



Cisco Crosswork Network Controller 6.0 Traffic Engineering and Optimization

First Published: 2023-11-14

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	1
Segment Routing Path Computation Element (SR-PCE)	2
What is Segment Routing?	2
SR-TE Policy PCC and PCE Configuration Sources	4
PCC-Initiated SR-TE Policy Example	5
What is Resource Reservation Protocol (RSVP)?	5
RSVP-TE Tunnel PCC and PCE Configuration Sources	6
PCC-Initiated RSVP-TE Tunnel Example	7
Get a Quick View of Traffic Engineering Services	7
View TE Event and Utilization History	9
View Traffic Engineering Device Details	11
Configure Traffic Engineering Settings	11
Configure TE Timeout Settings	11
Configure How Device Groups Are Displayed for Traffic Engineering	12
Configure Historical Data Settings	12
Resolve SR-TE Policies and RSVP-TE Tunnels	13

CHAPTER 2

SR-MPLS and SRv6	15
View SR-MPLS and SRv6 Policies on the Topology Map	15
View SR-MPLS and SRv6 Policy Details	17
Visualize IGP Path and Metrics	19
Find Multiple Candidate Paths (MCPs)	20
Visualize Underlying Paths Associated with a Defined Binding-Segment ID (B-SID) Label	23
Visualize Native SR Paths	24
Visualize Native Path Device Prerequisites	26

Configure TE Link Affinities	28
Create Explicit SR-MPLS Policies	29
Create Dynamic SR-MPLS Policies Based on Optimization Intent	30
Create SR-TE Policies (PCC Initiated)	31
Modify SR-MPLS Policies	32

CHAPTER 3**Resource Reservation Protocol (RSVP) 33**

View RSVP-TE Tunnels on the Topology Map	33
View RSVP-TE Tunnel Details	35
Create Explicit RSVP-TE Tunnels	36
Create Dynamic RSVP-TE Tunnels Based on Optimization Intent	37
Create RSVP-TE Tunnels (PCC Initiated)	38
Modify RSVP-TE Tunnels	39

CHAPTER 4**Flexible Algorithms 41**

Configure Flexible Algorithm Affinities	41
Visualize Flexible Algorithm Topologies	42
View Flexible Algorithm Details	43

CHAPTER 5**Tree Segment Identifier (Tree-SID) Multicast Traffic Engineering 47**

Visualize Tree-SID Policies	47
View a Point-to-Multipoint Tree on the Topology Map	48
Create Static Tree-SID Policies	52
Static Tree-SID Policy Configuration Example through Crosswork UI	53
Modify a Tree-SID Policy	55
Tree-SID Important Notes	55



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. TE improves network performance by forcing traffic to take predetermined routes and by using available resources effectively.

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.

The following table lists what Traffic Engineering SR policies and RSVP tunnels can be visualized and provisioned through the Crosswork UI.

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

Table 1: Supported TE Technologies

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



Note Users must be assigned admin roles or have certain Device Access Group permissions to provision and access some features. For more information on Role-based Access Control (RBAC) and user roles, see the "[Cisco Crosswork Network Controller Administration Guide](#)".

- [Segment Routing Path Computation Element \(SR-PCE\)](#), on page 2
- [What is Segment Routing?](#), on page 2
- [SR-TE Policy PCC and PCE Configuration Sources](#), on page 4
- [What is Resource Reservation Protocol \(RSVP\)?](#), on page 5
- [RSVP-TE Tunnel PCC and PCE Configuration Sources](#), on page 6
- [Get a Quick View of Traffic Engineering Services](#), on page 7
- [View TE Event and Utilization History](#), on page 9
- [View Traffic Engineering Device Details](#), on page 11
- [Configure Traffic Engineering Settings](#), on page 11
- [Resolve SR-TE Policies and RSVP-TE Tunnels](#), on page 13

Segment Routing Path Computation Element (SR-PCE)

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

Cisco SR-PCE (formerly Cisco XR Traffic Controller (XTC)) runs on the Cisco IOS XR operating system. SR-PCE provides stateful PCE functionality that helps control and reroute TE tunnels to optimize the network. PCE describes a set of procedures by which a Path Computation Client (PCC) can report and delegate control of headend tunnels that are sourced from the PCC to a PCE peer. The PCC and PCE establish a Path Computation Element Communication Protocol (PCEP) connection that SR-PCE uses to push updates to the network.

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



Note Features may not work as expected if the SR-PCE version is not supported. It is important to refer to the [Cisco Crosswork Network Controller Release Notes](#) for SR-PCE version support and compatibility.

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

What is Segment Routing?

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

to the destination, and instructs the routers in the provider core network to follow the specified path instead of the shortest path calculated by the IGP. The destination is unaware of the presence of the tunnel.

Segments

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

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

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

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

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

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

Segment Routing Policies

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

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

- [SR-MPLS and SRv6](#), on page 15
- [Flexible Algorithms](#), on page 41
- [Tree Segment Identifier \(Tree-SID\) Multicast Traffic Engineering](#), on page 47



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

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

Dynamic SR Policy

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

Explicit SR Policy

When you configure an explicit policy, you specify an explicit path which consists of a list of prefix or adjacency SIDs, each representing a node or link along on the path.

Disjointness

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

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

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



Note

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

SR-TE Policy PCC and PCE Configuration Sources

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

- Path Computation Client (PCC) initiated—Policies configured on a PCC (see [PCC-Initiated SR-TE Policy Example, on page 5](#)). This policy type displays as **Unknown** in the UI.
- Path Computation Element (PCE) initiated—Policies configured on a PCE or created dynamically by Crosswork. PCE Initiated policy types can be one of the following:
 - **Dynamic**
 - **Explicit**
 - **Circuit-Style**

- **Bandwidth on Demand**
- **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, which is associated with establishing and managing TE tunnels at the headend and tail end.
- Link-management, which manages link resources to do resource-aware routing of TE Label-Switched Path (LSP) and to program MPLS labels.
- Fast Reroute (FRR), which 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, which could be a sequence of IP prefixes or a sequence of autonomous systems, in the Explicit Route Object (ERO). The explicit path can be administratively specified, or automatically computed using an algorithm such as constrained shortest path first (CSPF).

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

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

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



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

RSVP FRR

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

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

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

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

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**.

