

Cisco Crosswork Network Controller 7.1 Network Bandwidth Management

First Published: 2025-03-28

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

© 2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **Network Bandwidth Management** 1

Network bandwidth management feature packs 1

CHAPTER 2 **SR Circuit-Style Manager (CSM)** 3

SR Circuit-Style Manager (CSM) 3

Important considerations for circuit-style SR-TE policies in Crosswork Network Controller 4

Access requirements for circuit-style SR-TE policy provisioning 4

Attribute constraints 4

Unsupported configurations 7

Path computation behavior 7

Path reversion 7

Set up CS SR-TE policy visualization workflow 8

Enable SR Circuit-Style Manager 9

Configure circuit-style SR policies 10

View circuit-style SR-TE policy information 12

Trigger CSM to recalculate a circuit-style SR-TE Policy 17

Effects of surpassing bandwidth reservation limits 18

CSM path failure management 22

CHAPTER 3 **Local Congestion Mitigation (LCM)** 25

Local Congestion Mitigation 25

Requirements for LCM congestion evaluation 26

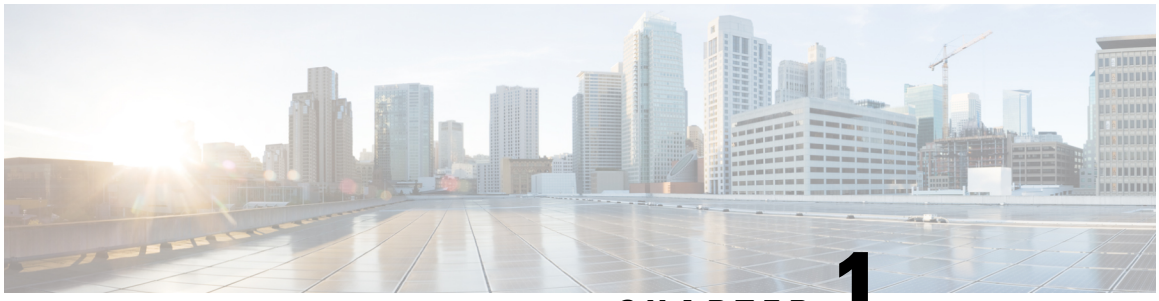
Enable strict SID for LCM usage 26

Requirements for LCM congestion mitigation 27

Important considerations when using LCM 28

BGP-LS speaker placement for multiple AS networks with a dedicated IGP instance between ASBRs 30
LCM calculation workflow 30
Example: Mitigate congestion on local interfaces 32
Configure LCM 38
Additional information on LCM configuration options 38
Configure link affinities 41
Example: Cisco IOS-XR affinity configuration 42
Add affinities in Crosswork Network Controller 42
Add individual interface thresholds 43
Monitor LCM operations 45
Temporarily exclude an interface from LCM 47

CHAPTER 4 **Bandwidth on Demand (BWoD) 49**
Important considerations when using BWoD 49
PCC-initiated BWoD SR-TE policies 50
Provision an SR-TE policy to maintain intent-based bandwidth requirements example 51
Configure Bandwidth on Demand 57
Troubleshoot BWoD 58



CHAPTER 1

Network Bandwidth Management

- Network bandwidth management feature packs, on page 1

Network bandwidth management feature packs

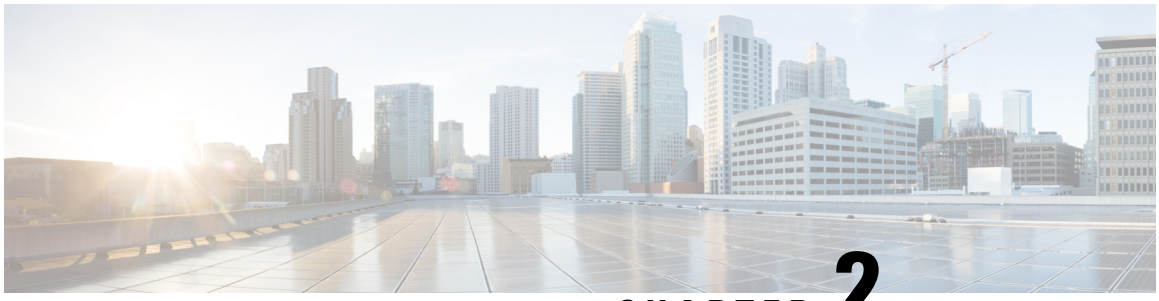
For service providers, managing bandwidth problems used to be a reactive and manual process. The pressure to solve it is huge. Network congestion leads to poor end-customer experiences. Congested links, high latency, and other network impairments lead to a poor perception of the services carried across your network or result in an inability to meet the service level agreements (SLAs) you have with your customers. In the worst-case scenario, your network issues lead to SLA or contract violations and the loss of your brand equity.

