Cisco Crosswork Network Controller 7.0 Traffic Engineering and Optimization

First Published: 2024-03-21

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

Traffic Engineering in Cisco Crosswork Network Controller

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

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

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

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

This section contains the following topics:

Supported SR-TE Policies and RSVP Tunnels

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

Table 1: Supported TE Technologies

TE Technology	Crosswork Network Controller	
	Visualize	**Provision**
SR-MPLS	✓	✓
SRv6	✓	✓
RSVP	✓	✓
Flexible Algorithm	✓	✗
Tree-SID	✓	✓
Circuit Style	✓	✓

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

For a list of known limitations, important notes, and what networking technologies are supported, see the Cisco Crosswork Network Controller Release Notes.

What is Segment Routing?

Segment routing for traffic engineering takes place through a tunnel between a source and destination pair. Segment routing for traffic engineering uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment. Segments are an identifier for any type of instruction. For example, topology segments identify the next hop toward a destination. Each segment is identified by the segment ID (SID) consisting of an unsigned 32-bit integer. Each segment is an end-to-end path from the source to the destination and instructs the routers in the provider core network to follow the specified path instead of the shortest path calculated by the IGP. The destination is unaware of the presence of the tunnel.

Segments

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

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

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

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

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

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

Segment Routing Policies

Segment routing for traffic engineering uses a "policy" to steer traffic through the network. An SR policy path is expressed as a list of segments that specifies the path, called a segment ID (SID) list. Each segment is an end-to-end path from the source to the destination, instructing the network routers to follow the specified path instead of the shortest path calculated by the IGP. If a packet is steered into an SR policy, the head-end pushes the SID list on the packet. The rest of the network executes the instructions embedded in the SID list.

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

- SR-MPLS and SRv6 , on page 15
- Flexible Algorithm, on page 43
- Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering, on page 49
- SR Circuit Style

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

Dynamic SR Policy

A dynamic path is based on an optimization objective and a set of constraints. The head-end computes a solution, resulting in a SID list or a set of SID lists. When the topology changes, a new path is computed. If the head-end does not have enough information about the topology, the head-end might delegate the computation to a path computation engine (PCE).

Explicit SR Policy

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

Disjointness

Crosswork uses the disjoint policy to compute two lists of segments that steer traffic from two source nodes to two destination nodes along disjoint paths. The disjoint paths can originate from the same head-end or different head-ends. Disjoint level refers to the type of resources that should not be shared by the two computed paths. The following disjoint path computations are supported:

- **Link** – Specifies that links are not shared on the computed paths.
- **Node** – Specifies that nodes are not shared on the computed paths.
- **SRLG** – Specifies that links with the same Share Risk Link Group (SRLG) value are not shared on the computed paths.

- **SRLG-node** – Specifies that SRLG and nodes are not shared on the computed paths.

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

Note

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

Segment Routing Path Computation Element (SR-PCE)

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

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

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

Note Certain features may not function as expected if the SR-PCE version is not supported. To avoid any compatibility issues, refer to the Cisco Crosswork Network Controller Release Note for SR-PCE version support and compatibility.

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

SR-TE Policy PCC and PCE Configuration Sources

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

-
- Path Computation Element (PCE) initiated—Policies configured on a PCE or created dynamically by Crosswork. Examples of PCE Initiated policy types:
  - **Dynamic**
  - **Explicit**

- **Bandwidth on Demand** (can be either PCC or PCE)
- **Local Congestion Mitigation**

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

PCC-Initiated SR-TE Policy Example

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

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

What is Resource Reservation Protocol (RSVP)?

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

The RSVP-TE process contains the following functionalities:

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

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

RSVP-TE Explicit Routing (Strict, Loose)

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

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

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

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

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

RSVP FRR

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

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

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

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

RSVP-TE Tunnel PCC and PCE Configuration Sources

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

- Path Computation Client (PCC) initiated—RSVP-TE tunnels configured on a PCC (see PCC-Initiated RSVP-TE Tunnel Example, on page 7).
- Path Computation Element (PCE) or PCC initiated dynamically.

PCC-Initiated RSVP-TE Tunnel Example

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

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

Get a Quick View of Traffic Engineering Services

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

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

Figure 1: Quick View of Traffic Engineering Services

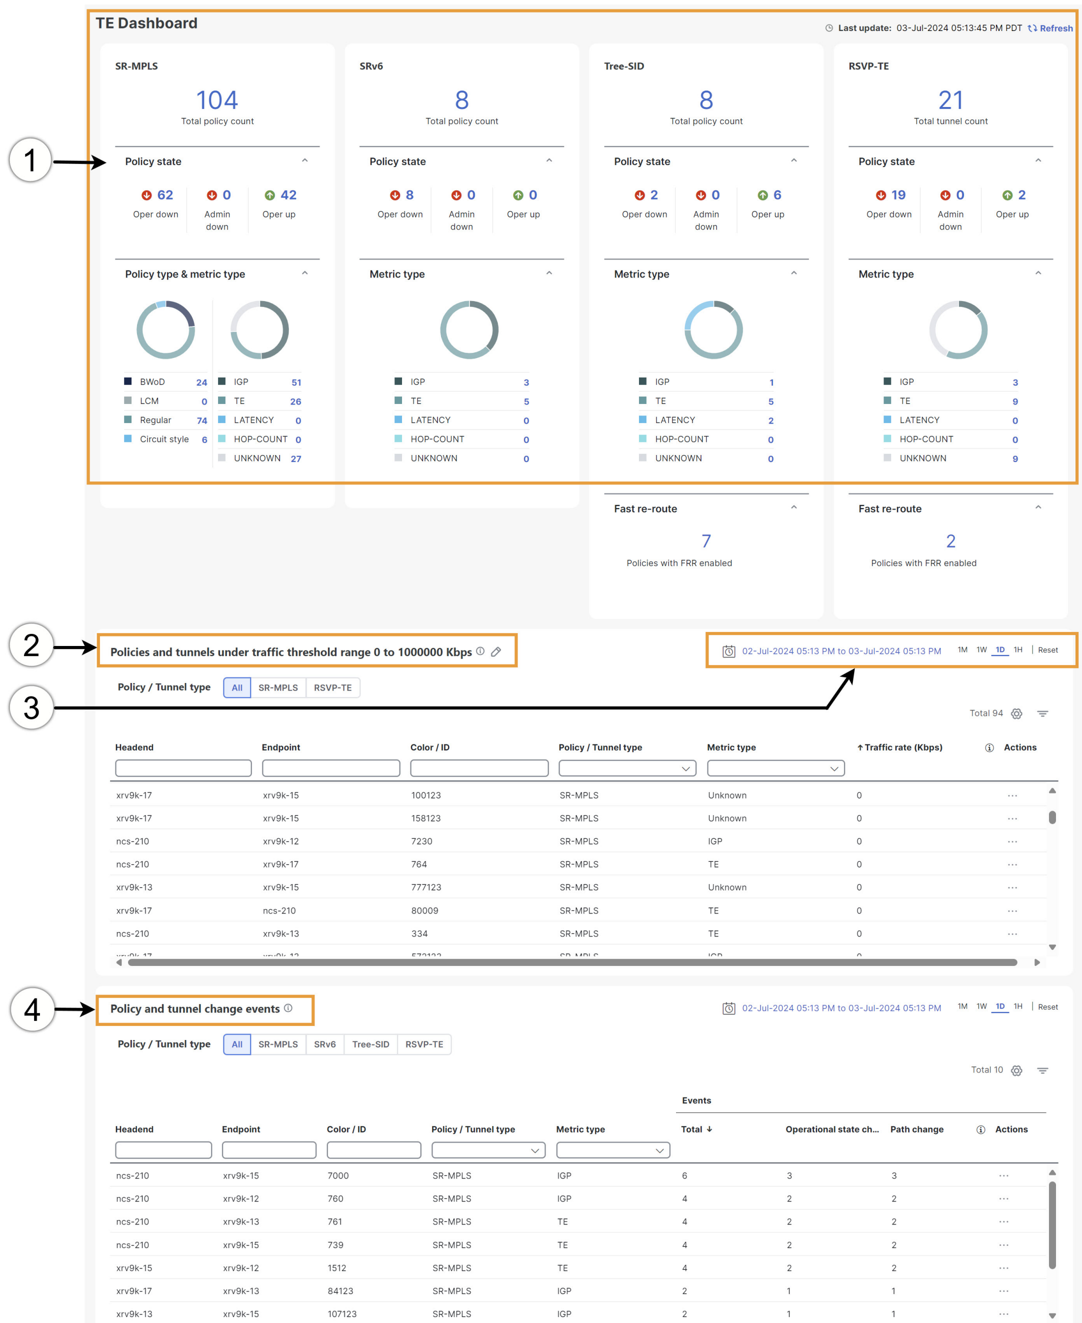

Note If you are viewing the HTML version of this guide, click the images to view them in full-size.

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

Note For a list of known limitations, see the Cisco Crosswork Network Controller Release Notes.

View TE Event and Utilization History

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

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

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

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

Figure 2: TE Event and Utilization History

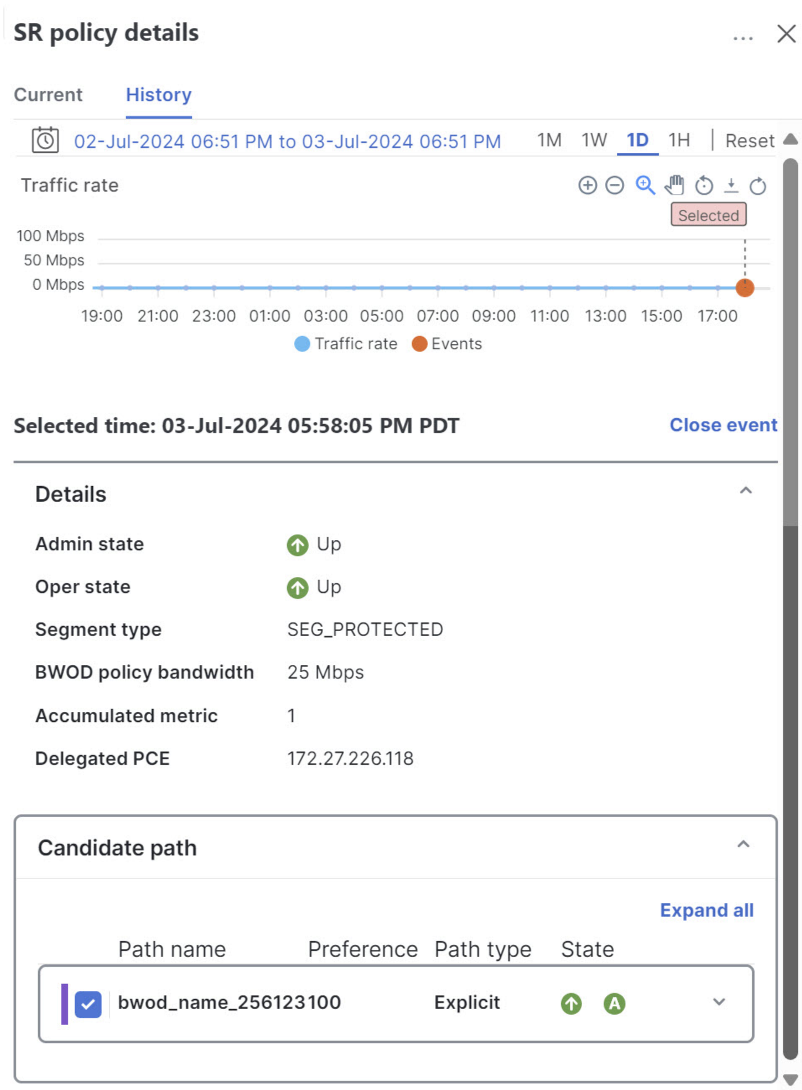

Additional Delay Data

When Crosswork Service Health is installed, Delay (avg) and Delay variance information is available. For more information, see "Enable SR PM Monitoring for Links and TE Policies," in the Cisco Crosswork Network Controller Service Health Monitoring Guide.

The extended TE link delay metric (minimum-delay value) can be used to compute paths for SR policies as an optimization metric or as an accumulated delay bound.

This can be used to monitor the end-to-end delay experienced by the traffic sent over an SR policy to ensure that the delay does not exceed the requested “upper-bound” and violate SLAs. You can verify the end-to-end delay values before activating the candidate-path or the segment lists of the SR policy in forwarding table, or to deactivate the active candidate-path or the segment lists of the SR policy in forwarding table.

