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 1: Traffic engineering UI: SR-MPLS and SRv6 policies

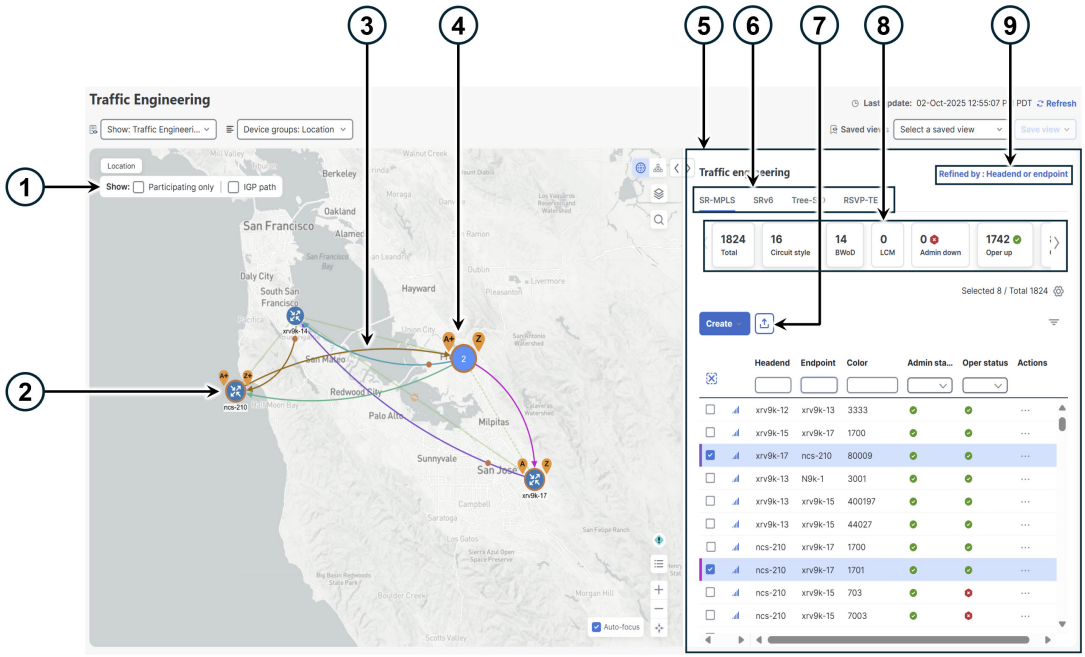

The table describes the callouts for the Traffic Engineering topology map for SR-MPLS and SRv6 policies.

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 20 • Create dynamic SR-MPLS policies based on optimization intent, on page 21 • Modify SR-MPLS policies, on page 24
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 2: View SR policy details

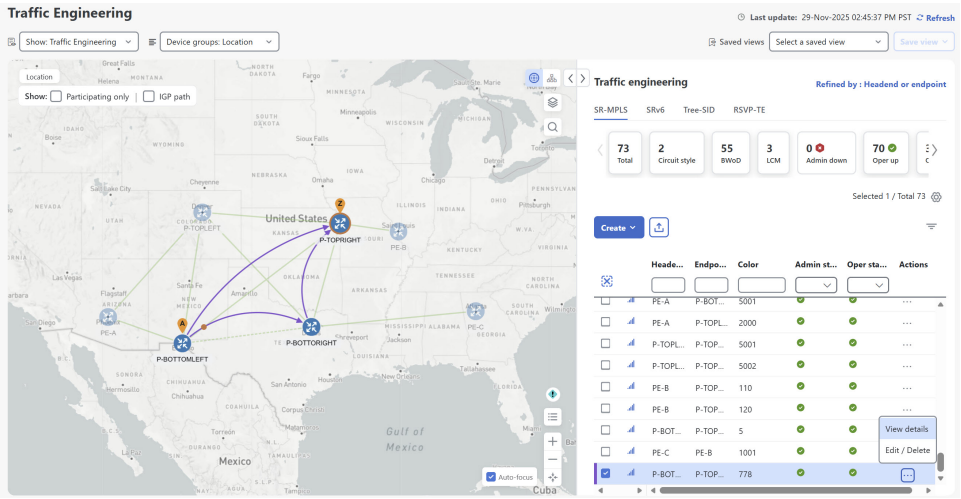

Step 2 A list of candidate paths appear along with policy details in the **SR policy details** page.

