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
- Flexible algorithm

- Tree Segment Identifier
- 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.

- **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.

```
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.

```
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.

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

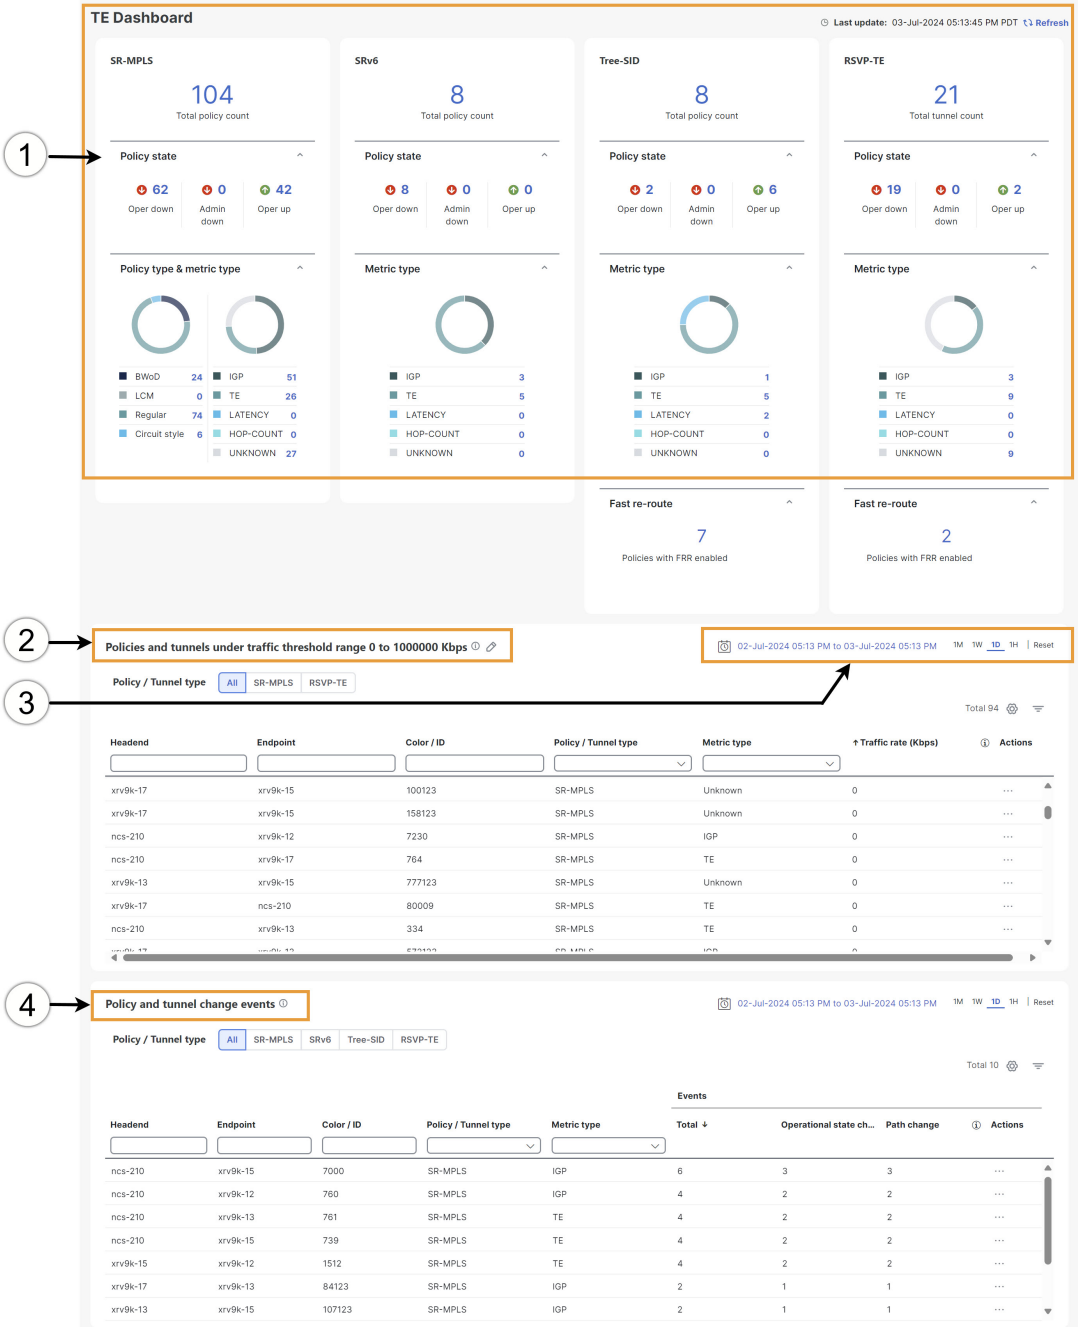

Note If you are viewing the HTML version of this guide, click the images to view them in full-size.