Get a Quick View of Traffic Engineering Services

TE Dashboard © Last Update: 07-Aug-2023 04:55:52 PM PDT | ↕

SR-MPLS

15

Total Policy Count

Policy State

Oper Down: 4, Admin Down: 0, Oper Up: 11

Policy Type & Metric Type

- BWsd: 1, LCM: 4, Regular: 9, Circuit Style1: 1
- IGP: 0, TE: 4, LATENCY: 1, HOPCOUNT: 0, UNKNOWN: 10

SRv6

0

Total Policy Count

Policy State

Oper Down: 0, Admin Down: 0, Oper Up: 0

Metric Type

- IGP: 0, TE: 0, LATENCY: 0, HOPCOUNT: 0, UNKNOWN: 0

Tree-SID

2

Total Policy Count

Policy State

Oper Down: 0, Admin Down: 0, Oper Up: 2

Metric Type

- IGP: 1, TE: 1, LATENCY: 0, HOPCOUNT: 0, UNKNOWN: 0

RSVP-TE

2

Total Tunnel Count

Policy State

Oper Down: 1, Admin Down: 0, Oper Up: 1

Metric Type

- IGP: 1, TE: 1, LATENCY: 0, HOPCOUNT: 0, UNKNOWN: 0

Fast Re-Route

0 Policies with FRR enabled

Fast Re-Route

0 Policies with FRR enabled

2 → Policies and Tunnels Under Traffic Threshold Range 0 to 1000 Kbps 📅 06-Aug-2023 16:55 to 07-Aug-2023 16:55 1M 1W 1D 1H | Reset

3 → Policy / Tunnel Type: All SR-MPLS RSVP-TE Total 17

Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	↑ Traffic Rate (Kbps)	Actions
cw-xrv57	cw-xrv58	4007	SR-MPLS	Unknown	0	...
cw-xrv51	cw-xrv52	3333	SR-MPLS	TE	0	...
cw-xrv51	cw-xrv54	10	RSVP-TE	TE	0	...
cw-xrv51	cw-xrv55	32321	RSVP-TE	IGP	0	...
cw-xrv59	cw-xrv61	312	SR-MPLS	LATENCY	0	...
cw-xrv50	cw-xrv54	16	SR-CS	TE	0	...
cw-xrv50	cw-xrv52	2022	SR-MPLS	TE	0	...
cw-xrv50	cw-xrv52	2222	SR-MPLS	TE	0	...

4 → Policy and Tunnel Change Events 📅 06-Aug-2023 16:55 to 07-Aug-2023 16:55 1M 1W 1D 1H | Reset


Policy / Tunnel Type: All SR-MPLS SRv6 Tree-SID RSVP-TE Total 7

Headend	Endpoint	Color / ID	Policy / Tunnel Type	Metric Type	Events			Actions
					Total ↓	Operational State	Change	
cw-xrv57	cw-xrv58	4005	SR-MPLS	Unknown	6	5	1	...
cw-xrv57	cw-xrv58	4006	SR-MPLS	Unknown	4	3	1	...
cw-xrv57	cw-xrv58	4007	SR-MPLS	Unknown	4	3	1	...
cw-xrv57	cw-xrv58	4004	SR-MPLS	Unknown	2	1	1	...
cw-xrv59	cw-xrv61	312	SR-MPLS	LATENCY	2	1	1	...
cw-xrv57	cw-xrv58	4003	SR-MPLS	Unknown	1	0	1	...
cw-xrv50	cw-xrv52	2022	SR-MPLS	TE	1	0	1	...

476133



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

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



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

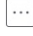
View TE Event and Utilization History

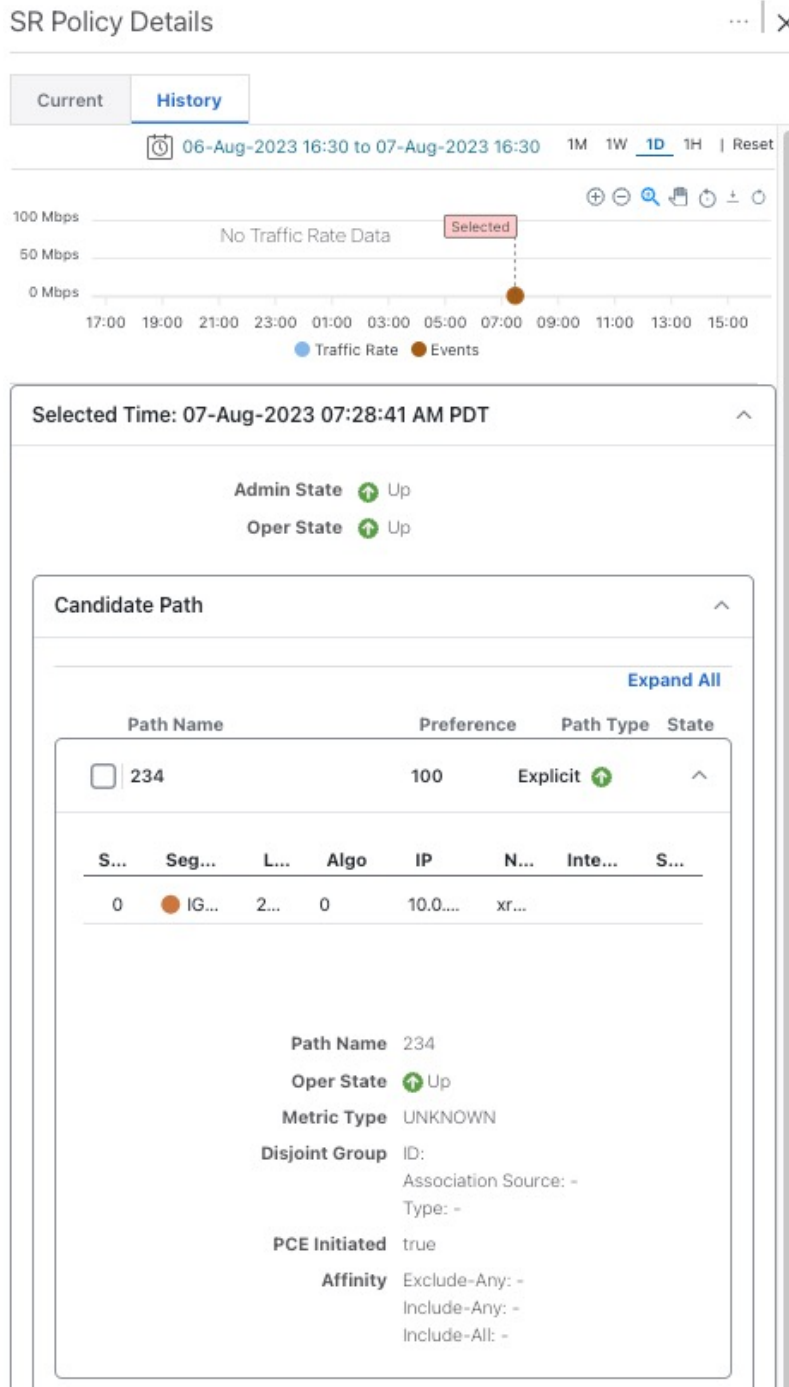
The historical data captures the traffic rate and change events for a policy or tunnel. To view the historical data:



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 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.

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.

Device Details

Details Links Traffic Engineering

General SR-MPLS SRv6 Tree-SID RSVP-TE Flex Algo

Selected 0 / Total 9

	Root Name	Root IP	Name	Tree ID	Label	Type	Programming St...	Fast Reroute	PCE Address	Admin Status	Oper Status	Actions
<input type="checkbox"/>	xrv9k-13	192.168.0.3	DAY_0_TREE_SID	-	35	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-17	192.168.0.7	MY_FIRST_TRR_SID	-	15200	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	R4_TREE_SID	-	22	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	TREE_SID_ERROR...	-	3233	Static	None	Disable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	TREE_SID_MULTL...	-	220	Static	None	Disable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-17	192.168.0.7	Test-TS1	-	1234	Static	None	Disable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-17	192.168.0.7	Test-TS3	-	8989	Static	None	Disable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-13	192.168.0.3	netflix	-	15202	Static	None	Enable	172.27.226.118			...
<input type="checkbox"/>	xrv9k-14	192.168.0.4	tree-sid-oper-error	-	4578	Static	None	Disable	172.27.226.118			...