Figure 3: Example of VPN Service when Monitoring is Enabled

View Traffic Engineering Device Details

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

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

Step 2 From the Traffic Engineering topology map, click on a device.

Step 3 From the **Traffic Engineering** tab, click on the policy type you are interested in. Each tab displays associated data for that device. From the browser, you can copy the URL and share with others.

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

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

Figure 4: Traffic Engineering Device Details

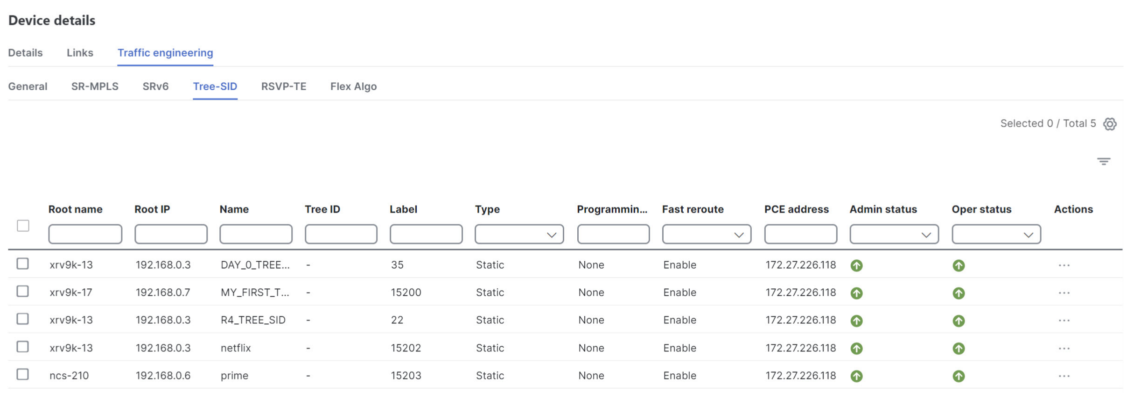

Configure Traffic Engineering Settings

Configure TE Timeout Settings

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

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

Figure 5: Traffic Engineering Timeout Settings

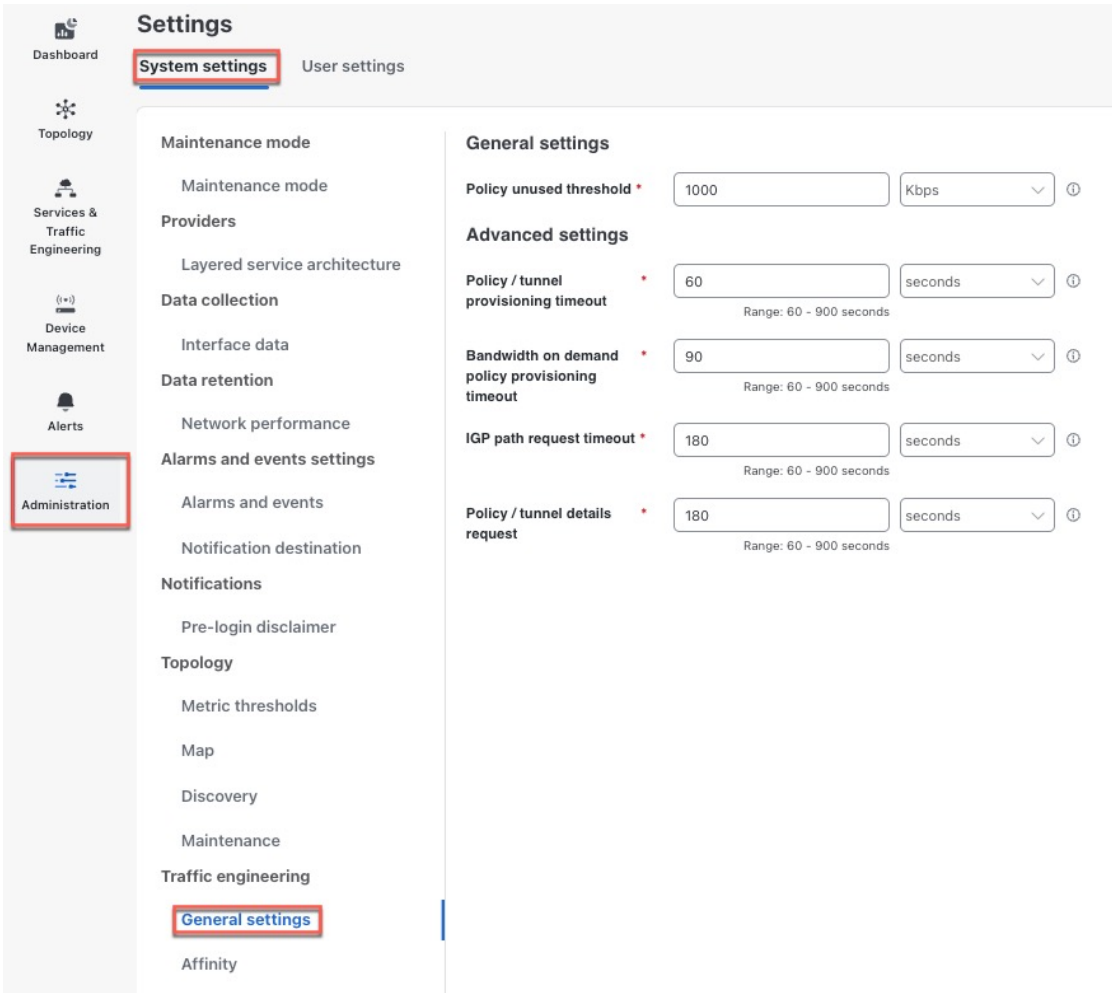

Configure How Device Groups Are Displayed for Traffic Engineering

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

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

Configure TE Data Retention Settings

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

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

Resolve SR-TE Policies and RSVP-TE Tunnels

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

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

CHAPTER 2

SR-MPLS and SRv6

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

View SR-MPLS and SRv6 Policies on the Topology Map

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

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

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

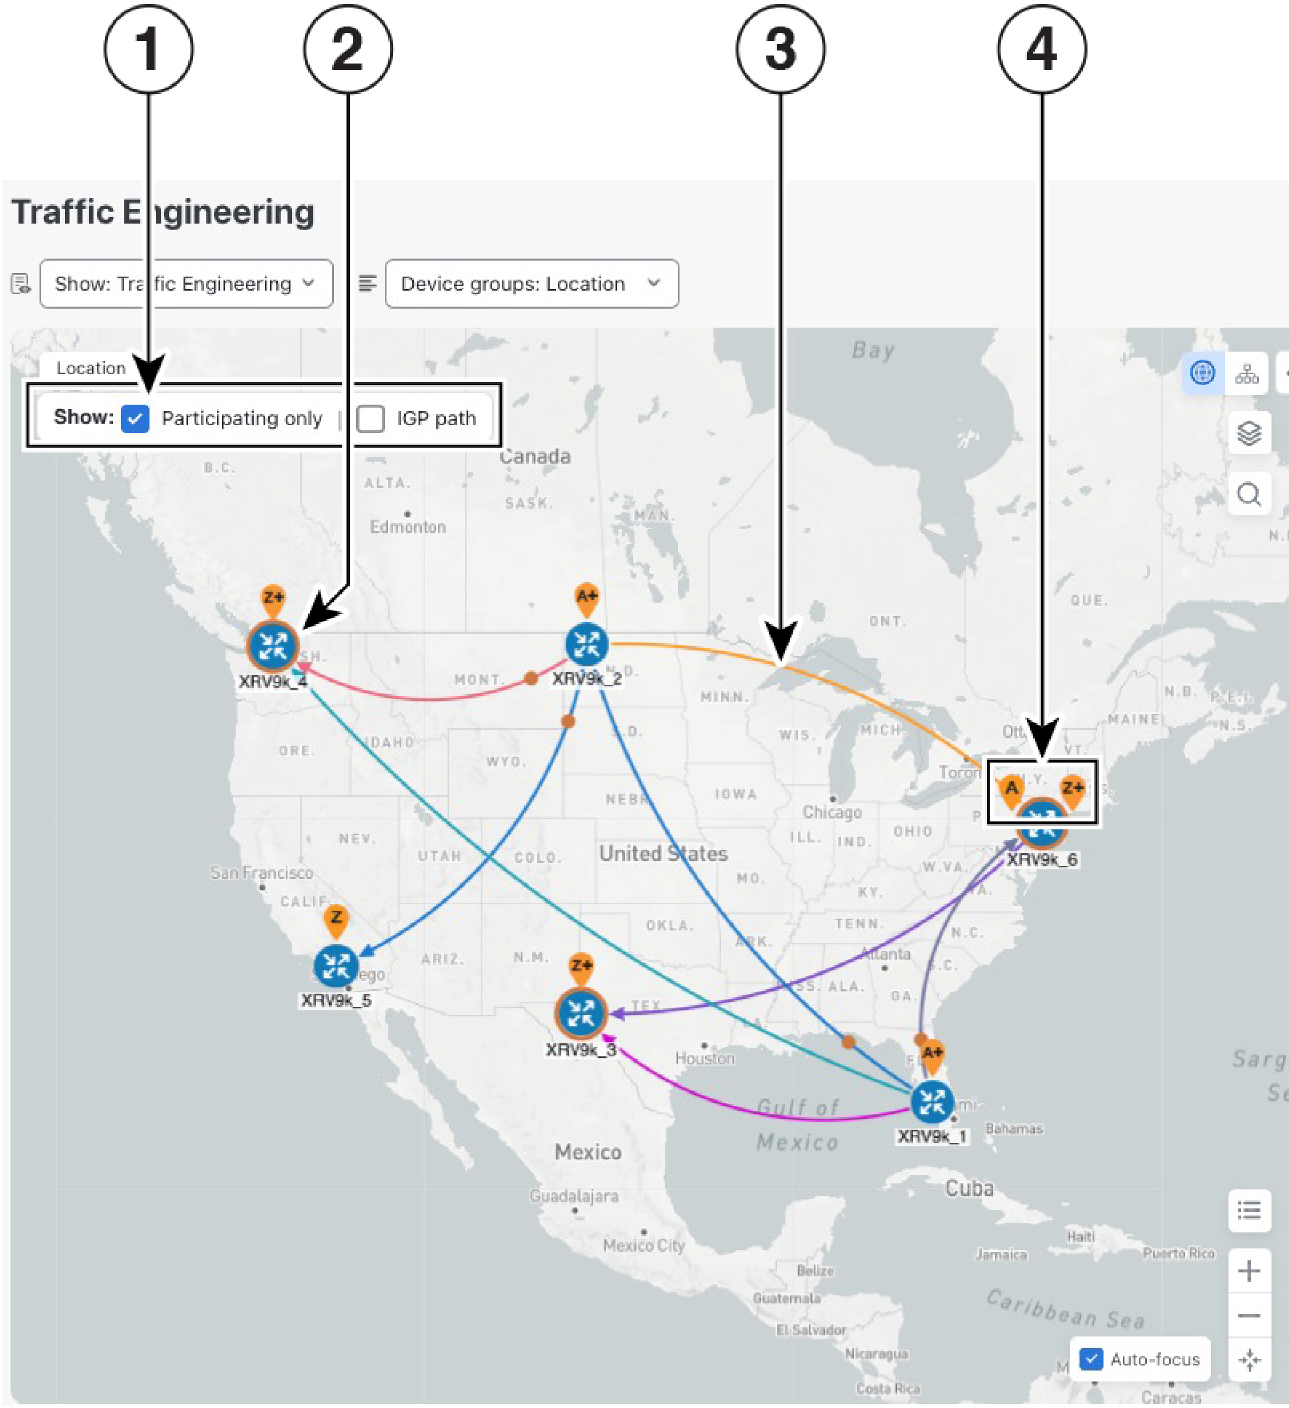

Callout No.	Description
1	Click the appropriate check box to enable the following options: • **Show: IGP path**—Displays the IGP path for the selected SR-TE policy. • **Show: Participating only**—Displays only links that belong to selected SR-TE policy. All other links and devices disappear.
2	A device with an orange ( ) outline indicates there is a node SID associated with that device or a device in the cluster.
3	When SR-TE policies are selected in the SR-MPLS or SRv6 tables, they show as colored directional lines on the map indicating source and destination. An adjacency segment ID (SID) is shown as an orange circle on a link along the path ( ).
4	**SR-MPLS and SRv6 Policy Origin and Destination**: If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The **A+** denotes that there is more than one SR-TE policy that originates from a node. The **Z+** denotes that the node is a destination for more than one SR policy.
5	The content of this window depends on what has been selected or filtered. In this example, the SR-MPLS tab is selected, and the SR Policy table is displayed.
6	Click on either the **SR-MPLS** or **SRv6** tabs to view the respective list of SR-TE policies.
7	Exports *all* data into a CSV file. You cannot export selected or filtered data.
8	The **Mini Dashboard** provides a summary of the operational SR-MPLS or SRv6 policy status. If filters are applied, the **Mini Dashboard** is updated to reflect what is displayed in the SR Policy and SRv6 Policy tables. In addition to the policy status, the **SR-MPLS Mini Dashboard** table displays the number of PCC and PCE initiated tunnels that are *currently* listed in the SR Policy table.
9	This option allows you to choose how the group filter (when in use) should be applied on the table data. For example, if **Headend only** was selected, then it would only display policies where the headend device of the policy is in the selected group. This filter allows you to see specific configurations and is useful when you have a large network. Filter options: • **Headend or Endpoint**—Show policies with either the headend or endpoint device in the selected group. • **Headend and Endpoint**—Show policies if both the headend and endpoint are in the group. • **Headend only**—Show policies if the headend device of the policy is in the selected group. • **Endpoint only**—Show policies if the endpoint device of the policy is in the selected group.

