Cisco Crosswork Network Controller 7.2 Network Bandwidth Management

First Published: 2026-01-30

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

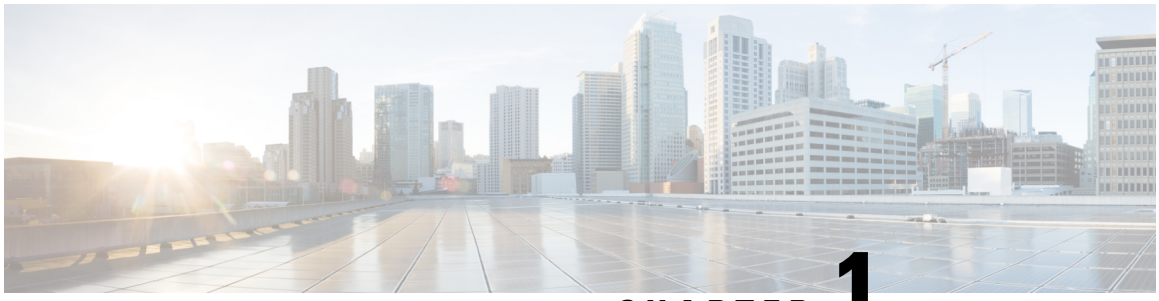

CHAPTER 1

Network Bandwidth Management

Network bandwidth management feature packs

Service providers face increasing pressure to deliver reliable, high-quality network services to their customers. Traditionally, bandwidth management was a manual and reactive process that made it difficult to respond quickly to emerging issues.

One of the most significant challenges is network congestion. Congested links, high latency, and other impairments negatively impact service quality, leading to poor end-customer experiences. Persistent network problems can make it difficult to meet service level agreements (SLAs), and in the worst cases, may result in SLA violations, contract breaches, and damage to brand reputation.

To overcome these challenges, network operators require automated tools that optimize bandwidth, reduce manual intervention, and ensure critical links always have sufficient capacity to avoid congestion. The Crosswork Network Controller addresses these needs with a suite of feature packs designed to streamline bandwidth management and traffic engineering.

- **Local Congestion Mitigation (LCM)** : A tactical solution for bandwidth management and congestion mitigation. LCM is ideal for directly addressing congestion issues on devices, without requiring a comprehensive traffic matrix or advanced planning.
- **SR Circuit-Style Manager (CSM)** : A strategic traffic engineering solution that enables you to reserve bandwidth in advance for critical services, avoiding congestion issues entirely for these high-priority services.
- **Bandwidth on Demand (BWoD)** : A solution that offers soft bandwidth guarantee services for SR policies, in contrast to the strict bandwidth guarantees provided by Circuit Style SR-TE services. Depending on the configuration, BWoD can either provide bandwidth reservation, or best-effort bandwidth paths for SR policies.

Note CSM and BWoD feature packs are mutually exclusive. Only one can be enabled at a time.

Feature pack requirements

- Ensure you have the correct licensing package to use feature packs.
- Users must be assigned administrator roles or specific Device Access Group permissions to access certain features or configurations. For more information on RBAC and user roles, see the Cisco Crosswork Network Controller Administration Guide .

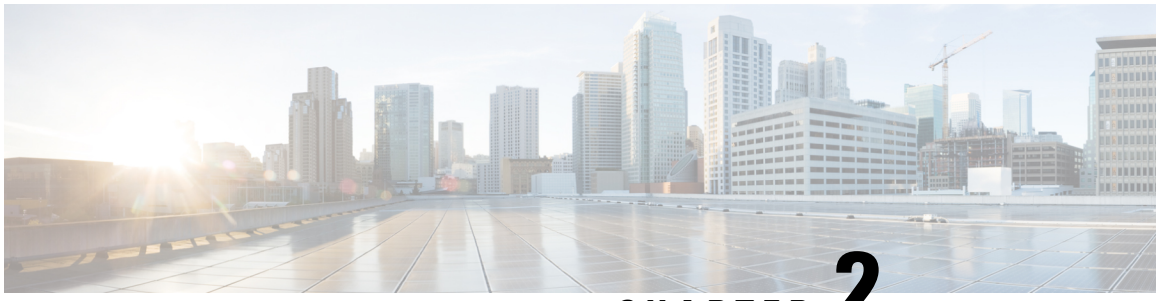

CHAPTER 2

SR Circuit-Style Manager (CSM)

Circuit-style manager

Circuit-style manager provides bandwidth-aware path computation and management for circuit-style SR-TE policy paths in large-scale networks. It is a network management tool that:

- performs centralized bookkeeping to track and allocate bandwidth resources across the network,
- computes policy paths that meet committed bandwidth requirements and service-level constraints,
- enables users to monitor bandwidth resource levels and identify areas where resources are running low, and
- manages and visualizes circuit-style SR-TE policies on the network topology map, ensuring proper failover, protection, and bidirectional operation.

Advantages of using circuit-style SR-TE policies

Key advantages of circuit-style SR-TE policies include:

- Ensuring reliable, bidirectional transport for high-priority and critical services, with guaranteed bandwidth and protected paths.
- Simplifying network operations with centralized bandwidth reservation and policy management without requiring extra protocols.
- Enabling rapid service recovery through automatically computed working, protect, and restore paths.

- Maintaining service-level agreements (SLAs) even as network loads change.
- Not requiring maintenance of network states at intermediate routers.
- Providing clear visibility and control of bandwidth resource allocation across the network.

Important considerations for circuit-style SR-TE policies

Circuit-Style SR-TE is subject to specific configuration requirements and operational constraints, including policy attribute compatibility.

Access requirements for circuit-style SR-TE policy provisioning

To provision a circuit-style SR-TE policy, you must have write-access to the head-end device based on Device Access Groups and assigned roles. Only circuit-style SR-TE admin users can modify circuit-style SR-TE configuration settings. For more information on Role-based Access Control (RBAC) and task permissions, see Cisco Crosswork Network Controller Administration Guide.

Attribute constraints

You set policy attribute values when you create a circuit-style SR-TE policy using either the device's command line interface (CLI) or through Crosswork Network Controller UI provisioning using Network Services Orchestration (NSO). To view a device CLI configuration example, see Configure circuit-style SR policies, on page 11.

This table outlines the requirements for each policy attribute and how changes affect them. All attributes listed function as constraints. Each attribute aligns with configuration elements that Crosswork Network Controller utilizes to manage the computation of circuit-style path hops. Each value serves as a constraint for path computation or optimization, either defining a necessary path property or eliminating potential path options.

Table 1: Attribute constraints

Attribute	Description
Policy path protection	The path protection constraint is required for both sides of a circuit-style SR-TE policy.

Attribute	Description

Bandwidth constraint	• The bandwidth constraint is required and must be the same on both sides of a circuit-style SR-TE policy. Changes to bandwidth can be applied to existing policies with the following outcomes:
	  • After configuring the new bandwidth on both sides, the system evaluates the path *without* recomputing it.
	  • If the new bandwidth is higher, the system checks the current path for sufficient resources. If all paths can support the new bandwidth, the same path is returned with the updated bandwidth value, indicating to the path computation client (PCC) that it was successful. If any path cannot support the new bandwidth, the old bandwidth value is returned, indicating failure. This evaluation is only retried if the bandwidth changes again.
	  • If the bandwidth is lower, the system returns the same path with the new bandwidth value to indicate to the PCC that it was successful.
	• When you view the policy details, the user interface shows both the requested and reserved bandwidth under each candidate path. These values can differ if the requested bandwidth is increased but there is insufficient available circuit-style pool bandwidth along one or more paths.

Candidate paths and roles	• The Working path is defined as the highest preference Candidate Path (CP). • The Protect path is defined as the CP of the second highest preference. • The Restore path is defined as the lowest preference CP. The headend must have `backup-ineligible` configured. • CPs of the same role in each direction must have the same CP preference.
Bi-directional paths	• All paths must be configured as co-routed. • Paths of the same role on both sides must have the same globally unique bi-directional association ID.

Attribute	Description

Disjointness	The disjoint policy is used to compute two lists of segments that steer traffic from two source nodes to two destination nodes along disjoint paths. The disjoint type refers to the resources the two computed paths should not share.
	• The supported disjoint path types are:
	  • **Link**: Links are not shared on the computed paths.
	  • **Node**: Nodes are not shared on the computed paths.
	  • **SRLG**: Links with the same Share Risk Link Group (SRLG) value are not shared on the computed paths. These links rely on a common resource, making them susceptible to the same potential failures. This setting specifies that the Working and Protect paths cannot use links that are part of the same SRLG.
	  • **SRLG-node**: SRLG and nodes are not shared on the computed paths.
	• The disjoint type used must be the same in both directions of the same policy.
	• Working and Protect paths on the same PCC must be configured with a disjointness constraint using the same disjoint association ID and disjointness type.
	• The disjointness association ID for a Working and Protect path pair in one direction must be unique when compared with the corresponding pair in the opposite direction.
	• The Restore path must not have a disjointness constraint set.
	• Crosswork Network Controller follows strict fallback behavior for all Working and Protect path disjointness computations. This means that if node type disjointness is configured but no path is available, the system makes no automatic attempt to compute a less restrictive link type disjoint path.

Metric type	Only the TE, IGP, hop count, and latency metric types are supported. The metric type must match Working, Protect and Restore paths in both directions.

Segment constraints	• All Working, Protect, and Restore paths must have the following segment constraints:
	  • protection unprotected-only
	  • adjacency-sid-only
	• To ensure persistency through link failures, configure static adjacency SIDs on all interfaces that might be used by circuit-style SR-TE policies.

Unsupported configurations

These configurations are not supported:

- Metric bounds
- SID-Algo contraints
- Partial recovery for devices running IOS XR 7.8.x
- Multiple circuit style SR-TE policies between the same nodes with the same color but different endpoint IP addresses
- State-sync configuration between PCEs of a high availability pair

Path computation and reversion behaviors

The successful provisioning and continuity of circuit-style SR-TE policies depend on both the process of path computation and the precise handling of recovery and restoration scenarios. Path computation determines how the system establishes candidate paths that meet bandwidth, protection, and constraint requirements. Path reversion logic governs the transition between working, protect, and restore paths in response to network events and recoveries.

Path computation behavior

The SR Circuit-Style Manager (CSM) computes paths for circuit style policies only after a complete bi-directional, path-protected set of candidate paths has been delegated, including Working and Protect paths on both sides.

- **Bandwidth availability and path delegation:** Path computation relies on bandwidth availability. If insufficient bandwidth prevents path establishment, the SR Circuit-Style Manager will retry every 30 minutes until a solution is found or circuit style SR-TE is disabled.
- **Restore path computation:** The Restore path is computed only after the Working and Protect paths go down. Use the configurable delay timer to set the wait period post-delegation, allowing topology and policy state changes to propagate before computation.
- **Path optimization and limitations:** Automatic path re-optimization is unavailable for topology or LSP state changes and periodic events. Path configurations must be manually adjusted as needed.
- **Supported path computation scenarios:** Path computation supports Intra/Inter-area and Intra/Inter IGP Domain scenarios. Inter-AS path computation is not supported, requiring manual configuration for such cases.

Path reversion

Reversion behavior

Reversion behavior is controlled by the configuration of the WTR lock timer option under the Protect and Revert paths (it is not relevant for the Working path):

- No lock configuration: Revert after a default 5-minute lock
- Lock with no duration specified: No reversion

- Lock duration: Revert after the specified number of seconds

Reversion logic

Path reversion depends on the initial state of the Working, Protect, and Restore paths and the events affecting each path. The scenarios in the following table provide examples of typical reversion behavior.

Table 2: Path reversion scenarios

Initial State	Events	Behavior
Working path is down, Protect path is up/active	Working path comes back up	**1.** Working path recovers to up/standby state. **2.** Each PCC moves the Working path to active after the WTR timer expires. **3.** Protect path moves to up/standby.
Working path is down, Protect path is down, Restore path is up/active	Working path comes back up, then Protect path comes back up	**1.** Working path recovers and goes to up/active state **2.** Restore path is removed **3.** Protect path recovers and goes to up/standby
Working path is down, Protect path is down, Restore path is up/active	Protect path comes back up, then Working path comes back up	On side A: The Working path failure is local (the first Adj SID in the SegList is **invalid**): **1.** Protect path recovers and goes to up/active. **2.** Restore path is removed. **3.** Working path recovers and goes to up/standby. **4.** Each PCC moves the Working path to active after the WTR timer expires, Protect path goes to up/standby. On side Z: Working path failure is remote (first Adj SID in SegList is **valid**): **1.** Protect path recovers but is not brought up, Restore path remains up/active. **2.** Working path recovers and goes up/active. **3.** Restore path is removed. **4.** Protect path goes to up/standby.

Set up the CS SR-TE policy visualization workflow

Complete these steps to ensure CS SR-TE policies appear with correct bandwidth details:

Procedure

Step 1 Enable SR CSM, on page 9.

Step 2 Configure circuit-style SR policies, on page 11 on the device.

Step 3 Verify that the CS SR-TE policies appear in the **Traffic Engineering** table.

Choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** > **Circuit-style**.

Step 4 From the topology map, click a participating CS SR-TE node and verify that the reserved bandwidth pool settings you defined in the first step are configured properly.

Choose **Links** > ***link-type-entry*** > **Traffic engineering** > **General**. See Review circuit-style SR-TE policy bandwidth utilization , on page 13.

Enable SR CSM

To manage and visualize circuit-style SR-TE policies on the topology map, you must first enable the SR-CSM and set bandwidth settings. When enabled, the CSM computes the best failover bidirectional paths with the requested bandwidth and other constraints defined in the circuit-style SR policy configuration between two nodes.

Complete these steps to enable SR Circuit-Style Manager.

Before you begin

Both CSM and BWoD cannot be enabled at the same time. If BWoD is enabled, you must disable it before enabling CSM.

Note It is also recommended to delete the respective policies from the network before disabling the associated function pack (BWoD or CSM). If policies remain in the disabled function pack, it may cause issues with new policy delegation and increase processing time.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Circuit Style SR-TE** > **Configuration** > **Basic**.

Step 2 Toggle the **Enable** switch to **True**.

Figure 1: Basic Circuit Style SR-TE configuration

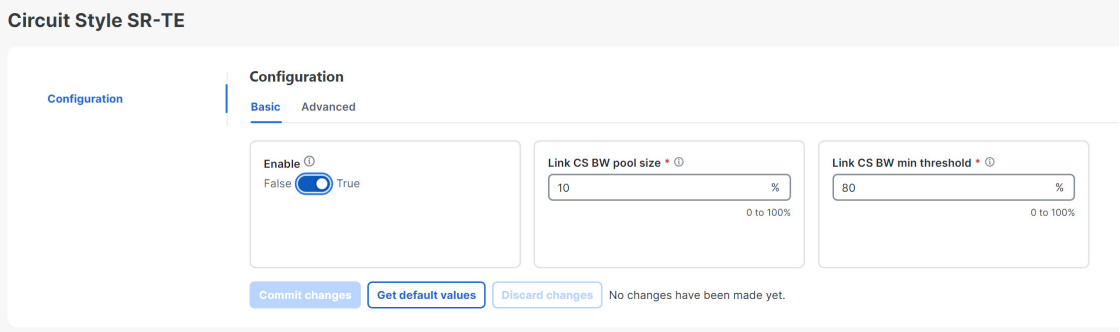

Step 3 Enter the required bandwidth pool size and threshold information. Additional field information is listed in this table. See Effects of surpassing bandwidth reservation limits, on page 19.

Field	Description
Link CS BW pool size	The percentage of each link's bandwidth reservable for circuit-style SR-TE policies.
Link CS BW min threshold	The Link CS BW Pool utilization percentage beyond which Crosswork Network Controller will generate a threshold-crossing event notification.

Step 4 Click **Commit changes** to save the Basic configuration.

Step 5 Click the **Advanced** tab to display additional CS-SR configuration values. Additional field information is listed in this table.

Figure 2: Circuit style SR-TE configuration - advanced tab

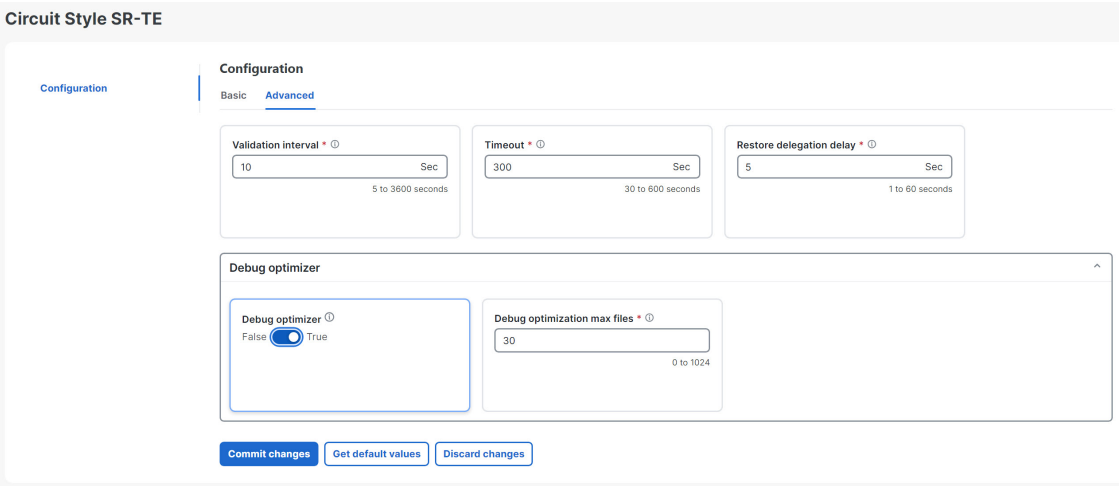

Field	Description
Validation interval	This is the interval that the CSM policy will wait before the bandwidth that is reserved for an undelegated policy is returned to the circuit-style SR-TE policy bandwidth pool.

Field	Description
Timeout	The duration until which the CSM will wait for the delegation request, before generating a threshold-crossing alarm.
Restore delegation delay	The duration until which the CSM will pause before processing a restore path delegation.
Debug optimizer	Toggle the switch to **True** to turn on the Debug Optimizer for all CS-SR policies. The Debug Optimizer will write log files to the Crosswork Network Controller file system whenever it calculates routes up to the maximum number of files you specify.
Debug optimization max files	Enter the maximum number of log files the Debug Optimizer will write. Once the maximum is reached, the Optimizer will overwrite existing files.

Step 6 When you are finished entering Advanced configuration values, click **Commit changes** to save the configuration.

What to do next

Configure circuit-style SR policy configurations either manually on the device (see Configure circuit-style SR policies) or through Crosswork Network Controller.

Configure circuit-style SR policies

A circuit-style SR policy configuration must include the destination endpoint, the amount of requested bandwidth, and the bidirectional attribute (see Important considerations for circuit-style SR-TE policies, on page 4 for additional requirements or notable constraints). The configuration should also include a Performance Measurement Liveness (PM) profile. A PM profile enables proper detection of candidate path liveness and effective path protection. PCCs do not validate past the first SID, so without PM, the path protection will not occur if the failure in the circuit-style SR policy candidate path is not the first hop in the segment list. For more information, see Configuring SR Policy Liveness Monitoring.

Procedure

Step 1 If applicable, enable the hardware module on the device for PM configuration.

Example:

```
hw-module profile offload 4
reload location all
```

Step 2 Configure the PM profile.

Example:

```
performance-measurement
  liveness-profile sr-policy name CS-active-path
    probe
      tx-interval 3300
    !
    npu-offload enable   !! Required for hardware Offload only
    !
  !
```

```
  liveness-profile sr-policy name CS-protect-path
    probe
      tx-interval 3300
    !
    npu-offload enable   !! Required for hardware Offload only
    !
  !
!
```

Step 3 Configure the Circuit Style SR policy with the PM profile. All configurations shown in the example are required in order for CSM to manage the circuit-style SR-TE policy. Entries that are defined by the user are italicized. See Important considerations for circuit-style SR-TE policies, on page 4 for additional requirements or notable constraints.