Network operators need a toolset to help automate bandwidth optimization, steer traffic with little operator intervention, and ensure that critical links always have sufficient bandwidth to avoid congestion. Cisco Crosswork offers the following feature packs:

- Local Congestion Mitigation (LCM) is a tactical solution for bandwidth management and congestion mitigation. It is best applied when you are attempting to solve congestion issues directly, on the devices themselves, without a full-scale traffic matrix or advanced planning.
- SR Circuit-Style Manager (CSM) is a strategic traffic engineering solution that permits you to reserve bandwidth in advance for critical services, avoiding congestion issues entirely for these high-priority services.
- Bandwidth on Demand (BWoD) is a solution, which provides soft bandwidth guarantee services for SR policies as opposed to strict bandwidth guarantees provided by Circuit Style SR-TE services. Depending on the configuration, BWoD may provide bandwidth reservation, or best-effort bandwidth paths for SR policies. CSM and BWoD feature packs are mutually exclusive. Only one can be enabled at a time.

  Feature pack access requirements

  - Ensure you have the correct licensing package to use feature packs.
  - Users must be assigned admin roles or have certain Device Access Group permissions to access some 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)

- SR Circuit-Style Manager (CSM), on page 3
- Important considerations for circuit-style SR-TE policies in Crosswork Network Controller , on page 4
- Set up CS SR-TE policy visualization workflow, on page 8
- Enable SR Circuit-Style Manager , on page 9
- Configure circuit-style SR policies, on page 10
- View circuit-style SR-TE policy information, on page 12
- Trigger CSM to recalculate a circuit-style SR-TE Policy, on page 17
- Effects of surpassing bandwidth reservation limits, on page 18
- **CSM path failure management**, on page 22

SR Circuit-Style Manager (CSM)

The SR Circuit-Style Manager (CSM) provides a bandwidth-aware Path Computation Element (PCE) for computing circuit-style Segment Routing Traffic Engineering (CS SR-TE) policy paths.

The CSM performs centralized bookkeeping of bandwidth resources across the network and computes paths that meet strict bandwidth requirements while adhering to user-specified constraints such as disjointness and latency minimization. Additionally, the CSM allows users to monitor bandwidth resource levels and identify where resources are running low. It manages and visualizes circuit-style SR-TE policies on the network topology map, ensuring that the best failover bidirectional paths are computed with the requested bandwidth and other constraints

Advantages of using circuit-style SR-TE policies

Circuit-style SR-TE policies:

- are bidirectional, co-routed and designed to carry high-priority or critical services traffic over packet-based networks that require committed bandwidth with protected paths,
- ensure reliability and no impact on service-level agreements (SLAs) due to changing network loads, and
- do not require the maintenance of network states at intermediate routers, nor does it require complex multiprotocols to be implemented.

Important considerations for circuit-style SR-TE policies in Crosswork Network Controller

This section outlines the scope of support for circuit-style SR-TE policies within Crosswork Network Controller. It includes the necessary requirements and constraints on policy attributes for each circuit-style SR-TE policy and describes the processing logic used during path reversions.

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 the "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 Cisco Crosswork Network Controller UI provisioning using Cisco Network Services Orchestration (NSO). To view a device CLI configuration example, see Configure circuit-style SR policies, on page 10.

This table outlines the requirements for each policy attribute and ow 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, and
- state-sync configuration between PCEs of a high availability pair

Path computation behavior

The system 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 <value>: 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 CS SR-TE policy visualization workflow

Complete these steps to view circuit style Segment Routing Traffic Engineering (CS SR-TE) policies on the topology map.

Procedure

Step 1 Enable SR Circuit-Style Manager , on page 9

Step 2 Configure circuit-style SR policies, on page 10 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.

Enable SR Circuit-Style Manager

To manage and visualize circuit-style SR-TE policies on the topology map, you must first enable the SR Circuit-Style Manager (CSM) and set bandwidth reservation 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

Note Both CSM and BWoD cannot be enabled at the same time. If BWoD is enabled, you must disable it before enabling CSM. 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 **Traffic Engineering** > **Circuit Style SR-TE** > **Configuration**.

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

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 18.

	Field	Description
Basic	Link CS BW Pool Size	The percentage of each link's bandwidth reserved for circuit-style SR-TE policies.
	Link CS BW Min Threshold	The Link CS BW Pool utilization percentage beyond which a threshold crossing event notification will be generated.
Advance	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.
	Timeout	The duration until which the CSM will wait for the delegation request, to generate a notification.
	Restore Delegation Delay	The duration until which the CSM will pause before processing a restore path delegation.

Step 4 To save the configuration, click **Commit Changes**.

What to do next

Configure circuit-style SR policy configurations either manually on the device (see Configure circuit-style SR policies, on page 10) or through Cisco 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 in Crosswork Network Controller , 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.

This section provides *guidance* on how to manually configure a circuit-style SR policy and a PM profile on a device.

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 in Crosswork Network Controller , 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
     !
    !

   !
  !

