Cisco Crosswork Network Controller 7.2 Traffic Engineering and Optimization

First Published: 2025-08-04

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1

Traffic Engineering in Cisco Crosswork Network Controller

Traffic Engineering in Cisco Crosswork Network Controller

Traffic engineering (TE) is a method of optimizing and steering traffic in a network to achieve specific operational goals or provide customized services, such as guaranteed bandwidth routes for prioritized traffic. By directing traffic along predetermined routes, traffic engineering enhances network performance, ensuring efficient use of available resources. One of the biggest advantages of using Crosswork Network Controller is its ability to visualize SR-TE policies and RSVP-TE tunnels on a topology map, which simplifies the provisioning and management of these policies. Furthermore, advanced traffic analysis capabilities such as the Deterministic Demand Matrix provides detailed visibility into SRv6 traffic demands and patterns, offering crucial insights for optimizing network performance and making informed traffic engineering decisions.

At the heart of effective traffic engineering are underlay networks, which form the foundational infrastructure of data transmission. Underlay transport policies govern the physical devices, routing protocols, and resource allocation within these networks, ensuring that communication is efficient, secure, and reliable. These transport policies are used in conjunction with VPN services to define, meet, and maintain SLAs between the service provider and the customer. Key technologies involved in traffic engineering include:

- SR-TE (SR-MPLS, SRv6, Flex-Algo, CS-SR)

- RSVP-TE

Supported SR-TE policies and RSVP tunnels

Crosswork Network Controller traffic engineering supports the visualization and provisioning of a variety of SR-TE policies and RSVP tunnels. It simplifies service provisioning by exposing YANG model-based forms in its UI and providing APIs for integration with external systems, while Cisco NSO acts as the underlying provisioning engine.

Additionally, Crosswork Network Controller can discover and visualize existing services that it did not create (such as brownfield service implementations) using telemetry and interaction with the SR-PCE. These services are marked as unmanaged in Crosswork Network Controller. To modify these services, administrators can use a combination of device CLI, NSO's service models or APIs, the Crosswork Network Controller UI tool set, and, in some cases, scripts to migrate pre-existing services from being unmanaged to being managed.

Operators can collaborate with Cisco CX Professional Services or use resources and articles on Cisco DevNet to customize or expand the capabilities of Crosswork Network Controller. This may include developing custom function packs tailored to specific use cases.

The table shows the TE technologies supported by Crosswork Network Controller:

Table 1: Supported TE Technologies

TE Technology	Crosswork Network Controller	
	Visualize	**Provision (PCE-initiated)**
SR-MPLS	✓	✓
SRv6	✓	✓
RSVP	✓	✓
Flexible Algorithm	✓	✗
Tree-SID	✓	✓[1]
Circuit Style	✓	✓

[1] Only static Tree-SID policies are supported. While dynamic Tree-SID policies can only be provisioned manually on devices or via APIs, they can be visualized in Crosswork Network Controller UI.

Note Crosswork Network Controller supports the use of role-based access control (RBAC) to limit not only what functions a user can perform, but also on which devices they are allowed to perform those functions. For more details, see the Cisco Crosswork Network Controller Administration Guide.

What is segment routing?

Segment routing for traffic engineering operates through a tunnel established between a source and destination pair. It uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment. Segments serve as identifiers for various types of instructions. For example, topology segments identify the next hop toward a destination. Each segment is identified by a segment ID (SID), which is an unsigned 32-bit integer. Routers within the provider's core network read and process these SIDs to forward packets along the intended path, calculated by the IGP, whereas the destination is unaware of the presence of the tunnel. Segments are stacked in the packet, and each router processes the top SID in the stack to determine the next hop.

Segment types and their roles

Interior gateway protocol (IGP) distributes two types of segments: prefix segments and adjacency segments. Each router (node) and each link (adjacency) in the network is associated with a segment identifier (SID).

- **Prefix SID**: A prefix SID is tied to an IP prefix and is manually configured from the segment routing global block (SRGB) range of labels. It is distributed by IS-IS (Intermediate System to Intermediate System) or OSPF (Open Shortest Path First). The prefix segment directs traffic along the shortest path to its destination. A **node SID** is a special type of prefix SID that identifies a specific node, typically configured on the loopback interface using the node's loopback address as the prefix. Prefix SIDs are **globally unique** within the segment routing domain.

- **Adjacency SID**: An adjacency segment is represented by an adjacency SID, a label that identifies a specific adjacency, such as an egress interface to a neighboring router. Adjacency SIDs are distributed by IS-IS or OSPF and are **locally unique** to each router. The adjacency segment steers traffic to a particular adjacency.

By combining prefix (node) and adjacency SIDs in a specific order, any desired path through the network can be constructed. Segments are arranged in a stack within the packet header. At each hop, the router reads the top segment in this stack to decide the next forwarding step.

- If the segment contains the identity of another node, the router uses Equal-Cost Multi-Path (ECMP) to forward the packet to the next hop.

- If the segment is meant for the current router, the router removes (pops) the top segment and processes the next segment in the stack.

Segment routing policies

Segment routing for traffic engineering uses a "policy" to steer traffic through a specific path in the network. An SR policy defines the path as an ordered list of segments, known as a segment ID (SID) list, where each SID represents a specific instruction, such as forwarding to a particular node or adjacency. When a packet is matched to an SR policy at the head-end router (the entry point of the policy), the router attaches (pushes) the SID list onto the packet header. As the packet traverses the network, each router reads and processes the top SID in the list, executing the corresponding forwarding instruction. This ensures that the packet follows the explicit path specified by the SR policy, regardless of the underlying IGP shortest path.

Crosswork Network Controller supports the visualization and some provisioning of these SR-TE policies:

- SR-MPLS and SRv6, on page 19

- Flexible algorithm, on page 57

- Tree Segment Identifier, on page 69
- SR Circuit Style

Dynamic vs. explicit SR policies

An SR-TE policy can use one or more candidate paths. A candidate path can be a single segment list (SID list) or a group of weighted SID lists.

- **Dynamic SR policy**: Dynamic paths are a type of candidate path computed based on optimization objectives (such as minimizing TE or IGP metrics) and constraints (like affinity or protection requirements). The head-end router typically computes these dynamic paths locally. However, if it does not have complete topology information, it can delegate the computation to a Segment Routing Path Computation Element (SR-PCE). When the topology changes, a new path is computed. This dynamic path calculation produces a sequence of interface IP addresses representing the specific links (adjacencies) along the path. Traffic engineering then maps each of these interface IP addresses to an adjacency Segment Identifier (adj-SID) label. Routes are learned and forwarded using these adjacencies over the Traffic Engineering tunnel.
- **Explicit SR policy**: The path is explicitly specified by the network operator as a fixed list of prefix or adjacency SIDs, each representing a node or link along the path. The path does not change automatically with topology changes unless manually reconfigured. This policy is used when precise control over the traffic path is required, such as for traffic engineering or compliance with specific routing policies.

Disjoint path computation

Crosswork Network Controller 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 the two computed paths should not share. The following disjoint path computations are supported:

- **Link** – Links are not shared on the computed paths.
- **Node** – Nodes are not shared on the computed paths. ensuring complete independence of routing devices.
- **SRLG** – Links with the same Share Risk Link Group (SRLG) value (representing a common risk) are not shared on the computed paths.

Important The SRLG value is displayed only for IPv4 links and is not shown for IPv6 links.

Note SRLGs are also utilized in flexible algorithm definitions to exclude links from specific topologies. For more details, see Configure and visualize flexible algorithm SRLG exclusion, on page 63.

- **SRLG-node** – SRLG and nodes are not shared on the computed paths, offering the highest level of fault isolation.

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

Note

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

Segment routing path computation element

Cisco Segment Routing Path Computation Element (SR-PCE) is a network control service that:

- uses network telemetry and topology data to analyze and compute optimal traffic engineering (TE) tunnels,
- provides stateful PCE functionality to control and reroute TE tunnels for network optimization,
- enables a Path Computation Client (PCC) to report and delegate control of headend tunnels to a PCE peer, and
- communicates with Path Computation Clients (PCCs) through the Path Computation Element Communication Protocol (PCEP) to delegate tunnel control and push real-time updates to the network.

Cisco SR-PCE is delivered as part of the Cisco IOS XR operating system, and can run on both physical devices and virtual routers within virtual machines. Crosswork Network Controller discovers all devices in the IGP domain, including those that do not establish PCEP peering with SR-PCE. Note that PCEP peering is required to deploy TE tunnels.

Note

To avoid any compatibility issues, refer to the Cisco Crosswork Network Controller Release Note for SR-PCE version support and compatibility.

For SR-PCE and HA configuration, see **Cisco SR-PCE providers** in the Cisco Crosswork Network Controller 7.2 Administration Guide.

SR-TE policy PCC and PCE configuration sources

SR-TE policies that are configured using the UI or API are the only types of policies that you can modify or delete in Crosswork Network Controller. SR-TE policies discovered and reported by Crosswork Network Controller may have been configured from these sources:

- **Path Computation Client (PCC) initiated**: Policies configured directly on a PCC (refer to PCC-initiated SR-TE policy example, on page 8). These policies display as Unknown in the UI because they are not provisioned or managed by Crosswork Network Controller. However, Bandwidth on Demand (BWoD) and Circuit Style (CS) policies are exceptions. These are not labeled as Unknown because Crosswork Network Controller recognizes and categorizes them based on their attributes and purpose, even if they are PCC-initiated.

Note Circuit Style policies are always PCC-initiated.

- **Path Computation Element (PCE) initiated**: Policies configured on a PCE or created dynamically by Crosswork Network Controller. Examples of PCE-initiated policy types include:
  - Dynamic
  - Explicit
  - Bandwidth on Demand (can be either PCC or PCE)
  - Local Congestion Mitigation
  - SR Circuit Style Manager

Resource reservation protocol

Resource Reservation Protocol-Traffic Engineering (RSVP-TE) is a signaling protocol that:

- allows systems and clients to request and reserve network resource reservations,
- creates, maintains, and deletes those reservations along explicit data paths, and
- ensures critical applications receive the necessary bandwidth and network resources to meet the desired QoS standards.

RSVP-TE capabilities

RSVP-TE provides these capabilities.

- **Endpoint control**: Establishes and manages TE tunnels at the headend and tail end of the network connection.
- **Link-management**: Enables efficient routing of TE label-switched paths (LSP) by assigning Multi-Protocol Label Switching (MPLS) labels.
- **Fast Reroute (FRR)**: Manages LSPs that need protection and assigns backup tunnel information for quick recovery in case of faults.

RSVP-TE explicit routing (Strict, Loose)

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

- **Strict paths**: Each network node and its preceding node in the ERO are directly connected and must be adjacent.

  For example. Node A → Node B → Node C, where each hop must be over a single direct physical link.

- **Loose paths**: A node in the ERO must be present in the path but is not required to be directly connected to its preceding node. When encountering a loose hop during ERO processing, the node that processes the loose hop can update the ERO with one or more nodes along the way to the next specified node. The advantage of a loose path is that the entire path does not need to be specified or known when creating the ERO.

  For example, Node A → Node X (any path, possibly multiple hops, between A and X is acceptable).

  A disadvantage of a loose path is the potential for forwarding loops that can occur during transient changes in the underlying routing protocol.

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

RSVP FRR (Fast Reroute)

RSVP FRR provides rapid LSP restoration when a failure is detected:

- If a router’s link or neighboring device fails, the router detects this via an interface-down notification.
- If an interface goes down, the router switches LSPs going out of that interface onto their respective backup tunnels, if available.

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

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

- The TE tunnel headend signals FRR requirements along the path.
- Points of Local Repair (PLRs) evaluate the FRR requirements based on the availability of backup tunnels at the PLR. These nodes are able to switch LSPs to backup tunnels if needed. If a suitable backup tunnel is available, the PLR selects it and signals the backup tunnel information to the headend.
- When an FRR event is triggered (for example, a link or node failure), the PLR sends PATH messages through the backup tunnel to the merge point (MP), where the backup tunnel rejoins the original LSP.
- The MP sends RESV messages back to the PLR using the RSVP-Hop object included by the PLR in its PATH message. This mechanism ensures that the original LSP is not torn down by the MP during the failover process.
- The PLR notifies the headend about the failure by sending a PATH-ERROR message, indicating that FRR is in use for the affected LSP. This prompts the headend to establish a new LSP for the TE tunnel while maintaining traffic flow using make-before-break techniques. Once the new LSP is operational, the headend tears down the failed path.

RSVP-TE tunnel PCC and PCE configuration sources

RSVP-TE tunnels discovered and visualized by Crosswork Network Controller might be configured from these sources:

- Path Computation Client (PCC) initiated—RSVP-TE tunnels configured directly on a PCC. See PCC-initiated RSVP-TE tunnel example, on page 9.
- Dynamically initiated—RSVP-TE tunnels dynamically computed and established by a Path Computation Element (PCE) or requested by a PCC.

Sample policy and device configurations

This section provides samples of policy and device configurations related to Traffic Engineering and Optimization functions.

To ensure that Traffic Engineering and telemetry functions operate successfully within Crosswork Network Controller, you must properly configure the devices. For more details about device configuration for other Crosswork Network Controller functions, refer to the "Onboard Devices" chapter in the Cisco Crosswork Network Controller 7.2 Administration guide.

Crosswork Network Controller can discover and visualize pre-existing services that it did not create, such as brownfield service implementations. The details of these service configurations appear in the topology screen when a policy is selected from the table. However, these policies are marked as unmanaged in Crosswork Network Controller. To modify these services, administrators can use a combination of device CLI, NSO service models or APIs, the Crosswork Network Controller UI tool set, and, in some cases, scripts to migrate pre-existing services from unmanaged to managed.

PCC-initiated SR-TE policy example

This example demonstrates the configuration of an SR-TE policy on the headend router. The policy uses a dynamic path that is computed by the headend router based on specified affinity constraints. In this example, a policy named **SampleSRTE** is created with the following attributes: a color value of 100, a candidate preference of 100, a metric of type **TE**, and an affinity constraint to exclude links assigned the color **red**.

Refer to the 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 sampleSRTE
      color 100 end-point ipv4 1.1.1.2
      candidate-paths
        preference 100
          dynamic
            metric
              type te
            !
          !
          constraints
            affinity
              exclude-any
```

```
                name RED
              !
            !
           !
         !
       !
```

Policy source-address configuration to support multiple loopback IP addresses

In order to support multiple loopback IP addresses, these policy configurations must be included on any PCC device that will act as the headend or origination point for a policy.

Global configuration for all policies

```
Router# segment-routing traffic-eng candidate-paths all source-address ipv4 ip-address
```

Configuration for a specific policy

```
Router# segment-routing traffic-eng policy policy-name source-address ipv4 ip-address
```

PCC-initiated RSVP-TE tunnel example

This example shows a device configuration for a PCC-initiated RSVP-TE tunnel. For detailed descriptions and supported RSVP-TE tunnel configuration commands for your specific device, review the relevant documentation (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
  !
```

Affinity map configurations

Affinity maps allow network operators to associate human-readable names (such as "red," "low delay," or "high bandwidth") with specific bit positions that represent link attributes. If an affinity mapping is not defined in the Crosswork Network Controller UI, the affinity name is displayed as "UNKNOWN". To configure the affinity attribute for visualization purposes as part of an SR-TE policy, Tree-SID, RSVP-TE tunnel, or any other policy supported by Crosswork Network Controller, the affinity map configured on the device must also be recreated in Crosswork Network Controller. Start by collecting the affinity mappings configured on the device, and then define the same mappings in the Crosswork Network Controller UI with matching names and bit positions.

SR-TE affinity map configuration on a device

This example is a sample SR-TE affinity mapping configuration on a device. For more information, see Configure TE link affinities, on page 37.

```
RP/0/RP0/CPU0:c12#sh running-config segment-routing traffic-eng affinity-map
   segment-routing
    traffic-eng
      affinity-map
```

```
        name red bit-position 1
        name blue bit-position 5
        name green bit-position 4
      !
    !
  !
```

Flexible algorithm affinity map configuration on a device

This example is a sample flexible algorithm affinity mapping configuration on a device. For more information, see Configure flexible algorithm affinities, on page 58.

```
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
  !
```

Note Similar to affinity maps, flexible algorithms can also use Shared Risk Link Group (SRLG) exclusion constraints to prune links from the topology. These exclusions are configured on the device and then discovered by Crosswork Network Controller. For detailed configuration and visualization steps, see Configure and visualize flexible algorithm SRLG exclusion, on page 63.

The Traffic Engineering dashboard

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

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

Figure 1: Traffic Engineering Dashboard

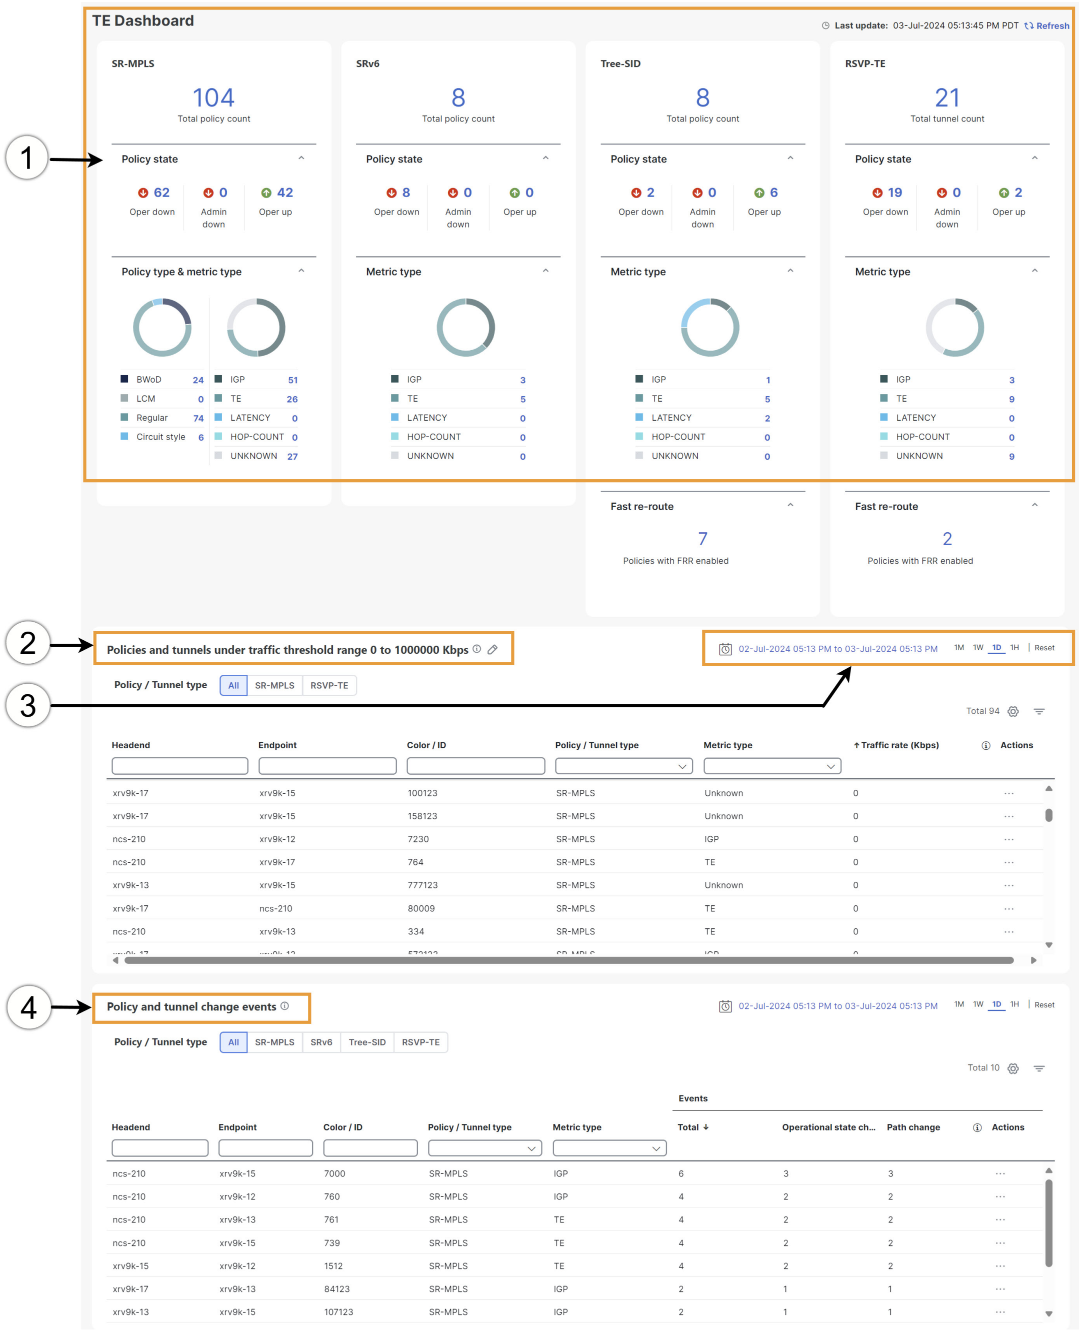

TE Dashboard

Last update: 03-Jul-2024 05:13:45 PM PDT Refresh

SR-MPLS

104
Total policy count

Policy state

62 Oper down | 0 Admin down | 42 Oper up

Policy type & metric type

BWoD	24
LCM	0
Regular	74
Circuit style	6

IGP	51
TE	26
LATENCY	0
HOP-COUNT	0
UNKNOWN	27

SRv6

8
Total policy count

Policy state

8 Oper down | 0 Admin down | 0 Oper up

Metric type

IGP	3
TE	5
LATENCY	0
HOP-COUNT	0
UNKNOWN	0

Tree-SID

8
Total policy count

Policy state

2 Oper down | 0 Admin down | 6 Oper up

Metric type

IGP	1
TE	5
LATENCY	2
HOP-COUNT	0
UNKNOWN	0

Fast re-route

7
Policies with FRR enabled

RSVP-TE

21
Total tunnel count

Policy state

19 Oper down | 0 Admin down | 2 Oper up

Metric type

IGP	3
TE	9
LATENCY	0
HOP-COUNT	0
UNKNOWN	9

Fast re-route

2
Policies with FRR enabled

Policies and tunnels under traffic threshold range 0 to 1000000 Kbps

02-Jul-2024 05:13 PM to 03-Jul-2024 05:13 PM 1M 1W 1D 1H | Reset

Policy / Tunnel type: All | SR-MPLS | RSVP-TE

Total 94

Headend	Endpoint	Color / ID	Policy / Tunnel type	Metric type	↑ Traffic rate (Kbps)	Actions
xrv9k-17	xrv9k-15	100123	SR-MPLS	Unknown	0	...
xrv9k-17	xrv9k-15	158123	SR-MPLS	Unknown	0	...
ncs-210	xrv9k-12	7230	SR-MPLS	IGP	0	...
ncs-210	xrv9k-17	764	SR-MPLS	TE	0	...
xrv9k-13	xrv9k-15	777123	SR-MPLS	Unknown	0	...
xrv9k-17	ncs-210	80009	SR-MPLS	TE	0	...
ncs-210	xrv9k-13	334	SR-MPLS	TE	0	...

Policy and tunnel change events

02-Jul-2024 05:13 PM to 03-Jul-2024 05:13 PM 1M 1W 1D 1H | Reset

Policy / Tunnel type: All | SR-MPLS | SRv6 | Tree-SID | RSVP-TE

Total 10

					Events			
Headend	Endpoint	Color / ID	Policy / Tunnel type	Metric type	Total ↓	Operational state ch...	Path change	Actions
ncs-210	xrv9k-15	7000	SR-MPLS	IGP	6	3	3	...
ncs-210	xrv9k-12	760	SR-MPLS	IGP	4	2	2	...
ncs-210	xrv9k-13	761	SR-MPLS	TE	4	2	2	...
ncs-210	xrv9k-15	739	SR-MPLS	TE	4	2	2	...
xrv9k-15	xrv9k-12	1512	SR-MPLS	TE	4	2	2	...
xrv9k-17	xrv9k-13	84123	SR-MPLS	IGP	2	1	1	...
xrv9k-13	xrv9k-15	107123	SR-MPLS	IGP	2	1	1	...

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

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

View TE event and utilization history

The TE utilization history captures the traffic rate and event changes for a policy or tunnel. Traffic rate history is not captured for SRv6 or Tree-SID policies. To view traffic engineering events and utilization history, complete these steps:

Before you begin

Ensure that LSP utilization collection is enabled and data retention is configured. See Configure TE data retention settings, on page 17.

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering**.

Step 2 In the **Actions** column for the desired policy or tunnel, choose ⋯ > **View Details** > **History**. The History page displays historical data for that device which could include the traffic rate, delay, delay variance, and loss metrics.