View SR-MPLS and SRv6 Policy Details

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

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

Figure 7: View SR Policy Details

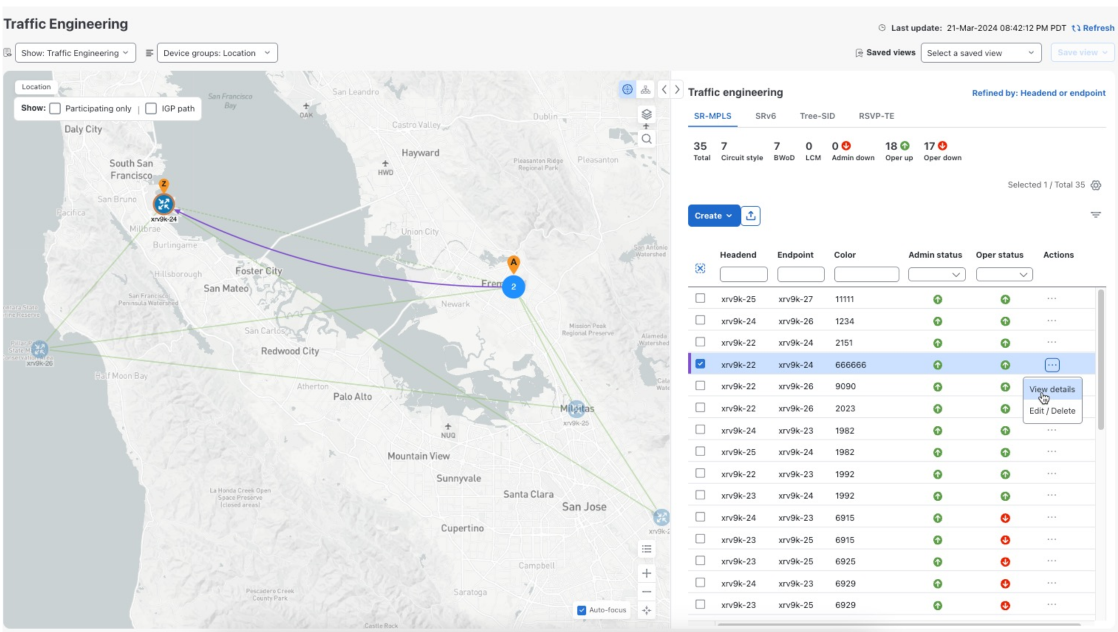

Step 2 View SR-MPLS or SRv6 policy details. From the browser, you can copy the URL and share with others.

Figure 8: SR Policy Details - Headend, Endpoint, and Summary

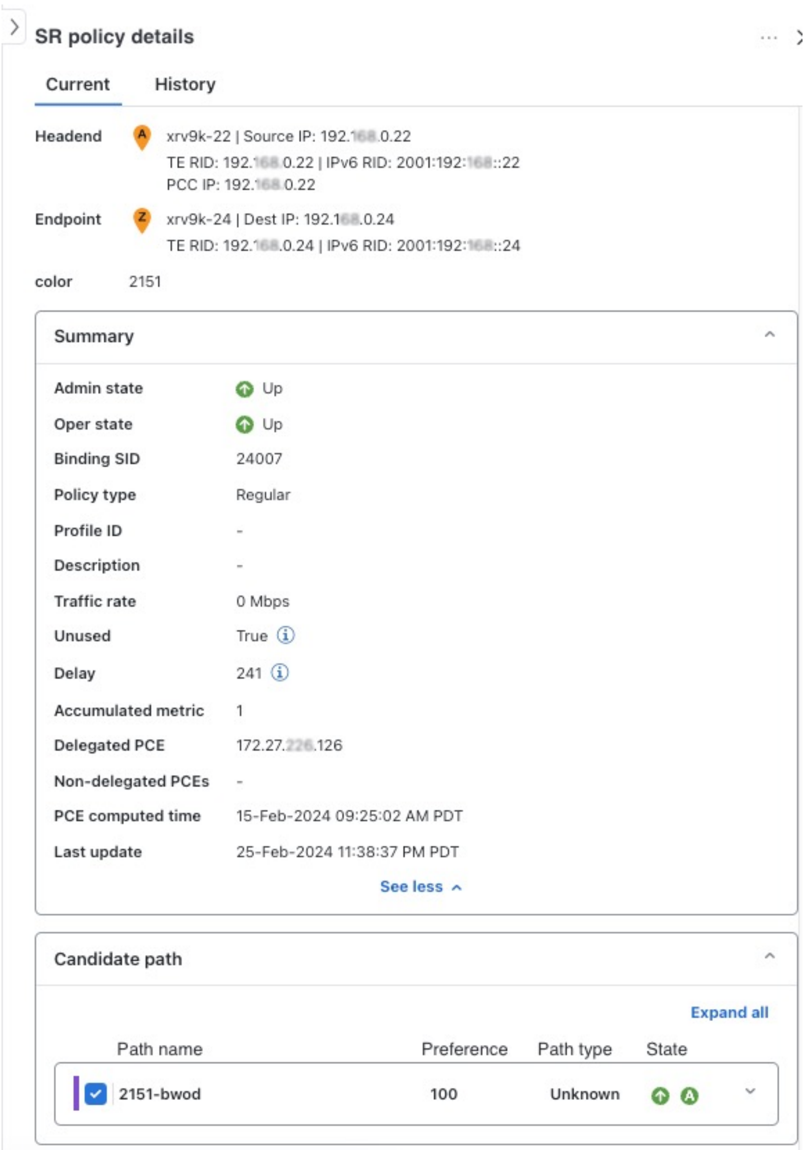

Note The Delay value is calculated for all policies every 10 minutes. Hover your mouse over the "i" icon (next to the Delay value) to view the last time the value was updated.

Visualize IGP Path and Metrics

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

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

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

Step 3 Click > **Metrics** tab.

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

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

Figure 9: View Physical Path and Metrics

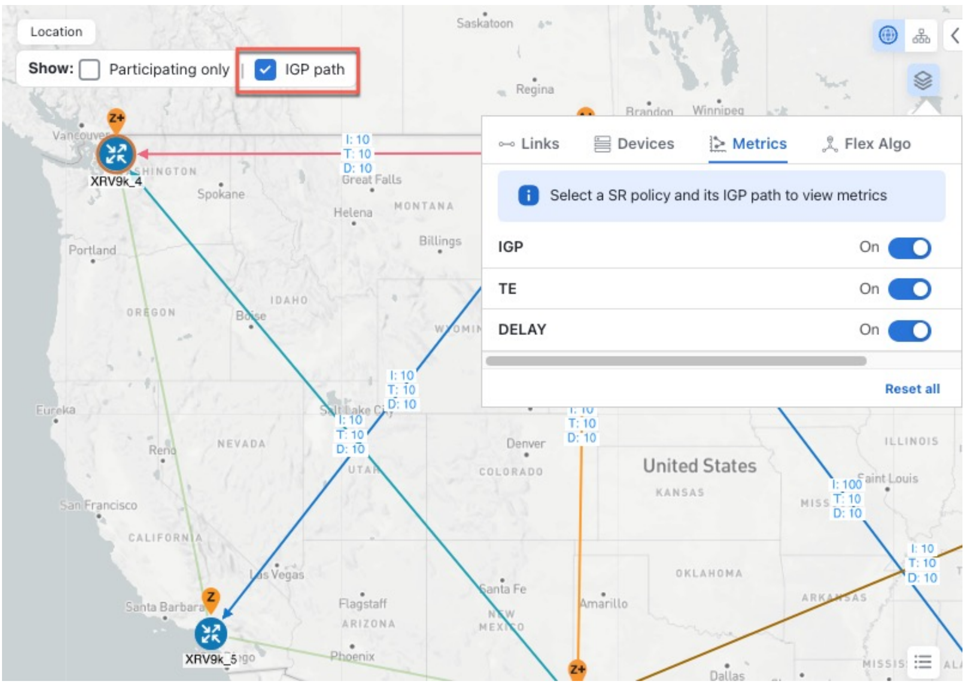

Find Multiple Candidate Paths (MCPs)

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

Important Notes

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

Before you begin

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

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

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

a) Check the check box next to the SR-TE policy that has MCPs configured.

b) View the SR-TE policy that is highlighted on the topology map.

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

Figure 10: SR-TE policy on the Topology Map

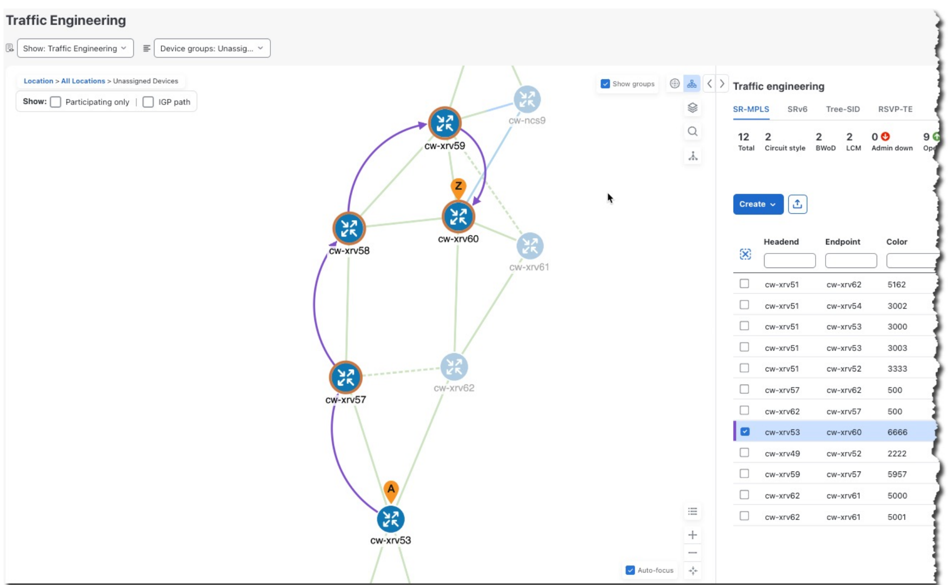

Step 3 View the list of candidate paths.