 !
!
```

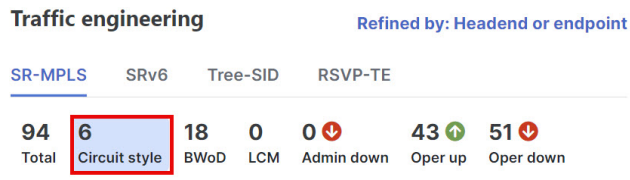
View circuit-style SR-TE policy information

Complete these steps to view circuit-style SR-TE these policy details such as the endpoints, bandwidth constraints, IGP metrics, and candidate (Working and Protect) paths.

Procedure

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

Figure 1: Select Circuit style tab

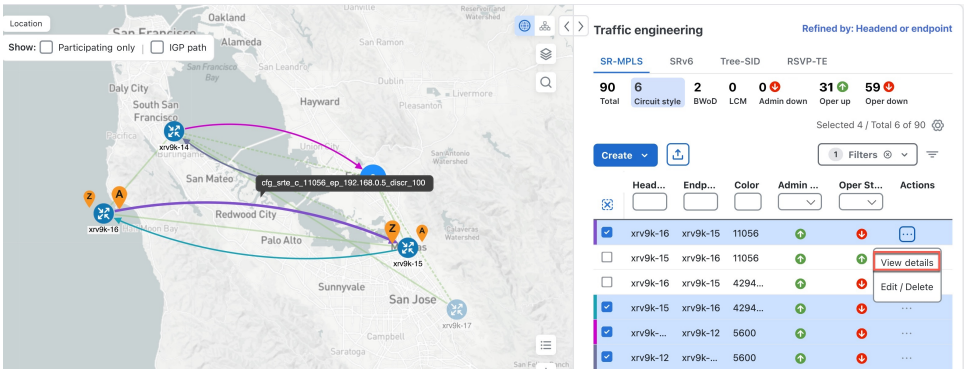


The table lists all circuit-style SR-TE policies.

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

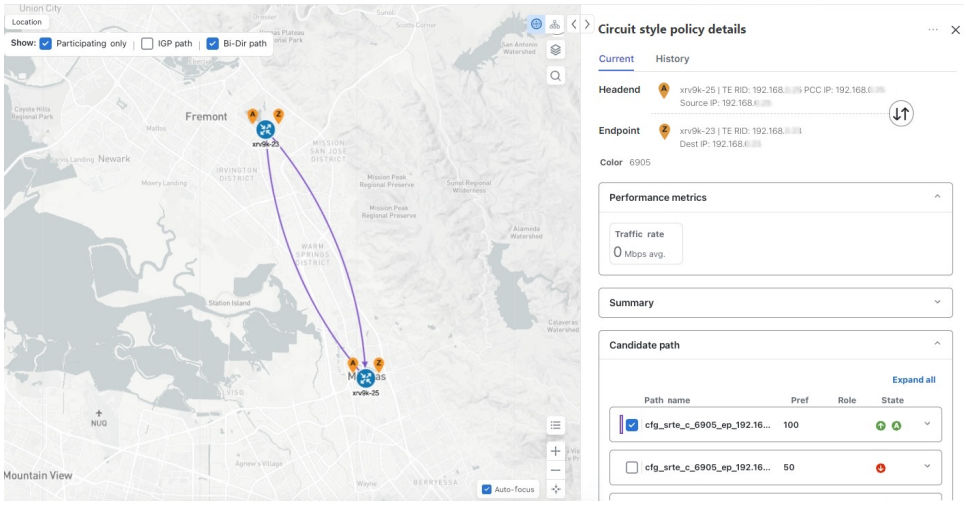
Note
You cannot edit or remove circuit-style SR-TE policy configurations that have been created directly on the device.

Figure 2: 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 3: CS-SR policy details summary



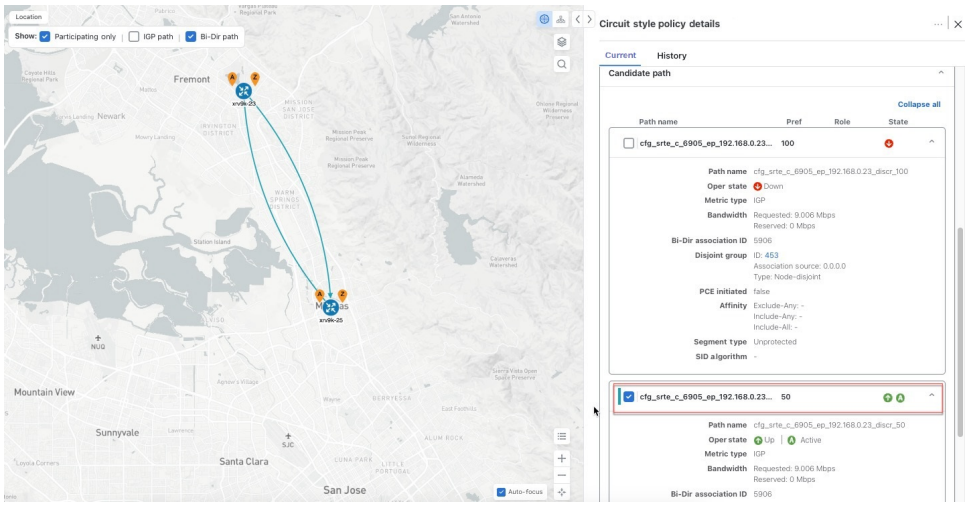
Note
The Bandwidth Constraint value can differ from the bandwidth you requested if the value is increased and insufficient resources exist to satisfy demand on all Working and Protect candidate paths.

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 22). 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 4: Candidate path on topology map

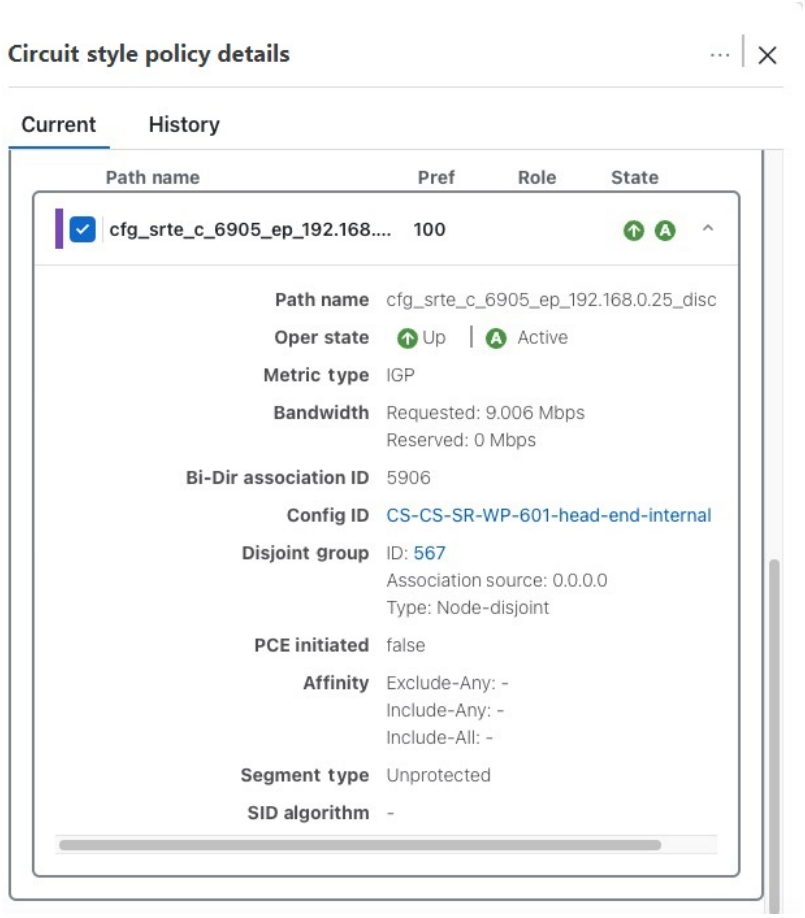


Note
- First preference paths are shown as purple links.

- 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 Cisco 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**. For example:

Figure 5: View Candidate path details

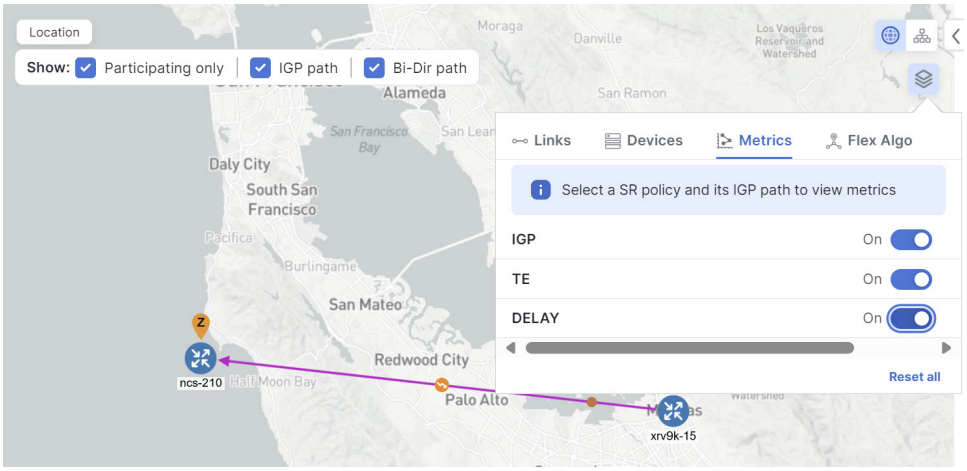


Here is a sample of a circuit-style policy configuration. See Configure circuit-style SR policies, on page 10.