Example:

```
segment-routing
  traffic-eng
    policy cs1-cs4
      performance-measurement
        liveness-detection
          liveness-profile backup name CS-protect  !! Name must match liveness profile defined for
Protect path
          liveness-profile name CS-active  !! Name must match liveness profile defined for Active
path
      !
      !
      bandwidth 10000
      color 1000
      end-point ipv4 192.168.20.4
      path-protection
      !
      candidate-paths
        preference 10
          dynamic
            pcep
          !
            metric
             type igp
            !
           !
          backup-ineligible
          !
          constraints
            segments
              protection unprotected-only
              adjacency-sid-only
            !
          !
          bidirectional
            co-routed
            association-id 1010
            !
          !
        preference 50
          dynamic
            pcep
            !
            metric
              type igp
              !
            !
          constraints
```

```
        segments
          protection unprotected-only
          adjacency-sid-only
          !
        disjoint-path group-id 3
        type node, srlg, or link
        !
        bidirectional
          co-routed
          association-id 1050
          !
         !
      preference 100
        dynamic
          pcep
         !
        metric
         type igp
         !
        !
        constraints
          segments
            protection unprotected-only
            adjacency-sid-only
           !
          disjoint-path group-id 3
          type node, srlg, or link
          !
          bidirectional
            co-routed
            association-id 1100
            !
          !
        !
      !
    !
!
```

Review circuit-style SR-TE policy bandwidth utilization

Complete these steps to verify that the reserved bandwidth pool settings (defined when enabling CSM, see Enable SR CSM, on page 9) are correctly configured. You can also view current circuit-style SR-TE bandwidth utilization and how much is still available for use.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** and click the **Circuit style** mini dashlet. The Traffic engineering table lists all circuit-style SR-TE policies.

Step 2 Check the check box next to the circuit-style SR-TE policy you are interested in.

Step 3 From the topology map, click a participating circuit-style SR-TE policy node.

Step 4 From the **Device details** page, choose **Links** > ***link-type-entry*** > **Traffic engineering** > **General**.

Under Circuit Style bandwidth pool, you can see

- the reserved bandwidth pool size,
- the amount of bandwidth currently being used, and
- amount of bandwidth, allocated to circuit-style SR-TE policies, that is still available.

Figure 3: CS SR policy bandwidth pool

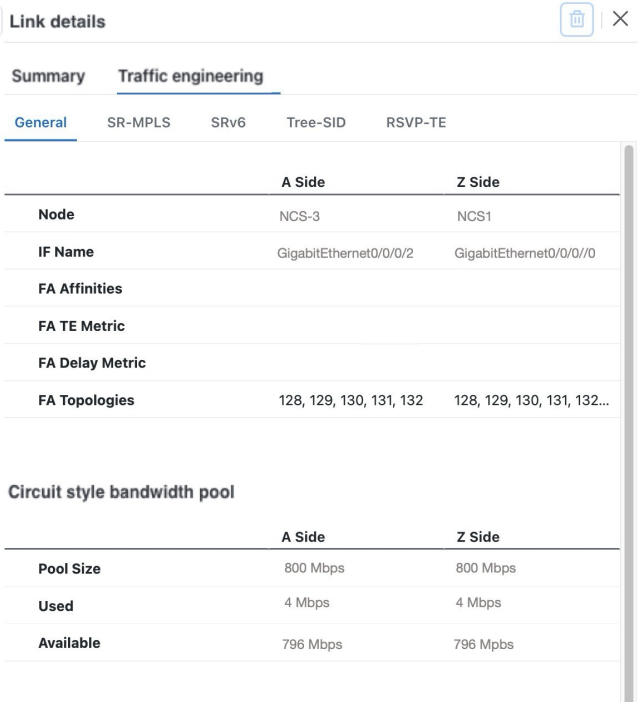

This example shows the reserved bandwidth pool size as 800 Mbps for NCS-3 and NCS1. The configured settings were earlier defined as 80% for the bandwidth pool size. Since the interface is 1 Gbps, we can confirm that CSM has correctly allocated 80% of the bandwidth for circuit-style SR-TE policies for these interfaces.

View circuit-style SR-TE policy information

Complete these steps to view circuit-style SR-TE policy information:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** and click **Circuit style**.

Figure 4: Circuit style dashlet

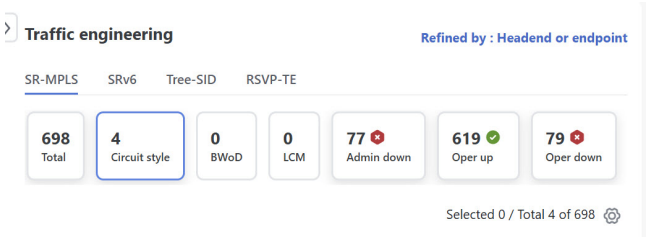

The table lists all circuit-style SR-TE policies.

Step 2 From the **Actions** column, choose ··· > **View Details** for one of the circuit-style SR-TE policies.

Note
You cannot edit or remove circuit-style SR-TE policy configurations that have been created directly on the device.

Figure 5: View circuit-style SR-TE policy details

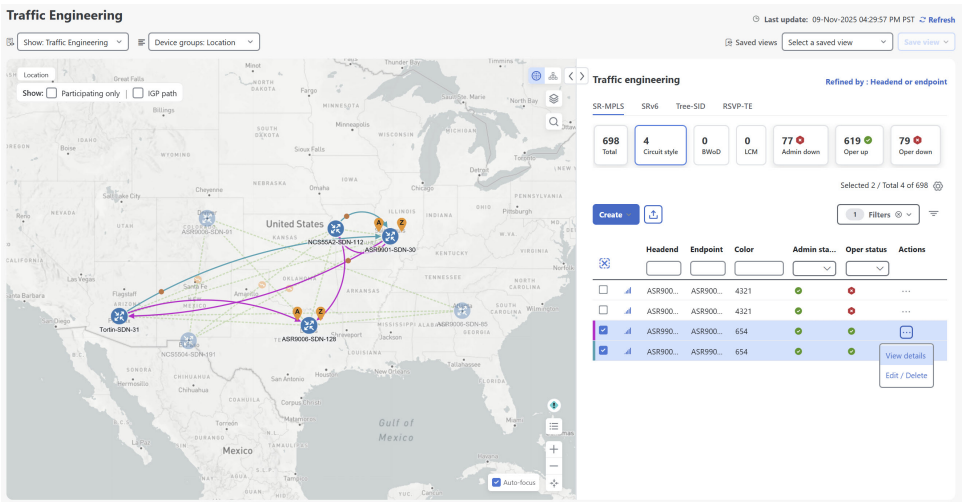

The **Circuit style policy details** page is displayed on the side panel. By default, the candidate path with an "active" state is displayed in the topology map. An active state is designated with a green "A" icon under **State**, indicating it is currently the operational active path. The map also has the **Bi-Dir path** checkbox checked by default, showing the bidirectional paths. The **Candidate path** list displays the candidate path with an active status (path that takes traffic) and other candidate paths.

Figure 6: CS-SR policy details summary

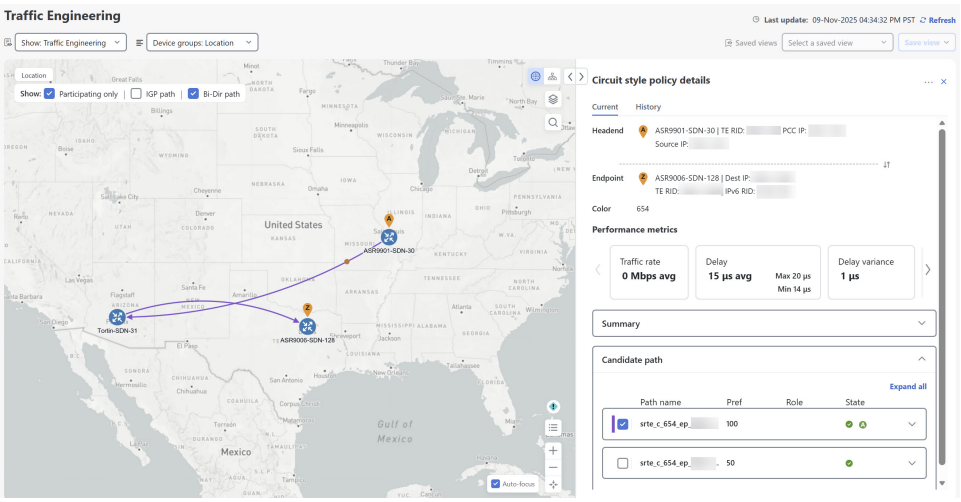

Note
The bandwidth constraint value can differ from the bandwidth you requested if the value is increased and insufficient resources exist to satisfy demand on all Working and Protect candidate paths.

Step 3 View candidate path configuration details.

a) The **Circuit style policy details** window allows you to drill down to view more information about the candidate paths. You can also copy the URL and share this information with others.

The Working path (highest preference path) with an operational state (Oper state) "Up" will always have an active state indicating that it takes traffic (see CSM path failure management , on page 23). If the Working path goes down, the Protect path is activated. In this example, the Protect path (with preference 50) is active and displayed on the topology map. Click **Expand all** to view more information about both paths.

Figure 7: Candidate path on topology map

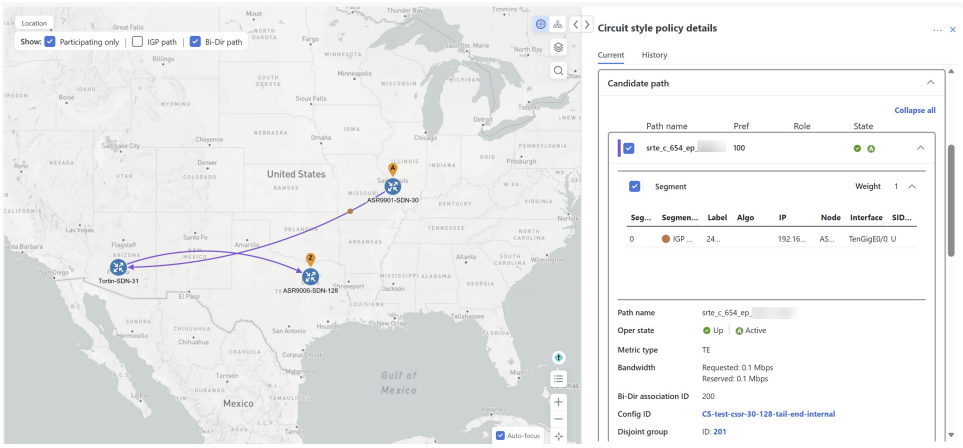

Note

- First preference paths are shown as purple links.
- Second preference paths are shown as blue links.

- Third preference paths are shown as pink links.

If the circuit-style SR-TE policy configuration was done through Crosswork Network Controller, you have the option to view the circuit-style SR-TE policy configuration. To see the configuration, click the link next to **Config ID**.

Figure 8: Config ID in candidate path details

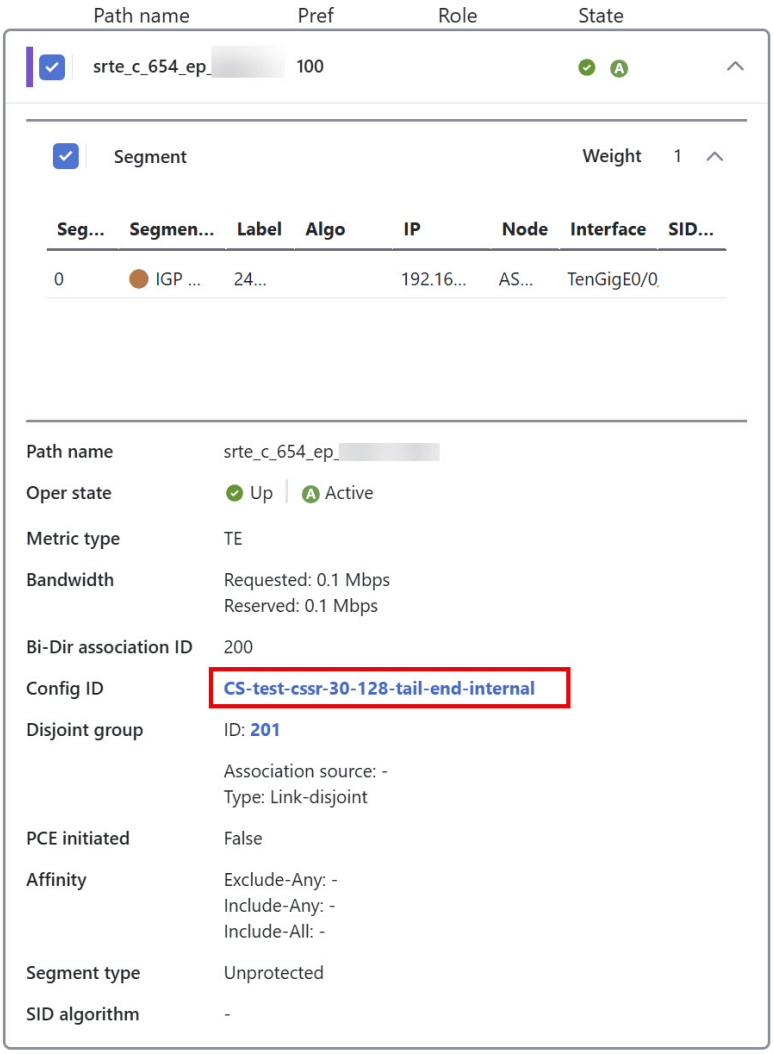

Here is a sample of a circuit-style policy configuration. For information on configuring CS-SR policies, see Configure circuit-style SR policies, on page 11.

Figure 9: Circuit-style policy configuration example

Step 4 To view the physical path and metrics between endpoints of the selected circuit-style SR-TE policies, click to turn applicable metrics on and check the **IGP path** checkbox.

Figure 10: IGP metrics

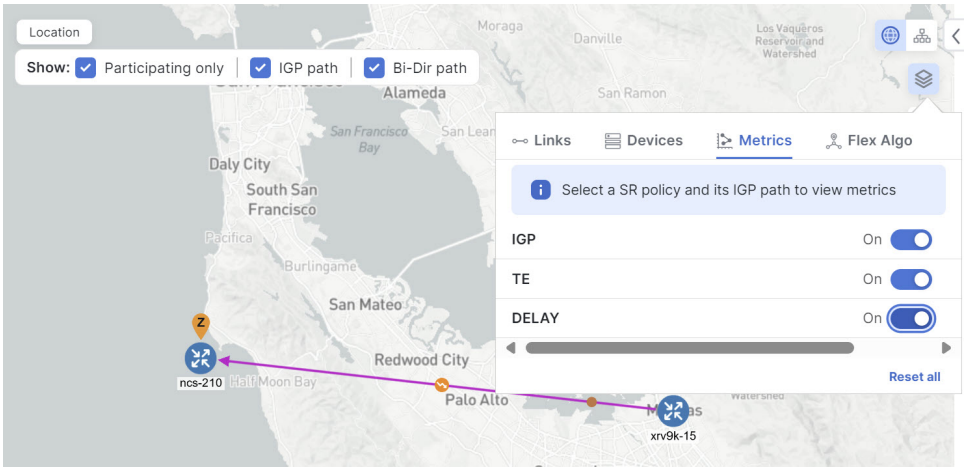

Trigger CSM to recalculate a circuit-style SR-TE policy

Circuit-style SR-TE policies are static in nature, meaning once the paths are computed, they will not be automatically reoptimized based on topology or operational status changes that may affect their paths. You can reoptimize a Working and Protect path (not a Restore path) after the policy's operational status went from down to up or if bandwidth size and requirement changes have been configured.

Complete these steps to manually trigger recalculation for a circuit-style SR-TE policy.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** and click the **Circuit style** mini dashlet. The **Traffic engineering** table lists all circuit-style SR-TE policies.

Step 2 From the **Actions** column, choose ··· > **View Details** for the circuit-style SR-TE policies you want CSM to recalculate a path for again.

Step 3 From the top-right corner, choose ··· > **Reoptimize**.

Effects of surpassing bandwidth reservation limits

CSM discovers and updates the available and reservable bandwidth in the network. It maintains an accounting of all bandwidth reservations provided for CS SR-TE policies to ensure that the total reserved bandwidth on

all interfaces remains at or below the network-wide resource pool (bandwidth pool size). When the bandwidth exceeds either the configured pool or the threshold, the system responds by:

- generating threshold-crossing event notifications for visibility and operational response,
- denying the policy establishment or bandwidth increase until resources become available,
- repeatedly retrying path computations at regular intervals (for example, every 30 minutes) until a solution is found or the policy is disabled.

Example: Bandwidth utilization surpasses defined threshold

In this example, we assume the reserved bandwidth settings are as follows:

- Link CS Bandwidth Pool Size: 10%
- Link CS Bandwidth Minimum Threshold: 10%

In this example, the bandwidth pool size for the 10 Gbps ethernet interfaces is 1Gbps and the alarm threshold is set at 100 Mbps (10% of pool size).

1. A circuit-style SR-TE policy from node 5501-02 to node 5501-01 (r02 - r01) is created with a bandwidth of 100 Mbps.

   Figure 11: CS-SR policy 10 mbps up

2. Later, the requested bandwidth configured for the policy is increased to 500 Mbps. CSM determines the additional bandwidth along the existing path is available and reserves it.

Figure 12: CS-SR policy 500 mbps up

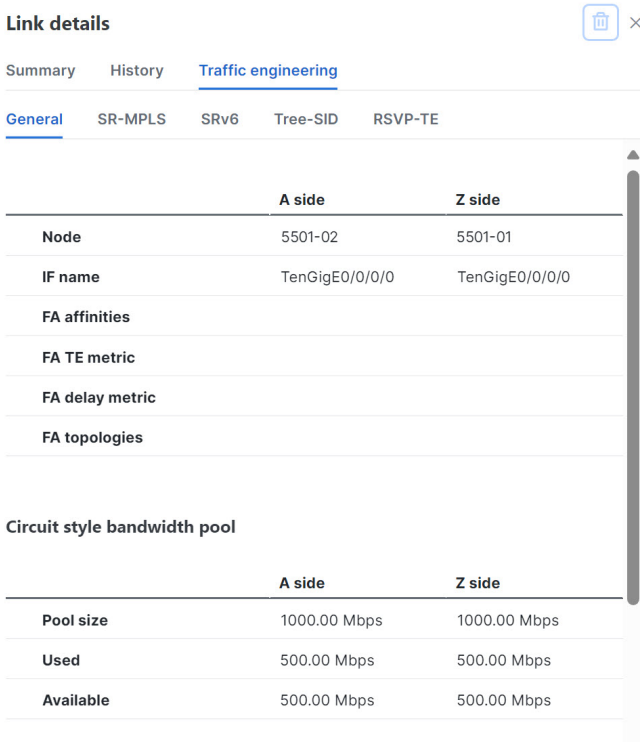

3. Since the bandwidth utilization (500 Mbps) with the updated policy is above the configured pool utilization threshold (100 Mbps), an event is triggered.

Figure 13: Threshold alerts

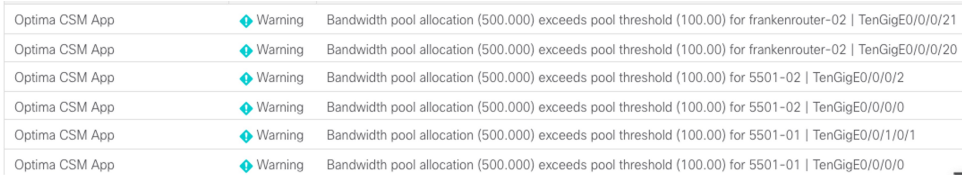

Example: Bandwidth pool size and utilization exceeded

In this example, we assume the reserved bandwidth settings are as follows:

- Link CS Bandwidth Pool Size: 10%
- Link CS Bandwidth Minimum Threshold: 90%

In this example, the bandwidth pool size for the 10 Gbps ethernet interfaces is 1Gbs and the alarm threshold is set for 900 Mbps.

