Cisco Crosswork Optimization Engine 2.0 User Guide

First Published: 2021-04-19

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

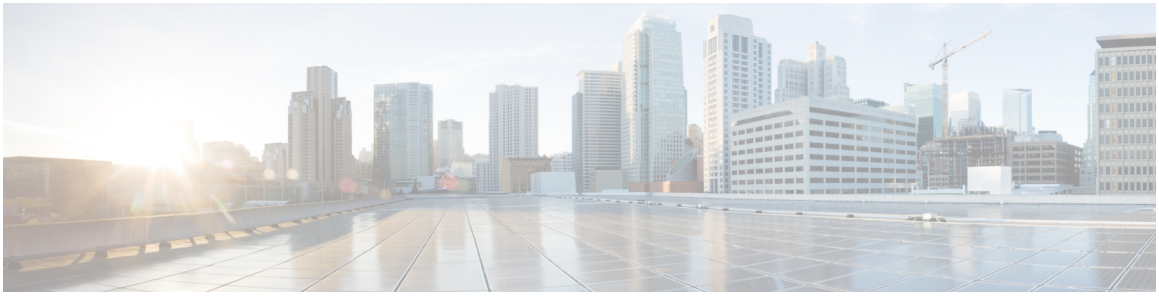

CONTENTS

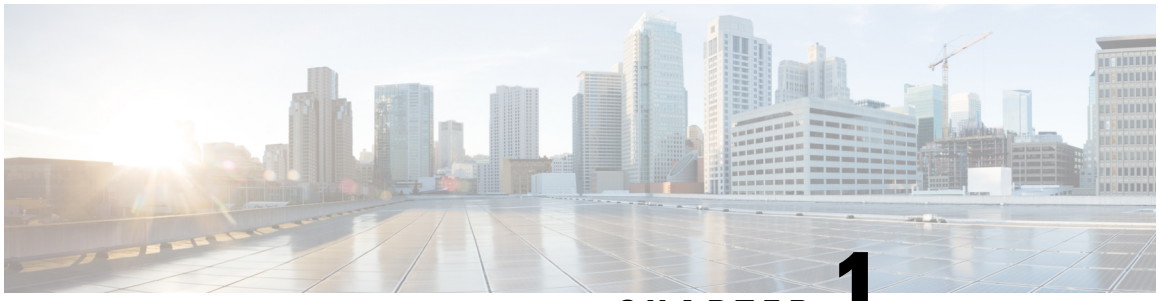

CHAPTER 1

Getting Started

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 Administrator's Guide.

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)

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 as well as 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 policies and RSVP-TE tunnels
- A UI that allows the network operator to perform the following tasks:
  - Provision SR-TE policies and RSVP-TE tunnels and modify or remove them using an intuitive workflow
  - Preview an SR-TE policy or RSVP-TE tunnel before deploying it to the network
  - Continuously track SR-TE policy dynamic path computations to maintain SLA objectives (with correct licensing)
  - Visualize SR-TE policies and RSVP-TE tunnels created directly on the network devices providing a comprehensive view of the active network configuration
- 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 all of the capabilities allowed by Crosswork Optimization Engine. However, either due to licensing or the configuration of the role associated with your user account, you may not be able to access all of the features and functions.

For licensing and ordering information, see the Cisco Crosswork Optimization Engine Ordering Guide (accessible to Cisco Partners)or contact your Cisco Sales representative.

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. For licensing and ordering information, see the Cisco Crosswork Optimization Engine Ordering Guide (accessible to Cisco Partners) or contact your Cisco Sales representative.

Crosswork Optimization Engine and the Crosswork Network Controller Solution

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 fulfilment, 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 Crosswork Optimization Engine as part of the Crosswork Network Controller solution, some options may not be 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 uses the combination of telemetry and data collected from the Cisco Segment Routing Path Computation Element (SR-PCE) to analyze and compute optimal TE tunnels.

Cisco SR-PCE (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 head-end tunnels sourced from the PCC to a PCE peer. The PCC and PCE establish a Path Computation Element Communication Protocol (PCEP) connection that SR-PCE uses to push updates to the network.

Crosswork discovers all devices 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.

About Segment Routing

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

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

Segments

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

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

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

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

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

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 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 head-end computes a solution, resulting in a SID list or a set of SID lists. When the topology changes, a new path is computed. If the head-end does not have enough information about the topology, the head-end might delegate the computation to a path computation engine (PCE).

Explicit SR Policy

When you configure an explicit policy, you specify an explicit path which consists of a list of prefix or adjacency SIDs, each representing a node or link along on the path. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the network to follow the specified path instead of the shortest path calculated by the IGP. If a packet is steered into an SR policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

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.
- Configuring affinity and disjointness at the same time is not supported.

Related Links

Provision SR-TE Policies, on page 35
Configure Link Affinities, on page 38

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:

- End-point control, which is associated with establishing and managing TE tunnels at the headend and tailend.
- 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 end-point 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 Explice Route Object (ERO). The explicit path could be administratively specified, or automatically computed using an algorithm such as constrained shortest path first (CSPF).

The explicit path specified in the ERO may 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.

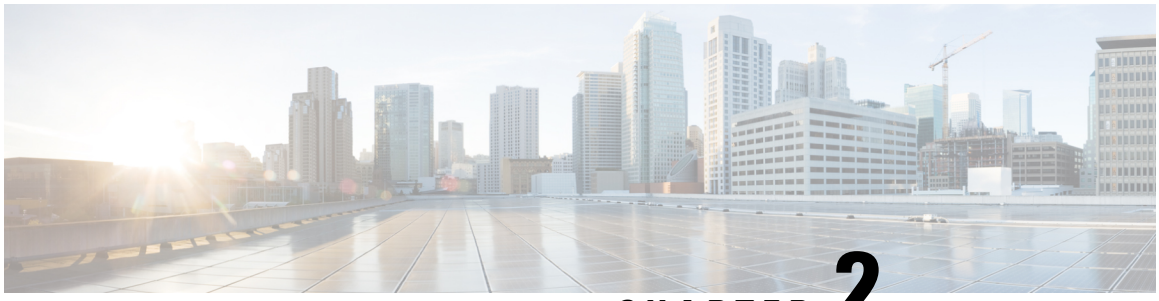

CHAPTER 2

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

The Home page displays the dashboard which provides an at-a-glance operational summary of the network being managed, including reachability and operational status of devices. Each dashlet represents different types of data belonging to the same category.

Figure 1: Crosswork Home page

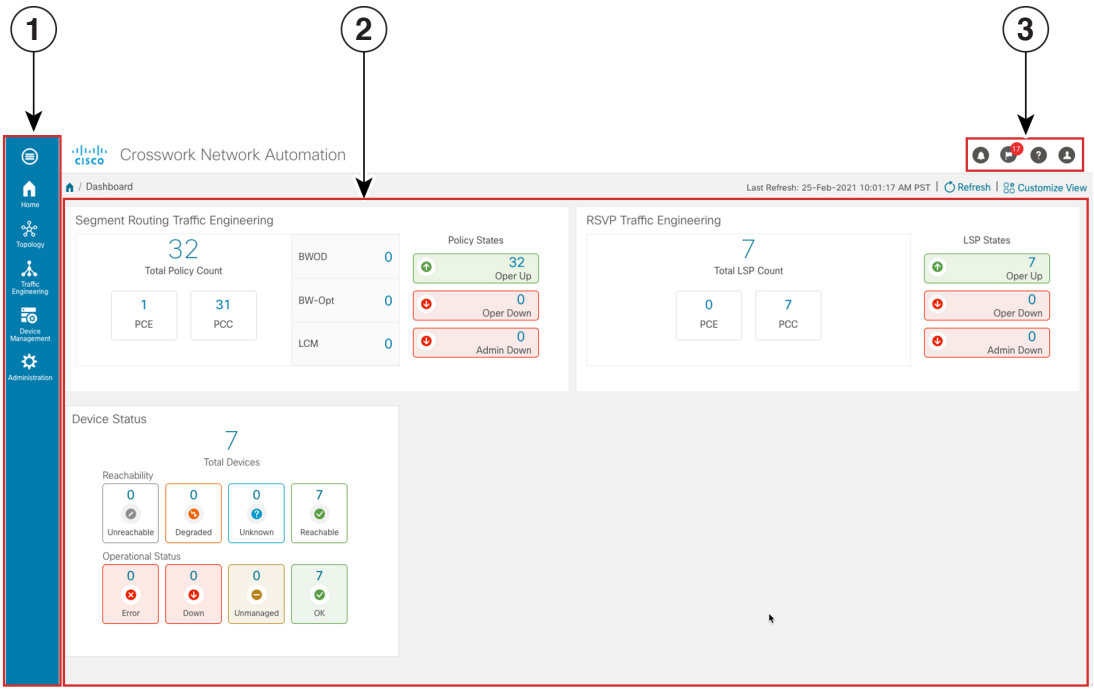

Callout No.	Description
1	**Main Menu**: The main menu allows you to navigate to installed Cisco Crosswork applications and device management and administrative tasks. Menu options may look slightly different depending on what Cisco Crosswork applications are installed.
2	**Dashlets**: Information varies depending on what Cisco Crosswork applications are installed. • 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: 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**.

Figure 2: Cisco Crosswork UI and Topology Map

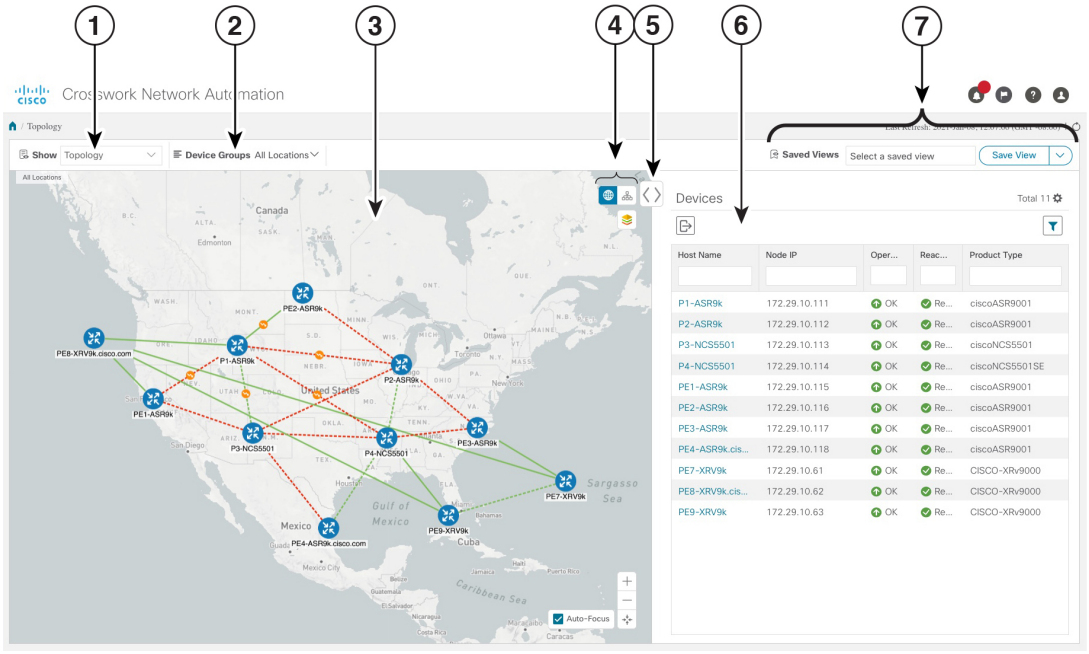

Callout No.	Description
1	**Topology Map View**: From the **Show** drop-down list, click the option that displays the data that you would like to see on the map. If **Topology** is selected, devices and links in the network are displayed. If **Traffic Engineering** is selected, TE tunnel information is displayed.
2	**Device Groups**: From the drop-down list, click the group of devices that you want to focus on the topology map. All other device groups will be hidden.

Callout No.	Description
3	**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. **Devices:** • 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. The number in a blue circle (2) indicates the number of devices in the cluster. Displaying devices in this manner helps prevent overlap and clutter on the map. **Links:** • A solid line indicates a *single link* 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 *aggregated* link that represents more than one link, or the use of multiple protocols (for example, IPv4 and IPv6) on the same physical link. • To view link information details, click on the link.
4	: 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. : 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. : The Display Preferences window allows you to change display settings for devices, links, utilization, and TE tunnel metrics.
5	**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.
6	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-TE policy, or RSVP-TE tunnel.
7	**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 using the topology map.

Step 1 From the main menu choose **Topology**.

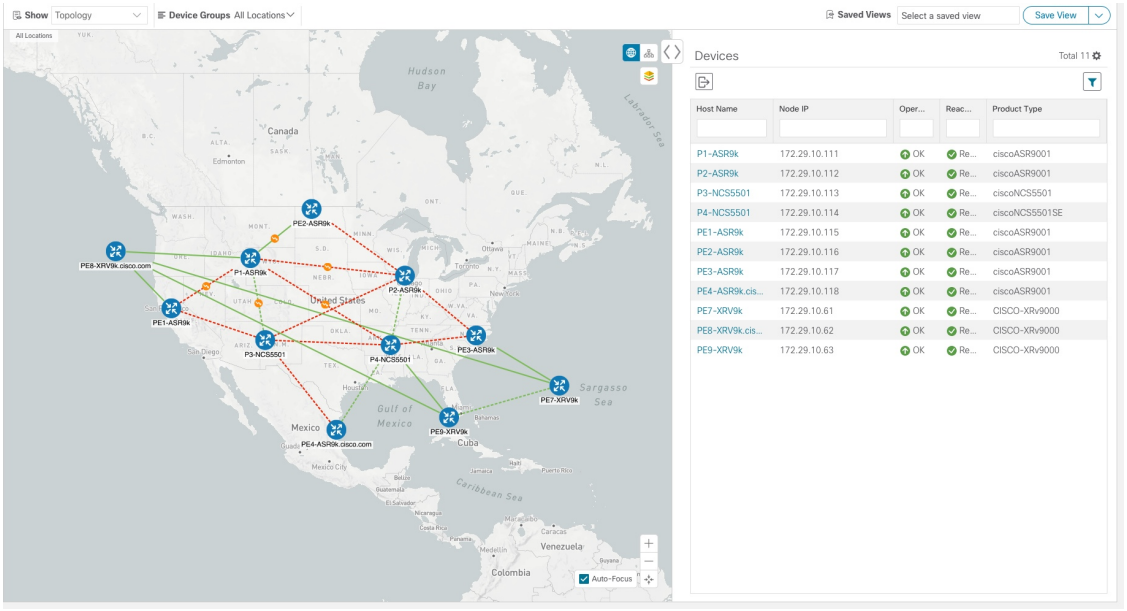

Step 2 To quickly view the host name, reachability state, IP address and type of device, hover the mouse over the device icon.

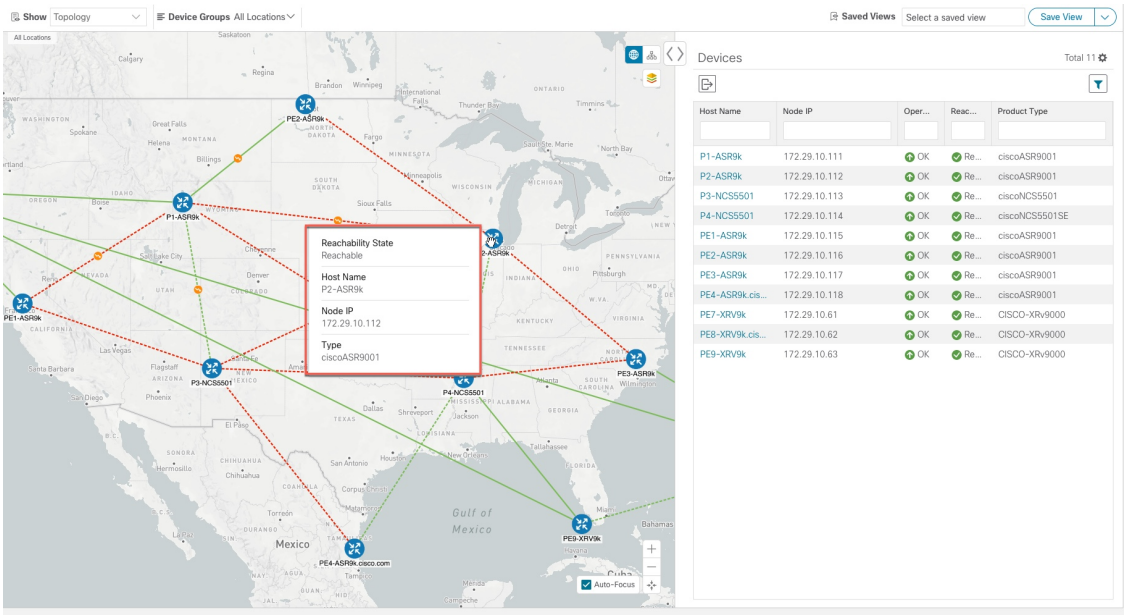

Step 3 To view more device details, click on the device icon.