Figure 6: 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 7: 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 CSM to recalculate a CS SR-TE policy.

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering** > **Traffic Engineering** > **SR-MPLS** and click **Circuit style**. 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. CSM maintains an accounting of all bandwidth reservations provided for CS SR policies to ensure that the total reserved bandwidth on all interfaces remains at or below the network-wide resource pool (bandwidth pool size).

This topic provides examples of how CSM handles policies that exceed the bandwidth pool size or bandwidth alarm threshold set on the CSM Configuration page.

- Example: Bandwidth utilization surpasses defined threshold, on page 18
- Example: Bandwidth pool size and utilization exceeded, on page 20

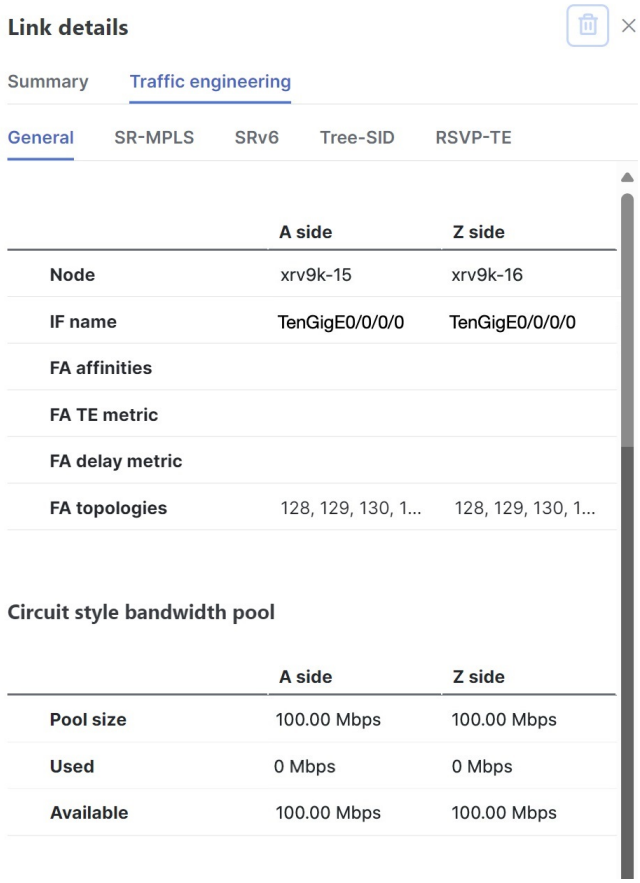
Example: Bandwidth utilization surpasses defined threshold

- 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).

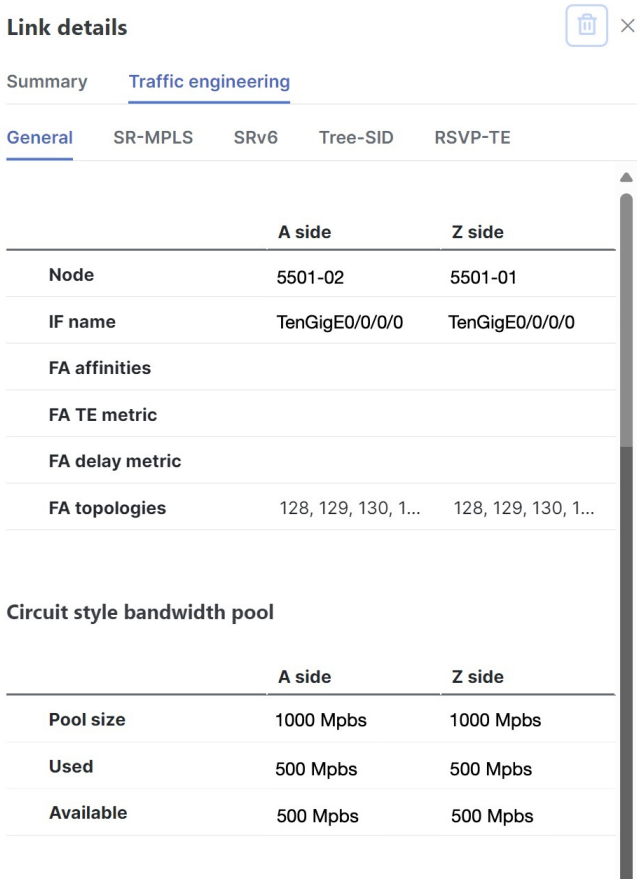
1. A circuit-style SR-TE policy from node 5501-02 to node 5501-01 (r02 - r01) is created with a bandwidth of 100 Mbps.

Figure 8: 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 9: 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 10: Threshold alerts



Example: Bandwidth pool size and utilization exceeded

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

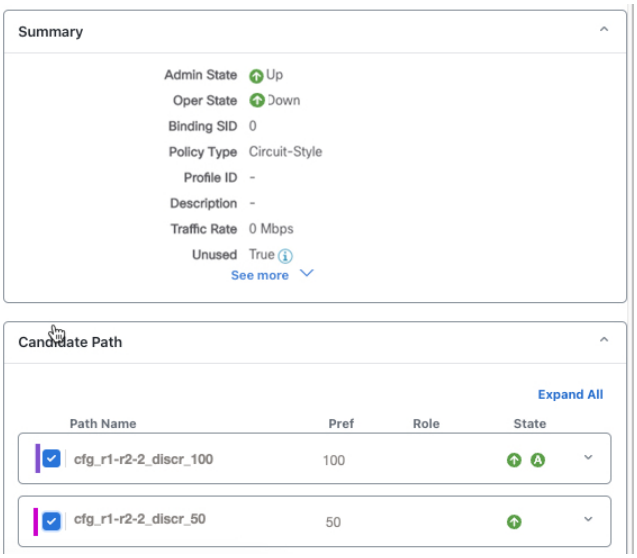
In this example, the bandwidth pool size for the 10 Gbps ethernet interfaces is 1Gbs and the alarm threshold is set for 900 Mbps.

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. The only paths available between these two nodes are the paths computed for the first CS policy.