1. An existing circuit-style SR-TE policy from node 5501-02 to node 5501-01 (**r02 - r01**) uses a bandwidth of 500 Mbps.
2. Later, a new policy requiring a bandwidth of 750 Mbps with a path from node 5501-02 to node 5501-01 to 5501-2 (**r02 - r01- r2**) is requested. Since the existing policy and this new policy together exceed the

bandwidth pool size, and alarm threshold of 1 Gbps (750 Mbps + 500 Mbps = 1250 Mbps), the following behaviors occur:

- The new CS-SR policy **r02 - r01 - r2** has been created, but remains operationally down because CSM cannot compute a path for the new policy. CSM will try again every 30 minutes to find a path that meets the bandwidth requirements.

Figure 14: CS-SR policy exceeds bandwidth pool size

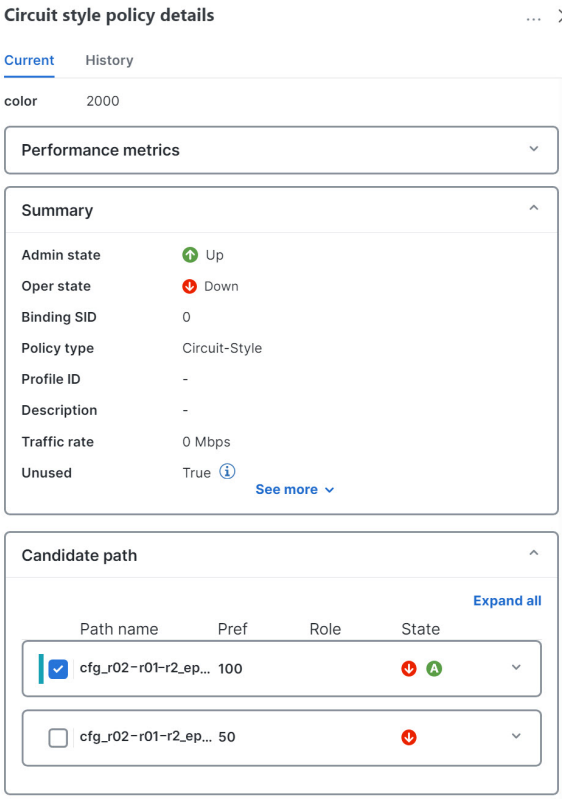

- Alerts are triggered.

Figure 15: Threshold alerts

3. Later, the circuit-style SR-TE policy **r02 - r01- r2** is updated and only requires 10 Mbps. The following behaviors occur:

   - Since the total bandwidth required for the two policies (10 Mbps + 500 Mbps = 510 Mbps) now requires less than the bandwidth pool size (1Gbps), circuit-style SR-TE policy **r02 - r01 - r2** receives a path computed by CSM and becomes operationally up.

Figure 16: Updated CS-SR policy operational

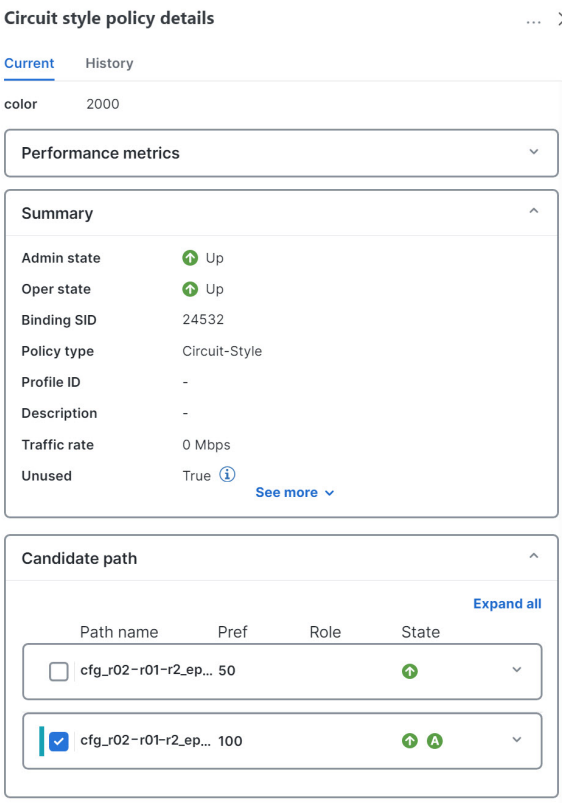

- Since the second circuit-style SR-TE policy with the reduced bandwidth is now provided a path by CSM , alerts are cleared.

Figure 17: Cleared alerts

CSM path failure management

Crosswork Network Controller computes paths for circuit-style SR-TE policies only after a complete bidirectional, path-protected set of candidate paths has been delegated. Three types of candidate paths are used during path failures:

- **Working** — This candidate path has the highest preference value.
- **Protect** — This candidate path has the second-highest preference value. If the Working path goes down, the Protect path (with the lower preference value) is activated. After the Working path recovers, the Protect path remains active until the default lock duration expires.

- **Restore** — This candidate path has the lowest preference value. Crosswork Network Controller computes the Restore path only after the Working and Protect paths are down. You can control how long after Restore paths are delegated from both sides to wait before the path is computed (see Enable SR CSM, on page 9). This delay allows topology and policy state changes to fully propagate to Crosswork Network Controller in cases where these changes triggered the Restore path delegation.

You can configure Performance Measurement (PM) to address path failures effectively and switch from the Working path to the Protect path. For more information, see Configure circuit-style SR policies, on page 11.

Examples

Note Illustrations are for demonstration purposes only and may not always reflect the exact UI or data described in the workflow content. If you are viewing the HTML version of this guide, click the images to view them in full size.

The following image shows that the Working and Protect paths of the circuit-style SR-TE policy are operational. The **active** path is indicated by the "A" icon.

Figure 18: Initial candidate paths

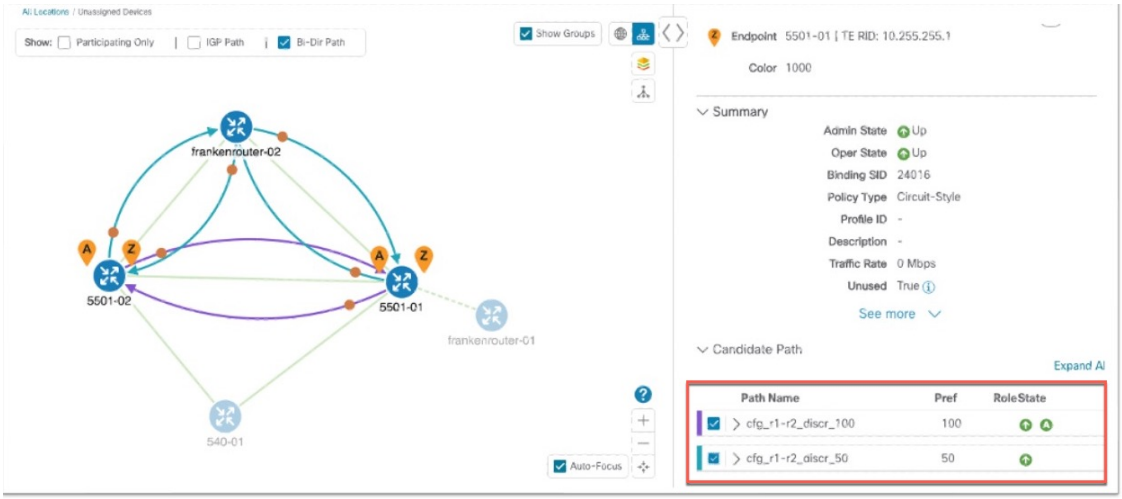

When a Working path having an active status goes down, the Protect path immediately becomes "active." When the Working path recovers, the Protect path moves to up/standby, and the Working path (with preference 100 in the example) becomes active.

Figure 19: Protected path becomes active

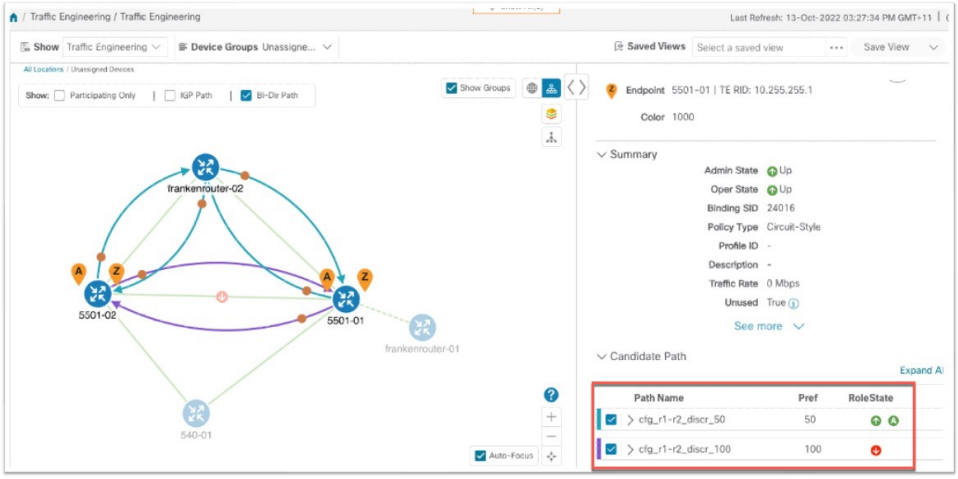

When both the Working and Protect paths go down, CSM calculates a Restore path, which becomes active. The Restore path only appears in this specific scenario. Note that the Restore path has the lowest preference value of 10 in the example. If the Working or Protected paths become operational again, the Restore path will no longer be visible on the topology map and will be removed from the **Candidate path** list.

Figure 20: Restore path

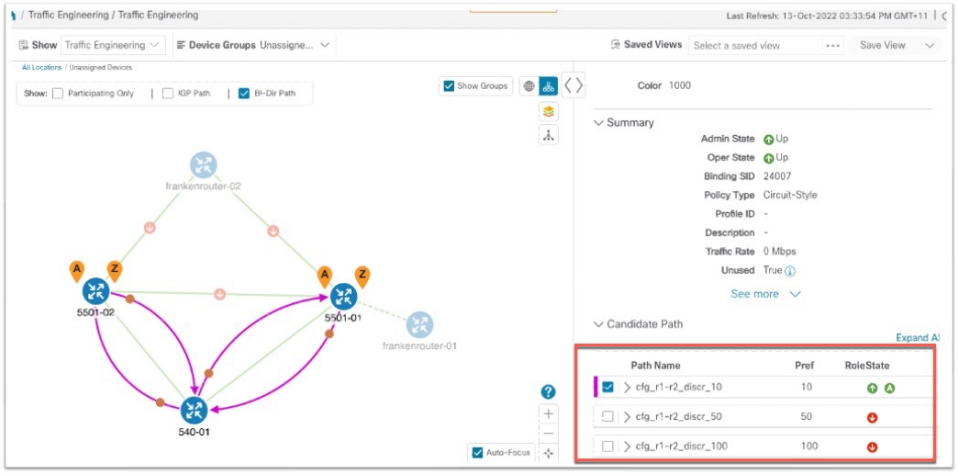

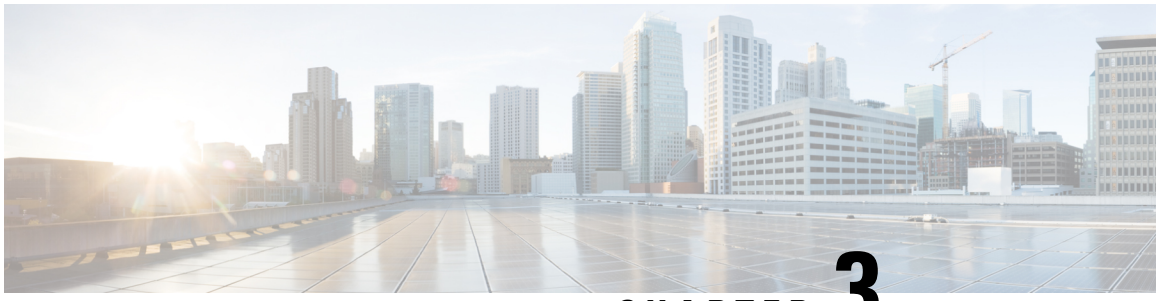

CHAPTER 3

Local Congestion Mitigation (LCM)

Local congestion mitigation

Local Congestion Mitigation is a network optimization technique that:

- monitors congestion as defined by the interface thresholds you specify and detects congestion on a configurable cadence (as opposed to a triggered event) by monitoring interface utilization and traffic thresholds,
- computes shortest paths for tactical policies to divert minimal traffic from congested interfaces to alternate paths with sufficient bandwidth,
- aims to keep as much traffic as possible on the original IGP path while mitigating congestion,
- provides localized mitigation recommendations in surrounding interfaces (local interface-level optimization) within a domain, eliminating the need to simulate edge-to-edge traffic flows in the network through a full traffic matrix,
- allows users to visually preview LCM recommendations before committing Tactical Traffic Engineering (TTE) SR policy deployments (feature available in Manual mode),
- can automatically deploy Multiple Segment List (MSL) policies for devices that are fully gRPC MSL compliant based on specified thresholds (feature available in Automated mode),
- supports automatic deletion of down, failed, or uncommitted LCM TTE policies to reduce network failure risks (see **Auto Repair Solution** and **Adjacency Hop Type**) in LCM configuration options, on page 55),

- collects TTE- SR policy and interface counters via SNMP and does not require Segment Routing Traffic Matrix (SR-TM), and
- is designed for scalability and applicability in large networks with multiple IGP areas, because of its simpler path computation and limitation to specific network elements.

Refer to Example: Mitigate congestion on local interfaces, on page 49 to see how to use LCM in your network.

LCM - device and network requirements

These requirements ensure that LCM has full visibility into network traffic and the capability to steer traffic effectively.

Enable traffic monitoring for LCM

For LCM to properly evaluate congestion, LCM requires traffic statistics from interface and headend SR-TE policy traffic measurements.

To ensure LCM is receiving these traffic statistics:

Procedure

Step 1 **Enable SNMP or gNMI**: Enable SNMP or gNMI on the devices whose traffic you want to monitor, including the headend device. For more information on how to configure these protocols, see the specific device platform configuration guide, for example, Configuring SNMP support.

Step 2 **Ensure device reachability**: Confirm all monitored devices are reachable from the Crosswork Data Gateway.

Enable strict SID label for LCM usage

All devices in the LCM domain must have strict SID enabled. Complete these steps in this example to configure strict SID on devices running Cisco IOS XR and XE.

To ensure LCM is receiving these traffic statistics

Procedure

Step 1 **Enable strict SID labels for all devices in the LCM domain**

Cisco IOS XR with ISIS

```
router isis core
interface Loopback0
address-family ipv4 unicast
prefix-sid absolute 16003
prefix-sid strict-spf absolute 16503
!
address-family ipv6 unicast
```

```
!
!
```

Cisco IOS XR with OSPF

```
router ospf 100
area 0
mpls traffic-eng
segment-routing mpls
interface Loopback0
passive enable
prefix-sid absolute 16002
prefix-sid strict-spf absolute 16502
!
```

Cisco IOS XE

```
segment-routing mpls
!
connected-prefix-sid-map
address-family ipv4
<ipv4-address> absolute 16010 range 1
exit-address-family
address-family ipv4 strict-spf
<ipv4-address> absolute 16510 range 1
exit-address-family
!
!
```

Step 2 **Enable segment routing with consistent SRGB**

Enable segment routing on the headend device and confirm that ***all*** devices

- use the same default SRGB range or a specified custom range.
- have the maximum SID depth explicitly configured if there are devices along the path that impose restrictions on the label stack depth.

```
segment-routing
global-block 16000 80000
traffic-eng
maximum-sid-depth 8
```

Step 3 **Configure SR-TE policies with strict SID labels**

If there are existing SR policies, the headend device must be configured to use strict SPF SID labels.

For PCC-initiated or computed SR policies

```
segment-routing
traffic-eng
policy srte_c_8000_ep
color 8000 end-point ipv4 <ipv4-address>
candidate-paths
preference 100
dynamic
metric
type igp
!
!
constraints
segments
sid-algorithm 1
```

For PCE computed or delegated SR policies

```
policy srte_c_8001_ep_198.19.1.4
color 8001 end-point ipv4 198.19.1.4
candidate-paths
preference 100
dynamic
pcep
!
metric
type igp
```

PCE configuration for returning paths with strict SID only

```
pce
segment-routing
strict-sid-only
```

Enable traffic steering using autoroute

The headend device must support PCE-initiated SR-TE policies with autoroute steering, a feature supported in Cisco IOS XR devices and provisioned via Cisco Crosswork Network Controller 7.2 using gRPC policy provisioning. Note that LCM will not operate correctly if the headend is a Cisco NCS device and there is L2VPN traffic in the network.

To enable traffic steering into SR-TE policies with autoroute:

Procedure

Step 1 Configure the headend device with `include ipv4 all` and `force-sr-include` under the appropriate PCC profile.

Example configuration

```
segment-routing
traffic-eng
pcc
profile 10      ! Profile ID must match the value in LCM Configuration > Basic > Profile ID
autoroute
include ipv4 all
force-sr-include
```

Note
The profile ID configured under the PCC profile must match the profile ID option set on the LCM Configuration page.

The profile ID identifies the PCC profile associated with the SR-TE policy provisioned by the PCE. This ID can be any integer from 1 to 65535, but it must match the profile ID used by the PCE to instantiate the policy. If the values do not match, the policy will not be activated. For example, if the PCE provisions a policy with profile ID 10, you must configure `segment-routing traffic-eng pcc profile 10 autoroute force-sr-include` on the headend router to enable autoroute announcement for that policy.

Step 2 Refer to the specific device platform configuration guide for more details (for example, *Segment Routing Configuration Guide, Cisco IOS XE 17 (Cisco ASR 920 Series)*.

Equal cost multi-path support

The headend device must support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies. To confirm that a device can support SR-TE policies with ECMP, verify:

- **Segment Routing is enabled**: Ensure Segment Routing is configured with a SRGB that matches the SRGB used by both the headend and tailend routers for SR-TE policies.

  Verify with `show segment-routing mpls state`

- **BGP-LS is enabled**: Confirm that BGP-LS is configured to advertise and receive link-state information from the headend and tailend routers.

  Verify status with `show bgp link-state link-state`

  Verify link-state information with `show bgp link-state link-state database`

- **ECMP is enabled**: Ensure ECMP is configured to load-balance traffic across multiple equal-cost paths.

  Verify ECMP routes with `show ip route`

  Verify the ECMP load-balancing algorithm with `show ip cef`

Prepare devices for gRPC policy management

To maximize LCM efficiency and improve performance, LCM can optionally provision SR-TE policies using weighted Multi-Segment Lists (MSL). This approach allows LCM solutions to typically consist of a single policy containing multiple weighted segment lists, enabling traffic detours without the need for parallel policies.

Weighted MSL LCM policies require gRPC policy provisioning instead of the legacy PCE-initiated method and are supported only on Cisco devices running IOS XR version 25.3.1 or later. This section details the additional configurations and requirements necessary to support these advanced capabilities across all participating devices.

- Enable gRPC on all devices in the target domain.
- Advertise SR MSL policies to BGP-LS peers and PCE neighbors.
- Prevent reporting MSL policies in PCEP.
- Add gRPC protocol connectivity to devices.
- Create and assign "grpc_msl" or "GRPC_MSL" tag to devices.

Note Ensure that devices with the gRPC port enabled are accessible from every Crosswork Network Controller node to support MSL policy deployment.