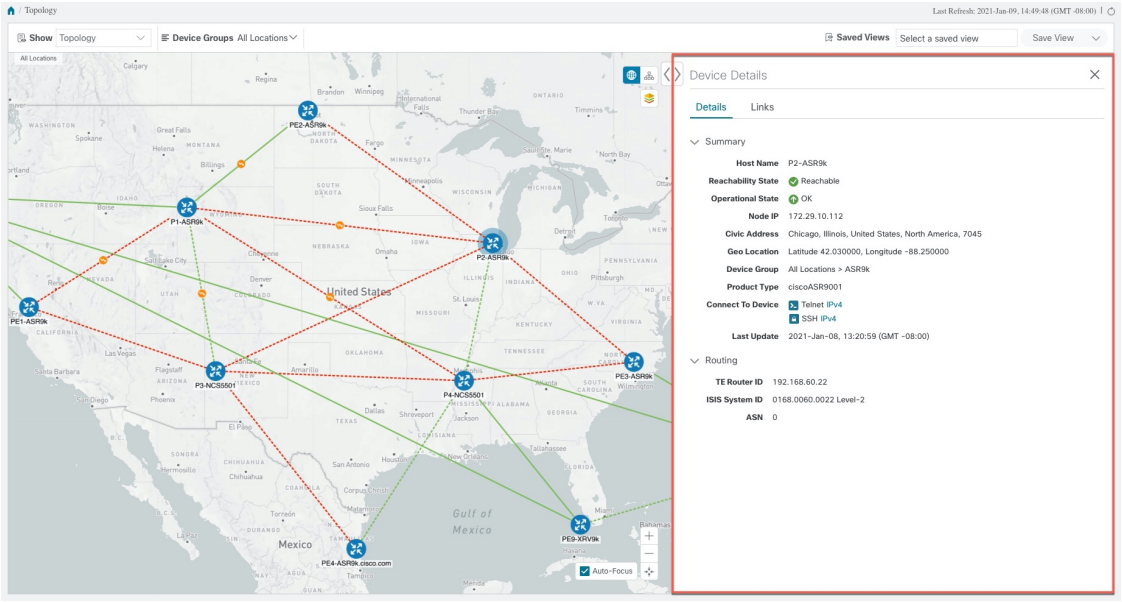

In a multiple IGP setup, you can also view all the IGP, IS-IS, and OSPF processes. See the following examples:

Figure 3: Multiple IGP: OSPF Processes

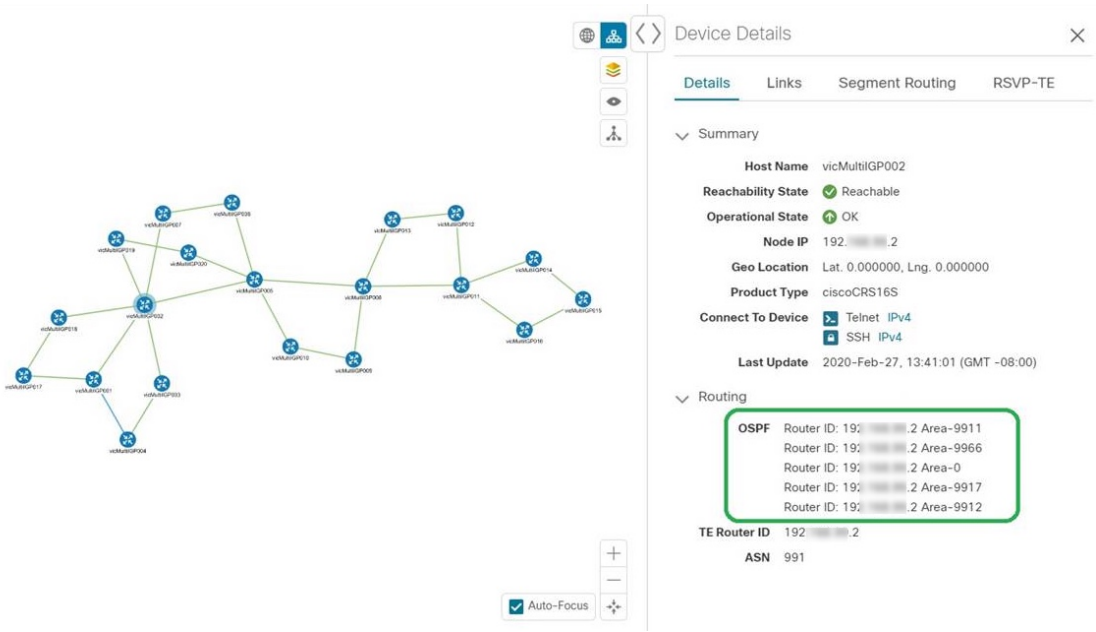

Figure 4: Multiple IGP: ISIS Processes

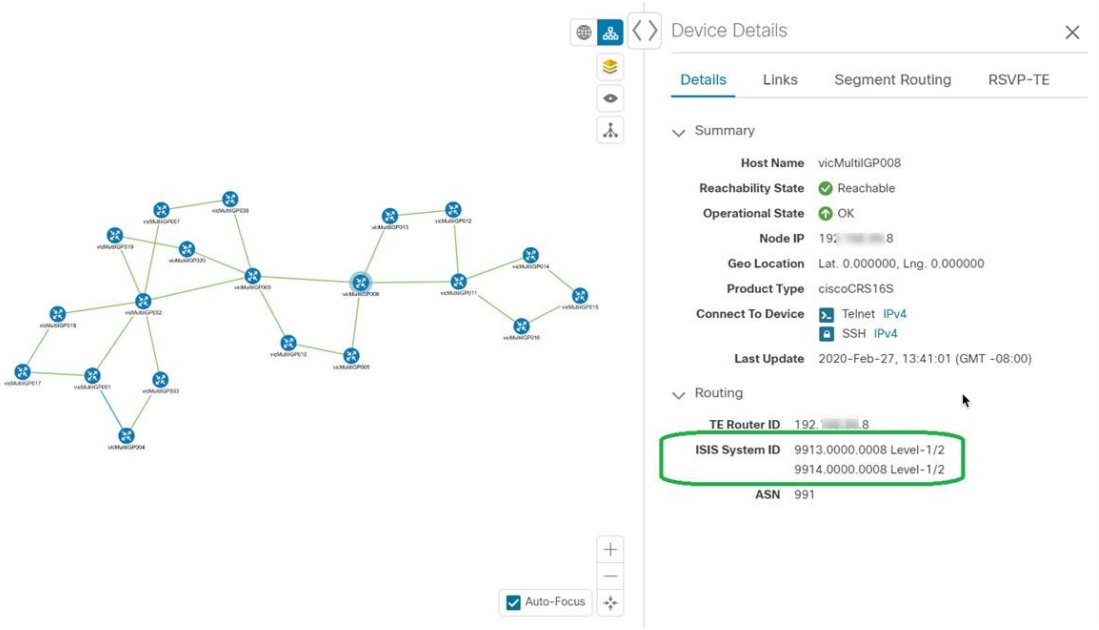

Figure 5: Multiple IGP: OSPF and ISIS Processes

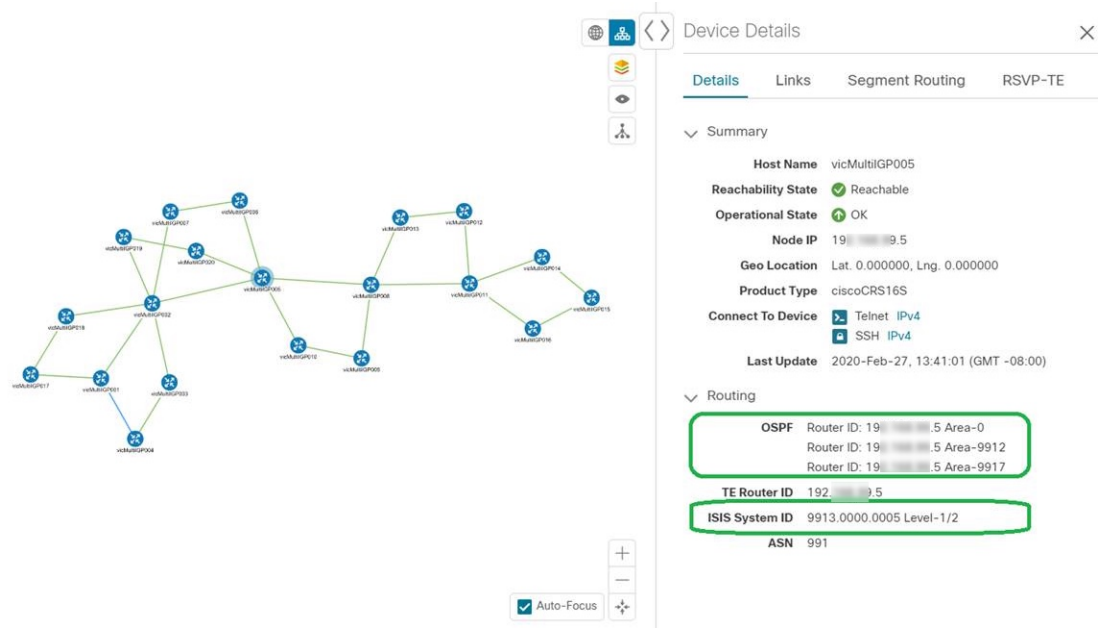

Step 4 To view links on the device, click the **Links** tab and expand the right panel to see all the link details.

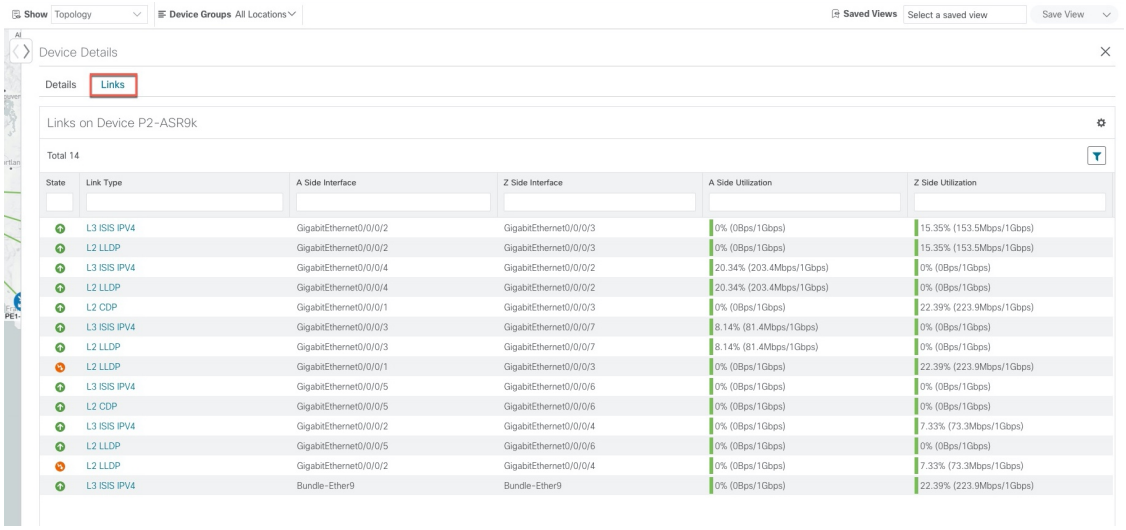

Step 5 Collapse the side panel and close the **Device Details** window.

Step 6 Click on a dashed line. A dashed line indicates an aggregated link that represents more than one link, or the use of multiple protocols (for example, IPv4 and IPv6) on the same physical link.. The links are displayed.

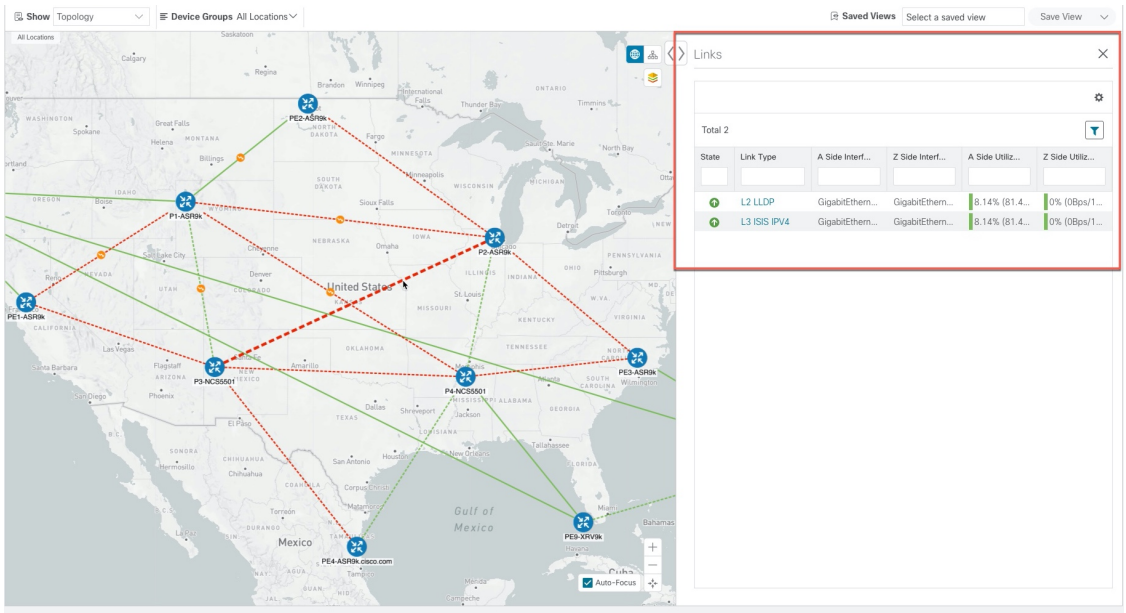

Use Device Groups to Filter Your Topology View

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 device groups they belong to. By default, all devices initially appear in the **Unassigned Devices** group.

This example walks you through how Device Grouping works in the geographical and logical maps.

Step 1 From the main menu, choose **Topology**. By default, only devices that have Geo Location set will appear on the geographical map.

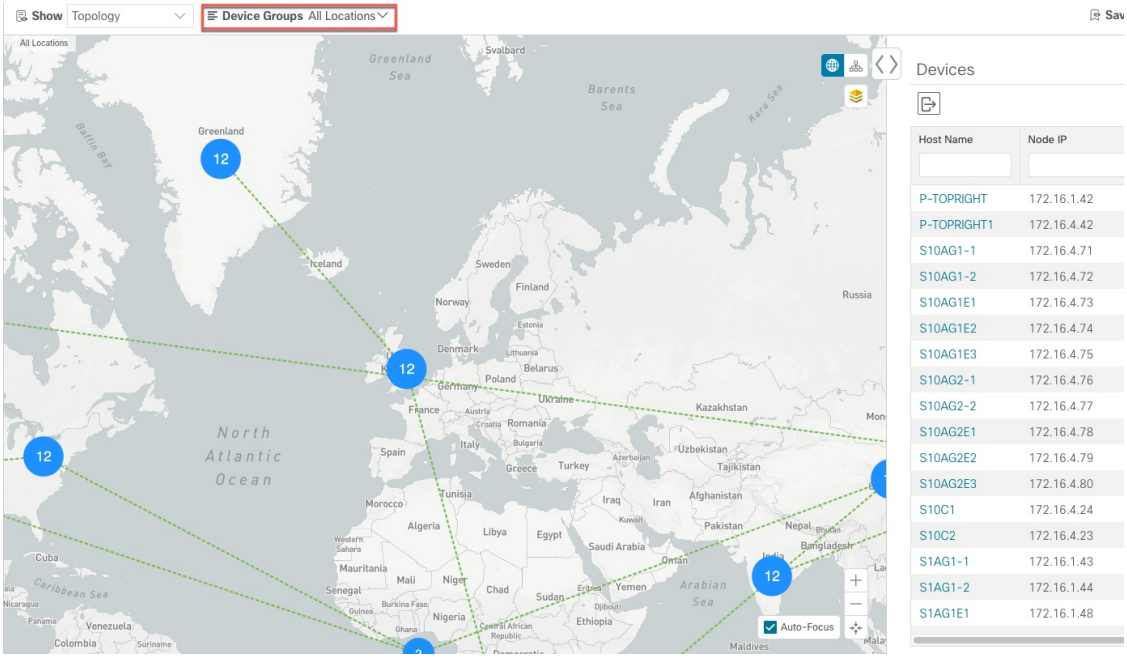

Step 2 From the **Device Group** drop-down list select a group (US West). Only the devices in that group and related links are displayed on the geographical map. Note that the Devices table has also been filtered to list only those devices in the group.

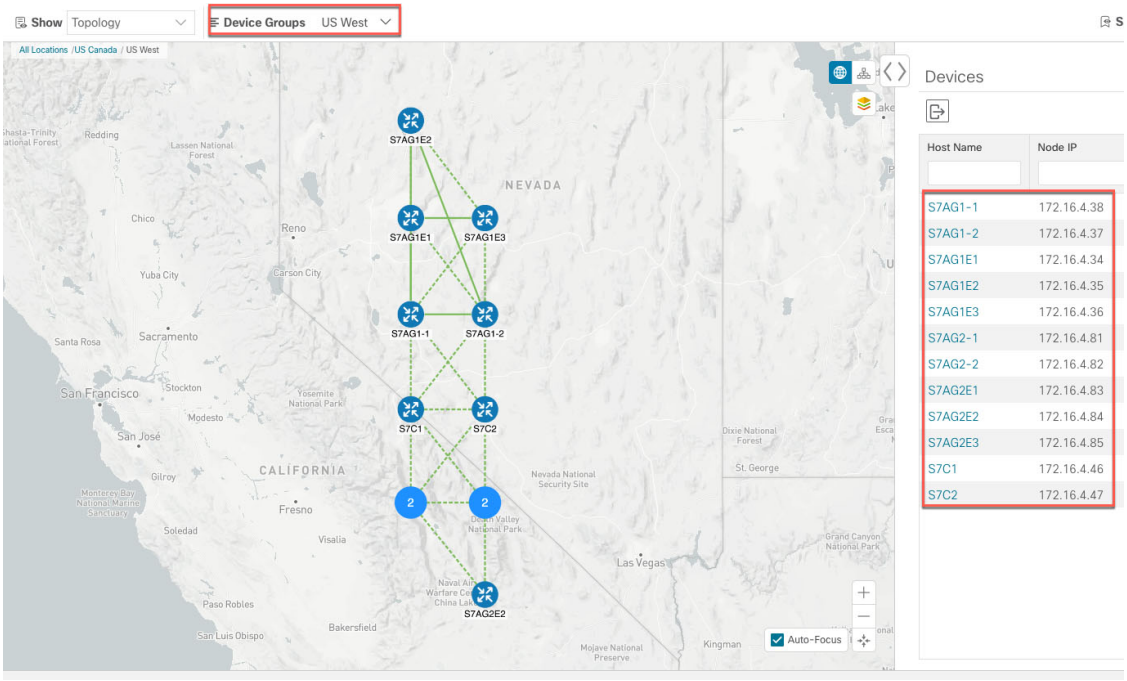

Step 3 Click  .

Step 4 From the **Device Group** drop-down list, select **All Locations** and check **Show Groups** if it is not already checked. Note that you can see all device groups in this view. Device groups can be seen in this way only within the logical map.

Note If **Show Groups** checkbox is de-selected, all the device groups are expanded , and could lead to a a cluttered map.

