E2	100	🟢	🟢	...

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 with 'Device Groups' set to 'Australia'. The network map now displays only devices within the Australia region, including S1AG1E1, S1AG1E2, S1AG1E3, S1AG1-1, S1AG1-2, S1C1, S1AG2-1, S1AG2E1, S1AG2E2, S1AG2E3, S1C2, S1AG2-2, and S1AG2E1. The 'SR POLICY' table on the right is filtered to show only policies relevant to the selected group.

Headend	Endpoint	Color	Admin St...	Oper Sta...	Actions	
<input type="checkbox"/>	S1AG1E1	SSAG1E2	63212	🟢	🟢	...
<input type="checkbox"/>	S1C2	S2C1	100	🟢	🟢	...
<input checked="" type="checkbox"/>	S1AG2E1	S1AG1E2	100	🟢	🟢	...
<input type="checkbox"/>	S5C1	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 [Set Display Behavior of Device Groups for TE Tunnels](#), on page 23.

The screenshot shows the Traffic Engineering interface with a dialog box overlaid. The dialog box contains the 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." Below the text are two buttons: "Switch Device Group" and "Don't Switch". The background interface shows the same network map and policy table as in Step 8.

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.

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 displays the Cisco Crosswork Network Automation interface. On the left, a network topology map shows various geographical locations: United Kingdom, Australia, South Africa, China, US Canada, Japan, and India. A path is highlighted through the network. On the right, the 'Traffic Engineering' dashboard is visible, showing a summary of SR-MPLS policies (15 PCE Init, 6 PCC Init, 0 Admin Down, 17 Oper Up, 4 Oper Down) and a table of selected SR-MPLS policies.

Headend	Endpoint	Color	Admin St...	Oper Stat...	Actions	
<input checked="" type="checkbox"/>	S1AG1E1	S5AG1E2	63212	↑	↑	...
<input type="checkbox"/>	S8AG1-1	S6C1	5522	↑	↑	...
<input type="checkbox"/>	S8AG1-1	S3AG1-1	102	↑	↑	...
<input type="checkbox"/>	SSC1	S7C1	22332	↑	↑	...
<input type="checkbox"/>	S10C1	S3C1	5123	↑	↓	...
<input type="checkbox"/>	S10AG2E3	S8AG2E1	3215	↑	↑	...
<input type="checkbox"/>	S2AG1-1	S2AG2-2	106	↑	↑	...
<input type="checkbox"/>	S10AG1-1	S10AG2E2	434	↑	↑	...
<input type="checkbox"/>	S2AG1E3	S2C1	6325	↑	↑	...
<input type="checkbox"/>	S10AG1-1	S10AG1E1	6325	↑	↑	...
<input type="checkbox"/>	SSC2	P-TOPRIGHT	100	↑	↑	...
<input type="checkbox"/>	S1C2	S2C1	100	↑	↑	...
<input type="checkbox"/>	S2AG1-1	S2AG2E3	100	↑	↑	...
<input checked="" type="checkbox"/>	S1AG2E1	S1AG1E2	100	↑	↑	...
<input type="checkbox"/>	SSC1	S1C2	111	↑	↓	...

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 Cisco Crosswork Network Automation interface with a filtered network topology. The topology map displays several nodes labeled with ASR and NCS identifiers (e.g., PE4-ASRtk.cisco.com, P3-NCS501, PE1-ASRtk). On the right, the 'Traffic Engineering' dashboard shows a summary of SR-MPLS policies (4 PCE Init, 1 PCC Init, 0 Admin Down, 5 Oper Up, 0 Oper Down) and a table of selected SR-MPLS policies. A red box highlights the 'Filters Applied (1)' dropdown and the filtered table.

Headend	Endpoint	Color	Admi...	Oper ...	Actions	
<input type="checkbox"/>	PE1-AS...	P3-NCS...	345	↑	↑	...
<input type="checkbox"/>	PE4-AS...	PE7-XR...	123	↑	↑	...
<input type="checkbox"/>	PE7-XR...	P4-NCS...	234	↑	↑	...
<input type="checkbox"/>	PE4-AS...	PE2-AS...	2258	↑	↑	...

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.

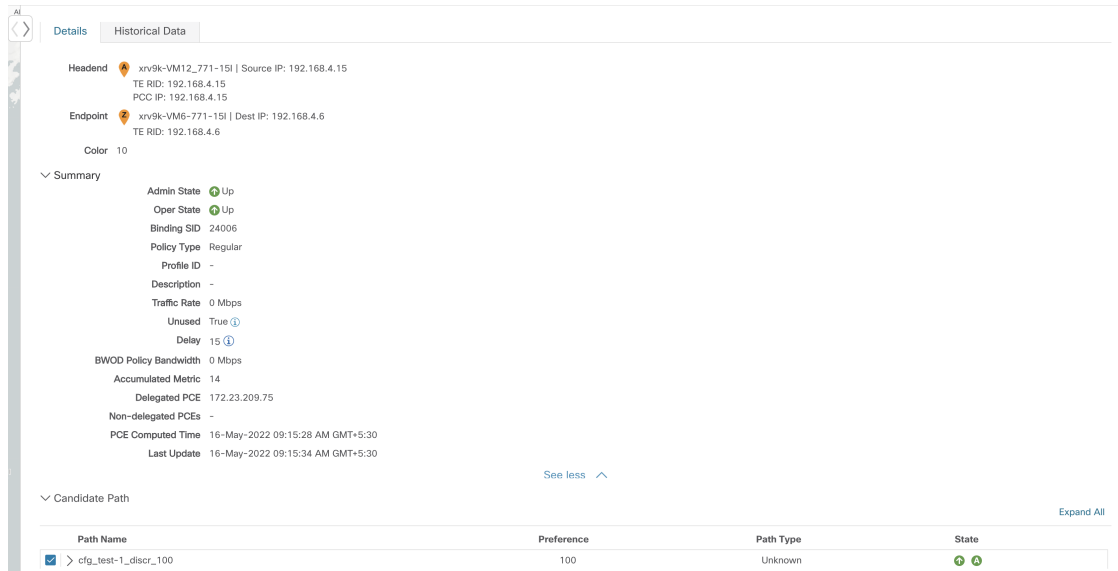
In this example, you see that the active path is going from **iosxrv-2 > iosxrv-3 > iosxrv-5**.



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 status 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:**

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

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.

SR Policy Details

PCE Computed Time 26-Aug-2021 03:31:10 PM PDT
Last Update 26-Aug-2021 03:39:23 PM PDT

Candidate Path Collapse All

Path Name	Preference	Path Type
<input type="checkbox"/> <input checked="" type="checkbox"/> ▼ cfg_test_mcp_diff_paths_discr_10000	10000	Unknown

Segm...	Segment Type	Label	Algo	IP	Node	Interface	Sid T...
0	IGP Adj SID	24002	0	10.0.0.9	iosxrv-2		P
1	IGP Adj SID	24012	0	10.0.0.25	iosxrv-3		P

Path Name cfg_test_mcp_diff_paths_discr_10000
Policy Type Unknown
Metric Type TE
Disjoint Group ID:
 Association Source: -
 Type: -
PCE Initiated false
Affinity Exclude-Any: -
 Include-Any: -
 Include-All: -

Segm...	Segment Type	Label	Algo	IP	Node	Interface	Sid T...
0	Node SID	18115	0	192.168.0.5	iosxrv-5		

Path Name cfg_test_mcp_diff_paths_discr_5000
Policy Type Unknown
Metric Type IGP
Disjoint Group ID:
 Association Source: -
 Type: -
PCE Initiated false
Affinity Exclude-Any: -
 Include-Any: -
 Include-All: -

- 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**.

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

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 callout box points to a specific hop in this path, labeled '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 orange. Below the table, the details for this path are shown, including its name, policy type, metric type, disjoint group, 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.



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**.

SR Policy Details

Details | Historical Data

Headend A cw-asr23 | Source IP: 3.3.3.100
TE RID: 3.3.3.100 | IPv6 RID: fb00:3:3::100
PCC IP: 3.3.3.100

Endpoint Z cw-xrv52 | Dest IP: 3.3.3.52
TE RID: 3.3.3.52 | IPv6 RID: fb00:3:3::52

Color 2222

Summary

- Admin State ↑ Up
- Oper State ↑ Up
- Binding SID 24011
- Policy Type Regular
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True i

See more v

Candidate Path Collapse All

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> v 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	cw...

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**.

The screenshot displays the SR Policy Details window. On the left, a map of the United States shows a path starting at cw-xrv51 (Chicago), going to cw-xrv55 (Denver), then to cw-xrv54 (Minneapolis), and finally to cw-xrv52 (St. Louis). On the right, the details panel shows the following information:

- Headend:** cw-xrv51 | Source IP: 3.3.3.51, TE RID: 3.3.3.51 | IPv6 RID: fb00:3:3:51, PCC IP: 3.3.3.51
- Endpoint:** cw-xrv52 | Dest IP: 3.3.3.52, TE RID: 3.3.3.52 | IPv6 RID: fb00:3:3:52, Color: 3333
- Summary:** Admin State: Up, Oper State: Up, Binding SID: 15790, Policy Type: Regular, Profile ID: -, Description: -, Traffic Rate: 0 Mbps, Unused: True
- Candidate Path:**

Path Name	Preference	Path Type	State
cfg_bsid-policy1_discr_100	100	Unknown	Up
- Path Name:** cfg_bsid-policy1_discr_100
- Oper State:** Up | Active
- Metric Type:** TE
- Disjoint Group ID:**

Visualizing Native SR Paths

Crosswork Optimization Engine allows you to visualize the Native SR paths. Since this feature uses multipaths, all ECMP paths will be shown between the source and destination. 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.



Note This is applicable only for SR-MPLS policies.

To create a path query, do the following:

Before you begin


Confirm that device requirements are met. See [Visualize Native Path Device Prerequisites](#), on page 50.