To leverage automated mode and support multiple segment lists with SR-TE, complete these steps:

Procedure

Step 1 **Enable gRPC for SR-TE policy reporting**

On devices running IOS XR version > 25.3.1, enable gRPC to allow policy services and communication.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config grpc
grpc
  segment-routing
    traffic-eng
      policy-service
    !
  !
  port 57400
  no-tls
```

Step 2 **Advertise SR MSL policies to BGP-LS peers and PCE neighbors**

To provide full visibility and support network orchestration, SR MSL policies must be advertised via BGP-LS both to peers and to the PCE neighbor. This involves enabling reporting of SR MSL policies into the link-state database and configuring BGP sessions with the PCE neighbor in the link-state address family.

a) **Enable reporting of SR MSL policies to BGP-LS peers**

Configure your router to report both active and inactive SR MSL policies into the link-state database. This allows policies to be advertised via BGP-LS to controllers or peers. Use the following configuration snippet to enable reporting of all configured SR MSL policies

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config segment-routing traffic-eng distribute link-state
segment-routing
  traffic-eng
    distribute link-state
      report-candidate-path-inactive
    !
  !
!
```

b) **Advertise SR MSL policies to PCE neighbor via BGP-LS**

Establish a BGP session with the PCE neighbor and configure the link-state address family to ensure the PCE can receive and learn all SR MSL policies from the router.

Note
The link-state address family must be configured on both the headend and the PCE for successful exchange.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config router bgp
router bgp 60
  neighbor <NEIGHBOR_IP>   ! PCE neighbor
    remote-as 60
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
    !
    address-family ipv6 unicast
    !
    address-family link-state link-state. ! Enable BGP-LS for SR MSL policy advertisement
    !
  !
```

Step 3 Prevent reporting MSL policies in PCEP

Since PCEP does not fully support MSL policies (it only advertises a single segment list, which can cause operational issues), it is recommended to remove the report-all command from the PCC configuration on the headend router. Use this configuration to prevent SR MSL policies from being reported via PCEP.