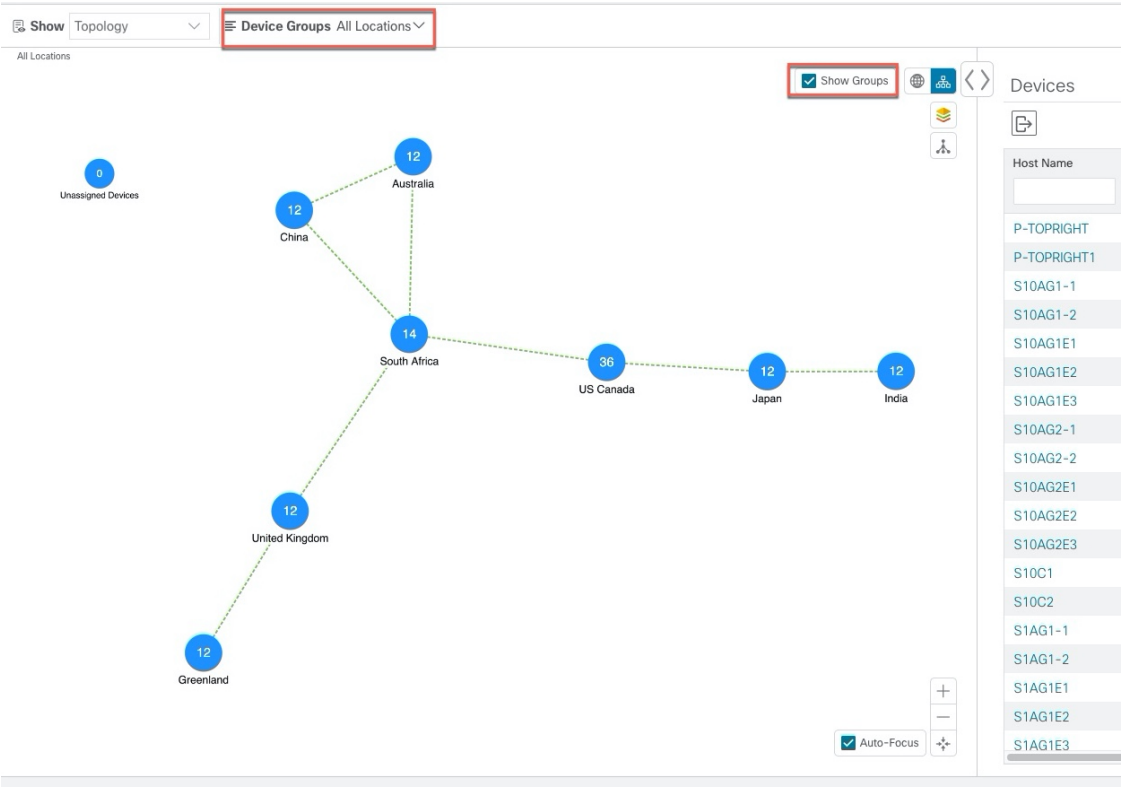

Step 5 Click the US West group. Again, only devices that belong to this group are shown in the topology map and the Devices table.

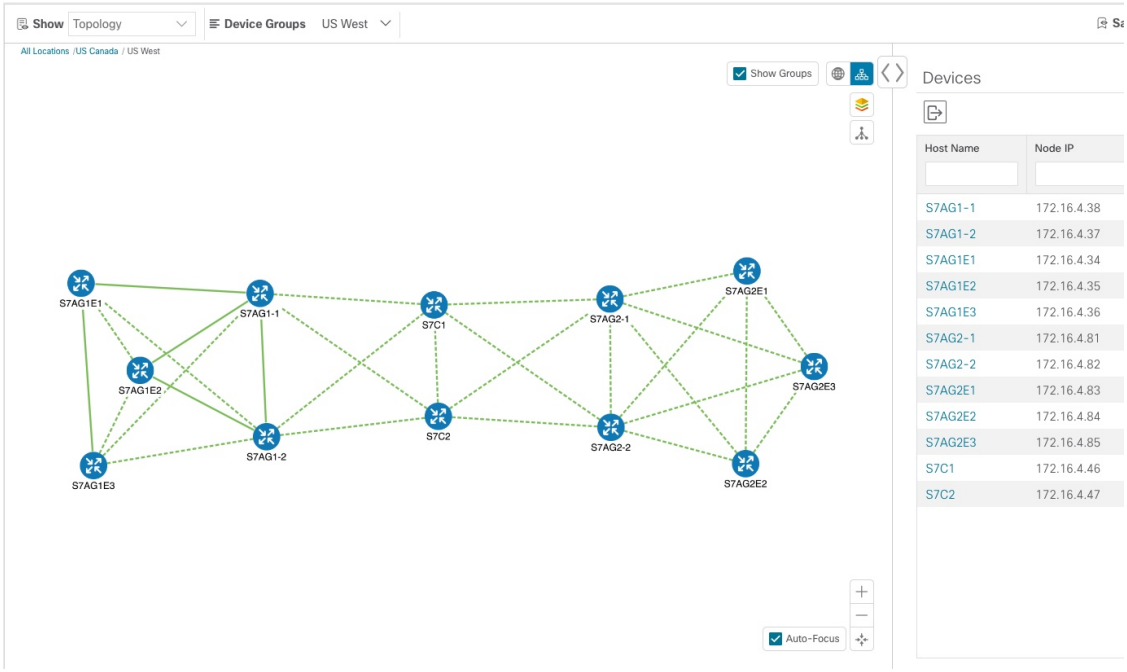

Step 6 Filter devices in the Device table by entering S7C in the hostname. The Device table displays only devices that match the filtering criteria. However, filtering the Device table does not filter the devices visually on the topology map. The only way to visually filter devices on the geographical or logical maps is to use device groups.

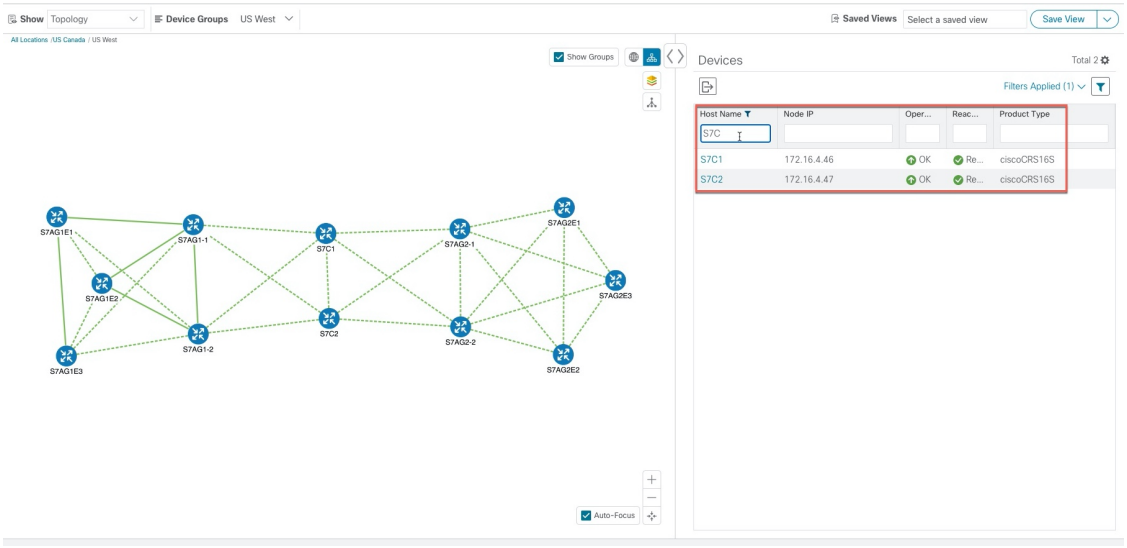

Create and Modify Device Groups

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

Step 2 From the Device Groups tree, click … next to a group.

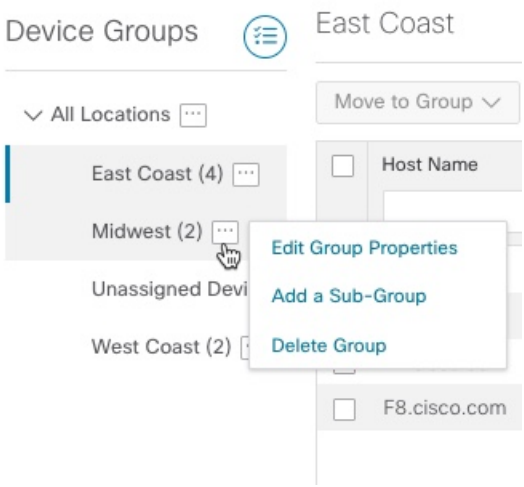

Step 3 Choose to add, delete, or edit (rename or move) a group. If you delete a group, all devices that belong to that group are moved to the Unassigned Devices group.

Note Devices can belong to only one device group.

Step 4 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 matches the rule will be placed in the group.

Note Dynamic rules do not apply to devices that already belong in groups. You must move them to Unassigned Devices if you want to include them as part of the devices that the dynamic rule will consider.

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 .

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 Check the **Enable Rule** checkbox. After the rule is enabled, the system checks devices every minute and will either create or assign them into groups.

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 corresponding group hierarchy.

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

a) Select **…** next to All Locations and click **Add a Sub-Group**.
b) Enter the New Group details and click **Save**.
c) Select **…** next to the unassigned created dynamic group and select **Edit Group Properties**.
d) Click **Change Parent Group** and select the 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 19
- Set Display Behavior of Device Groups for TE Tunnels , on page 19

Customize the Display of Links and Devices

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

- For devices, you can choose whether 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.
- For links, you can choose whether 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.
- For TE tunnels, you can choose whether to show IGP, TE, and delay (latency) metrics. By default, these metrics are not enabled.

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 **Admin** > **Settings** > **User Settings** and select one of the behavior options.

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

Save Topology Views for Easy Access

When you rearrange the devices and links on a map, your changes are not normally saved. When you open the map later, your map settings are lost.

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 edit (modify, rename, or delete) the view.

Step 1 To create a custom view:

a) Customize the current map view until it contains only the information you want and until the layout meets your needs.
b) When you have the view the way you want it, click **Save View**.
c) Enter a unique name for the new custom view and click **Save**.

Step 2 To delete a custom view:

a) Click the **Saved Views** field.
b) Find the custom view you want to delete and click  .

Step 3 To edit a custom view:

a) Click the **Saved Views** field.
b) Click the custom view you want to edit. The custom view appears.
c) Make any changes to the current view and click **Save View**. This overwrites the previously saved view.

Step 4 To rename or save a view with another name:

a) Click the **Saved Views** drop-down list.
b) Select the appropriate option.

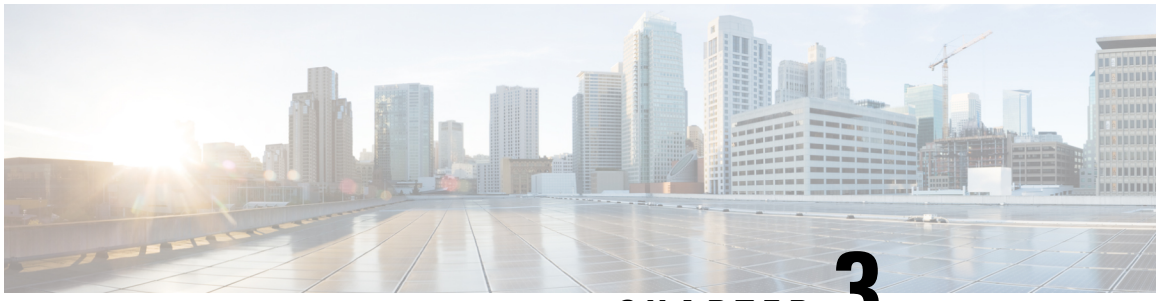

CHAPTER 3

Monitor SR-TE Policies and RSVP-TE Tunnels

This section contains the following topics:

View SR-TE Policies and RSVP-TE Tunnels on the Topology Map

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

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

Note Throughout this document, 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**.

Figure 6: Traffic Engineering UI

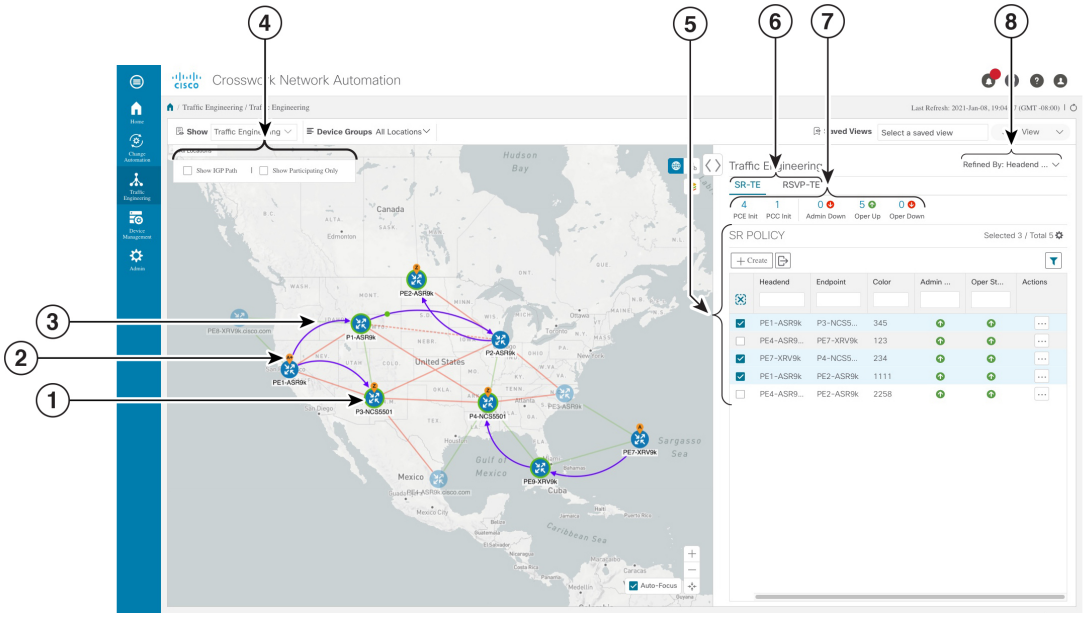

Callout No.	Description
1	SR-TE Policies—A device with a green ( ) outline indicates there is a node SID associated with that device or a device in the cluster. RSVP-TE Tunnels—A device with a solid orange outline ( )indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered. **Note** RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.
2	**SR-TE Policy and RSVP-TE Tunnel 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 or RSVP-TE tunnel that originates from a node. The **Z+** denotes that the node is a destination for more than one TE tunnel.
3	**SR-TE Policies and RSVP-TE Tunnels**: When SR-TE policies or RSVP-TE tunnels are selected in the SR-TE Policy or RSVP-TE Tunnel tables, they show as purple directional lines on the map indicating source and destination. An adjacency segment ID (SID) is shown as a green dot on a link along the path ( ).
4	Click the appropriate check box to enable the following options: • **Show IGP Path**—Displays the IGP path for the selected SR-TE policy. This option is not available when viewing RSVP TE tunnels. • **Show Participating Only**—Displays only links that belong to selected TE tunnels. All other links and devices disappear.

Callout No.	Description
5	The content of this window depends on what has been selected or filtered. In this example, the SR-TE 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 TE tunnels, you can do the following: • Visualize SR-TE Policies and RSVP-TE Tunnels Example, on page 24 • Provision SR-TE Policies, on page 35 • Provision RSVP-TE Tunnels, on page 41
6	Click on either the **SR-TE** or **RSVP-TE** tabs to view the respective list of TE tunnels.
7	The **Mini Dashboard** provides a summary of the operational TE tunnel status and the number of PCC and PCE initiated tunnels that are *currently* listed in the SR Policy or RSVP-TE tables. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the SR Policy or RSVP-TE 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.

Figure 7: RSVP-TE Tunnels

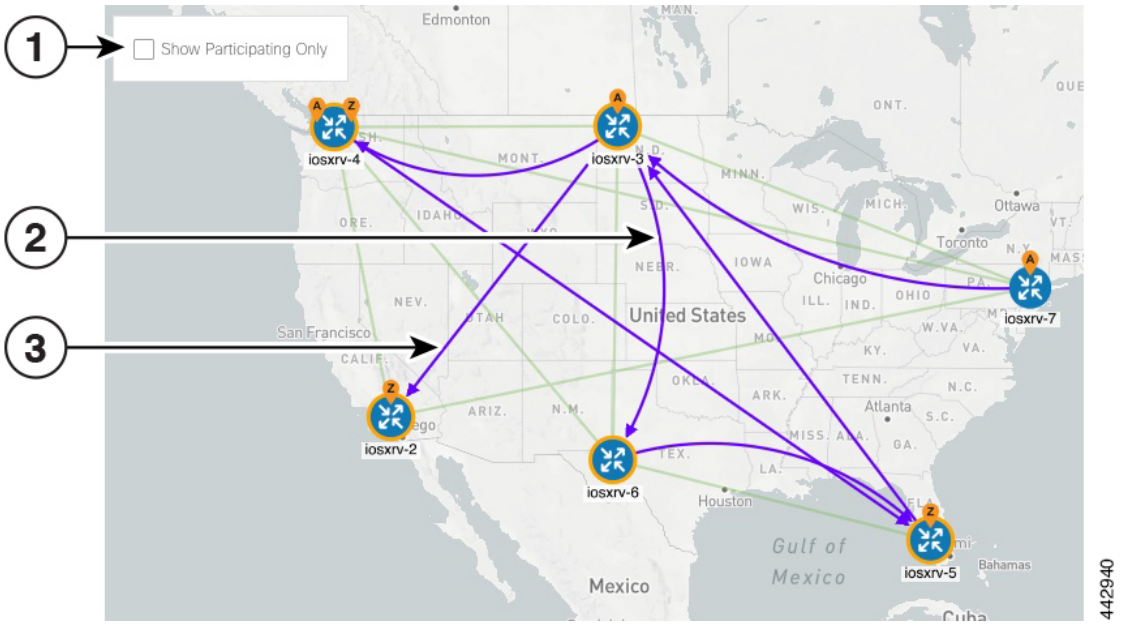

The display of RSVP-TE tunnels is similar *except* for the following:

Callout No.	Description
1	The **Show IGP Path** option is not available.

Callout No.	Description
2	Record Route Object (RRO) paths are shown as straight lines.
3	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.

Visualize SR-TE Policies and RSVP-TE Tunnels Example

This example walks you through a number of TE tunnel visualization features that are available from the topology map. The topology map displays TE tunnels provisioned using the UI along with tunnels discovered from the network by SR-PCE. From there you can drill down to details and visualization of participating TE tunnels.

In this example, we assume that devices and SR-TE policies have been added and device groups have been created. SR-TE policies are not yet highlighted in the map.