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

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

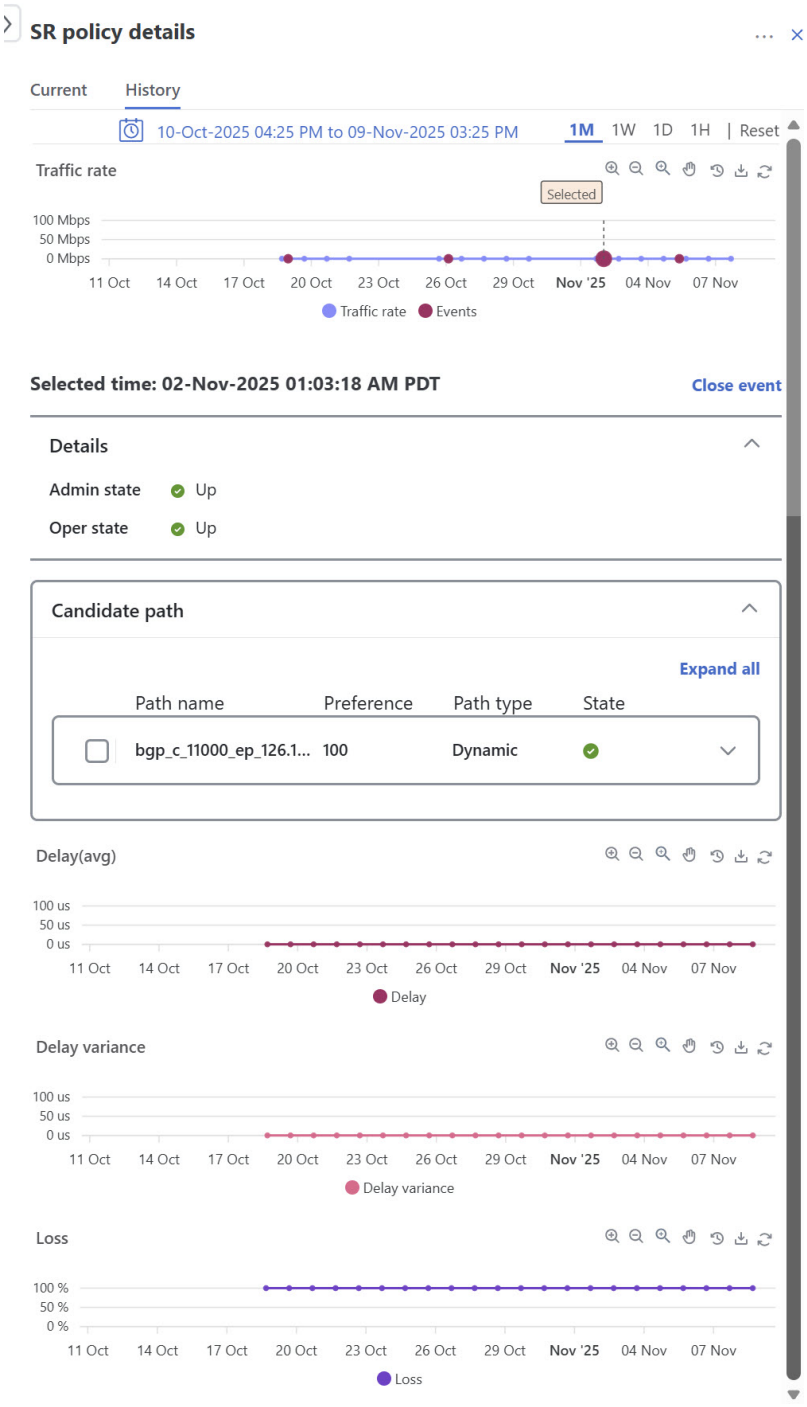

View TE device details

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

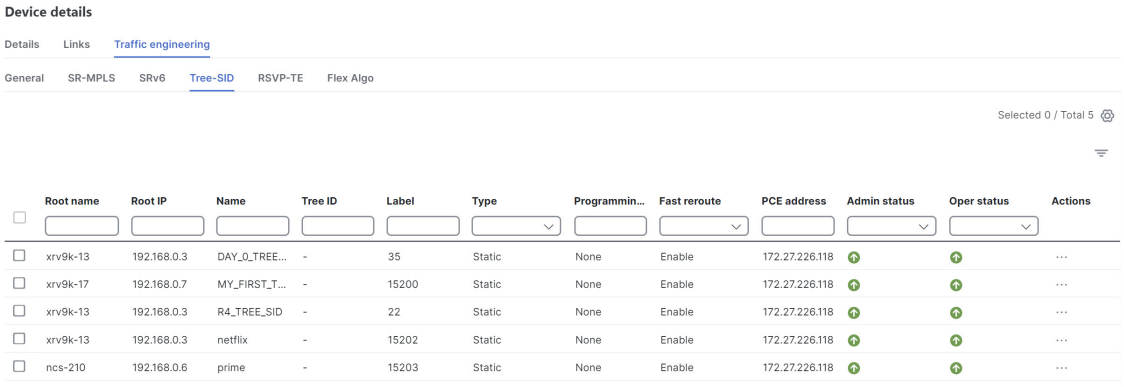

Note