```
RP/0/RP0/CPU0:L4-NCS560#sh running-config segment-routing traffic-eng pcc
segment-routing
  traffic-eng
    pcc
      source-address ipv4 192.100.0.4
      pce address ipv4 100.100.0.1
        precedence 25
      !
      pce address ipv4 100.100.0.2
        precedence 50
      !
      !  Remove the following line to prevent reporting MSL policies to PCE
      !  report-all
      redundancy pcc-centric
      profile 1981
        autoroute
          include ipv4 all
          force-sr-include
        !
      !
    !
  !
!
```

Step 4 Add gRPC protocol connectivity

In Device Management, ensure all devices in the target domain have gRPC protocol connectivity configured.

a) You can add gRPC protocol in the device credential profile by navigating to **Device Management** > **Network Devices**.

b) Edit each device as needed to add gRPC connectivity details.

Figure 21: Edit devices for gRPC connectivity details

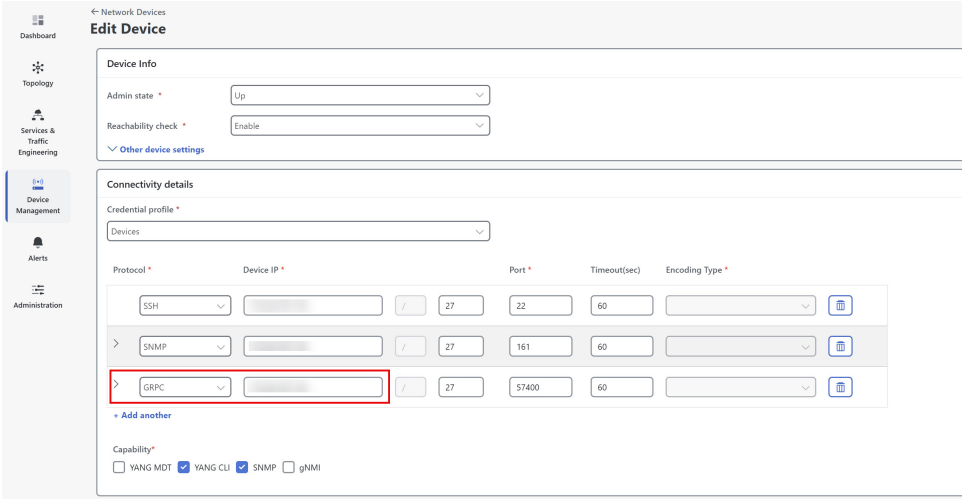

Step 5 Create and assign "grpc_msl" or "GRPC_MSL" tag to devices

Note

LCM will only deploy an MSL policy if this tag is present. In Automated mode, PCE-initiated policies are not supported. In manual mode, if the tag is missing, LCM will deploy a PCE-initiated policy.

a) Choose **Administration** > **Tag Management** >  . This displays the **Add tags** pane.
b) Choose the tag category from the **Select tag category** drop-down list or type a new category's name in the text field and click **Add**.
c) In the **Add tags for <category name>**, create a "grpc_msl" or "GRPC_MSL" tag and press Enter.
d) Click **Save**.
e) Navigate to **Device Management** > **Network Devices** and select the devices you would like to tag.
f) Click  . This displays the **Edit tags** pane.
g) In the **Associate tag** field, type "grpc_msl" or "GRPC_MSL" tag that you created.
h) Click on tag in the search result list to associate it with the device.
i) Click **Save**.

Important considerations when using LCM

Review these important considerations to ensure proper setup, optimal operation, and secure management of LCM domains.

User roles and permissions

- Ensure user roles have LCM task permissions for a domain before configuring LCM and committing recommendations. For more details on RBAC and user roles, see the Cisco Crosswork Network Controller Administration Guide.
- Device Access Group (DAG) access is **not** supported by LCM. Users with LCM permissions can configure and commit LCM recommendations regardless of whether or not they have DAG access for any devices in that domain.

Supported network features and limitations

- Do not steer LDP-labeled traffic into LCM autoroute TTE SR policies. LCM does not support LDP-labeled traffic.
- Avoid using LCM on networks with Tree SID policies, as incomplete traffic measurements can distort calculations.

Domain management and device support

- Limit domains to a maximum of 2000 devices for efficient LCM operation. A domain is identified by the IGP process and domain ID from the PCC router's configuration (link-state instance-id) used for BGP-LS advertisement.
- LCM recommended solutions utilize resources within a single domain only.
- If domain interfaces or links are removed or go down (LINK_DOWN state), either intentionally or unintentionally, LCM configuration and the Domain UI card (see Configure LCM, on page 55) remain available until links are aged out, providing up to 4 hours for recovery.

- Manually remove links from the UI if you need to force domain removal before the automatic aging period. The domain remains "ready for deletion" until the last link is removed.

Traffic evaluation and statistics

- LCM evaluates network utilization on a regular, configurable cadence of 1 minute or more, with a default of 10 minutes. The cadence can be set lower to improve responsiveness but is typically equal to or greater than the SNMP polling interval.
- The traffic statistics collection interval affects how quickly LCM can respond to topology changes and LSP deployments that affect interface and LSP traffic measurements. Be aware that LCM can take up to twice the statistics collection interval plus the LCM evaluation interval for recommendations to fully reflect changes. During this period, LCM recommendations may evolve as the traffic measurements are updated and eventually fully converge in Crosswork Network Controller.

ECMP handling and optimization eligibility

- LCM uses ECMP across parallel TTE SR policies, assuming roughly equal splitting of traffic. Actual ECMP behavior depends on traffic patterns and aggregation. LCM can be configured to detect and notify about excessive uneven ECMP splitting.
- To mitigate the effects of uneven ECMP, the overprovisioning factor is used in LCM. For more information, see Configure LCM.
- Do not steer traffic from existing SR-TE policies into LCM TTE SR policies. Ensure existing non-LCM SR-TE policies do not use regular Algo-0 prefix SIDs. Any combination of Algo-1 Strict, Flexible Algorithm, or adjacency SIDs is recommended to prevent this traffic from being steered into LCM TTE SR policies.

High Availability (HA) and SR-PCE behavior

- After an HA switchover, you can manually add missing interfaces that were previously monitored or update domain configuration options once the system stabilizes. Missing interfaces may occur if added after the last cluster data synchronization.
- When an SR-PCE goes down, LCM enters a dormant stage, and remains so until all SR-PCEs are reconnected and their associated topologies are fully synchronized with the topology service. LCM does not have visibility into the state of the SR-PCE redundancy set.

BGP-LS speaker placement for multiple AS networks with a dedicated IGP instance between ASBRs

Interdomain latency-optimized SR policy path computation refers to the process of finding the best routes between different networks (autonomous systems) using SR-PCE, aiming to minimize data travel delay or latency. This approach is especially important when egress peer engineering (EPE) is not supported.

- **Dedicated IGP instance:** A dedicated Interior Gateway Protocol (IGP) instance can be configured between autonomous system border routers (ASBRs) across different autonomous system numbers (ASNs) to support this computation.
- **Topology reporting:** Identifying ASBRs that report the topology via BGP-LS (Border Gateway Protocol Link State) is essential for accurate topology discovery.

- **BGP-LS configuration:** At least one ASBR in each AS participating in the dedicated inter-AS IGP (for example, Domain 100) must have BGP-LS enabled to report the IGP between each ASBR.
- **BGP-LS identifier:** Each ASBR must use the same BGP-LS identifier to report the domain.
- **Multiple ASBR support:** Multiple ASBRs per AS can report BGP-LS topology, providing flexibility in topology reporting.

Figure 22: BGP-LS session reporting domain 100

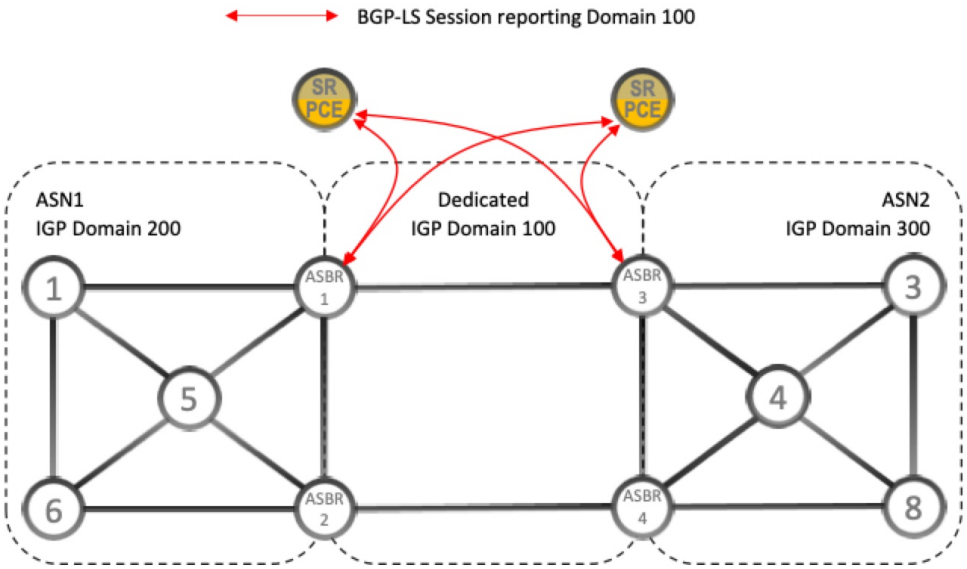

LCM calculation workflow

Summary

This example guides you through the process from congestion detection to the calculations performed by LCM before recommending tactical tunnel deployment. These calculations are conducted on a per-domain basis, enhancing scalability and enabling faster calculations for larger networks.

Figure 23: LCM configuration workflow example

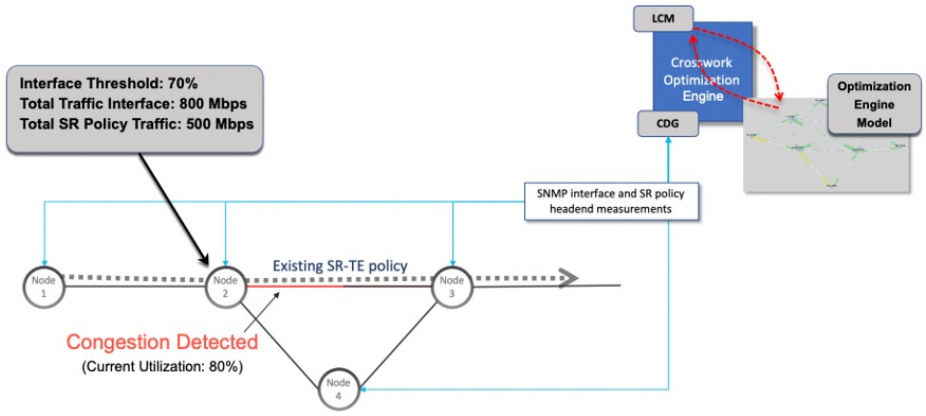

Workflow

1. **Analysis of network condition:** LCM first analyzes the Optimization Engine Model, which is a real-time topology and traffic representation of the physical network, on a regular cadence. In this example, LCM detects congestion after a congestion check interval when Node 2 utilization exceeds the 70% utilization threshold.

2. **Calculation of eligible traffic:** LCM calculates the amount of traffic eligible for diversion. LCM only diverts traffic that is not already routed on an existing SR policy or RSVP-TE tunnel (for example, unlabeled, IGP routed, or carried via FlexAlgo-0 SIDs). Traffic within an SR-TE policy is excluded from the LCM calculation and continues to travel over the original programmed path.

   LCM calculates the traffic eligible for diversion by subtracting the sum of traffic statistics for all SR-TE policies that flow over the interface from the total interface traffic.

   Total interface traffic – SR policy traffic and RSVP-TE tunnels = Eligible traffic that can be optimized

   This process must account for any ECMP splitting of SR policies to ensure the proper accounting of SR policy traffic. In this example, the total traffic on congested Node 2 is 800 Mbps, and the total traffic of all SR policies routed over Node 2 is 500 Mbps. So, the total traffic that LCM can divert is 800 Mbps – 500 Mbps = 300 Mbps

3. **Traffic diversion calculation:** LCM determines the amount of traffic that must be sent over alternate paths by subtracting the threshold equivalent traffic from the total interface traffic. In this example, LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path.

   800 Mbps – 700 Mbps (70% threshold) = 100 Mbps

4. **Over-provisioning factor (OPF):** The OPF represents a percentage deducted from the congestion threshold during solution computation to provide utilization headroom and account for uneven ECMP traffic distribution. For example, with a congestion threshold of 80% and an OPF of 3%, the optimizer uses an effective threshold of 77% when computing solutions. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see Configure LCM, on page 55.

5. **Determination of TTE SR policies:**

   - **Multiple parallel tactical SR-TE policies:**

LCM calculates the number of TTE SR policies needed and their paths based on the traffic split ratio. The ratio of how much LCM-eligible traffic can stay on the shortest path to the amount that must be detoured will determine the number of TTE SR policies needed on the shortest versus alternate paths.

In this example, LCM must divert one-third of the total eligible traffic (100Mbps out of 300Mbps) away from the congested link. Assuming perfect ECMP, LCM estimates that three tactical SR-TE policies are required to achieve this traffic split: one tactical SR-TE policy will take the diversion path and two tactical SR-TE policies will take the original path. There is sufficient capacity in the path between Node 2 and Node 4. Therefore, LCM recommends deploying three TTE SR policies (each expected to route approximately 100Mbps) from Node 2 to Node 3 via SR-PCE:

- 2 TTE SR policies to take a direct path to Node 3 (200 Mbps)
- 1 TTE SR policy takes hop via Node 4 (100 Mbps)

These recommendations are displayed in the **LCM operational dashboard**

Figure 24: LCM recommendation example

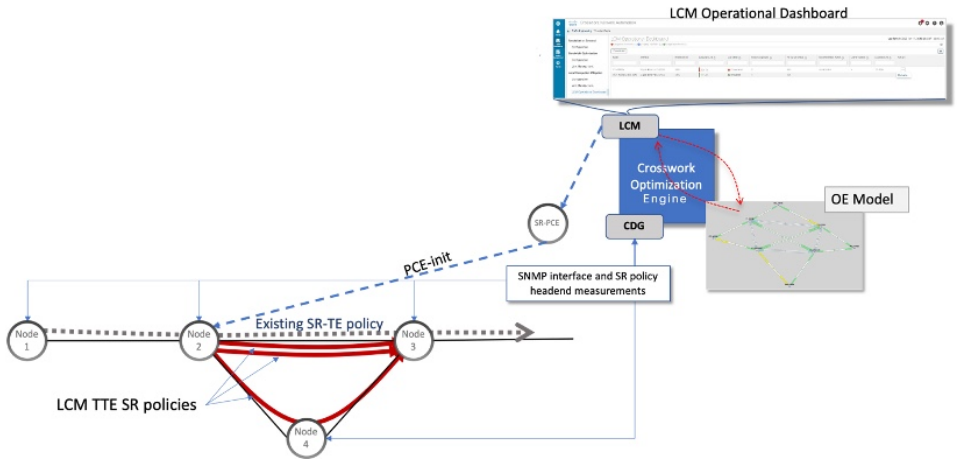

- **Single TTE SR policy with multiple weighted segment lists:**

  LCM calculates a single tactical SR-TE policy per congested interface, which includes multiple weighted segment lists. Each segment list corresponds to a distinct path, and the weights determine the proportion of total traffic steered along each path. In this approach, instead of deploying multiple parallel SR-TE policies that split traffic roughly equally, LCM uses weights to precisely control traffic distribution across the shortest path and one or more detour paths. This reduces the number of policies needed and eliminates downstream ECMP effects caused by parallel policies.

  In our example, instead of deploying multiple parallel tactical SR-TE policies, LCM creates a single tactical SR-TE policy. This policy includes two weighted segment lists to control traffic distribution precisely.

  - One segment list corresponds to the shortest IGP path from Node 2 to Node 3, weighted to carry approximately 200 Mbps (two-thirds of the eligible traffic).
  - The other segment list corresponds to the detour path via Node 4, weighted to carry approximately 100 Mbps (one-third of the eligible traffic).

  This approach simplifies policy management and provides finer control over traffic engineering compared to multiple parallel policies. As traffic patterns change, LCM can dynamically adjust the

weights of the segment lists within this single policy without adding or deleting policies, simplifying management.

6. **Monitoring and adjustments:** LCM continuously monitors the deployed TTE policies and recommends modifications or deletions as needed in the **LCM operational dashboard**. LCM recommends deleting deployed TTE SR policies if the mitigated interface remains uncongested after their removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Monitor LCM operations

The LCM dashboards and their monitoring roles include:

- **LCM domain dashboard**: Displays all discovered domains with key information such as domain identifiers, LCM status, configuration descriptions, operation modes, and urgency levels. It also provides links to configure LCM if not set up and to view TTE policy recommendations when congestion is detected.
- **LCM operational dashboard**: Displays congested interfaces based on configured utilization thresholds.
- **LCM operational history**: Displays detailed, time-stamped snapshots after each key LCM event, such as evaluation, commit, pause or resume, and mitigated or degraded states, and provides a chronological, visual record of congestion management activities.

For information on how to use LCM in your network, see the Example: Mitigate congestion on local interfaces, on page 49 topic.

LCM domains dashboard

The LCM Domain dashboard (**Services & Traffic Engineering** > **Local Congestion Mitigation**) displays all the domains discovered by Crosswork Network Controller. A *domain* is an identifier assigned to an IGP process.

Figure 25: LCM domains dashboard

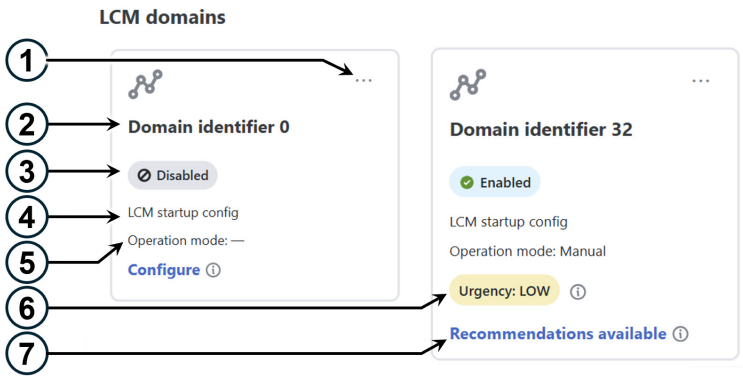

Callout No.	Description
1	**Main menu**: Allows you to navigate to these pages: • Operational dashboard • Operational history • Interface thresholds • Configuration
2	**Domain identifier**: The domain ID is taken from the router configuration (`link-state instance-id`) that is used to advertise IGP with BGP-LS.
3	**LCM status**: Indicates whether LCM is enabled for the domain or if the domain can be deleted.
4	**LCM configuration description**: The description is defined on the LCM Configuration page. The default description is "LCM startup config".
5	**Operation mode**: Indicates if LCM is running in Automatic or Manual mode. The default is Manual mode. • **Automated mode**—LCM automatically deploys TE tunnel recommendations based on thresholds that a user configures. Automated Mode is only supported on fully gRPC MSL compliant domains. • **Manual mode**—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
6	**Urgency**: Indicates the importance of the recommendation deployment or action. • **Low**: Indicates that LCM instantiated policies can be removed because they are no longer needed or that no changes are required. • **Medium**: Indicates new or modified recommendations. • **High**: Indicates network failures and recommendations should be deployed. This is a candidate that can be addressed automatically if the **Auto repair solution** advanced option was enabled. See Configure LCM, on page 55. **Dormant**: This status appears when the domain is inactive. LCM does not perform any operations on dormant domains.
7	**Configure**: This link appears if LCM has not yet been configured. Click **Configure** to go to the LCM Configuration page. **Recommendations available**: This link appears if LCM has detected congestion and has TTE policy recommendations. To view LCM recommendations, click the link to go to the **LCM operational dashboard**. **Delete**: Indicates that the domain card can be removed from LCM monitoring.

LCM operational dashboard