Configure Traffic Engineering Settings

Configure TE Timeout Settings

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



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

Configure How Device Groups Are Displayed for Traffic Engineering

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

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

Configure Historical Data Settings

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

The screenshot displays the configuration interface for the Cisco Crosswork Network Controller. The 'System Settings' tab is active, and the 'Historical Data' option is selected in the left-hand navigation menu. The 'Historical Data Settings' section is visible on the right, containing the following configuration options:

- LSP Traffic Rate:** Enabled (On)
- LSP State Change:** Enabled (On)
- LSP Path Change:** Enabled (On)
- Retention Interval:** Set to 2 days. A note below the input field indicates 'Range: 1 to 30 days'.

Table 2: Available Historical Data Setting Options

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

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 will be able to view policy details, but not modify them since they were not included as part of the last data synchronization. Crosswork will display an alarm when it finds orphan TE policies (**Administration > Alarms**).

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 \(Crosswork Optimization Engine APIs > 6.0 Release APIs\)](#).



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](#).

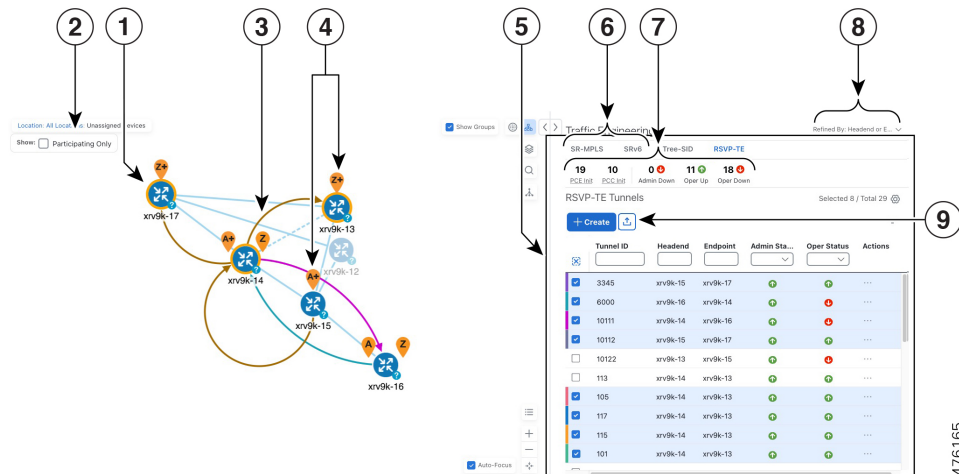
- [View SR-MPLS and SRv6 Policies on the Topology Map, on page 15](#)
- [View SR-MPLS and SRv6 Policy Details, on page 17](#)
- [Visualize IGP Path and Metrics, on page 19](#)
- [Find Multiple Candidate Paths \(MCPs\), on page 20](#)
- [Visualize Underlying Paths Associated with a Defined Binding-Segment ID \(B-SID\) Label, on page 23](#)
- [Visualize Native SR Paths, on page 24](#)
- [Configure TE Link Affinities, on page 28](#)
- [Create Explicit SR-MPLS Policies, on page 29](#)
- [Create Dynamic SR-MPLS Policies Based on Optimization Intent, on page 30](#)
- [Create SR-TE Policies \(PCC Initiated\), on page 31](#)
- [Modify SR-MPLS Policies, on page 32](#)

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 1: Traffic Engineering UI : SR-MPLS and SRv6 Policies



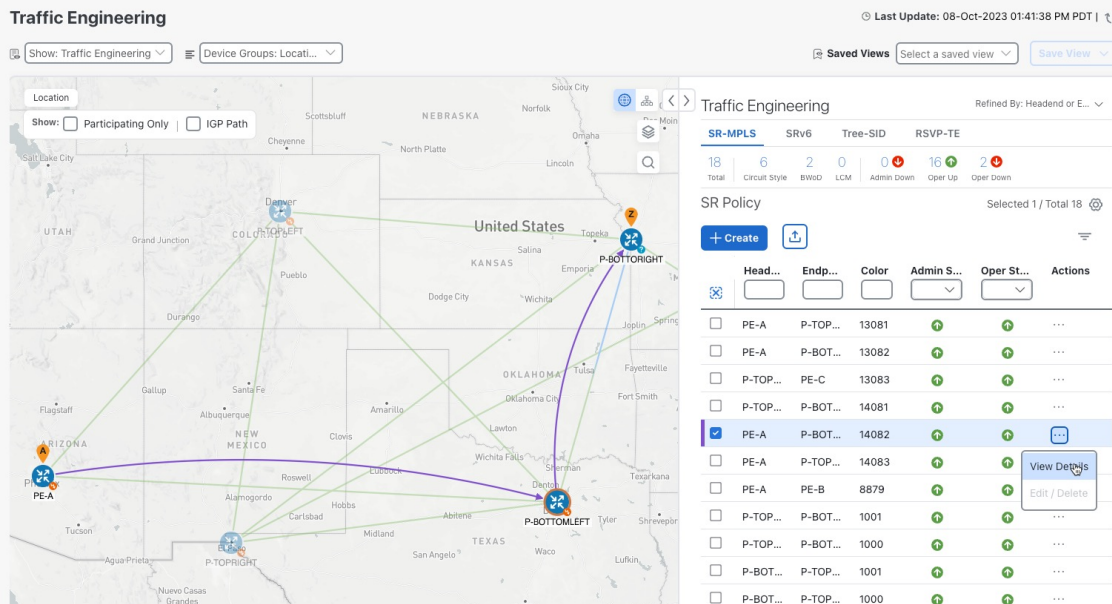
Callout No.	Description
1	A device with an orange (🔗) outline indicates there is a node SID associated with that device or a device in the cluster.
2	Click the appropriate check box to enable the following options: <ul style="list-style-type: none"> • Show IGP Path—Displays the IGP path for the selected SR-TE policy. • Show Participating Only—Displays only links that belong to selected SR-TE policy. All other links and devices disappear.
3	When SR-TE policies are selected in the SR-MPLS or SRv6 tables, they show as 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	The Mini Dashboard provides a summary of the operational SR-MPLS or SRv6 policy status. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the SR Policy and SRv6 Policy tables. In addition to the policy status, the SR-MPLS Mini Dashboard table displays the number of PCC and PCE initiated tunnels that are <i>currently</i> listed in the SR Policy table.

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

View SR-MPLS and SRv6 Policy Details

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

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



The screenshot displays the Traffic Engineering dashboard. On the left, a map of the United States shows network paths between various cities. On the right, a table lists SR Policies. The table has columns: Headend, Endp..., Color, Admin S..., and Oper St... The policy with ID 14082 is selected, and a 'View Details' button is visible in the Actions column.