   - CSM cannot compute a path for the new circuit-style SR-TE policy r02 - r01 - r2 and remains operationally down. CSM will try again every 30 minutes to find a path that meets the bandwidth requirements.

     Figure 11: CS-SR policy exceeds bandwidth pool size



   - Alerts are triggered.

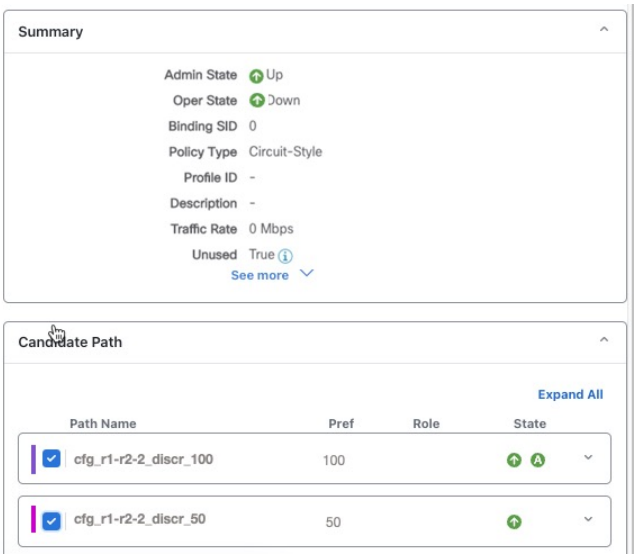
     Figure 12: Threshold alerts



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 13: 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 14: Cleared alerts



CSM path failure management

Cisco Crosswork 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 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 Circuit-Style Manager , on page 9). This delay allows topology and policy state changes to fully propagate to Crosswork in cases where these changes triggered the Restore path delegation.

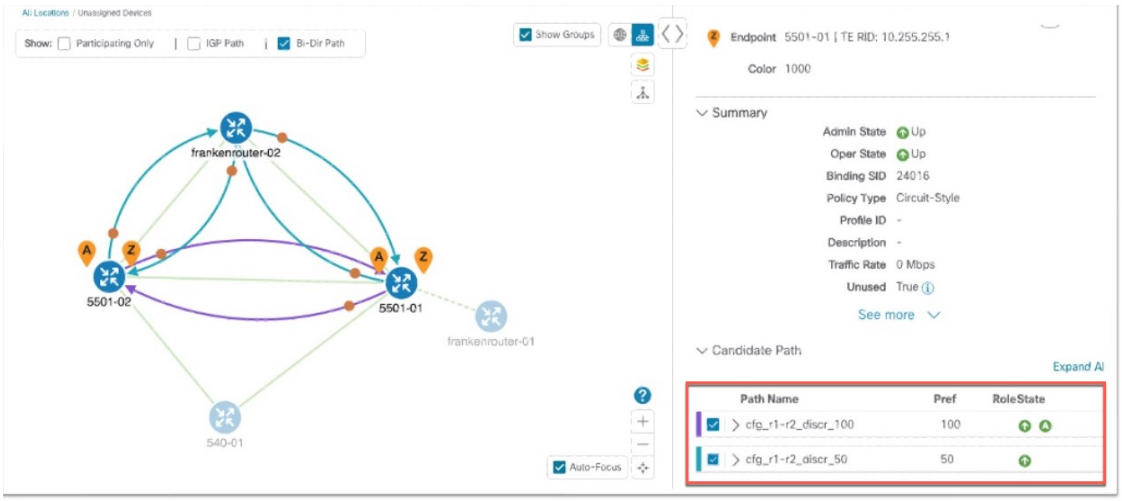
You can configure Performance Measurement (PM) to address path failures effectively and switch from the Working path to the Protect path. For more information, see Configure circuit-style SR policies, on page 10.

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.

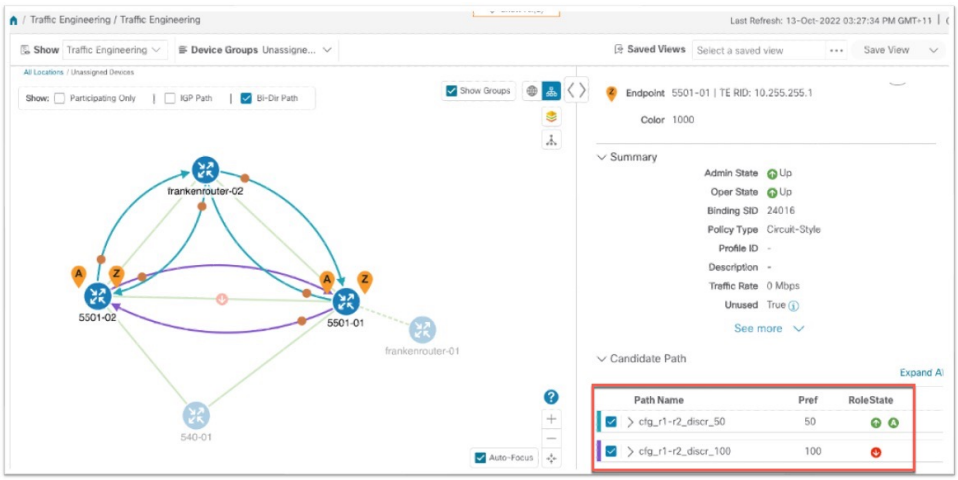
The following image shows that the Working and Protect paths of the circuit-style SR-TE policy are operational. The *active* path is indicated by the "A" icon.

Figure 15: 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.

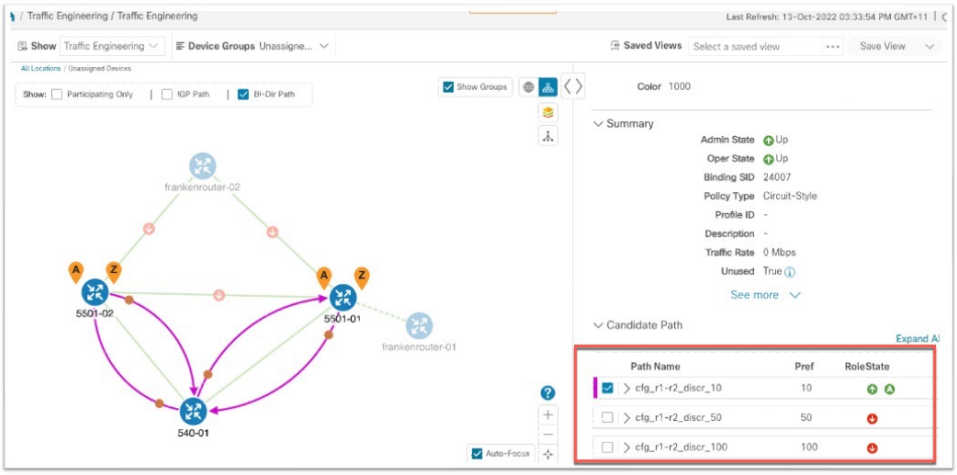
Figure 16: Protected Path Becomes Active

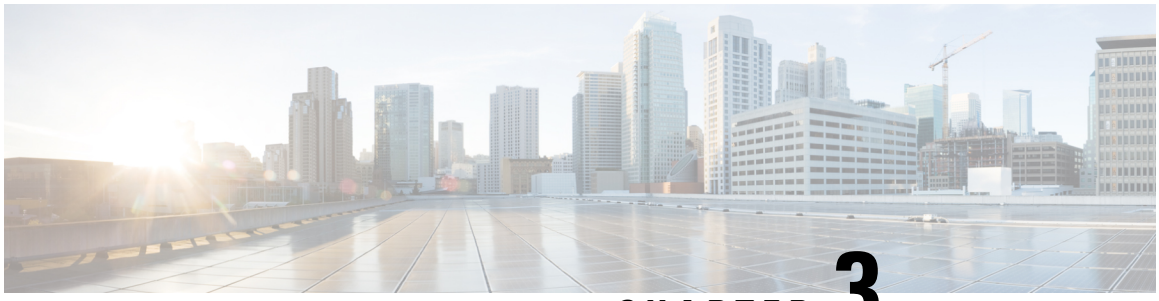