The LCM operational dashboard helps you preview the traffic engineering policies and path recommendations to mitigate congestion. Congested interfaces are those that exceed the configured utilization threshold.

Access LCM operational dashboard

To access LCM operational dashboard:

1. From the main menu, choose **Services & Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-ID*** > ··· > **Operational dashboard** .

   Figure 26: LCM operational dashboard

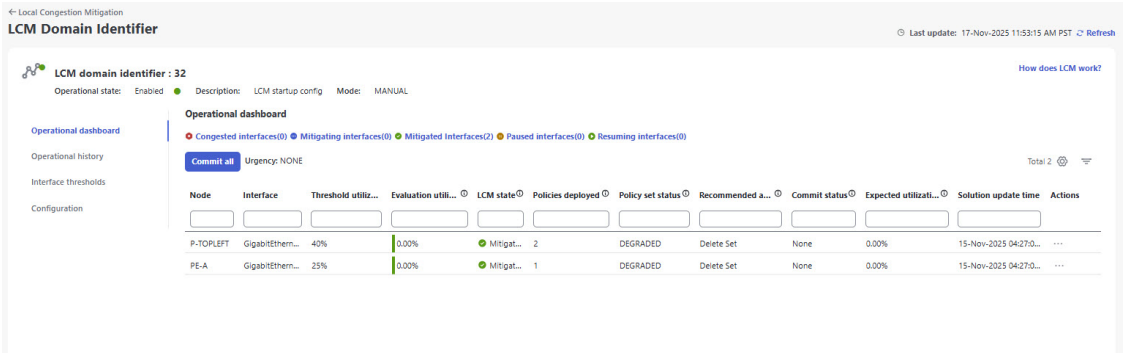

   Each interface lists details such as current utilization, recommended action, status, expected utilization after committing recommendations, and so on.

2. Hover the mouse pointer over ⓘ to view a description of the type of information each column provides. In the Actions column, you have these options:

   - **Preview solution**: This option opens the **Recommended TTE policies preview** page displaying the policies that LCM suggests for mitigating congestion. For headends that support MSL policies, you may see an MSL-based policy as a recommended solution.
   - **View deployed policies**: This option opens the **Traffic Engineering - LCM deployed policy** page displaying the various policies deployed to mitigate congestion. You can further click **View details** for each policy in the Actions menu to see the different paths and segments, along with detailed policy information.
   - **Resume/Pause**: Pause or resume an interface.

LCM states and dashboard behavior

Mitigated state: When LCM is in the Mitigated state, the dashboard displays the number of policies currently deployed to address previous congestion. It also provides recommended actions; for example, a "Delete Set" recommendation suggests removing these policies because congestion is no longer expected. The Commit status will show as **None** if the recommendation has not been committed, or **Committed** if it has been applied.

Congested state: When LCM detects congestion, the state changes to Congested. The dashboard shows recommended actions to remediate the congestion; for example, a "Create Set" recommendation advises

deploying new policies. You can preview the new solution by selecting the Actions menu (> **Preview solution**) before committing

Note If LCM cannot find a solution (**Recommended action - no solution**), it may be due to constraints set in the **LCM configuration**. For more information, see Configure LCM, on page 55.

Paused state: When LCM is in the Paused state, two events appear in the Operational history: one for the evaluation of policies for removal, and another for the commit of that removal. Similarly, when resuming policies, an evaluation event precedes the new commit event. If LCM is in a congestion check suspension interval, the operational history will display a resume event but will wait until the interval has passed before computing a new solution.

Recommendations are deployed as a set and require clicking **Commit all** to apply changes.

LCM operational history

The LCM operational history provides a powerful way to review and analyze LCM historical actions within your network. Events are generated at key points: during evaluation, when policies are committed or removed, when interfaces are paused or resumed, and when congestion states change. By capturing detailed, time-stamped snapshots after every key LCM events, it offers a chronological, visual record of congestion management activities. You can filter the data by nodes and interfaces to quickly locate relevant events. Historical records are retained based on a configurable setting (default is 30 days). This history enables auditing, troubleshooting, and operational review by showing:

- Date & time of each event
- Type of LCM event (commit, degraded, evaluation, mitigated, pause, and resume)
- Recommended actions or next steps based on the event outcome
- Configuration updates committed indicating whether the recommended changes were deployed to the network
- Total LCM policies deployed
- Number of congested interfaces
- Number of mitigating/mitigated interfaces
- Number of paused/resuming interfaces
- Filtering options to narrow down event list to specific nodes and interfaces.

Access LCM operational history

To view LCM operational history:

1. From the main menu, choose **Services & Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-ID*** > > **Operational history**.

Figure 27: LCM operational history

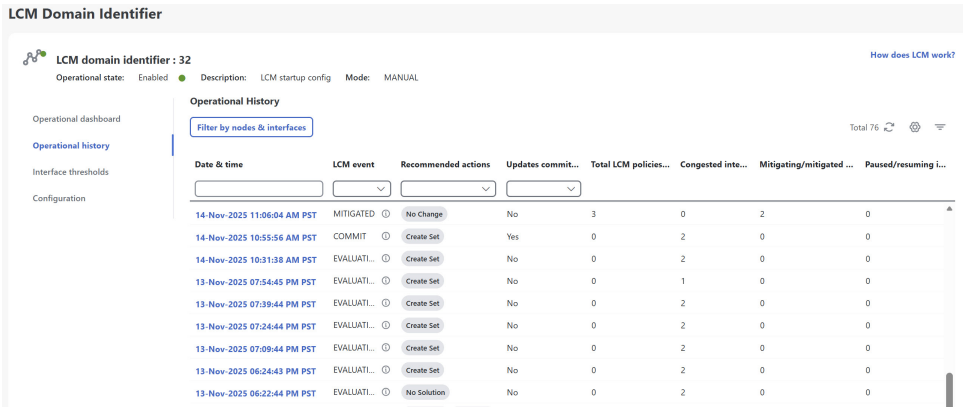

The table rows show the operational history of LCM events. If an event is paused, it will remain in the paused state until there is user intervention. You can review paused policies in the operational dashboard and resume them as needed.

2. Cick on any event in the table to see what the dashboard looked like at that specific point in time and as a result of that event. To better understand the information provided by the LCM operational history, let's click on the second row event (14-Nov-2025 10:55:56 AM PST) in the image above. It opens event details, where you can see that the interfaces were congested at 44% utilization.

Figure 28: Operational history event snapshort

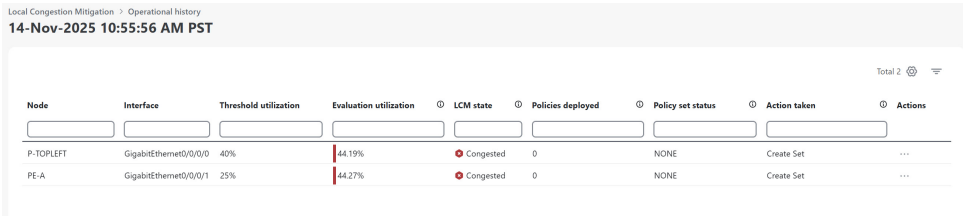

3. Click ··· in the Actions column and select **View proposed policies** to visualize policies in a generated solution, giving you better insights of the past events.

Figure 29: Operational history event snapshort - actions

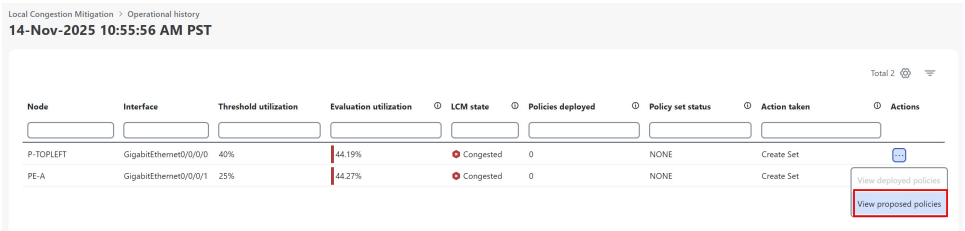

The Proposed policies page displays the policies that were proposed to mitigate congestion for that event.

Figure 30: Operational history event snapshot - proposed policies

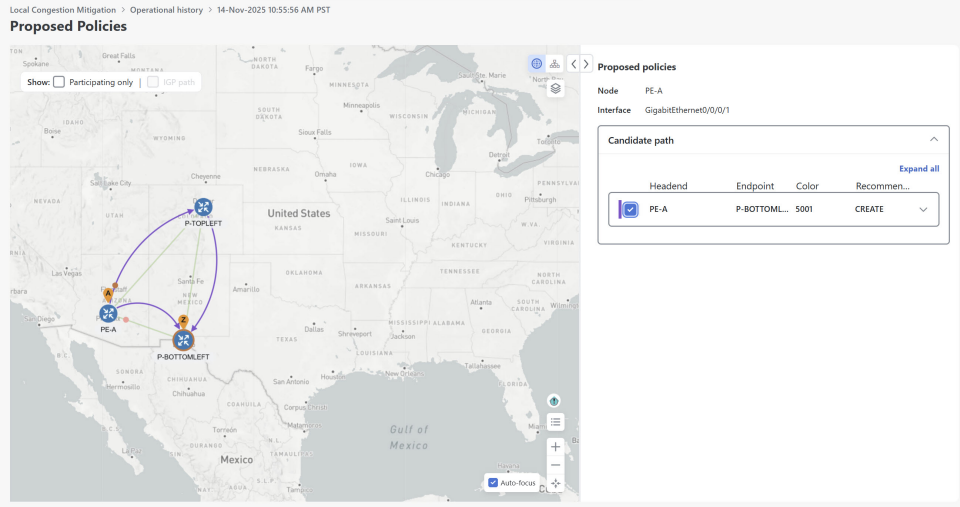

4. In the Candidate Path area, click **Expand All** to view the proposed policies along with segment list details. For headends that support MSL policies, you may see MSL-based policies as a proposed solution.

Figure 31: Proposed policy with multiple segment list

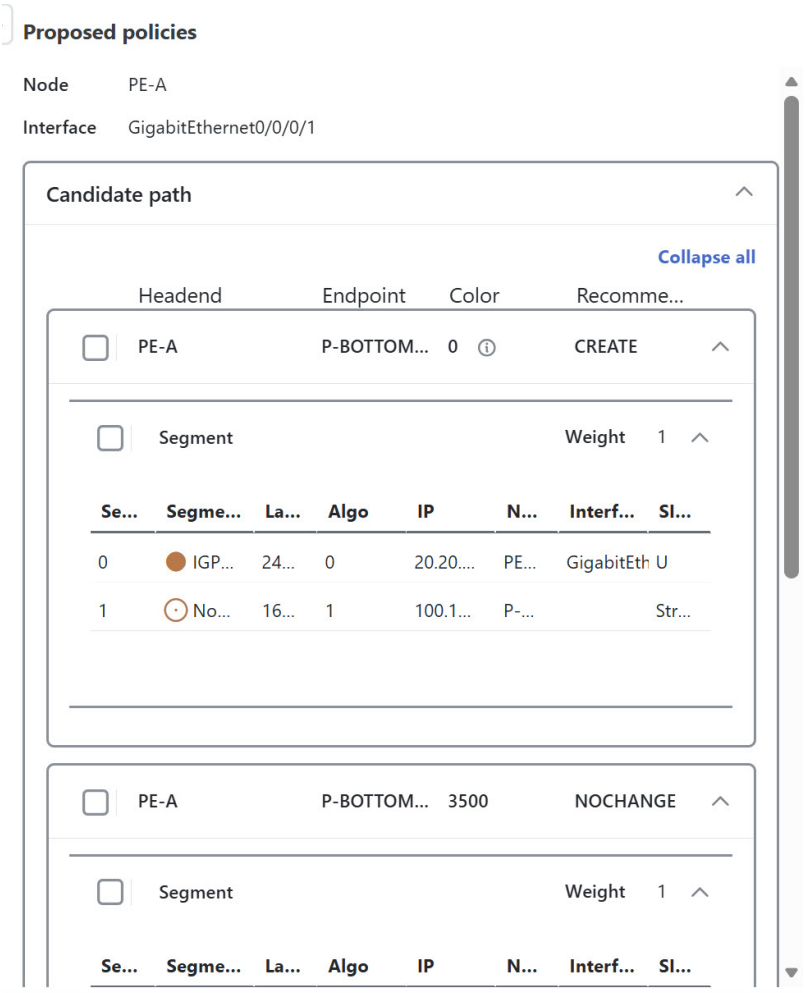

LCM solution events

The LCM generation of events are tied closely to the states of evaluation, user actions, and congestion check intervals.

Table 3: LCM solution events and descriptions

LCM event	Description
Evaluation	Indicates that a new recommendation is available, after LCM detects congestion and computes mitigation policies. At this stage, you can pause or commit the solution. Once the congestion check suspension interval has passed, depending on the action taken, the state will change to either mitigated if congestion is resolved, or degraded if issues persist. **Note** Policy preview colors may change after commit.
Commit	Indicates that the recommendations has been committed (deployed) to mitigate congestion. **Note** Policy preview colors may change if used by another policy.
Degraded	Indicates that the mitigation solution has not fully resolved congestion on the interface, or that congestion has worsened despite the committed policies.
Mitigated	Indicates that the committed recommendation has successfully resolved congestion after the congestion check suspension interval has passed. The interface is no longer congested.
Pause	Indicates that a request to pause the solution has been received. The interface is temporarily excluded from mitigation calculations. Pausing triggers two events in operational history: one indicating the **evaluation** of the policies for removal, and another for the **commit** of that removal. LCM waits until the congestion check suspension interval has passed before computing the solution. The user can later resume the interface, which also generates a resume event.
Resume	Indicates that a request to resume the solution has been received. The interface is re-included in mitigation calculations. In the Operational history, you will see an evaluation event preceding the commit event. If LCM is in a congestion check suspension interval, the operational history will display a resume event but will wait until the interval has passed before computing the solution.

Temporarily exclude an interface from LCM

You can temporarily pause LCM from including an interface for mitigations in either Automated or Manual modes. When an interface is paused, it will no longer be considered as part of a recommendation, and any existing solutions that the interface participates in will be removed. Pausing operations in Automated mode may be necessary in many use cases, such as the following:

- Where deployed solutions do not result in the intended resolution
- When there is uneven ECMP traffic

- When there are policies that do not carry traffic
- When an interface is continuously throttling between different solutions

LCM may automatically pause an interface when certain anomalies are detected, for example, when there is:

- No LCM SR policy traffic
- Excessive imbalance in LCM policy traffic
- Excessive LCM oscillations or removals per hour

In these circumstances, the user may perform a corrective action, and manually resume the interface.

Pause and resume an interface

From the Actions column of the LCM Operational Dashboard, select **Pause** for the interface you would like to exclude from LCM calculations. To include the interface in LCM calculations again, select **Resume**.

Note Pausing multiple interfaces at the same time may result in requests timing out. However, each request will be queued and displayed on the dashboard.

Figure 32: Pause interface

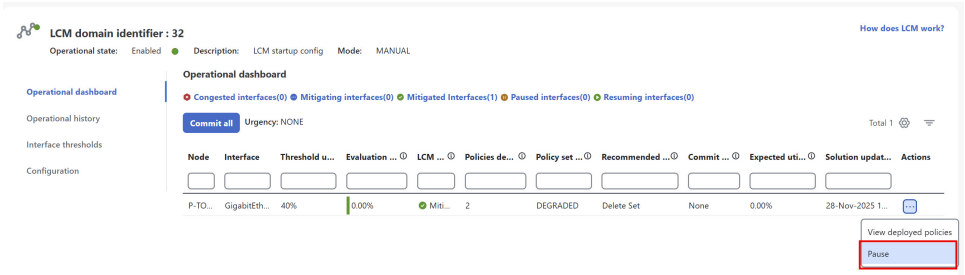

Mitigate congestion automatically

The LCM Automated mode allows the system to operate without user intervention by automatically making changes in the network. It enables continuous, closed-loop detection and mitigation of network congestion in domains where all devices are fully gRPC MSL compliant. To meet this requirement, each device in the target domain must support gRPC protocol connectivity and have the gRPC MSL tag applied to indicate compliance. When Automated mode is enabled, the system monitors network congestion and automatically applies or removes Multiple Segment Lists (MSL) policies to mitigate congestion based on configured thresholds, reducing the need for manual intervention.

Note In Automated mode, LCM can deploy only Multiple Segment Lists (MSL) policies; PCE-initiated policies are not supported.

LCM also applies auto repair soluions in automated mode. It not only mitigates congestion but also identifies network issues and proactively addresses them.

Automated mode includes safeguards to detect excessive solution oscillation. If an interface experiences repeated deployment and removal of mitigation policies beyond the configured threshold, the system automatically pauses automated actions for that interface to maintain network stability.

Before you begin

- Complete all requirements in the Prepare devices for gRPC policy management, on page 31 to ensure all devices in the target domain are gRPC-enabled.
- Ensure you have created the "grpc_msl" or "GRPC_MSL" tag and assigned it to the compliant devices in Tag Management.

Note LCM will only deploy an MSL policy if this tag is present. In Automated mode, PCE-initiated policies are not supported. In manual mode, if the tag is missing, LCM will deploy a PCE-initiated policy.

To enable Automated mode, complete these steps:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Local Congestion Mitigation** > **LCM-Domain-Card**. Click ··· and then choose **Configuration**.

Step 2 In the Advanced tab toggle the **Operation mode** option to **Automated**. If any device in the domain is missing the "grpc_msl" or "GRPC_MSL" tag or is not GRPC-configured, the system will prevent enabling Automated mode and display a list of non-compliant devices.

Figure 33: LCM configuration - Automated mode

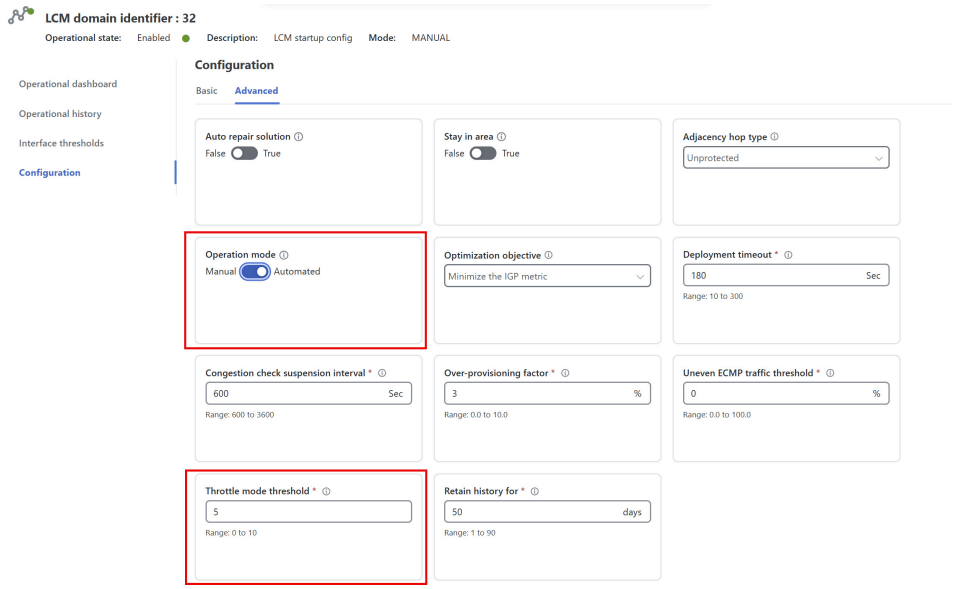

Step 3 In the **Throttle mode threshold** field, the maximum number of times LCM can apply or remove a solution on an interface within one hour. If policy changes on any interface exceed this threshold, LCM will pause automated actions for that interface and stop suggesting new solutions to maintain network stability. Paused events remain in this state until a user intervenes. You can review and resume paused policies from the operational dashboard as needed. Setting the value to 0 disables oscillation detection.

Step 4 Click **Commit changes** to save your configuration. After committing the configuration changes, LCM will display recommendations on the **Operational dashboard** if congestion occurs on any monitored interfaces.

The LCM Operational dashboard displays the current interface and mitigation status that is associated with the domain. Automated mode will continuously monitor congestion and deploy or remove MSL policies as needed based on configured thresholds.

You can also review deployed policies and paused interfaces as needed.

Figure 34: LCM operational dashboard

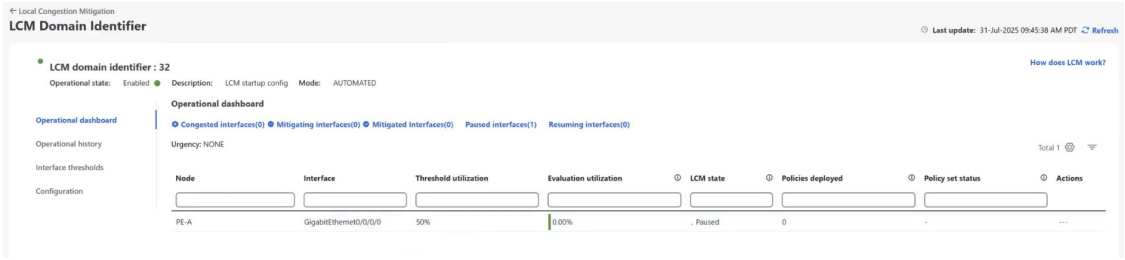

Example: Mitigate congestion on local interfaces

In this example, we will enable LCM and monitor the congestion mitigation recommendations to deploy TTE SR policies when an interface's utilization exceeds a defined threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion. The example covers the following steps:

1. View the uncongested network topology.
2. Set utilization thresholds for individual interfaces.
3. Enable and configure LCM in manual mode, which allows you to review recommended TTE policies before deciding whether to deploy them.
4. After LCM detects congestion, view the recommendations on the Operational dashboard.
5. Visually preview the recommended LCM TTE policies on the topology map.
6. Commit and deploy all recommended LCM TTE policies to mitigate congestion.
7. Verify that the LCM TTE policies have been successfully deployed.

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

Procedure

Step 1 View initial topology and utilization prior to LCM configuration. In this example, note that the node cw-xrv54 has a utilization of 7.17%.

Figure 35: Initial utilization

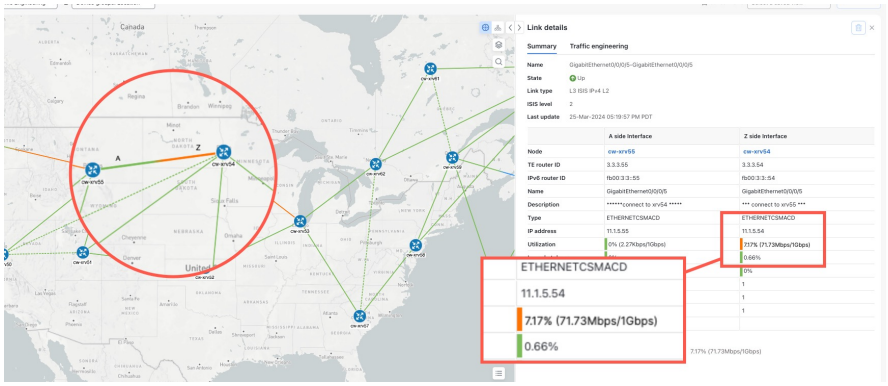

Step 2 Define any individual interface thresholds.

LCM allows you to configure a **global** utilization threshold that can be used for all interfaces. When traffic utilization surpasses the threshold, LCM will try to find bypass policies to remediate the congestion. You set the global utilization threshold on the **LCM configuration** page. However, if you want to define different thresholds for individual interfaces, we recommend defining them on the **Customized interface threshold** page before enabling LCM.

a) In this example, we will define an individual interface threshold. Go to the **Customized interface thresholds** page (**Services & Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-Identifier*** > ··· > **Interface thresholds**). You can add interfaces individually or upload a CSV file with a list of nodes and interfaces with custom utilization thresholds. For more information, see Add individual interface thresholds, on page 60.

See the following example and note the defined threshold for cw-xrv54 with interface GigabitEthernet0/0/0/1 is 20%.

Note
The utilization thresholds in this example are extremely low and best used for lab environments.

Figure 36: Customized interface thresholds

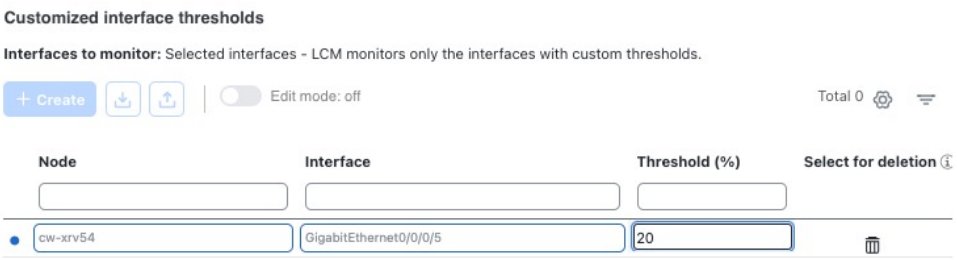

Note
By default, LCM monitors all interfaces. This includes any individual thresholds that are imported to this page. The rest of the interfaces will be monitored using the global **Utilization threshold** defined on the **LCM Configuration** page.

b) After adding interfaces and defining thresholds, click **Save**.

Step 3 Enable LCM and configure the global utilization thresholds.