- 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 10: Path Details




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 11: Path Details



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. 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:


```
grpc
  port 50000
  no-tls
  address-family dual
  !
mpls oam
!
```



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-7.2.1#config
```

```
RP/0/RP0/CPU0:xrvr-7.2.1(config)#router static
```

```
RP/0/RP0/CPU0:xrvr-7.2.1(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-7.2.1(config-static)#commit
```




CHAPTER 5

Visualize 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.



Note You cannot filter Flexible Algorithms on multiple domains.

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 Throughout this section, the navigation is documented as **Traffic Engineering > Traffic Engineering**. However, when using the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering**.

- [Configure Flexible Algorithm Affinities, on page 53](#)
- [Visualize Flexible Algorithm, on page 54](#)
- [Find Flexible Algorithms for Links and Devices, on page 56](#)

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 > Traffic Engineering > Affinity > Flex-Algo Affinities** tab.

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

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

Example:

Name ?	Bit Position (0-255) ?	Actions
<input type="text"/>	<input type="text"/>	
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 56.

Visualize Flexible Algorithm

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 82.

Before you begin

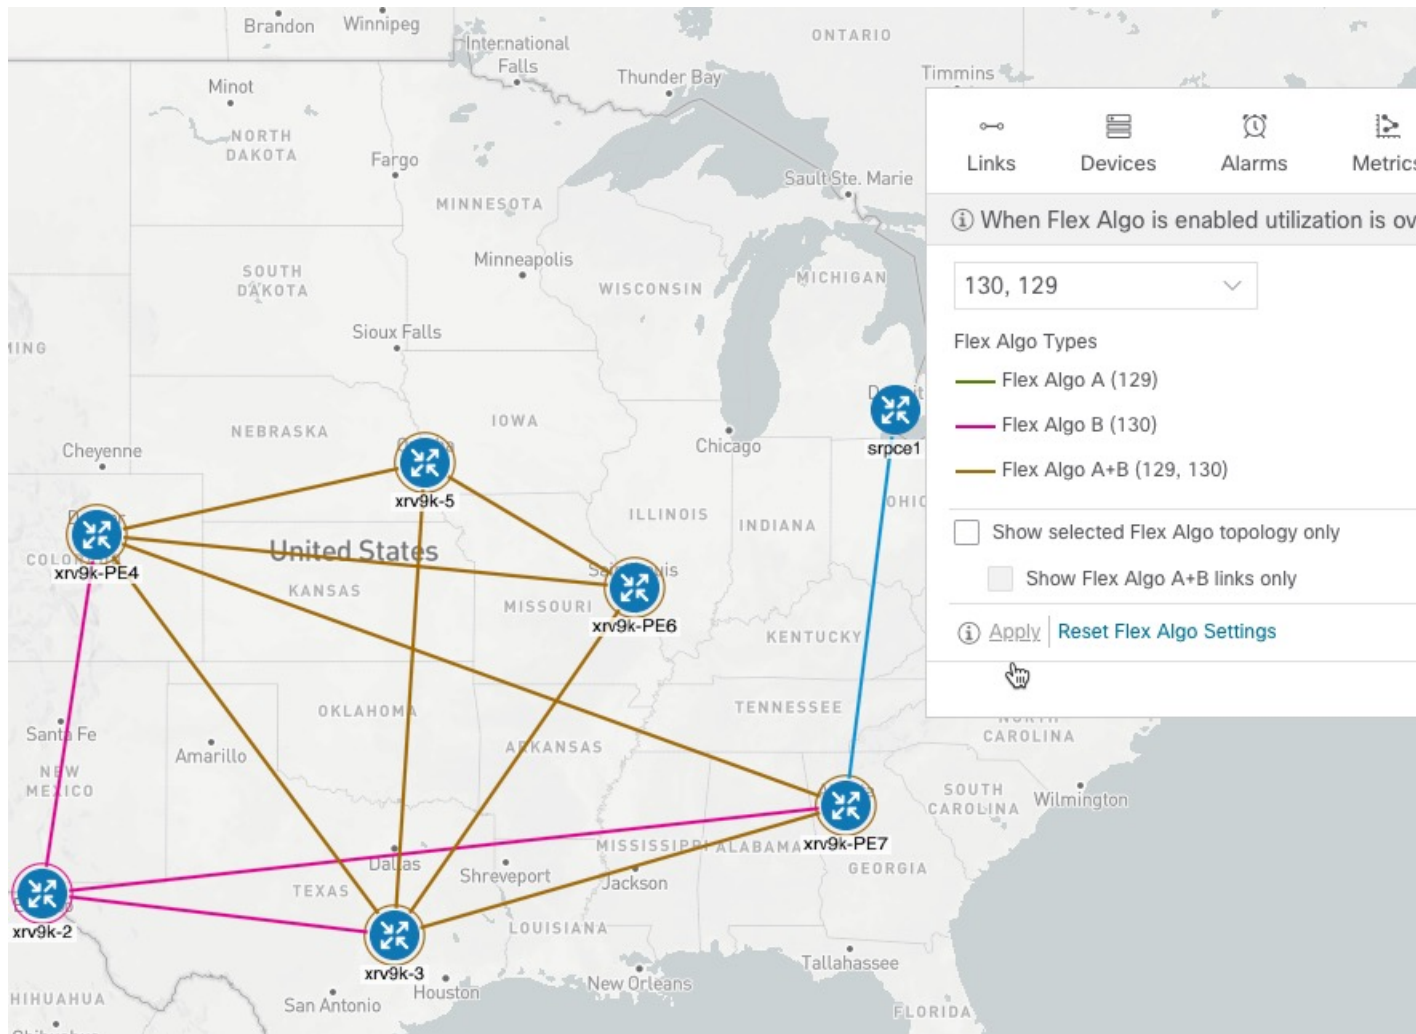
- You must understand and configure Flexible Algorithms in your network. See the SR Flexible Algorithm configuration documentation for your specific device to view descriptions and supported configuration commands (for example: [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#)).
- 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 56.



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:

Device Details ×

< Alarms SR-MPLS SRv6 Tree-SID RSVP-TE Flex-Algo >

∨ IGP: Domain ID: 1001, ISIS System ID: 0000.0000.0005, Level: 2 Expand All

∨ Algo 128

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

∨ Algo 129

Participating Yes

Elected Definition Metric Type: IGP

Exclude-Any Affinity:

Include-Any Affinity:

Include-All Affinity:

Advertised Yes

Priority: 229

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 6

Visualize Tree-SID Policies

Crosswork Optimization Engine lets you visualize the Tree-SID policies implemented in your network. This provides the ability to view details of the Tree-SID root, transit and leaf nodes, bud nodes and allows you to easily confirm that Tree-SID is implemented correctly in your network. The P2MP SR policy also prevents transient loop and packet loss when updating the path of a P2MP SR policy.

The Root node encapsulates the multicast traffic, replicates it, and forwards it to the transit nodes. Transit nodes replicate the multicast traffic and forward it to the Leaf nodes. The Bud node, is a node that acts as a leaf (egress) node as well as a mid-point (transit) node toward the downstream sub-tree. Leaf nodes decapsulate the multicast traffic and forward it to the multicast receivers.

To configure Tree-SID in your network, see the SR Tree-SID configuration documentation for your specific device (for example: [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#)).

COE supports visualizing the following types of Tree-SID policies:

- **Static:** A Static Tree-SID policy is configured via the PCE. All the paths are explicitly mentioned in static Tree-SID policy. The Tree-SID name is assigned during the configuration and it does not have an ID.
- **Dynamic:** A Dynamic Tree-SID policy is configured on service end-points, and require a day 0 configuration on PCE and the SEPs.



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



Note When using Crosswork Optimization Engine to visualize Tree-SID policies, always choose **Traffic Engineering > Traffic Engineering**. If you are using Crosswork Network Controller solution to visualize these policies, the navigation path is **Traffic Engineering & Services > Traffic Engineering**.

- [View a Point-to-Multipoint Tree on the Topology Map, on page 60](#)
- [Limitations for Tree-SID Policies, on page 61](#)
- [Tree SID Configuration Example, on page 64](#)

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 network 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 two leaves. Also, bud nodes have a separate leaf node path and are displayed on the Topology map.

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

The screenshot shows the Cisco Crosswork Network Automation interface. The main panel displays a network topology map with nodes and links. The right-hand panel shows details for the selected Tree-SID path, including a table of nodes and their roles.

Leaf Node Name	Leaf Node IP	Collapse All		
<input checked="" type="checkbox"/> xrv9k-VM7_3_0_732_cco	192.168.4.7			
Role	Name	IP	Egress Link	Remote IP
Root	xrv9k-VM5-...	192.168.4.5	10.0.2.26	10.0.2.25
Bud	xrv9k-VM3-...	192.168.4.3	10.0.2.41	10.0.2.42
Leaf	xrv9k-VM7...	192.168.4.7	-	-
<input checked="" type="checkbox"/> xrv9k-VM3-771-151	192.168.4.3			
Role	Name	IP	Egress Link	Remote IP
Root	xrv9k-VM5-...	192.168.4.5	10.0.2.26	10.0.2.25
Leaf	xrv9k-VM3-...	192.168.4.3	-	-
<input checked="" type="checkbox"/> xrv9k-VM8	192.168.4.9			

Before you begin

The following configurations are required for the Tree-SID policy and nodes:

- Transit node: PCEP is required.
- Bud node, Egress node, and Ingress node: PCEP, active BGP MVPN session, BGP autodiscovery segment-routing and MDT default segment-routing, MDT partitioned segment-routing

To visualize a multicast tree in the network 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 29](#)

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

Summary

- Admin State Up
- Oper Status Up
- Label 18
- Type Static
- Programming State None
- Metric Type TE
- Constraints Exclude-Any: -
- Include-Any: -
- Include-All: -
- SR-PCE Address 172.23.209.75

[See more](#)

Tree-SID path

Leaf Node Name	Leaf Node IP			
xrv9k-VM9-732	192.168.4.10		Collapse All	
	Name	IP	Egress Link	Remote IP
Root	xrv9k-VM2_751_151	192.168.4.2	10.0.2.18	10.0.2.17
Bud	xrv9k-VM7_3_0_732_cco	192.168.4.7	10.0.2.90	10.0.2.91
Leaf	xrv9k-VM9-732	192.168.4.10	-	-

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.

Limitations for Tree-SID Policies

Limitation

- Only visualization of Tree-SID policies is supported. You cannot create, edit or delete Tree-SID policies from the UI.
- Tree-SID policies are only supported on devices running Cisco IOS XR software.
- Tree-SID policies are not deleted from the UI when the PCE in HA mode is down and also when the PCE is removed from the Crosswork UI.
- PCE HA is not supported.
- Tree-SID policies are not supported in Label Switch Multicast (LSM) routing. In cases where LSM is enabled, IGP updates and traffic utilization data are not supported.
- Tree-SID policy is not created back in PCE, after "no vrf" under pim or "no multicast" on PCC.
- Ignore the local-hop-address if FRR=true, show value of 'next-hop-address'.
- FIB platform update fails on PE routers, when IGP configuration is updated.
- LCM will not operate in portions of the network carrying Tree-SID LSPs.
- The RestConf API is not supported.
- Tree-SID policy details do not show IPv6 router ID or Srv6 core information.

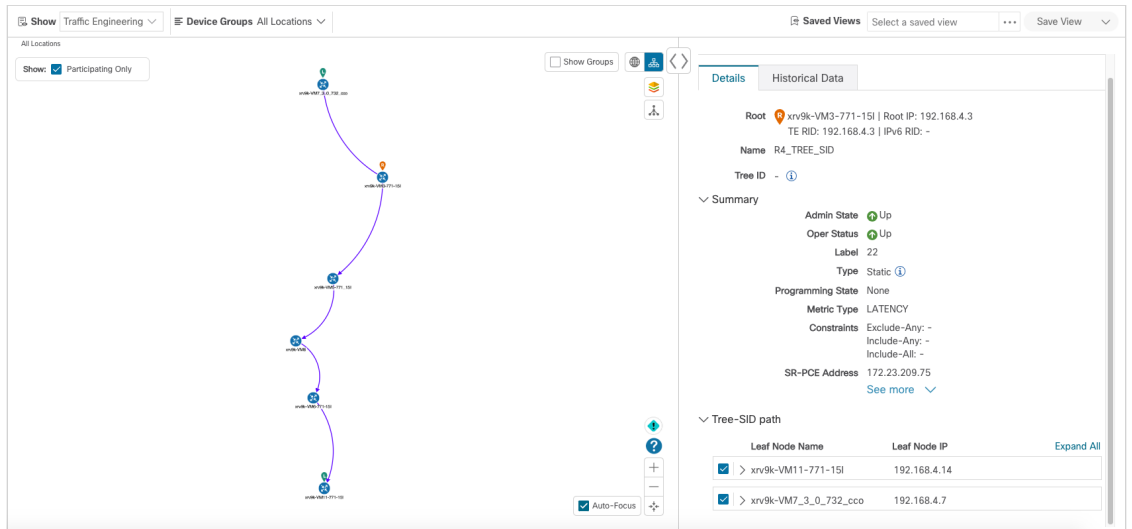
Visualization of Tree-SID Paths with Missing Nodes

Following are the scenarios with missing Tree-SID nodes on the topology:

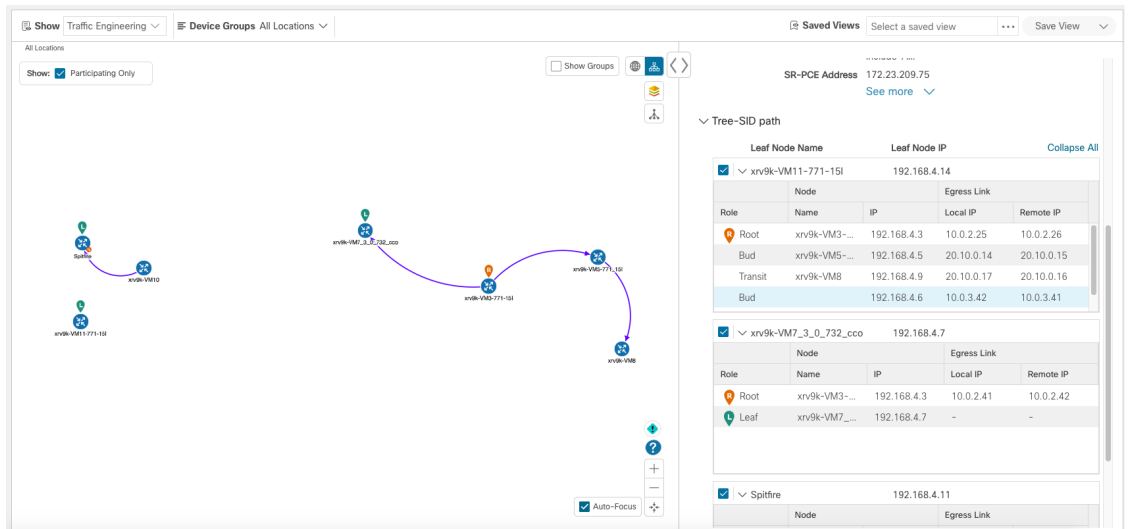
- You cannot visualise Tree-SID policy incase the source or root nodes are not configured on the PCE. The details for such Tree-SID policies are not populated, and the policy will be oper-down.

Root...	Root...	Name	Tree...	Label	Ad...	Op...	Actions
<input type="checkbox"/>	xrv9K...	192.1...	MY_F...	-	18	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	MY_S...	-	19	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	MY_S...	-	31	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	MY_S...	-	30	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	MY_T...	-	17	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	NEW...	-	28	✔	✔ ...
<input type="checkbox"/>	-	-	R13_...	-	32	✔	✘ ...
<input type="checkbox"/>	Source address is not configured/available for this Tree-SID policy		-	-	22	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	TREE...	-	37	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	TREE...	-	40	✔	✔ ...
<input type="checkbox"/>	xrv9K...	192.1...	TREE...	-	34	✔	✔ ...

- You cannot visualize Tree-SID policy path with specific leaf nodes missing. The details for such Tree-SID policies are populated with leaf node path missing. The other Tree-SID paths in the policy are displayed in the network, if any.



- You can visualize the Tree-SID policy with missing transit or bud node with partial overlay. The details for the Tree-SID policies are populated without the hostname.



Note The below warnings are displayed in case of missing transit or bud nodes:

1. Historical data tab: Some device links are missing from the policy path because they are not present in the current device inventory and topology database.
2. Topology UI: Topology map reflects the current state of the network. Current device/link states do not necessarily impact Traffic Engineering policies , tunnels, or services.

- If the root or source node is removed from the UI, the root hostname will be empty and the Tree-SID policy will be oper-up with no paths available on the topology map. The root router IP is displayed from the earlier Tree-SID discovery.

Tree SID Configuration Example

To visualize Segment Routing Tree-SID, some configurations are required on the SR-PCE and on the devices involved in the Tree-SID paths. Following are some example configurations for each of the steps required in your network:

- [Static Tree-SID Policy Configuration Example, on page 64](#)
- [Dynamic Tree-SID Policy Configuration Example with VRF, on page 65](#)
- [Dynamic Tree-SID Policy Configuration Example without VRF, on page 70](#)

The following day 0 configuration is required:

Enabling the MVPN address family on all SEPs and on PCE.

Enabling p2mp on PCE.

Static Tree-SID Policy Configuration Example



Note See Tree-SID 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 steps show examples of Static Tree-SID device configuration:

Step 1 Configure the Path Computation Element Protocol (PCEP) and Path Computation Client (PCC) on all nodes involved in the Tree-SID path (root, transit/bud, and leaf).

Example:

```
pce
address ipv4 <pce-loopback0-IP>

api
user admin
password encrypted xxxx
!
!

segment-routing
traffic-eng
```

```

p2mp

  endpoint-set MY_FIRST_TREE_SID_EPS

  ipv4 <leaf or pcc1-loopback0-IP>

  ipv4 <leaf or pcc2-loopback0-IP>

  !

```

Step 2 Configure P2MP SR static policy on the SR-PCE with end-points.

Example:

```

policy MY_FIRST_TREE_SID

  source ipv4 <root or pcc3-loopback0-IP>

  color 20 endpoint-set MY_FIRST_TREE_SID_EPS

  treesid mpls 18

  candidate-paths

    preference 100

    dynamic

    metric

    type te

  !

  !

  !

  !

  !

```

Dynamic Tree-SID Policy Configuration Example with VRF

To add more dynamic policies to Tree-SID policy, create VRF on both root and leaf devices. Mention the corresponding VRF, neighbor under BGP router config on PCE, root and leaf devices. VRF under multicast routing, router pim and create route-policy for each different VRFs as mentioned below in examples.

Follow the steps for Dynamic Tree-SID policy:

Pre-req route-policies (configure on both PCE, Root and leaf) devices

Under PCE

```

route-policy PASS

  pass

end-policy

```

```

!
Under Root and Leaf
route-policy bgp_in
    pass
end-policy
!
route-policy PIM-RPF
    set core-tree sr-p2mp
end-policy
!
route-policy bgp_out
    pass
end-policy
!
route-policy PASS_ALL
    pass
end-policy
!
route-policy TREESID-CORE
    set core-tree sr-p2mp
end-policy
!

```

Step 1 Configure under pce → segment-routing traffic engineering -> p2mp-> label range <>, multi-path disable.

Example:

```

label-range min 15400 max 60000
    fast-reroute lfa
    multipath-disable

```

Step 2 Under router bgp - configure address family ipv4 mvpn at top level and under neighbor node IP <root> and <leaf> level as well with address family ipv4 mvpn.

Example:

```

router bgp 1

```

```

.....

```



```

        address-family ipv4 mvpn
        route-reflector-client
    !
        neighbor <root or pcc3-loopback0-IP>
        remote-as 1
        update-source Loopback0
        address-family ipv4 unicast
        route-policy PASS in
        route-policy PASS out
    !
        address-family ipv4 mvpn
    !
!
neighbor <leaf or pcc1-loopback0-IP>
    remote-as 1
    update-source Loopback0
    address-family ipv4 unicast
    route-policy PASS in
    route-policy PASS out
!
    address-family ipv4 mvpn
!
!

```

Step 3 Configure headend and end-points.

Note You can add end-point routers as neighbor under PCE → route BGP configuration. Router-IDs need to be updated with each PCC loopbackIP in topology.

- a) Create interface Loopback<80>

Example:

```

interface Loopback80
    ipv4 address 80.80.10.1 255.255.255.252
    ipv6 address 2001:192:168:80::1/128
!

```

- b) Create VRF <vrf-name-80>

Example:

```
vrf L3VPN_NM-MVPN-80
  address-family ipv4 unicast
    import route-target
      80:80
    !
  export route-target
    80:80
  !
  !
  !
```

- c) Mention <vrf-name-80> under routing BGP configuration

Example:

```
vrf L3VPN_NM-MVPN-80
  rd 80:80
  address-family ipv4 unicast
    redistribute connected
  !
  address-family ipv4 mvpn
  !
  neighbor 80.80.10.1 <leaf or pcc1-vrf-IP>
    remote-as 80
  address-family ipv4 unicast
    route-policy PASS_ALL in
    route-policy PASS_ALL out
  !
  !
  !
  !
```

- d) Mention <vrf-name-80> under multicast-routing configuration

Example:

```
multicast-routing
  address-family ipv4
    interface Loopback0
      enable
    !
    mdt source Loopback0
    mdt static segment-routing
  !
vrf L3VPN_NM-MVPN-80
  address-family ipv4
    interface all enable
    bgp auto-discovery segment-routing
  !
  mdt default segment-routing mpls color 80
  !
  !
  !
```

e) Create route-policy <vrf-name-80>

Example:

```
route-policy L3VPN_NM-MVPN-80
  if destination in (232.0.0.80) then
    set on-demand-color 80
  pass
endif
end-policy
!
```

f) Under segment routing traffic eng -> configure ODN color <80>

Example:

```
on-demand color 80
  dynamic
  pcep
  !
  metric
```

```

        type te
    !
    !
    !

```

Step 4 Configure Leaf

Note Follow step a to d to configure headend and endpoint for root nodes.

Example:

```

router pim

address-family ipv4

    rpf topology route-policy PIM-RPF

!

vrf L3VPN_NM-MVPN-80

address-family ipv4

    rpf topology route-policy TREESID-CORE

    mdt c-multicast-routing bgp

!

!

!

```

Dynamic Tree-SID Policy Configuration Example without VRF

To add dynamic policies to Tree-SID policy, without VRF on both root and leaf devices, follow the steps below:



Note The configuration for PCE are same as for Dynamic Tree-SID with VRF. See [Dynamic Tree-SID Policy Configuration Example with VRF, on page 65](#)

Step 1 Configure Root

- a) Mention <leaf-node-IP or pcc1-IP> as neighbor under router BGP configuration.
- b) Mention unique RTs under multicast-routing configuration.

Note The RTs should be unique between Root and Leaf set.

Example:

```

multicast-routing

    address-family ipv4

        import-rt 12:12

        export-rt 12:12

        mdt source Loopback0

        interface all enable

        bgp auto-discovery segment-routing

    !

    mdt default segment-routing mpls color 12 fast-reroute lfa

    mdt data segment-routing mpls 5 threshold 0

    !

```

- c) Under segment routing traffic eng -> configure ODN color <unique one>.

Step 2

Configure Leaf

- a) Mention <root-node-IP or pcc3-IP> as neighbor under router BGP configuration.
- b) Mention unique RTs under multicast-routing configuration.

Note The RTs should be unique between Root and Leaf set.

Example:

```

multicast-routing

    address-family ipv4

        import-rt 12:12

        export-rt 12:12

        mdt source Loopback0

        interface all enable

        bgp auto-discovery segment-routing

    !

    mdt default segment-routing mpls color 12 fast-reroute lfa

    mdt data segment-routing mpls 5 threshold 0

    !

```

- c) Configure router PIM, route-policy TREESID_CORE.



CHAPTER 7

Visualize RSVP-TE Tunnels



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

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

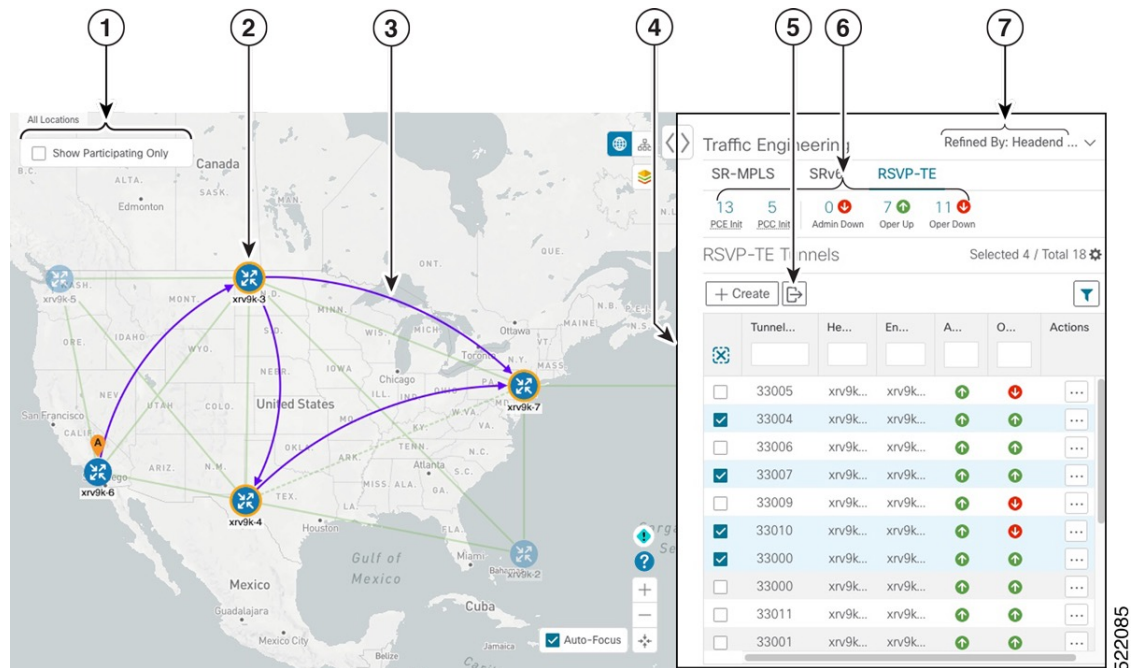
This section contains the following topics:


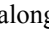
- [View RSVP-TE Tunnels on the Topology Map, on page 73](#)
- [View RSVP-TE Tunnel Details, on page 75](#)
- [View Traffic Engineering Device Details, on page 77](#)

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 12: Traffic Engineering UI - RSVP-TE Tunnels



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	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. Note RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.
3	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. <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. Note If both RRO and ERO paths are available, the RRO path is displayed by default. <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>

Callout No.	Description
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 88 • Create Explicit RSVP-TE Tunnels, on page 86 • Modify RSVP-TE Tunnels, on page 89 • View RSVP-TE Tunnel Details, on page 75 • View Device and Link Details, on page 12
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.
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.

View RSVP-TE Tunnel Details

The screenshot displays the Traffic Engineering interface. On the left, a map of North America shows several RSVP-TE tunnels (represented by blue icons with 'x') and their paths. A purple tunnel is highlighted. On the right, the 'RSVP-TE Tunnels' table is shown, listing 15 tunnels. The tunnel with ID 33000 is selected, and a 'View Details' button is visible over its row.

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

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

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 **Device Details** page, click on the traffic engineering tab 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.

View Traffic Engineering Device Details

All Locations

Show Groups
 Auto-Focus

Device Details

[Details](#) | [Links](#) | [Alarms](#) | **SR-MPLS** | [SRv6](#) | [Tree-SID](#) | [RSVP-TE](#)

[Policies](#) | **Prefixes** [Expand All](#)

∨ IGP: Domain ID: 1000, ISIS System ID: 0000.0000.0005, Level: 2

SRGB 16000 - 23999
SRLB 105000 - 105999
MSD 10

Prefixes	Label	Algo
192.168.0.5	18115	0



CHAPTER 8

Provision SR-MPLS Policies



- Note** When using Crosswork Optimization Engine within the Crosswork Network Controller solution:
- The navigation is **Traffic Engineering & Services > Traffic Engineering > Traffic Engineering**.
 - SRv6 policies can only be provisioned through the NSO menu.

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

This section contains the following topics:

- [SR-TE Policy Configuration Sources, on page 79](#)
- [Create Explicit SR-MPLS Policies, on page 80](#)
- [Configure Link Affinities, on page 81](#)
- [Create Dynamic SR-MPLS Policies Based on Optimization Intent, on page 82](#)
- [Modify SR-MPLS Policies, on page 83](#)

SR-TE Policy Configuration Sources

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

- PCC initiated—Policies configured on a PCC (see [PCC-Initiated SR-TE Policy Example, on page 80](#)). This policy type displays as **Unknown** in the UI.
- PCE initiated—Policies configured on a PCE or created dynamically by Crosswork Optimization Engine. SR-MPLS explicit or dynamic policies that are configured using the UI are the only types of SR-TE policies that you can modify or delete in Crosswork Optimization Engine. PCE Initiated policy types can be one of the following:
 - **Dynamic**
 - **Explicit**
 - **Bandwidth on Demand**
 - **Bandwidth Optimization**

- Local Congestion Mitigation


PCC-Initiated SR-TE Policy Example

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

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


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.

-
- Step 1** From the main menu, choose **Traffic Engineering > Traffic Engineering > SR-MPLS** tab.
- Step 2** From the **SR Policies** table, click **+ Create**.
- 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. You can continue editing, incase the preview does not appear or click **Cancel**.
- 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 scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see the [Cisco Crosswork Infrastructure and Applications Administration Guide](#).

Configure Link Affinities

Affinity names defined on devices are not collected by Crosswork Optimization Engine. The affinity mapping is only used for visualization in Crosswork Optimization Engine. For this reason, you should collect affinities on the device, then define affinity mapping in Crosswork Optimization Engine 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".

Affinity of an SR-TE policy or RSVP-TE tunnel is used to specify the link attributes for which the SR-TE policy or RSVP-TE tunnel has affinity for. It determines which links are suitable to form a path for the SR-TE policy or RSVP-TE tunnel. It is a 32-bit value, with each bit position (0 - 31) representing a link attribute. Affinity mapping is used to map each bit position or attribute to a color. This makes it easier to refer to link attributes.

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 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 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:

TE Link Affinities
Flex-Algo Affinities

+ Create
⌵

Name ?	Bit Position (0-31) ?	Actions
<input style="width: 100%;" type="text"/>	<input style="width: 100%;" 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 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 Link Affinities, on page 81](#) or [Configure Flexible Algorithm Affinities, on page 53](#).

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

Step 2 From the **SR Policy** table, click **+ Create**.

Step 3 Under **Policy Details**, enter or select the required SR-MPLS policy values. Hover the mouse pointer over ? to view a description of each field.

Tip If you have set up device groups, you can select the device group from the **Device Groups** drop-down menu. Then navigate and zoom in on the topology map to click the device for headend or endpoint selection.

Step 4 Under **Policy Path**, click **Dynamic Path** and enter a path name.

Step 5 Under **Optimization Objective**, select the metric you want to minimize.

Step 6 Define any applicable constraints and disjointness.

- Note**
- Affinity constraints and disjointness cannot be configured on the same SR-MPLS policy. Also, there cannot be more than two SR-MPLS policies in the same disjoint group or subgroup. The configuration will not be allowed during Preview.
 - If there are existing SR-MPLS policies belonging to a disjoint group that you define here, all SR-MPLS policies that belong to that same disjoint group are shown during Preview.

Step 7 Under **Segments**, select whether or not public 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 the [Cisco Crosswork Infrastructure and Applications Administration Guide](#).

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 9

Provision RSVP-TE Tunnels



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

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

This section contains the following topics:

- [RSVP-TE Tunnel Configuration Sources, on page 85](#)
- [Create Explicit RSVP-TE Tunnels, on page 86](#)
- [Configure Link Affinities, on page 86](#)
- [Create Dynamic RSVP-TE Tunnels Based on Optimization Intent, on page 88](#)
- [Modify RSVP-TE Tunnels, on page 89](#)

RSVP-TE Tunnel Configuration Sources

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

- PCC initiated—RSVP-TE tunnels configured on a PCC (see [PCC-Initiated RSVP-TE Tunnel Example, on page 85](#)).
- 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
!
```