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

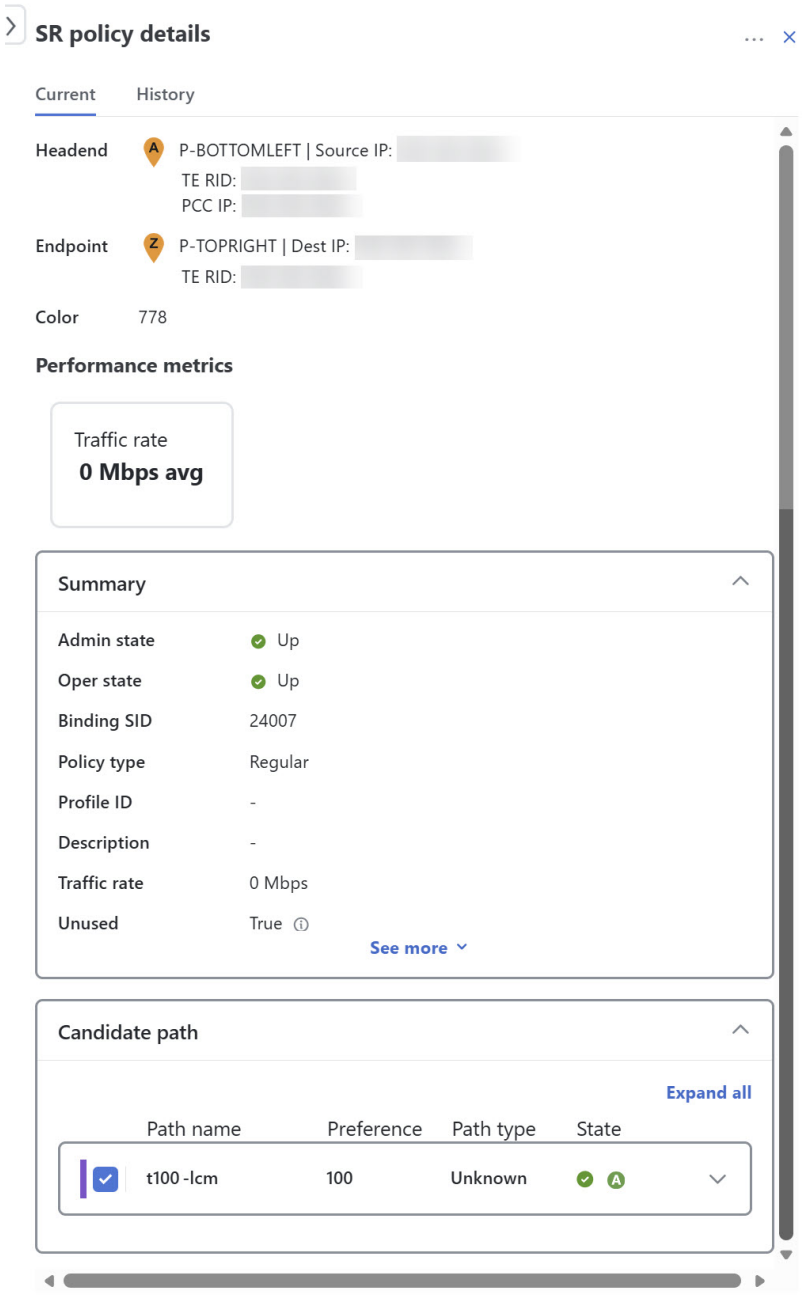

Step 3 In the candidate path area, click **Expand all** to view additional details on the different paths and segments.

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 4: SR policy details with multiple segment list