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

Figure 11: Candidate Path in SR Policy Details

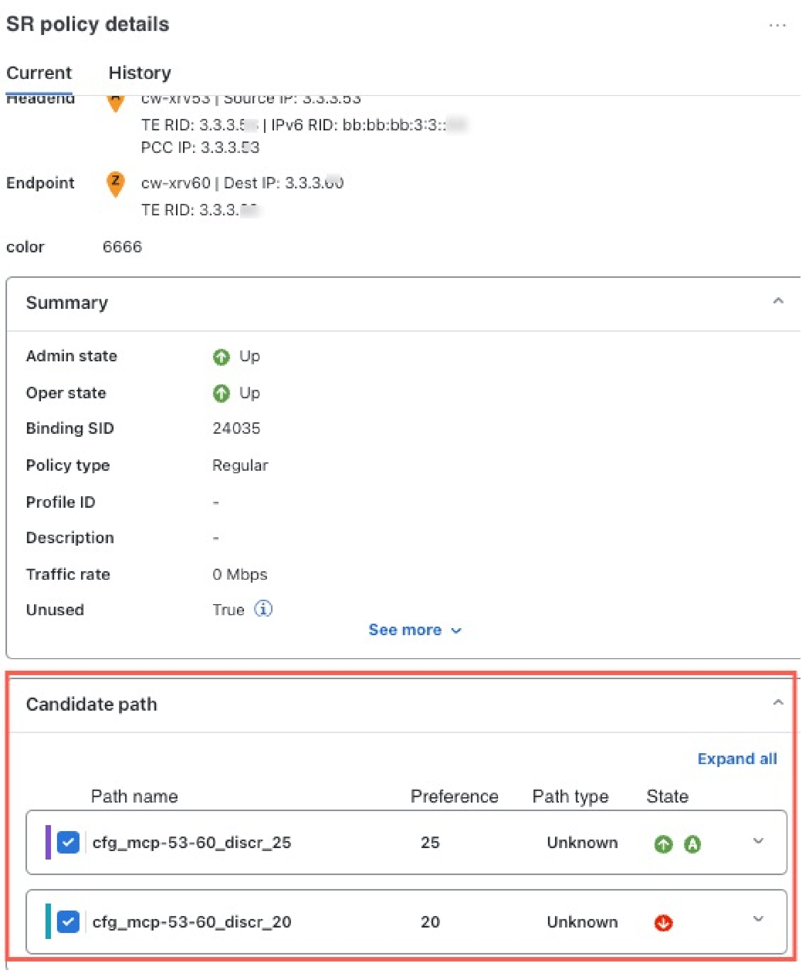

Step 4 You can expand individual paths or click **Expand all** to view details of each path.

Step 5 Visualize the candidate path on the topology map.

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

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

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

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

Figure 12: Candidate Path on the Topology Map

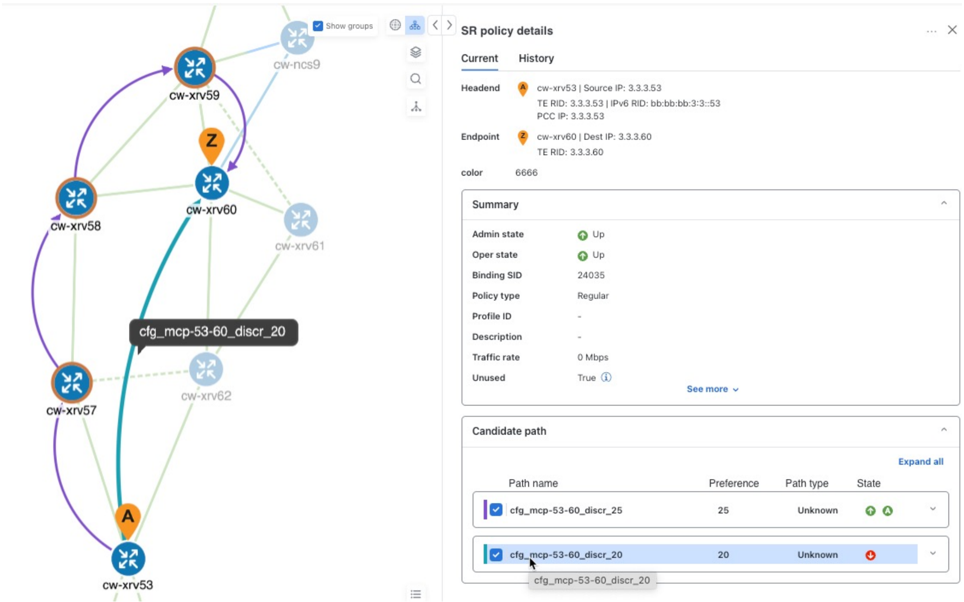

Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label

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

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

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

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

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

Figure 13: B-SID Label

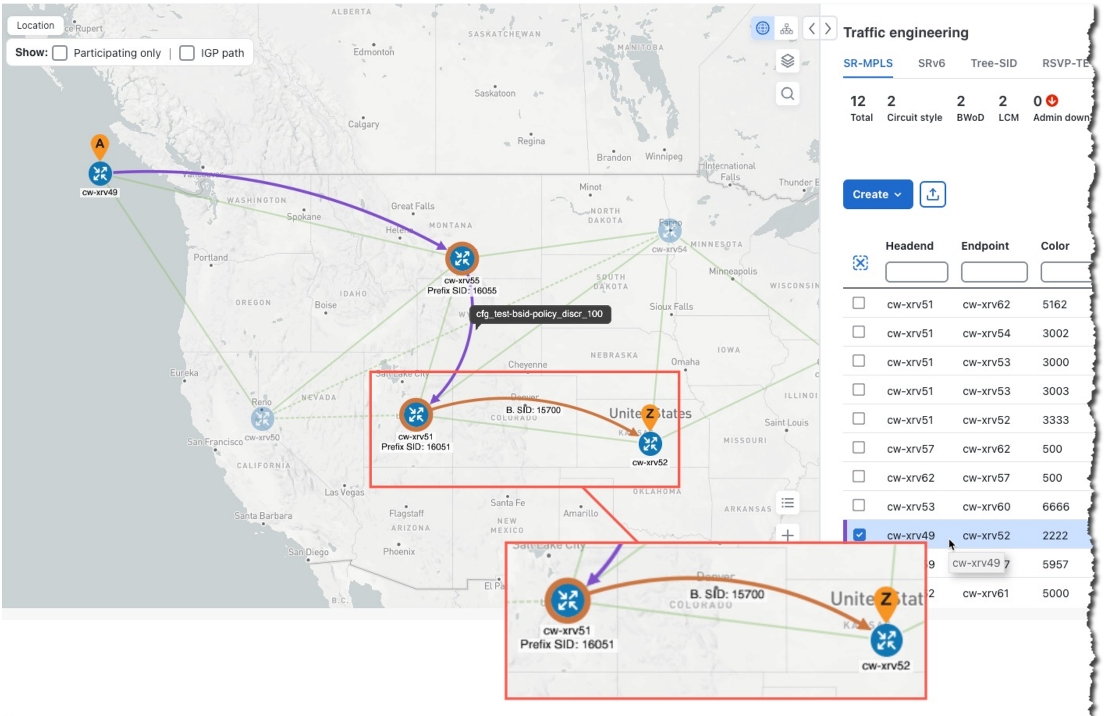

Step 3 From the **SR policy details** window, click ⋯ > **View details**.

Figure 14: View Details

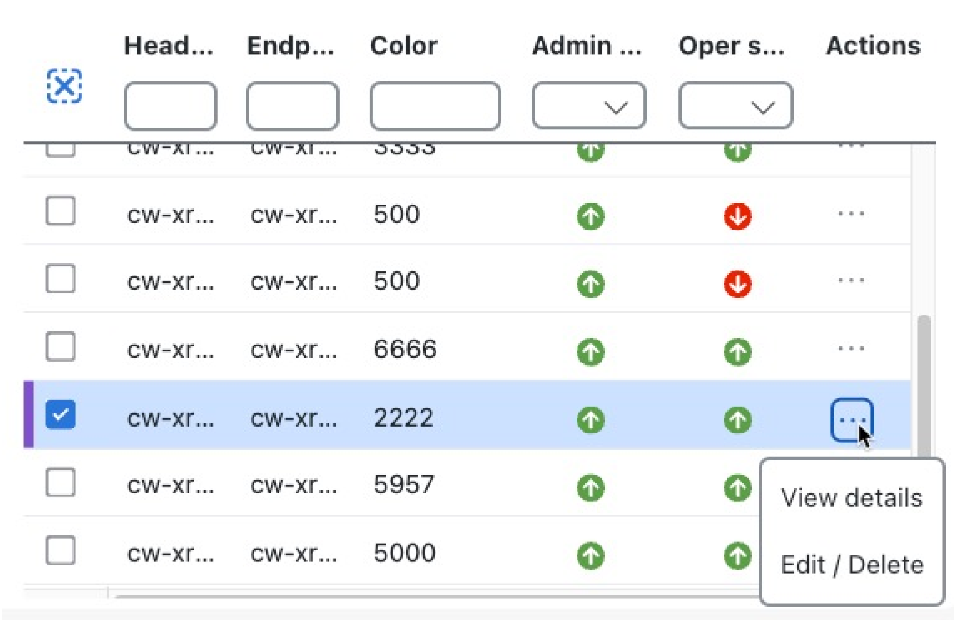

Step 4 Expand the active path and click the B-Sid Label ID to see the underlying path.

Figure 15: B-Sid Label ID

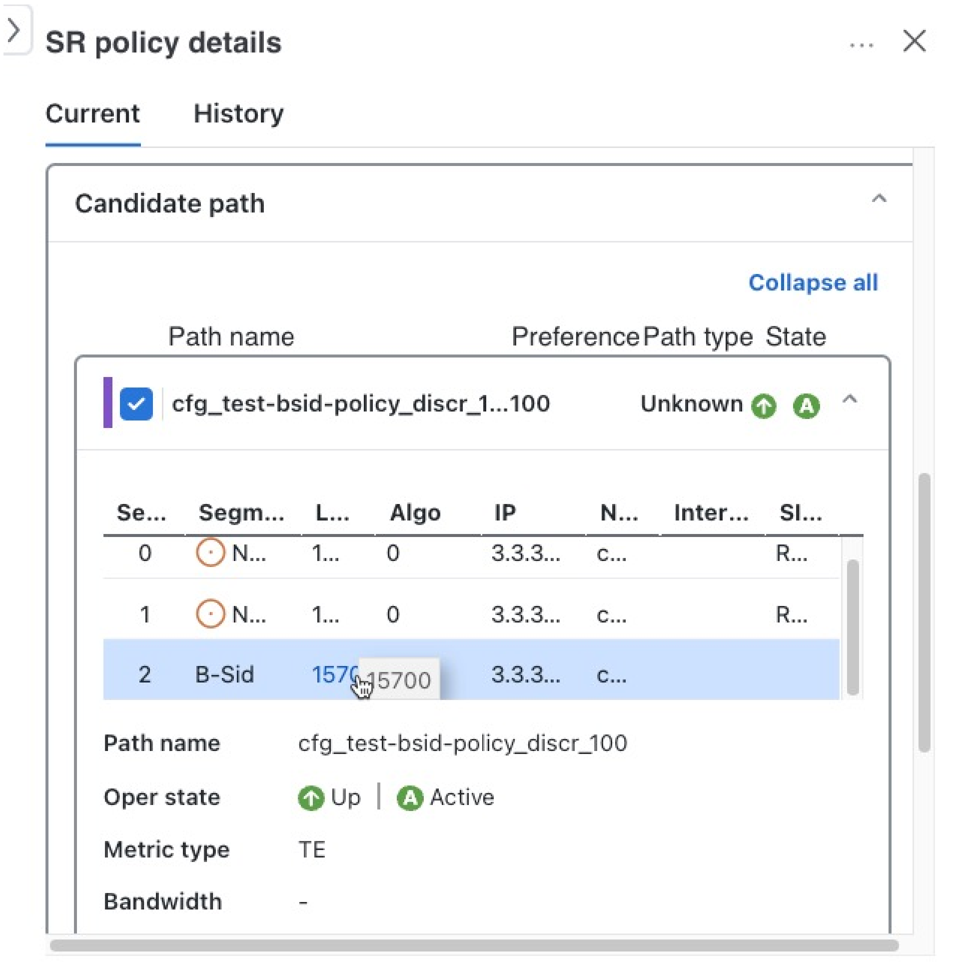

In this example, the underlying path actually goes from **cw-xrv51** > **cw-xrv54** > **cw-xrv53** > **cw-xrv52**.