When both the Working and Protect paths go down, CSM calculates a Restore path, which becomes active. The Restore path only appears in this specific scenario. Note that the Restore path has the lowest preference

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 17: Restore Path





CHAPTER 3

Local Congestion Mitigation (LCM)

- Local Congestion Mitigation, on page 25
- Requirements for LCM congestion evaluation , on page 26
- Requirements for LCM congestion mitigation , on page 27
- Important considerations when using LCM, on page 28
- LCM calculation workflow, on page 30
- Example: Mitigate congestion on local interfaces, on page 32
- Configure LCM, on page 38
- Add individual interface thresholds, on page 43
- Monitor LCM operations, on page 45
- Temporarily exclude an interface from LCM, on page 47

Local Congestion Mitigation

Local Congestion Mitigation (LCM) searches for congestion on a configurable cadence (as opposed to a triggered event). It provides localized mitigation recommendations in surrounding interfaces (local interface-level optimization) within a domain. LCM computes the shortest paths for one or more tactical policies to divert minimal traffic on a congested interface to alternate paths with sufficient bandwidth. It attempts to keep as much of the traffic on the original IGP path. With LCM, you can do the following:

- Monitor congestion as defined by the interface thresholds you specify.
- Visually preview LCM recommendations on your network before you decide whether to commit the Tactical Traffic Engineering (TTE) SR policy deployment.