Figure 2: SR policy details history page

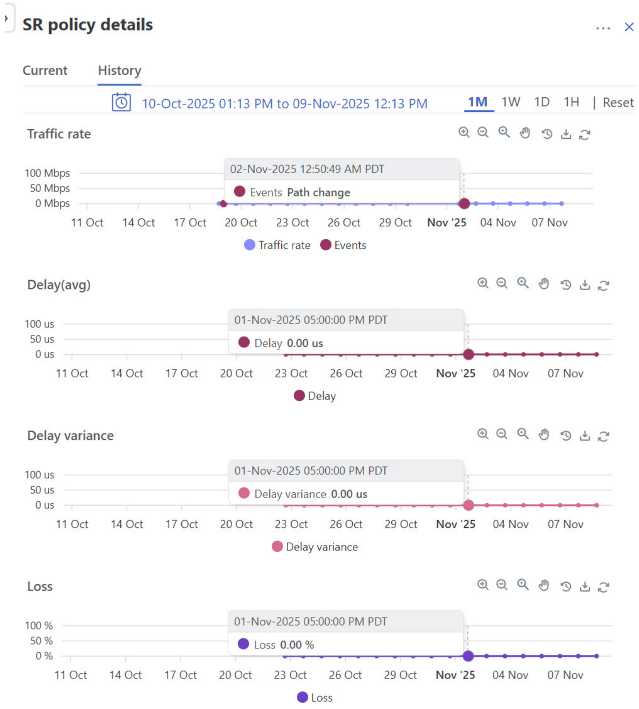

SR policy details

Current History

10-Oct-2025 01:13 PM to 09-Nov-2025 12:13 PM

1M 1W 1D 1H | Reset

Traffic rate

02-Nov-2025 12:50:49 AM PDT

Events Path change

Delay(avg)

01-Nov-2025 05:00:00 PM PDT

Delay 0.00 us

Delay variance

01-Nov-2025 05:00:00 PM PDT

Delay variance 0.00 us

Loss

01-Nov-2025 05:00:00 PM PDT

Loss 0.00 %

You can view delay, delay variance, and loss metrics for SR-MPLS and RSVP-TE policies only when Cisco Crosswork Service Health is installed and SR-PM collection is enabled. To enable monitoring, see *Enable SR PM Monitoring for Links and TE Policies* in the Cisco Crosswork Network Controller Service Health Monitoring Guide.

- The extended TE link delay metric (minimum-delay value) is used in Segment Routing (SR) policies either as an optimization metric for path calculation or as an accumulated delay bound. This metric allows monitoring of the actual end-to-end delay experienced by traffic sent over an SR policy, ensuring that the delay remains within a defined upper bound to meet Service Level Agreement (SLA) requirements.
- The Loss metric provides detailed visibility into packet loss across core network links. You can set loss severity thresholds, in the **Administration** > **Settings** > **System settings** > **Topology** > **Metric thresholds** page.
- Use the measured end-to-end delay values to determine whether to activate a candidate path or segment list. Only SR policies that meet the delay or SLA criteria should be placed into the forwarding table. If the measured delay of a candidate path exceeds the defined threshold, that path should be deactivated to prevent SLA violations.

Step 3 Click a specific event to view detailed information about path or state changes.

Figure 3: TE event and utilization history

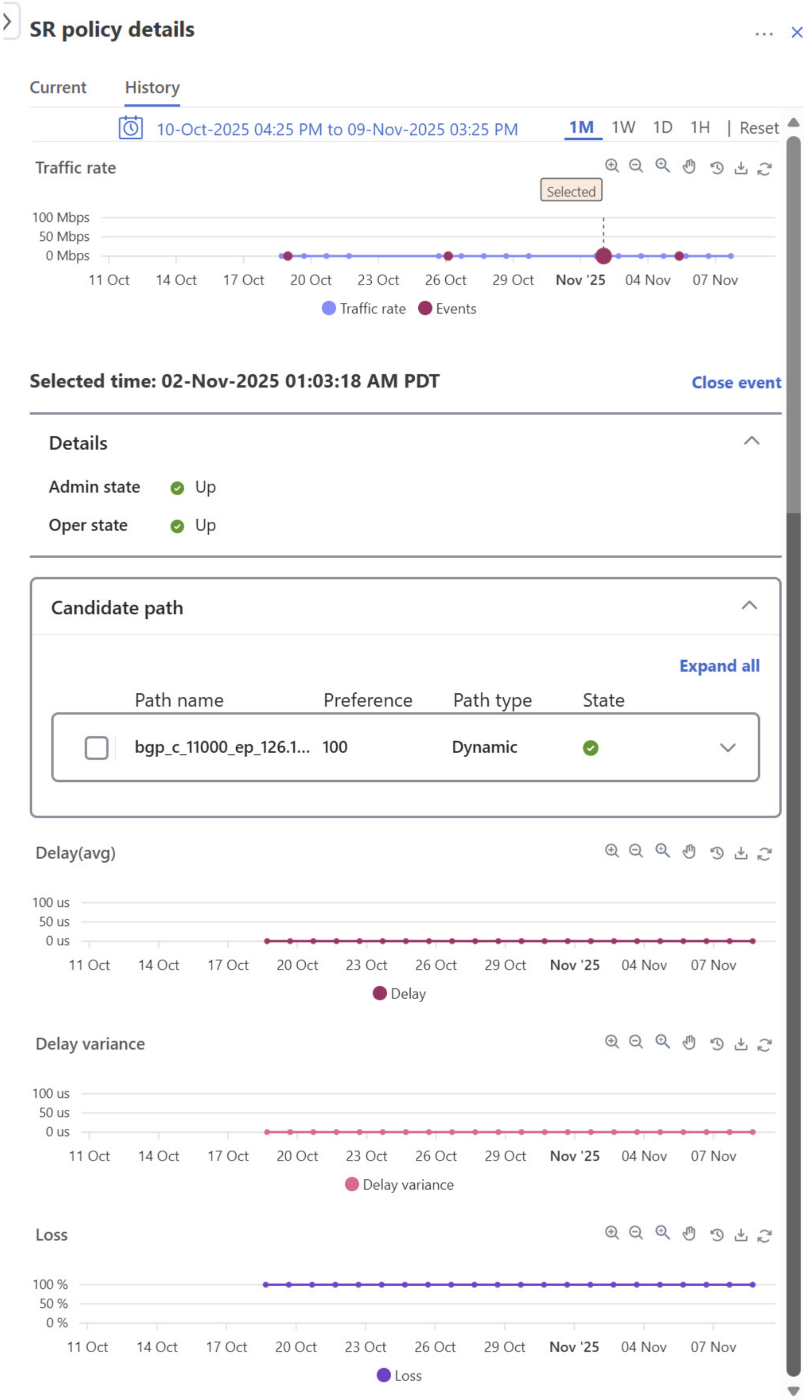

SR policy details

Current | History

10-Oct-2025 04:25 PM to 09-Nov-2025 03:25 PM **1M** 1W 1D 1H | Reset

Selected time: 02-Nov-2025 01:03:18 AM PDT Close event

Details

Admin state	✓ Up
Oper state	✓ Up

Candidate path

Expand all

	Path name	Preference	Path type	State
□	bgp_c_11000_ep_126.1...	100	Dynamic	✓

View TE device details

Follow these steps to view traffic engineering device details (SR-MPLS, SRv6, RSVP-TE, and Flexible Algorithm information).

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering**.

Step 2 In the topology map, select a device.

Step 3 Under **Device details** , choose **Traffic engineering** > ***policy-tunnel-type***. Each tab displays associated policy or tunnel data for that device.

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

Figure 4: Traffic engineering device details

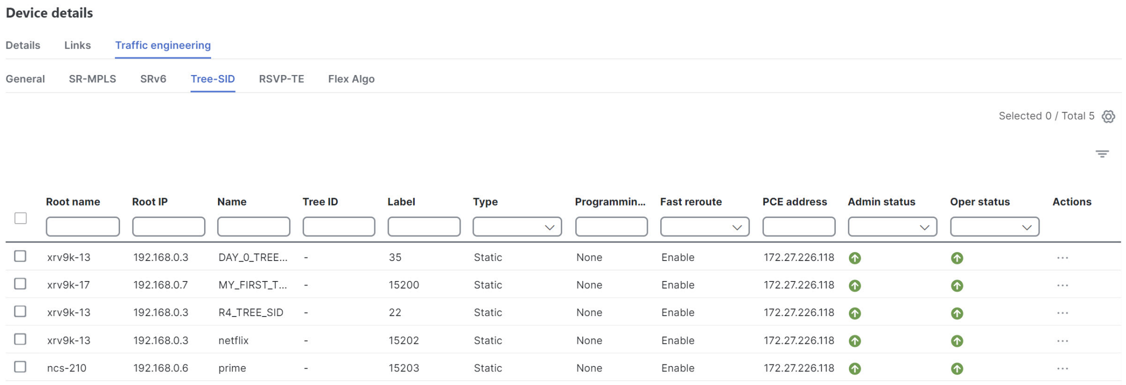

Device details

Selected 0 / Total 5

□	Root name	Root IP	Name	Tree ID	Label	Type	Programmin...	Fast reroute	PCE address	Admin status	Oper status	Actions
□	xrv9k-13	192.168.0.3	DAY_0_TREE...	-	35	Static	None	Enable	172.27.226.118	⬆	⬆	···
□	xrv9k-17	192.168.0.7	MY_FIRST_T...	-	15200	Static	None	Enable	172.27.226.118	⬆	⬆	···
□	xrv9k-13	192.168.0.3	R4_TREE_SID	-	22	Static	None	Enable	172.27.226.118	⬆	⬆	···
□	xrv9k-13	192.168.0.3	netflix	-	15202	Static	None	Enable	172.27.226.118	⬆	⬆	···
□	ncs-210	192.168.0.6	prime	-	15203	Static	None	Enable	172.27.226.118	⬆	⬆	···

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

Step 4 (Optional) To share this information you can copy the URL and send the link to others.

Receive traffic engineering notifications

Crosswork Network Controller provides robust support for real-time notifications, helping operators monitor the network and respond quickly to events that may impact service.

The system uses the YANG Data Modeling Language (RFC 7950) to define notifications for a wide range of network changes, including the creation, update, or deletion of nodes, links, interfaces, and LSPs. These notifications are delivered via RESTCONF and enable users to stay informed about their network's topology and policy state in near real-time.

Additionally, Crosswork Network Controller supports notifications for SR policies whose IGP paths are impacted by link down events. For example, when a network link goes down, the system identifies which SR policies use that link in their IGP paths and sends RESTCONF notifications to subscribed users with details about the affected link and policies. For existing policies, these notifications are sent right away; for newly configured policies, notification delivery may take up to 10 minutes.

For detailed instructions on subscribing to notifications, as well as examples of notification types and their JSON/YANG structures, see the Topology and Traffic Engineering Notifications page in the Crosswork Network Controller API documentation.

Configure TE settings

Configure TE timeout settings

If your setup includes many nodes, policies, or interfaces, a timeout may occur during policy deployment. To configure timeout settings for the provisioning and retrieval of data for SR-TE policies, RSVP-TE tunnels, Bandwidth on Demand and IGP paths, complete these steps:

Procedure

Step 1 Choose **Administration** > **Settings** > **System settings** > **Traffic engineering** > **General settings**.

Step 2 Enter the timeout duration options. For more information, click ⓘ.

Note

Timeouts affect how quickly an action is completed if SR-PCE responds slowly. You can modify these settings for large-scale topologies or to address slow SR-PCE responses caused by latency or high load.

Figure 5: Traffic engineering timeout settings

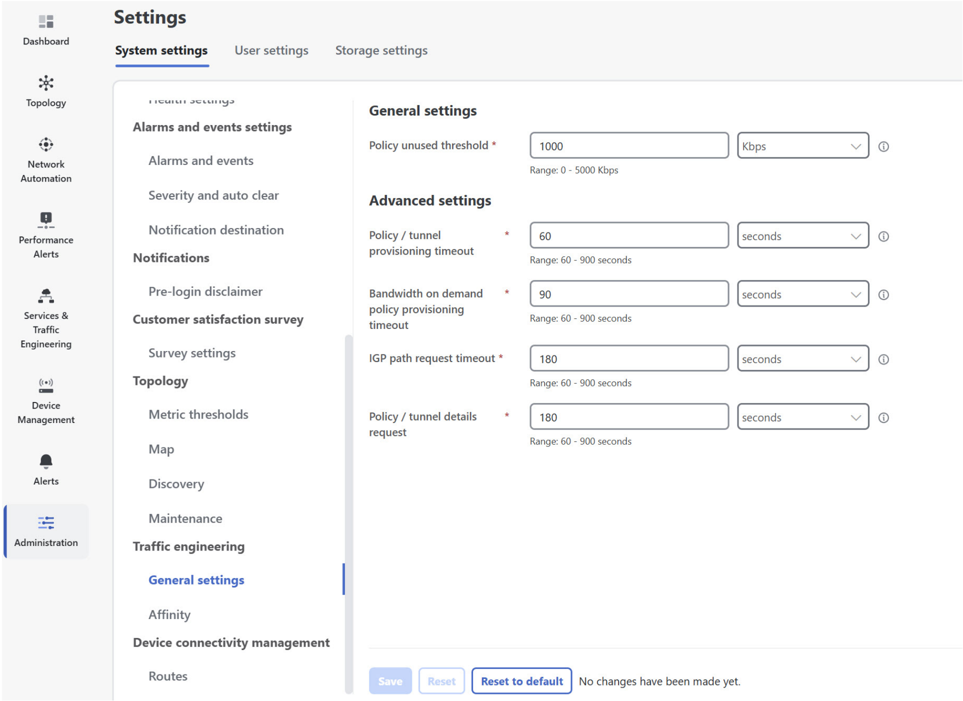

General settings

Policy unused threshold * 1000 Kbps
Range: 0 - 5000 Kbps

Advanced settings

Policy / tunnel provisioning timeout * 60 seconds
Range: 60 - 900 seconds

Bandwidth on demand policy provisioning timeout * 90 seconds
Range: 60 - 900 seconds

IGP path request timeout * 180 seconds
Range: 60 - 900 seconds

Policy / tunnel details request * 180 seconds
Range: 60 - 900 seconds

No changes have been made yet.

Configure device group display for TE

You can configure how the topology map displays devices when a device group is selected. If a device that is part of an SR policy, service, or RSVP-TE tunnel does not belong to the selected group, its status can be shown separately.

To adjust this setting, choose **Administration** > **Settings** > **User settings** > **Switch device group** and select one of the behavior options.

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

Configure TE data retention settings

To view LSP utilization history in the Historical tab, LSP utilization collection must be enabled and the retention period for collected data must be specified.

To configure TE data retention settings, complete these steps:

Procedure

Step 1 Choose **Administration** > **Settings** > **System settings** > **Data retention** > **Network performance**.

Step 2 In the **Collect metrics data** section, select the items (**LSP utilization**, **LSP PM**, and **Link PM**) for which you would like metrics to be collected and retained.

Step 3 Optionally, edit the default data retention periods according to your organization requirements.

Note
If you reduce the retention period, all data older than the new period is deleted. For example, if the daily retention interval is set to 24 hours, and then reduced to 7 days, all data older than 7 days will be deleted.

Step 4 Click **Save**.

Resolve orphaned SR-TE policies and RSVP-TE tunnels

Orphaned TE policies refer to any PCE-initiated SR-TE policies (including SRv6, SR-MPLS, and Tree-SID) or RSVP-TE tunnels that were created in the Crosswork Network Controller after the last cluster data synchronization. These orphaned policies can occur following a switchover in a High Availability setup or after a backup and restore operation. After a switchover, the system automatically checks for any orphaned TE policies or tunnels.

When orphaned TE policies or tunnels are detected, you can view their details but cannot modify them, as they were not part of the last data synchronization. The Crosswork Network Controller will raise an alarm when it identifies orphaned TE policies, which can be viewed under **Alerts** > **Alarms and Events**.

To help manage these orphans, Crosswork Network Controller provides APIs to list and clear them:

- To retrieve a list of orphaned SR-TE policies or RSVP-TE tunnels, use the following APIs with the parameter `is-orphan=True` and the default action `GET`:
  - `cisco-crosswork-optimization-engine-sr-policy-operations:sr-datalist-oper`
  - `cisco-crosswork-optimization-engine-rsvp-te-tunnel-operations:rsvp-te-datalist-oper`
- To make orphaned policies or tunnels manageable again, perform a `SAVE` action for the corresponding URL and policy type.

For more information, refer to the API documentation on Devnet (**API Reference** > **Crosswork Optimization Engine**).

CHAPTER 2

SR-MPLS and SRv6

SR-MPLS and SRv6

SR-MPLS (Segment Routing with MPLS) and SRv6 (Segment Routing with IPv6) are technologies that enable segment routing. In this approach, the source node selects a path and encodes it in the packet header as an ordered list of segments. In SR-MPLS, these segments are represented as MPLS labels. In SRv6, they are encoded directly into the IPv6 header. Each segment is identified by a Segment ID (SID), which can represent any type of instruction. For example, a SID may identify the next hop toward a destination and guides packets along the specified end-to-end path calculated by the IGP.

SR-TE policies can use one or more candidate paths. Each candidate path may consist of a single SID list or a group of weighted SID lists. When a packet is directed into an SR-TE policy, the head end adds the SID list to the packet, and the rest of the network executes the instructions embedded in that list.

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

SR-MPLS and SRv6 policies on the topology map

The Traffic Engineering (TE) topology map in Crosswork Network Controller is a powerful visualization tool designed to empower network operators with deep insights into their network's structure and active TE policies. By providing a clear and intuitive graphical representation of Segment Routing Traffic Engineering (SR-TE) policies and RSVP-TE tunnels, it significantly reduces the complexity of provisioning and managing these traffic engineering mechanisms. This visualization enhances operators' understanding of traffic flows, resource utilization, and overall network performance, enabling more effective traffic management, optimization, and troubleshooting.

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

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

Figure 6: Traffic engineering UI: SR-MPLS and SRv6 policies

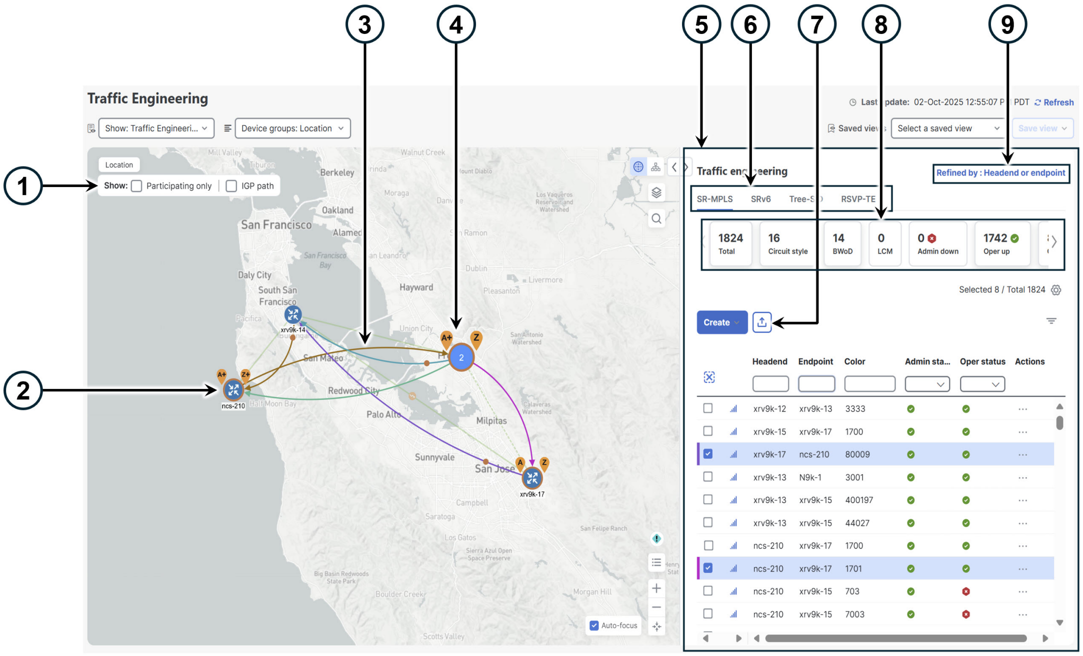

Traffic Engineering

Show: Traffic Engineeri... | Device groups: Location

Location

Show: □ Participating only | □ IGP path

Last update: 02-Oct-2025 12:55:07 PM PDT Refresh

Saved views: Select a saved view | Save view

Traffic engineering

Refined by : Headend or endpoint

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

1824 Total | 16 Circuit style | 14 BWoD | 0 LCM | 0 Admin down | 1742 Oper up

Selected 8 / Total 1824

Create

	Headend	Endpoint	Color	Admin sta...	Oper status	Actions
□	xrv9k-12	xrv9k-13	3333	✓	✓	...
□	xrv9k-15	xrv9k-17	1700	✓	✓	...
☑	xrv9k-17	ncs-210	80009	✓	✓	...
□	xrv9k-13	N9k-1	3001	✓	✓	...
□	xrv9k-13	xrv9k-15	400197	✓	✓	...
□	xrv9k-13	xrv9k-15	44027	✓	✓	...
□	ncs-210	xrv9k-17	1700	✓	✓	...
☑	ncs-210	xrv9k-17	1701	✓	✓	...
□	ncs-210	xrv9k-15	703	✓	✖	...
□	ncs-210	xrv9k-15	7003	✓	✖	...

The table describes the callouts for the Traffic Engineering topology map for SR-MPLS and SRv6 policies.

Callout no.	Description
1	Select the appropriate check box to enable these options: • **Show IGP path:** Displays the IGP path for the selected SR-TE policy. • **Show Participating only:** Displays only links that belong to selected SR-TE policy. All other links and devices are hidden.
2	**Device outlines:** A device with an orange ( ) outline indicates there is a node SID associated with that device or a device in the cluster.

Callout no.	Description
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. • **A+** denotes multiple SR-TE policies originating from a node. • **Z+** denotes multiple SR-TE policies terminating at a node.
5	**Window content:** The window content depends on the selected or filtered items. In this example, the SR-MPLS tab shows SR Policy table. Depending on the map selection, you can create, modify, or view policies. • Create explicit SR-MPLS policies, on page 38 • Create dynamic SR-MPLS policies based on optimization intent, on page 39 • Modify SR-MPLS policies, on page 42
6	**Tabs:** Click the **SR-MPLS** or **SRv6** tab to view the corresponding list of SR-TE policies.
7	**Export function:** Exports all data into a CSV file. You cannot export selected or filtered data.
8	**Mini dashlets:** Summarizes the operational SR-MPLS or SRv6 policy status and displays the number of PCC and PCE-initiated tunnels listed in the SR Policy table. When you select a dashlet, filters are applied and the policy table updates to display data corresponding to the filtered dashlet.
9	**Group filter:** Controls how group filters apply to table data. For example, if **Headend only** is selected, the table only displays policies where the policy's headend device is in the selected group. This filter helps you efficiently manage policies in large networks. Filter options: • **Headend or Endpoint:** Show policies with either the headend or endpoint device in the selected group. • **Headend and Endpoint:** Show policies if both the headend and endpoint are in the group. • **Headend only:** Show policies if the headend device of the policy is in the selected group. • **Endpoint only:** Show policies if the endpoint device of the policy is in the selected group.

View SR-MPLS and SRv6 policy details

You can view SR-MPLS or SRv6 TE policy level, single segment lists, multiple segment lists, and any path computation constraints configured on a per-candidate path basis.

Procedure

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

Figure 7: View SR policy details

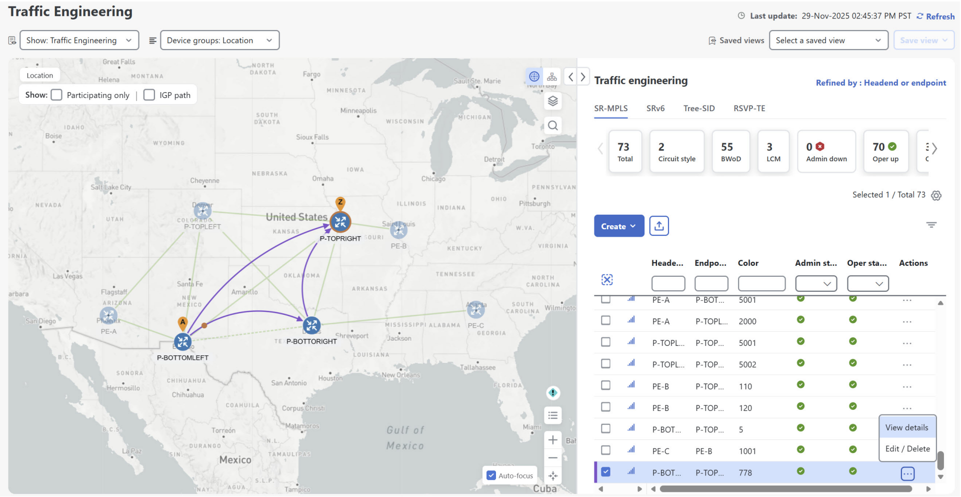

Traffic Engineering

Last update: 29-Nov-2025 02:45:37 PM PST Refresh

Show: Traffic Engineering | Device groups: Location | Saved views: Select a saved view | Save view

Location — Show: □ Participating only | □ IGP path

☑ Auto-focus

Traffic engineering — Refined by : Headend or endpoint

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

73 Total | 2 Circuit style | 55 BWoD | 3 LCM | 0 Admin down | 70 Oper up

Selected 1 / Total 73

Create

		Heade...	Endpo...	Color	Admin st...	Oper sta...	Actions
□		PE-A	P-BOT...	5001	✓	✓	⋯
□		PE-A	P-TOPL...	2000	✓	✓	⋯
□		P-TOPL...	P-TOP...	5001	✓	✓	⋯
□		P-TOPL...	P-TOP...	5002	✓	✓	⋯
□		PE-B	P-TOP...	110	✓	✓	⋯
□		PE-B	P-TOP...	120	✓	✓	⋯
□		P-BOT...	P-TOP...	5	✓	✓	View details
□		PE-C	PE-B	1001	✓	✓	Edit / Delete
☑		P-BOT...	P-TOP...	778	✓	✓	⋯

Step 2 A list of candidate paths appear along with policy details in the **SR policy details** page.

If the delay value is displayed, it is calculated for all policies every 10 minutes. Click the ⓘ icon next to the delay value to see the last update time.