Figure 16: B-SID Path

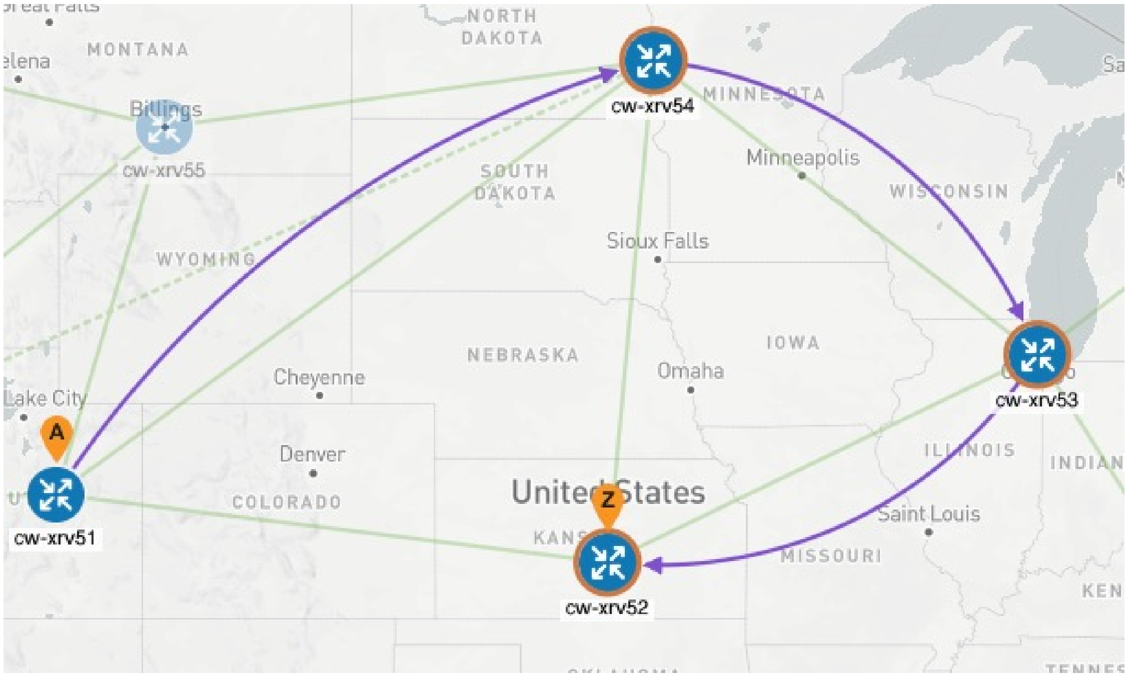

Visualize Native SR Paths

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

Before you begin

Confirm that device requirements are met. See Visualize Native Path Device Prerequisites, on page 27.

To create a path query, do the following:

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Path Query.**The Path Query dashboard appears.

Step 2 Click **New query.**

Step 3 Enter the device information in the required fields to find available Native SR IGP Paths and click **Get paths**.

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

Figure 17: New Path Query

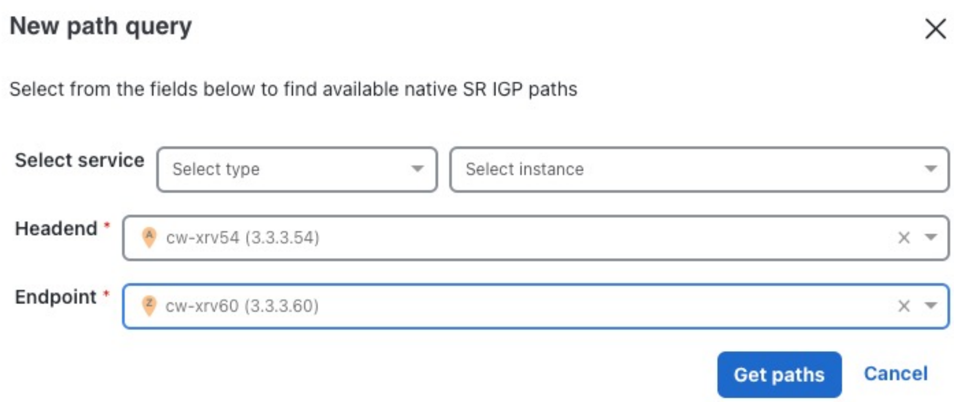

Step 4 Click **View results** when it becomes available on the Running Query ID pop-up. The Path Details window appears with corresponding available paths details while the topology map displays the available Native SR IGP paths on the left.