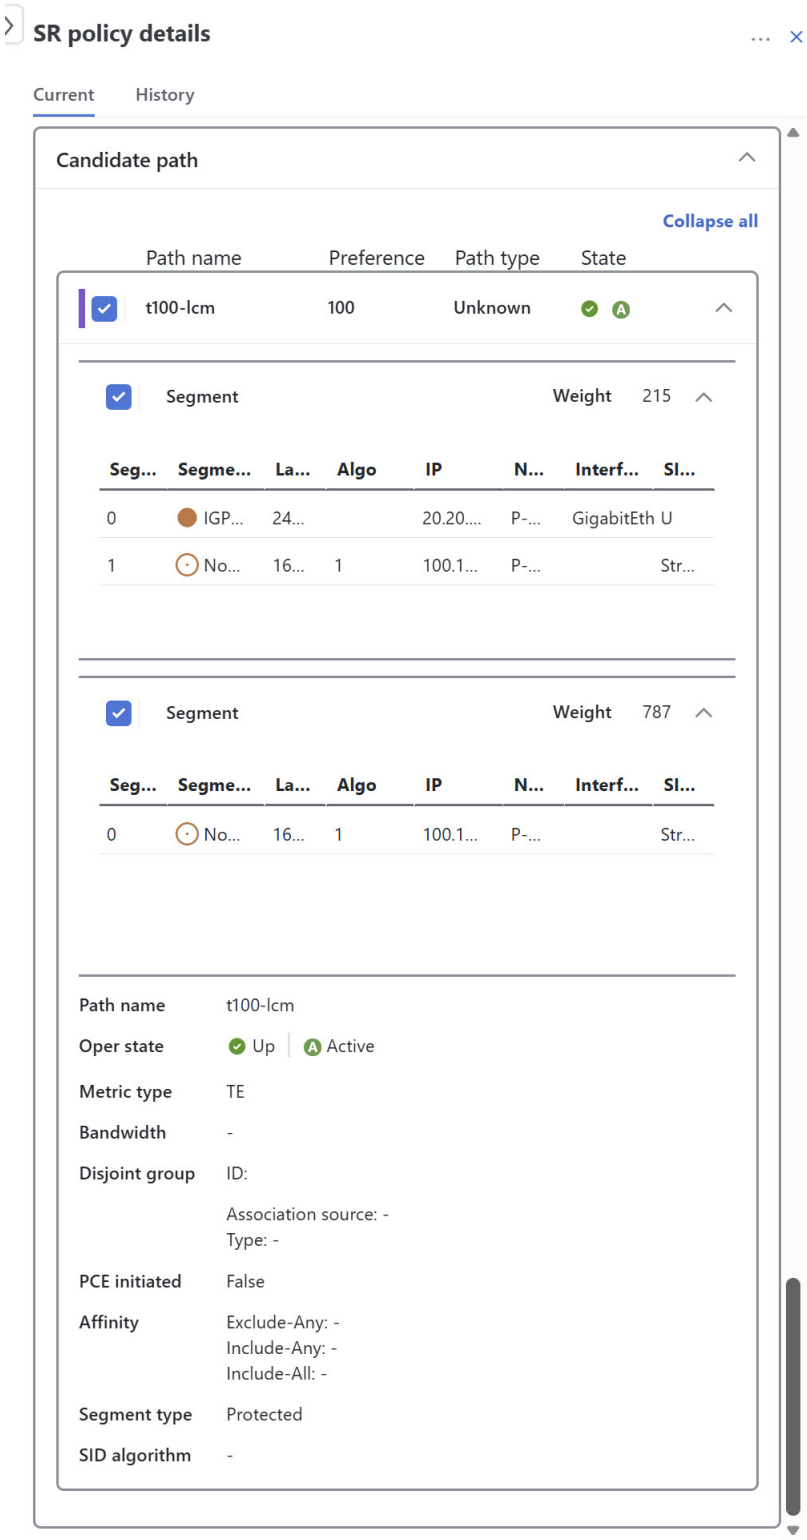

View IGP path and metrics

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 5: View physical path and metrics