Figure 8: SR policy details—headend, endpoint, summary, and candidate path

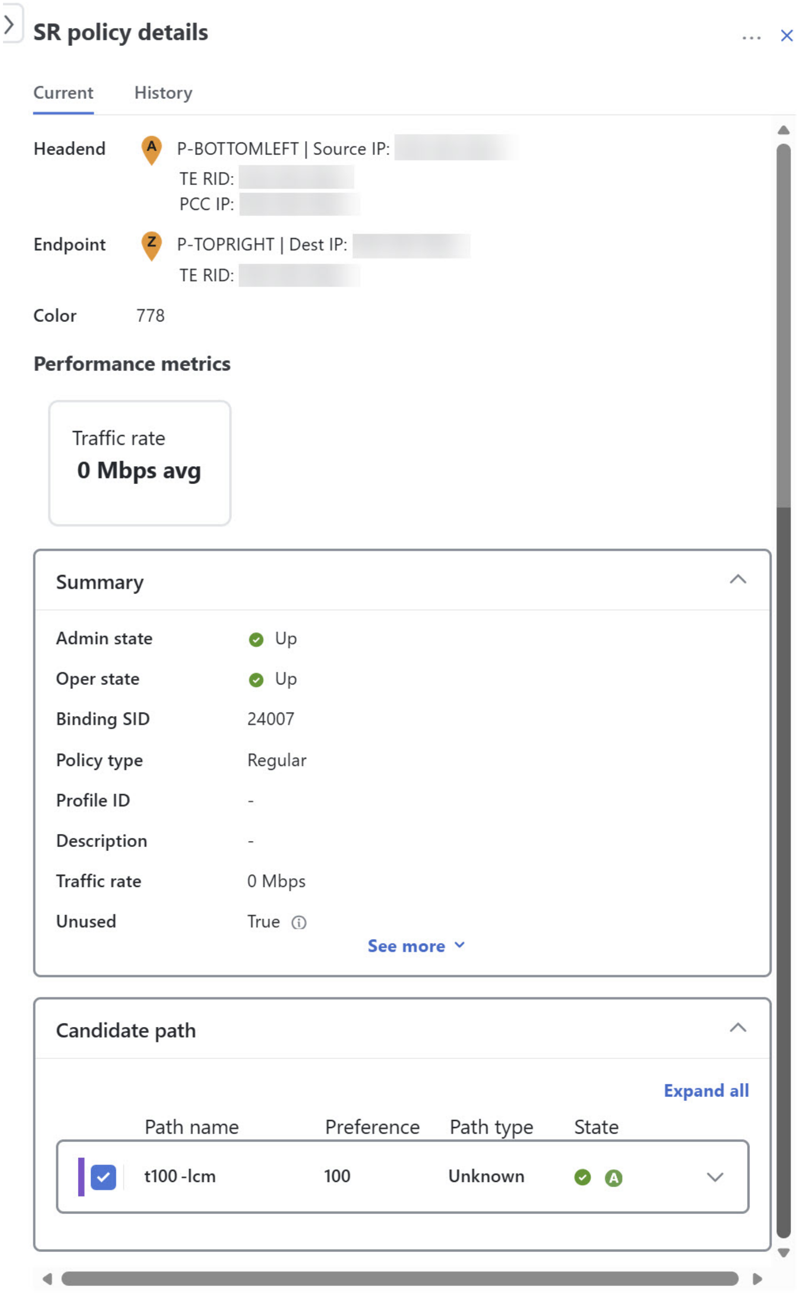

SR policy details

Headend: A P-BOTTOMLEFT | Source IP: [redacted]
TE RID: [redacted]
PCC IP: [redacted]

Endpoint: Z P-TOPRIGHT | Dest IP: [redacted]
TE RID: [redacted]

Color: 778

Performance metrics

Traffic rate
0 Mbps avg

Summary

Admin state	✓ Up
Oper state	✓ Up
Binding SID	24007
Policy type	Regular
Profile ID	-
Description	-
Traffic rate	0 Mbps
Unused	True ⓘ

Candidate path

	Path name	Preference	Path type	State
☑	t100 -lcm	100	Unknown	✓ A

Step 3 In the candidate path area, click **Expand all** to view additional details on the different paths and segments.

For headends that support MSL policies, you may see a policy with a single candidate path and multiple segment lists. You can also view the weight associated with each segment list, along with other details such as segment type, node names, labels, algo, and SID type.

Note
Local Congestion Mitigation (LCM) automatically deploys multiple segment lists policies for devices that are gRPC MSL compliant based on specified thresholds. This feature is available when LCM is in Automated mode.

For detailed instructions on how to enable gRPC-MSL support, refer to the Prepare devices for gRPC policy management section in the *Cisco Crosswork Network Controller 7.2 Network Bandwidth Management guide*.

Figure 9: SR policy details with multiple segment list

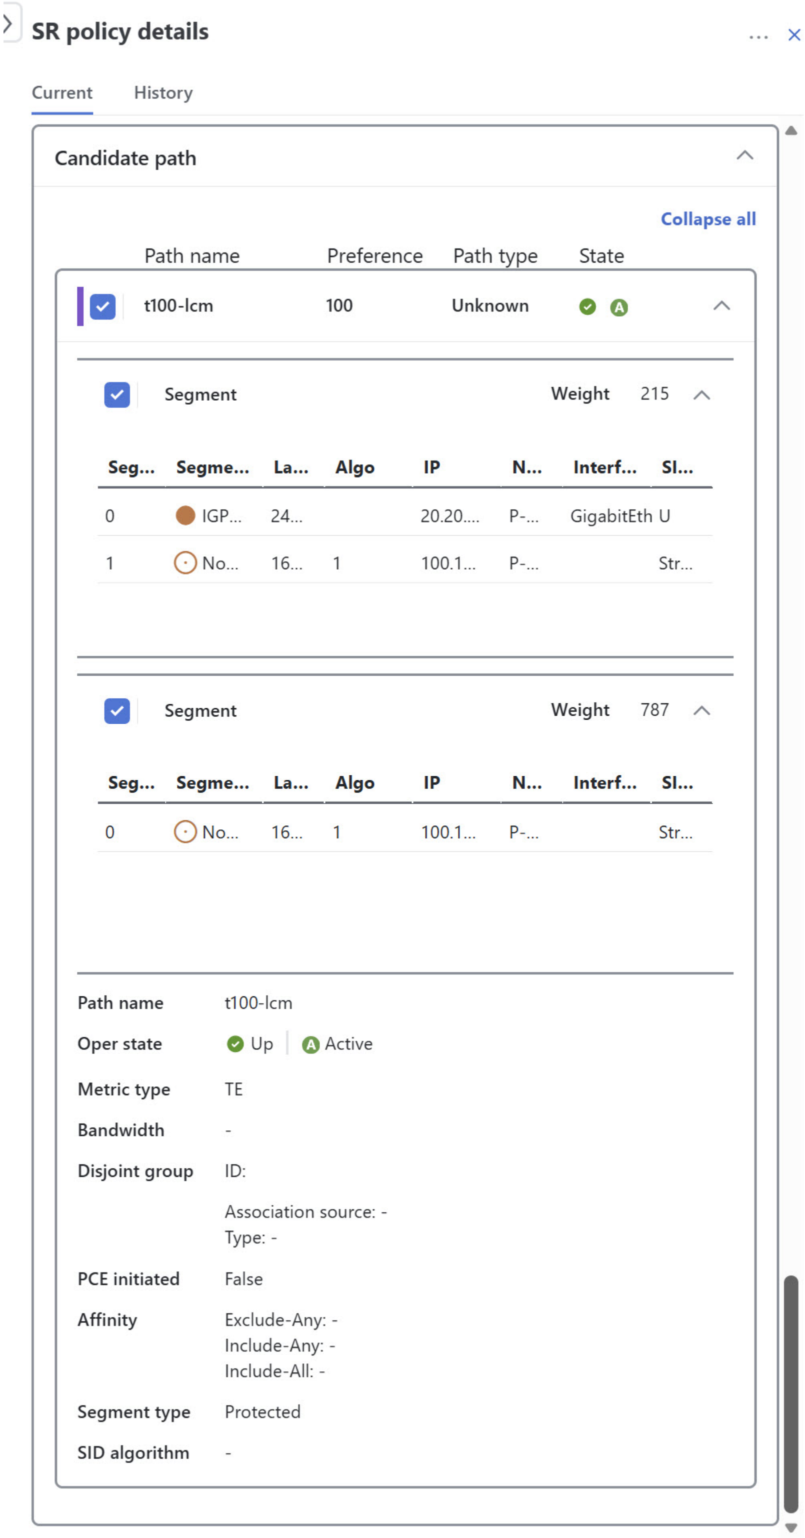

SR policy details

Current History

Candidate path

Collapse all

	Path name	Preference	Path type	State
☑	t100-lcm	100	Unknown	✓ A

☑ Segment Weight 215

Seg...	Segme...	La...	Algo	IP	N...	Interf...	Sl...
0	● IGP...	24...		20.20....	P-...	GigabitEth	U
1	⊙ No...	16...	1	100.1...	P-...		Str...

☑ Segment Weight 787

Seg...	Segme...	La...	Algo	IP	N...	Interf...	Sl...
0	⊙ No...	16...	1	100.1...	P-...		Str...

Path name	t100-lcm
Oper state	✓ Up | A Active
Metric type	TE
Bandwidth	-
Disjoint group	ID: Association source: - Type: -
PCE initiated	False
Affinity	Exclude-Any: - Include-Any: - Include-All: -
Segment type	Protected
SID algorithm	-

View IGP path and metrics

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

Procedure

Step 1 In the **SR policy** table, select the SR-TE (SR-MPLS and SRv6) policies that interest you.

Step 2 Select the **Show IGP path** check box to view metrics. The IGP paths for the selected SR-MPLS policies appear as straight lines instead of segment hops.

Step 3 In a dual-stack topology, select the **Participating only** checkbox to view metrics on participating links.

Step 4 Click > **Metrics** tab.

Step 5 Toggle applicable metrics to **ON**.

Figure 10: View physical path and metrics

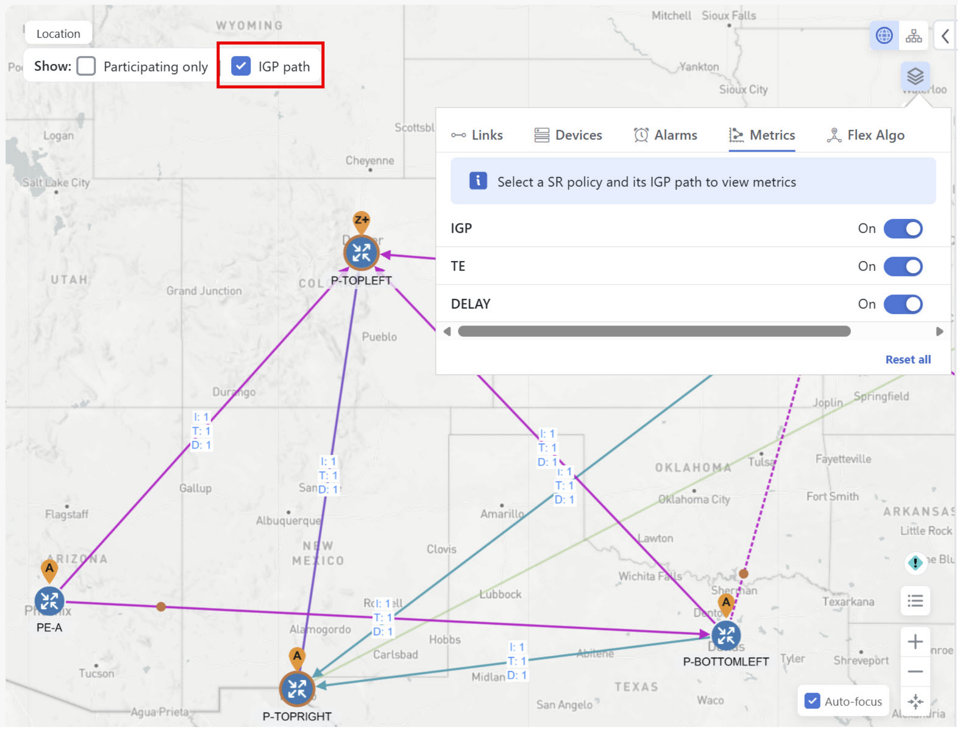

Show: □ Participating only ☑ IGP path

Select a SR policy and its IGP path to view metrics

IGP On

TE On

DELAY On

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 want to switch the active path, you must manually configure the device to activate the desired candidate path.

Important notes

- Only PCC-initialized SR-TE policies with MCPs are supported.
- Crosswork Network Controller does not distinguish dynamic paths from explicit paths. The Policy Type field value displays as 'Unknown'.
- Active explicit paths are visible, but inactive candidate explicit paths are not shown in the UI.

Before you begin

A policy must be pre-configured with Multiple Candidate Paths (MCPs) on devices before they can be visualized on the Crosswork Network Controller Traffic Engineering topology map.

Procedure

Step 1 From the main menu, choose or **SRv6** tab.

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

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

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

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

Figure 11: SR-TE policy on the topology map

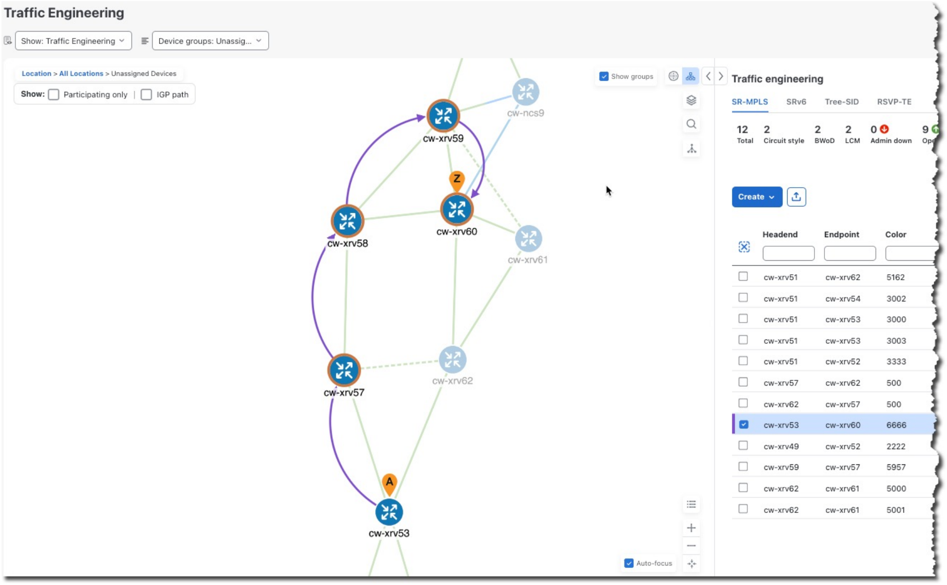

	Headend	Endpoint	Color
□	cw-xrv51	cw-xrv62	5162
□	cw-xrv51	cw-xrv54	3002
□	cw-xrv51	cw-xrv53	3000
□	cw-xrv51	cw-xrv53	3003
□	cw-xrv51	cw-xrv52	3333
□	cw-xrv57	cw-xrv62	500
□	cw-xrv62	cw-xrv57	500
☑	cw-xrv53	cw-xrv60	6666
□	cw-xrv49	cw-xrv52	2222
□	cw-xrv59	cw-xrv57	5957
□	cw-xrv62	cw-xrv61	5000
□	cw-xrv62	cw-xrv61	5001

Step 3 View the list of candidate paths.

a) In the **Actions** column of the SR-MPLS or SRv6 Policy table, click ··· > **View details**.

A list of candidate paths appear along with policy details in the **SR policy details** window. The green A under the State column indicates the active path.

Figure 12: Candidate path in SR policy details

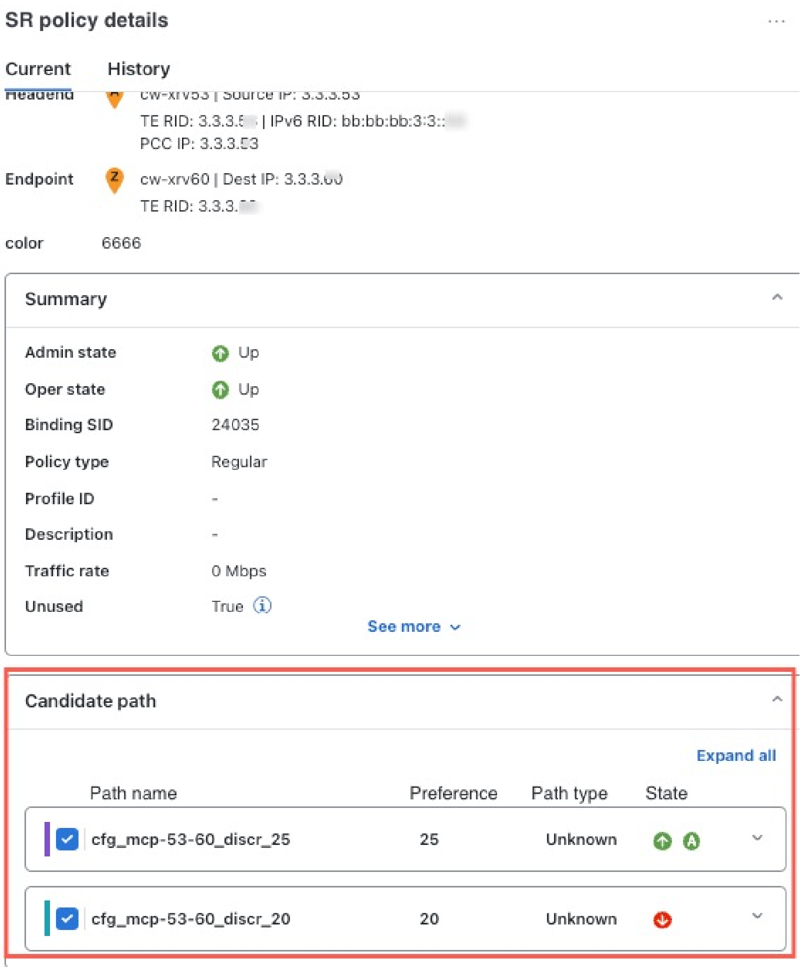

SR policy details

Current History

Headend cw-xrv53 | Source IP: 3.3.3.53
TE RID: 3.3.3.5[redacted] | IPv6 RID: bb:bb:bb:3:3::[redacted]
PCC IP: 3.3.3.53

Endpoint cw-xrv60 | Dest IP: 3.3.3.[redacted]
TE RID: 3.3.3.[redacted]

color 6666

Summary

Admin state	Up
Oper state	Up
Binding SID	24035
Policy type	Regular
Profile ID	-
Description	-
Traffic rate	0 Mbps
Unused	True

See more

Candidate path

Expand all

	Path name	Preference	Path type	State
☑	cfg_mcp-53-60_discr_25	25	Unknown	↑ A
☑	cfg_mcp-53-60_discr_20	20	Unknown	↓

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) Select the check box next to any candidate path.

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

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

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

Figure 13: Candidate path on the topology map

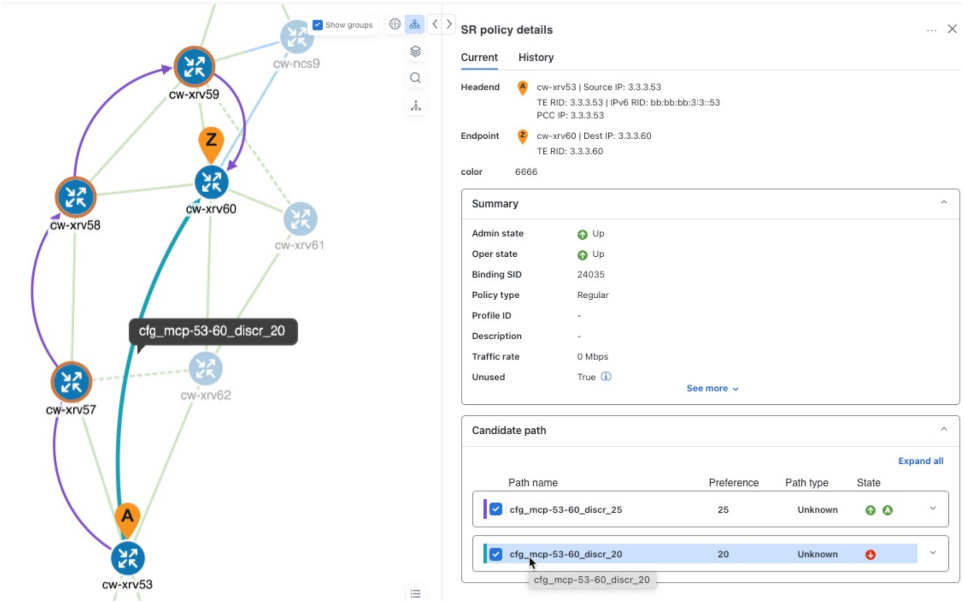

SR policy details

Current History

Headend: cw-xrv53 | Source IP: 3.3.3.53
TE RID: 3.3.3.53 | IPv6 RID: bb:bb:bb:3:3::53
PCC IP: 3.3.3.53

Endpoint: cw-xrv60 | Dest IP: 3.3.3.60
TE RID: 3.3.3.60

color 6666

Summary

Admin state	Up
Oper state	Up
Binding SID	24035
Policy type	Regular
Profile ID	-
Description	-
Traffic rate	0 Mbps
Unused	True

See more

Candidate path

Expand all

	Path name	Preference	Path type	State
☑	cfg_mcp-53-60_discr_25	25	Unknown	↑ A
☑	cfg_mcp-53-60_discr_20	20	Unknown	↓

cfg_mcp-53-60_discr_20

View underlying paths associated with a Binding-Segment ID (B-SID) label

Crosswork Network Controller provides a powerful tool for network visualization, enabling users to explore the underlying paths of B-SID hops configured on devices. Whether the paths are set manually or configured through the Crosswork Network Controller, this feature allows for detailed inspection of SR-MPLS and SRv6 policy paths. By assigning a B-SID label, such as **15700** in this example, users can easily identify and trace the path of a policy hop.

To view the B-SID underlying path for an SR-MPLS or SRv6 policy, complete these steps:

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering**.

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

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

Figure 14: B-SID label

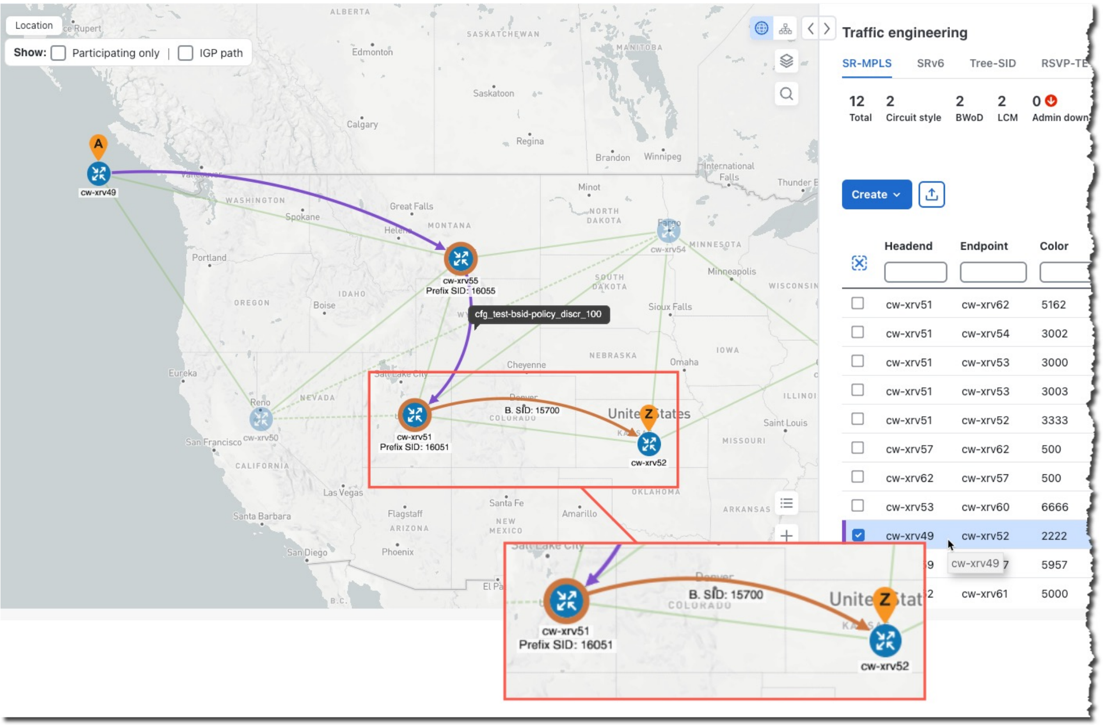

Traffic engineering

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

12 Total | 2 Circuit style | 2 BWoD | 2 LCM | 0 Admin down

Create ⌄

	Headend	Endpoint	Color
□	cw-xrv51	cw-xrv62	5162
□	cw-xrv51	cw-xrv54	3002
□	cw-xrv51	cw-xrv53	3000
□	cw-xrv51	cw-xrv53	3003
□	cw-xrv51	cw-xrv52	3333
□	cw-xrv57	cw-xrv62	500
□	cw-xrv62	cw-xrv57	500
□	cw-xrv53	cw-xrv60	6666
☑	cw-xrv49	cw-xrv52	2222
	9	cw-xrv49 7	5957
	2	cw-xrv61	5000

Step 3 In the **SR policy details** page, click ⋯ > **View details**.