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

Figure 8: Topology Map Example

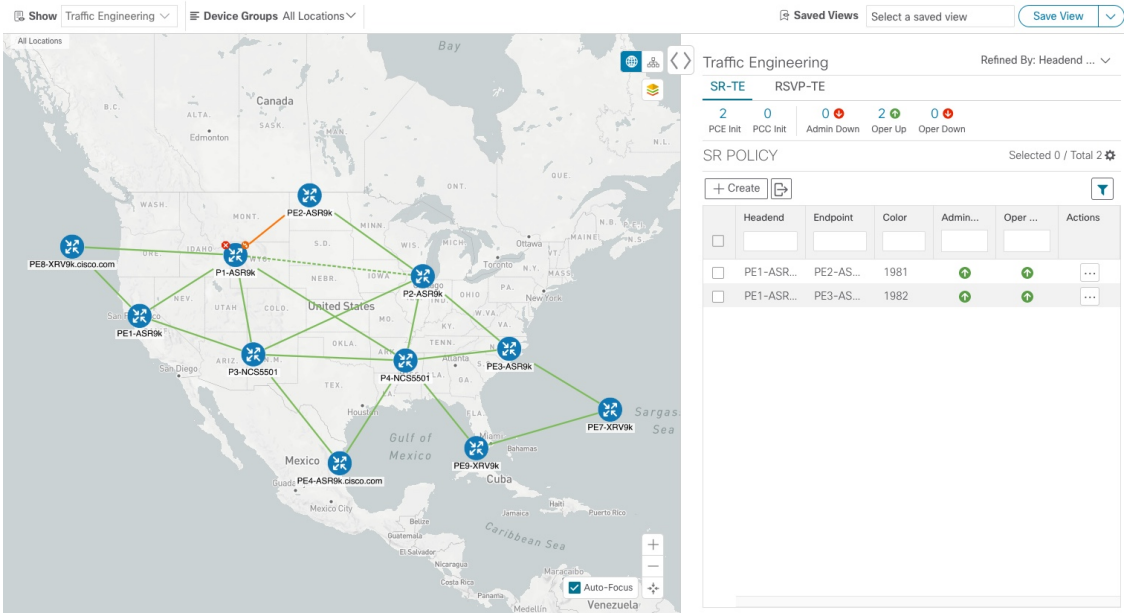

Step 1 From the **SR Policy** table, check the checkbox next to the SR-TE policies you are interested in. In this example, there are two SR policies selected.

Figure 9: SR-TE Policy Selection

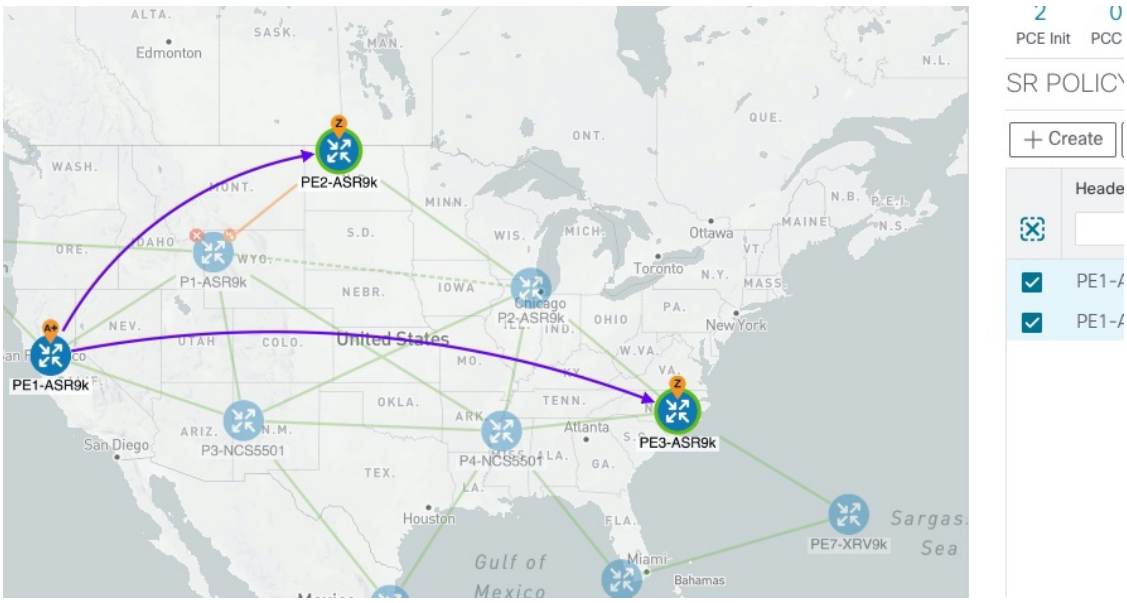

After SR-TE policy selection, the map displays the following:

- SR-TE policies appear as purple links with arrows that indicate the path direction.
- The PE1-ASR9k node is an origin for the both selected policies. PE2-ASR9k and PE3-ASR9k are destinations for the selected policies. SR-TE 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+** would denote that the device is a destination for more than one policy.

  Note If both **A** and **Z** are displayed in a device cluster, at least one device in the cluster is a source and another is a destination.

-  indicates that PE2-ASR9k and PE3-ASR9k have node SIDs.

Step 2 From the **SR Policy** table, *hover* over a selected policy. The path name of that policy is highlighted on the topology view. You will also see prefix SID information.

Figure 10: Hover over an SR-TE Policy for Details

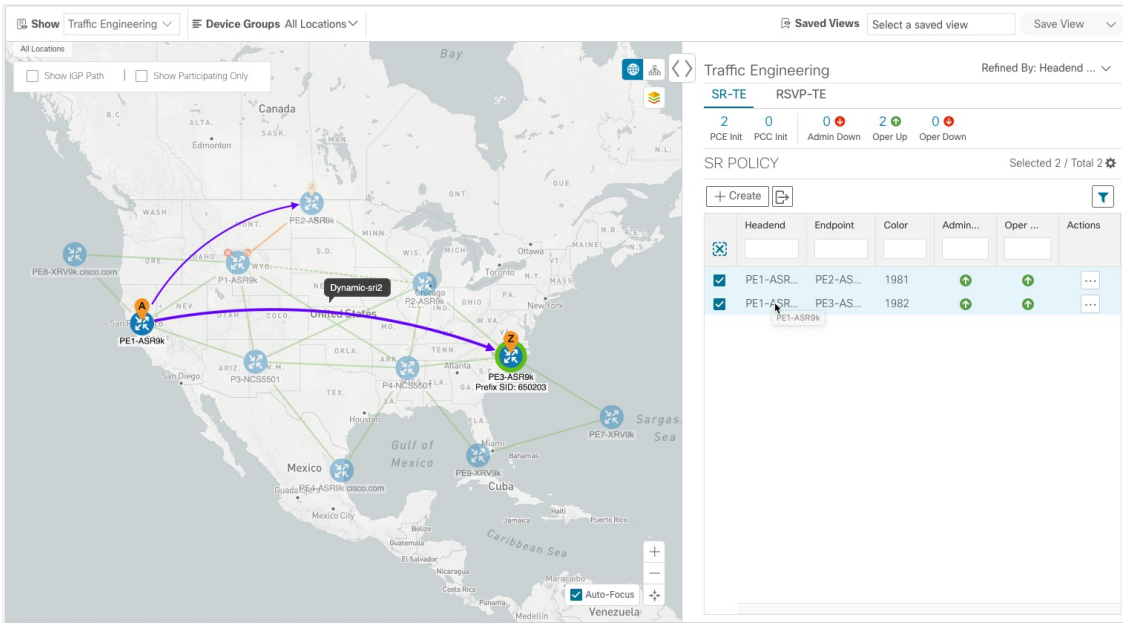

Step 3 To see the physical path between endpoints, check the **Show IGP Path** check box (available only with SR-TE policies). The IGP paths for the selected SR-TE policies are displayed, with straight lines, instead of the segment hops.

Figure 11: IGP Paths

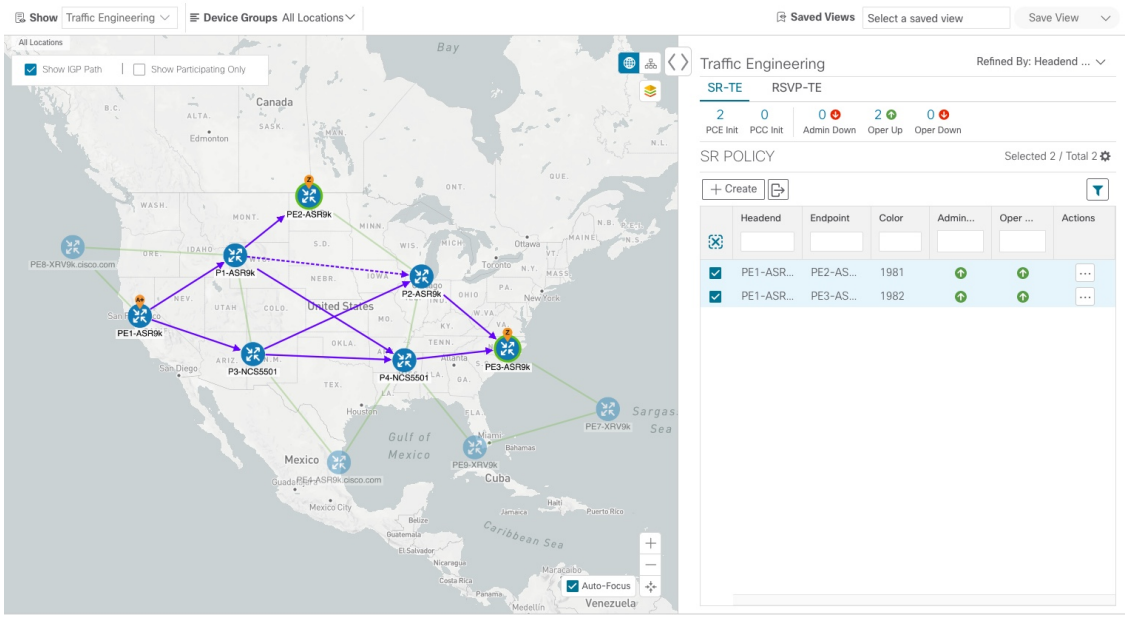

Step 4 Check the **Show Participating Only** check box. All non-participating links and devices disappear. Only participating policies are displayed.

Figure 12: Participating SR-TE Policies

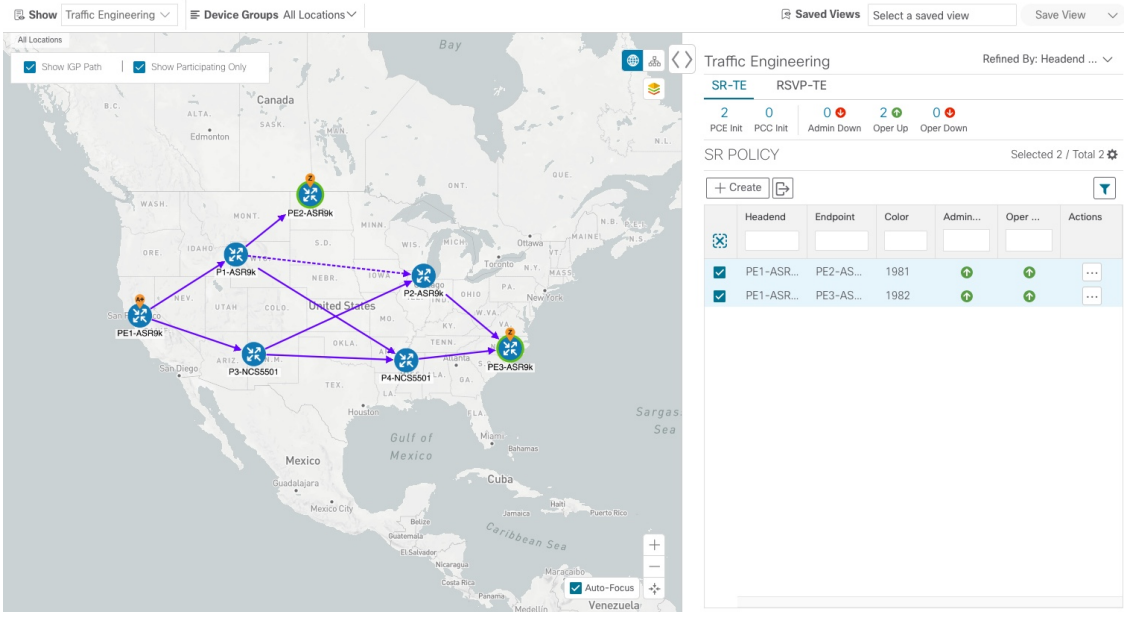

Step 5 To view the IGP, TE or Delay metrics for each tunnel along a policy's path, do the following:

a) For SR-TE policies only, confirm that the **Show IGP Path** checkbox is checked.

b) Click .

c) Click the **Metrics** tab.

d) Toggle applicable metrics to **ON**.

The metric details are displayed for each policy on the map.

Figure 13: IGP, Delay, and TE Metrics

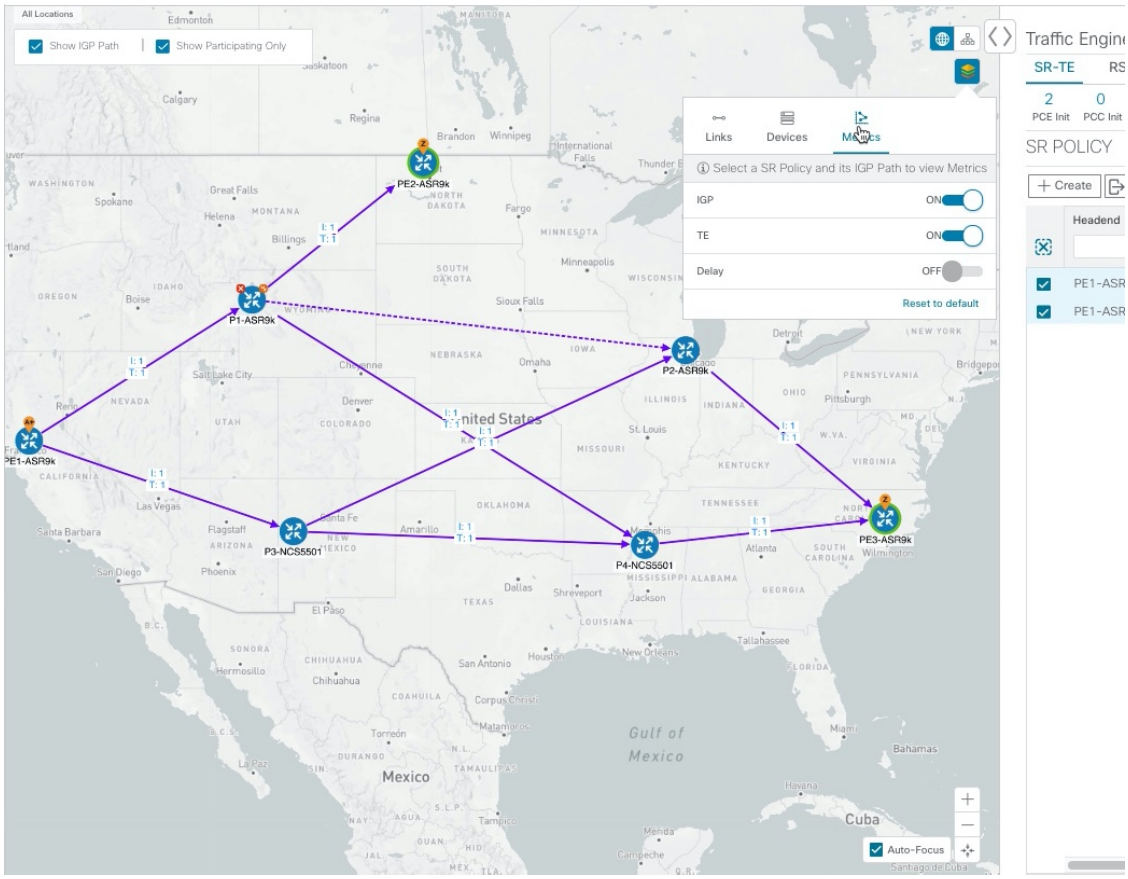

Step 6 Click ⁂ to display the logical view.

Figure 14: Logical Map

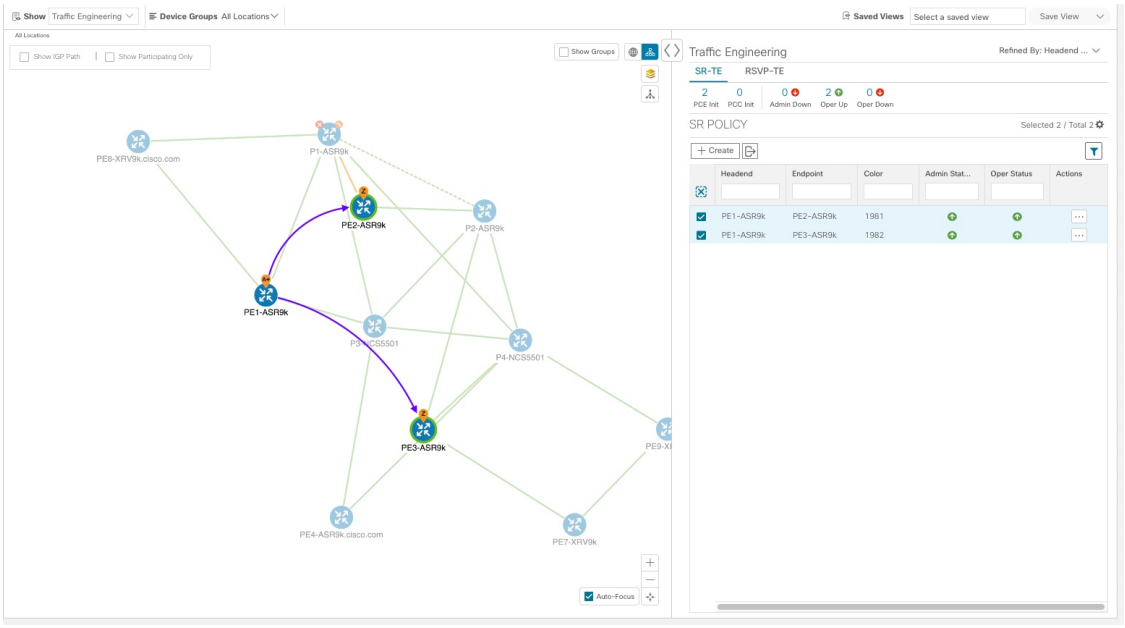

You are able to see the same information (aside from geographical location) that is available on the geographical topology map. You also have the ability to move devices and links on the map to make it easier to view. Click **Save View** to save the current view and retrieve it later.

Step 7 To view SR-TE policy details such as disjoint groups, metric type, segment hop information, and so on, click ⋯ under the **Actions** column from the table.
The **SR Policy Details** window is displayed in the side panel. Note that only the selected policy is now highlighted on the topology map.