	Head...	Endp...	Color	Admin S...	Oper St...	Actions
<input type="checkbox"/>	PE-A	P-TOP...	13081	+	+	...
<input type="checkbox"/>	PE-A	P-BOT...	13082	+	+	...
<input type="checkbox"/>	P-TOP...	PE-C	13083	+	+	...
<input type="checkbox"/>	P-TOP...	P-BOT...	14081	+	+	...
<input checked="" type="checkbox"/>	PE-A	P-BOT...	14082	+	+	View Details
<input type="checkbox"/>	PE-A	P-TOP...	14083	+	+	Edit / Delete
<input type="checkbox"/>	PE-A	PE-B	8879	+	+	...
<input type="checkbox"/>	P-TOP...	P-BOT...	1001	+	+	...
<input type="checkbox"/>	P-TOP...	P-BOT...	1000	+	+	...
<input type="checkbox"/>	P-BOT...	P-TOP...	1001	+	+	...
<input type="checkbox"/>	P-BOT...	P-TOP...	1000	+	+	...

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

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.

Figure 2: SR Policy Details - Headend, Endpoint, and Summary

SR Policy Details ... ✕

Clea

Current | History

Headend A PE-A | Source IP: 100.100.100.5
TE RID: 100.100.100.5
PCC IP: 100.100.100.5

Endpoint Z P-BOTTORIGHT | Dest IP: 100.100.100.4
TE RID: 100.100.100.4

Color 14082

Summary ^

- Admin State** ↑ Up
- Oper State** ↑ Up
- Binding SID** 24027
- Policy Type** Regular
- Profile ID** -
- Description** -
- Traffic Rate** 0 Mbps
- Unused** True i
- Delay** 2 i
- Accumulated Metric** 0
- Delegated PCE** 172.20.118.119
- Non-delegated PCEs** 172.20.118.63
- PCE Computed Time** 01-Jan-2023 09:02:46 PM PDT
- Last Update** 05-Oct-2023 01:21:08 PM PDT

[See less](#) ^


Figure 3: SR Policy Details - Candidate Path

The screenshot displays the 'Candidate Path' configuration interface. At the top, there is a 'Collapse All' button. Below it is a table with columns: Path Name, Preference, Path Type, and State. The selected path is 'bwod_name_101236' with a preference of 100 and path type 'Explicit'. Below the table, a detailed view of the path is shown with the following properties:


Path Name	bwod_name_101236
Oper State	Up Active
Metric Type	UNKNOWN
Bandwidth	Requested: 0 Mbps Reserved: 0 Mbps
Disjoint Group	ID: Association Source: - Type: -
PCE Initiated	true
Affinity	Exclude-Any: - Include-Any: - Include-All: -
Segment Type	Protected
SID Algorithm	-

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, with 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 in order to view metrics.

Find Multiple Candidate Paths (MCPs)



Find Multiple Candidate Paths (MCPs)

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

Important Notes

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

Before you begin


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

Step 1 From the main menu, choose **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.

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

In this example, you see that the active path is going from **xrv9k-12 > xrv9k-14 > xrv9k-16**.



Step 3 View the list of candidate paths.

- a) From the SR-TE Policy table **Actions** column, click > **View Details**. A list of candidate paths appear along with policy details in the **SR Policy Details** window. The green A in the state column indicates the active path.

Find Multiple Candidate Paths (MCPs)

SR Policy Details ... | X

Current **History**

Headend A xrv9k-12 | Source IP: 192.168.0.2
TE RID: 192.168.0.2 | IPv6 RID: 2001:192:168::2
PCC IP: 192.168.0.2

Endpoint Z xrv9k-16 | Dest IP: 192.168.0.6
TE RID: 192.168.0.6 | IPv6 RID: 2001:192:168::6

Color 2023

Summary

Admin State Up

Oper State Up

Binding SID 24012

Policy Type Regular

Profile ID -

Description -

Traffic Rate 0 Mbps

Unused True ⓘ

[See more](#) ▾

Candidate Path

[Collapse All](#)

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> cfg_sr2023_discr_100	100	Unknown	Up A
<input type="checkbox"/> cfg_sr2023_discr_50	50	Unknown	Down
<input type="checkbox"/> cfg_sr2023_discr_25	25	Unknown	Down

Path Name cfg_sr2023_discr_25

Oper State Down

Metric Type TE

Bandwidth -

Disjoint Group ID: -
Association Source: -
Type: -

PCE Initiated false

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

Segment Type Unprotected

SID Algorithm -

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

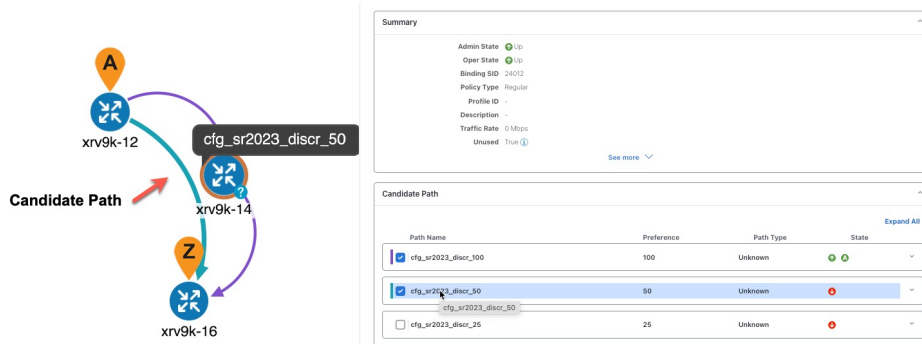
Step 5 Visualize the candidate path on the topology map.

- 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 **xrv9k-12** > **xrv-16**.



Summary

Admin State: Up
Oper State: Up
Binding SID: 24012
Policy Type: Regular
Profile ID: -
Description: -
Traffic Rate: 0 Mbps
Unused: True

See more

Candidate Path

Path Name	Preference	Path Type	State
cfg_sr2023_discr_100	100	Unknown	Up
cfg_sr2023_discr_50	50	Unknown	Down
cfg_sr2023_discr_25	25	Unknown	Down

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

Cisco Crosswork 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:

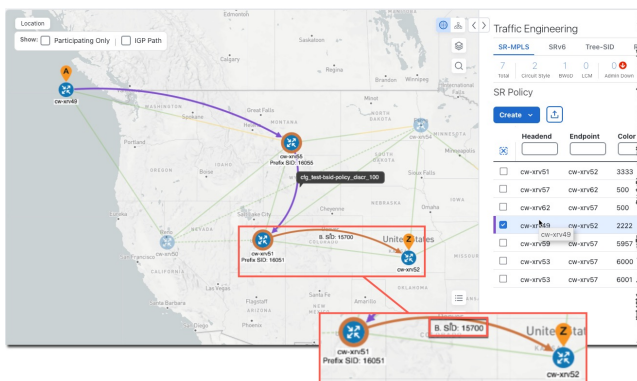
Step 1

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**.



Step 3

From the **Actions** column, click **...** > **View Details**.

Step 4

From the **SR Policy Details** window, expand the active path name to view more information. In this example, the underlying path actually goes from **cw-xrv51 > cw-xrv54 > cw-xrv53 > cw-xrv52**.