Create Explicit RSVP-TE Tunnels

This task creates RSVP-TE tunnels using an explicit (fixed) path consisting of a list of prefix 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:
- 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.
 - 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.
-

Configure Link Affinities

Affinity names defined on devices are not collected by Crosswork Optimization Engine. The affinity mapping is only used for visualization in Crosswork Optimization Engine. For this reason, you should collect affinities on the device, then define affinity mapping in Crosswork Optimization Engine 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".

Affinity of an SR-TE policy or RSVP-TE tunnel is used to specify the link attributes for which the SR-TE policy or RSVP-TE tunnel has affinity for. It determines which links are suitable to form a path for the SR-TE policy or RSVP-TE tunnel. It is a 32-bit value, with each bit position (0 - 31) representing a link attribute. Affinity mapping is used to map each bit position or attribute to a color. This makes it easier to refer to link attributes.

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 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 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 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 Link Affinities, on page 81](#).

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 10

Use Local Congestion Mitigation (LCM) to Mitigate Network Congestion Locally



Note

- Functionality described within this section is only available as part of the Advanced RTM license package.
- Throughout this section, the navigation is documented as **Traffic Engineering > Traffic Engineering**. However, when using Crosswork Optimization Engine within the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering**.

- [Local Congestion Mitigation Overview, on page 91](#)
- [LCM Important Notes, on page 92](#)
- [LCM Calculation Workflow, on page 94](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 96](#)
- [Configure LCM, on page 105](#)
- [Add Individual Interface Thresholds, on page 107](#)
- [Monitor LCM Operations, on page 109](#)

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 105](#).

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 96](#) 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.

- 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 105](#)) 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

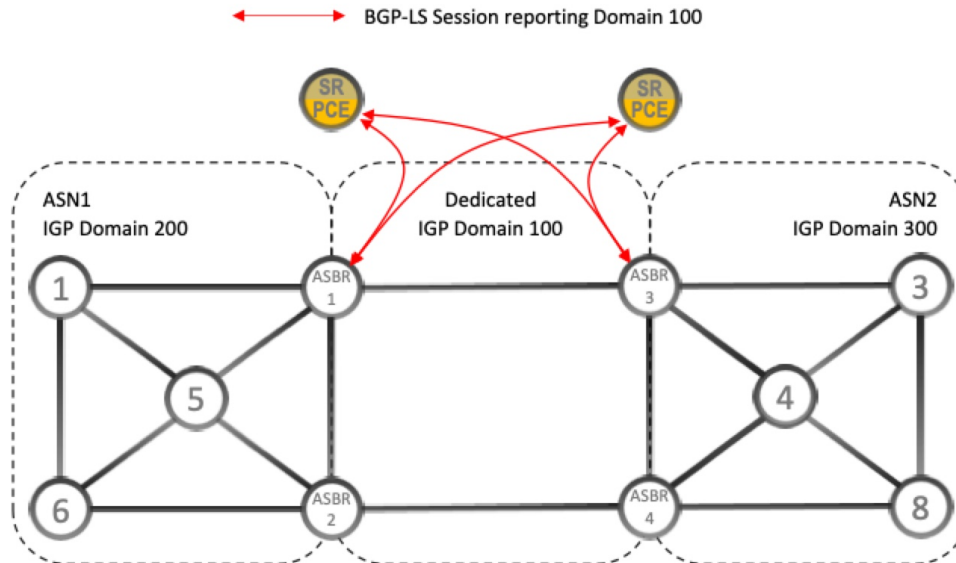
To support interdomain latency-optimized SR policy path computation by an SR-PCE (or other use cases where egress peer engineering (EPE) is not supported), a dedicated IGP instance may be configured between autonomous system border routers (ASBRs) in different ASNs. In these cases, it is important to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

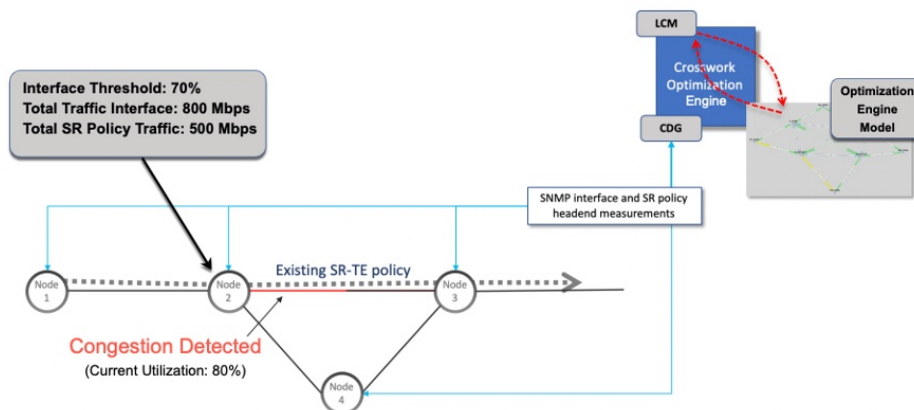
Figure 13: BGP-LS Session Reporting Domain 100



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 14: 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

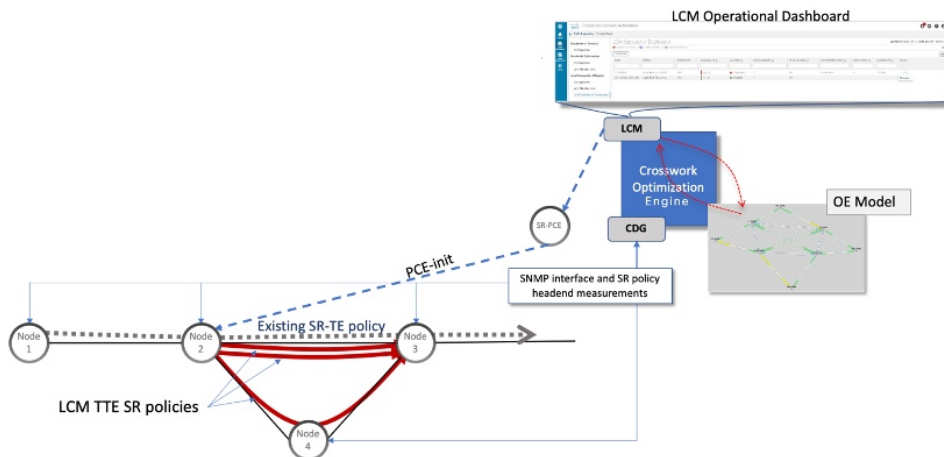
LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see [Configure LCM, on page 105](#).

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.

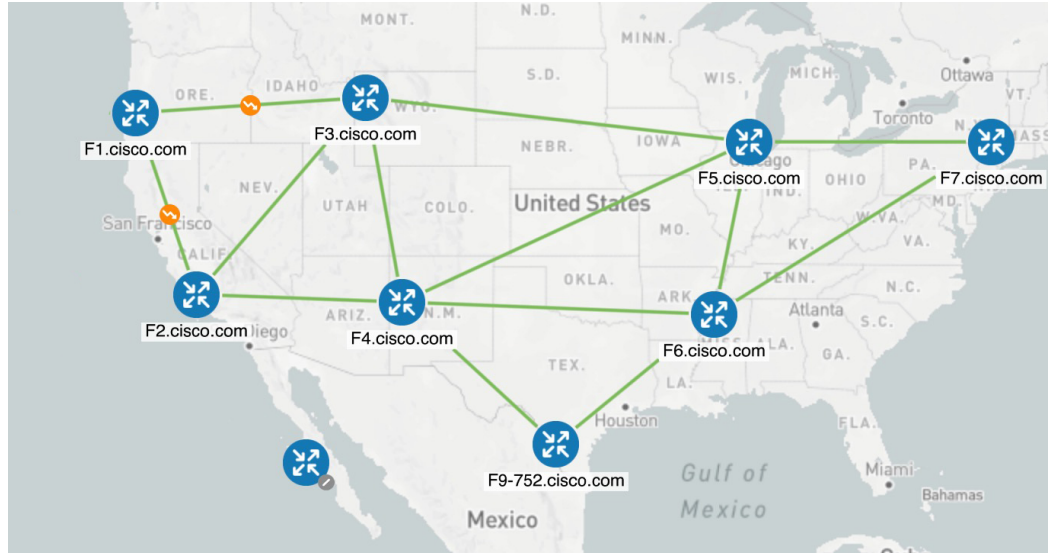
In this example, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE SR policies when utilization on a device's interface surpasses a defined utilization threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion.