Figure 15: SR-TE Policy Details

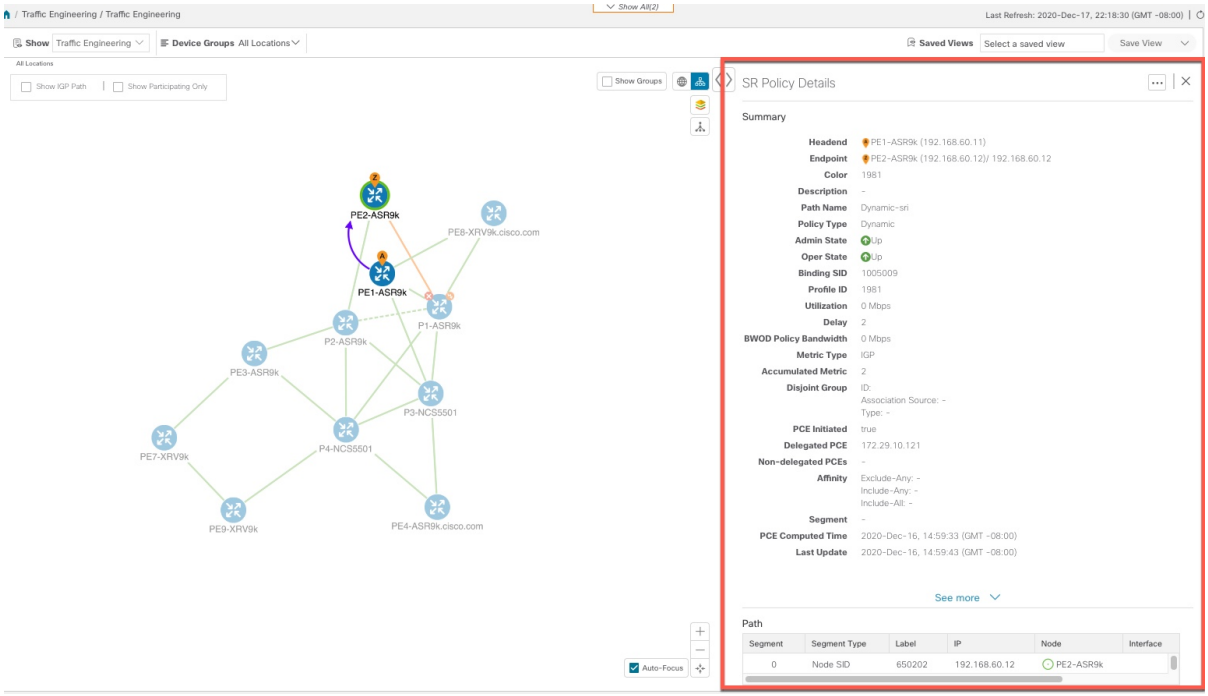

Step 8 Close (✕) the current view to return to the **SR Policy** table.

Step 9 To understand how device groups are displayed with the selection of SR-TE policies or RSVP-TE tunnels, uncheck any SR-TE policies that might be selected and check **Show Groups**.

Figure 16: Show Groups

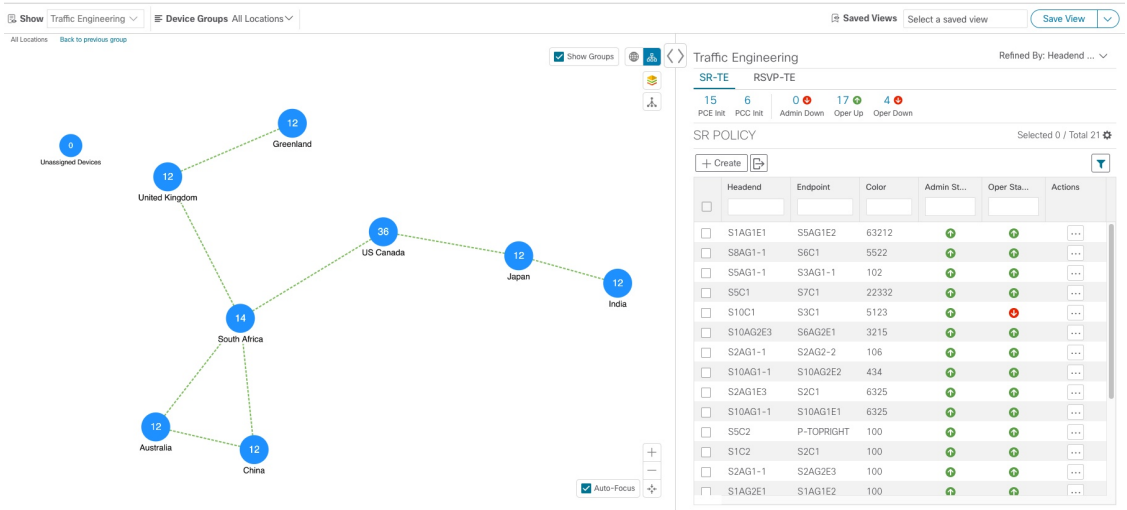

Step 10 Selecting a specific group from the **Device Groups** drop-down list, will only display that group in the map and . In this example, **Australia** is selected and the associated SR-TE policy is selected and displayed.

Figure 17: Device Group Selection

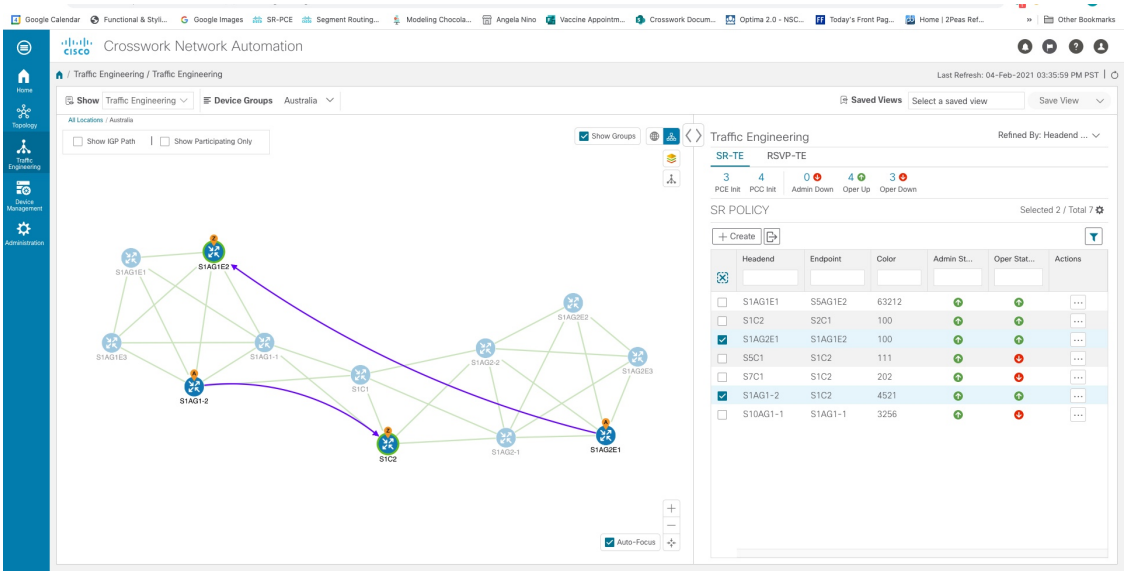

Step 11 If you select a policy where participating devices are not part of the selected group, then a dialog appears giving you an option to switch the group view. This is the default behavior. If this window does not appear, then the administrator has configured the display to automatically switch view or stay in the current view. For more information, see Set Display Behavior of Device Groups for TE Tunnels , on page 19.

Figure 18: Switch Device Group Dialog

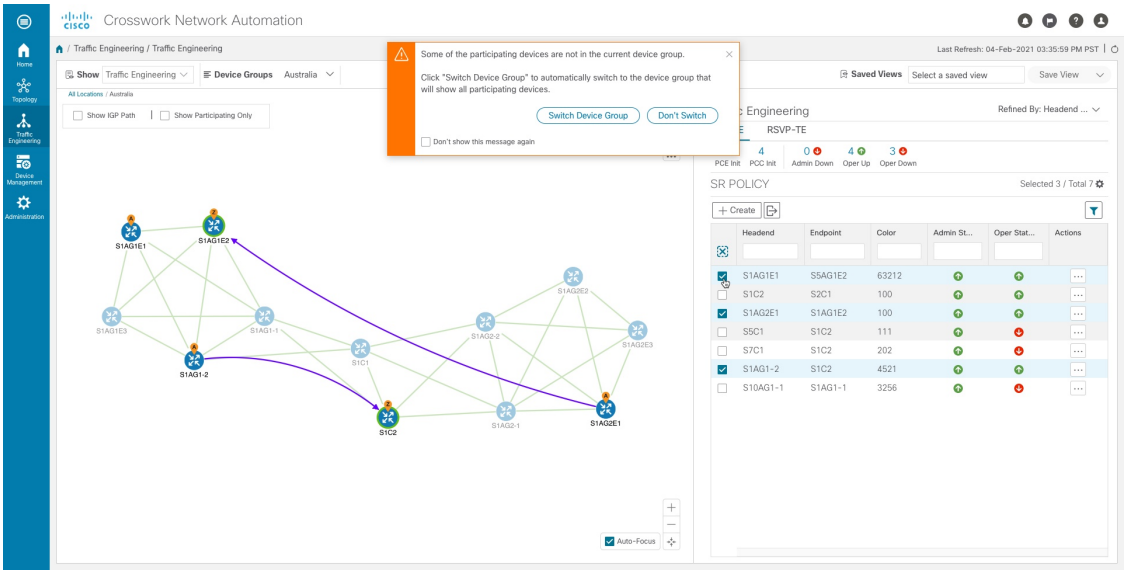

Step 12 If you select **Switch Device Group**, then the group will change and you will see all participating devices for the SR policies you have selected.
To go back to the previous group view, click **Back** (this link appears later in the yellow text area indicated in the following figure).

Figure 19: Result of Switching Device Group

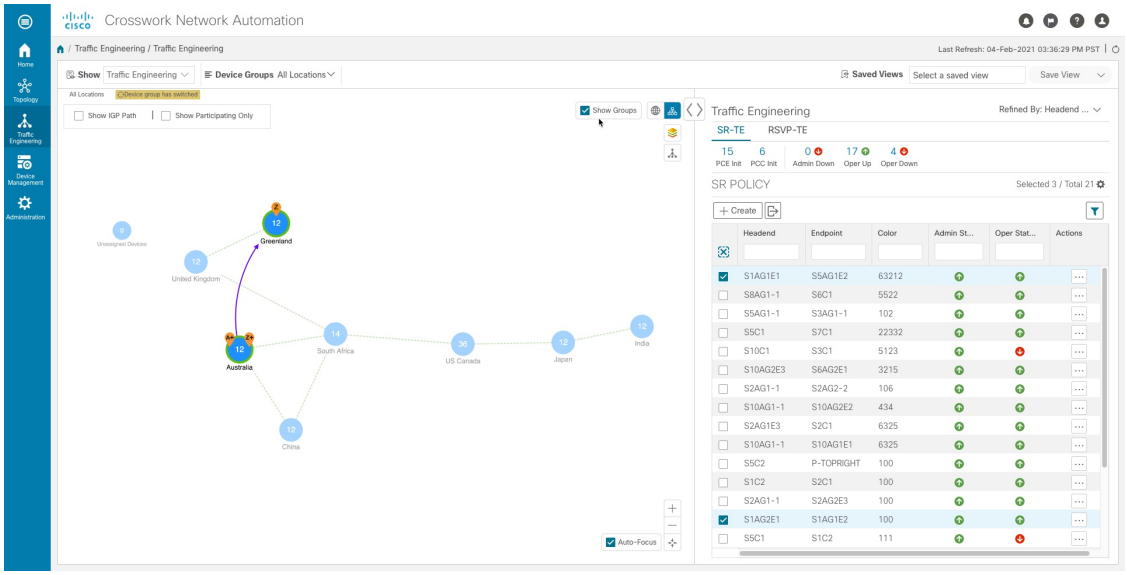

Step 13 You can also use the Mini Dashboard to drill down and focus on certain SR-TE policies.

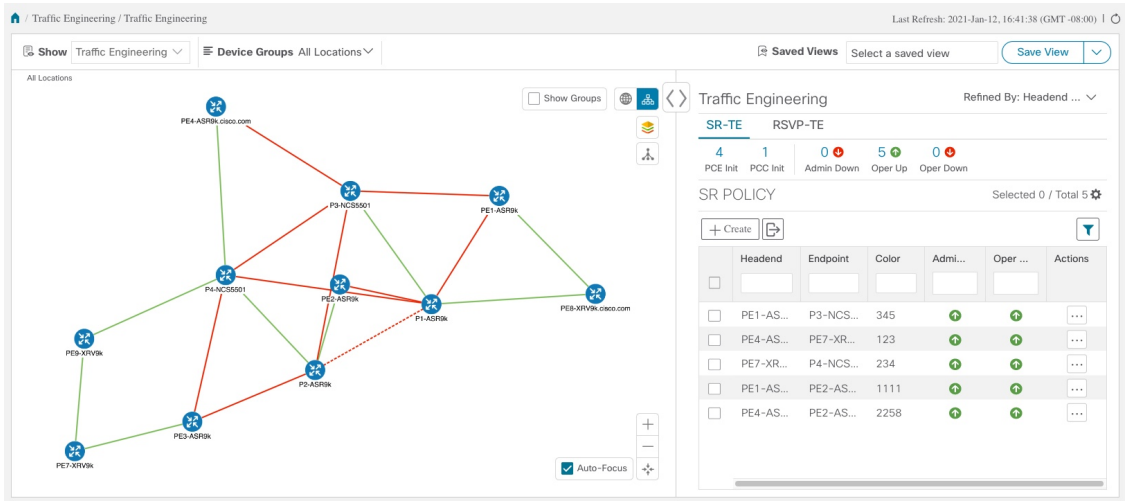

To filter the SR Policy table to show only PCE-initiated policies, click the value for PCE Init from the SR-TE Mini Dashboard. Note that the **Filters Applied** text appears.

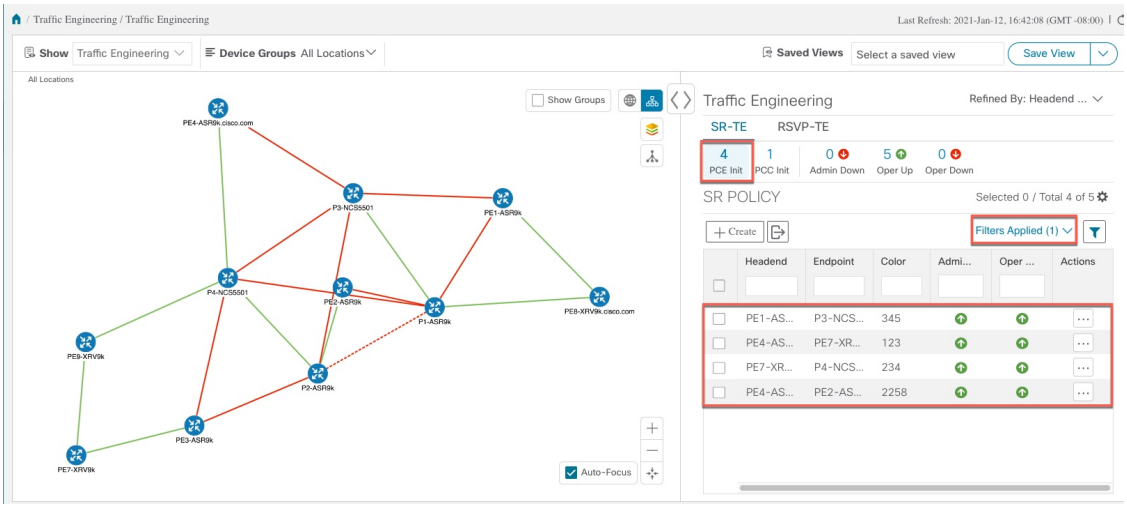

Step 14 Change the layout of the nodes. To save the layout and the filtered list of SR policies, click **Save View**.

Note You cannot save a custom view with any SR-TE policies selected.

Step 15 To remove filter criteria, click **Filters Applied** > **Clear All Filters**. You can also select individual filters if more than one filter has been applied.

Configure Timeout Settings

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

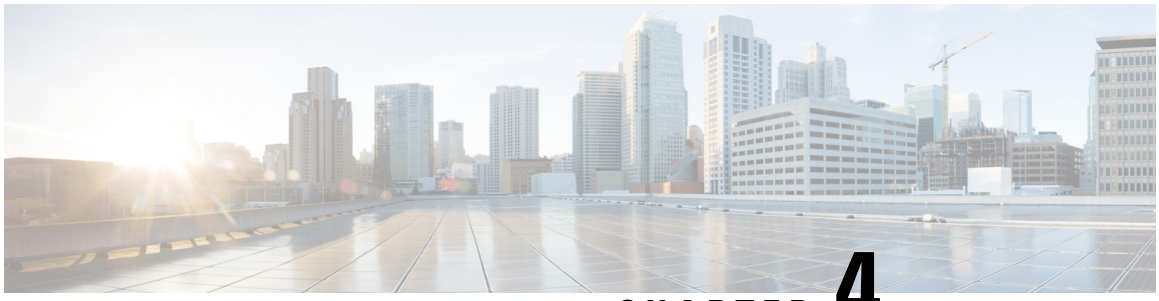

CHAPTER 4

Provision SR-TE Policies

This section contains the following topics:

SR-TE Policy Support

Table 1: Supported Features

Capability	Notes
PCE-initiated policies (provisioned or discovered by Crosswork)	—
PCC-initiated polices (discovered by Crosswork)	—
SR-TE On-Demand Next Hop (ODN) policies discovered by Crosswork	—
Single consistent Segment Routing Global Block (SRGB) configured on routers throughout domain covered by Crosswork	If index SIDs are used and there are different SRGB bases along a path of a policy, the label can change along the path.
Prefix SID	—
Adjacency SID	—
EPE adjacency SID	—
Protected and Unprotected adjacency SIDs	—
Regular and Strict prefix SIDs	—

Capability	Notes
SR-TE policy optimization objective min-metric (IGP, TE, and Latency)	—
SR-TE policy path constraints (affinity and disjointness)	Only 2 SR-TE policies per disjoint group or sub-id are supported
Binding SID for explicit or dynamic policies	—
Profile ID	—

Table 2: Unsupported Features and Limitations

Description	Notes
Provisioning multiple candidate paths via Crosswork	These paths are not discovered if configured on PCC. Crosswork does not support configuration of these paths.
Weighted Equal-Cost Multipath (WECMP)	—
Multiple segment lists per candidate path	• This configuration is not supported • These segment lists will not be discovered if configured on a PCC.
Visualization of multiple candidate paths	Only the current active path can be seen in the UI.
Binding SIDs as Segment List Hops	—
SR IGP Flexible Algorithm (Flex Algo)	—
Anycast SIDs	—
Hop count metric type for policies	Cisco Crosswork does not support provisioning with this metric type and does not discover this metric type if configured on the PCC
Routers that are not SR-capable	The assumption is that all routers discovered by Cisco Crosswork are SR-capable
SR-TE policies with Loopback IPs other than TE router ID for headend/endpoint and prefix SIDs in segment list	
SR-TE policy provisioned with IPv6 endpoints/hops	—
SRv6	Only 2 SR-TE policies per disjoint group/sub-id

Description	Notes
SR-TE policy optimization objective min-metric with margin	Not supported for policies provisioned by Cisco Crosswork. Margin is not discovered for PCC-initiated policies.
SR-TE policy constraints (resource exclusion or metric bound)	Not supported for policies provisioned by Cisco Crosswork. Constraints are not discovered for PCC-initiated policies.