- Enable LCM to delete any down, failed, or uncommitted LCM TTE policies when there is an imminent risk of network failures based on LCM solution configurations. For more information, see the advanced configuration options (**Auto Repair Solution** and **Adjacency Hop Type**) in Configure LCM, on page 38.

LCM enables a wider application of the solution in different network topologies, including those with multiple IGP areas, because of its simpler path computation and limitation to specific network elements. Focusing on the problem locally within a domain eliminates the need to simulate edge-to-edge traffic flows in the network through a full traffic matrix. It allows for better scaling of large networks. Also, LCM collects TTE SR policy and interface counters via SNMP and does not require SR-TM.

Note Take a look at the Example: Mitigate congestion on local interfaces, on page 32 to see how to use LCM in your network.

Requirements for LCM congestion evaluation

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,

- 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
- confirm that all the devices are reachable from the Crosswork Data Gateway, and
- enable strict SID labels on all devices within an LCM domain. See Enable strict SID for LCM usage, on page 26.

Enable strict SID for LCM usage

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

Note Provisioning SR policies with `sid-algorithm 1` (strict SPF SIDs) is not supported in the Crosswork Network Controller UI. Additionally, ensure that non-LCM SR policies do not use SIDs with algorithm 0 in their segment lists.

Procedure

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

Example:

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
  !
!
```