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 15: 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 16: Initial Utilization

The screenshot shows the 'Link Details' window for the link 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 section 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 107](#).

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 17: 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 105](#).

Figure 18: LCM Configuration Page

Configuration


Basic Advanced

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

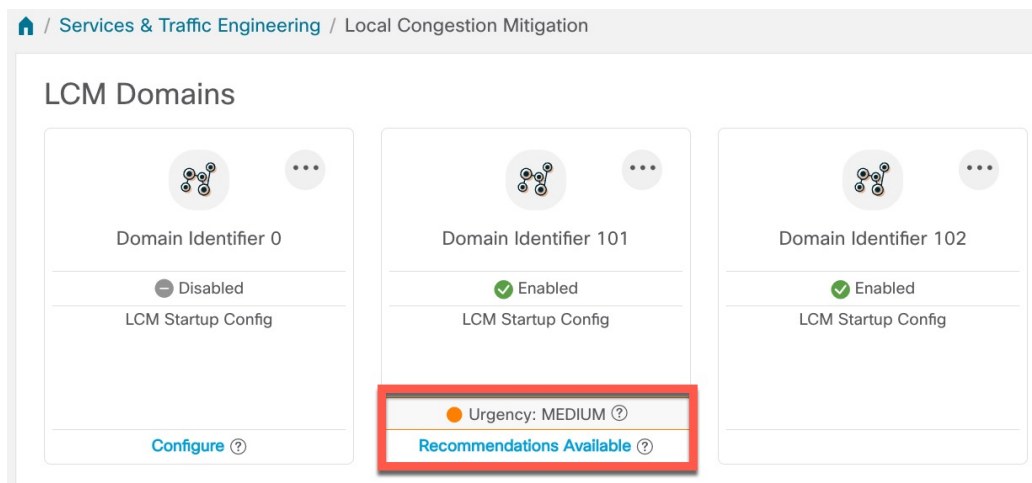
Figure 19: Observed Congestion




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

- 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 20: Congested Detected and LCM Recommendations



- (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 109](#).

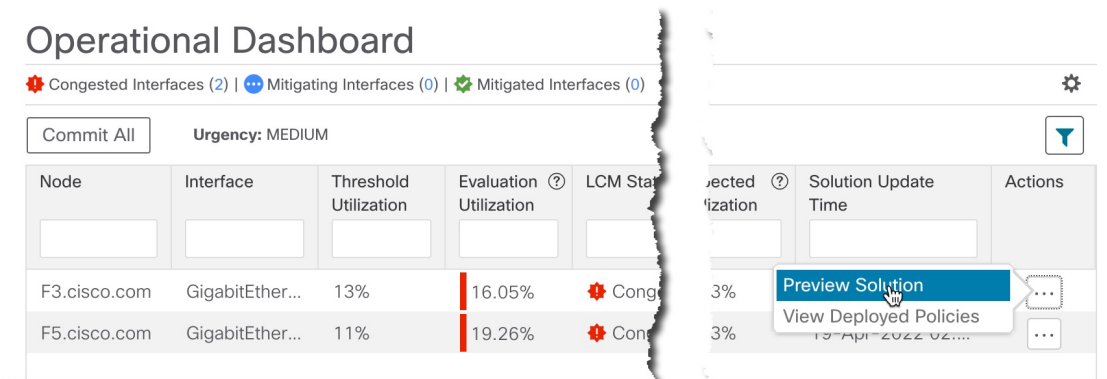
Figure 21: LCM Operational Dashboard

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 22: 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 23: 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 24: 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.

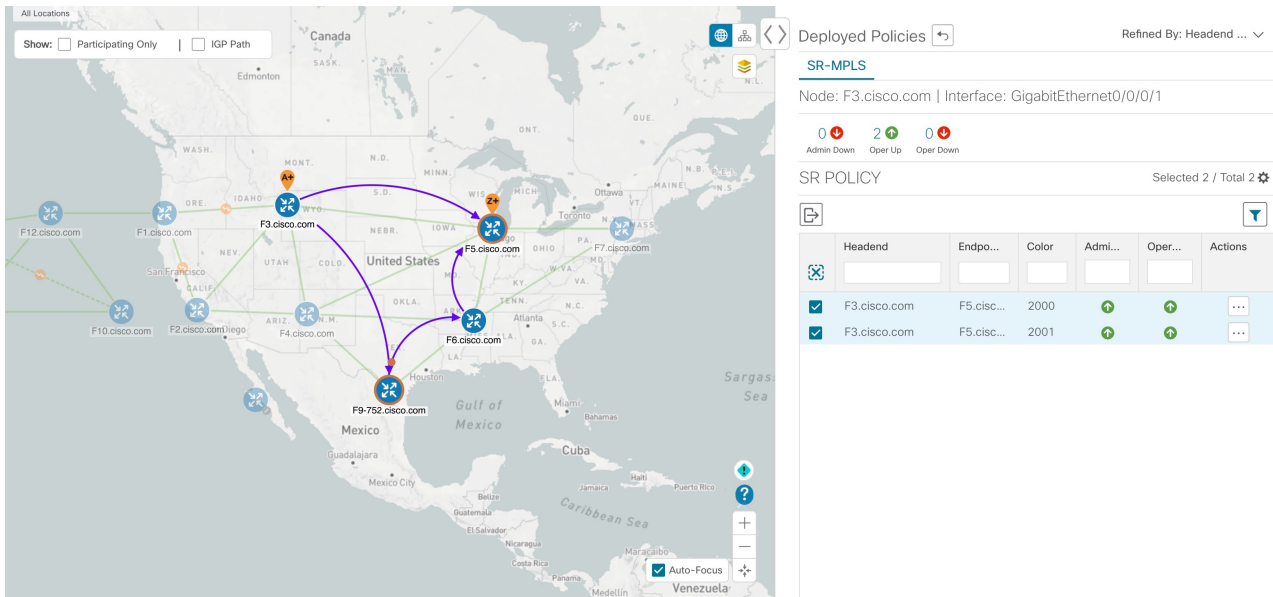
- b) Return to the **Operational Dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

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 25: View TTE Deployment Policies on Topology Map



The screenshot shows a network topology map of the United States and Mexico. A panel on the right displays the details for the SR-MPLS policy. The panel includes a table of deployed policies:

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.

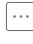
From the **Actions** column of one of the deployed policies click  and choose **View Details**. Note that the **Policy Type** is **Local Congestion Mitigation**.

Figure 26: 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.

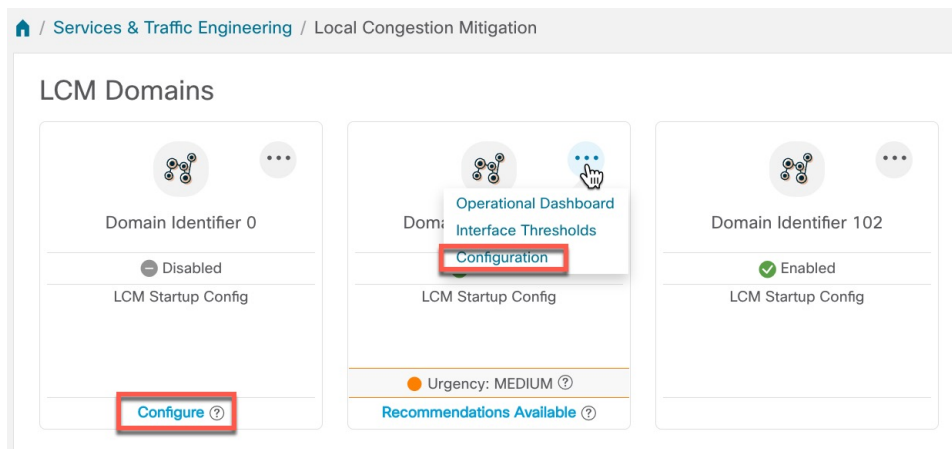
In this scenario we observed how to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands but at the same time gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations 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 94. 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**. For more information, see [View Traffic Engineering Device Details, on page 32](#).

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 *Crosswork Infrastructure and Applications 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 27: Customized Interface Thresholds

Interface	Threshold (%)	Select to Delete	
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"/>

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 28: Add First Interface



The screenshot shows the 'Customized Interface Thresholds' page. At the top, there are icons for adding, deleting, and refreshing, along with an 'Edit Mode: OFF' toggle. Below this is a table with four columns: 'Node', 'Interface', 'Threshold (%)', and 'Select for Deletion'. The first row of the table is highlighted with a red border, and a red box highlights the input fields for 'Node', 'Interface', and 'Threshold (%)'. The 'Select for Deletion' column has a trash icon. A filter icon is visible in the top right corner.

- 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 96](#) 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 29: LCM Domains Dashboard



Callout No.	Description
1	<p>Main Menu: Allows you to navigate to the following pages:</p> <ul style="list-style-type: none"> Operational Dashboard Add Individual Interface Thresholds Configure LCM
2	<p>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.</p>

Callout No.	Description
3	LCM Status: Indicates whether LCM had been enabled for the domain.
4	LCM Configuration Description: The description is defined in the Configure LCM page. The default description is "LCM Startup Config".
5	<p>Urgency: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following:</p> <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 105.
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 30: LCM Operational Dashboard

Operational Dashboard

2 Congested Interfaces (2) | 0 Mitigating Interfaces (0) | 1 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 105](#).

- 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 11

Use Bandwidth Optimization (BWOpt) to Optimize the Network



Note

- Functionality described within this section is only available as part of the Advanced RTM license package.
- Throughout this section, the navigation is documented as **Traffic Engineering > Traffic Engineering**. However, when using Crosswork Optimization Engine within the Crosswork Network Controller solution, the navigation is **Traffic Engineering & Services > Traffic Engineering**.