SR Policy Configuration Sources

SR 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 Policy Example, on page 37).
- PCE initiated—Policies configured on a PCE or created dynamically by Crosswork Optimization Engine. An SR-TE policy that is configured using Crosswork Optimization Engine is the only type of SR-TE policy that Crosswork Optimization Engine can modify or delete.

PCC-Initiated SR Policy Example

The following example shows a configuration of an SR 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-TE Policies

This task creates SR-TE 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**.

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

Step 3 Enter the required SR-TE 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 **Explicit Path** and enter a path name.

Step 5 Add segments that will be part of the SR-TE policy path.

Step 6 Click **Preview**.

Step 7 If you want to commit the policy path, click **Provision**.

Step 8 Validate the SR-TE policy creation:

a. Confirm that the new SR-TE policy appears in the SR-TE 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-TE 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 <admin guide link>

Configure Link Affinities

Affinities defined on devices are not collected by Crosswork Optimization Engine. The affinity mapping name is only used for visualization in Crosswork Optimization Engine. For this reason, you should manually collect affinities on the device interface, then define affinity mapping in Crosswork Optimization Engine with the same name and bits that are used on the device interface. Crosswork Optimization Engine will only send bit information to SR-PCE during provisioning.

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.

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

a) Enter the name (color) and the bit it will be assigned to.

b) Click to save the mapping.

Step 3 To edit an affinity mapping, click .

a) Make the necessary changes. If you want to cancel your changes, click .

b) Click to save the changes.

Step 4 To delete an affinity mapping, click

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-TE Policies Based on Optimization Intent

This task creates an SR-TE policy with a dynamic path. SR-PCE computes a path for the policy based on metrics and path constraints (affinity 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.

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-TE policy. For more information, see Configure Link Affinities, on page 38.

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

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

Step 3 Enter the required SR-TE 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-TE policy. Also, there cannot be more than two SR-TE policies in the same disjoint group or subgroup. If there are existing SR-TE policies belonging to a disjoint group that you define here, all SR-TE 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 Click **Preview**. The path is highlighted on the map.

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

Step 10 Validate the SR-TE policy creation:

a. Confirm that the new SR 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 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 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 <admin guide link>

Modify SR-TE Policies

To view, edit, or delete a policy, do the following:

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

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

Step 3 Locate the SR policy you are interested in and click ⋯.

Step 4 Choose **View** or **Edit/Delete**.

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

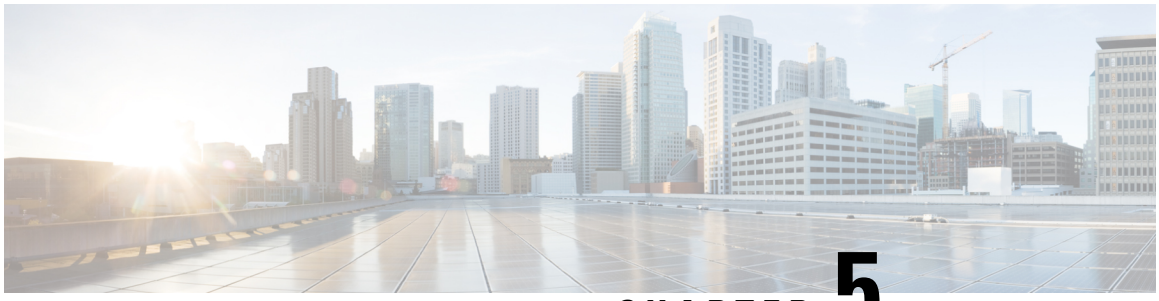

CHAPTER 5

Provision RSVP-TE Tunnels

This section contains the following topics:

RSVP-TE Tunnel Support

Table 3: Supported Features

Capability	Notes
PCE-initiated tunnels (provisioned or discovered by Crosswork Optimization Engine)	—
PCC-initiated tunnels (discovered by Crosswork Optimization Engine)	—
ERO strict hops	—
ERO loose hops (PCC-initiated only)	—
FRR protection on tunnels provisioned by Crosswork Optimization Engine	—
Path optimization objective min-metric (IGP,TE, or Latency)	—
Path constraints (affinity and disjointness)	Only 2 RSVP tunnels per disjoint group or sub-id is supported
Binding Label for explicit and dynamic tunnels	—
Signaled Bandwidth	—
Setup/Hold Priority	—

Table 4: Unsupported Features and Limitations

Description	Notes
Configuring loose hop ERO in COE	Only strict hops can be configured. If strict hops are not configured for every hop along the path and those hops are not remote interface IPs or loopback IPs, unexpected behavior may occur. For example, a tunnel may remain operationally down, hops may be modified, and so on.
Named tunnels configured on PCCs	These tunnels are not discovered byCrosswork Optimization Engine.
Tunnels with Loopback IPs other than TE router ID for headend or endpoint and path hops	—
Display of active FRR protected paths in the topology map.	Crosswork Optimization Engine discovers FRR tunnels which are displayed in the topology map, but will not associate an actively protected tunnel with the FRR tunnel being used. The path in the topology map will not include FRR protected paths when protection is active.
P2MP tunnels	—

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 42).
- Dynamically created.

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 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 Policy Path, click **Explicit Path** and enter a path name.

Step 6 Add segments that will be part of the RSVP-TE path.

Step 7 Click **Preview**. The path is highlighted on the map.

Step 8 If you want to commit the tunnel path, click **Provision**.

Step 9 Validate the RSVP-TE tunnel creation:

a. Confirm that the new RSVP-TE tunnel appears in the RSVP-TE Tunnels table. You can also click the check box next to the policy to see it highlighted in the map.

Note The newly provisioned RSVP-TE tunnel may take some time, depending on the network size and performance, to appear in the **RSVP-TE Tunnels** table. The **RSVP-TE Tunnels** table is refreshed every 30 seconds.

b. View and confirm the new RSVP-TE tunnel details. From the **RSVP-TE** table, click ··· (in the same row as the RSVP-TE tunnel), and select **View**.

Note On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. Please contact a Cisco representative to fine tune the timers involved.

Configure Link Affinities

Affinities defined on devices are not collected by Crosswork Optimization Engine. The affinity mapping name is only used for visualization in Crosswork Optimization Engine. For this reason, you should manually collect affinities on the device interface, then define affinity mapping in Crosswork Optimization Engine with the same name and bits that are used on the device interface. Crosswork Optimization Engine will only send bit information to SR-PCE during provisioning.

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.

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

a) Enter the name (color) and the bit it will be assigned to.

b) Click to save the mapping.

Step 3 To edit an affinity mapping, click .

a) Make the necessary changes. If you want to cancel your changes, click .

b) Click to save the changes.

Step 4 To delete an affinity mapping, click

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 38.

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, edit, 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 delete RSVP-TE tunnels that have been created with the UI.
- After updating the RSVP-TE tunnel details, you can preview the changes on the map before saving it.

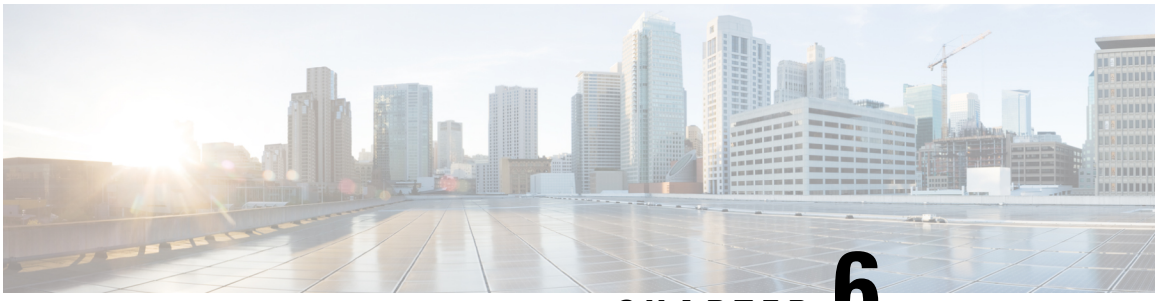

CHAPTER 6

Mitigate Network Congestion

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

Cisco Crosswork can proactively monitor network bandwidth utilization and mitigate congestion to help alleviate the difficult task of tracking and reacting to traffic utilization changes that go above a specified threshold by using one of two tools: Local Congestion Mitigation and Bandwidth Optimization.

Bandwidth Optimization (BWOpt) provides closed-loop traffic engineering by automatically rerouting intent-based traffic dynamically throughout the network in response to congestion. For more information, see Use BWOpt to Optimize the Network, on page 57

Local Congestion Mitigation (LCM) searches for congestion on a configurable cadence (as opposed to a triggered event) and provides localized mitigation recommendations within surrounding interfaces (local interface-level optimization). You are able to visually preview these recommendations on your network before you decide whether to commit the Tactical Traffic Engineering (TTE) SR policy deployment. LCM performs the collection of TTE SR policy and interface counters via SNMP and does not require the use of SR-TM. For more information, see Use LCM to Mitigate Congestion Locally, on page 47

Note 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 eliminates the need for simulating edge-to-edge traffic flows in the network through a full traffic matrix.

Use LCM to Mitigate Congestion Locally

Local Congestion Mitigation (LCM) checks the capacity locally, in and around the congested area, at an interface level. 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.

TTE tunnel recommendations are listed in the **LCM Operational Dashboard**. From the dashboard, you can visually preview the TTE SR policy recommendations before deployment. TTE SR policy deployment to resolve congestion is not automated. You must approve and commit LCM recommended actions. LCM also recommends removal of previous TTE SR policies (instantiated by LCM) if they are no longer needed.

LCM Important Notes

Consider the following information when using LCM:

- LCM evaluates network utilization on a regular, configurable cadence of 10 minutes or more. The cadence is typically set to be greater than or equal to the SNMP traffic polling interval.
- 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.
- Traffic that can be optimized must not be carried on existing SR-TE policies.

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:

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

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

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

LCM Calculation Workflow

This example walks you from congestion detection to the calculations LCM performs prior to recommending tactical tunnel deployment.

Figure 20: LCM Configuration Workflow Example

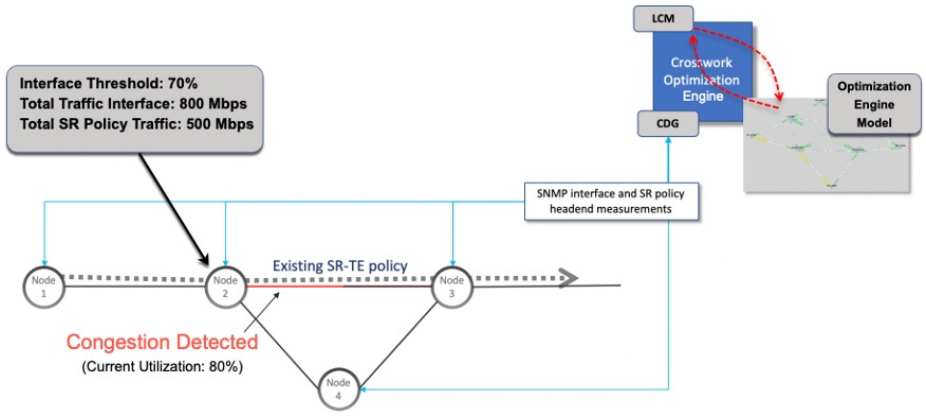

Step 1 LCM first analyzes the Optimization Engine Model (a realtime topology and traffic representation of the physical network) on a regular cadence.

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 estimates how much traffic is eligible to divert.

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

Eligible traffic is computed by taking the interface traffic stats that account for all traffic on the interface and subtracting the sum of traffic stats 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.

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 1/3 of the total eligible traffic (100Mbps out of 300Mbps) away from congested link. Assuming a perfect ECMP, LCM estimates 3x tactical SR-TE policies in total to create this traffic split: 1 tactical SR-TE policy will take the diversion path and 2 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 3 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**.

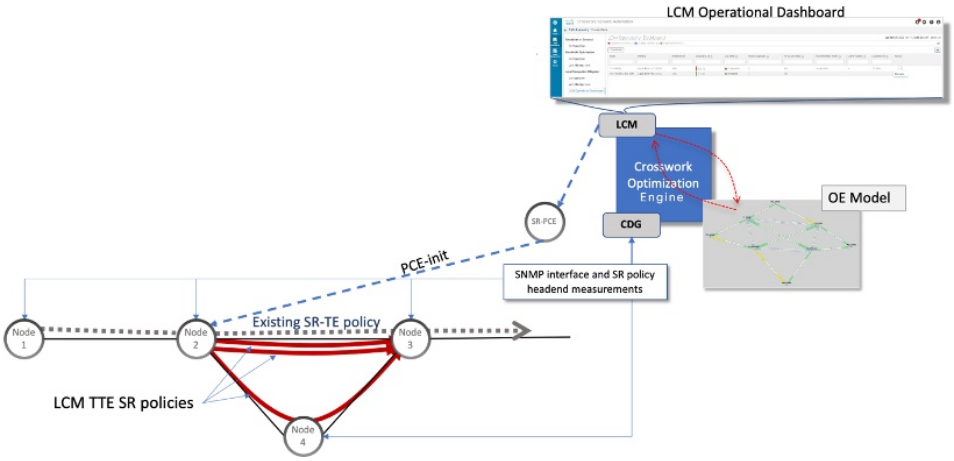

Step 6 Assuming you deploy these TTE SR policies, LCM continues to monitor the deployed TTE policies and will recommend modifications or deletions as needed in the **LCM Operational Dashboard**. TTE SR policy removal recommendations will occur if the mitigated interface would not be congested if these policies were removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Mitigate Congestion on Local Interfaces Example

In this example, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE SR polices when utilization surpasses a defined threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion. The following image shows the initial topology before congestion occurs.

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

a) Click on the link between PE1-ASR9k and P1-ASR9k to view link details. Note that there is currently no congestion (0% utlization).

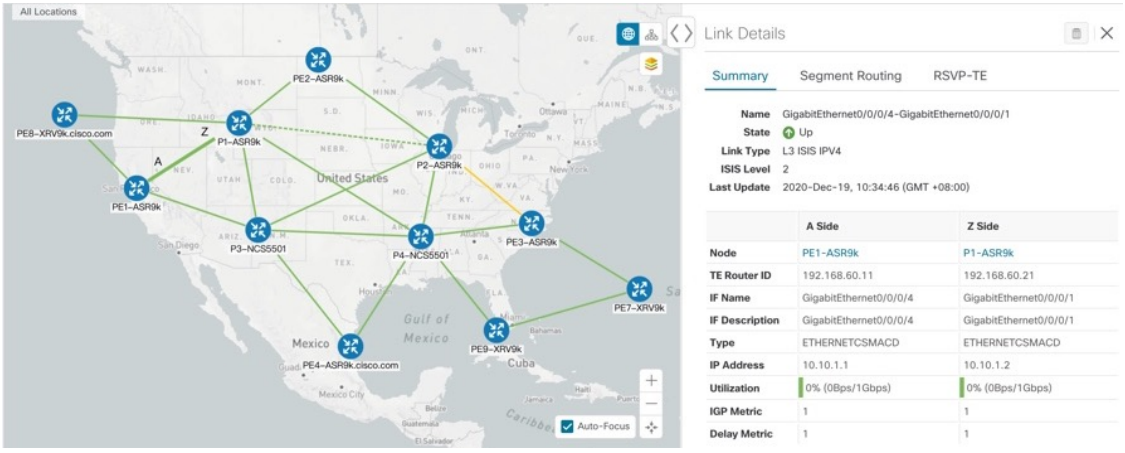

Step 2 Enable LCM and configure the global utilization threshold.

a) From the main menu, choose **Traffic Engineering > Local Congestion Mitigation** > **Configuration**. In this case, the threshold is set at 25%.

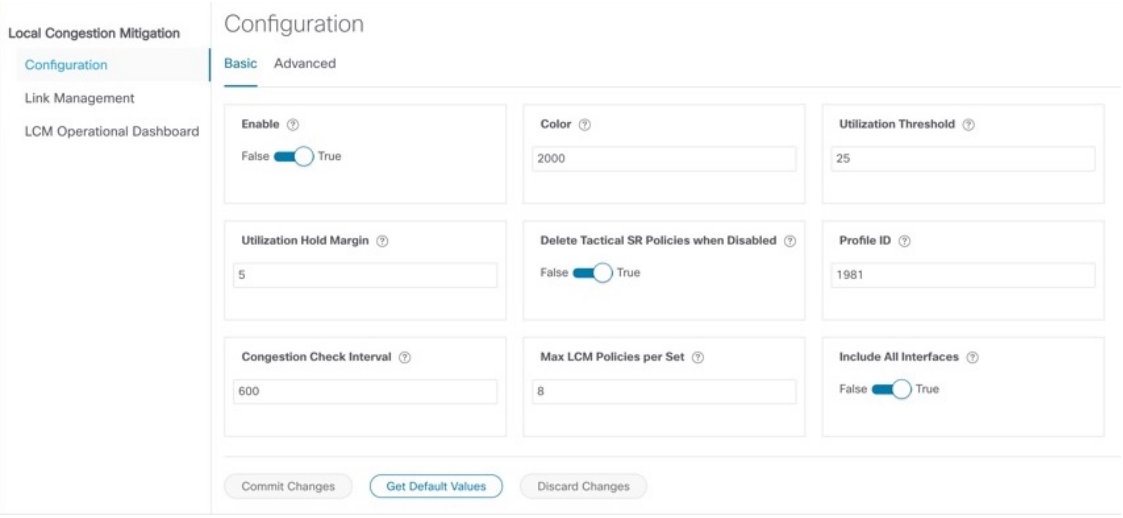

If you want to set separate threshold for individual interfaces, toggle the Include All Interfaces to **False**.

b) (optional) Define any specific thresholds for individual links by uploading a CSV file (**Traffic Engineering > Local Congestion Mitigation>Link Management**).