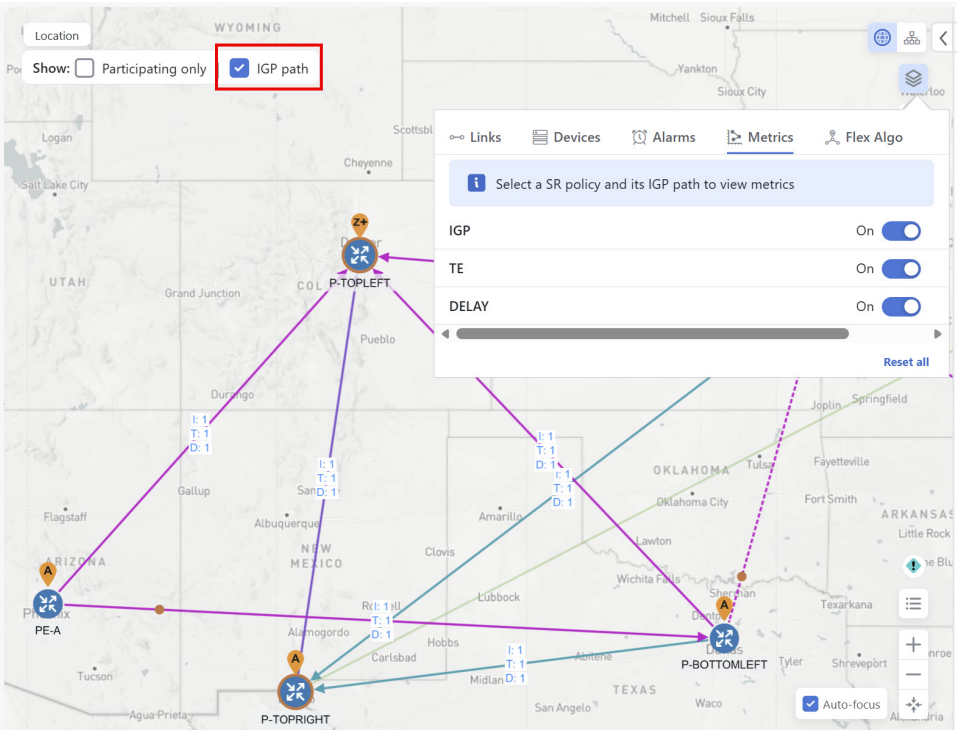

Find Multiple Candidate Paths (MCPs)

Visualizing MCPs gives you insight into which paths might be a better alternative to the currently active 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 6: SR-TE policy on the topology map

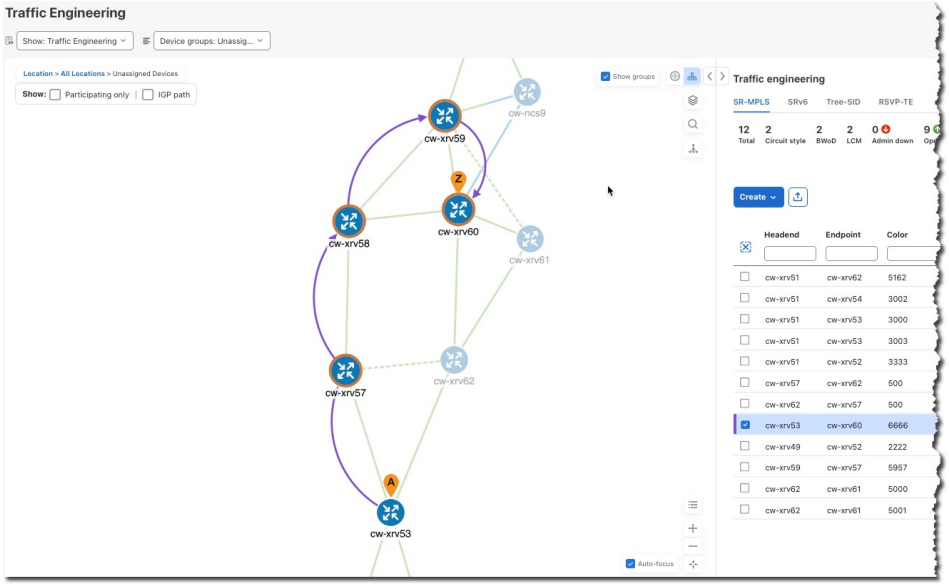

Step 3 View the list of candidate paths.

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 7: Candidate path in SR policy details