a) From the main menu, choose **Services & Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-Identifier*** and click **Configuration**. Toggle the **Enable** switch to **True** and configure other LCM options. In this example, the global threshold is set at 80%, and the **Interfaces to monitor** > **All interfaces** option is selected. In the **Advanced** tab, Operation mode is set to **Manual**. For more information on all the available options, see Configure LCM, on page 55.

Figure 37: LCM Configuration page

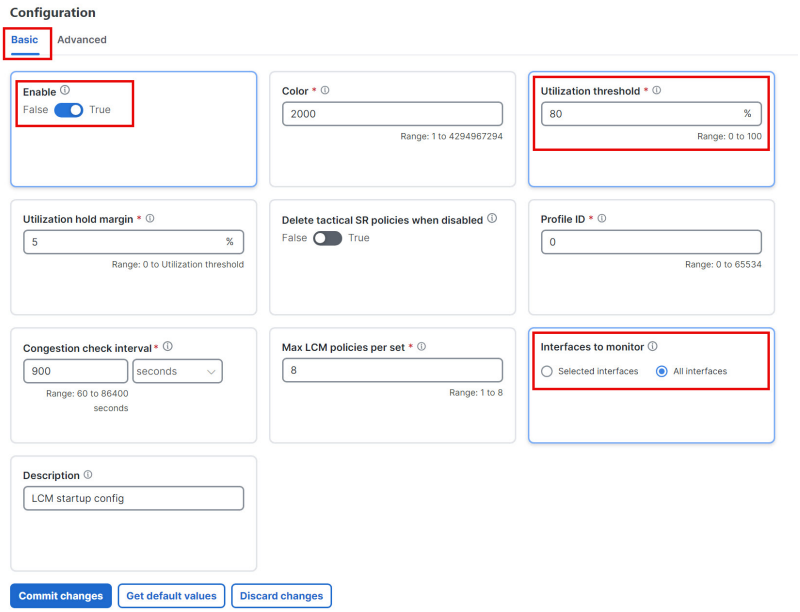

b) Click **Commit changes** to save your configuration. After committing the configuration changes, LCM will display recommendations on the **Operational dashboard** if congestion occurs on any monitored interfaces. LCM will not commit or deploy new TTE policies automatically when Manual mode is enabled. Later, you will be able to preview the recommended TTE policies and decide whether or not to commit and deploy them onto your network.

Step 4 After some time, congestion occurs, surpassing the custom LCM threshold defined at 20% for node cw-xrv54 with interface GigabitEthernet0/0/0/5.

Figure 38: Observed congestion

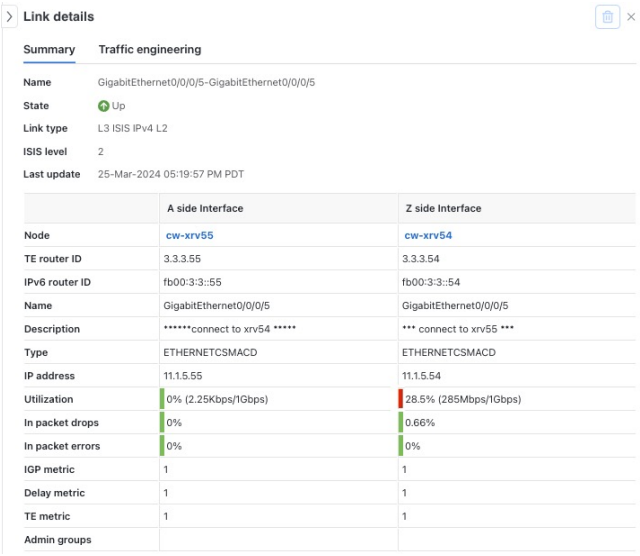

Step 5 View TTE SR policy recommendations in the LCM Operational Dashboard.

a) Navigate to **Services & Traffic Engineering > Local Congestion Mitigation**. When congestion is detected, the domain displays the urgency type and recommendations that are available. Click the question mark icons to display more information about the urgency type and when the most recent recommendation was given.

Figure 39: Congested detected and LCM recommendations

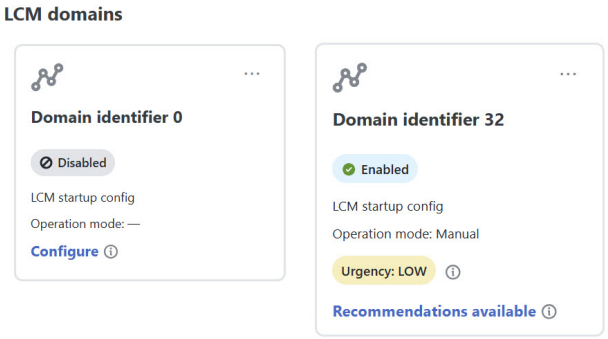

b) (Optional) View LCM events.

From the top-right corner of the Crosswork Network Controller UI, click > **Events** tab to view LCM events. You can also monitor this window to view LCM events as they occur. You should see events for LCM recommendations, commit actions, and any exceptions.

c) Open the Operational dashboard by navigating to **Services & Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-Identifier*** > ⋯ > **Operational dashboard**.

The dashboard shows that cw-xrv54 utilization has surpassed 20% and is now at 29.46%. In the **Recommended action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended action - Create set**) to address the congestion on the interface. For more information, see Monitor LCM operations, on page 39.

Note

If LCM cannot find a solution (**Recommended action - No solution**), it may be due to constraints enabled when configuring LCM (Configure LCM, on page 55).

d) Before committing TTE policies, you can preview the deployment of each TTE policy solution set. Click ⋯ in the **Actions** column and choose **Preview solution**.

Figure 40: Preview solution

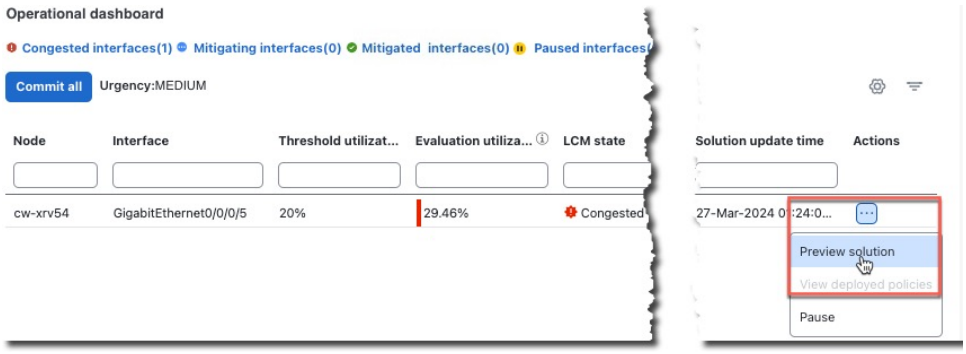

The resulting window displays the node, interface, and the recommended action for each TTE policy. From the **Preview** window, you can select the individual TTE policies and view different aspects and information as you would normally on the topology map. You can expand each policy to view individual segments. After reviewing the potential implications on your network, you can decide whether or not to deploy the bypass policies that LCM recommends.

The following figure shows the recommended TTE policies for node cw-xrv54.

Figure 41: LCM TTE deployment preview

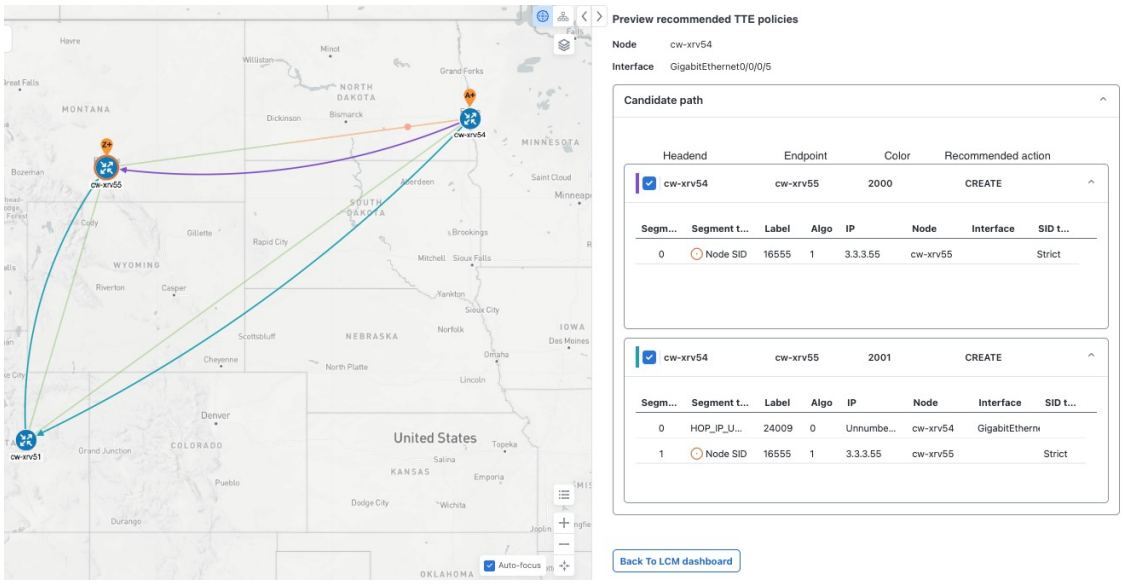

e) After you are done viewing the recommended TTE policies on the map, go back to the **Operational dashboard** and click **Commit all**. The **LCM state** column changes to **Mitigating**.

Figure 42: Mitigating state

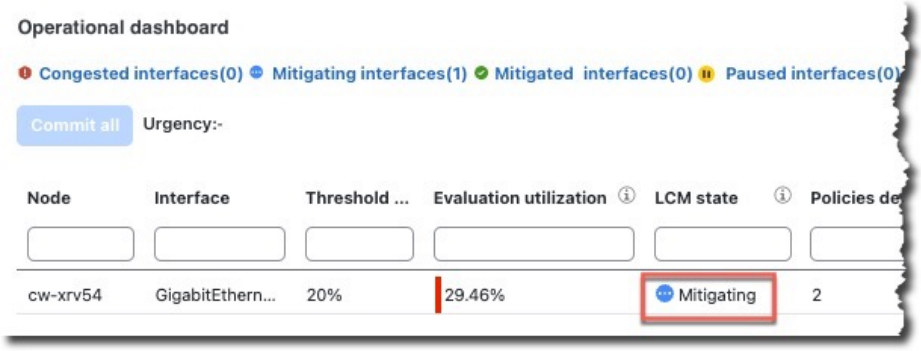

Note
All LCM recommendations per domain must be committed to mitigate congestion and produce the expected utilization as shown in the **Operational dashboard**. The mitigating solution is based on **all** LCM recommendations being committed because of dependencies between solution sets.

Step 6 Validate TTE SR policy deployments.

a) Click > **Events** tab. Note which LCM events are listed in the **Events** window.

Note
Crosswork Network Controller will report network events detected based on the policies and features you have enabled. For example, if a link drop causes an SR-TE policy to go down or if LCM detects congestion an event is displayed. These alerts are reported in the UI and, if desired, can be forwarded to third-party alerting/monitoring tools.

b) Return to the **Operational dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

Note
The LCM state change will take up to 2 times longer than the SNMP cadence.

c) Confirm the TTE policy deployment by viewing the topology map.

Click ... in the **Actions** column and choose **View deployed policies**. The deployed policies are displayed in focus within the topology map.

Step 7 Remove the TTE SR policies based on the LCM recommendation.

a) After some time, the deployed TTE SR policies may no longer be needed. This occurs if the utilization continues to stay under the threshold without the LCM-initiated TTE tunnels. If this is the case, LCM generates new recommended actions to delete the TTE SR policy sets.
b) Click **Commit all** to remove the previously deployed TTE SR policies.
c) Confirm the removal by viewing the topology map and SR Policy table.

In this scenario, we observed how LCM can be used to alleviate network traffic congestion. LCM automates tracking and calculations, reducing manual effort while still giving you control over whether to implement its congestion mitigation recommendations. You can preview recommendations and assess their potential impact on your network before deploying them. As traffic patterns change, LCM continuously monitors the

deployed TTE SR-TE policies and determines if they remain necessary. If a policy is no longer needed, LCM will recommend deleting them.

Configure LCM

To enable and configure LCM:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering > Local Congestion Mitigation > Domain-identifier-card**. Click ··· and then click **Configuration**.

Figure 43: LCM configuration

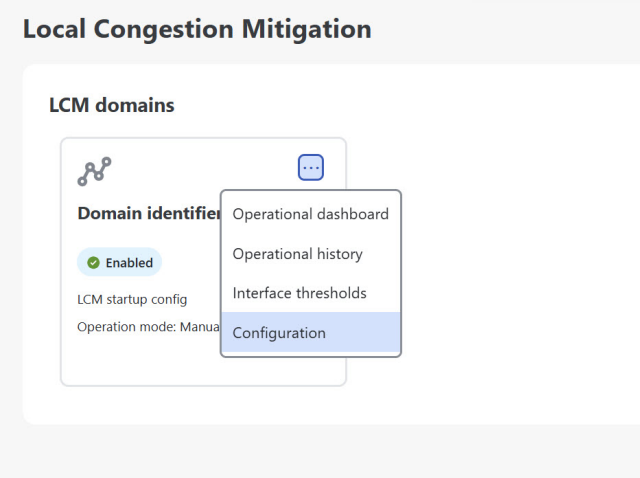

Step 2 Set **Enable** to **True**.

Step 3 Enter the required information. Refer to the LCM configuration options, on page 55 section to view the description of each field.

Step 4 To save your configuration, click **Commit changes**. If congestion occurs on any monitored interfaces, LCM will display recommendations on the LCM Operational dashboard. Note that LCM will not automatically commit or deploy new TTE policies. You can preview the recommended TTE policies and then decide whether to commit and deploy them to your network.

Note
If LCM is enabled, but cannot find a solution (**Recommended action - No solution**), it may be due to constraints enabled on this page.

LCM configuration options

These tables provide information on the LCM configuration options available in the UI.

- Basic LCM configuration
- Advanced LCM configuration

Basic LCM options

Table 4: Basic configuration options

Option	Description
Enable	Enables or disables the LCM function pack.
Color	Assigns color values sequentially to SR policies starting this value.
Utilization threshold	Sets the utilization percent at which LCM will consider an interface to be congested. This value applies to all interfaces unless you specify thresholds to individual interfaces on the **Customized interface thresholds** page.
Utilization hold margin	This value dampens the removal of deployed tactical SR policies. For example, if the utilization threshold is 90% and the utilization hold margin is 5%, then the tactical SR policy is removed from the network only when total interface utilization falls below 85% (90 - 5) without the tactical policy in place.
Delete tactical SR policies when disabled	Deletes all deployed tactical SR policies when LCM is disabled.
Profile ID	This is a required configuration to enable traffic steering onto LCM policies. Autoroute (steers traffic into the tactical SR-TE policies LCM creates) is applied to SR-TE policies through the proper **Profile ID** option that is set here to align with the configuration on the PCC associating that Profile ID with autoroute feature.
Congestion check interval (seconds)	Determines the interval at which LCM will evaluate the network for congestion. Under a steady state, when there are no recommendation commits, it uses this interval to re-evaluate the network to determine if changes are required. For example, if the interval is set to 600 seconds (10 minutes), LCM will evaluate the network every 10 minutes for new congestion and determine whether a new recommendation or modifications to existing recommendations are needed. Examples of modifications can include removal or updates to individual policies that were previously recommended. This option is typically set to greater than or equal to the SNMP polling cadence but can be set as low as 60 sec to improve responsiveness within the bounds imposed by the traffic collection interval.
Max LCM policies per set	The maximum number of tactical policies used to mitigate a single interface.

Option	Description
Interfaces to monitor	By default, this is set to **Selected interfaces**, and you will need to add thresholds to individual interfaces by importing a CSV file on the **Customized interface thresholds** page (**Services & Traffic Engineering** > **Local Congestion Mitigation** > **domain-identifier** > ··· > **Interface thresholds**). Only interfaces defined on the **Customized interface thresholds** page will be monitored. If set to **All interfaces**, LCM will monitor the interfaces with custom thresholds that are uploaded on the **Customized interface thresholds** page and the rest of the interfaces using the **Utilization threshold** value configured on this page.
Description	Description for the domain identifier.

Advanced LCM options

Table 5: Advanced configuration options

Option	Description
Auto repair solution	If set to **True**, LCM will automatically delete any down, failed, or uncommitted LCM TTE policies. This option is mainly to address a failure in a policy. If this option is disabled, and the **Urgency** status of the recommendation shown in the LCM Operational Dashboard is **High**, then the recommended solution is a candidate for the **Auto repair solution**. This means that a network failure will most likely occur if the solution is not deployed.
Stay in area	Restricts bypass LSP paths to stay within the mitigated area for OSPF or levels for ISIS.
Adjacency hop type (seconds)	If set to **Protected**, LCM will create SR policies using protected adjacency SIDs. This allows for Topology-Independent Loop-Free Alternate (TI-LFA) to compute a path for any adjacency failures. This option should only be set to **Protected** if all nodes in the same IGP area as LCM is operating are strict SPF SID capable.
Operation mode	• **Automated mode**—This option allows LCM to automatically deploy TE tunnel recommendations based on thresholds that a user configures. • **Manual mode**—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
Optimization objective	LCM calculates tactical SR policies based on the metric type chosen to minimize.
Deployment timeout	Enter the maximum number of seconds allowed to confirm deployment of tactical SR policies.
Congestion check suspension interval (seconds)	This interval determines the time to wait (after a **Commit all** is performed) before resuming congestion detection and mitigation. Since this interval should allow time for network model convergence, set the interval to no less than twice the SNMP collection cadence.

Option	Description
Over-provisioning factor (OPF)	This option reduces the congestion threshold by a set percentage during calculations to provide utilization margin and account for uneven ECMP traffic distribution. For example, with a congestion threshold of 80% and an OPF of 3%, the optimizer uses an effective threshold of 77% when computing solutions. The default value is 0.
Uneven ECMP traffic threshold	The percentage of sensitivity to detect uneven amounts of traffic across solution bypass tunnels.
Throttle mode threshold	In Automated mode, enter the number of times LCM throttles between solutions per hour until the interface is automatically paused.
Retain history for	The duration for which data is collected and retained before being deleted. The default value is 30 days. The deletion process occurs every 24 hours or when there is a change to the configured retention time.
Debug optimizer	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 opt max plan files**.
Maximum segment hops	Prior to using this option, you must create device tag groups to which you want to assign certain MSD values. For information on creating tags and assigning them to devices, see the Cisco Crosswork Network Controller Administration Guide. When calculating bypass TTE policies, LCM uses the effective Maximum SID Depth (MSD) value (as entered here) for specified device tags. You can assign up to five device tags with specific MSD values. A **0** value will not result in a solution. Setting a **0** value is equivalent to LCM monitoring and indicating when there is congestion in the network without providing a recommendation. The system learns from SR-PCE the MSD for each platform advertising the hardware limit in the IGP and BGP-LS. It represents the hardware limit that can be imposed exclusive of any service/transport/special labels. Therefore, you may want to use this new option to assign less than the advertised MSD value that LCM can use for bypass TTE policy calculation. To view the MSD value for a device, navigate to the **Traffic Engineering** topology map and click on the device. From the **Device details** page, click **SR-MPLS** > > **Prefixes** > **Expand all**.
Affinity	You can configure LCM to include or exclude links by using affinities to route data based on specific criteria. For example, if an affinity is excluded, LCM will try to alleviate a congested link by diverting traffic using paths that do not have that affinity. Affinities must already be configured on devices and then mapped using the Crosswork Network Controller UI in order to see the list of affinity names. See Example: Cisco IOS-XR affinity configuration, on page 59 and Configure link affinities, on page 59.

Configure link affinities

Link affinities are attributes or tags associated with links. Link affinities help in directing traffic along preferred paths based on specific criteria, such as bandwidth availability, latency, or cost. The affinity configuration on interfaces simply turns on some bits. It is a 32-bit value, with each bit position (0–31) representing a link attribute. Affinity mappings can be colors representing a certain type of service profile (for example, low delay, high bandwidth, and so on). Crosswork Network Controller sends bit information to the SR-PCE during provisioning.

If you have any affinities you wish LCM to account for when provisioning policy paths, follow these steps:

Procedure

Step 1 Configure affinities on your devices. See Example: Cisco IOS-XR affinity configuration, on page 59.

Step 2 Add affinities in Crosswork Network Controller, on page 59.

Step 3 Configure LCM, on page 55 using the advanced affinity option.

Example: Cisco IOS-XR affinity configuration

There are different ways to apply affinity configurations on a device.

See Segment Router configuration documentation for your specific device to view descriptions and supported configuration commands.

Cisco IOS-XR affinity configuration example