Note A sample CSV template is available for download.

Step 3 View TTE SR policy recommendations in the LCM Dashboard.

a) After some time, congestion occurs surpassing the configured LCM threshold. Note that the link is Orange, indicating higher utilization.

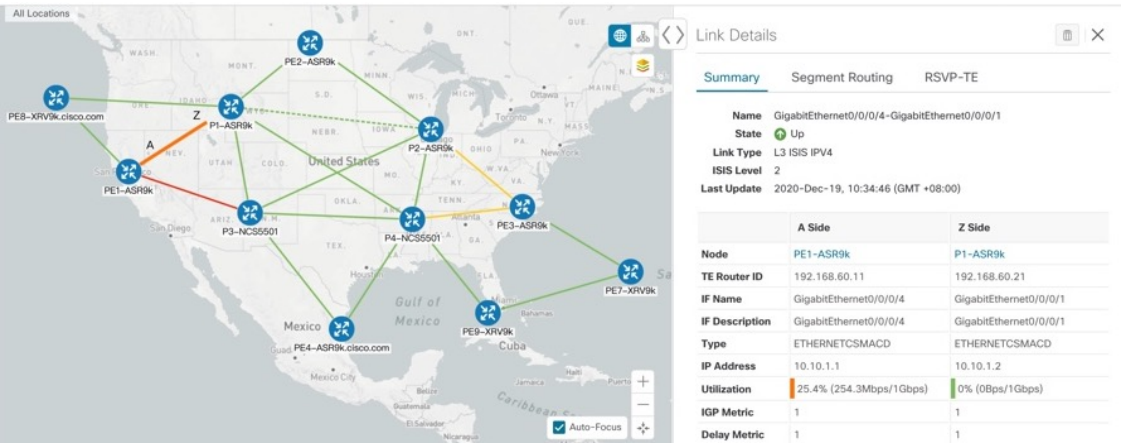

b) Click to view the new event. You can also monitor this window to view LCM events as they occur. You should see events for LCM recommendations, commit actions, and any exceptions.

c) Open the **LCM Operational Dashboard** (**Traffic Engineering > Local Congestion Mitigation> LCM Operational Dashboard**).

The dashboard shows that the utilization has surpassed 25%. In the Recommended Action colum, there is a recommendation to deploy 2 TTE policy solution sets to address the congestion on each interface. The Expected Util column shows the expected utilization of each of the interface if the recommended action is committed.

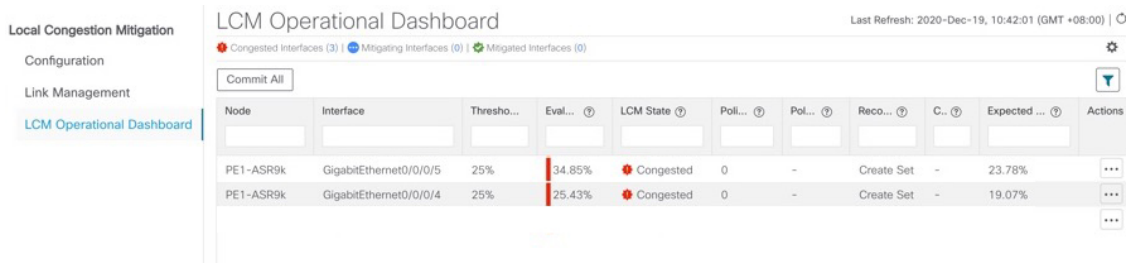

d) To preview the TTE deployment of each TTE policy solution set, click ⋯ and choose **Preview**. The window displays the node, interface, and the recommended action for each TTE policy. The following figure shows the recommended TTE policies for the interface GigabitEthernet0/0/0/4.

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.

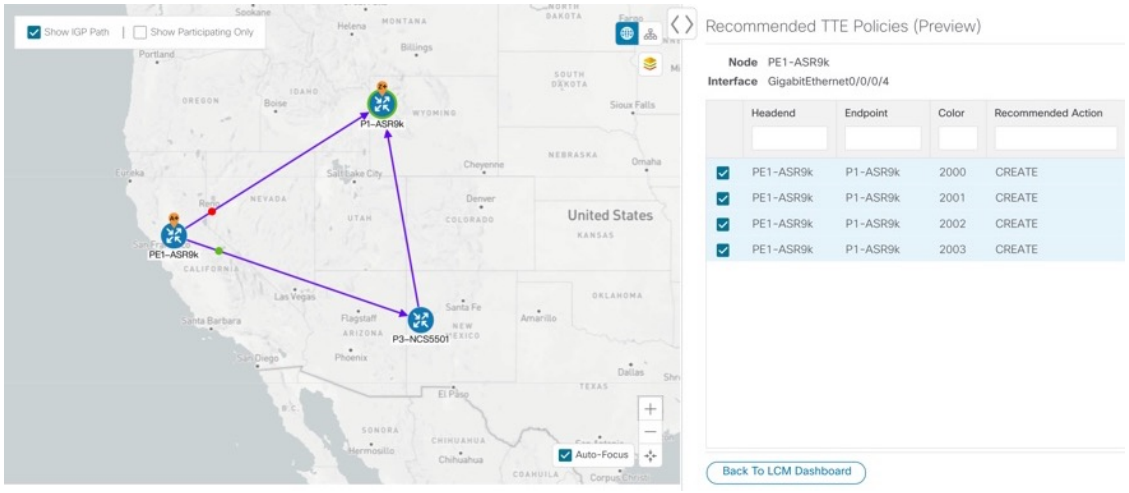

e) If you are satisfied with the LCM recommendations, click **Commit All**. The LCM **Status** column changes to **Mitigating**.

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

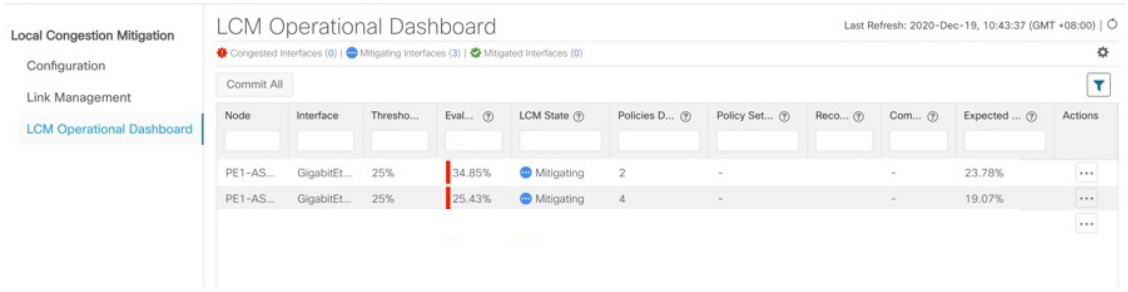

Step 4 Validate TTE SR policy deployments.

a) Click to open the **Events** window and note which LCM events are listed in this window.

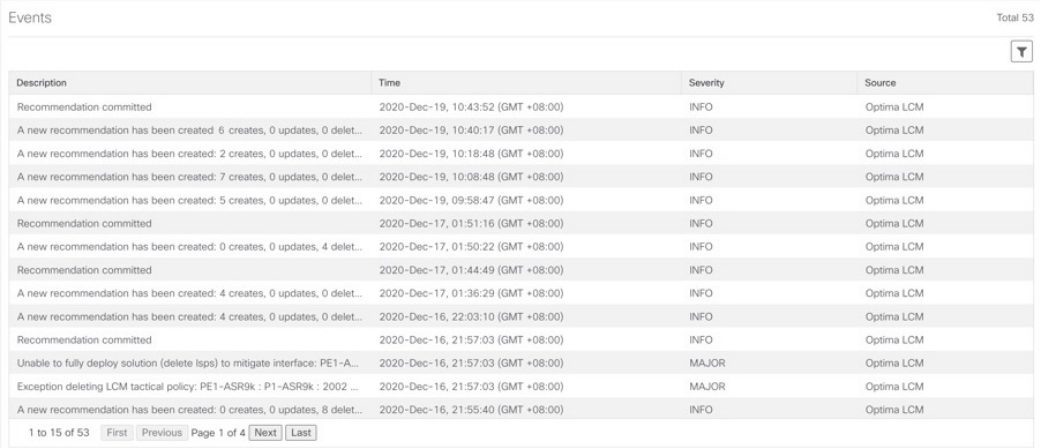

b) Return to the LCM Dashboard to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

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

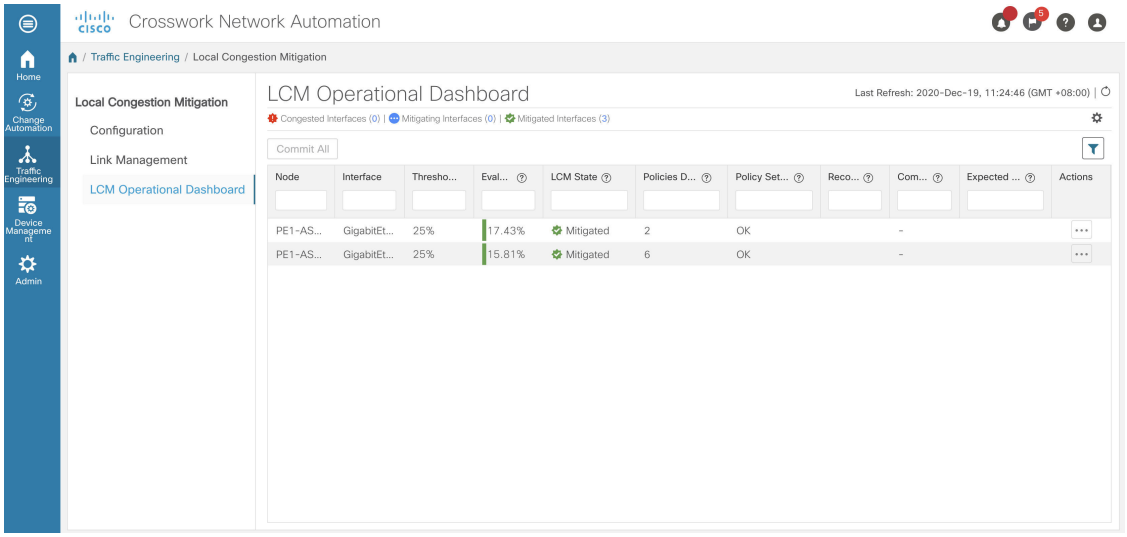

c) Confirm the TTE policy deployment by viewing the topology map and the **SR Policy** table (**Traffic Engineering** > **Traffic Engineering** > **SR-TE** tab).

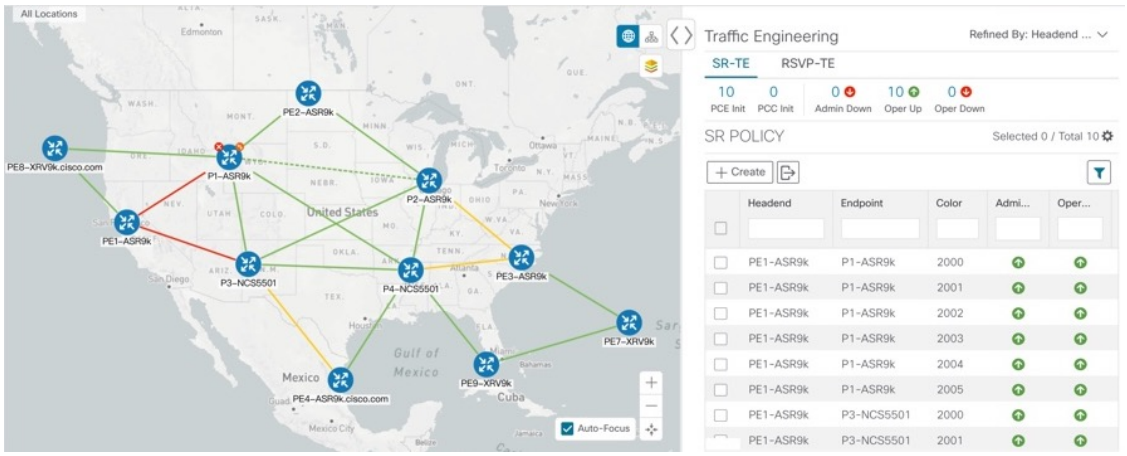

Tip To help narrow the search for the SR-TE policies that were just deployed, click ⚙ from the SR Policy table and click the checkbox to include **Policy Type**. Then filter the policy type as **Local Congestion Mitigation**. While it shows all SR-TE policies of this type, the SR-TE policy list should be easier to sort through.

d) Select one of the new SR-TE policies and view the SR policy details (click ⋯ and choose **View**).

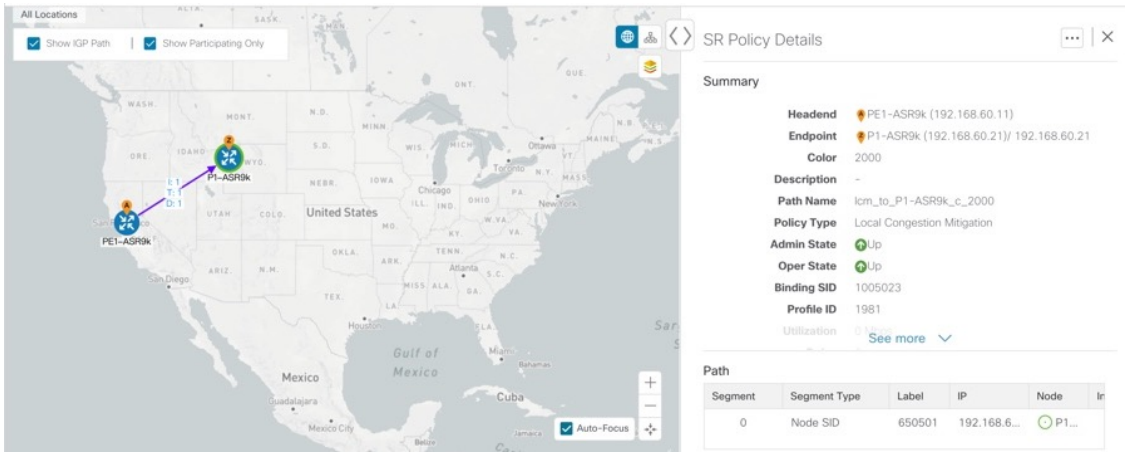

Step 5 Remove the TTE SR policies upon LCM recommendation.

a) After some time, the deployed TTE SR policies might no longer be needed. This occurs if the utilization will continue to be under the threshold without the LCM-initiated TTE tunnels. In this case, LCM generates new recommended actions to delete the TTE SR policy sets. Click **Commit All** to remove the deployed TTE SR policies.

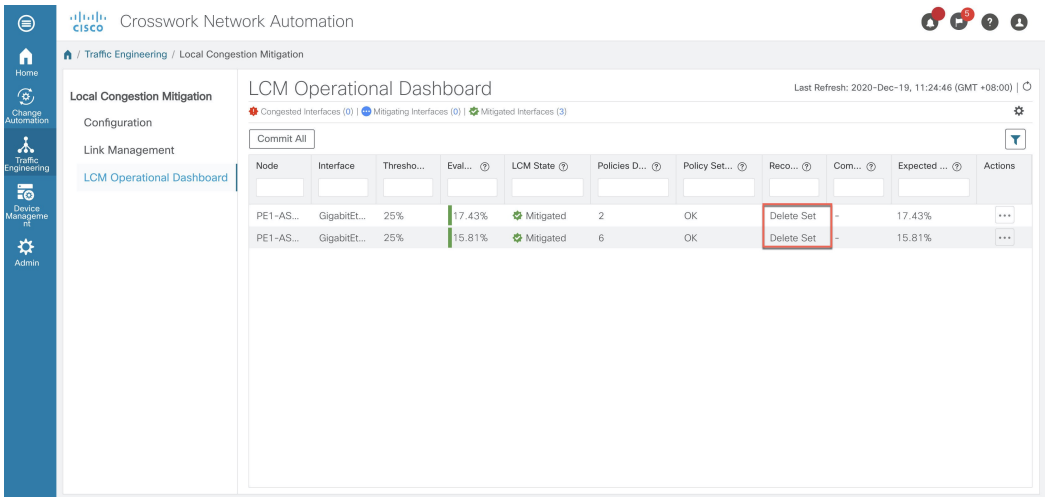

b) Click **Commit All** to remove the SR policies.

c) Confirm the removal by viewing the topology map and SR Policy table.

Related Topics

Add Individual Interface Thresholds, on page 62

Configure LCM

To enable and configure LCM:

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

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.

The following list describes additional field information:

- **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 (5 minutes), LCM will evaluate the network every 5 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. Since network changes may take time for the information to stabilize and propagate to LCM, set the interval to no less than twice the SNMP collection cadence.
- **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.

Step 4 Click **Commit Changes**.

Monitor LCM Operations

View the LCM Dashboard (**Traffic Engineering > Local Congestion Mitigation** > **LCM Operational Dashboard**) to monitor LCM operations. The LCM 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. Hover the mouse pointer over ⓘ to view a description of what type of information each column provides. From this dashboard, you can also preview and deploy TTE policy recommendations.

In addition to the LCM Operational Dashboard, you can click to view LCM events.

Use BWOpt to Optimize the Network

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:

- 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.
- BWOpt uses simulated traffic based on measured SRTM 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 21: Bandwidth Optimization Configuration

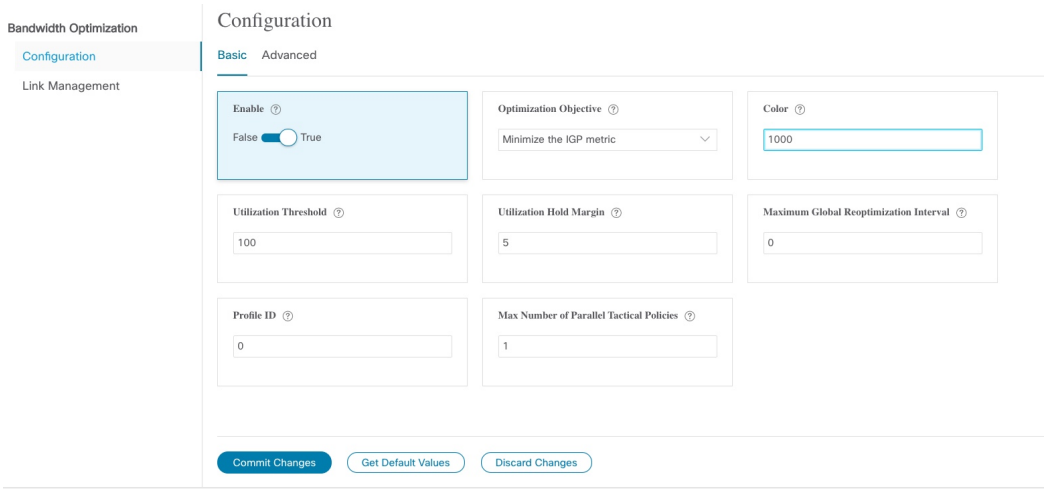

Below is a network with various devices and links that span the United States. Note that there are no SR-TE policies listed in the **SR Policies** table.