Figure 18: Path Details

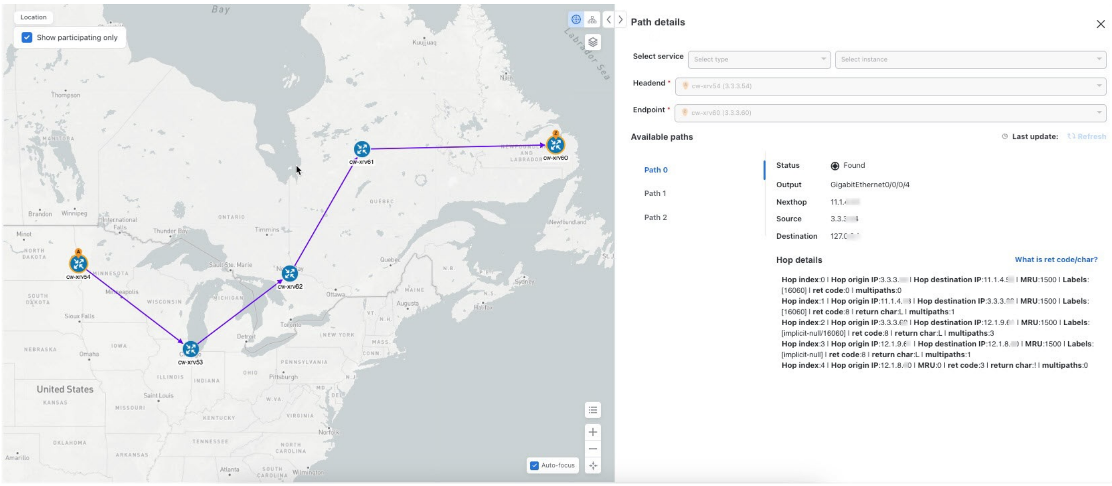

Visualize Native Path Device Prerequisites

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

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

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

```
address-family ipv6
default-route mgmt
!
!
!

or

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

Note

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

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

Note

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

- JSON
- BYTES
- PROTO
- ASCII
- JSON IETF

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

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

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

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

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

Configure TE Link Affinities

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

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

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

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

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

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

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

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

Figure 19: Mapping Affinities

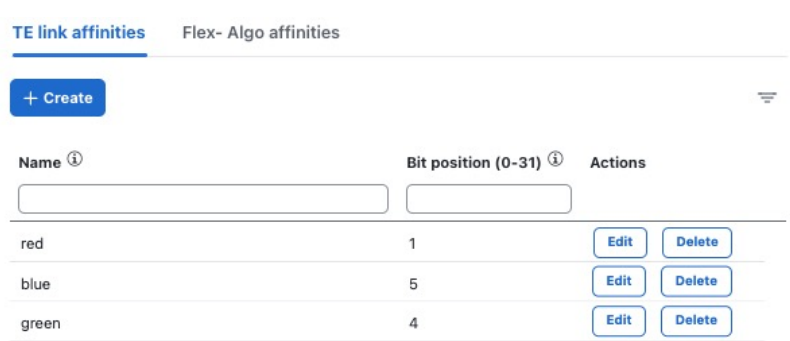

Step 4 Click **Save** to save the mapping. To create another mapping, you must click **+ Create** and save the entry.

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

Create Explicit SR-MPLS Policies

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

Tip If you plan to use affinities, collect affinity information from your devices, and then map them in Cisco Crosswork before creating an explicit SR-MPLS policy. For more information, see Configure TE Link Affinities, on page 29.

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

Step 2 Click **Create** > **PCE Init** .

Note If you would like to provision a PCC initiated policy using Cisco Network Services Orchestrator (NSO) via the Crosswork UI, see Create SR-TE Policies (PCC Initiated), on page 32.

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

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

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

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

Step 6 Click **Preview** and confirm that the policy you created matched your intent.

Step 7 If you want to commit the policy path, click **Provision** to activate the policy on the network or exit to abort the configuration process.

Step 8 Validate the SR-MPLS policy creation:

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

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

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

Note On a setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see Configure TE Timeout Settings, on page 12.

Create Dynamic SR-MPLS Policies Based on Optimization Intent

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

Tip For visualization purposes, you can optionally collect affinity information from your devices and then map them in Cisco Crosswork before creating a dynamic SR-MPLS policy. For more information, see Configure TE Link Affinities, on page 29 or Configure Flexible Algorithm Affinities, on page 43.

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

Step 2 Click **Create** > **PCE Init**.

Note If you would like to provision a PCC initiated policy using Cisco Network Services Orchestrator (NSO) via the Crosswork UI, see Create SR-TE Policies (PCC Initiated), on page 32.

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

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

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

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

Step 6 Define any applicable constraints and disjointness.

Note

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

Step 7 Under **Segments**, select whether or not protected segments should be used when available.

Step 8 If applicable, enter a SID constraint in the **SID Algorithm** field. Cisco Crosswork will try to find a path with this SID. If a path with the SID constraint cannot be found, the provisioned policy will remain operationally down until the conditions are met.

Note

- Flexible Algorithm: The values correspond to the Flexible Algorithm that are defined on the device and the 128-255 range is enforced by Cisco IOS XR.
- Algorithm 0: This is a Shortest Path First (SPF) algorithm based on link metric. This shortest path algorithm is computed by the Interior gateway protocol (IGP).
- Algorithm 1: This is a Strict Shortest Path First (SSPF) algorithm based on link metric. The algorithm 1 is identical to algorithm 0 but requires that all nodes along the path honor the SPF routing decision. Local policy does not alter the forwarding decision. For example, a packet is not forwarded through locally engineered path.

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

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

Step 11 Validate the SR-MPLS policy creation:

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

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

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

Note On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see Configure TE Timeout Settings, on page 12.

Create SR-TE Policies (PCC Initiated)

This task creates explicit or dynamic SR-MPLS or SRv6 policies using Cisco Network Services Orchestrator (NSO) via the Crosswork UI.

Before you begin

If you want to create explicit PCC initiated SR-MPLS or SRv6 policies, you must create a Segment IDs list (**Services & Traffic Engineering** > **Provisioning (NSO)** > **SR-TE** > **SID-List**). An explicit (fixed) path consists of a list of prefix or adjacency Segment IDs, each representing a node or link along on the path.

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Provisioning (NSO)**.

Step 2 From SR-TE > Policy, click . Crosswork displays the **Create SR-TE** > **Policy** window.

Note You may also click to import an existing SR-TE policy.

Step 3 Enter the policy constraints and required values.

You must populate the following options:

Table 2: SR-TE Policy Configuration

Expand this:	To specify this:
name	Enter a name for this SR-TE policy.
head-end	• You can click + to select a node or manually enter the node name.
tail-end	Manually enter the node name.
color	Enter a color. For example: 200.
path	**a.** Click + and enter a preference value. For example: 123 **b.** Select one of the following and toggle switch to enable: • explicit-path—Click + to add previously configured SID lists. • dynamic-path—Select the metric you want to minimize and define any applicable constraints and disjointness.
srv6	If you are creating an SRv6 policy, toggle **Enable srv6**.

Step 4 When you are finished, click **Dry Run** to validate your changes and save them. Crosswork will display your changes in a pop-up window.

If you want to configure a service that has requirements that do not match those we describe in this example, contact Cisco Customer Experience.

Step 5 When you are ready to activate the policy, click **Commit Changes**.

Modify SR-MPLS Policies

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

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

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

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

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

CHAPTER 3

Resource Reservation Protocol (RSVP)

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

View RSVP-TE Tunnels on the Topology Map

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

Figure 20: Traffic Engineering UI - RSVP-TE Tunnels

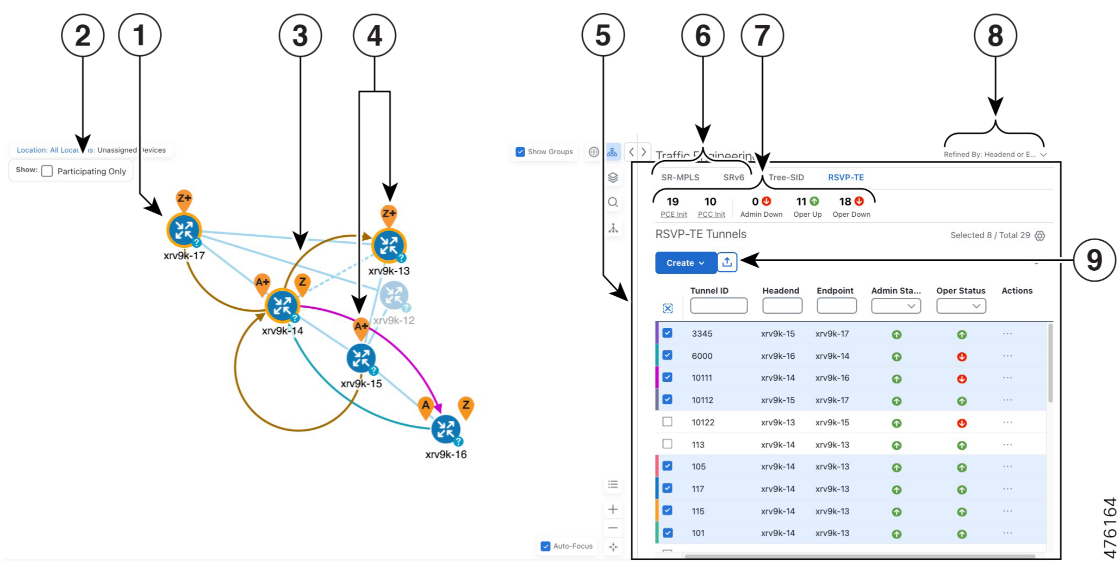

Callout No.	Description
1	Click **Show Participating Only** to display links belonging to the selected RSVP-TE tunnels. All other links and devices disappear.
2	A device with a solid orange outline ( ) indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered. **Note** RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.
3	When RSVP-TE tunnels are selected in the RSVP-TE Tunnel table, they show as colored directional lines on the map indicating source and destination. • Record Route Object (RRO) paths are shown as straight lines. • Explicit Route Object (ERO) paths are shown as curved lines. **Note** If both RRO and ERO paths are available, the RRO path is displayed by default. • An adjacency segment ID (SID) is shown as a green dot on a link along the path ( ). If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The **A+** denotes that there is more than one RSVP-TE tunnel that originates from a node. The **Z+** denotes that the node is a destination for more than one RSVP-TE tunnel.
4	**SR-MPLS and SRv6 Policy Origin and Destination**: If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. The **A+** denotes that there is more than one SR-TE policy that originates from a node. The **Z+** denotes that the node is a destination for more than one SR policy.
5	The content of this window depends on what has been selected or filtered. In this example, the RSVP-TE tab is selected and the RSVP-TE Tunnels table is displayed. Depending on what is selected on the topology map, or whether you are in the process of viewing and managing RSVP-TE tunnels, you can do the following: • Create Dynamic RSVP-TE Tunnels Based on Optimization Intent, on page 40 • Create Explicit RSVP-TE Tunnels, on page 40 • Modify RSVP-TE Tunnels, on page 42 • View RSVP-TE Tunnel Details, on page 37
6	Click the **RSVP-TE** tab.
7	The **Mini Dashboard** provides a summary of the operational RSVP-TE tunnel status and the number of PCC and PCE initiated tunnels that are *currently* listed in the RSVP-TE tables. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the RSVP-TE table.

Callout No.	Description
8	This option allows you to choose how the group filter (when in use) should be applied on the table data. For example, if **Headend only** was selected, then it would only display policies where the headend device of the policy is in the selected group. This filter allows you to see specific configurations and is useful when you have a large network. Filter options: • **Headend or Endpoint**—Show policies with either the headend or endpoint device in the selected group. • **Headend and Endpoint**—Show policies if both the headend and endpoint are in the group. • **Headend only**—Show policies if the headend device of the policy is in the selected group. • **Endpoint only**—Show policies if the endpoint device of the policy is in the selected group.
9	Exports *all* data into a CSV file. You cannot export selected or filtered data.

View RSVP-TE Tunnel Details

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

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

Figure 21: RSVP-TE > View details

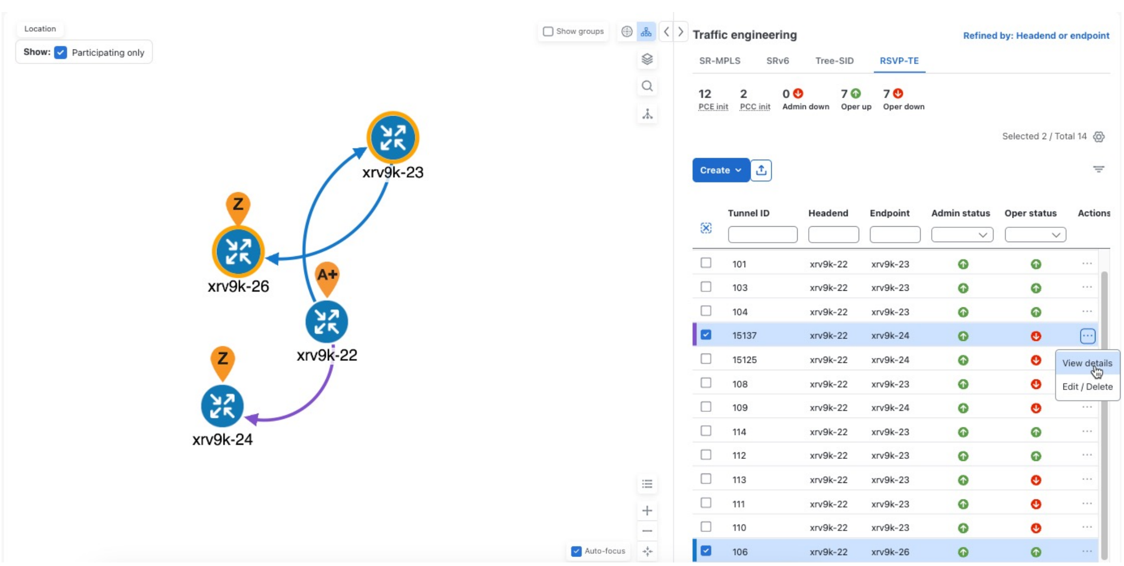

Step 2 View RSVP-TE tunnel details. From the browser, you can copy the URL and share with others.

Note

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

Figure 22: RSVP-TE Tunnel Details

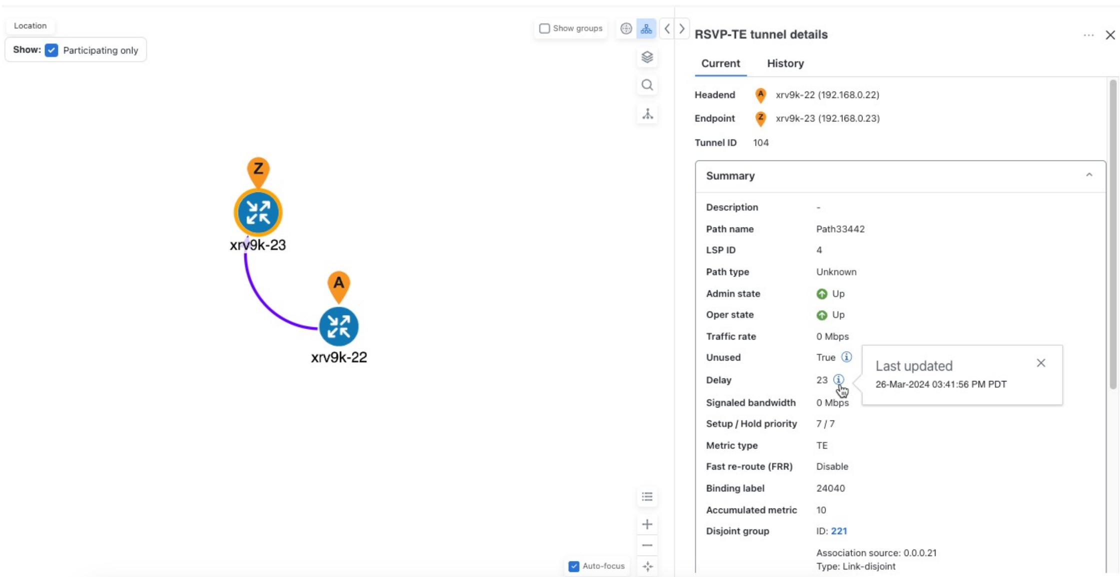

Figure 23: RSVP-TE Tunnel Details (close-up)

Create Explicit RSVP-TE Tunnels

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