Figure 15: View details

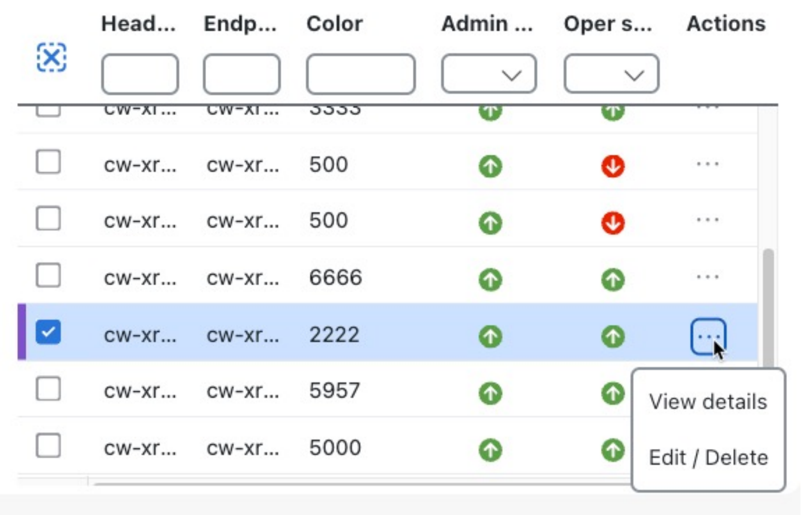

	Head...	Endp...	Color	Admin ...	Oper s...	Actions
□	cw-xr...	cw-xr...	3333	⬆	⬆	⋯
□	cw-xr...	cw-xr...	500	⬆	⬇	⋯
□	cw-xr...	cw-xr...	500	⬆	⬇	⋯
□	cw-xr...	cw-xr...	6666	⬆	⬆	⋯
☑	cw-xr...	cw-xr...	2222	⬆	⬆	⋯
□	cw-xr...	cw-xr...	5957	⬆	⬆	
□	cw-xr...	cw-xr...	5000	⬆	⬆	

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

Figure 16: B-Sid label ID

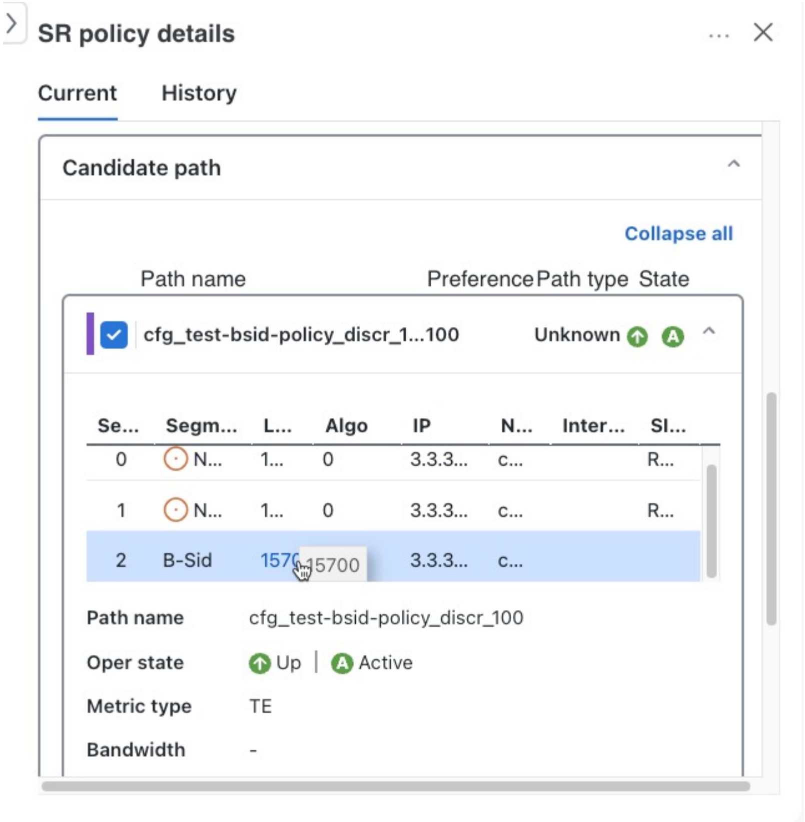

Path name	Preference	Path type	State
cfg_test-bsid-policy_discr_1...100		Unknown	● A

Se...	Segm...	L...	Algo	IP	N...	Inter...	Sl...
0	N...	1...	0	3.3.3...	c...		R...
1	N...	1...	0	3.3.3...	c...		R...
2	B-Sid	157(15700		3.3.3...	c...		

Path name cfg_test-bsid-policy_discr_100

Oper state Up | Active

Metric type TE

Bandwidth -

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

Figure 17: B-SID path

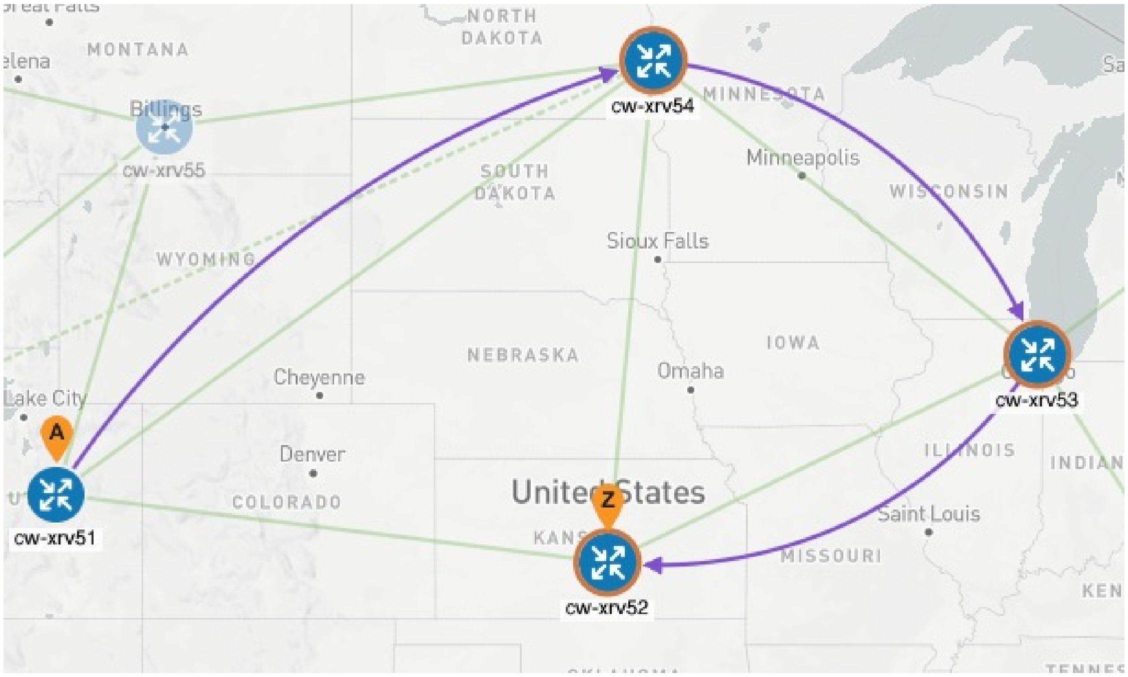

SR policies with multiple segment lists

Multiple segment list (MSL) is a feature of SR policy that

- allows a single SR policy to contain multiple distinct segment lists, each with configurable weights,
- enables granular and efficient distribution of network traffic over several alternative paths, and
- improves flexibility and control for traffic engineering strategies.

Operational benefits

Operational benefits of using multiple segment lists include:

- **Greater precision**: Assign weights to individual segment lists for fine-grained traffic distribution beyond fixed equal splits.
- **Simplified management**: Reduce operational complexity by eliminating the need to deploy multiple parallel policies.
- **Adaptive traffic control**: Easily adjust traffic distribution by updating segment list weights, without the need to add or remove policies.
- **Improved network stability**: Lower the risk of disrupting existing ECMP flows and enables smoother adaptation to traffic shifts.

Requirements and limitations

- MSL policies are automatically generated and deployed by Local Congestion Mitigation (LCM) only for GRPC_MSL compliant devices in automated mode.
- You cannot create or edit MSL policies directly using the Crosswork Network Controller UI or API. Once deployed by LCM, you can only view these policies and their corresponding segments.
- Supported for XR devices with the required gRPC and BGP-LS configurations.
- Device must be tagged "grpc_msl" or "GRPC_MSL" for deployment.
- In Automated mode, PCE-initiated policies are not supported. In manual mode, if the tag is missing, LCM deploys a PCE-initiated policy.

For detailed instructions on how to enable gRPC-MSL support, refer to the Prepare devices for gRPC policy management section in the *Cisco Crosswork Network Controller 7.2 Network Bandwidth Management guide*.

Traditional SR policies vs. SR policies with multiple segment list

Traditionally, to divert traffic and mitigate congestion, multiple parallel SR policies (up to a maximum of eight) were deployed. Traffic was split equally among these policies, which meant the smallest possible diversion was determined by the number of policies in use. This approach limited the granularity of traffic distribution and required deploying a large number of policies to achieve finer control over traffic splitting.

With MSL, a single SR policy can define several segment lists mapped to unique paths, each with an assigned weight. The weight determines how much traffic is allocated to each path, enabling precise traffic splitting and optimal utilization of network resources.

For example, consider a scenario where interface A→B has a utilization of 50%, a congestion threshold of 40%, and all network links are 100 Mbps with the same IGP metric.

- **Traditional approach**: Five parallel SR policies (each carrying 10 Mbps) are created to split the traffic: four policies would keep 40 Mbps on the shortest path, and one policy would reroute 10 Mbps via a detour path.
- **MSL approach**: LCM deploys only a single policy with two weighted segment lists: one for the shortest path (A→B) and another for the detour path. The segment weights are set so that 40 Mbps remains on the main path and 10 Mbps is diverted to the detour. If traffic patterns shift over time, the solution can be easily adjusted by modifying the segment list weights, without creating or removing additional policies.

The screenshot illustrates how multiple segment lists appear within a single SR policy. You can see the individual segment lists, their associated weights, and other key details, providing clear visibility into how traffic is distributed across different paths.

Figure 18: Single SR policy with multiple segment lists

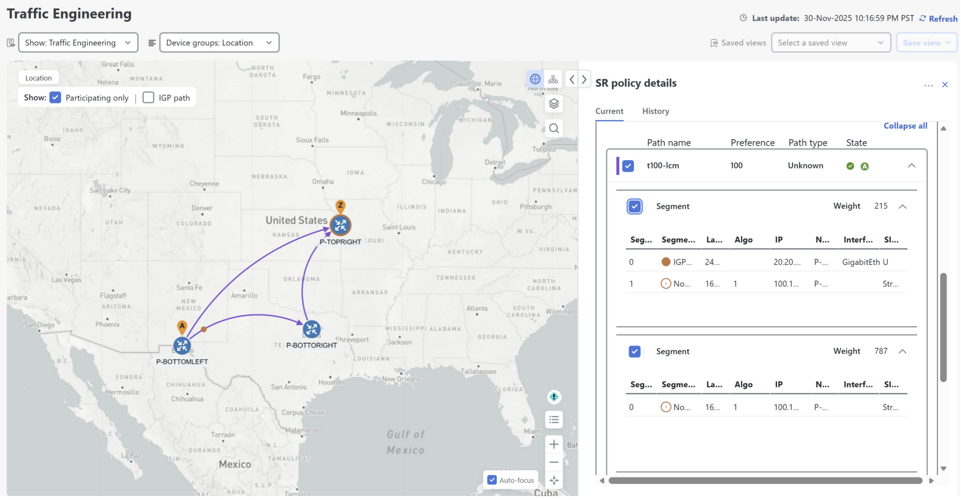

Path name	Preference	Path type	State
☑ t100-lcm	100	Unknown	✓ A

☑ Segment — Weight 215

Seg...	Segme...	La...	Algo	IP	N...	Interf...	Sl...
0	● IGP...	24...		20.20....	P-...	GigabitEth	U
1	⊙ No...	16...	1	100.1...	P-...		Str...

☑ Segment — Weight 787

Seg...	Segme...	La...	Algo	IP	N...	Interf...	Sl...
0	⊙ No...	16...	1	100.1...	P-...		Str...

Native SR paths

A native SR path is a segment routing path that

- uses the network's native IGP (such as OSPF or IS-IS) with segment routing extensions to forward traffic,
- uses all available Equal Cost Multi-Path (ECMP) routes between a source and destination, and
- supports operations, administration, and maintenance (OAM) activities to monitor segment-routed label-switched paths (LSPs) and isolate forwarding problems for troubleshooting.

Device prerequisites to view native SR paths

To enable successful visualization of native SR paths, ensure devices satisfy these requirements:

Run the required software version

- Ensure your device operates on Cisco IOS XR 7.3.2 or higher.
- To verify the version, run the `show version` command.

Enable and confirm gRPC configuration

- Enable gRPC on your device and verify the configuration. For more information about enabling gRPC on PCE, see Requirements for adding SR-PCE providers in the Cisco Crosswork Network Controller 7.2 Administration guide.
- Run the `show run grpc` command to confirm gRPC configuration.

  Example gRPC configuration:

```
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, upload valid security certificates to the device.

Enable and verify GNMI capability

- Confirm that your device supports and has GNMI enabled.
- To verify, navigate to **Device Management > Network Devices**, click the device IP address, and check that GNMI is listed under **Connectivity details**.
- The encoding type depends on the device's capabilities, the supported data model, and the expected transmission method between the device and Crosswork Network Controller. Devices that support GNMI may offer these encoding types:
  - **JSON**: Human-readable and widely supported by most devices.
  - **BYTES**: Encodes data in binary format for efficient transmission.
  - **PROTO**: A compact, efficient binary format used with gRPC.
  - **ASCII**: A plain-text format that is human-readable but less commonly used compared to JSON.
  - **JSON IETF**: A standardized variant of JSON that adheres to IETF YANG specifications.

Configure the static route to the CDG router

- Add a static route from your device to the southbound CDG IP address.

  Example configuration:

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

Visualize native paths

This task guides you to create and run a path query, view past and ongoing queries, and see paths in the topology map.

Follow these steps to create a path query.

Procedure

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

Step 2 Click **New query**.

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

Note

Path queries may take a moment to complete. When the Running Query ID pop-up appears, select **View past queries** to return to the path query dashboard. If path queries exist in the list, you can view their details while your new query runs in the background. The blue running icon in the Query State column indicates a query in progress. A green query state means you can view the completed query.

Figure 19: New path query

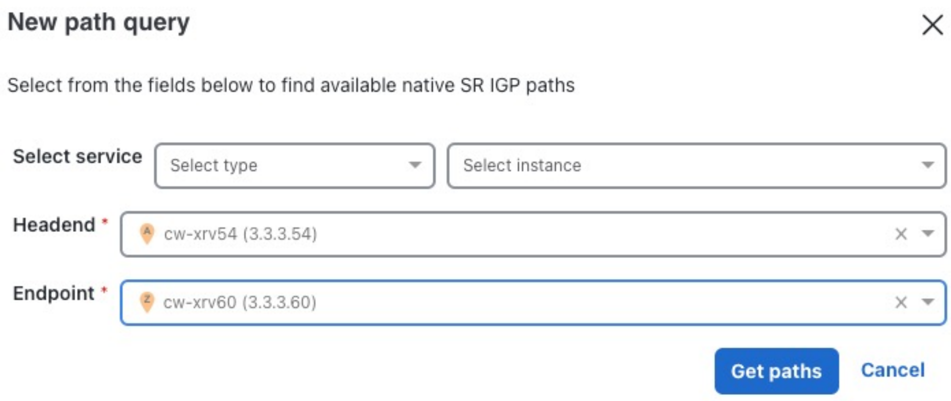

New path query

Select from the fields below to find available native SR IGP paths

Select service: Select type | Select instance

Headend *: cw-xrv54 (3.3.3.54)

Endpoint *: cw-xrv60 (3.3.3.60)

Step 4 Click **View results** when it becomes available on the Running Query ID pop-up. The Path details page appears, showing the available path details. The topology map displays Native SR IGP paths on the left.

Figure 20: Path details

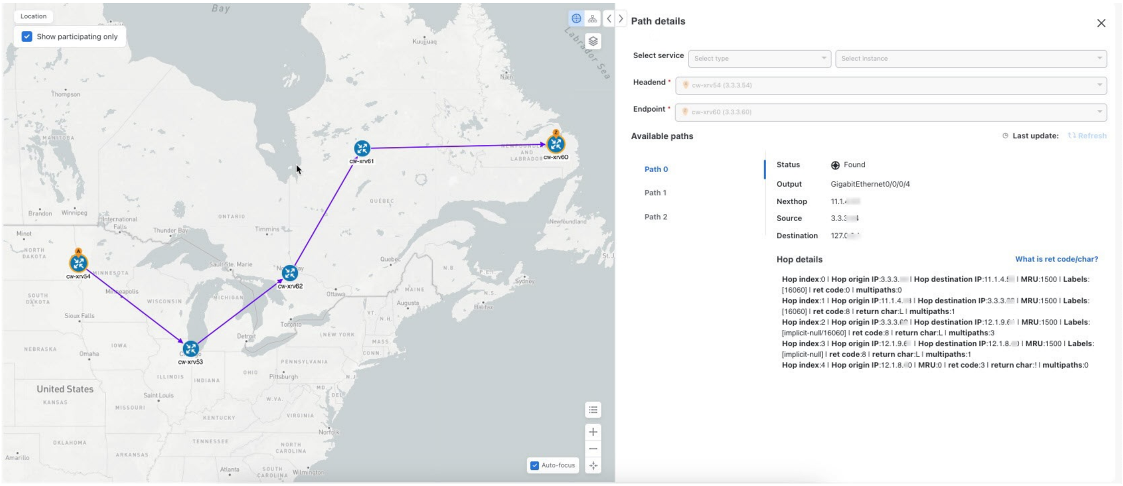

Path details

Select service | Select type | Select instance

Headend * cw-xrv54 (3.3.3.54)

Endpoint * cw-xrv60 (3.3.3.60)

Available paths | Last update: | Refresh

Path 0
Path 1
Path 2

Status: Found
Output: GigabitEthernet0/0/0/4
Nexthop: 11.1.[redacted]
Source: 3.3.3[redacted]4
Destination: 127.0[redacted]

Hop details | What is ret code/char?

Hop index:0 | **Hop origin IP**:3.3.3.[redacted] | **Hop destination IP**:11.1.4.[redacted] | **MRU**:1500 | **Labels**: [16060] | **ret code**:0 | **multipaths**:0
Hop index:1 | **Hop origin IP**:11.1.4.[redacted]3 | **Hop destination IP**:3.3.3.[redacted] | **MRU**:1500 | **Labels**: [16060] | **ret code**:8 | **return char**:L | **multipaths**:1
Hop index:2 | **Hop origin IP**:3.3.3.[redacted] | **Hop destination IP**:12.1.9.[redacted] | **MRU**:1500 | **Labels**: [implicit-null/16060] | **ret code**:8 | **return char**:L | **multipaths**:3
Hop index:3 | **Hop origin IP**:12.1.9.[redacted] | **Hop destination IP**:12.1.8.[redacted]0 | **MRU**:1500 | **Labels**: [implicit-null] | **ret code**:8 | **return char**:L | **multipaths**:1
Hop index:4 | **Hop origin IP**:12.1.8.[redacted]0 | **MRU**:0 | **ret code**:3 | **return char**:! | **multipaths**:0

Configure TE link affinities

Affinity mapping allows you to define link attributes (affinities) in Crosswork Network Controller with the same names and bit positions that are used on the device. This helps maintain consistency and improve visualization. The affinity is represented as a 32-bit value, where each bit (0–31) corresponds to a link attribute such as service profile colors (for example, low delay, high bandwidth, and so on).

Note

- If an affinity mapping is not defined in the UI, the affinity name will appear as "UNKNOWN".
- Before removing an affinity, remove any associated TE tunnels to avoid orphan tunnels. If an affinity is removed while still associated with a TE tunnel, it will show as "UNKNOWN" in the **SR policy / RSVP-TE tunnel details** window.

On devices, affinity maps are configured by setting bits for each affinity name. The following example shows an SR-TE 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
                !
                !
                !
```

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)

To map the affinity names to the bits, complete these steps:

Procedure

Step 1 Choose **Administration** > **Settings** > **System settings** > **Traffic engineering** > **Affinity** > **TE link affinities**. Alternatively, you can define affinities while provisioning an SR-TE policy, Tree-SID, or RSVP-TE tunnel by clicking **Manage mapping** under the **Constraints** > **Affinity** field.

Step 2 Click **+ Create** to add a new affinity mapping.

Step 3 Enter the name and the bit position corresponding to the device configuration. For example, using the configuration described earlier:

Figure 21: Mapping affinities

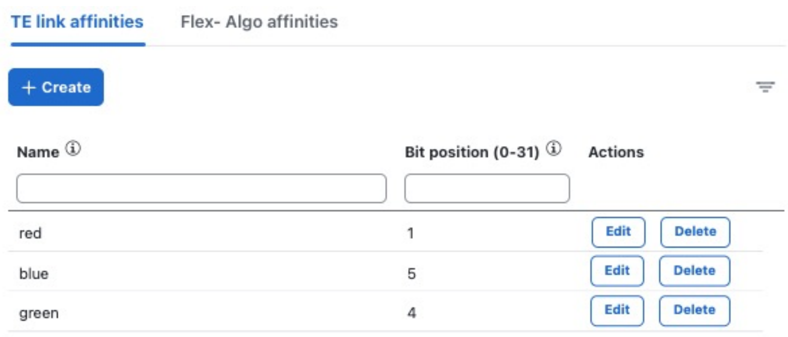

Name	Bit position (0-31)	Actions
red	1	Edit Delete
blue	5	Edit Delete
green	4	Edit Delete

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

Policy deployment considerations

Before provisioning policies, it's important to consider these options to ensure smooth deployment.

- In a scaled setup with high nodes, policies, or interfaces, you may encounter timeouts during policy deployment. To address this, configuring timeout options is recommended. See Configure TE timeout settings, on page 16.
- For enhanced visualization, you can collect affinity information from your devices and map them in Crosswork Network Controller prior to provisioning an SR policy, Tree-SID, or RSVP-TE tunnel. See Affinity map configurations, on page 9 for sample configurations and Configure and visualize flexible algorithm SRLG exclusion, on page 63 for details on SRLG exclusion.

Create explicit SR-MPLS policies

This task explains how to create SR-MPLS policies using an explicit (fixed) path. Each path consists of a list of prefix or adjacency Segment IDs (SID list), with each ID representing a node or link on the path.

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS**.

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

Note
If you would like to provision a PCC initiated policy using Network Services Orchestrator (NSO) via the Crosswork Network Controller UI, see Create SR-TE policies (PCC-initiated), on page 41.

Step 3 Under **Policy details**, enter or select the required SR-MPLS policy values. Hover over the ⓘ icon 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 to include in the SR-MPLS policy path.

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

Step 7 To activate the policy on the network, click **Provision**.

Step 8 Validate the SR-MPLS policy creation:

a. Confirm that the new SR-MPLS policy appears in the **Traffic engineering** table. Select the policy in the table to highlight it on the map.

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

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

Note
If the setup includes many nodes, policies, or interfaces, a timeout may occur during policy deployment. To configure timeout options, see Configure TE timeout settings, on page 16.

Create dynamic SR-MPLS policies based on optimization intent

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 metrics to minimize in-path computation: IGP, TE, or latency. SR-PCE automatically re-optimizes the path as necessary, responding to topology changes. When a link or interface fails, the network locates an alternate path that meets the criteria specified in the policy and raises an alarm. If no valid path is found, an alarm occurs, and the packets are dropped.

To create SR-MPLS policies with a dynamic path, complete these steps:

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS**.

Step 2 Click **Create > PCE Init** . To provision a PCC-initiated policy using NSO via the Crosswork Network Controller UI, see Create SR-TE policies (PCC-initiated), on page 41.

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

Note
If you have set up device groups, you can select the device group from the **Device groups** drop-down list. Navigate and zoom in on the topology map to choose 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 or specify any required disjointness.

Affinity considerations

- Affinity constraints and disjointness cannot be configured on the same SR-MPLS policy. No more than two SR-MPLS policies can exist in any disjoint group or subgroup. The configuration will not be allowed during Preview.
- If there are existing SR-MPLS policies in a disjoint group, all policies from that group are displayed during Preview.
- If SRLG exclusion constraints are defined for a Flexible Algorithm, any links belonging to the specified excluded SRLGs will be automatically filtered out from the IGP routing calculations, optimizing alternate paths by considering only the available links. For more information, see Configure and visualize flexible algorithm SRLG exclusion.

Step 7 Under **Segments** , select whether to use protected segments when available. Enter any applicable SID constraint. Crosswork Network Controller will try to find a path with this SID. If it cannot find a path with the SID constraint, the provisioned policy remains operationally down until the conditions are met.

SID information

- Flexible Algorithm—The values correspond to the Flexible Algorithm defined on the device. Cisco IOS XR enforces the 128-255 range.
- Algorithm 0—This is a Shortest Path First (SPF) algorithm based on link metric. The Interior Gateway Protocol (IGP) computes this shortest path.
- Algorithm 1—This is a Strict Shortest Path First (SSPF) algorithm based on link metric. It is identical to algorithm 0, but requires all nodes along the path to 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 8 Click **Preview** to highlight the path on the map. Click **Provision** to activate the policy on the network.

Step 9 Validate the SR-MPLS policy creation.

a. Confirm that the new SR-MPLS policy appears in the **Traffic engineering** table. Select the policy to highlight it on the map.

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

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

Create SR-TE policies (PCC-initiated)

Before you begin

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 list of prefix or adjacency Segment IDs, each representing a node or link along on the path.

This task creates explicit or dynamic SR-MPLS or SRv6 policies using NSO via the Crosswork Network Controller UI.

Procedure

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

Step 2 Click +. This 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.