Figure 22: Example: Current Network

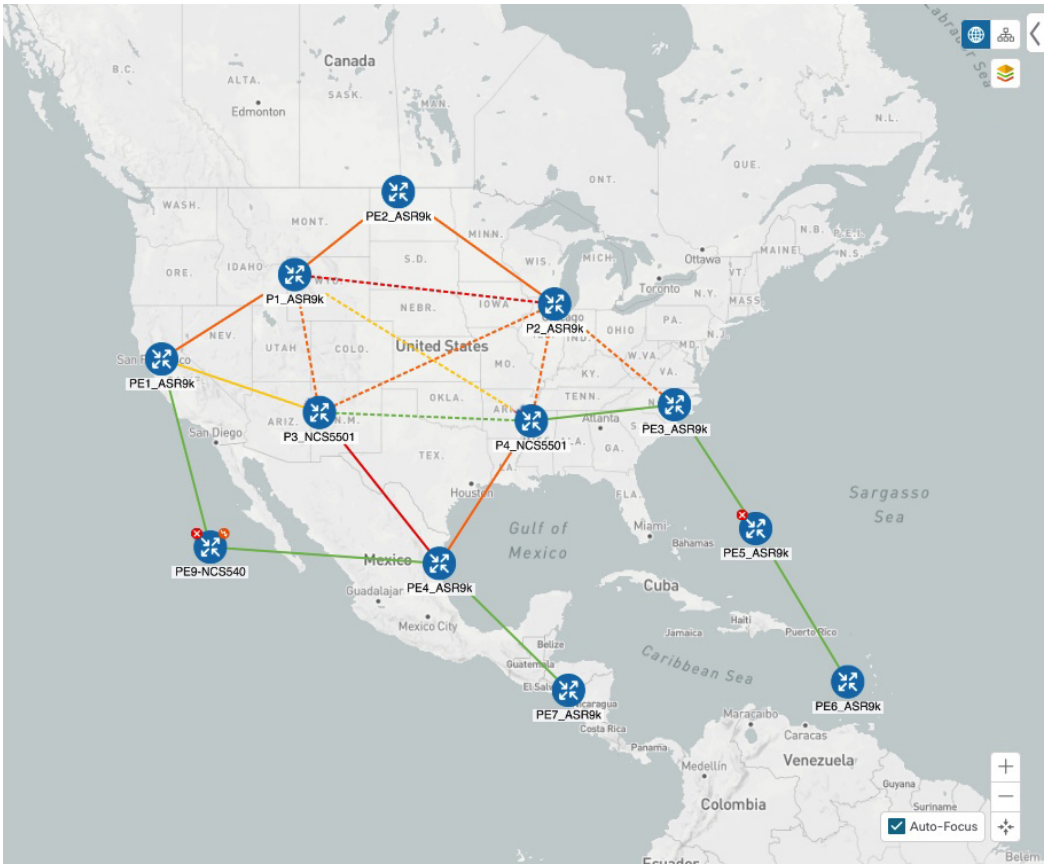

Suppose the link between P3_NCS5501 and P4_NCS5501 goes down. Traffic moves towards other links causing congestion and exceeds the configured utilization threshold.

Figure 23: Example: Link Down Between P3 and P4 Nodes

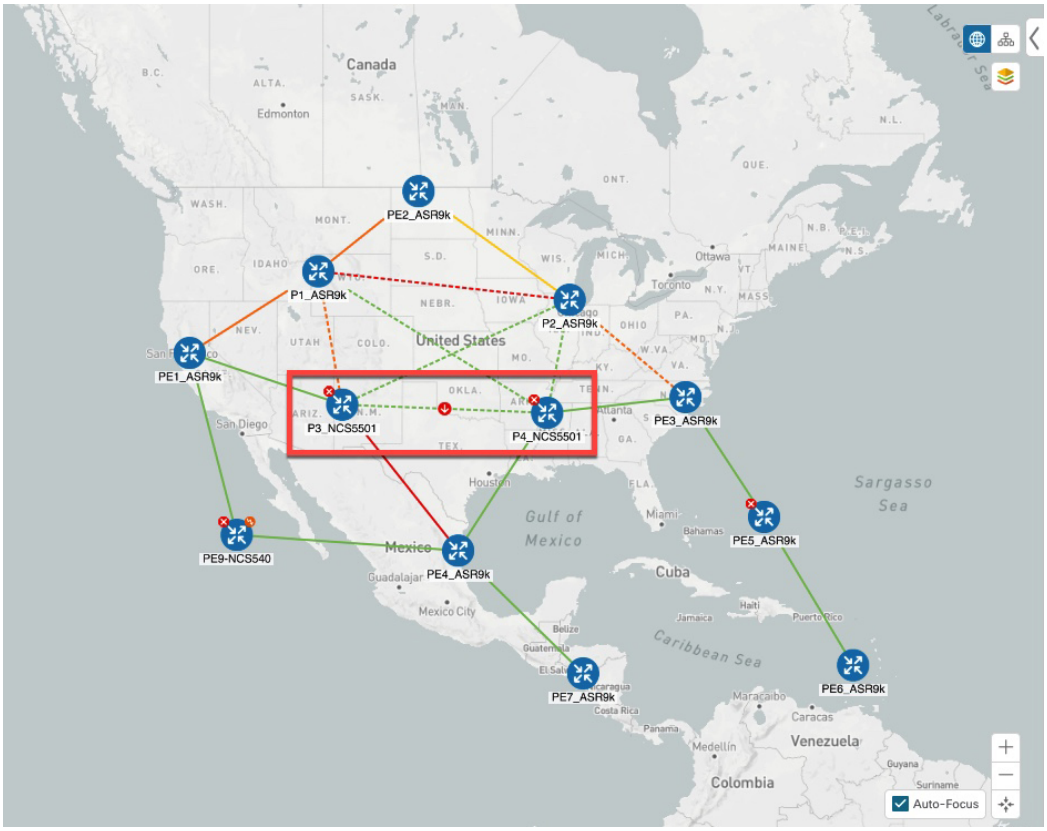

BWOpt recognizes the congestion and immediately calculates and deploys a tactical SR-TE policy. This new tactical SR-TE policy is listed in the **SR Policies** window.

Figure 24: Example: Tactical SR Policy Deployed

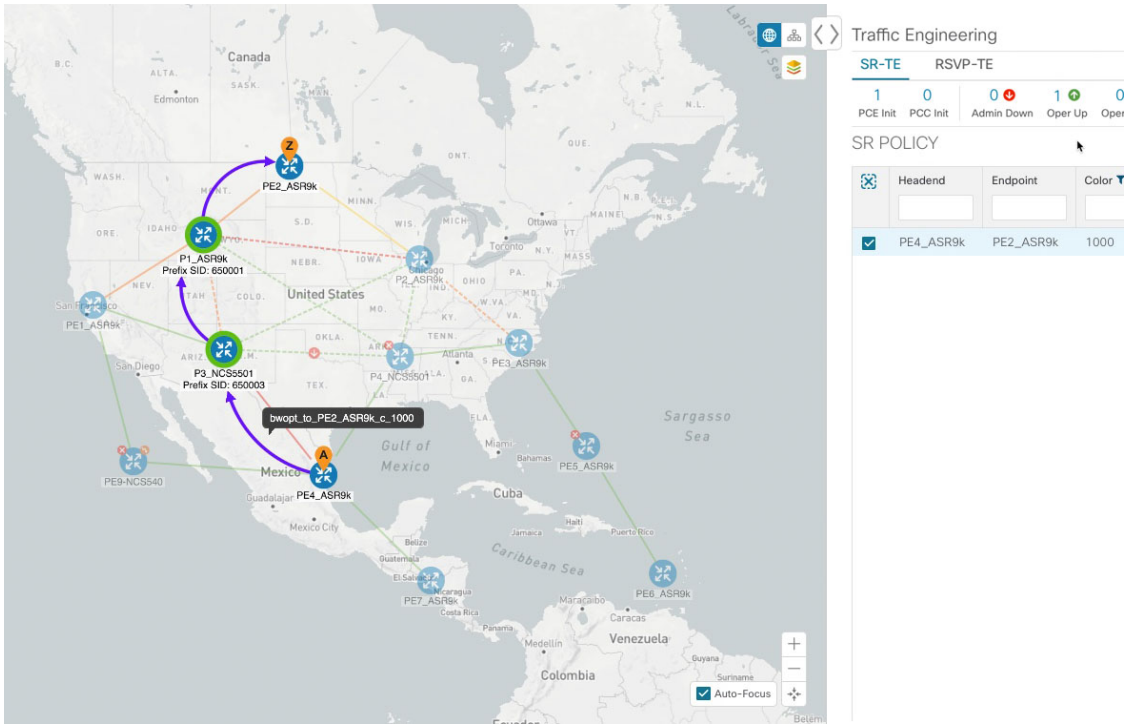

BWOpt continually monitors the network. When the links between P3_NCS5501 and P4_NCS5501 are back up, BWOpt will detect that the congestion (based on the defined criteria) has been mitigated. When the congestion falls under the set utilization threshold minus the utilization hold margin, the tactical 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 Bandwith 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.

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.

Table 5: 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.

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 LCM or Bandwidth Optimization, do the following:

Step 1 From the main menu, choose one of the following:

- **Local Congestion Mitigation** > **Link Management**
- **Bandwidth Optimization** > **Link Management**

Step 2 Click .

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

Step 4 Click **Cancel**.

Step 5 Open and edit the configuration file (sampleLcmLinkManagement.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 **Link Management** window.

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

Step 9 Click **Import**.

Step 10 Confirm that the information appears correctly in the **Link Management** window.

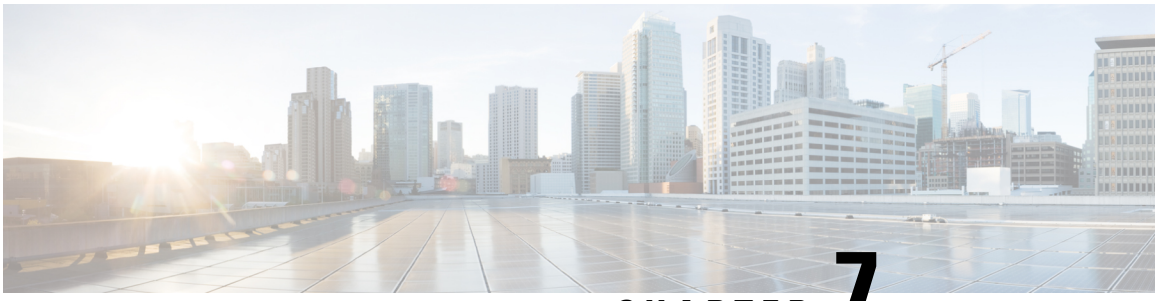

CHAPTER 7

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-init and PCC-init 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

Consider the following information when using 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.

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

Figure 25: Initial BWoD Topology Example

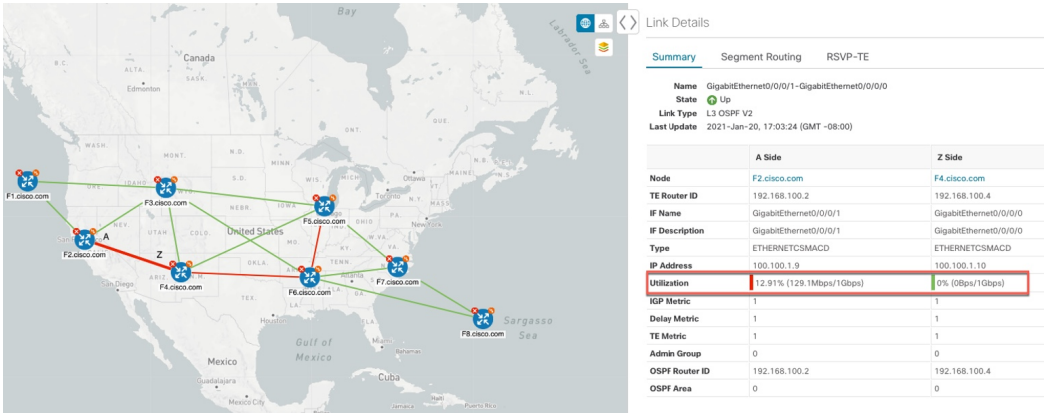

In this scenario we are using the above topology. The goal is to create a path from F2.cisco.com to F7.cisco.com that can accommodate 920 Mbps of traffic while keeping the utilization at 80%. The above example highlights the utilization on nodes F2.cisco.com and node F4.cisco.com to show that the link 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.

a) From the main menu, choose **Traffic Engineering** > **Bandwidth on Demand** > **Configuration**.
b) 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 .
c) Click **Commit Changes**.

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

a) From the main menu, choose **Traffic Engineering** > **SR-TE** tab and click **+Create**.
b) Enter the required SR-TE policy details.
c) In the **Policy Path** field, click **Bandwidth on Demand** and enter a unique name for the BWoD path. In this case, **bwOD-Policy-new**.
d) From the **Optimization Objective** drop-down list, select **Traffic Engineering (TE) Metric**.
e) 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.

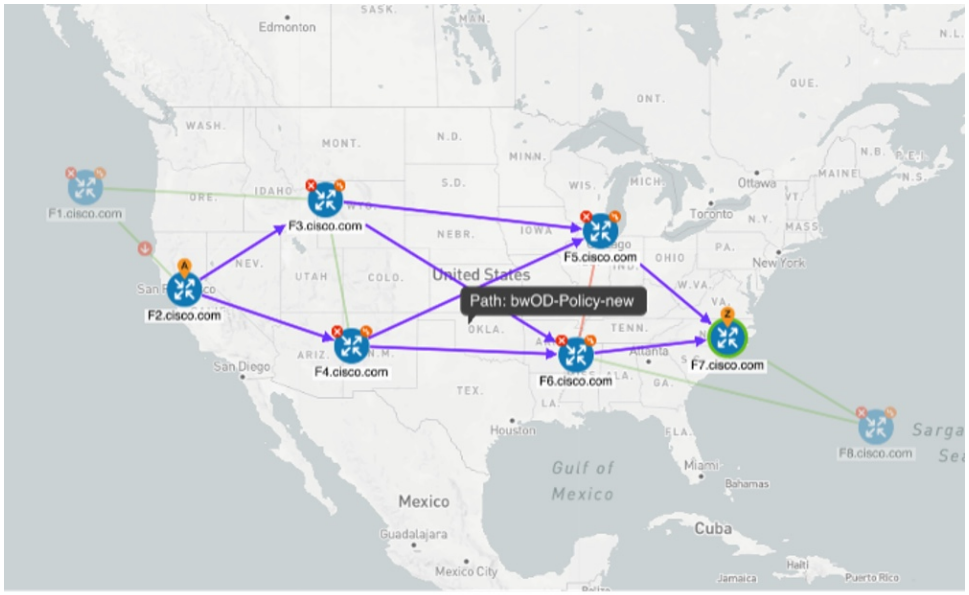

In the above example, BWoD cannot find a single path because of utilization and capacity limitations across several links. In this case, BWoD splits the path to obtain bandwidth and utilization requirements.

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.

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

b) Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View**). 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 26: PCC-Initiated BWoD SR-TE Policies

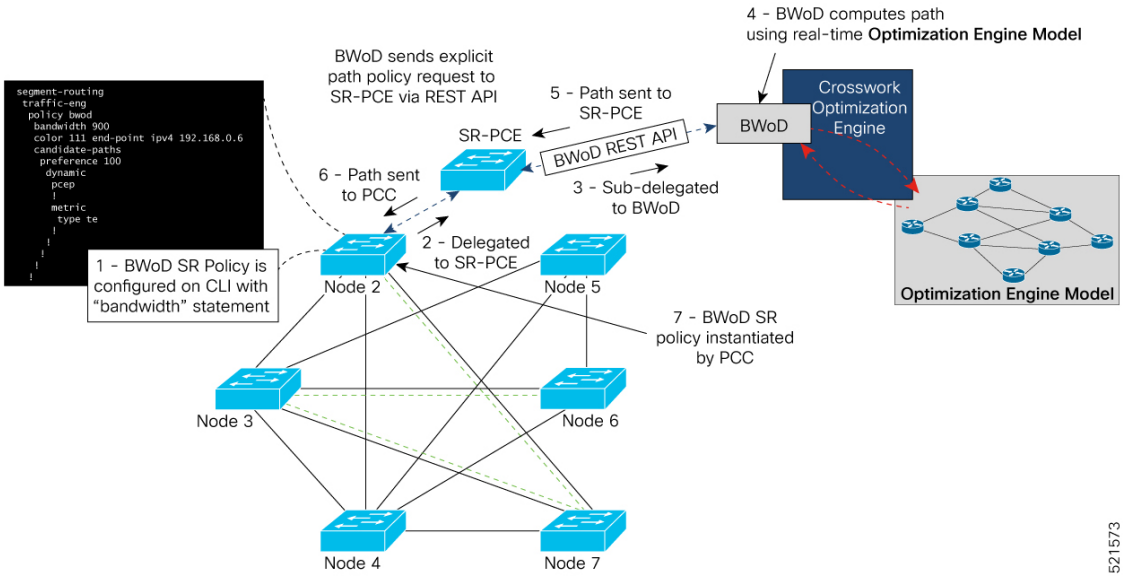

Callout No.	Description

1	A BWoD policy is configured on a PCC via the CLI. For example:
	`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`
	`       !`
	`      !`
	`     !`
	`    !`
	`   !`

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 6: Errors

Error Event Message	Possible Causes and Recommended Corrective Action
OptimaModelError	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. 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.
NATSTimedOutError	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.
Traceback or other errors found in the log file	Please contact your Cisco service representative.