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **RSVP-TE** tab.

Step 2 Click **Create > PCE Init**.

> **Note** If you would like to provision a PCC initiated tunnel using NSO via the Crosswork UI, see Create RSVP-TE Tunnels (PCC Initiated), on page 42.

Step 3 Under **Tunnel details**, enter the required RSVP-TE tunnel values. Hover the mouse pointer over ⓘ to view a description of each field.

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

Step 4 Under **Tunnel path**, click **Explicit path** and enter a path name.

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

Step 6 Click **Preview**. The path is highlighted on the map.

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

Step 8 Validate the RSVP-TE tunnel creation:

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

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

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

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

Create Dynamic RSVP-TE Tunnels Based on Optimization Intent

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

Tip If you plan to use affinities, collect affinity information from your devices and map them in Cisco Crosswork before creating a dynamic RSVP-TE tunnel. For more information, see Configure TE Link Affinities, on page 29.

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **RSVP-TE** tab.

Step 2 Click **Create** > **PCE Init**.

Note If you would like to provision a PCC initiated tunnel using NSO via the Crosswork UI, see Create RSVP-TE Tunnels (PCC Initiated), on page 42.

Step 3 Under **Tunnel details**, enter the required RSVP-TE tunnel values. Hover the mouse pointer over ⓘ to view a description of each field.

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

Step 4 Under **Tunnel path**, click **Dynamic path** and enter the Path Name.

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

Step 6 Define any applicable constraints and disjointness.

Note Affinity constraints and disjointness cannot be configured on the same RSVP-TE tunnel. Also, there can be up to two RSVP-TE tunnels in the same disjoint group or subgroup. If there are existing RSVP-TE tunnels belonging to a disjoint group that you define here, all RSVP-TE tunnels belonging to that same disjoint group are shown during Preview.

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

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

Step 9 Validate the RSVP-TE tunnel creation:

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

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

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

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

Create RSVP-TE Tunnels (PCC Initiated)

This task creates explicit or dynamic RSVP-TE tunnels using Cisco Network Services Orchestrator (NSO) via the Crosswork UI.

Before you begin

If you want to create explicit PCC initiated RSVP-TE tunnels, you must create a Segment IDs list (**Services & Traffic Engineering** > **Provisioning (NSO)** > **SR-TE** > **SID-List**). An explicit (fixed) path consists of a list of prefix or adjacency Segment IDs, each representing a node or link along on the path.

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Provisioning (NSO)**.

Step 2 From **RSVP-TE** > **Tunnel**, click +. Crosswork displays the **Create RSVP-TE** > **Tunnel** window.

Note You may also click to import an existing RSVP-TE tunnel.

Step 3 Enter the policy constraints and required values.

Step 4 When you are finished, click **Dry run** to validate your changes and save them. Crosswork will display your changes in a pop-up window.

Step 5 When you are ready to activate the policy, click **Commit changes**.

Modify RSVP-TE Tunnels

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

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **RSVP-TE** tab.

Step 2 Locate the RSVP-TE tunnel you are interested in and click ….

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

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

CHAPTER 4

Flexible Algorithm

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

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

This section contains the following topics:

Configure Flexible Algorithm Affinities

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

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

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

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

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

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

Step 1 From the main menu, select **Administration** > **Settings** > **Traffic engineering > Affinity > Flex-Algo affinities** tab.

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

Step 3 Enter the name and the bit it will be assigned.

Step 4 Click **Save** to save the mapping. To view all Flexible Algorithm affinities for a link, see View Flexible Algorithm Details , on page 45.

Visualize Flexible Algorithm Topologies

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

Note To apply a Flexible Algorithm constraint when dynamically provisioning an SR-MPLS policy, see Create Dynamic SR-MPLS Policies Based on Optimization Intent, on page 31.

Before you begin

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

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

Step 1 From the main menu, select **Services & Traffic Engineering** > **Traffic Engineering**.

Step 2 From the topology map, click .

Step 3 Click the **Flex Algo** tab.

Step 4 From the drop-down list, select up to two Flexible Algorithm IDs.

Step 5 View the Flexible Algorithm Types and confirm that the selection is correct. Also, note the color assignments for each Flexible Algorithm.

Step 6 (Optional) Check the **Show selected Flex Algo topology only** check box to isolate the Flexible Algorithms on the topology map. When this option is enabled, SR policy selection is disabled.

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

Step 7 Click **Apply**. You must click **Apply** for any additional changes to Flexible Algorithm selections to be reflected on the topology map.

Figure 24: Flexible Algorithm on Map

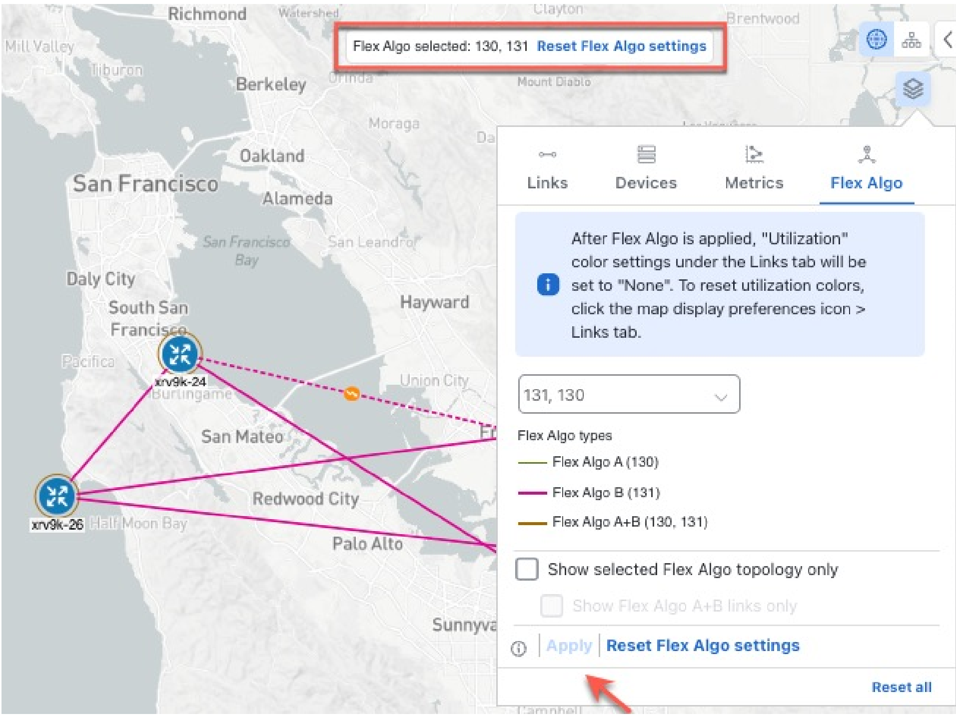

Note If a selected Flexible Algorithm is defined with criteria but there are no link and node combinations that match it (for example, a defined affinity to include all nodes or links with the color blue), then the topology map will be blank. If a selected Flexible Algorithm is not configured on a node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

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

View Flexible Algorithm Details

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

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

Step 2 To view device Flexible Algorithm details:

a) From the topology map, click on a device.

b) In the **Device details** window, navigate to the **Traffic engineering** > **Flex Algo** tab. For example:

Figure 25: Flex Algo Device Details

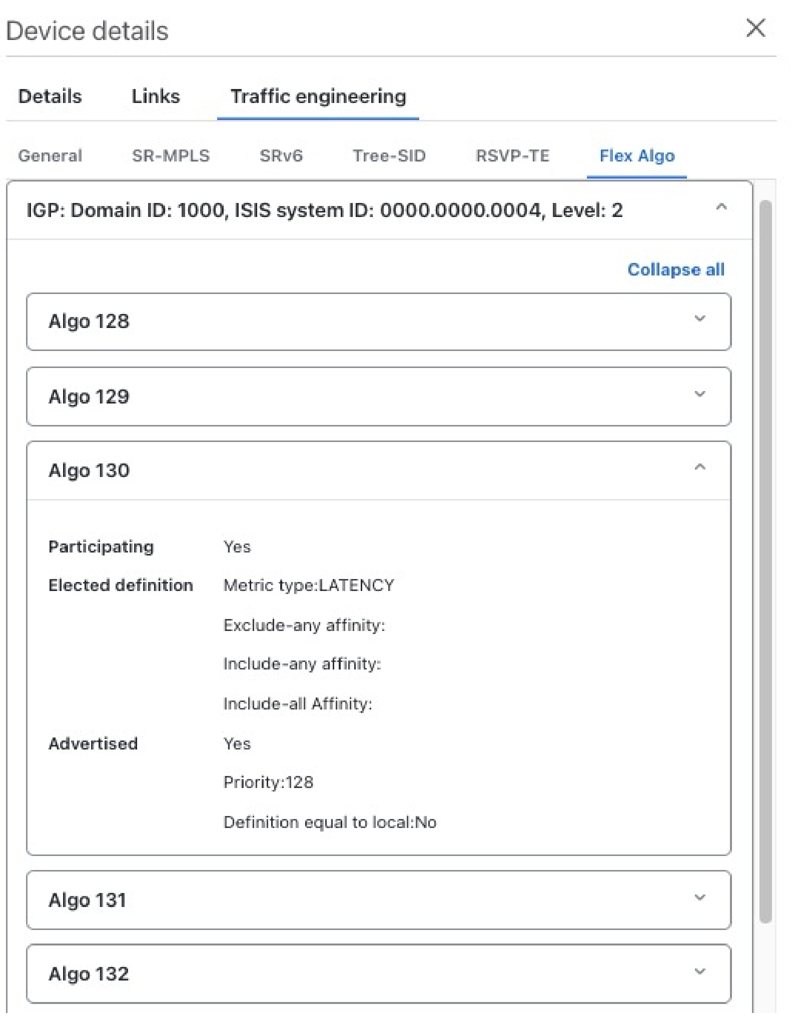

Step 3 To view whether a link is part of a Flexible Algorithm topology:

a) From the topology map, click on a link. If a list of links appear, click on a link type.

b) From the **Link details** window, click the **Traffic engineering** tab. If the link is a member, then the **FA topologies** row displays what Flexible Algorithm each source and destination device belong to.

Figure 26: Flex Algo Link Details

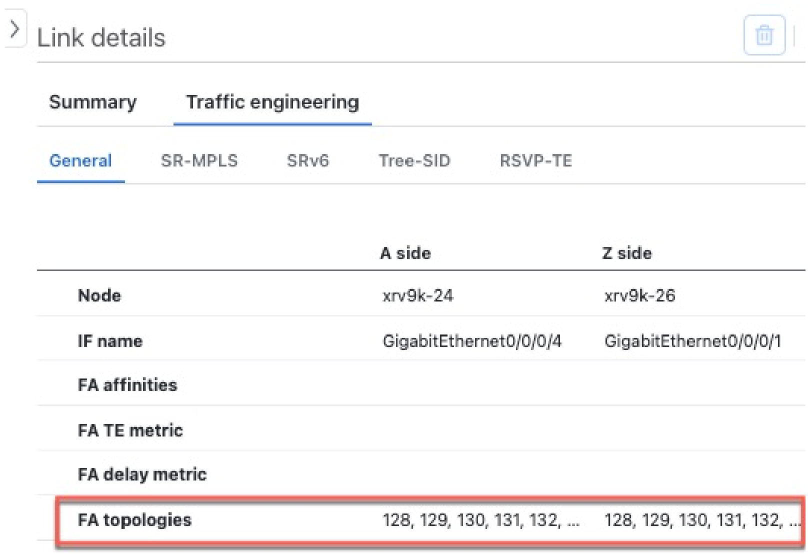

Note

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

CHAPTER 5

Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering

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

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

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

This section contains the following topics.

Visualize Tree-SID Policies

Crosswork UI provides the ability to view details of the Tree-SID root, transit, leaf nodes, and bud nodes in the UI and allows you to easily confirm that Tree-SID is implemented correctly in your network (see View a Point-to-Multipoint Tree on the Topology Map, on page 50).

The Tree-SID policy has the following nodes:

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

You can visualize the following Tree-SID policies:

- **Static:** A Static Tree-SID policy is configured via SR-PCE, directly using SR-PCE CLI or from the Crosswork UI. You can refer to the Tree-SID configuration documentation for your specific device for more information and examples of the supported configuration commands. (for example, Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers

- **Dynamic:** A Dynamic Tree-SID policy is not explicitly configured; it is configured as part of an L3VPN/ mVPN service.

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

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

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

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

Figure 27: Create a new Tree-SID Policy (static)

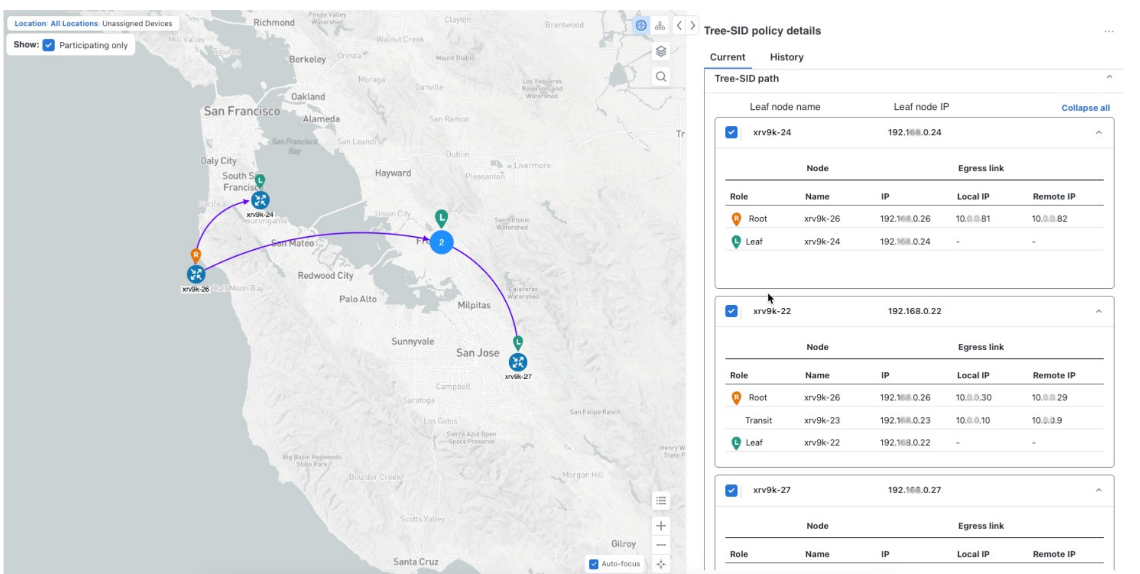

Before you begin

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

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

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

Figure 28: Tree-SID Policies (static) on the Topology Map

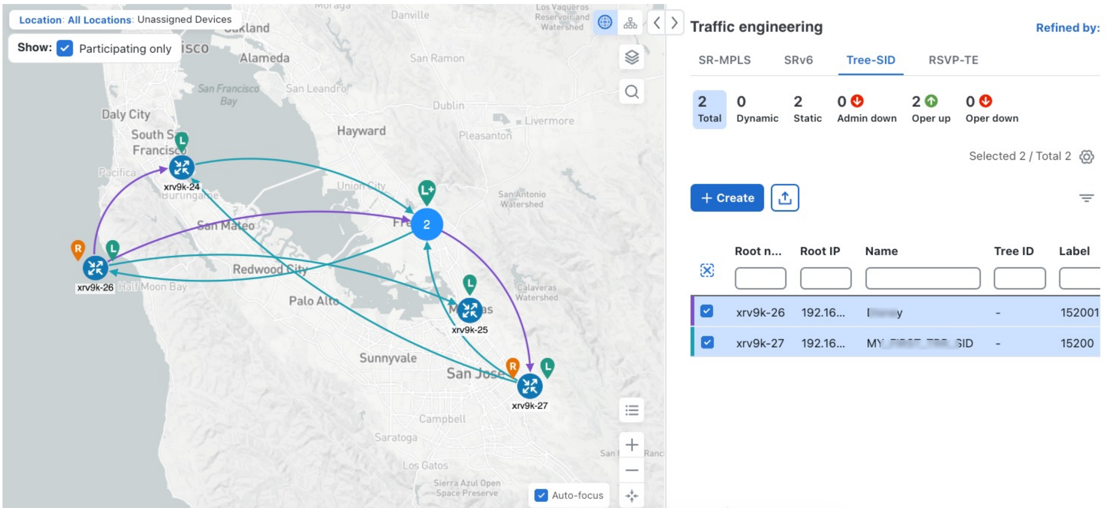

Note Any change in end-points is captured as an event in the historical data tab. For information on Tree-SID Historical Data, see View TE Event and Utilization History, on page 9

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

Example:

Note

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

Figure 29: Tree-SID Details Summary

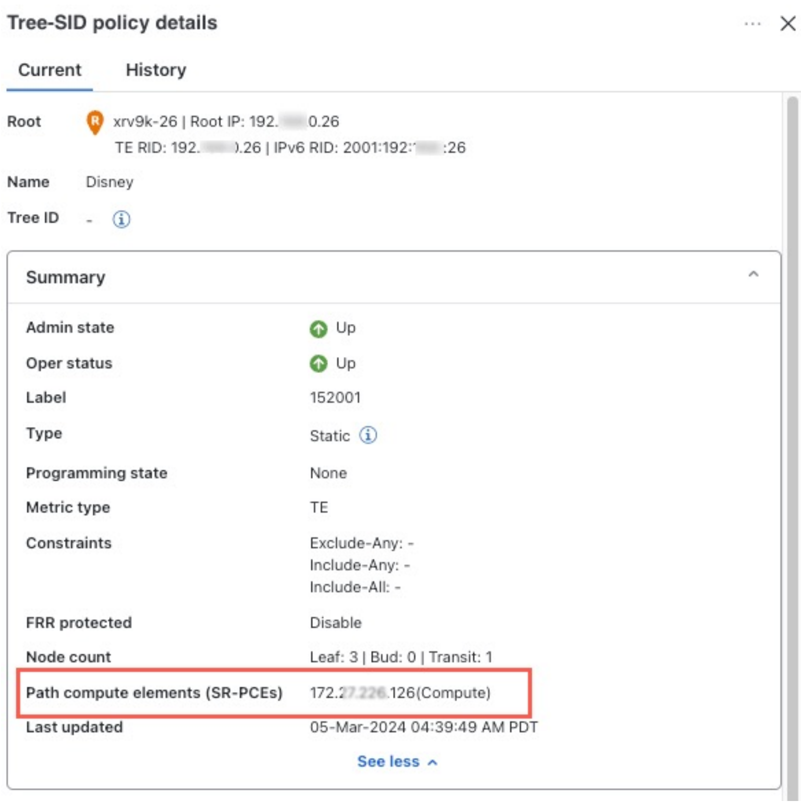

Figure 30: Tree-SID Path Details

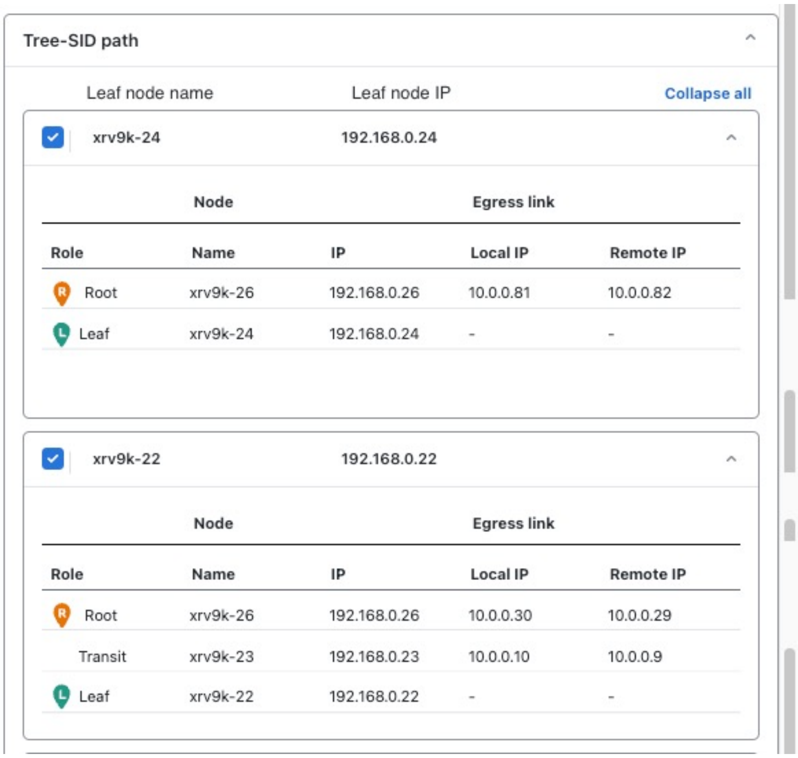

Create Static Tree-SID Policies

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

Tip If you plan to use affinities, collect affinity information from your devices and map them in Cisco Crosswork before creating a static Tree-SID policy. For more information, see Configure TE Link Affinities, on page 29.

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

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

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

Figure 31: Create Static Tree-SID Policy

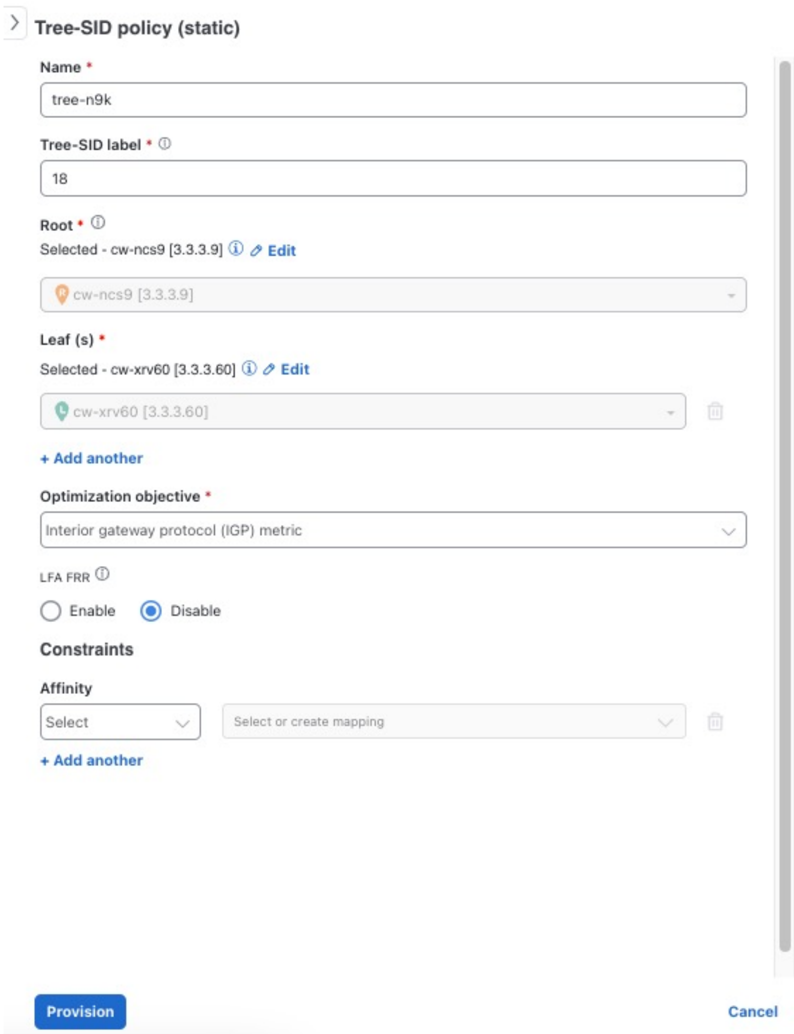

Step 3 To commit the policy, click **Provision**.

Step 4 Validate the Tree-SID policy creation:

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

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

Figure 32: Newly Added Tree-SID Policy on Topology Map

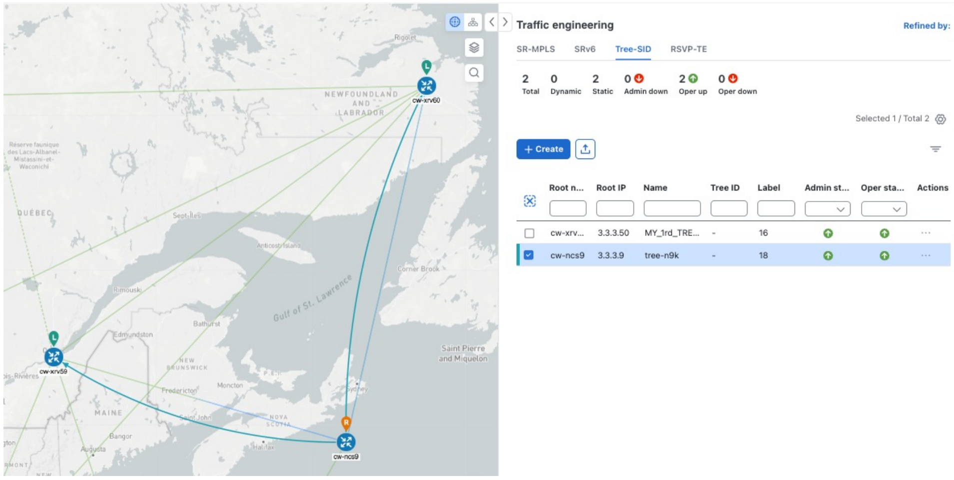

b. View and confirm the new Tree-SID policy details. From the Actions column, click ··· and select **View details**.

Figure 33: Tree-SID Policy Details

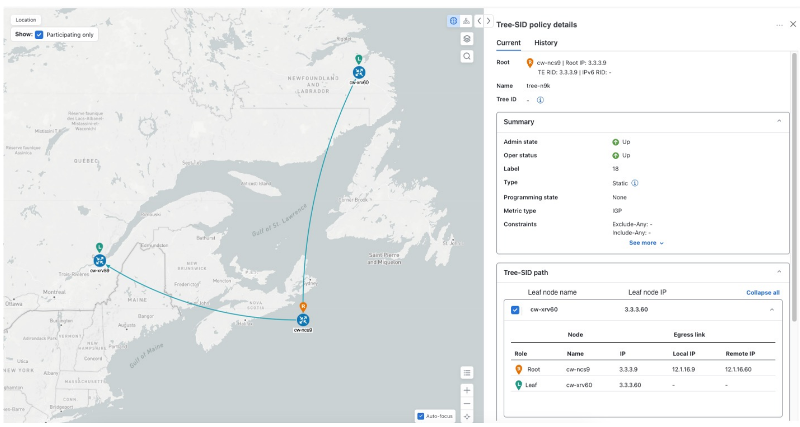

Static Tree-SID Policy Configuration Example through Crosswork UI

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

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


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

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

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

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

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