SR Policy Details

Current History

Headend cw-xrv51 | Source IP: 3.3.3.51
TE RID: 3.3.3.51 | IPv6 RID: fe00:3:3::51
PCC IP: 3.3.3.51

Endpoint cw-xrv52 | Dest IP: 3.3.3.52
TE RID: 3.3.3.52 | IPv6 RID: fe00:3:3::52

Color: 3333

Summary

Admin State Up
Oper State Up
Binding SID: 15700
Policy Type: Regular
Profile ID: -
Description: -
Traffic Rate: 0 Mbps
Unused: True

See more

Candidate Path

Collapse All

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> cfg_bsid-51-52_discr_100	100	Unknown	

Segment	Segment Type	Label	Algo	IP	Node	Interface	SID Type
0	Node SID	16054	0	3.3.3.54	cw-xrv54		Reg
1	Node SID	16053	0	3.3.3.53	cw-xrv53		Reg
2	Node SID	16052	0	3.3.3.52	cw-xrv52		Reg

Path Name: cfg_bsid-51-52_discr_100
Oper State: Up | Active
Metric Type: TE
Bandwidth: -
Disjoint Group ID: Association Source: -
Type: -
PCE Initiated: false
Affinity: Exclude-Any: -
Include-Any: -
Include-All: -
Segment Type: Unprotected
SID Algorithm: -

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 26](#).

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** On the Path Query dashboard, click **New Query**.
- Step 3** Enter the device information in the required fields to find available Native SR IGP Paths.
- Step 4** Click **Get Paths**. The Running Query ID pop-up appears.

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 turn green, completed, it can be viewed.

Step 5 Click **View Result** when it becomes available on the Running Query ID pop-up. The Path Details panel appears with corresponding available paths details while the defined topology map appears with the available Native SR IGP Paths on the left.

Example:

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

Figure 4: Path Details

The screenshot displays the 'Path Query' dashboard. On the left, a map of the United States shows a network topology with nodes labeled 'TOP-LEFT', 'TOP-RIGHT', 'BOTTOM-LEFT', and 'BOTTOM-RIGHT', and paths connecting them. On the right, a 'New Path Query' dialog box is open, allowing users to select a service and instance, and define the headend and endpoint. The dialog includes fields for 'Select Service' (with 'Select Type' and 'Select Instance' dropdowns), 'Headend' (with a search icon and text 'Enter host name, or select node on the map..'), and 'Endpoint' (with a search icon and text 'Enter host name, or select node on the map..'). Buttons for 'Get Paths' and 'Cancel' are at the bottom right of the dialog. The top right of the dashboard shows the last update time: '© Last Update: 09-Nov-2023 06:02:31 PM IST | ↻'.

Step 6 From the main menu, choose **Services & Traffic Engineering > Path Query** to return to the Path Query dashboard.

Step 7 From the **Actions** column, click **View Details**.

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

Step 8 From the available paths, click **Path 0** to expand and view the active path.

Example:

Figure 5: Path Details

Path Details

Select Service: L3vpn-Service | NSS-L3-Shared-540-internal

Headend: PE-C (100.100.10.7)

Endpoint: PE-B (100.100.10.8)

Available Paths

Path 0

Status: Found

Output: GigabitEthernet0/0/0/0

Nexthop: 20.20.20.30

Source: 100.100.10.7

Destination: 127.0.0.0

Hop Details [What is ret code/char?](#)

Hop Index:0 | Hop Origin IP:100.100.10.7 | Hop Destination IP:20.20.20.30 | MRU:1500 | Labels:[16006] | ret code:0 | multipaths:0

Hop Index:1 | Hop Origin IP:20.20.20.30 | Hop Destination IP:20.20.20.21 | MRU:1500 | Labels:[16006] | ret code:8 | return char:L | multipaths:1

Hop Index:2 | Hop Origin IP:20.20.20.21 | Hop Destination IP:20.20.20.26 | MRU:1500 | Labels:[implicit-null] | ret code:8 | return char:L | multipaths:1

Hop Index:3 | Hop Origin IP:20.20.20.26 | MRU:0 | ret code:3 | return char: | multipaths:0

Visualize Native Path Device Prerequisites

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

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

```


tpa
vrf default
address-family ipv4
default-route mgmt
!
address-family ipv6
default-route mgmt
!
!
!
or

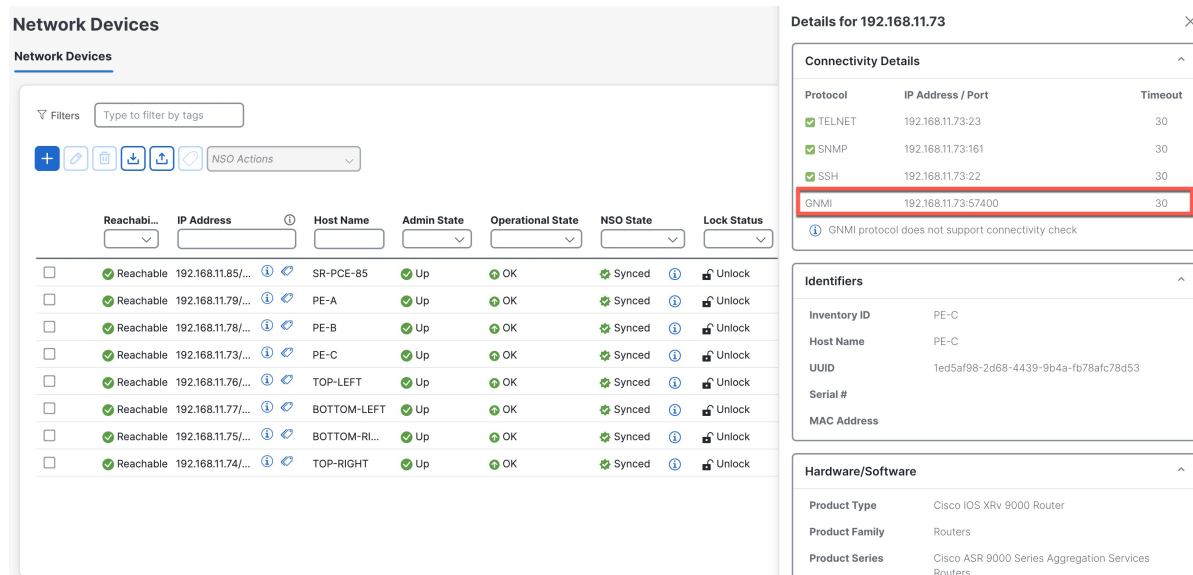
linux networking
vrf default
address-family ipv4
default-route software-forwarding
!
address-family ipv6
default-route software-forwarding
!
!
!
```

**Note**

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

3. Devices should have GNMI capability enabled and configured.

- From Device Management > Network Devices, click  icon for the device you are interested.
- Confirm that GNMI is listed under Connectivity Details.



The screenshot shows the 'Network Devices' management interface. On the left, a table lists several devices with columns for Reachability, IP Address, Host Name, Admin State, Operational State, NSO State, and Lock Status. The device 192.168.11.73 is highlighted. On the right, the 'Details for 192.168.11.73' panel is open, showing 'Connectivity Details' with a table of protocols and their configurations. The GNMI entry is highlighted with a red box.