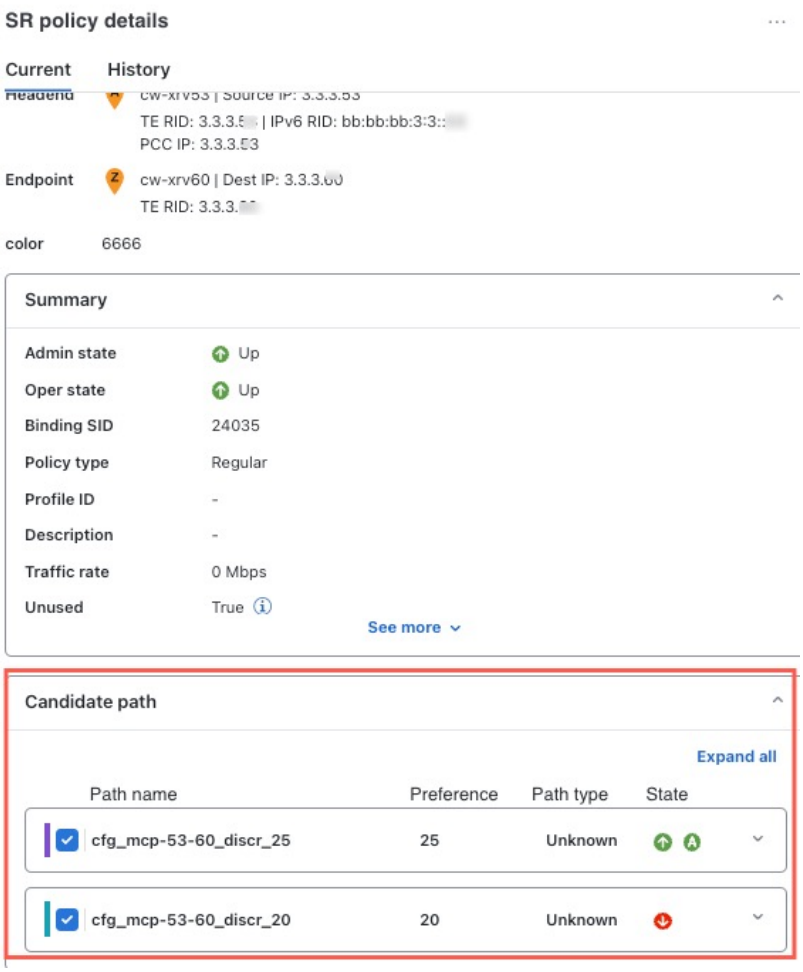

Step 4 You can expand individual paths or click **Expand all** to view details of each path.

Step 5 Visualize the candidate path on the topology map.

a) 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 8: Candidate path on the topology map