- [Bandwidth Optimization Overview, on page 113](#)
- [BWOpt Important Notes, on page 113](#)
- [Automated Network Congestion Mitigation Example, on page 115](#)
- [Configure Bandwidth Optimization, on page 118](#)
- [Add Individual Interface Thresholds, on page 118](#)
- [Troubleshoot Bandwidth Optimization, on page 119](#)

Bandwidth Optimization Overview

Bandwidth Optimization (BWOpt) provides closed-loop tactical traffic engineering (TTE) for segment routed policies by *automatically* detecting and mitigating congestion in your network. It achieves this through a real-time view of the network topology overlaid with a demand matrix built through telemetry-based Segment Routing Traffic Matrix (SRTM). The intent is to optimize bandwidth resource utilization by setting utilization thresholds on links. BWOpt uses the threshold interface utilization requested by the user and compares it to the actual utilization in the network. When interface congestion is detected by BWOpt, it attempts to reroute intent-based traffic from hot spots through the use of TTE SR policies which are deployed to the network via SR-PCE. As network conditions (topology and/or traffic) change over time, BWOpt continues to monitor interface utilization and manage any TTE SR policies deployed, including changing their paths and/or removing them from the network when deemed no longer necessary.

BWOpt Important Notes

Consider the following information when using BWOpt:

- You must have the Advanced RTM license package to use BWOpt.
- You cannot enable Bandwidth Optimization if LCM is enabled.
- BWOpt will not shift traffic in existing SR-TE policies that it did not create. This may prevent it from being able to mitigate congestion if most of the traffic on the congested link is in non-BWOpt SR-TE policies.
- BWOpt relies on the PCC's autoroute feature to steer traffic into the tactical SR-TE policies it creates. Autoroute is applied to these policies through the proper **Profile ID** option set in BWOpt (to align with configuration on the PCC associating that Profile ID with autoroute feature). This is critical to tactical SR policies shifting traffic away from congested links.
- Enable BWOpt on single-level IGP domains only.
- 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**.
- BWOpt leverages on the use of Segment Routing Traffic Matrix (SR-TM). SR-TM has the following limitations:
 - IPv6 is not supported.
 - Management, bundle, subinterfaces and tunnel interfaces are not supported as external interfaces.
 - Non-default Virtual, Routing, and Forwarding (VRF) is not supported as external interfaces.
 - SR-TM only takes SR labeled traffic into consideration. It will not account for Label Distribution Protocol (LDP) traffic.



Note For more information on SR-TM, see [Segment Routing Traffic Matrix](#).

- BWOpt uses simulated traffic based on measured SR-TM data to determine link utilizations and when to mitigate congestion. The simulated interface utilization that BWOpt monitors should closely align with the SNMP-based interface utilization that is displayed in the UI. However, due to various factors, including SNMP polling cadence and rate averaging techniques, they may differ at times. This can result in scenarios like a link appearing to be congested in the UI and BWOpt not reacting.
- BWOpt only creates tactical SR-TE policies on PCCs that are sources of SRTM telemetry data. Only these nodes (typically provider edge routers) provide the telemetry-based data needed to create simulated traffic demands in the internal model representing the traffic from that node to other PE nodes in the network.
- Only solutions that produce interface utilization below the threshold (set across all interfaces) will be deployed. If BWOpt is unable to mitigate congestion across the entire network, it will not deploy any tactical SR-TE policies and a “Network Congested. BWOpt unable to mitigate.” alarm is raised. This alarm goes away when congestion either subsides on its own or can be addressed successfully through BWOpt tactical SR-TE policy deployments.
- BWOpt temporarily pauses operation whenever the system is unavailable due to a restart or a rebuild of the topology from Topology Services. When this occurs, an alarm indicating this condition is set by BWOpt. During this time, BWOpt will not evaluate congestion in the network. All currently deployed tactical SR policies are maintained, but will not be modified or deleted. As soon as the model becomes available, the alarm is cleared and BWOpt will resume normal operation.

Automated Network Congestion Mitigation Example

This example demonstrates how Bandwidth Optimization (BWOpt) automatically mitigates network congestion by rerouting intent-based traffic without user intervention. In this example, the optimization intent is set to minimize the IGP metric.

The following BWOpt options are set (**Traffic Engineering > Bandwidth Optimization > Configuration**):

Figure 31: Bandwidth Optimization Configuration


The screenshot shows the configuration page for Bandwidth Optimization. The left sidebar has 'Bandwidth Optimization' selected, with sub-items 'Configuration' and 'Link Management'. The main area is titled 'Configuration' and has two tabs: 'Basic' (selected) and 'Advanced'. The configuration is organized into several sections:

- Enable:** A toggle switch is set to 'True'.
- Optimization Objective:** A dropdown menu is set to 'Minimize the IGP metric'.
- Color:** A text input field contains the value '1000'.
- Utilization Threshold:** A text input field contains the value '100'.
- Utilization Hold Margin:** A text input field contains the value '5'.
- Maximum Global Reoptimization Interval:** A text input field contains the value '0'.
- Profile ID:** A text input field contains the value '0'.
- Max Number of Parallel Tactical Policies:** A text input field contains the value '1'.

At the bottom of the configuration area, there are three buttons: 'Commit Changes' (highlighted in blue), 'Get Default Values', and 'Discard Changes'.

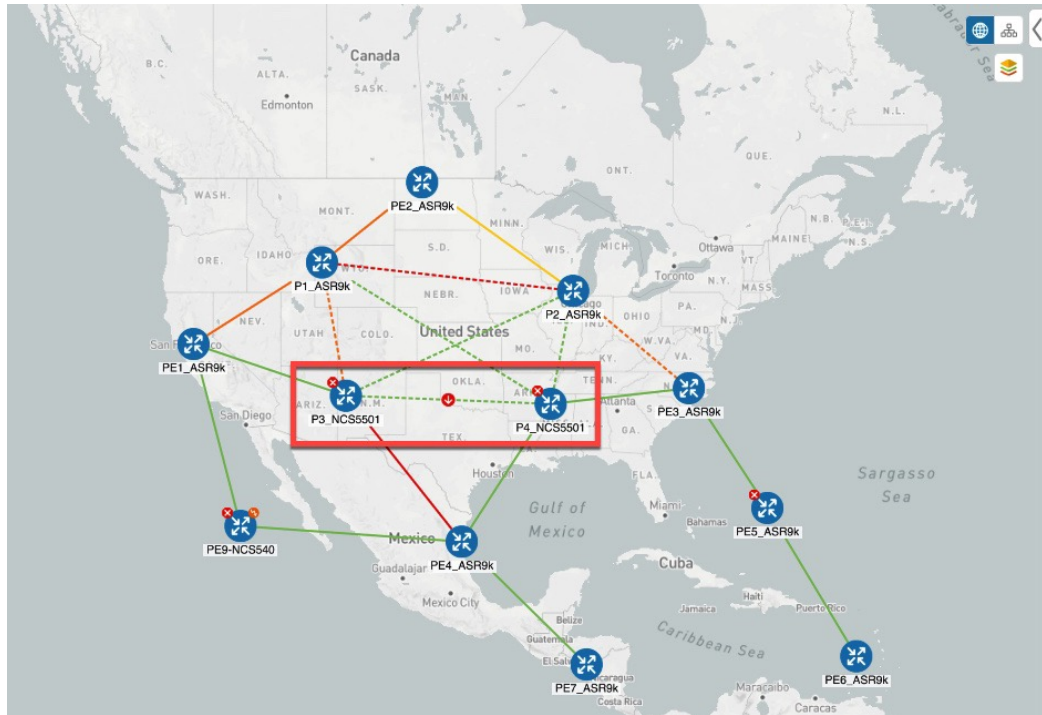
Below is a network with various devices and links that span the United States. Note that there are no SR-TE policies listed in the **SR Policies** table.

Figure 32: Example: Current Network



Suppose the link between P3_NCS5501 and P4_NCS5501 goes down. Traffic moves towards other links causing congestion and exceeds the configured utilization threshold.

Figure 33: Example: Link Down Between P3 and P4 Nodes



BWOpt recognizes the congestion and immediately calculates and deploys a tactical SR-TE policy. This new tactical SR-TE policy is listed in the **SR Policies** window.

Figure 34: Example: Tactical SR Policy Deployed

The diagram shows the same network topology as Figure 33, but with a tactical SR-TE policy deployed. A purple path is highlighted, showing traffic being rerouted from P3_NCS5501 to P2_ASR9k and then to P4_NCS5501. A tooltip for the link between P2_ASR9k and P4_NCS5501 shows the policy name: "bwopt_to_PE2_ASR9k_c_1000". On the right side, the "Traffic Engineering" window is open, showing the "SR POLICY" configuration.


SR-TE	RSVP-TE
1	0
0	0
0	1
0	0

PCE Init: 0, PCC Init: 0, Admin Down: 0, Oper Up: 1, Oper: 0

SR POLICY	Headend	Endpoint	Color
<input checked="" type="checkbox"/>	PE4_ASR9k	PE2_ASR9k	1000

BWOpt continually monitors the network. When the links between P3_NCS5501 and P4_NCS5501 are back up, BWOpt will detect that the congestion (based on the defined criteria) has been mitigated. When the

congestion falls under the set utilization threshold minus the utilization hold margin, the tactical SR-TE policy is automatically removed from the network.

You can also click  to view events relating to instantiation and removal of tactical SR-TE policies created by BWOpt.

Configure Bandwidth Optimization




Note Bandwidth Optimization (BWOpt) is only available as part of the Advance License package.

After BWOpt is enabled, it monitors all interfaces in the network for congestion based on the configured utilization threshold. When the utilization threshold is exceeded, it automatically deploys tactical polices and moves traffic away from the congested links. When congestion is alleviated, BWOpt automatically removes the tactical SR policy.

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

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

Note LCM and Bandwidth Optimization cannot be enabled at the same time.

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

Step 4 Click **Commit Changes**. BWOpt begins to monitor network congestion based on the threshold and optimization intent that was configured.

Add Individual Interface Thresholds

Networks have many different links (10G, 40G, 100G) that require different thresholds to be set. To assign specific threshold values for individual interfaces when using Bandwidth Optimization, do the following:

Step 1 From the main menu, choose **Traffic Engineering > Bandwidth Optimization > Interface Thresholds**

Step 2 Click **Import .CSV File**.

Step 3 Click the **Download sample configuration file** link.

Step 4 Click **Cancel**.