Protocol	IP Address / Port	Timeout
TELNET	192.168.11.73-23	30
SNMP	192.168.11.73-161	30
SSH	192.168.11.73-22	30
GNMI	192.168.11.73-57400	30

Below the connectivity table, a note states: "GNMI protocol does not support connectivity check".

The 'Identifiers' section shows: Inventory ID: PE-C, Host Name: PE-C, UUID: 1ed5af98-2d68-4439-9b4a-fb78afc78d53, Serial #: , MAC Address: .

The 'Hardware/Software' section shows: Product Type: Cisco IOS XRv 9000 Router, Product Family: Routers, Product Series: Cisco ASR 9000 Series Aggregation Services Routers.

**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:xrivr-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**.

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

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

Example:

TE Link Affinities		Flex-Algo Affinities
+ Create		
Name ⓘ	Bit Position (0-31) ⓘ	Actions
<input type="text"/>	<input type="text"/>	
green	4	Edit Delete
blue	5	Edit Delete
red	1	Edit Delete

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

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

Create Explicit SR-MPLS Policies

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




Tip If you plan to use affinities, collect affinity information from your devices, and then map them in Cisco Crosswork before creating an explicit SR-MPLS policy. For more information, see [Configure TE Link Affinities, on page 28](#).

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

Step 2 Under **SR Policies**, 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 31](#).

Step 3 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 SR Policy table. You can also click the check box next to the policy to see it highlighted in the map.

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

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

Note On a setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see [Configure TE Timeout Settings, on page 11](#).

Create Dynamic SR-MPLS Policies Based on Optimization Intent

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




Tip 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 28](#) or [Configure Flexible Algorithm Affinities, on page 41](#).

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

Step 2 From the **SR Policy** table, 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 31](#).

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

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

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

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

Step 6 Define any applicable constraints and disjointness.

Note

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

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

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

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


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

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

Step 11 Validate the SR-MPLS policy creation:

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

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

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

Note On a scaled setup with high node, policy, or interface counts, a timeout may occur during policy deployment. To configure timeout options, see [Configure TE Timeout Settings, on page 11](#).

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 3: SR-TE Policy Configuration

Expand this:	To specify this:
name	Enter a name for this SR-TE policy.
head-end	<ul style="list-style-type: none"> 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	<p>a. Click  and enter a preference value. For example: 123</p> <p>b. Select one of the following and toggle switch to enable:</p> <ul style="list-style-type: none"> 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 SR Policy table, locate the SR-MPLS policy you are interested in and click .

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

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



CHAPTER 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, on page 33](#)
- [View RSVP-TE Tunnel Details, on page 35](#)
- [Create Explicit RSVP-TE Tunnels, on page 36](#)
- [Create Dynamic RSVP-TE Tunnels Based on Optimization Intent, on page 37](#)
- [Create RSVP-TE Tunnels \(PCC Initiated\), on page 38](#)
- [Modify RSVP-TE Tunnels, on page 39](#)


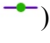
View RSVP-TE Tunnels on the Topology Map

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

Figure 6: Traffic Engineering UI - RSVP-TE Tunnels

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


476164

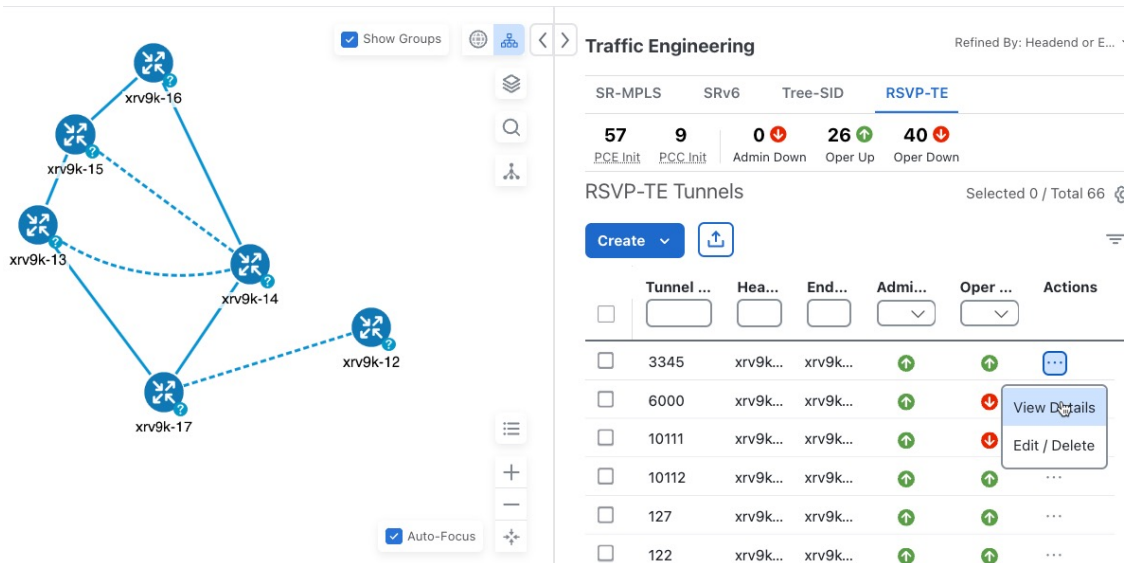
Callout No.	Description
1	Click Show Participating Only to display links that only belong to the selected RSVP-TE tunnels. All other links and devices disappear.
2	<p>A device with a solid orange outline () indicates that it is a strict hop. A dashed orange outline indicates that a loose hop was discovered.</p> <p>Note RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.</p>
3	<p>When RSVP-TE tunnels are selected in the RSVP-TE Tunnel table, they show as colored directional lines on the map indicating source and destination.</p> <ul style="list-style-type: none"> Record Route Object (RRO) paths are shown as straight lines. Explicit Route Object (ERO) paths are shown as curved lines. <p>Note If both RRO and ERO paths are available, the RRO path is displayed by default.</p> <ul style="list-style-type: none"> An adjacency segment ID (SID) is shown as a green dot on a link along the path () <p>If both A and Z are displayed in a device cluster, at least one node in the cluster is a source and another is a destination. The A+ denotes that there is more than one RSVP-TE tunnel that originates from a node. The Z+ denotes that the node is a destination for more than one RSVP-TE tunnel.</p>
4	<p>SR-MPLS and SRv6 Policy Origin and Destination: If both A and Z are displayed in a device cluster, at least one node in the cluster is a source and another is a destination. The A+ denotes that there is more than one SR-TE policy that originates from a node. The Z+ denotes that the node is a destination for more than one SR policy.</p>
5	<p>The content of this window depends on what has been selected or filtered. In this example, the RSVP-TE tab is selected and the RSVP-TE Tunnels table is displayed. Depending on what is selected on the topology map, or whether you are in the process of viewing and managing RSVP-TE tunnels, you can do the following:</p> <ul style="list-style-type: none"> Create Dynamic RSVP-TE Tunnels Based on Optimization Intent, on page 37 Create Explicit RSVP-TE Tunnels, on page 36 Modify RSVP-TE Tunnels, on page 39 View RSVP-TE Tunnel Details, on page 35
6	Click the RSVP-TE tab.
7	The Mini Dashboard provides a summary of the operational RSVP-TE tunnel status and the number of PCC and PCE initiated tunnels that are <i>currently</i> listed in the RSVP-TE tables. If filters are applied, the Mini Dashboard is updated to reflect what is displayed in the RSVP-TE table.