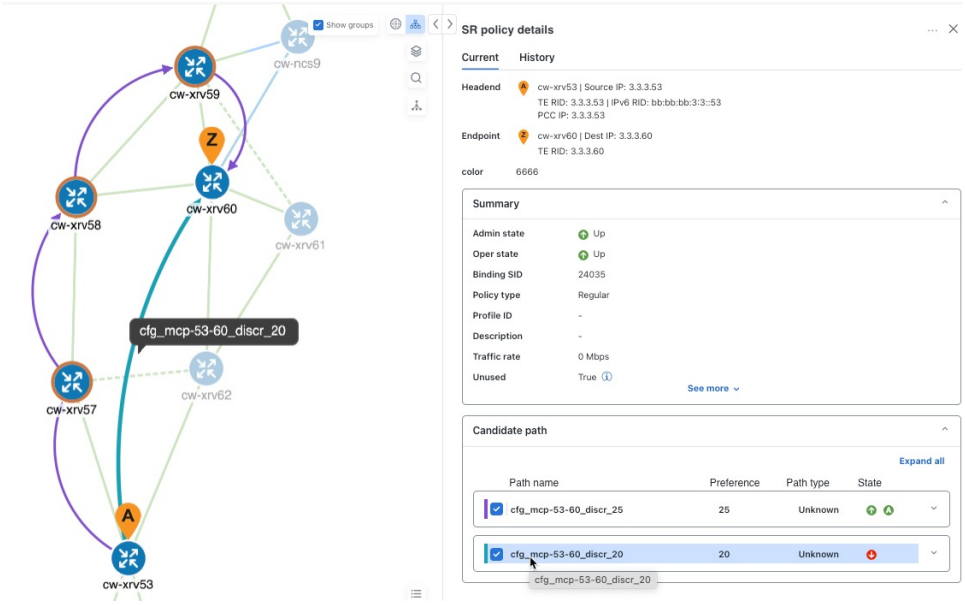

View underlying paths associated with a Binding-Segment ID (B-SID) label

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 9: B-SID label

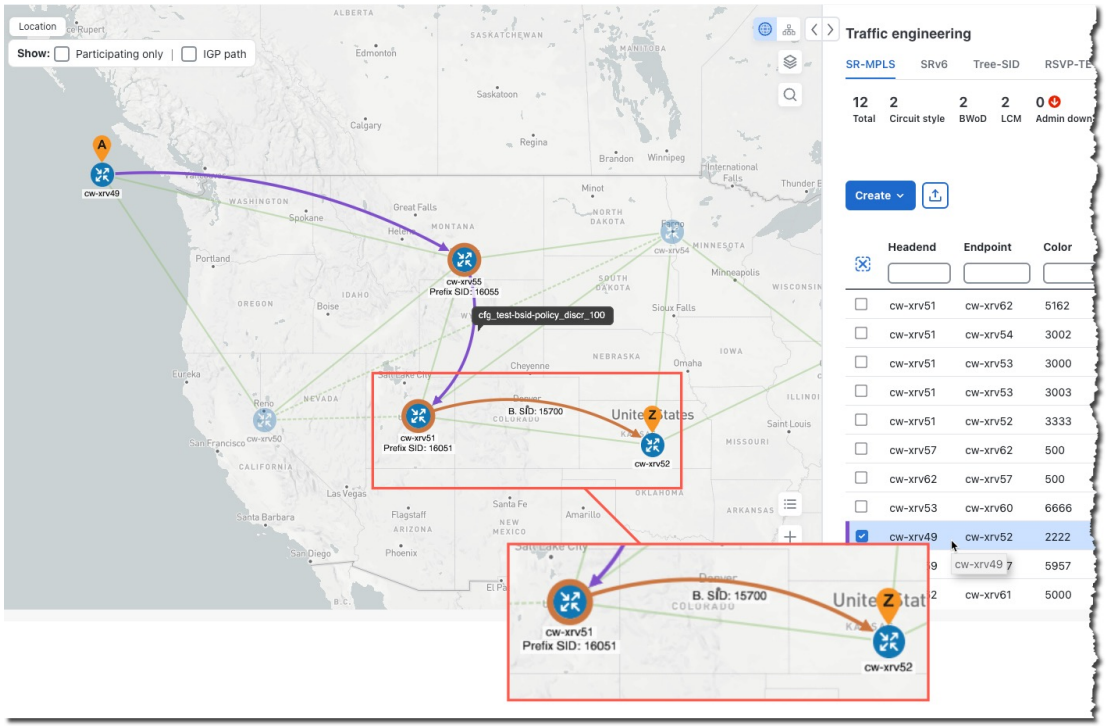

Step 3 In the **SR policy details** page, click ⋯ > **View details**.

Figure 10: View details

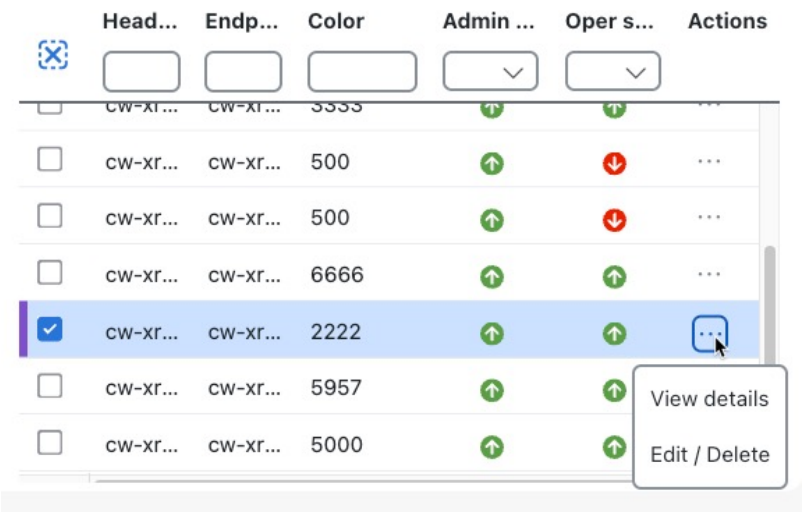

Step 4 Expand the active path and click the B-Sid Label ID to view the underlying path.