Table 2: SR-TE policy configuration

For	Complete these steps
name	Enter a name for this SR-TE policy.
head-end	Click + to select a node or manually enter the node name.
tail-end	Enter the IP address of the tail-end router.
color	Enter a the SR policy color to identify the traffic. For example: 200.
path	a. Click + and enter a preference value. For example: 123 b. For explicit-path, click + to add previously configured SID lists. c. For dynamic-path, select the metric you want to minimize and define any applicable constraints (for example, affinity, SRLG exclusion) and disjointness.
srv6	If you are creating an SRv6 policy, enable srv6 and enter the locator details.

Step 4 Click **Dry Run** to validate your changes and save them. Crosswork Network Controller will display your changes in a pop-up window.

If your requirements differ from those described 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, complete these steps:

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** tab.

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

Step 3 Choose **View details** or **Edit/Delete**. After updating the SR-MPLS policy details, you can preview the changes on the map before saving it.

Create an ODN template

SR-TE ODN integrates segment routing with on-demand traffic engineering, allowing the network to dynamically create traffic-engineered paths as needed. You can configure an ODN template for each color that represents a specific SLA or traffic requirement. This setup enables a service head-end router to automatically generate an SR-TE policy for a BGP next-hop whenever necessary. The head-end is configured with an ODN template linked to a specific color to optimize the traffic path when a prefix with that color is detected.

Complete the steps to create an ODN template.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Provisioning (NSO)** > **SR-TE** > **ODN-Template**.

Step 2 Click + and enter the unique name for the ODN template. Click **Continue**.

Step 3 To apply a pre-configured custom template, click + and enter its name. Crosswork Network Controller displays the variables that can be substituted or parameterized as defined in the template. Under Iteration, specify the number of times to apply the custom-template.

Step 4 In the **head-end** area, specify the list of source routers:

a) In the **name** field, enter the source device or the router where the tunnel begins. Click **Continue**.

b) To apply a pre-configured custom template, click + and enter its name. Crosswork Network Controller displays the variables that can be substituted or parameterized as defined in the template. Under Iteration, specify the number of times to apply the custom-template.

Step 5 Specify these policy options:

a) In the **maximum-sid-depth** field, enter the maximum SID depth supported by the router.

b) In the **pce-group** field, enter the PCE group to assign to the template.

c) For **color**, specify an SR policy color to identify the traffic.

d) In the **bandwidth** field enter the requested bandwidth value in kbps.

e) Enter the **source-address** of the policy.

Step 6 In the SRv6 area:

a) in the **locator-name** field, enter the required SRv6 node. The locator name should match what is configured on the router.

b) from the **behavior** list, choose how IPv6 packets should be treated or processed by the network.

c) from the **binding-sid-type** list, choose the type of binding segment ID.

Step 7 In the **performance-measurement** area, create delay and liveness profiles.

- **Delay profile** - Allows scheduling probes and configuring metric advertisement parameters. You can configure different profiles for different types of delay measurements. To enable performance measurement, you require a catalog of profiles.
- **Liveness profile** - Allows network to confirm that a specific path, segment, or node is operational and capable of forwarding packets. These checks maintain network availability and reliability.

If you choose	Complete these steps
delay	**a.** In the **profile** field, enter the delay profile name. **b.** In the **logging** area, toggle **Enable logging** to enable system logging for delay measurement. **c.** Check **delay-exceeded** to log messages in syslog when the delay exceeds the threshold.
liveness	**a.** In the **profile** field, enter the liveness profile name. **b.** From the **invalidation-action** list, select the action to be taken when Performance Management liveness session is invalidated. Selecting **none** results in no action being taken. If logging is enabled, the failure is logged, but the SR Policy operational state remains unchanged. **down** (default) means the candidate path is immediately deactivated. **c.** In the **logging** area, toggle **Enable logging** to enable system logging for liveness detection. **d.** Check **session-state-change** to log messages in syslog when the state of the session changes.

To return packets to head-end, in the **reverse-path-label** field, enter the MPLS label to be used for the reverse path.

Step 8 In the **dynamic** area, define settings for dynamic path computation.

a) Select **pce** to delegate dynamic path computation to PCE.

b) In the **flex-alg** field, enter the SID algorithm constraint. This setting allows operators to customize IGP shortest path computation based on their needs. When a constraint, such as SRLG exclusion is applied, any links belonging to the specified excluded SRLGs will be automatically filtered out from the IGP routing calculations. Alternate paths are optimized by considering only the available links, thereby enhancing network resilience and service availability. For more information, see Configure and visualize flexible algorithm SRLG exclusion, on page 63.

c) In the **metric-type** list, choose the metric for use in path computation.

d) In the **metric-margin** area, specify the absolute value or relative percent to configures the on-demand dynamic path metric margin.

e) The **affinity** area specifies a relationship between policy path and link colors. SR-TE finds a path that includes or excludes links that have specific colors or combinations of colors. To compute the path with link color constraints, create a **rule**with the required **action** and **color**.

f) The **segments** area specifies an SID constraint to find a path with that SID. Enter an SID algorithm number to configure path segment constraints.

g) Select **disjoint-path** to compute a path that is disjoint from another path in the same disjoint-group. The disjoint paths can originate from the same head-end or different head-ends.

1. From the **type** list, select the type of disjoint path.
2. In the **group-id** field, enter the group id of the disjoint group.
3. In the **sub-id** field, enter the subgroup ID of the disjoint group.
4. In the **source** field, enter the association source. This is applicable only on XE devices and required when setting association group.

Step 9 Click **Dry run** to validate and save your changes. When you are ready to activate the policy, click **Commit changes**.

CHAPTER 3

Resource Reservation Protocol (RSVP)

Resource Reservation Protocol

Resource Reservation Protocol is a signaling protocol that:

- enables systems and local clients to request resource reservations from the network,
- establishes explicit paths for data traffic, and
- ensures that necessary bandwidth and network resources are allocated for critical applications to meet desired Quality of Service (QoS) standards.

Resource Reservation Protocol-Traffic Engineering (RSVP-TE) processes protocol messages from other systems, handles resource requests from local clients, and generates protocol messages. It manages the creation, maintenance, and deletion of resource reservations for data flows.

This section describes the RSVP-TE features that Crosswork Network Controller supports.

View RSVP-TE tunnels on the topology map

A Traffic Engineering topology map is a visualization tool that displays the network topology with clear emphasis on participating links and devices. The map distinguishes path types by representing Record Route Objects (RRO) as straight lines and Explicit Route Objects (ERO) as curved lines, while also marking adjacency segment IDs for precise identification. Device clusters are labeled to indicate tunnel sources and destinations, enhancing clarity. Additionally, the map features a mini dashboard that summarizes tunnel statuses and counts, providing a concise overview of network traffic engineering elements.

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

From the Traffic engineering table, select the RSVP-TE tunnels you want to view on the map. RSVP-TE tunnels appear as colored lines indicating their source and destination.

Figure 22: Traffic Engineering UI - RSVP-TE tunnels

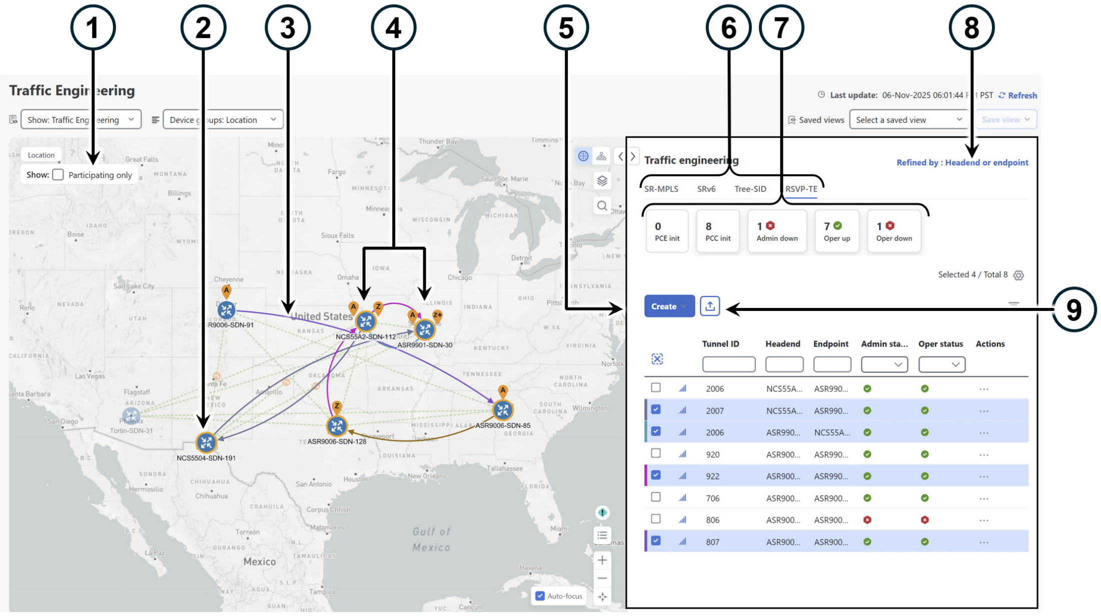

		Tunnel ID	Headend	Endpoint	Admin sta...	Oper status	Actions
□		2006	NCS55A...	ASR990...	✓	✓	...
☑		2007	NCS55A...	ASR990...	✓	✓	...
☑		2006	ASR990...	NCS55A...	✓	✓	...
□		920	ASR900...	ASR990...	✓	✓	...
☑		922	ASR900...	ASR990...	✓	✓	...
□		706	ASR900...	ASR900...	✓	✓	...
□		806	ASR900...	ASR900...	☒	☒	...
☑		807	ASR900...	ASR900...	✓	✓	...

The table describes the callouts for the Traffic Engineering topology map for RSVP-TE tunnels.

Callout no.	Description
1	**Show Participating only:** Displays only links belonging to the selected RSVP-TE tunnels. All other links and devices are hidden.
2	**Device outlines:** • A solid orange outline ( ) indicates a device with a strict hop. • A dashed orange outline indicates a device with a loose hop. **Note** RSVP-TE tunnels cannot be configured with loose hops when provisioning in the UI.

Callout no.	Description
3	**Tunnel visualization:** When you select RSVP-TE tunnels in the RSVP-TE Tunnel table, the map displays colored directional lines showing the source and destination. • Record Route Object (RRO) paths appear as straight lines. • Explicit Route Object (ERO) paths appear as curved lines. **Note** If both RRO and ERO paths are available, the RRO path is displayed by default. • An adjacency segment ID (SID) is shown as a green dot on a link along the path (—●—). If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. • **A+** denotes multiple RSVP-TE tunnels originating from a node. • **Z+** denotes multiple RSVP-TE tunnels terminating at a node.
4	**RSVP-TE tunnel origin and destination:** If both **A** and **Z** are displayed in a device cluster, at least one node in the cluster is a source, and another is a destination. • **A+** denotes multiple RSVP-TE tunnels originating from a node. • **Z+** denotes multiple RSVP-TE tunnels terminating at a node.
5	**Window content:** The window content depends on the selected or filtered items. In this example, the RSVP-TE tab shows RSVP-TE Tunnels table. Depending on the map selection, you can create, modify, or view RSVP-TE tunnels. • Create dynamic RSVP-TE tunnels based on optimization intent, on page 52 • Create explicit RSVP-TE tunnels, on page 52 • Modify RSVP-TE tunnels, on page 54 • View RSVP-TE tunnel details, on page 48
6	**Tabs:** Click the **RSVP-TE** tab to access RSVP-TE data.
7	**Mini dashlets:** Summarizes the operational RSVP-TE tunnel status and displays the number of PCC and PCE-initiated tunnels listed in the RSVP-TE table. When you select a dashlet, filters are applied and the policy table updates to display data corresponding to the filtered dashlet.

Callout no.	Description
8	**Group filter:** Controls how group filters apply to table data. For example, if **Headend only** is selected, the table only displays policies where the policy's headend device is in the selected group. This filter helps users efficiently manage policies in large networks. Filter options: • **Headend or Endpoint:** Show policies with either the headend or endpoint device in the selected group. • **Headend and Endpoint:** Show policies if both the headend and endpoint are in the group. • **Headend only:** Show policies if the headend device of the policy is in the selected group. • **Endpoint only:** Show policies if the endpoint device of the policy is in the selected group.
9	**Export function:** Exports all data into a CSV file. You cannot export selected or filtered data.

View RSVP-TE tunnel details

You can view RSVP-TE tunnel details, including the binding label, delegated PCE, metric type, ERO/RRO, and delay. To view RSVP-TE tunnel details, complete these steps:

Before you begin

To ensure end-to-end delays on RSVP-TE tunnels, all inter-domain RSVP-TE tunnels must be explicit. This means that every interface along the path must be specified as an adjacency hop.

Procedure

Step 1 Click 📶 to view performance metrics for the RSVP-TE tunnel. This data helps assess the policy health and indicates if any of the metrics violated SLAs defined in the Heuristic package. Delay and Delay Variance metrics for RSVP-TE policies are visible only when Crosswork Service Health is installed and SR-PM collection is enabled. The Loss metric provides detailed visibility into packet loss across core network links. You can set loss severity thresholds, in the **System settings** > **Metric Thresholds** page.

Note
The RSVP policy under **Transport** > **RSVP** tab in VPN Services and the **Traffic Engineering** > **RSVP** tab represent the same Traffic Engineering policy. Both pages display RSVP Performance Measurement (PM) metrics with identical values. However, the Threshold label appears only in the VPN Services – Transport tab when Service Health monitoring is enabled and the device has delay measurement configured for the policies. If data retention is enabled, historical data and trends are available in the **History** tab.

Figure 23: Performance metrics

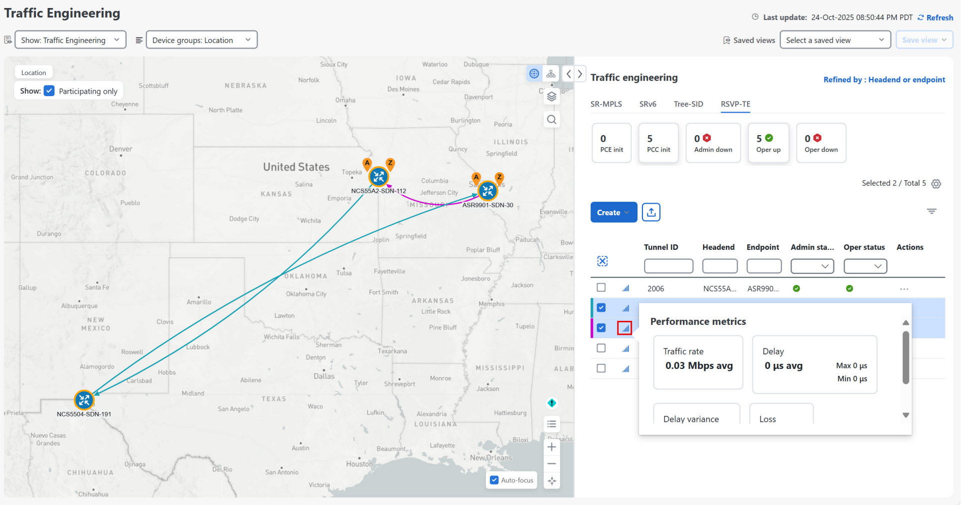

Traffic Engineering

Last update: 24-Oct-2025 08:50:44 PM PDT Refresh

Show: Traffic Engineering | Device groups: Location | Saved views | Select a saved view | Save view

Traffic engineering — Refined by : Headend or endpoint

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

0 PCE init	5 PCC init	0 Admin down	5 Oper up	0 Oper down

Selected 2 / Total 5

Create

		Tunnel ID	Headend	Endpoint	Admin sta...	Oper status	Actions
□		2006	NCS55A...	ASR990...	✓	✓	...
☑							
☑							
□							
□							

Performance metrics

Traffic rate
0.03 Mbps avg

Delay
0 µs avg
Max 0 µs
Min 0 µs

Delay variance

Loss

Step 2 Alternatively, from the **Actions** column, choose ⋯ > **View details** for one of the RSVP-TE tunnels.

Figure 24: RSVP-TE > View details

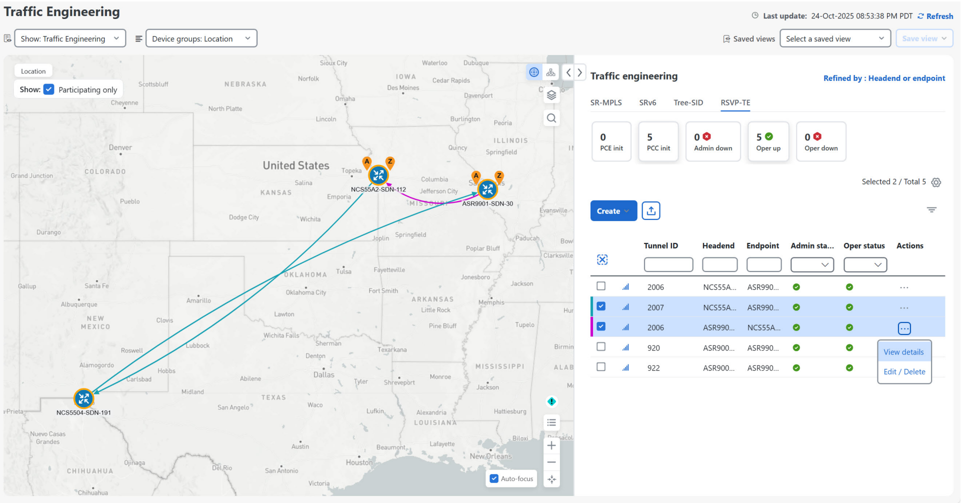

Traffic Engineering

Last update: 24-Oct-2025 08:53:38 PM PDT Refresh

Show: Traffic Engineering | Device groups: Location | Saved views | Select a saved view | Save view

Traffic engineering — Refined by : Headend or endpoint

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

0 PCE init	5 PCC init	0 Admin down	5 Oper up	0 Oper down

Selected 2 / Total 5

Create

		Tunnel ID	Headend	Endpoint	Admin sta...	Oper status	Actions
□		2006	NCS55A...	ASR990...	✓	✓	...
☑		2007	NCS55A...	ASR990...	✓	✓	...
☑		2006	ASR990...	NCS55A...	✓	✓	...
□		920	ASR900...	ASR990...	✓	✓	
□		922	ASR900...	ASR990...	✓	✓	

View details
Edit / Delete

Step 3 View RSVP-TE tunnel details.

Figure 25: RSVP-TE tunnel details

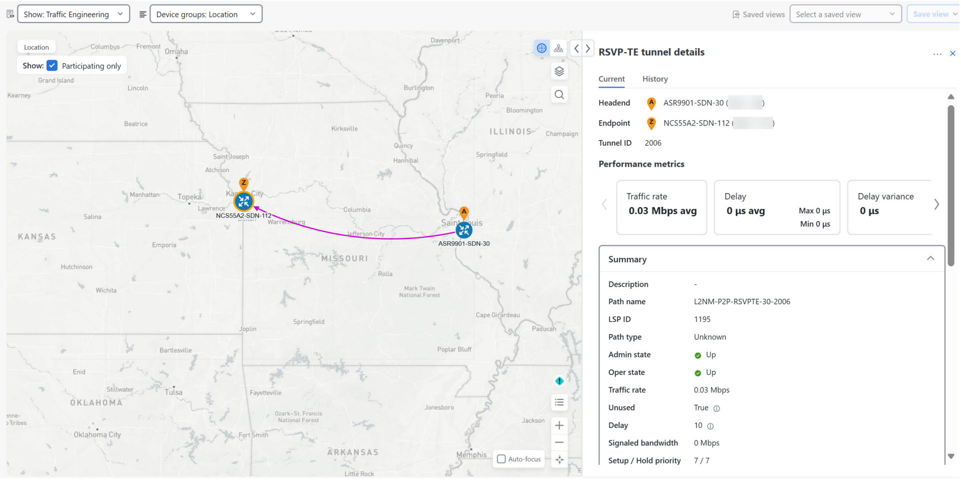

RSVP-TE tunnel details

Current History

Headend A ASR9901-SDN-30 ([redacted])

Endpoint Z NCS55A2-SDN-112 ([redacted])

Tunnel ID 2006

Performance metrics

Traffic rate	Delay		Delay variance
0.03 Mbps avg	0 µs avg	Max 0 µs Min 0 µs	0 µs

Summary

Description	-
Path name	L2NM-P2P-RSVPTE-30-2006
LSP ID	1195
Path type	Unknown
Admin state	● Up
Oper state	● Up
Traffic rate	0.03 Mbps
Unused	True ⓘ
Delay	10 ⓘ
Signaled bandwidth	0 Mbps
Setup / Hold priority	7 / 7

Figure 26: RSVP-TE tunnel details (close-up)

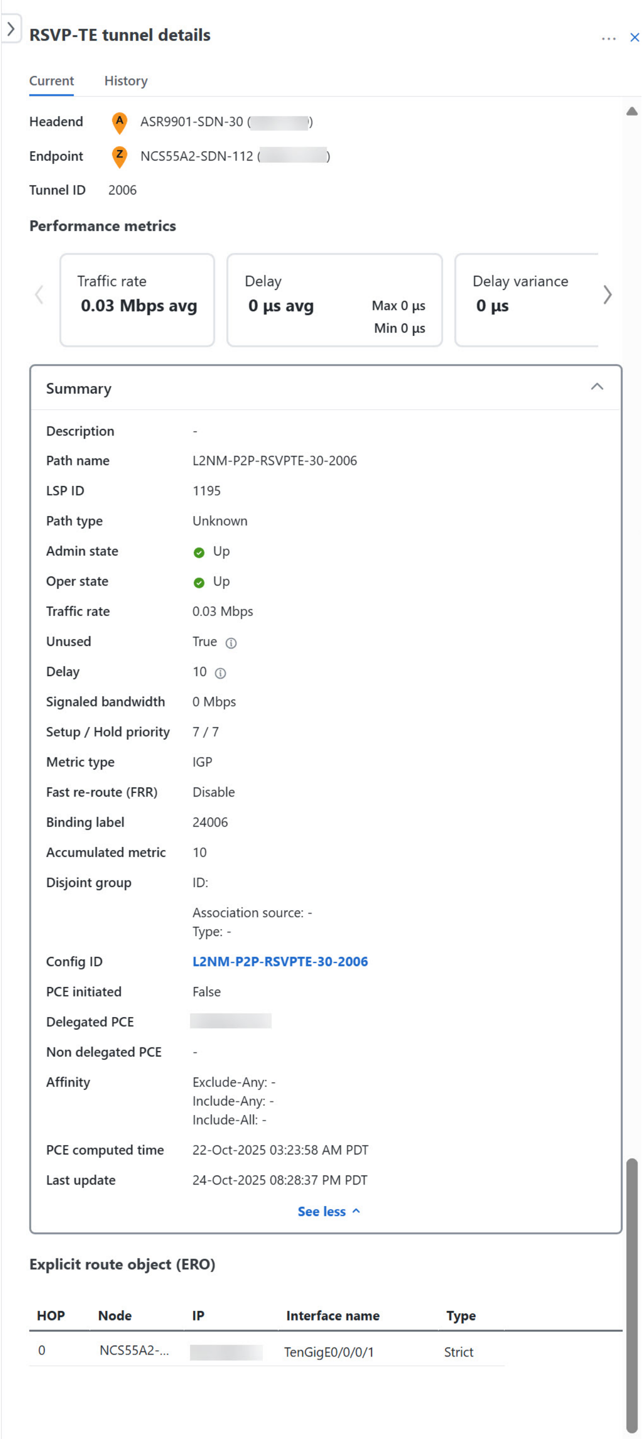

RSVP-TE tunnel details

Headend	A ASR9901-SDN-30 ([redacted])
Endpoint	Z NCS55A2-SDN-112 ([redacted])
Tunnel ID	2006

Performance metrics

Traffic rate	Delay	Delay variance
0.03 Mbps avg	0 µs avg Max 0 µs Min 0 µs	0 µs

Summary

Description	-
Path name	L2NM-P2P-RSVPTE-30-2006
LSP ID	1195
Path type	Unknown
Admin state	✓ Up
Oper state	✓ Up
Traffic rate	0.03 Mbps
Unused	True ⓘ
Delay	10 ⓘ
Signaled bandwidth	0 Mbps
Setup / Hold priority	7 / 7
Metric type	IGP
Fast re-route (FRR)	Disable
Binding label	24006
Accumulated metric	10
Disjoint group	ID: Association source: - Type: -
Config ID	L2NM-P2P-RSVPTE-30-2006
PCE initiated	False
Delegated PCE	[redacted]
Non delegated PCE	-
Affinity	Exclude-Any: - Include-Any: - Include-All: -
PCE computed time	22-Oct-2025 03:23:58 AM PDT
Last update	24-Oct-2025 08:28:37 PM PDT

Explicit route object (ERO)

HOP	Node	IP	Interface name	Type
0	NCS55A2-...	[redacted]	TenGigE0/0/0/1	Strict

Step 4 If the delay value is displayed, it is calculated for all policies every 10 minutes. Click the ⓘ icon next to the delay value to see the last update time.

Create explicit RSVP-TE tunnels

This task creates RSVP-TE tunnels using an explicit (fixed) path, which consists of a list of prefix or adjacency Segment IDs (SID list). Each SID represents a node or link along the path. To create explicit RSVP-TE tunnels, complete these steps:

Before you begin

If your setup includes many nodes, policies, or interfaces, a timeout may occur during policy deployment. To configure timeout options, see Configure TE timeout settings, on page 16.

Procedure

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

Step 2 Click **Create > PCE Init**. To provision a PCC-initiated tunnel using NSO in the Crosswork Network Controller UI, see Create RSVP-TE tunnels (PCC-initiated), on page 54.