If you are viewing the HTML version of this guide, click the image to view it in full-size.

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

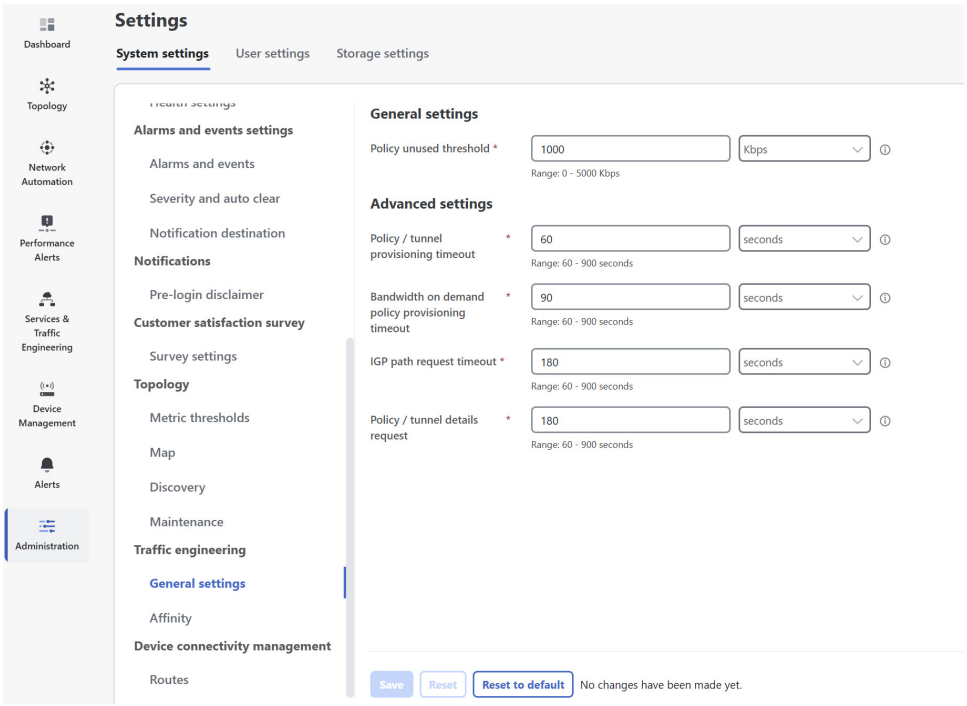

Configure device group display for TE

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**).