```
segment-routing
traffic-eng
interface GigabitEthernet0/0/0/1
affinity
name red
name blue
affinty-map
name red bit-position 1
name blue bit-position 5
```

Add affinities in Crosswork Network Controller

Crosswork Network Controller does not collect affinity names on devices. To make it easier to use link affinities, define affinity mapping in Crosswork Network Controller with the same name and bits that are used on the device. If affinity names are not mapped, the affinity name is displayed as "UNKNOWN" in the UI.

To add affinities, complete these steps:

Before you begin

Configure and note down the affinities on your devices.

Procedure

Step 1 From the main menu, choose **Administration** > **Settings** > **Traffic engineering** > **Affinity** > **TE link affinities**. You can also define affinities while configuring LCM (click **Manage mapping** under the **Constraints** > **Affinity** field).

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

Step 3 Enter the name and the assigned bit position.

Figure 44: Affinity

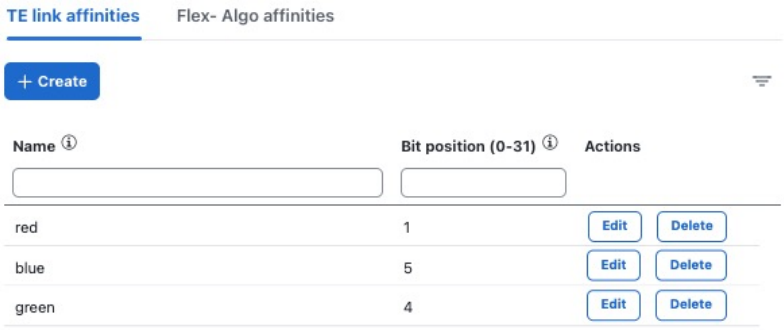

Step 4 Click **Save**. To create another mapping, you must click **+ Create** and save the entry.

Add individual interface thresholds

Networks have many different links (10G, 40G, 100G) that require different thresholds to be set. The **Customized interface thresholds** page allows you to manage and assign individual thresholds to nodes and interfaces.

Figure 45: Customized interface thresholds

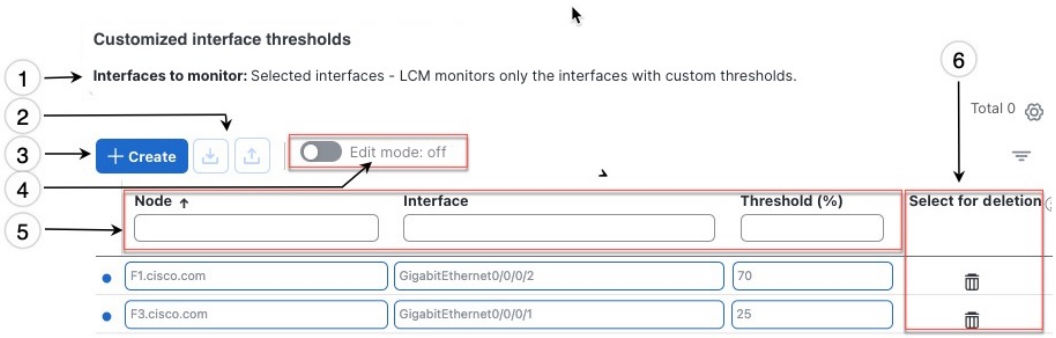

Callout No.	Description
1	**Interfaces to monitor**: Displays the option that is currently configured on the LCM Configuration page.

Callout No.	Description
2	**Import CSV file**: All interfaces currently in the table will be replaced with the data in the CSV file you import. **Export CSV file**: *All* interfaces are exported to a CSV file. You cannot filter data for export.
3	**+ Create**: Click this button to add new interface threshold rows.
4	**Edit mode**: When **Edit mode** is **ON**, you can edit multiple fields in one session, then click **Save**.
5	**Filter**: By default, this row is available for you to enter text in which to filter content.
6	**Select for deletion**: Click to delete the row. When **Edit mode** is **ON**, you can check multiple rows to delete, then click **Save**.

To assign specific threshold values for individual interfaces, complete these steps:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Local Congestion Mitigation** > **Domain-identifier** > ··· > **Interface thresholds**. Choose how you would like to add the interfaces.

- **Import CSV file**: Edit a CSV file to include a list of interfaces and thresholds, then later import the file into LCM.
- **Add new interface**: Manually add individual interfaces and thresholds.

Step 2 If you import a CSV file:

a) Click the **Download sample configuration file** link.
b) Click **Cancel**.
c) Open and edit the configuration file (LCMLinkManagementTemplate.csv) you just downloaded. Replace the sample text with your specific node, interface, and threshold information.
d) Rename and save the file.
e) Navigate back to the **Customized interface thresholds** page.
f) Click **Import CSV file** and navigate to the CSV file you just edited.
g) Click **Import**.

Step 3 If you manually add individual interfaces:

a) Click the first empty row and enter the appropriate node, interface, and threshold values.

Figure 46: Add first interface

b) Click **+ Create** to add more interfaces.

Step 4 Confirm that the information appears correctly on the **Customized interface thresholds** page.

Note
To update the table, you can either turn on Edit Mode or import a CSV file that replaces all current data in the table.

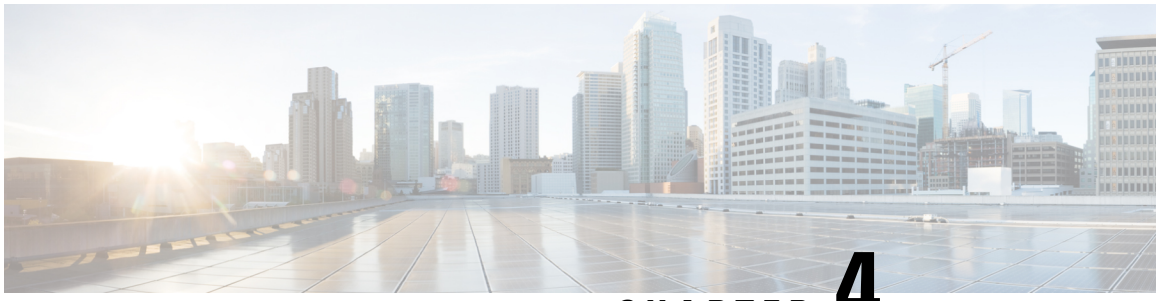

CHAPTER 4

Bandwidth on Demand (BWoD)

Bandwidth on Demand

Bandwidth on Demand (BWoD) is a bandwidth-aware Path Computation Element (PCE) that:

- integrates with SR-PCE to compute SR policy paths meeting requested bandwidth requirements,
- creates paths between endpoints in the network based on specific user-defined intents and dynamically maintains these bandwidth guarantees as network conditions change, and
- supports both PCC-initiated (PCE-delegated) and PCE-initiated policies, providing soft bandwidth guarantees over SR policies.

Key features of BWoD

Key features of BWoD include:

- **Intent-based path computation:** BWoD creates SR policy paths based on user-defined intents such as minimizing IGP cost, TE metrics, or latency.
- **Continuous monitoring and reoptimization:** BWoD continuously monitors network conditions and automatically re-optimizes BWoD paths to ensure that total BWoD traffic on any interface does not exceed a configured threshold percentage.

Note Functionality described within this section is only available with certain licensing options.

Limitations of BWoD

BWoD does not track total interface utilization; therefore, interfaces can still become congested if combined BWoD and non-BWoD traffic exceed capacity. In addition, BWoD does not enforce the total amount of traffic

entering a BWoD SR policy. BWoD policies may traverse Equal Cost Multi-Path (ECMP) paths assuming even traffic distribution, but actual ECMP distribution can be uneven, especially with large flows.

Important considerations when using BWoD

Consider the following information when using BWoD:

- Write-access to the head-end device is required based on Device Access Groups and assigned user roles. Only BWoD admin users can modify BWoD configuration settings. See the Cisco Crosswork Network Controller Administration Guide.
- If BWoD cannot find a path for a policy that guarantees its requested bandwidth, BWoD will attempt to find a *best effort* path if this option is enabled.
- BWoD disables itself when an unexpected error is encountered to avoid network disruption.
- BWoD temporarily pauses operation whenever the Optimization Engine model is unavailable due to an Optimization Engine restart or a rebuild of the topology from Topology Services. Any requests to BWoD during this time are rejected. When the model becomes available and BWoD receives two traffic updates from the Optimization Engine, BWoD will resume normal operation.
- If the Policy Violation advanced field is set to **Strict**, then the SR Policy Traffic option should be set to **Max Measured Requested**.
- After a switchover in a High Availability setup, BWoD policies created after the last cluster data synchronization will not be manageable and are considered orphaned TE policies. Crosswork Network Controller will display an alarm when it finds orphan TE policies (**Administration > Alarms**). You can use APIs to help clear these orphan policies so that they are manageable. For more information, see API documentation on Devnet.

PCC-initiated BWoD SR-TE policies

PCC-initiated BWoD SR-TE policies are traffic engineering policies that:

- allow devices to configure bandwidth requirements locally,
- delegate path computation to an external SR-PCE, and
- continuously optimize and monitor traffic-engineered paths based on bandwidth constraints.

BWoD automatically connects to all SR-PCE providers configured in Crosswork Network Controller, maintaining a persistent connection to the SR-PCE BWoD REST API, which registers as a PCE for bandwidth-constrained SR-TE policies. If bandwidth constraints cannot be fully met, BWoD computes best-effort paths and issues events accordingly. BWoD also monitors and re-optimizes paths to maintain bandwidth guarantees across the network.

How PCC-initiated BWoD SR-TE policies work

Summary

The key components involved in the process are:

- **Path Computation Client (PCC)**: Configures and initiates BWoD SR-TE policies and delegates path computation to the SR-PCE.
- **SR Path Computation Element (SR-PCE)**: Receives delegated policies from the PCC, coordinates bandwidth requirements, and interacts with BWoD functionality.
- **BWoD module**: Performs constraint-based path computation to satisfy bandwidth constraints, returns segment lists, and updates policy status as needed.

The following figure shows the PCC-initiated workflow for BWoD:

Workflow

Figure 47: PCC-Initiated BWoD SR-TE policies

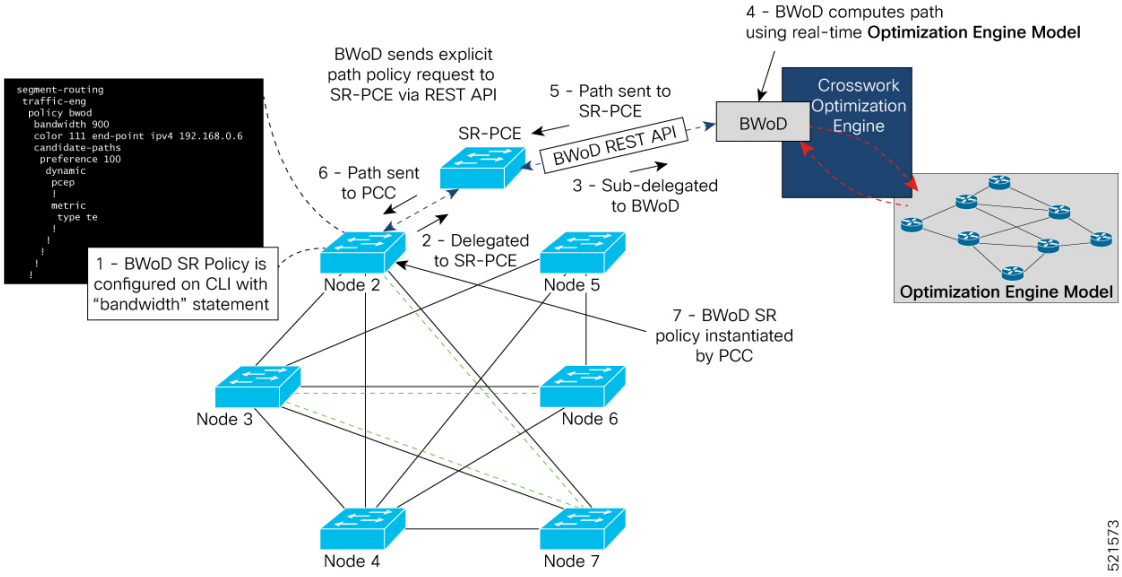

Table 6: PCC-Initiated BWoD SR-TE Policies

Callout No.	Workflow

1	The user configures a BWoD SR-TE policy on the PCC via the CLI, specifying bandwidth, endpoint, candidate paths, and constraints. For example:
	`segment-routing`
	`  traffic-eng`
	`    policy bwod`
	`      bandwidth 900`
	`      color 100`
	`      end-point ipv4 1.1.1.2`
	`      candidate-paths`
	`        preference 100`
	`          dynamic`
	`            pcep`
	`            !`
	`            metric`
	`              type te`
	`              !`
	`            !`
	`          constraints`
	`            affinity`
	`              exclude-any`
	`                name RED`
	`                !`
	`            !`
	`        !`
	`    !`
	`!`

2	Upon committing the configuration, the PCC delegates the path compute to SR-PCE.
3,4	SR-PCE further delegates the policy to BWoD, which attempts to compute a path that meets the requested bandwidth constraints.
5,6	If a bandwidth-compliant path is found, the segment list is returned to SR-PCE, which forwards it over PCEP to the PCC, and the PCC instantiates it. If BWoD is unable to compute a bw-compliant path for the policy or doing so will force an existing BWoD policy to not have a bw-compliant path, best effort paths may be computed by BWoD, which attempts to minimize violations. This occurrence will also trigger BWoD to issue an event to the Events UI indicating which BWoD policies are now on best-effort paths.
7	A BWoD SR-TE policy is instantiated.

Configure bandwidth on demand

Before you begin

Both CSM and BWoD cannot be enabled at the same time. If you have CSM enabled, you must disable it before enabling BWoD.

Note It is recommended to delete the respective policies from the network before disabling the associated function pack (BWoD or CSM). If policies remain in the disabled function pack, this may cause issues with new policy delegation and increase processing time.

Complete these steps to configure bandwidth on demand and create BWoD SR policies. As long as BWoD is enabled, you can create multiple BWoD SR policies.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Bandwidth on Demand** > **Configuration**.

Step 2 Toggle the **Enable** switch to **True**.

Step 3 Configure the required options. Refer to the BWoD configuration options section to view the description of each field.

Step 4 Click **Commit changes** to save the configuration.

Step 5 To create BWoD SR policies, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** tab and click **Create** > **PCE init**.

Step 6 In addition to entering the required SR policy details, click the **Bandwidth on demand** option and enter the required bandwidth.

Step 7 If applicable, enter a Flexible Algorithm constraint in the **SID Algorithm** field. The values correspond to the Flexible Algorithm that are defined on the device and the 128-255 range is enforced by Cisco IOS XR. Crosswork Network Controller will try to find a path with this SID. If a path with the SID constraint cannot be found, the provisioned policy will remain operationally down until the conditions are met.

Step 8 Click **Preview** to view the proposed SR policy.

Step 9 Click **Provision** to commit the SR policy.

BWoD configuration options

These tables provide information on the BWoD configuration options available in the UI.

- Basic BWoD configuration
- Advanced BWoD configuration

Basic BWoD options

Table 7: Basic configuration options

Option	Description
Enable	Enables or disables the BWoD function pack.

Option	Description
Primary Objective	Sets the primary objective when optimizing policies. • **Maximize available bandwidth**: Computes an SR policy path maximizing the overall available bandwidth in the network. This setting generally attempts to maximize usable network capacity at the expense of potentially longer paths. • **Metric minimization**: Computes an SR policy path minimizing the metric selected. This setting generally results in the shortest available paths for a metric type.
Link Utilization	Sets the congestion constraint (in percentage). When searching for paths for delegated policies, the Bandwidth on Demand will avoid any path that may exceed this utilization threshold, ensuring traffic is not routed through congested links. Additionally, when an interface exceeds this value, the system will trigger optimization to redistribute traffic and alleviate congestion.
Re-optimization interval (seconds)	Sets the minimum time (in seconds) before paths can be re-optimized if network conditions change. This acts as a countdown timer; the BWod policy will wait for this duration to expire before allowing re-optimization.
Metric re-optimization time	Sets the duration (in seconds) before paths can be re-optimized for metric improvements. If bandwidth requirements are met but a better IGP or TE path is available, BWod will wait for this timer to expire before re-optimizing. This helps prevent frequent path changes and unnecessary re-optimizations.

Advanced BWoD options

Table 8: Advanced configuration options

Option	Description
SR policy traffic	Determines how bandwidth optimization is calculated for each policy. • **Measured**: Uses the current measured traffic of BWoD provisioned SR policies for optimization caculations. • **Max measured requested**: Uses the maximum value between the current measured traffic on BWoD provisioned SR policies or the amount of bandwidth requested for optimization calculations.
Update throttle (seconds)	Sets the time (in seconds) to wait between updates. Set to 0 to disable throttling.

Option	Description
Optimizer event threshold (seconds)	Sends an alert to the UI if the optimizer runs longer than the specified time (in seconds). **Note** In large-scale environments, the optimizer run may exceed the default run time limit of 60 seconds. For deployments where policies are provisioned in batches of 100, setting the **Optimizer event threshold** to 240 seconds and the **Optimizer run time limit** to 300 seconds has been effective in internal testing. For environments where 1,000 or more policies are deployed, these values may need to be increased further. Adjust these parameters based on the scale of your deployment to ensure optimal performance.
Optimizer run time limit (seconds)	Sets the maximum allowed run time (in seconds) for the BWoD optimizer. **Note** In large-scale environments, the optimizer run may exceed the default run time limit of 60 seconds. For deployments where policies are provisioned in batches of 100, setting the **Optimizer event threshold** to 240 seconds and the **Optimizer run time limit** to 300 seconds has been effective in internal testing. For environments where 1,000 or more policies are deployed, these values may need to be increased further. Adjust these parameters based on the scale of your deployment to ensure optimal performance.
Policy violations	Determines how BWoD responds when new policies are provisioned and policy violations are detected.
Prefer strict SIDs	When enabled, prefers strict SIDs when deriving segment lists. Required for compatibility with LCM.
Debug optimizer	Enables logging of optimizer plan files to the Crosswork Network Controller file system. The number of saved files is limited by the **Debug opt max plan files** setting.

Provision an SR-TE policy to maintain intent-based bandwidth requirements example

This example demonstrates:

- how to enable and configure Bandwidth on Demand (BWoD)
- how to create BWoD policies
- how BWoD calculates paths, and
- how BWoD calculates new policies when the Policy violation option is set to Loose or Strict.

In this example, three BWoD policies will be created using the same headend (L1-NCS5501.cisco.sjc20.com) and endpoint (L5-8201-1-crosswork.cisco.com), at different bandwidths (700 and 1000 Mbps), with network utilization capped at 80%. All interfaces have the capacity of **1 Gbps**.

Figure 48: Initial BWoD topology

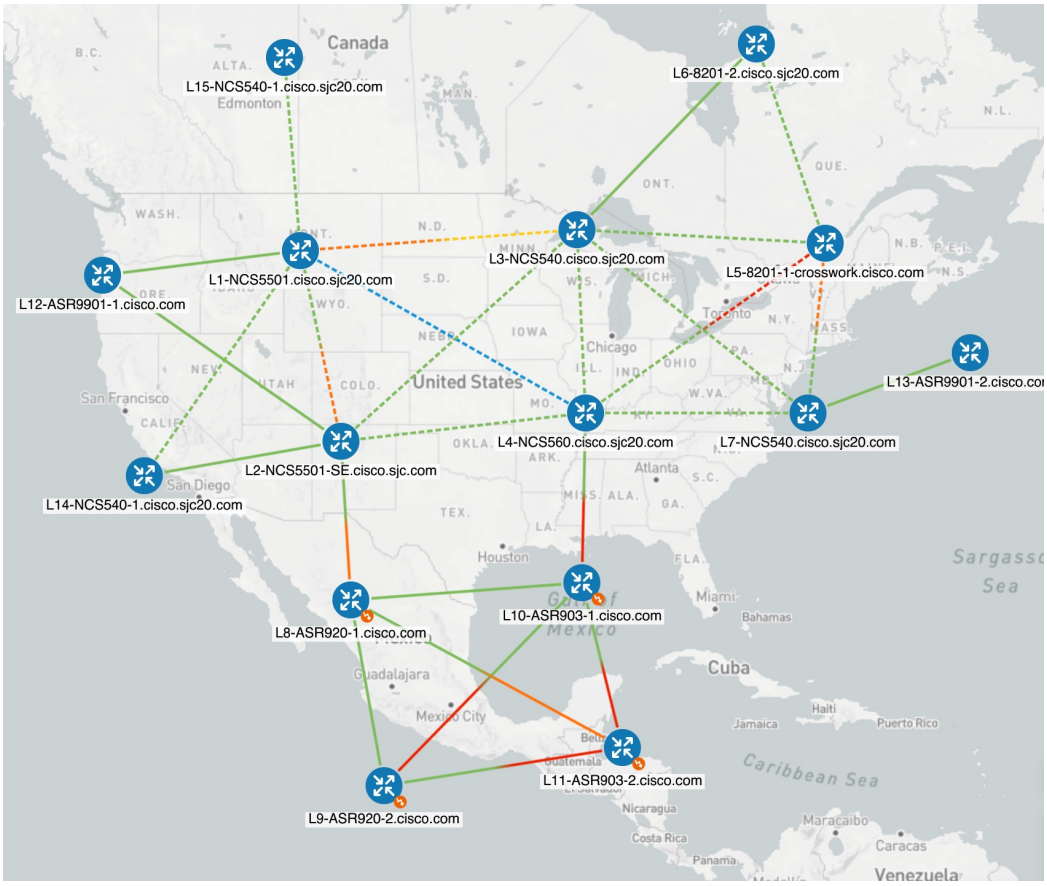

Procedure

Step 1 Enable and configure BWoD.

Note
Both CSM and BWoD cannot be enabled at the same time. If you have CSM enabled, you must disable it before enabling BWoD. It is recommended to delete the respective policies from the network before disabling the associated function pack (BWoD or CSM). If policies remain in the disabled function pack, this may cause issues with new policy delegation and increase processing time.

a) From the main menu, choose **Services & Traffic Engineering** > **Bandwidth on Demand** > **Configuration**.
b) Set Enable to **True**, enter **80** in the **Link utilization** field, and confirm that **Advance > Policy violations** is set to **Loose**. To find descriptions of other options, simply hover the mouse over ⓘ.
c) Click **Commit changes**.

Step 2 Create the first PCE-initiated BWoD SR-TE policy.

a) From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** tab and click **Create > PCE init**.
b) Enter the required policy details. In this example, we are creating a policy with these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70000**

Figure 49: Policy details

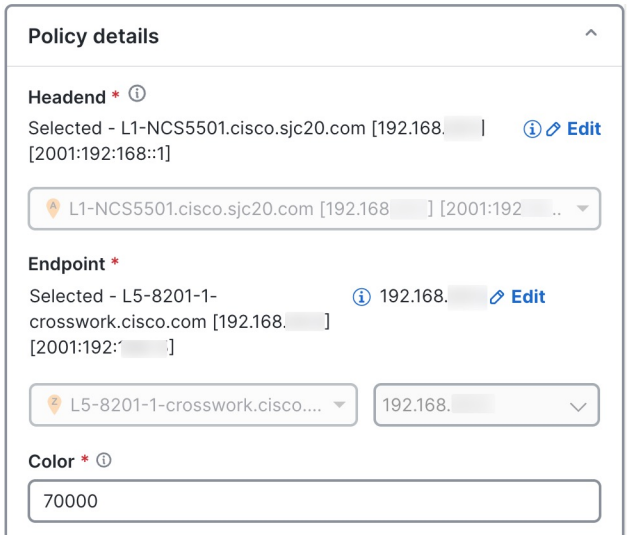

c) In the **Policy path** area, click **Bandwidth on demand**, and enter the required policy path details. In this example, we use these values:

- Path name: **bwod-70000**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **7000 Mbps**

Figure 50: Policy path details

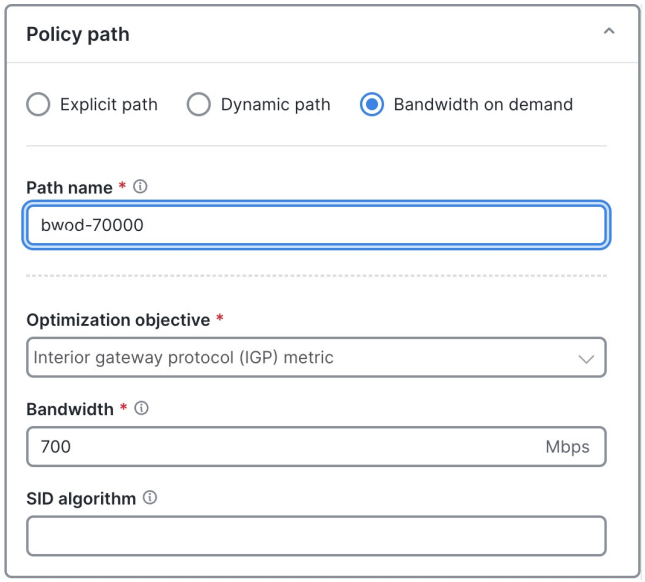

d) Click **Preview** . BWoD only takes into account current interface utilization that has been reserved by another BWoD policy. Otherwise, BWoD only considers the capacity of the interface in its calculations. In this example, all interfaces have the capacity of 1 Gbps. Since there are no existing BWoD policies, BWoD considers the capacity of all nodes and takes the shortest route.

Figure 51: First BWoD policy (bwod-70000)

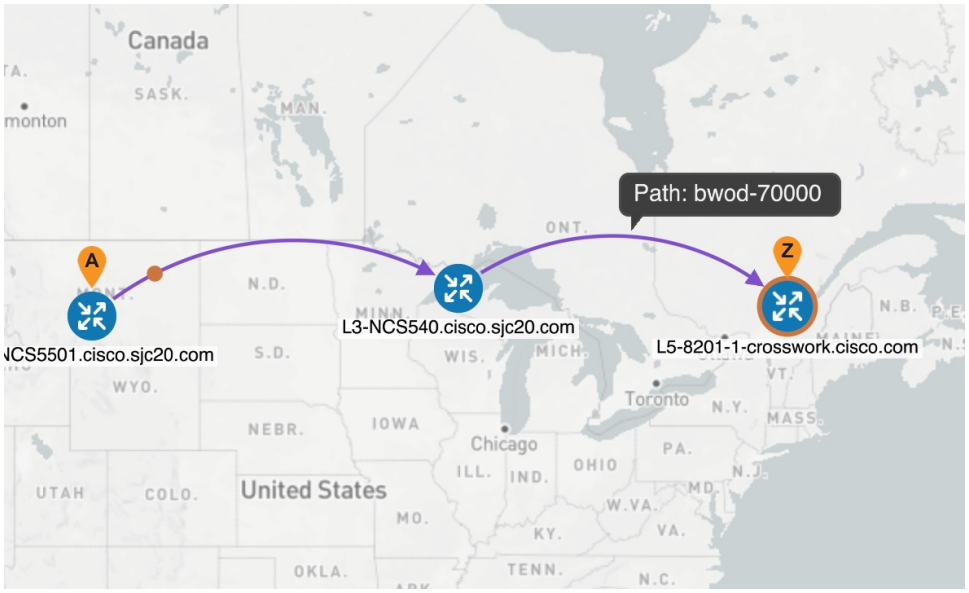

e) If you are satisfied with the proposed BWoD SR-TE policy deployment, click **Provision**.

Step 3 Verify that the new BWoD SR-TE policy has been created.

a) From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS**.

b) Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View details**).

Step 4 Create a second BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70001**
- Path name: **bwod-70001**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **700 Mbps**

BWoD considers the existing BWoD policy (bwod-70000) and its bandwidth requirement into its interface capacity calculations. So, a new path is created for the bwod-70001 policy.

Figure 52: New bwod-70001 policy

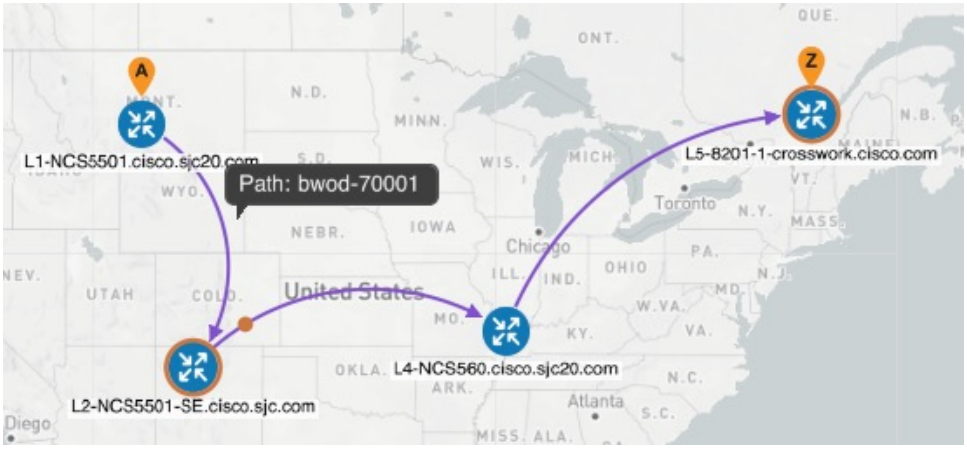

Step 5 Create a third BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70002**
- Path name: **bwod-70002**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **1000 Mbps**

Since BWoD takes into account all previous BWoD policy requirements and the BWoD policy violation option was set to **Loose** , BWoD creates a best effort path for the bwod-70002 policy. You will receive this message when you provision the new policy:

Figure 53: Best effort message

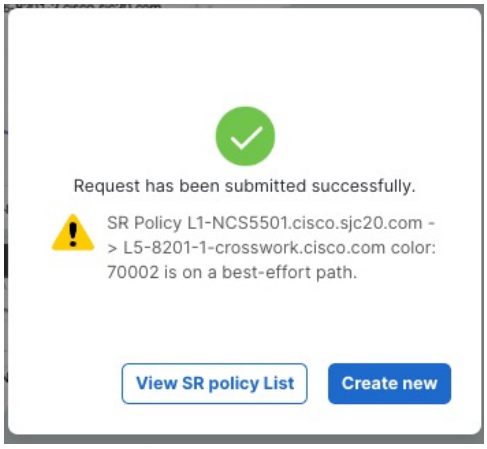

Note that existing paths for bwod-7000 and bwod-70001 are moved to accommodate the new bwod-70002 policy.

Figure 54: BWoD policies with Loose option

Step 6 Change the BWoD policy violation option to **Strict** (**Services & Traffic Engineering** > **Bandwidth on Demand** > **Configuration** > **Advanced**).

Step 7 Create a fourth BWoD policy. In this example, we use these values:

- Headend: **L1-NCS5501.cisco.sjc20.com**
- Endpoint: **L5-8201-1-crosswork.cisco.com**
- Color: **70003**
- Path name: **bwod-70003**
- Optimization objective: **Interior gateway protocol (IGP) metric**
- Bandwidth: **1000 Mbps**

Since the BWoD policy violation option is set to **Strict** , BWoD is not be able to overwrite existing BWoD policies, and the request for additional 1000 Mbps policy results in a "No solution found" message.

Figure 55: No solution found

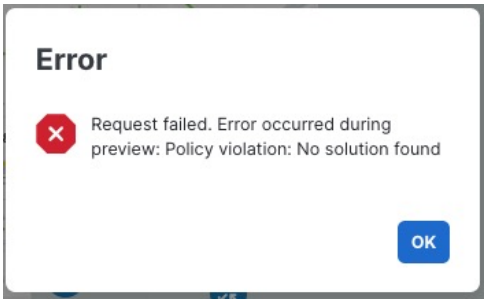

BWoD error messages

Some of the most common error conditions for BWoD and their possible corrective actions are listed below.

Table 9: Error event messages

Error event message	Possible causes and recommended corrective Aation
OptimaModelError	The network model used by BWoD from the Optimization Engine is corrupt or is missing key data needed to properly support BWoD. Possible causes include network discovery issues or synchronization problems between the Optimization Engine and Topology Services. Try restarting the Optimization Engine pod to rebuild the model. This error can also occur if the time required to discover a policy and add it to the model after it has been deployed exceeds the **Deployment Timeout** option set for BWoD. The default is 30 seconds, sufficient for small to medium-sized networks. However, larger networks may require additional time.
NATSTimedOutError	The deployment of a bandwidth policy through SR-PCE exceeds the **Deployment Timeout** option set for BWoD. Increase the **Deployment Timeout** option to allow for additional time for deployments in larger networks.
Traceback or other errors found in the log file	Contact your Cisco service representative.