```
Role: Ingress
Hops:
 Incoming: 505254 CC-ID: 1
 Outgoing: 505254 CC-ID: 1 (11.1.28.54)
 Outgoing: 505254 CC-ID: 1 (11.1.1.51)
```

Modify a Tree-SID Policy

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

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

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

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

Step 3 Choose **Edit/Delete**.

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

Tree-SID Important Notes

Limitation

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

Visualization of Tree-SID Paths with Missing Nodes

Missing Tree-SID nodes can cause the following issues:

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

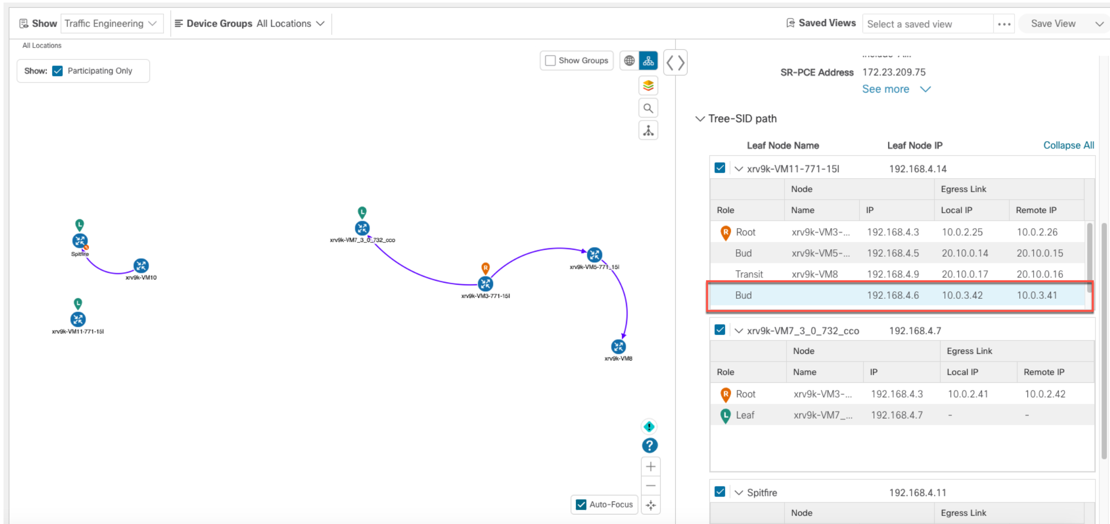
Show
Traffic Engineering
Device Groups
All Locations
Saved Views
Select a saved view
Save View
Show:
Participating Only
Show Groups
SR-PCE Address
172.23.209.75
See more
Tree-SID path
Leaf Node Name
Leaf Node IP
Collapse All
xrv9k-VM11-771-15I
192.168.4.14
Node
Egress Link
Role
Name
IP
Local IP
Remote IP
Root
xrv9k-VM3-...
192.168.4.3
10.0.2.25
10.0.2.26
Bud
xrv9k-VM5-...
192.168.4.5
20.10.0.14
20.10.0.15
Transit
xrv9k-VM8
192.168.4.9
20.10.0.17
20.10.0.16
Bud
192.168.4.6
10.0.3.42
10.0.3.41
xrv9k-VM7_3_0_732_cco
192.168.4.7
Root
xrv9k-VM3-...
192.168.4.3
10.0.2.41
10.0.2.42
Leaf
xrv9k-VM7_...
192.168.4.7
Spitfire
192.168.4.11
Auto-Focus