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

View RSVP-TE Tunnel Details

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

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



The screenshot displays the 'Traffic Engineering' section of the Cisco Crosswork Network Controller. On the left, a network diagram shows nodes xrv9k-12 through xrv9k-17 connected by solid and dashed lines. On the right, a table lists RSVP-TE tunnels. The table has columns for Tunnel ID, Headend, End, Admin, Oper, and Actions. A 'View Details' button is highlighted in the Actions column for the tunnel with ID 6000.

Tunnel ...	Hea...	End...	Admi...	Oper ...	Actions
3345	xrv9k...	xrv9k...	↑	↑	...
6000	xrv9k...	xrv9k...	↑	↓	View Details
10111	xrv9k...	xrv9k...	↑	↓	Edit / Delete
10112	xrv9k...	xrv9k...	↑	↑	...
127	xrv9k...	xrv9k...	↑	↑	...
122	xrv9k...	xrv9k...	↑	↑	...

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. Hover your mouse over the "i" icon (next to the Delay value) to view the last time the value was updated.

> RSVP-TE Tunnel Details
⋮ | ✕

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

Summary

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

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

Explicit Route Object (ERO)

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

Create Explicit RSVP-TE Tunnels


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

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

Step 2 From the right window, click **RSVP-TE**.

Step 3 Under **RSVP-TE Tunnels**, 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 38.

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 **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.

Create Dynamic RSVP-TE Tunnels Based on Optimization Intent

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



Tip If you plan to use affinities, collect affinity information from your devices and then map them in Cisco Crosswork before creating a dynamic RSVP-TE tunnel. For more information, see [Configure TE Link Affinities](#), on page 28.

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

Step 2 From the right window, click **RSVP-TE**.

Step 3 Under **RSVP-TE Tunnels**, 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 38.

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.

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**.
- Step 2** From the **Traffic Engineering** window select the **RSVP-TE** tab.
- Step 3** Locate the RSVP-TE tunnel you are interested in and click .
- Step 4** Choose **View** or **Edit/Delete**.
- Note**
- You can only modify or delete RSVP-TE tunnels that have been created with the UI or API.
 - After updating the RSVP-TE tunnel details, you can preview the changes on the map before saving it.
-



CHAPTER 4

Flexible Algorithms

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

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

This section contains the following topics:

- [Configure Flexible Algorithm Affinities, on page 41](#)
- [Visualize Flexible Algorithm Topologies, on page 42](#)
- [View Flexible Algorithm Details , on page 43](#)

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 > System Settings > Traffic Engineering > Affinity > Flex-Algorithm Affinities** tab.
- Step 2** To add a new Flexible Algorithm affinity mapping, click **+ Create**.
- Step 3** Enter the name and the bit it will be assigned. For example (using the above configuration):

Example:

Name ?	Bit Position (0-255) ?	Actions
<input type="text"/>	<input type="text"/>	
b33	33	Edit Delete
red	1	Edit Delete
blue	5	Edit Delete

- Step 4** Click **Save** to save the mapping. To view all Flexible Algorithm affinities for a link, see [View Flexible Algorithm Details](#), on page 43.

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 30.

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

Example:



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

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

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:

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

The screenshot shows the 'Device Details' window with the 'Traffic Engineering' tab selected. Under the 'Flex Algo' sub-tab, the configuration for 'Algo 128' is displayed. The configuration includes:

- IGP: Domain ID: 300, ISIS System ID: 0000.0000.0062, Level: 2
- Participating: Yes
- Elected Definition: Metric Type: LATENCY
- Exclude-Any Affinity:
- Include-Any Affinity:
- Include-All Affinity:
- Advertised: Yes
- Priority: 128
- Definition Equal to Local: Yes

Below the configuration for Algo 128, there are expandable sections for Algo 129, Algo 130, and Algo 131.

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

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

- From the topology map, click on a link.
- 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.

The screenshot shows the 'Link Details' window with the 'Traffic Engineering' tab selected. The 'Summary' section is active, and the 'General' sub-tab is selected. The following table displays the Flexible Algorithm topology information for the link:

	A Side	Z Side
Node	cw-xrv61	cw-xrv62
IF Name	GigabitEthernet0/0/0/3	GigabitEthernet0/0/0/3
FA Affinities		
FA TE Metric	1	23
FA Delay Metric		
FA Topologies	128, 129, 130, 131	128, 129, 130, 131

Note

- Application-Specific Link Attribute (ASLA) is supported on PCC and core routers that are Cisco IOS XR 7.4.1 or later versions.
 - Cisco Crosswork 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, which represents all of the tree segments and devices in it. The SDN controller is highly flexible, as it understands the segmentation and can construct routing paths using any kind of constraints that network architects can specify.

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

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

This section contains the following topics.

- [Visualize Tree-SID Policies, on page 47](#)
- [View a Point-to-Multipoint Tree on the Topology Map, on page 48](#)
- [Create Static Tree-SID Policies, on page 52](#)
- [Modify a Tree-SID Policy, on page 55](#)

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 48](#)).

The Tree-SID policy has the following nodes:

- Root node—Encapsulates the multicast traffic, replicates it, and forwards it to the transit nodes.
- Transit node—Acts as a leaf (egress) node as well as a mid-point (transit) node toward the downstream sub-tree.
- Leaf node—Decapsulates the multicast traffic and forward it to the multicast receivers.

- Bud Node—Has a separate leaf node path and is displayed separately in the topology map.

You can visualize the following Tree-SID policies:

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



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

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

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

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

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

The screenshot displays the Crosswork Traffic Engineering interface. On the left, a map shows a network topology with nodes and links. A purple path highlights a Tree-SID policy. The path starts at a root node (R) labeled 'xrv9k-26' in San Francisco, goes through transit nodes 'xrv9k-23' in San Jose and 'xrv9k-27' in San Jose, and ends at leaf nodes 'xrv9k-24' in San Jose and 'xrv9k-27' in San Jose. A bud node 'xrv9k-27' is also shown in San Jose.

On the right, the 'Tree-SID Policy Details' panel is open, showing the configuration for two policies: 'xrv9k-27' and 'xrv9k-24'.

Leaf Node Name	Leaf Node IP	Collapse All		
<input checked="" type="checkbox"/> xrv9k-27	192.168.0.27			
		Node	Egress Link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29
Transit	xrv9k-23	192.168.0.23	10.0.0.41	10.0.0.42
Leaf	xrv9k-27	192.168.0.27	-	-

Leaf Node Name	Leaf Node IP	Collapse All		
<input checked="" type="checkbox"/> xrv9k-24	192.168.0.24			
		Node	Egress Link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29
Transit	xrv9k-23	192.168.0.23	10.0.0.41	10.0.0.42
Bud	xrv9k-27	192.168.0.27	10.0.0.46	10.0.0.45
Leaf	xrv9k-24	192.168.0.24	-	-

Before you begin

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

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

Step 2 Select the Tree-SID policies you want to view on the topology map.