Step 3 Under **Tunnel details**, enter the required RSVP-TE tunnel values. To see a description for each field, hover over ⓘ.

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

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

Step 5 Add the segments for the RSVP-TE path.

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

Step 7 To commit the tunnel path, click **Provision**.

Step 8 Validate the RSVP-TE tunnel creation.

a. Confirm that the new RSVP-TE tunnel appears in the RSVP-TE Tunnels table. To highlight the policy on the map, select the check box next to the policy.

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

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

Create dynamic RSVP-TE tunnels based on optimization intent

This task creates an RSVP-TE tunnel with a dynamic path. SR-PCE computes a tunnel path based on the metrics and path constraints, such as affinity or disjointness, that you define. You can select one of three available metrics to minimize in-path computation: IGP, TE, or delay. SR-PCE automatically re-optimizes the path when the topology changes.

Before you begin

Policy deployment considerations

- If your setup includes many nodes, policies, or interfaces, a timeout may occur during policy deployment. To configure timeout options, see Configure TE timeout settings, on page 16.
- To improve visualization, you can optionally collect affinity information from devices. Map this information in Crosswork Network Controller before creating a dynamic SR-MPLS policy. See Configure TE link affinities, on page 37.

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering** > **RSVP-TE**.

Step 2 Click **Create > PCE init**. To provision a PCC-initiated tunnel using NSO in the Crosswork Network Controller UI, see Create RSVP-TE tunnels (PCC-initiated), on page 54.

Step 3 Under **Tunnel details**, enter the required RSVP-TE tunnel values. To see a description for each field, hover over ⓘ.

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

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

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

Step 6 Define any applicable constraints and any required disjointness.

Affinity considerations

- You cannot configure both affinity constraints and disjointness on the same RSVP-TE tunnel.
- There can be up to two RSVP-TE tunnels in the same disjoint group or subgroup. If RSVP-TE tunnels exist in a disjoint group you define here, all tunnels in that group are shown during preview.

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

Step 8 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
Depending on network size and performance, the new RSVP-TE tunnel may take some time to appear in the **Traffic engineering** table. The table refreshes every 30 seconds.

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

Create RSVP-TE tunnels (PCC-initiated)

To create explicit or dynamic RSVP-TE tunnels using the Crosswork Network Controller UI, complete these steps:

Before you begin

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

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Provisioning (NSO)** > **RSVP-TE** > **Tunnel**.

Step 2 Click + and enter a name for the tunnel. Click **Continue**.

Note

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

Step 3 Enter the required policy constraints and values.

Step 4 Click **Dry run** to validate and save your changes. Crosswork Network Controller 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

You can modify or delete RSVP-TE tunnels created using the UI or the API.

To view, modify, or delete an RSVP-TE tunnel, complete these steps:

Procedure

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

Step 2 Locate the RSVP-TE tunnel you want to modify, and click ⋯.

Step 3 Choose **View details** or **Edit/Delete**. After you update the RSVP-TE tunnel details, preview the changes on the map, and save them.

CHAPTER 4

Flexible Algorithm

Flexible algorithm

A flexible algorithm is a customizable IGP routing method that:

- enables operators to define path computation constraints based on specific metrics and link properties.
- allows confining the path to a particular logical plane in networks with multiple planes.
- supports user-defined meanings and intent for routing behaviors within a network.

A flexible algorithm filters or confines the IGP topology to meet specific transport characteristics or policies, rather than just computing a default shortest path to all destinations. Crosswork Network Controller enables visualization of network subsets that provide a specific set of transport characteristics. This visualization helps you deploy, maintain, and verify that the flexible algorithm’s intended behavior is realized in your network.

For example:

- You can use a flexible algorithm to improve service availability.
- You can define disjoint logical topologies to enhance network resiliency against failures.
- You can verify that two flexible algorithm topologies have no shared nodes or links, ensuring complete disjointness.
- If common nodes or link exist, you can identify them and update configurations to better meet your network goals.

Supported Flexible Algorithm metrics and constraints

The Crosswork Network Controller supports a limited set of metrics and constraints compared to those available in Cisco IOS XR. If you configure the Flexible Algorithm with unsupported metrics and constraints, the feature operates as expected on the routers. However, the SR policy path and topology visualization in the Crosswork Network Controller UI may be inaccurate.

Metrics

The Crosswork Network Controller supports these metrics:

- IGP metric
- ASLA traffic engineering default: required, or the link is pruned.
- ASLA min unidirectional link delay: required, or the link is pruned.

Constraints

In Application-Specific Link Attribute (ASLA), Crosswork Network Controller supports these constraints:

- include-any
- include-all
- exclude
- exclude srlg: Links belonging to specified Shared Risk Link Groups (SRLGs) can be excluded from Flexible Algorithm topology calculations and visualization.

Configure flexible algorithm affinities

Crosswork Network Controller does not automatically collect flexible algorithm affinity names from devices. For consistency with flexible algorithm affinity names, you can define affinities in Crosswork Network Controller with the same names and bit positions that are used on the device. This helps maintain consistency and improve visualization.

Note

- The affinity bit position ranges from 0–255 allowing more granular affinity definitions.
- If an affinity mapping is not defined in the UI, the affinity name will appear as "UNKNOWN".
- Crosswork Network Controller sends the bit information to the SR-PCE only during provisioning.

On devices, affinity maps are configured by setting bits for each affinity name. 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
!
```

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*).

To add flexible algorithm affinities in Crosswork Network Controller, complete these steps:

Procedure

Step 1 Choose **Administration** > **Settings** > **Traffic engineering** > **Affinity** > **Flex-Algo affinities**.

Step 2 Click **+ Create** to add a new affinity mapping.

Step 3 Enter the name and the bit position corresponding to the device configuration.

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

Step 5 To view all flexible algorithm affinities for a link, see View flexible algorithm details , on page 61.

Visualize flexible algorithm topologies

Crosswork Network Controller allows you to visualize flexible algorithm nodes and links on the topology map. You can view nodes and links that have been manually configured or dynamically provisioned using the UI.

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 39.

Before you begin

Configure flexible algorithms in your network. Refer to the SR flexible algorithm configuration documentation for your specific device to view descriptions and supported configuration commands (for example, see the *Segment Routing Configuration Guide for Cisco NCS 540 Series Routers*).

Note Visualization of flexible algorithms is not possible if the same flexible algorithm ID is used across different domains.

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering**.

Step 2 On the topology map, click .

Step 3 Click **Flex Algo**.

Step 4 From the drop-down list, select up to two flexible algorithm IDs.

Step 5 View the flexible algorithm types and confirm that your selections are correct. Note the color assignments for each flexible algorithm.

Step 6 (Optional) Select **Show selected Flex Algo topology only** to isolate the flexible algorithms on the topology map. When you enable this option, SR policy selection is disabled.

a) Select **Show Flex Algo A+B links only** to display links and nodes that participate in both flexible algorithms.

Step 7 After making changes to flexible algorithm selections, click **Apply** to display the updated topology map.

Figure 27: Flexible algorithm on map

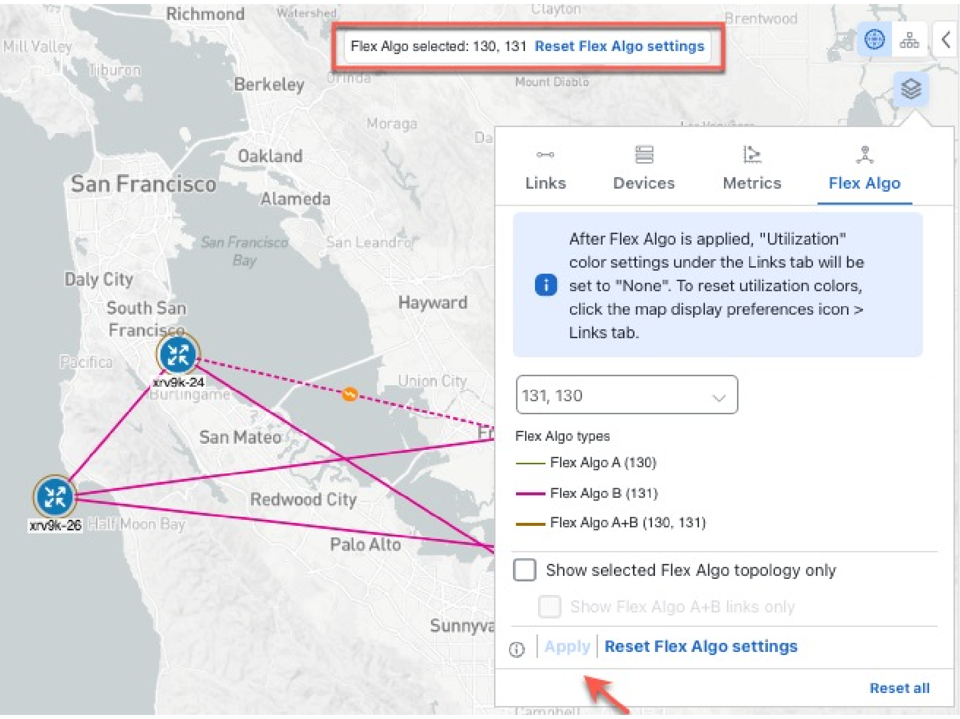

Flex Algo selected: 130, 131 Reset Flex Algo settings

After Flex Algo is applied, "Utilization" color settings under the Links tab will be set to "None". To reset utilization colors, click the map display preferences icon > Links tab.

131, 130

Flex Algo types

- Flex Algo A (130)
- Flex Algo B (131)
- Flex Algo A+B (130, 131)

☐ Show selected Flex Algo topology only

☐ Show Flex Algo A+B links only

Apply | Reset Flex Algo settings

Reset all

Note

If a selected flexible algorithm is defined with criteria, but no link and node combinations match those criteria (for example, an affinity to include all nodes or links with the color blue), the topology map appears blank. When a selected flexible algorithm is not configured on a node or link, the topology map shows the default blue color for that link or node.

Step 8 Click **Save View** to save the topology view and flexible algorithm selections.

View flexible algorithm details

Before you begin

Flexible algorithm considerations

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

To view device or link flex algorithm details, complete these steps:

Procedure

Step 1 Choose **Services & Traffic Engineering** > **Traffic Engineering**.

Step 2 View device flexible algorithm details.

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

b) In the **Device details** page, choose **Traffic engineering** > **Flex Algo**. Then, select the Flexible Algorithm for which you would like to view details.

The **Elected definition** displays all metrics and constraints defined for this flexible algorithm.

Figure 28: Flex Algo device details

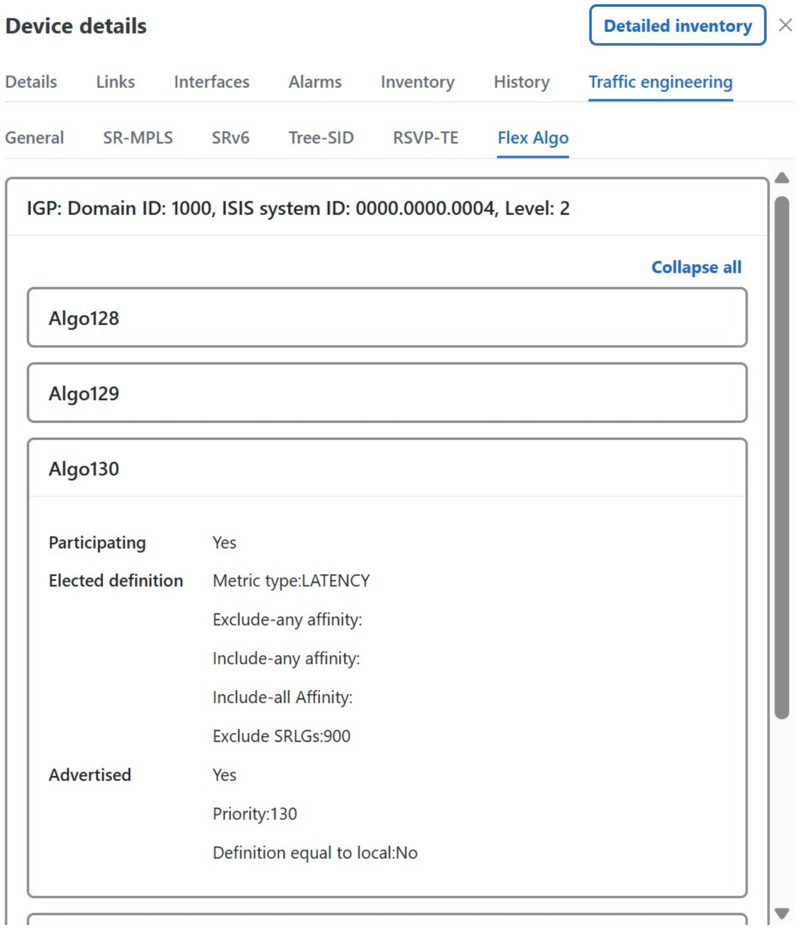

IGP: Domain ID: 1000, ISIS system ID: 0000.0000.0004, Level: 2

Participating	Yes
Elected definition	Metric type:LATENCY Exclude-any affinity: Include-any affinity: Include-all Affinity: Exclude SRLGs:900
Advertised	Yes Priority:130 Definition equal to local:No

Step 3 To view whether a link is part of a flexible algorithm topology:

a) On the topology map, click on a link.

b) In the **Link details** page, click the **Traffic engineering** tab.

- If the link is part of a flexible algorithm topology, the **FA topologies** row shows the flexible algorithm(s) to which the source and destination devices belongs.
- If SRLG exclusion constraints are defined for a flexible algorithm and the link is a member, the **FA SRLGs** row displays the configured SRLG values for the source and destination devices.

Figure 29: Flex Algo link details

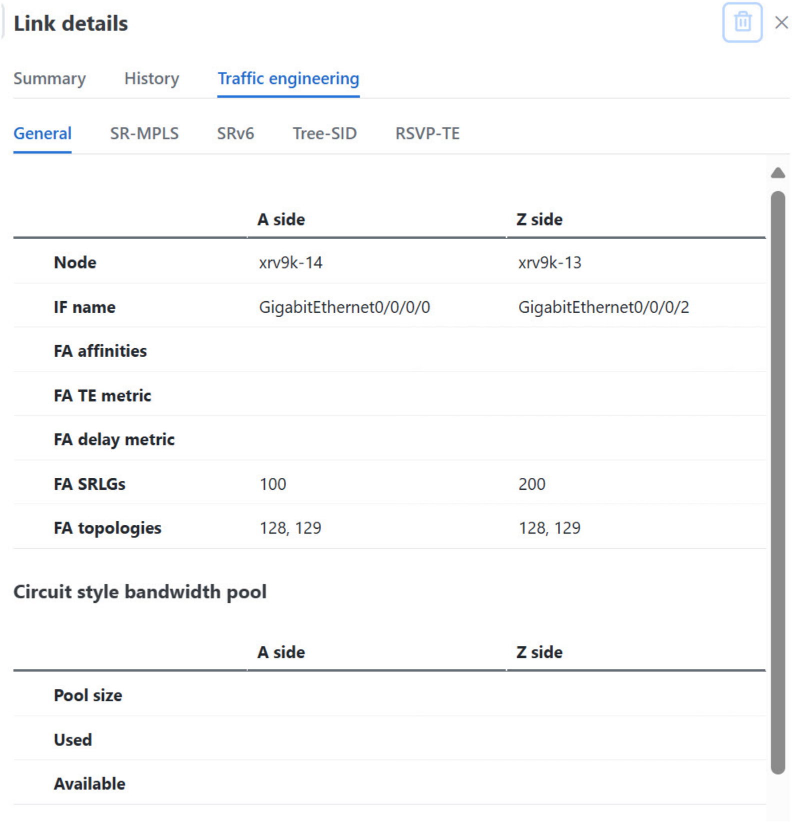

Link details

	A side	Z side
Node	xrv9k-14	xrv9k-13
IF name	GigabitEthernet0/0/0/0	GigabitEthernet0/0/0/2
FA affinities		
FA TE metric		
FA delay metric		
FA SRLGs	100	200
FA topologies	128, 129	128, 129

Circuit style bandwidth pool

	A side	Z side
Pool size		
Used		
Available		

Configure and visualize flexible algorithm SRLG exclusion

Flexible algorithms can be configured to exclude groups of links that share common risks, known as Shared Risk Link Groups (SRLGs). This allows operators to prune links from the flexible algorithm topology based on shared risk attributes, ensuring that paths avoid common failure points and enhance network resiliency. When you set an SRLG exclusion for a flexible algorithm, the excluded SRLGs are automatically filtered out from the topology visualization and routing calculations. This helps operators plan alternate, protected paths and avoid risks tied to certain links.

Note This feature is supported on all devices running Cisco IOS XR software version 25.1.1 or later.

To configure flexible algorithm SRLG exclusion:

Procedure

Step 1 Configure SRLGs on your devices. Create a group to identify specific SRLG and give a unique value. By associating an interface with an SRLG, you can apply policies or constraints based on these groups.

Note
The SRLG value is displayed only for IPv4 links and is not shown for IPv6 links.

Example:

Cisco IOS XR - SRLG definition

```
srlg
 interface GigabitEthernet0/0/0/4
  name groupA
 !
 name groupA  !! user-defined group names to identify specific SRLG
  value 900 !! user-assigned numerical identifiers for the SRLG groups
 !
 name groupB
  value 800
 !
```

Step 2 Configure the flexible algorithm definition with the SRLG exclusion constraint on your devices.

Example:

Cisco IOS XR - ISIS routing configuration