Figure 11: B-Sid label ID

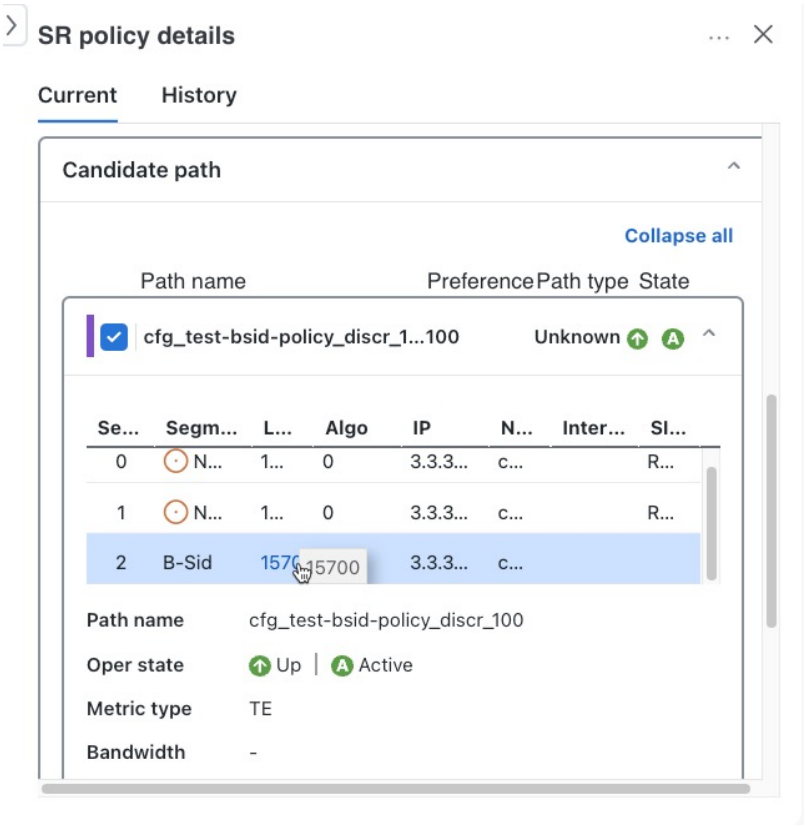

In this example, the underlying path actually goes from **cw-xrv51** > **cw-xrv54** > **cw-xrv53** > **cw-xrv52**.

Figure 12: B-SID path

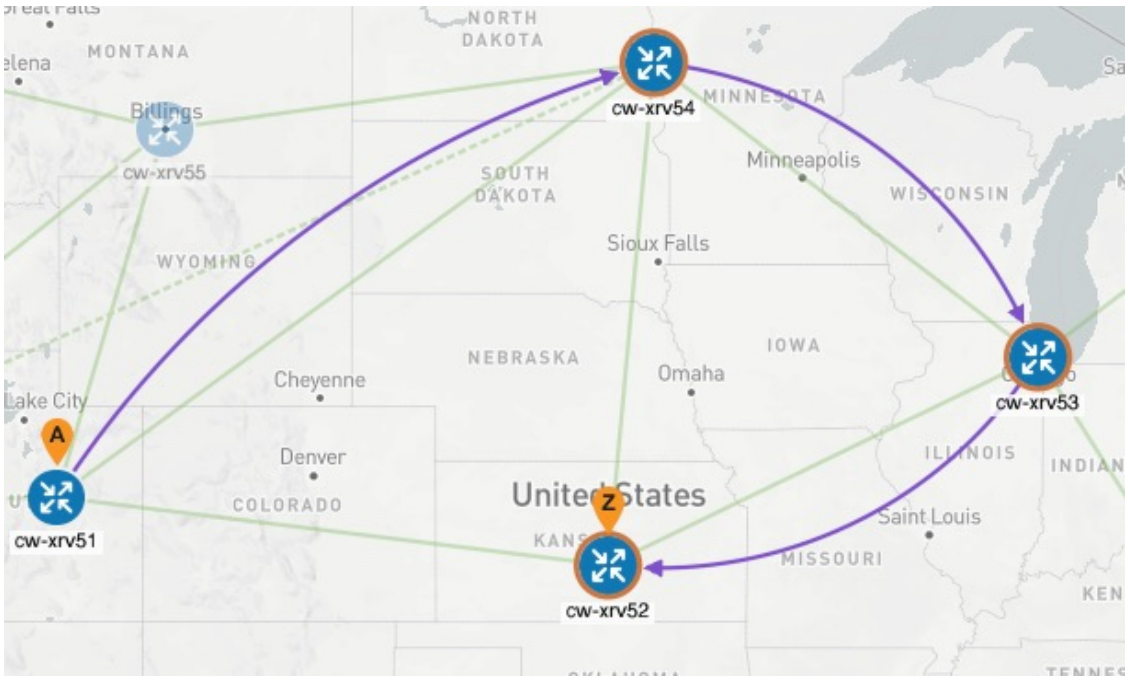

SR policies with multiple segment lists

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 13: Single SR policy with multiple segment lists