Step 5 Open and edit the configuration file (BWOptLinkManagementTemplate.csv) you just downloaded. Replace the sample text with your specific node, interface, and threshold information.

Step 6 Rename and save the file.

Step 7 Navigate back to the **Customized Interface Thresholds** window.

Step 8 Click **Import .CSV File** and navigate to the CSV file you just edited.

Step 9 Click **Import**.

Step 10 Confirm that the information appears correctly in the **Customized Interface Thresholds** window.

Troubleshoot Bandwidth Optimization

BWOpt disables itself and issues an alarm when specific error conditions occur that hinder its ability to manage congestion properly and may lead to instability. The following table defines some of these conditions and possible causes to investigate. Additional details can be obtained for each error condition by referring to the BWOpt logs.



Note You can navigate to **Administration > Collection Jobs** and then filter the list of active collection jobs for Optim in the App ID column.

Table 1: Errors

Error Event Message	Possible Causes and Recommended Corrective Action
Optima Engine model error	The network model used by BWOpt from the Optimization Engine is corrupt or is missing key data that is needed to properly support BWOpt. 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. This error can also occur if the time required to deploy a tactical policy through SR-PCE, discover it, and add it to the model exceeds the Deployment Timeout option set for BWOpt. The default is 30 seconds which should suffice for small to medium sized networks. However, larger networks may require additional time.
PCE Dispatch unreachable	The deployment of a tactical policy to the network is not confirmed successful before the Deployment Timeout is exceeded. Increase the Deployment Timeout option to allow for additional time for deployments in larger networks.
Unable to deploy a tactical SR policy	A tactical SR policy deployment to SR-PCE was unsuccessful. There could be a variety of reasons for this. BWOpt and/or PCE Dispatch logs can provide some guidance as to the details of the failure. Confirm basic SR policy provisioning capability to the PCC via one of the SR-PCE providers is working.



CHAPTER 12

Define and Maintain Intent-Based Bandwidth Requirements



Note Functionality described within this section is only available as part of the Advance RTM license package.

Bandwidth on Demand (BWoD) provides a bandwidth-aware Path Computation Element (PCE) to derive SR policy paths with requested bandwidth when available. Computed paths are deployed to the network through SR-PCE. BWoD continuously monitors link utilization to ensure no congestion occurs along the path. If conditions change in the network which causes link utilization to exceed the congestion threshold set by the user, BWoD automatically reoptimizes the policy path. BWoD supports bandwidth constraints for both PCE-initiated and PCC-initiated SR-TE policies.

BWoD utilizes a near real-time model of the network along with SNMP-based SR policy traffic measurements to ensure BWoD policies meet their bandwidth constraints. Users may fine tune the behavior of BWoD, affecting the path it computes, through the selection of application options including network utilization threshold (definition of congestion) and path optimization intent. BWoD works as a bandwidth-aware PCE for SR policies created through the UI and for SR policies created through CLI configuration on a headend with delegation to the SR-PCE. In the latter case, SR-PCE will subdelegate the SR policy with a bandwidth constraint to BWoD for path computation and relay the computed path returned by BWoD to the headend for instantiation.

- [BWoD Important Notes, on page 121](#)
- [Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example, on page 122](#)
- [Configure Bandwidth on Demand, on page 125](#)
- [Troubleshoot BWoD, on page 125](#)

BWoD Important Notes

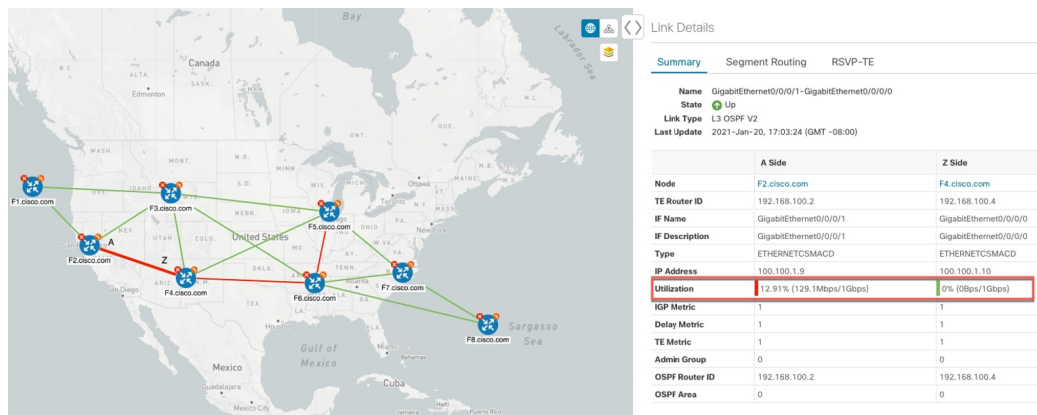
Consider the following information when using BWoD:

- 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**.


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

Figure 35: 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. If a single path cannot be found, BWoD may recommend splitting the path.

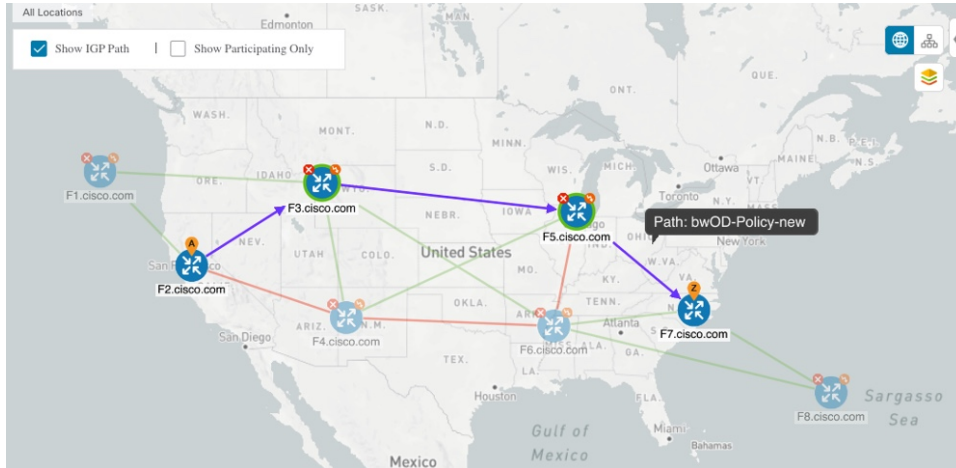
Step 1 Enable and Configure BWoD.

- 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.

f) Click **Preview**.



In the above example, BWoD finds a single path that is under utilized and can still accommodate the requested bandwidth without going above the utilization threshold.



In the above example, BWoD cannot find a single path because of utilization and capacity limitations across several links. In this case, BWoD splits the path to obtain bandwidth and utilization requirements.

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.

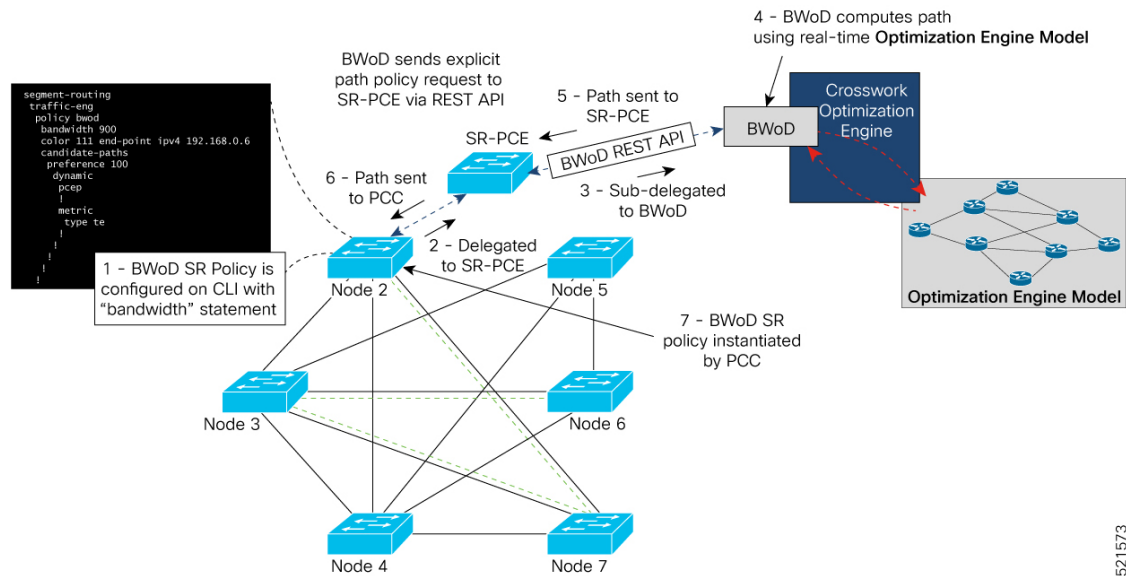
- From the main menu, choose **Traffic Engineering** > **SR-TE**.
- Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View**). Note that the Policy Type is **Bandwidth on Demand**.

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 36: PCC-Initiated BWoD SR-TE Policies



521573


Callout No.	Description
1	<p>A BWoD policy is configured on a PCC via the CLI. For example:</p> <pre>segment-routing traffic-eng policy bwod bandwidth 900 color 100 end-point ipv4 1.1.1.2 candidate-paths preference 100 dynamic pcep ! metric type te ! ! constraints affinity exclude-any name RED ! ! !</pre>

Callout No.	Description
2	The bandwidth statement is added to a PCE delegated SR policy to create a BWoD policy. Once committed, the PCC delegates the path compute to SR-PCE.
3, 4	SR-PCE then sub-delegates the policy to BWoD which attempts to compute a path that meets the bandwidth constraint.
5, 6	If a bandwidth-compliant path is found, the segment list is returned to SR-PCE which forwards it over PCEP to the PCC and the PCC instantiates it. If BWoD is unable to compute a bw-compliant path for the policy or doing so will force an existing BWoD policy to not have a bw-compliant path, best effort paths may be computed by BWoD which attempt to minimize violations. This occurrence will also trigger BWoD to issue an event to the COE events UI indicating which BWoD policies are now on best effort paths.
7	A BWoD SR-TE policy is instantiated.

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 2: 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.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

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)

© 2022 Cisco Systems, Inc. All rights reserved.