```
router isis CORE
 address-family ipv4 unicast
  advertise application flex-algo link-attributes srlg
 !
 flex-algo 129
  priority 129
  srlg exclude-any groupB
  advertise-definition
 !
```

This configuration advertises the relevant link and SRLG attributes via IS-IS, so other routers can also compute consistent, SRLG-aware paths. It ensures that for Flex-Algo 129, any links that are part of groupB are excluded from path computation and routing.

Note

- Ensure the flexible algorithm priority is set higher than 128 (for example, 129 in the example) for the exclude SRLG configuration to be correctly reflected and applied.
- The `advertise application flex-algo link-attributes srlg` configuration command is not required for OSPF routing protocol.

Step 3 Crosswork Network Controller automatically discovers these SRLG exclusion constraints from the network. The flexible algorithm for the device then reflects the exclusion. To view these details:

a) Choose **Services & Traffic Engineering** > **Traffic Engineering**.
b) On the topology map, click on a device. In the **Device details** page, choose **Traffic engineering** > **Flex Algo**.

   Algo129 displays Exclude SRLG as an elected definition.

Figure 30: Flex Algo device details with SRLG exclusion

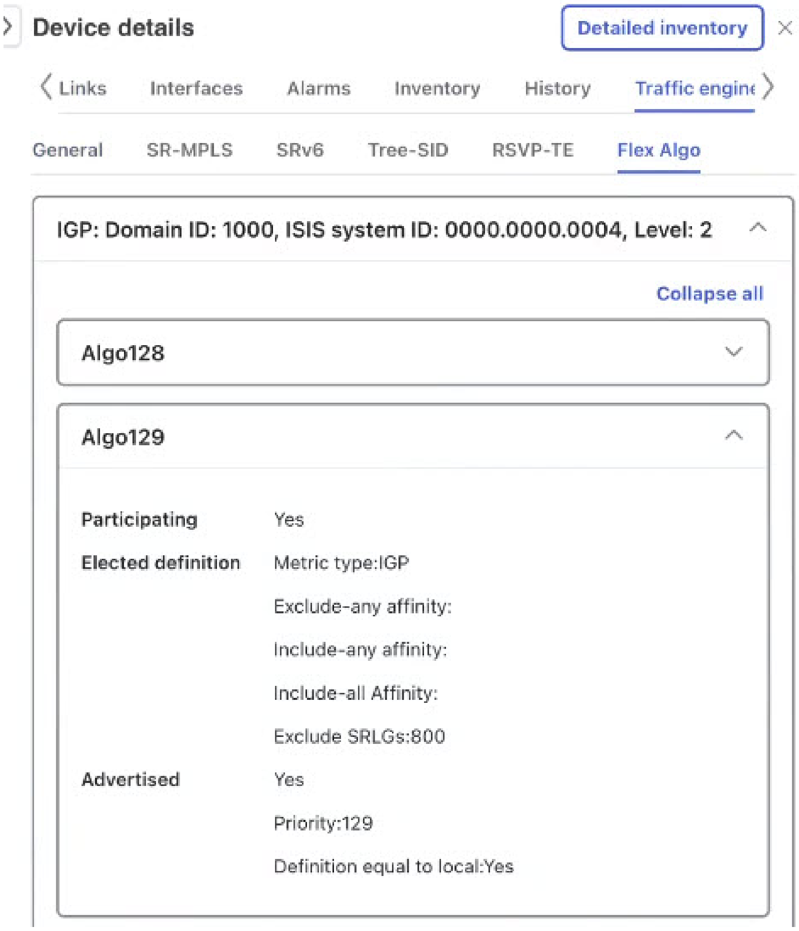

IGP: Domain ID: 1000, ISIS system ID: 0000.0000.0004, Level: 2

Algo128

Algo129

Participating	Yes
Elected definition	Metric type:IGP
	Exclude-any affinity:
	Include-any affinity:
	Include-all Affinity:
	Exclude SRLGs:800
Advertised	Yes
	Priority:129
	Definition equal to local:Yes

Step 4 To view whether a link is part of a flexible algorithm topology:

a) On the topology map, click on a link.

b) In the **Link details** page, click the **Traffic engineering** tab. If the link is a member, the **FA SRLGs** row displays the configured SRLG values for the source and destination devices.

Figure 31: Flex Algo link details with FA SRLGs

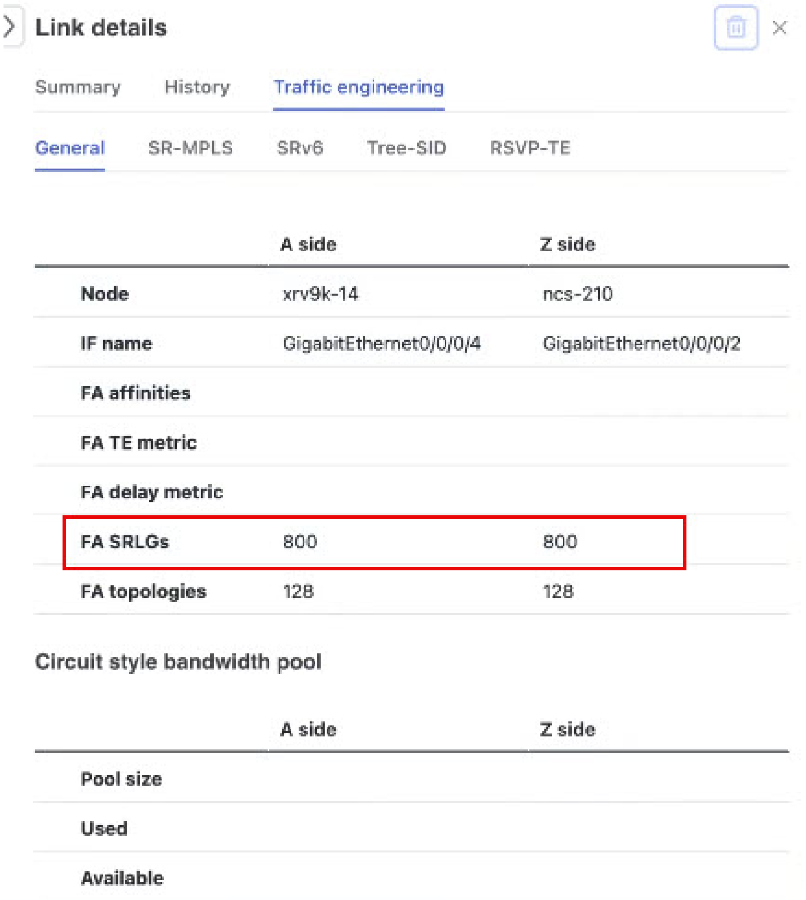

	A side	Z side
Node	xrv9k-14	ncs-210
IF name	GigabitEthernet0/0/0/4	GigabitEthernet0/0/0/2
FA affinities		
FA TE metric		
FA delay metric		
FA SRLGs	800	800
FA topologies	128	128

Circuit style bandwidth pool

	A side	Z side
Pool size		
Used		
Available		

Step 5 To visualize the flexible algorithm topology with SRLG exclusions:

a) On the topology map, click ≋ and then click **Flex Algo**.

b) From the drop-down list, select the flexible algorithm ID (for example, 129) to filter and click **Apply**.

The updated topology map displays the flexible algorithm path. Links that were excluded due to SRLG constraints are marked differently to indicate their exclusion from this specific topology.

Figure 32: Flexible algorithm on map with excluded link

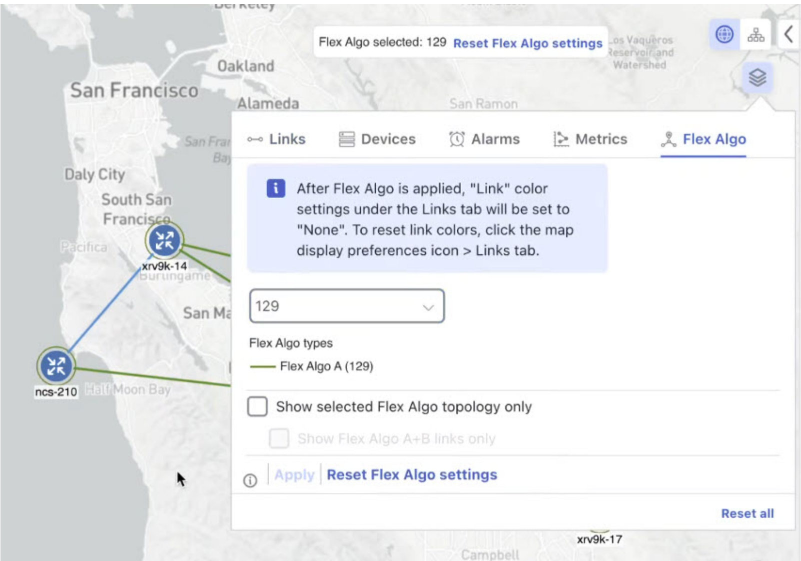

Flex Algo selected: 129 Reset Flex Algo settings

After Flex Algo is applied, "Link" color settings under the Links tab will be set to "None". To reset link colors, click the map display preferences icon > Links tab.

129

Flex Algo types

Flex Algo A (129)

☐ Show selected Flex Algo topology only

☐ Show Flex Algo A+B links only

CHAPTER 5

Tree Segment Identifier Multicast Traffic Engineering

Tree Segment Identifier

A Tree-SID is a method for implementing tree-like multicast flows over a segmented routing network that:

- assigns a single SID label to represent all segments and devices within a multicast tree,
- uses an SDN controller (such as a device running SR-PCE with PCEP) to calculate the tree structure, and
- designates a specific role to each node within the tree for forwarding multicast data.

The SDN controller constructs routing paths using any constraints defined by network architects.

A key use case for constraint-based Tree-SID is when routers are configured to deliver two point-to-multipoint (P2MP) streams carrying the same content over different paths. In this scenario, the multicast flow is forwarded twice, with each copy following a unique path through the network. The two copies use different nodes and links to reach the destination. This approach reduces packet loss because a network failure on one path does not impact the delivery of the multicast stream on the other path.

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).

Tree-SID node roles

Crosswork Network Controller provides the ability to view the details of Tree-SID root, transit, leaf, and bud nodes in the UI, allowing you to easily confirm proper Tree-SID implementation in your network (see View point-to-multipoint trees on the topology map, on page 70).

Tree-SID policies define four distinct node roles, each of which serves a specific function in multicast data forwarding:

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

Tree-SID policy types

Crosswork Network Controller allows you to visualize these Tree-SID policy types.

- **Static**: Configured via SR-PCE, either directly using SR-PCE CLI or through the Crosswork Network Controller UI. For more information and supported configuration commands, for your specific device, refer to the Tree-SID configuration documentation (for example, Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers.
- **Dynamic**: Not explicitly configured. It is configured as part of an L3VPN or multicast VPN (mVPN) service.

Note You can use either static or dynamic Tree-SID policies to enable fast reroute (FRR) capabilities in your multicast network.

View point-to-multipoint trees on the topology map

Before you begin

To visualize multicast trees 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).

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

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

Figure 33: Create a new Tree-SID policy (static)

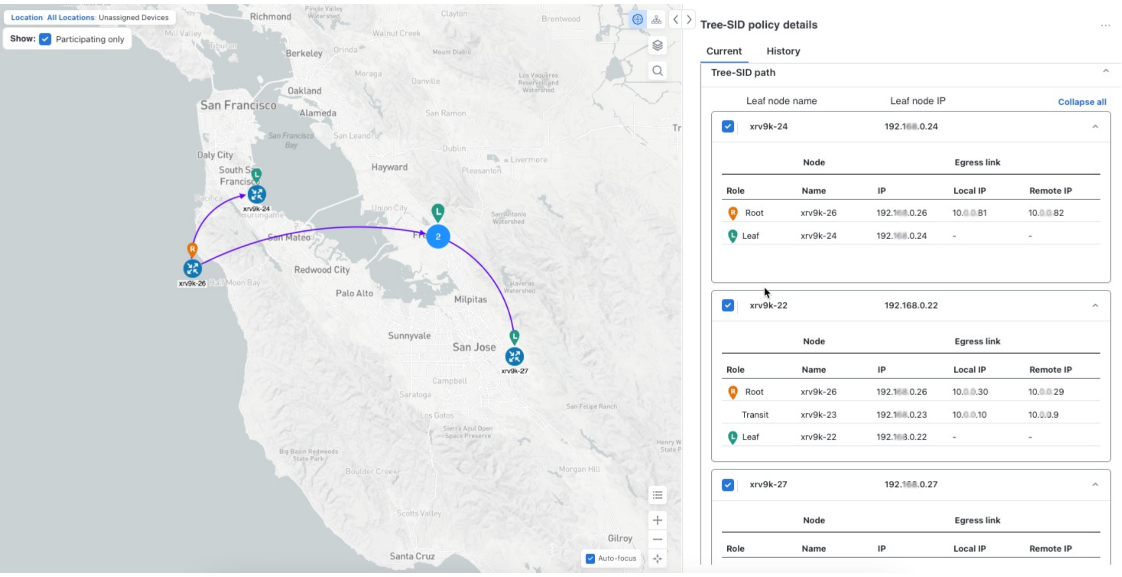

Location: All Locations: Unassigned Devices

Show: ☑ Participating only

☑ Auto-focus

Tree-SID policy details

Current | History

Tree-SID path

Leaf node name | Leaf node IP | Collapse all

☑ xrv9k-24 | 192.1[redacted].0.24

Node			Egress link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	192.1[redacted].0.26	10.[redacted].81	10.[redacted].82
Leaf	xrv9k-24	192.[redacted].0.24	-	-

☑ xrv9k-22 | 192.168.0.22

Node			Egress link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	192.1[redacted].0.26	10.[redacted].30	10.[redacted].29
Transit	xrv9k-23	192.1[redacted].0.23	10.[redacted].10	10.[redacted].9
Leaf	xrv9k-22	192.1[redacted].0.22	-	-

☑ xrv9k-27 | 192.1[redacted].0.27

Node			Egress link	
Role	Name	IP	Local IP	Remote IP

Procedure

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

Step 2 Select the Tree-SID policies you want to view on the topology map. Only two policies can be displayed on the topology map at a time.

Figure 34: Tree-SID policies (static) on the topology map

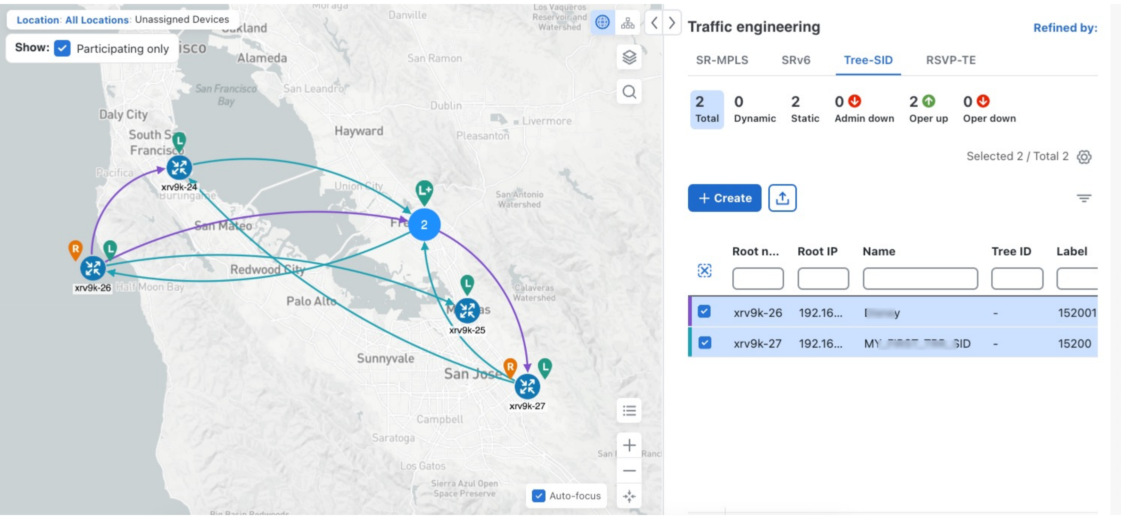

Location: All Locations: Unassigned Devices

Show: ☑ Participating only

☑ Auto-focus

Traffic engineering | Refined by:

SR-MPLS | SRv6 | Tree-SID | RSVP-TE

2 Total | 0 Dynamic | 2 Static | 0 Admin down | 2 Oper up | 0 Oper down

Selected 2 / Total 2

+ Create

	Root n...	Root IP	Name	Tree ID	Label
☑	xrv9k-26	192.16...	[redacted]y	-	152001
☑	xrv9k-27	192.16...	MY[redacted]SID	-	15200

Note

Any change in end-points is recorded as an event in the historical data tab. For information about Tree-SID historical data, see View TE event and utilization history, on page 12.

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

Example:

Figure 35: Tree-SID details summary

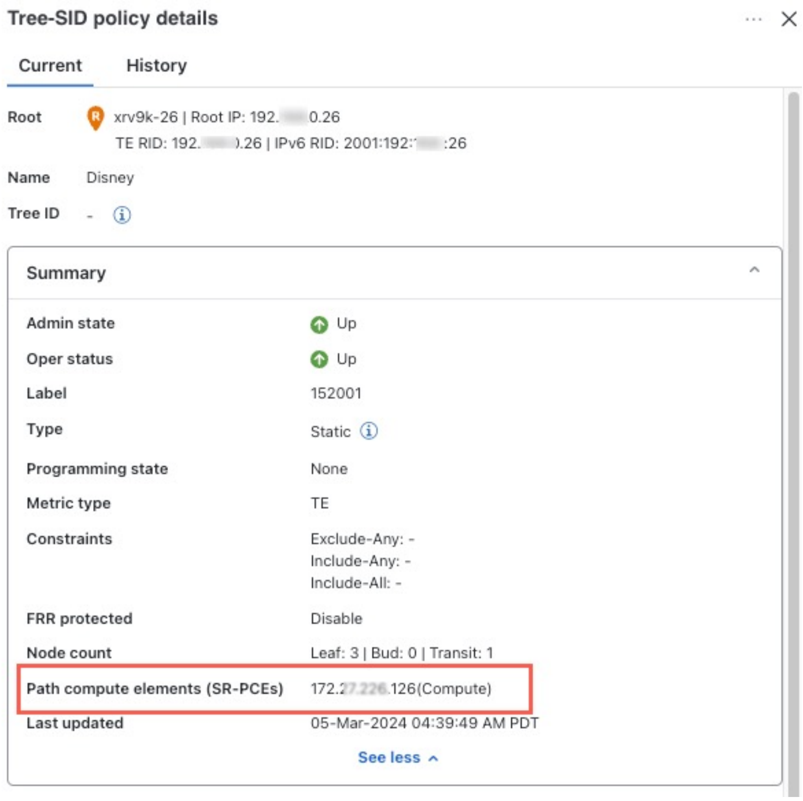

Tree-SID policy details

Root xrv9k-26 | Root IP: 192.[redacted]0.26
TE RID: 192.[redacted].26 | IPv6 RID: 2001:192:[redacted]: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
Node count	Leaf: 3 | Bud: 0 | Transit: 1
Path compute elements (SR-PCEs)	172.2[redacted].126(Compute)
Last updated	05-Mar-2024 04:39:49 AM PDT

Note

- A (Compute) label, next to the SR-PCE field identifies the SR-PCE used to create the policies.
- If a source node is unavailable, a warning icon and message appear next to the Oper Status field. Hover the mouse over the warning icon to see where the connection issue is located.

Figure 36: Tree-SID path details

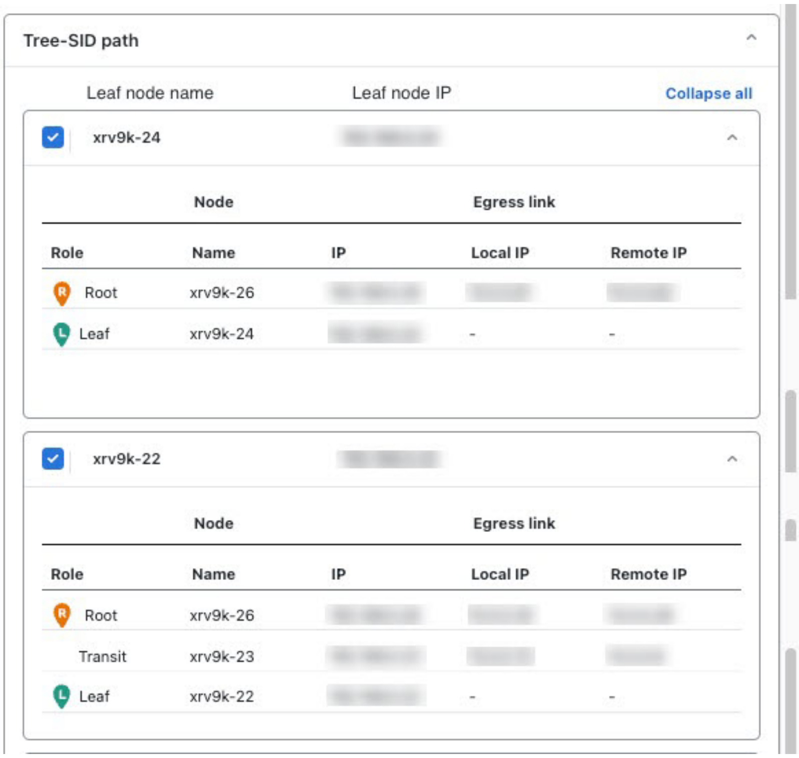

Tree-SID path

Leaf node name	Leaf node IP	Collapse all
☑ xrv9k-24	[redacted]	

Node			Egress link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	[redacted]	[redacted]	[redacted]
Leaf	xrv9k-24	[redacted]	-	-

Leaf node name	Leaf node IP
☑ xrv9k-22	[redacted]

Node			Egress link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-26	[redacted]	[redacted]	[redacted]
Transit	xrv9k-23	[redacted]	[redacted]	[redacted]
Leaf	xrv9k-22	[redacted]	-	-

Create static Tree-SID policies

Before you begin

If you plan to use affinities, collect affinity information from your devices. Map the affinities in Cisco Crosswork Network Controller before creating a static Tree-SID policy. For more information, see Configure TE link affinities, on page 37.

To create a static Tree-SID policy that represents a leaf or a root node, complete these steps:

Procedure

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

Step 2 Enter or select the required Tree-SID policy values. To view a description of a field, hover the mouse pointer over the ⓘ icon.

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

Figure 37: Create static tree-SID policy

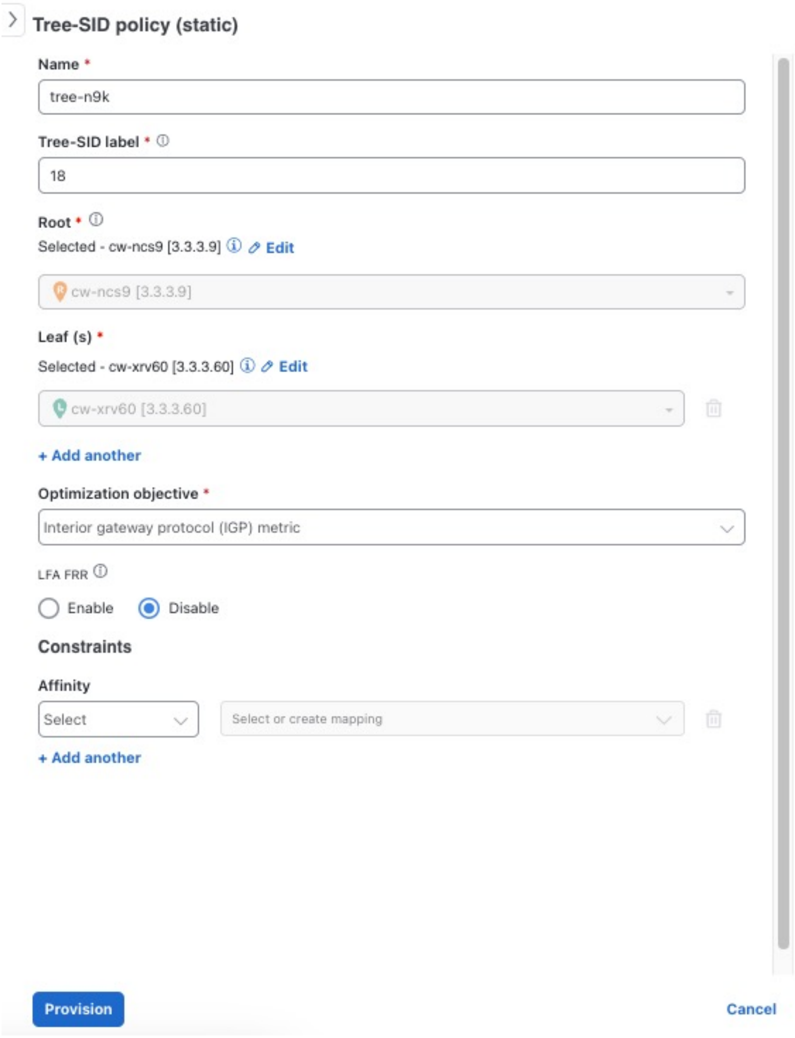

Tree-SID policy (static)

Name * tree-n9k

Tree-SID label * 18

Root * Selected - cw-ncs9 [3.3.3.9] Edit
cw-ncs9 [3.3.3.9]

Leaf (s) * Selected - cw-xrv60 [3.3.3.60] Edit
cw-xrv60 [3.3.3.60]

Optimization objective * Interior gateway protocol (IGP) metric

LFA FRR ○ Enable ● Disable

Constraints

Affinity Select | Select or create mapping

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

Step 4 Validate the Tree-SID policy creation.

a. Confirm that the new Tree-SID policy appears in the **Traffic engineering** table. To highlight the policy on the map, select the check box next to the policy.

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

Figure 38: Newly added tree-SID policy on topology map

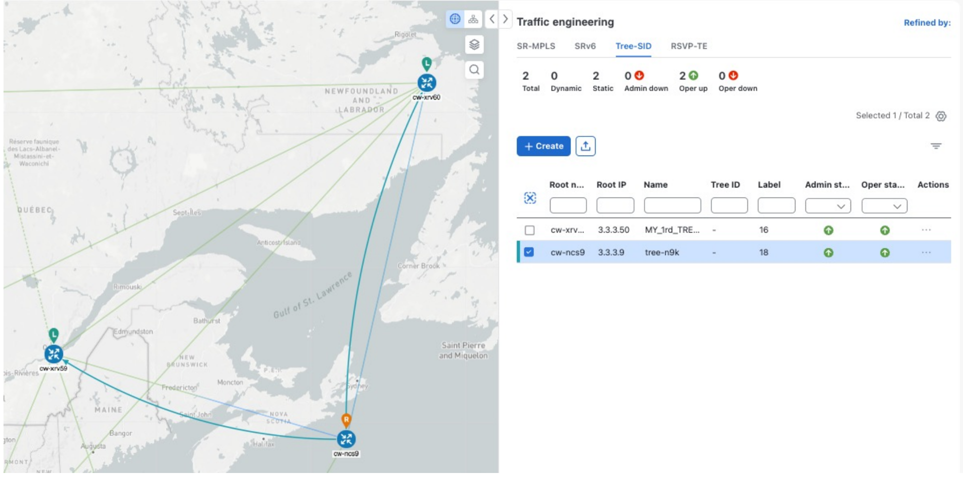

Traffic engineering

2	0	2	0	2	0
Total	Dynamic	Static	Admin down	Oper up	Oper down

Selected 1 / Total 2

	Root n...	Root IP	Name	Tree ID	Label	Admin st...	Oper sta...	Actions
□	cw-xrv...	3.3.3.50	MY_1rd_TRE...	-	16	↑	↑	⋯
☑	cw-ncs9	3.3.3.9	tree-n9k	-	18	↑	↑	⋯

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

Figure 39: Tree-SID policy details

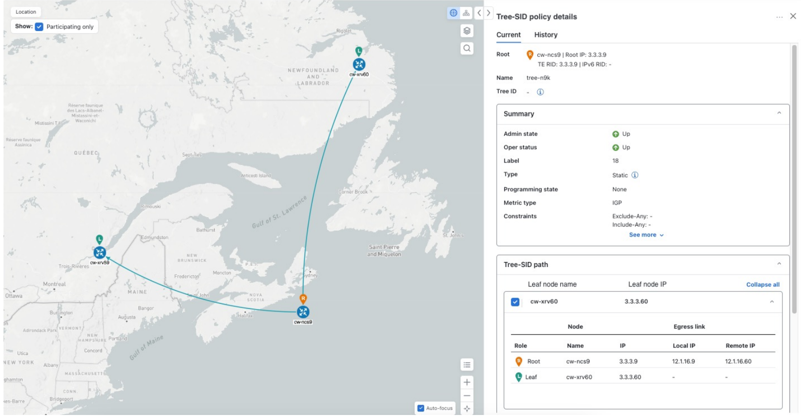

Tree-SID policy details

Root: cw-ncs9 | Root IP: 3.3.3.9
TE RID: 3.3.3.9 | IPv6 RID: -

Name: tree-n9k

Tree ID: -

Summary

Admin state	Up
Oper status	Up
Label	18
Type	Static
Programming state	None
Metric type	IGP
Constraints	Exclude-Any: - Include-Any: -

Tree-SID path

Leaf node name	Leaf node IP
cw-xrv60	3.3.3.60

Node			Egress link	
Role	Name	IP	Local IP	Remote IP
Root	cw-ncs9	3.3.3.9	12.1.16.9	12.1.16.60
Leaf	cw-xrv60	3.3.3.60	-	-

Static Tree-SID policy configuration

The following output shows a static Tree-SID policy configured from the Crosswork Network Controller 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 below output shows the same static Tree-SID policy as viewed 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

Before you begin

Tree SID modification considerations

- You cannot modify the name, label, or root of a Tree-SID policy.
- You can modify or delete only a static Tree-SID policy created using the Crosswork Network Controller UI or API. You cannot modify or delete static Tree-SID policies created using SR-PCE CLI.

To view, modify, or delete a Tree-SID policy, complete these steps:

Procedure

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

Step 2 Locate the Tree-SID policy you want to modify, and click ⋯.

Step 3 Choose **View details** or **Edit/Delete**. After you update the Tree-SID policy details, preview the changes on the map, and save them.

Tree-SID important considerations

Limitations

Working with Tree-SID policies has these limitations:

- 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 UI, but is not supported if configured directly on the SR-PCE CLI.
- Tree-SID policy details based on SRv6 are not supported.
- If you use a single instance of SR-PCE, note that all static Tree-SID policies configured from the UI are deleted if the SR-PCE restarts.
- IPv4 unnumbered interfaces are not supported.

Visualization of Tree-SID paths with missing nodes

If a node on a Tree-SID policy path is not added to the Crosswork device inventory, it will not appear in the Crosswork Network Controller topology information. As a result, one or more root-to-leaf paths may look broken in the topology map, even though the right panel will still display the full Tree-SID policy path.

Figure 40: Tree-SID path visualization

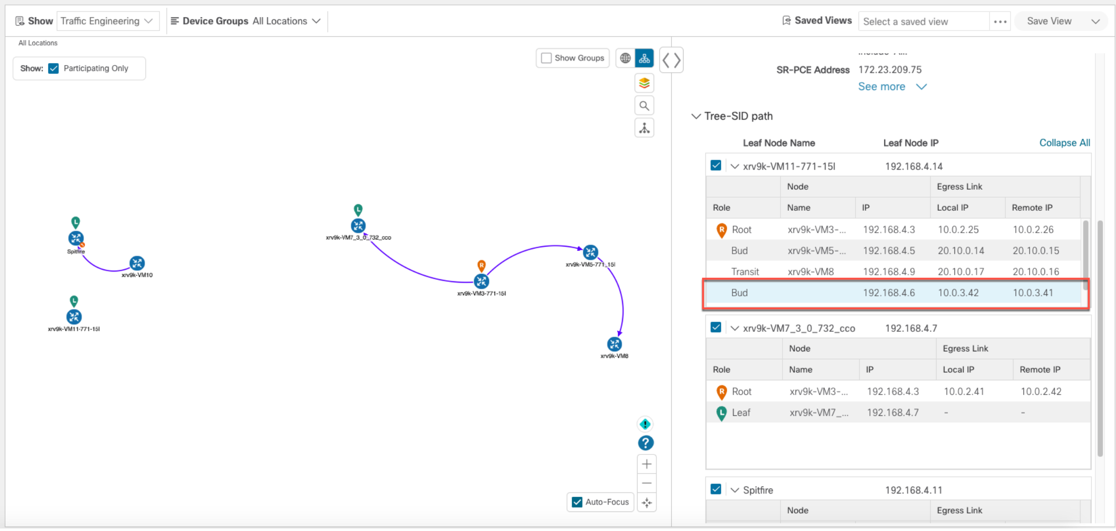

SR-PCE Address 172.23.209.75

See more

Tree-SID path

Leaf Node Name	Leaf Node IP
☑ xrv9k-VM11-771-15I	192.168.4.14

	Node		Egress Link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-VM3-...	192.168.4.3	10.0.2.25	10.0.2.26
Bud	xrv9k-VM5-...	192.168.4.5	20.10.0.14	20.10.0.15
Transit	xrv9k-VM8	192.168.4.9	20.10.0.17	20.10.0.16
Bud		192.168.4.6	10.0.3.42	10.0.3.41

Leaf Node Name	Leaf Node IP
☑ xrv9k-VM7_3_0_732_cco	192.168.4.7

	Node		Egress Link	
Role	Name	IP	Local IP	Remote IP
Root	xrv9k-VM3-...	192.168.4.3	10.0.2.41	10.0.2.42
Leaf	xrv9k-VM7_...	192.168.4.7	-	-

Leaf Node Name	Leaf Node IP
☑ Spitfire	192.168.4.11

Collapse All

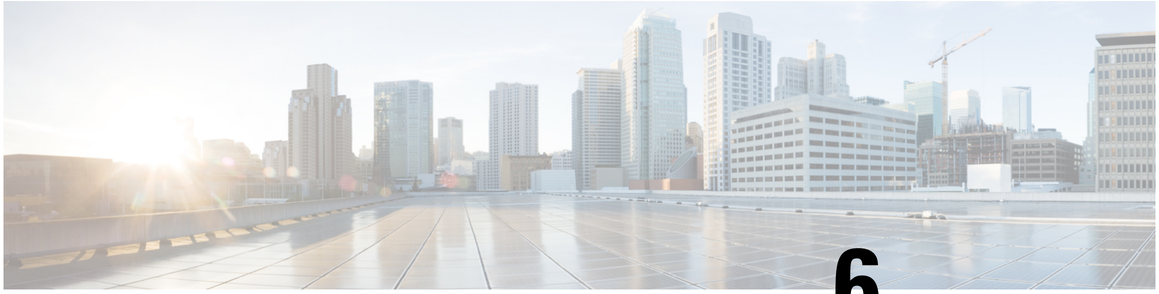

CHAPTER 6

Deterministic Demand Matrix

Demand matrix

The Demand Matrix is a representation of the aggregate SRv6 traffic flows within a network domain. Each flow (demand) corresponds to the total SRv6 traffic entering the domain at a specific node and exiting at another, destined for a specific Locator. Demands are associated with specific destination SRv6 Locators, which define how traffic is forwarded within the domain based on any relevant flex-algo assigned to the Locator. The Demand Matrix, also referred to as the Deterministic Demand Matrix (DDM), provides

- **detailed, near real-time visibility** into SRv6 traffic demands within IGP domains, and
- **granular insights** into traffic patterns, traffic volume, and where potential bottlenecks may arise.

It achieves this by leveraging per-locator, per-egress interface counters from network devices. Refer to the Segment Routing v6 Configuration Guide for Cisco 8000 Series Routers, Cisco IOS XR Releases guide for information on how SRv6 locator counters track external and internal traffic flows and calculate net external traffic using the demand matrix.

Advantages of using DDM

Using DDM offers several key benefits for network monitoring, analysis, and planning:

- **Enhanced visibility**: Gain clear insights into traffic flow demands across the network through the Crosswork Network Controller UI.
- **Per-Domain insights**: Easily manage and monitor individual IGP domains. View metadata about active/total nodes and traffic reporting status within each domain.
- **Effective Network Planning**: Filter demands to view those routed over specific links or interfaces, supporting targeted troubleshooting and planning by identifying traffic distribution and potential bottlenecks.

How DDM Works

At its core, DDM continuously collects per-locator, per-egress interface (LOC.INT.E) counters directly from network devices. It uses this counter data, along with network topology to deterministically compute traffic demands for each IGP domain. These computed demands are then used to visualize traffic patterns overlaid on the topology map.

DDM uses a user-defined cadence to determine how often the demand matrices are updated. When a computation interval is reached, DDM retrieves the latest network topology to perform demand computation. Once computed, these demands are used for path visualization and detailed link views. DDM also records domain metadata, such as the total and active nodes, providing context on domain health and activity.

- If a device misses sending a counter update, DDM uses the previously received counter records to continue computing demands, ensuring continuity.
- DDM has two-cadence grace period to maintain operational stability. This period ensures that late-arriving counter updates are processed, and that demands (such as those from a removed node or link) or domains stay visible in the UI for up to two cadences before being cleared.

DDM device and configuration requirements

For DDM to provide visibility into traffic demands, your network devices must meet these requirements and configurations.

Platform compatibility

Your devices must meet these requirements.

- Devices must be running Cisco IOS-XR version 7.10.1 or higher to support the LOC.INT.E counters. Most platforms support this by version 24.3.1.

Note Currently, there is no support for Cisco XE or multi-vendor environments.

- Enable gNMI for all participating devices, by adding gNMI credentials profile, and gNMI protocol connectivity. For instructions, see **Configure device for gNMI** in Cisco Crosswork Network Controller Administration guide.

Configure and activate performance policy

DDM relies on specific performance policies to collect necessary SRv6 locator from devices. To collect locator counters, you must configure and activate an **SRv6 traffic accounting** policy.

To create the policy, complete these steps:

Procedure

Step 1 From the main menu, choose **Device Management** > **Performance Policies**. Click **Create new policy**.

Step 2 Configure these policy parameters:

a) Choose **SRv6 traffic accounting** as the **policy type**.

b) Enter a **policy name**.

c) Select **devices** and choose Cisco IOS-XR devices with gNMI enabled.

d) Click **Next**.

Step 3 **Set polling frequency** to define how often devices should be polled for this policy. By default data will be collected according to the polling frequency.

Step 4 Click **Activate** to apply the policy. The policy appears under **Performance Policies**.

Figure 41: SRv6 traffic accounting performance policy

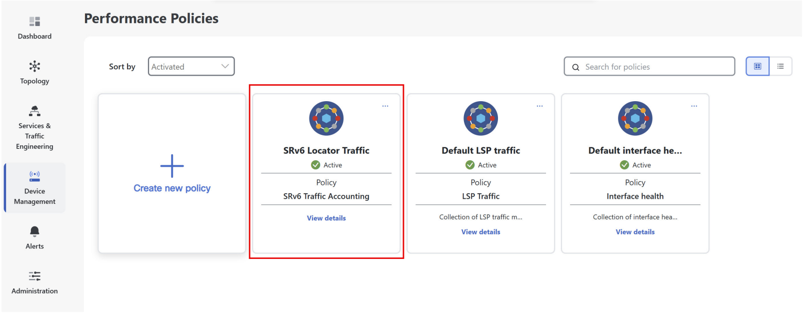

Configure device configurations

These configurations are required on your Cisco IOS-XR devices to enable SRv6 data collection and traffic steering for DDM integration.

Procedure

Step 1 **Enable SRv6 locator accounting**.

This configuration enables the router to perform detailed accounting for IPv6 traffic specifically related to SRv6 locators. By tracking traffic on a per-prefix and per-nexthop basis, operators gain granular visibility into the usage and flow of SRv6-enabled services.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config accounting
accounting
 prefixes
  ipv6
   mode per-prefix per-nexthop srv6-locators
  !
 !
!
```