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

The screenshot displays the Cisco Traffic Engineering interface. On the left, a map shows a network topology with nodes and connections. On the right, a panel titled 'Traffic Engineering' shows a summary of policies and a table of selected Tree-SID policies.

Traffic Engineering Summary:

SR-MPLS	SRV6	Tree-SID	RSVP-TE
2	0	2	0
Total	Dynamic	Static	Admin Down
		Oper Up	Oper Down

Tree-SID Policy Table:

Roos...	Roos...	Name	Tre...	Label	Admi...	Oper ...	Action
<input checked="" type="checkbox"/>	xrv9K...	192.1...	Disney	-	152001	+	+
<input checked="" type="checkbox"/>	xrv9K...	192.1...	MY_F...	-	15200	+	+

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


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

Traffic Engineering Last Update: 13-Nov-2023 04:13:51 PM PST | ↕


Show: Traffic Engineering | Device Groups: Locat... | Saved Views | Select a saved view | Save View

Tree-SID Policy Details


Current | History


Root  xrv9k-26 | Root IP: 192.168.0.26
TE RID: 192.168.0.26 | IPv6 RID: 2001:192:168::26

Name: Disney


Tree ID: - 

Summary

Admin State  Up

Oper Status  Up

Label: 152001

Type: Static 

Programming State: None




Metric Type: TE

Constraints: Exclude-Any: -
Include-Any: -
Include-All: -

FRR Protected: Disable

[See more](#) ↓

Tree-SID path

Leaf Node Name	Leaf Node IP			
<input checked="" type="checkbox"/> xrv9k-27	192.168.0.27	Collapse All		
Node				
Egress Link				
Role	Name	IP	Local IP	Remote IP
 Root	xrv9k-26	192.168.0.26	10.0.0.30	10.0.0.29
 Transit	xrv9k-23	192.168.0.23	10.0.0.41	10.0.0.42
 Leaf	xrv9k-27	192.168.0.27	-	-

xrv9k-24 | 192.168.0.24

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

Note When viewing Tree-SID Policy Details, if a Source Node is not available, 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.

The screenshot displays the 'Tree-SID Policy Details' window. It includes tabs for 'Current' and 'History'. The main content area shows the following details:

- Root:** xrv9k-12 | Root IP: 192.168.0.2
TE RID: 192.168.0.2 | IPv6 RID: 2001:192:168::2
- Name:** Test-TS2
- Tree ID:** - ⓘ

The 'Summary' section provides further configuration details:

- Admin State:** Up (green arrow icon)
- Oper Status:** Down (red arrow icon) - A tooltip indicates: 'source node 192.168.0.2 not connected via PCEP'
- Label:** 9999
- Type:** Static ⓘ
- Programming State:** None
- Metric Type:** TE
- Constraints:** Exclude-Any: -
Include-Any: -
Include-All: -
- FRR Protected:** Disable

A 'See more' link with a downward arrow is located at the bottom of the summary section.

Note A (Compute) label, next to the SR-PCE field, details the SR-PCE used to create the policies.

Tree-SID Policy Details

Current History

Summary

Admin State ↑ Up

Oper Status ↓ Down ⚠

Label 9999

Type Static i

Programming State None

Metric Type TE

Constraints Exclude-Any: -
Include-Any: -
Include-All: -

FRR Protected Disable

Node Count Leaf: 3 | Bud: 0 | Transit: 0

Path Compute Elements (SR-PCEs) 172.27.226.118(Compute)

Last Update 01-Aug-2023 07:23:41 PM CDT

[See less](#) ^

Create Static Tree-SID Policies

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



Tip If you plan to use affinities, collect affinity information from your devices and then map them in Cisco Crosswork before creating a static Tree-SID policy. For more information, see [Configure TE Link Affinities, on page 28](#).

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 i to view a description of the field.

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

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

Step 4 Under **Affinity**, define any applicable affinities. For more information on Affinities, see [Configure TE Link Affinities, on page 28](#)

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

Step 6 Validate the Tree-SID policy creation:

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

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

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

Static Tree-SID Policy Configuration Example through Crosswork UI

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

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

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

```

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

```

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

```

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

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

```

```

PLSP-ID: 205
Role: Ingress
Hops:
  Incoming: 505254 CC-ID: 1
  Outgoing: 505254 CC-ID: 1 (11.1.28.54)
  Outgoing: 505254 CC-ID: 1 (11.1.1.51)

```

Modify a Tree-SID Policy


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



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

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

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

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

Step 4 Choose **Edit**.

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

Tree-SID Important Notes

Limitation

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

Visualization of Tree-SID Paths with Missing Nodes

Missing Tree-SID nodes can cause the following to happen:

A node on a Tree-SID policy path may not be available in the Crosswork topology information for various reasons. For example, the node is not added to the Crosswork device inventory. This affects how the Tree-SID policy path is displayed on the topology map. The unavailable node will cause one or more root-to-leaf paths

to look broken on the map, as shown in the figure below. However, the path details in the right panel will still show the full path.

The screenshot shows a network management interface with a network map on the left and a detailed view of a Tree-SID path on the right. The network map shows several nodes connected by lines, with a path highlighted in purple. The detailed view on the right includes a table of nodes and their roles, with one row highlighted in red.

Leaf Node Name	Leaf Node IP	Egress Link	
<input checked="" type="checkbox"/> xrv9k-VM11-771-15l Role: Node Name: xrv9k-VM11-771-15l IP: 192.168.4.14	192.168.4.14	Collapse All	
<input checked="" type="checkbox"/> xrv9k-VM3-... Role: Root Name: xrv9k-VM3-... IP: 192.168.4.3	192.168.4.3	Local IP: 10.0.2.25	Remote IP: 10.0.2.26
<input checked="" type="checkbox"/> xrv9k-VM5-... Role: Bud Name: xrv9k-VM5-... IP: 192.168.4.5	192.168.4.5	Local IP: 20.10.0.14	Remote IP: 20.10.0.15
<input checked="" type="checkbox"/> xrv9k-VM8-... Role: Transit Name: xrv9k-VM8-... IP: 192.168.4.9	192.168.4.9	Local IP: 20.10.0.17	Remote IP: 20.10.0.16
<input checked="" type="checkbox"/> xrv9k-VM7-... Role: Bud Name: xrv9k-VM7-... IP: 192.168.4.6	192.168.4.6	Local IP: 10.0.3.42	Remote IP: 10.0.3.41
<input checked="" type="checkbox"/> xrv9k-VM7_3_0_732_cco Role: Node Name: xrv9k-VM7_3_0_732_cco IP: 192.168.4.7	192.168.4.7	Egress Link	
<input checked="" type="checkbox"/> xrv9k-VM3-... Role: Root Name: xrv9k-VM3-... IP: 192.168.4.3	192.168.4.3	Local IP: 10.0.2.41	Remote IP: 10.0.2.42
<input checked="" type="checkbox"/> xrv9k-VM7-... Role: Leaf Name: xrv9k-VM7-... IP: 192.168.4.7	192.168.4.7	-	-
<input checked="" type="checkbox"/> Splitfire Role: Node Name: Splitfire IP: 192.168.4.11	192.168.4.11	Egress Link	

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2023 Cisco Systems, Inc. All rights reserved.