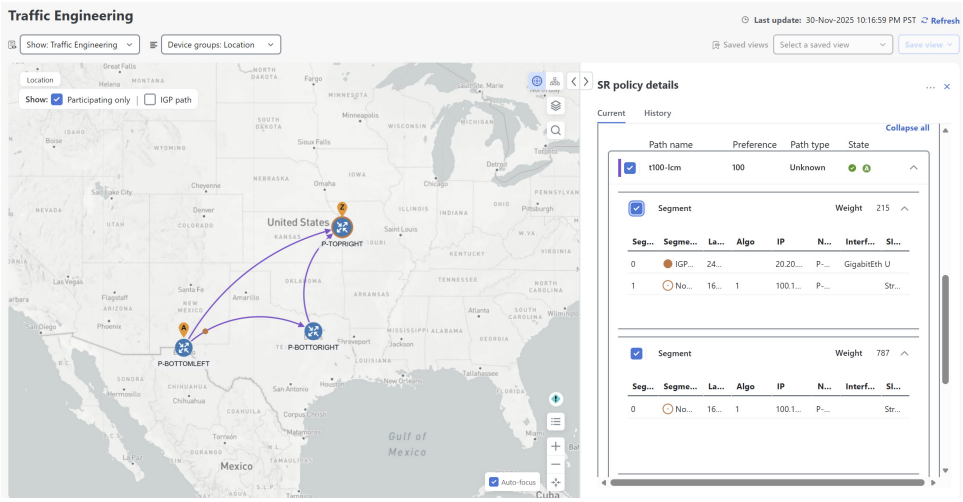

Native SR paths

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 14: New path query

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 15: Path details

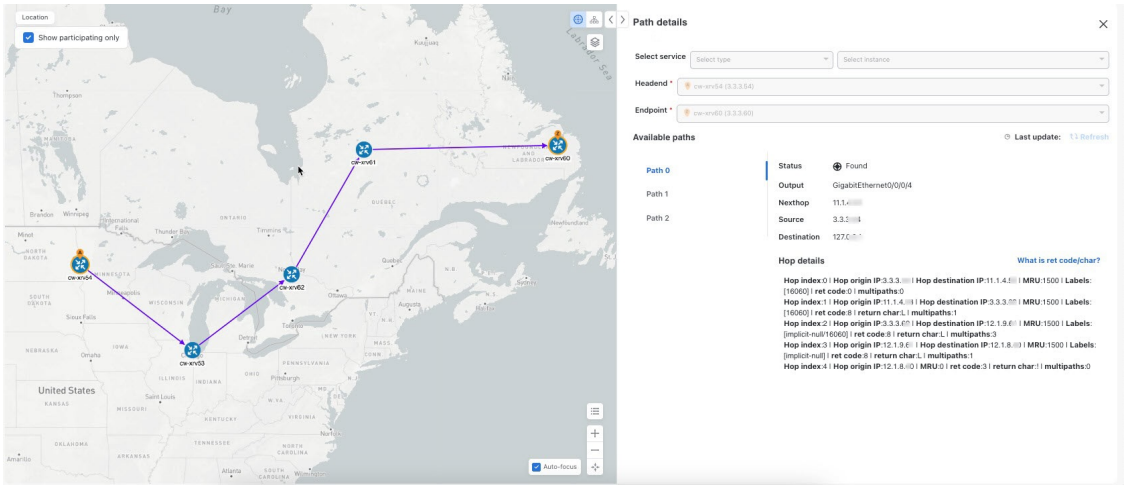

Configure TE link affinities

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 16: Mapping affinities

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.
- 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 for sample configurations and Configure and visualize flexible algorithm SRLG exclusion 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 23.

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.

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 23.

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 1: 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.

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**.