Step 2 **Enable SRv6 accounting data to telemetry**

This configuration sets up model-driven telemetry on the router to stream SRv6 accounting data to external collectors. By defining specific sensor paths, the router can push operational data related to SRv6 locator accounting, enabling real-time monitoring, analysis, and orchestration of SRv6 network performance and traffic patterns.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config telemetry model-driven
telemetry model-driven
 sensor-group cisco_models

  sensor-path
Cisco-IOS-XR-fib-common-oper:cef-accounting/vrfs/vrf[vrf-name='default']/afis/afi[afi-type=ipv6]/pfx/srv6locs/srv6loc

 !
!
```

Step 3 **Enable customer/VRF traffic steering to SRv6 locators via BGP**

This configuration enables an edge router to steer customer or VRF (Virtual Routing and Forwarding) IPv4 and IPv6 traffic into specific SRv6 locators using BGP.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config router bgp
router bgp 60
 bgp router-id <ROUTER_ID_IP>
 segment-routing srv6
  locator L1algo0
 !
 address-family ipv4 unicast
  network <ROUTER_ID_IP>/32
 !
 address-family vpnv4 unicast
  vrf all                ! If there are multiple VRF where traffic is ingressing, add srv6 locator
in vrf all.
   segment-routing srv6
    locator L1algo0
    alloc mode per-vrf
   !
  !
 !
 vrf ntt
  rd 200:200
  address-family ipv4 unicast
   segment-routing srv6   ! If there is only one VRF where traffic is ingressing, add srv6 locator
in this vrf alone, if there is no VRF, then add the locator in neighbor address family
    locator L1algo0
    alloc mode per-vrf
   !
   redistribute connected
  !
  neighbor <NEIGHBOR_IP>
   remote-as 61
   update-source GigabitEthernet0/0/0/0
   address-family ipv4 unicast
    route-policy PASS_ALL in
    route-policy PASS_ALL out
   !
  !
 !
```

Step 4 **Verify SRv6 traffic steering via CEF accounting**

To verify that IPv6 traffic is being steered into SRv6 locators, rather than MPLS labels, use the `sh cef ipv6` accounting command on the device.

```
sh cef ipv6 accounting
fccc:cc3e:3::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
 via fe80::2/128, Bundle-Ether1201
  path-idx 0
```

```
next hop fe80::2/128
Accounting: 200000/58400000 packets/bytes output  <<< Traffic packets for prefix fccc:cc3e:3::
```

Important considerations when using DDM

Before using DDM, review these important considerations applicable for Crosswork Network Controller version 7.2.

- **SRv6 policy accuracy**: The demand matrix may be inaccurate in the presence of SRv6 policies. Full incorporation of these policies is planned for future phases.
- **Summarized locators**: There is no current support for redistributed or summarized locators across IGP domain boundaries. While demands for these locators are computed, they will appear as unrouted.
- **Locator leaking**: Locator leaking is assumed to occur at IS-IS level boundaries.
- **High availability (HA)**: DDM does not support HA. This means there is no built-in redundancy or failover capability, and data loss may occur during upgrades or unexpected restarts.
- **User permissions**: Users must have admin permissions to modify global configurations and perform enable/disable actions on domains. Read-only users can only view demand matrices.
- **Locator accounting coverage**: For the most accurate and complete SRv6 demand matrix, Locator counter collection should include every device in the IGP domain. See Locator counter collection coverage for the effects of partial coverage.
- **Locator counter collection synchronization**: Collecting locator counters from nodes at different times during the collection interval can result in transient false demands. See Locator Counter Collection Synchronization for details and mitigation steps.

Manage demand matrix domains

Use these dashboards to manage demand matrix domains and visualize the demands in detail by drilling down to specific locators and links:

- **Domain dashboard**: Acts as the central hub for managing and monitoring all IGP domains. It displays all discovered domains and provides an overview of each domain's status and key metadata.
- **Global configurations**: Allows you to define global settings that influence how DDM operates across all monitored IGP domains.
- **Demand matrix**: Provides a detailed view of the computed demands for a selected domain and its visualization on the topology map.

Demand matrix domains dashboard

The Demand matrix domain dashboard (**Services & Traffic Engineering** > **Demand Matrix**) displays all the domains discovered by Crosswork Network Controller. A domain is an identifier assigned to an IGP process.

Note The domains will begin displaying the updated info on active and total nodes, per the specified cadence set in the Global configuration settings. The default cadence is 15 minutes.

Figure 42: Demand matrix domains dashboard

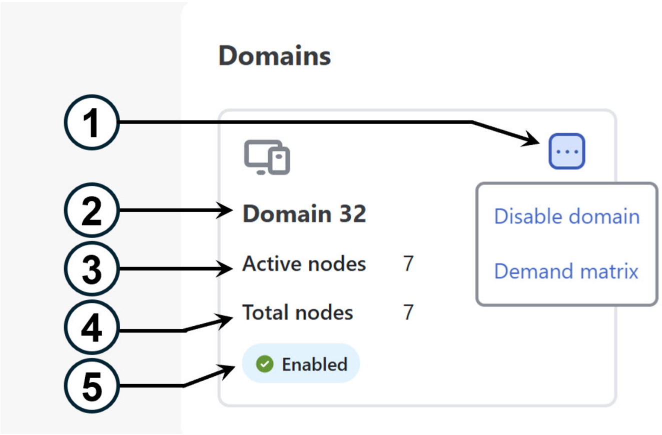

Callout No.	Description
1	• **Enable domain**: Enable the domain to start demand computation and visualization for that domain. By default, domains are in a disabled state. • **Disable domain**: Disable the domain to stop demand computation for that domain. You will not be able to see the demand matrix for this domain in a disabled state. • Demand matrix: View detailed traffic demand metrics for the domain on a topology map.
2	**Domain identifier**: The domain ID is taken from the router configuration (`link-state instance-id`) that is used to advertise IGP with BGP-LS.
3	**Active nodes**: Number of nodes in the domain that are actively sending SRv6 locator counter data to DDM.
4	**Total nodes**: Total nodes in the domain.
5	**Status**: Indicates whether DDM is enabled for the domain.

Configure global configurations

The Demand Matrix global configuration allows you to define system-wide settings that influence how the DDM operates across all monitored IGP domains. By adjusting these system-wide settings, you can influence the frequency of demand computations, filter out insignificant traffic, and manage debug logging as per your monitoring requirements. Configuration updates made here are also captured in the audit logs.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Demand matrix** and click on the **Global configuration** tab.

Figure 43: Demand Matrix global configuration

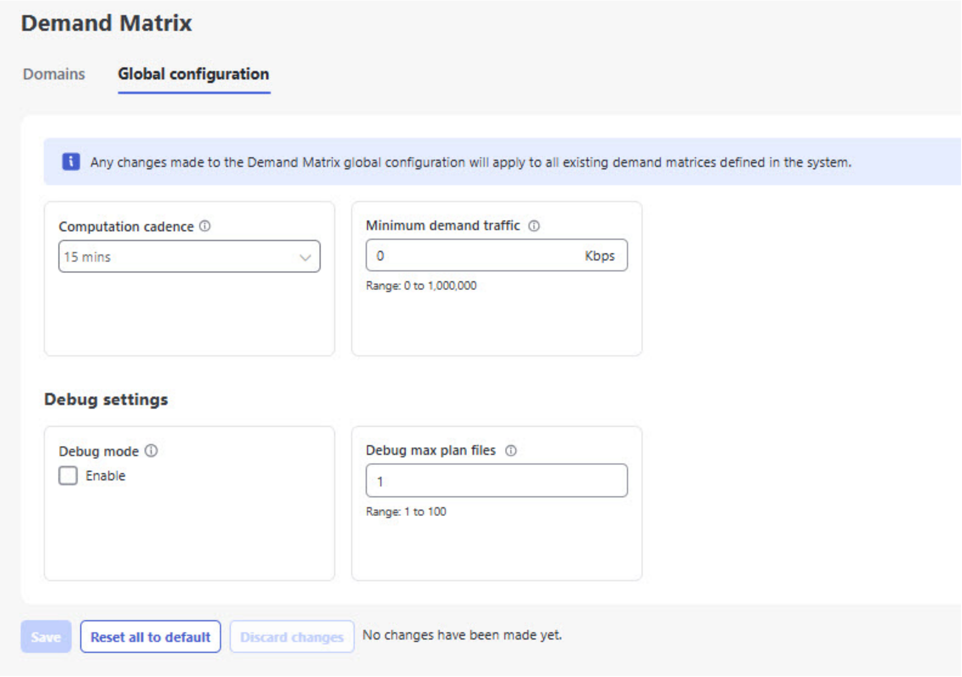

Demand Matrix

Domains | Global configuration

Any changes made to the Demand Matrix global configuration will apply to all existing demand matrices defined in the system.

Computation cadence: 15 mins

Minimum demand traffic: 0 Kbps
Range: 0 to 1,000,000

Debug settings

Debug mode
□ Enable

Debug max plan files: 1
Range: 1 to 100

No changes have been made yet.

Step 2 Adjust the required settings.

Note
Any changes made to the Demand Matrix global configuration will apply to all existing demand matrices defined in the system.

a) **Computation Cadence**: Select the frequency for demand matrix computation (15, 30, or 60 minutes). The default is 15 minutes. This cadence should be logically set to the same or a higher value than the polling frequency for the "SRv6 traffic accounting" performance policy.
b) **Minimum Demand Traffic**: Set a minimum traffic threshold (default is 0 Kbps). Demands with traffic below this value will be excluded from visualization, allowing focus on significant traffic flows.
c) **Debug mode**: Enable debug optimizer to log plan files to the Crosswork Network Controller file system. Files are saved to the maximum number of files you specify in **Debug max plan files**.
d) **Debug max plan files**: Set the maximum number of debug plan files you would like to save. The default is 1.

Step 3 Click **Save** to save your configuration.

View demand matrix

The Demand Matrix interface is your primary tool for monitoring SRv6 traffic demands, offering both a comprehensive overview of domain-wide traffic, the ability to drill down into specific demand paths, and to view demands routed through a particular link.

Note DDM operates strictly on a per-IGP domain basis. For traffic that crosses domain boundaries, it can accurately quantify the amount of traffic destined for an SRv6 locator in a neighboring domain. Within the current domain, you will see this traffic visualized as originating from the domain boundary router and extending to its destination within that domain. However, DDM cannot currently visualize the traffic's path beyond the current domain, meaning the full end-to-end path into the neighboring domain will not be displayed.

To view and monitor demand matrix, complete these steps:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Demand matrix** > ***Enabled Domain-ID*** > ··· > **Demand matrix**.

The Domain's demand matrix page is displayed with the different domain nodes in a topology map.

Figure 44: Demand matrix for a domain

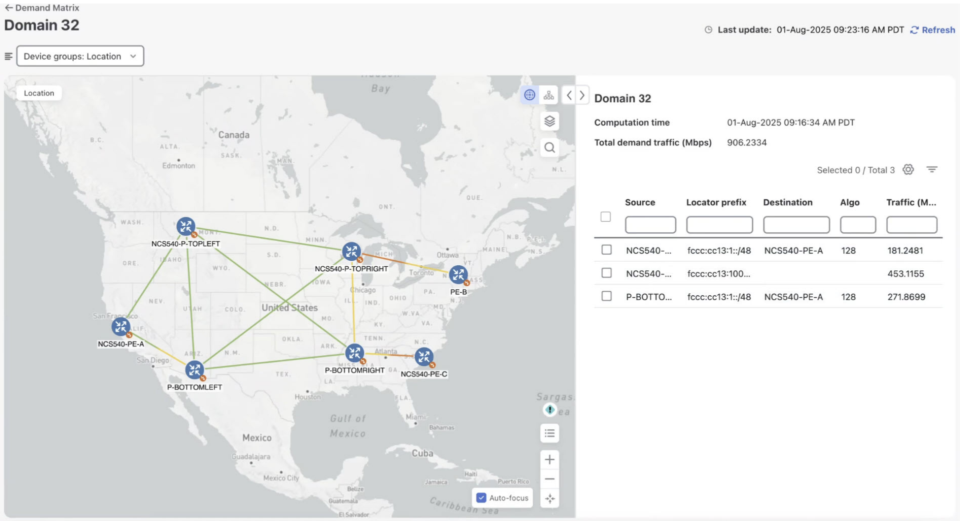

← Demand Matrix

Domain 32

Last update: 01-Aug-2025 09:23:16 AM PDT

Device groups: Location

Domain 32

Computation time 01-Aug-2025 09:16:34 AM PDT

Total demand traffic (Mbps) 906.2334

Selected 0 / Total 3

□	Source	Locator prefix	Destination	Algo	Traffic (M...
□	NCS540-...	fccc:cc13:1::/48	NCS540-PE-A	128	181.2481
□	NCS540-...	fccc:cc13:100...			453.1155
□	P-BOTTO...	fccc:cc13:1::/48	NCS540-PE-A	128	271.8699

On the right-hand side, you can view the computation time (reflecting the configured cadence) and the total traffic volume across all demands in the domain. A table lists all individual demands, detailing their source, destination, locator prefixes, and traffic volume.

Step 2 To visualize a specific demand's details, select it from the list. You can then see its IGP path from the source to the destination on the topology map, including details like the flexible algorithm used and traffic volume.

Figure 45: Visualize demand metric

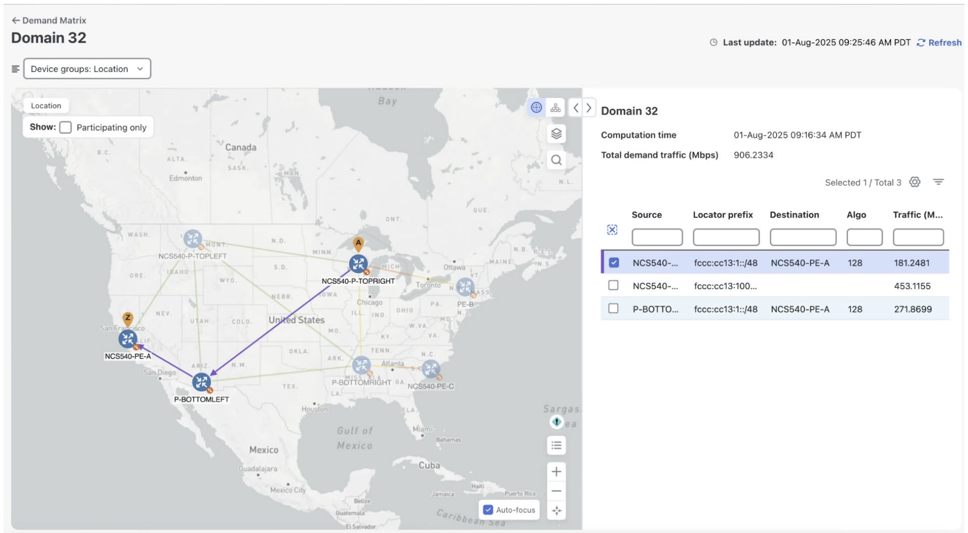

← Demand Matrix

Domain 32

Last update: 01-Aug-2025 09:25:46 AM PDT Refresh

Device groups: Location

Location

Show: □ Participating only

Domain 32

Computation time 01-Aug-2025 09:16:34 AM PDT

Total demand traffic (Mbps) 906.2334

Selected 1 / Total 3

	Source	Locator prefix	Destination	Algo	Traffic (M...
☑	NCS540-...	fccc:cc13:1::/48	NCS540-PE-A	128	181.2481
□	NCS540-...	fccc:cc13:100...			453.1155
□	P-BOTTO...	fccc:cc13:1::/48	NCS540-PE-A	128	271.8699

Step 3 You can filter metrics by Flex Algo. To visualize the flexible algorithm topology:

a) On the topology map, click and then click **Flex Algo**.

b) From the drop-down list, select the flexible algorithm ID (for example, 129) to filter.

c) (Optional) Select **Show selected Flex Algo topology only** to isolate the flexible algorithms on the topology map.

d) Click **Apply** to display the updated topology map.

Figure 46: Flex algo in demand matrix

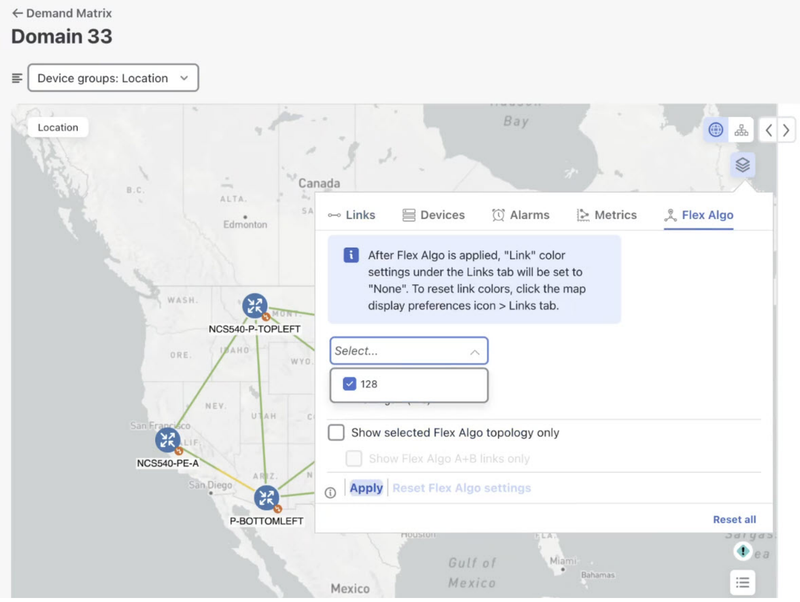

← Demand Matrix

Domain 33

Device groups: Location

Location

Links Devices Alarms Metrics Flex Algo

After Flex Algo is applied, "Link" color settings under the Links tab will be set to "None". To reset link colors, click the map display preferences icon > Links tab.

Select...

☑ 128

□ Show selected Flex Algo topology only

□ Show Flex Algo A+B links only

Apply Reset Flex Algo settings

Reset all

Step 4 To view metrics for the IGP path:

a) Select a demand metric from the table and click and then click **Metrics**.

b) Toggle applicable metrics to **On**.

Step 5 To view demands routed through a particular link:

a) On the topology map, click on a link.
b) From the **Links** page, select **View details** for the link. This displays the **Link details** page.
c) Click the **Traffic engineering** tab then click the **Demands** tab.

All demands routed through this link are displayed. Along with other link details, the table also shows the ECMP split percentage between these demands and the flex algos using that link.

Locator counter collection coverage

For the most accurate and complete SRv6 demand matrix, Locator counter (LOC.INT.E) collection should include every device in the IGP domain. This is the recommended deployment model for DDM. With partial coverage (for example, when some devices support Locator accounting and others do not), the demand matrix may be impacted in several ways, including:

- false demands
- missing demands
- double counting of traffic

In general, with partial LOC.INT.E coverage, false demands can arise at any node serving as a transit for traffic to a Locator (for example, when the node has non-zero LOC.INT.E counters for transiting traffic), if one or more neighboring nodes do not support LOC.INT.E counters and, as a result, the ingress counters from those neighbors are not considered.

Locator counter collection synchronization

Locator counters are collected from each device via gNMI streaming, with each device sending updates periodically. Because these updates represent different snapshots in time, the data across devices may not be perfectly synchronized within the collection interval. As a result, transient false demands with relatively low rates may appear on different nodes, especially when traffic is dynamic over short periods.

To address this, you can set a Minimum Demand Traffic threshold. Any detected demands below this value will be filtered out, reducing the impact of such false demands. However, use this setting carefully: if your network includes legitimate, low-rate demands that are important, setting the threshold too high could also filter out real demands.