Example:

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
  !
```

Example:

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 on the headend device and confirm that ***all*** devices

- are either using the same default segment routing global block (SRGB) range or a custom range you specify, and
- have the maximum SID depth explicitly configured ***if*** there are devices along the path that impose restrictions on the label stack depth.

Example:

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

Requirements for LCM congestion mitigation

For LCM to correctly calculate and mitigate congestion, the headend device must support autoroute steering and Equal Cost Multi-Path (ECMP) .

Autoroute steering

The headend device must support PCE-initiated SR-TE policies with autoroute steering. However, LCM will not work if the headend is a Cisco NCS device and there is L2VPN traffic in the network.

Devices should be configured with **`include ipv4 all`** and **`force-sr-include`** to enable traffic steering into SR-TE policies with autoroute.

For example:

```
segment-routing
traffic-eng
```

```
 pcc
  profile 10      !!The profile ID must match the value in the UI LCM Configuration > Basic
> Profile ID
    autoroute
     include ipv4 all
     force-sr-include
```

The *ID* parameter in this command identifies the PCC profile associated with the SR-TE policy that PCE has provisioned. The ID value can be any integer from 1 to 65535, but it must match the profile ID that PCE uses to instantiate the policy. If not, the policy will not be activated. For example, if 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. See the specific device platform configuration guide, for example, *Segment Routing Configuration Guide, Cisco IOS XE 17 (Cisco ASR 920 Series)*.

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

Equal Cost Multi-Path (ECMP)

The headend device must support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies. To verify that a device can support SR-TE policies using ECMP, check that the device has the following:

- Segment Routing is enabled and configured with a Segment Routing Global Block (SRGB) that matches the SRGB of the SR-TE policy headend and tailend routers. Use the `show segment-routing mpls state` command to verify the SRGB configuration on the device.
- BGP-LS is enabled and configured to advertise and receive link-state information from the SR-TE policy headend and tailend routers. Use the `show bgp link-state link-state` command to verify the BGP-LS status and the `show bgp link-state link-state database` command to verify the link-state information on the device.
- ECMP is enabled and configured to load-balance traffic across multiple equal-cost paths based on flows. Use the `show ip route` command to verify the ECMP routes and the `show ip cef` command to verify the ECMP load-balancing algorithm on the device.

Important considerations when using LCM

Consider the following information when using LCM:

- User roles must be granted with LCM task permissions for a domain in order to confgure LCM and commit LCM recommendations. For more information on RBAC and user roles, see the "Cisco Crosswork Network Controller Administration Guide".
- Device Access Group (DAG) access is ***not*** supported by LCM. Users that have been granted with LCM task permissions in a domain are able to configure and commit LCM recommendations regardless of whether or not they have DAG access for any devices in that domain.
- LCM does not support LDP-labeled traffic. LDP-labeled traffic *must not* be steered into LCM autoroute TTE SR policies.

- The use of LCM is not recommended on networks with Tree SID policies. Initial calculations are skewed because full traffic measurements are unavailable.
- LCM supports domains with up to 3000 devices. A *domain* is an identifier that is assigned to an IGP process. Domains are learned from the network. The domain ID is taken from the PCC router configuration (`link-state instance-id`) that you use to advertise IGP with BGP-LS.
- LCM recommended solutions use the resources within a single domain only.
- LCM evaluates network utilization on a regular, configurable cadence of 1 minute or more. The cadence is typically set to be greater than or equal to the SNMP traffic polling interval but can be set lower to improve responsiveness. The default cadence is 10 minutes.
- The traffic statistics collection interval affects how quickly LCM can respond to topology changes and LSP deployments that affect interface and LSP traffic measurements. It can take up to twice the traffic statistics collection interval plus the LCM evaluation interval for LCM recommendations to fully reflect these changes. During this period, LCM recommendations may evolve as the traffic measurements are updated and eventually fully converge in Crosswork.
- LCM leverages ECMP across parallel TTE SR policies and assumes roughly equal splitting of traffic. The degree to which actual ECMP splitting adheres to this assumption depends on the presence of large elephant flows and the level of traffic aggregation.

  You can configure LCM to detect excessive uneven ECMP splitting among parallel TTE SR policies and issue an event to notify. To mitigate the effects of uneven ECMP, the overprovisioning factor is used in LCM. For more information, see Configure LCM.
- LCM assumes traffic in an *existing* SR-TE policy is ineligible for optimization and should not be steered into LCM TTE SR policies. To enforce this assumption, any existing non-LCM SR-TE policies should 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.
- When domain interfaces and links are removed (intentionally or unintentionally), the following occurs:
  - If all links in a domain go down (LINK_DOWN state), LCM configuration and the Domain UI card (see Configure LCM, on page 38) will remain available until the links are aged out (after 4 hours). This behavior is intentional as it gives you time to recover domain interfaces and links if this was done by mistake.
  - If you want to force domain removal before links age out, then you can remove links manually from the UI. The domain will remain in a "ready for deletion" status until the last link is removed.
- After an HA switchover, you can manually add missing interfaces that were previously monitored or update domain configuration options after the system is stable. Missing interfaces and other configuration options occur if they were added after the last cluster data synchronization.
- When an SR-PCE goes down, **Local Congestion Mitigation** (LCM) enters a dormant stage. To exit this state, all SR-PCEs must be connected, and their associated topologies fully synchronized with the topology service. LCM will remain dormant until these conditions are met. It is important to note that LCM does not have visibility into the state of the SR-PCE redundancy set.

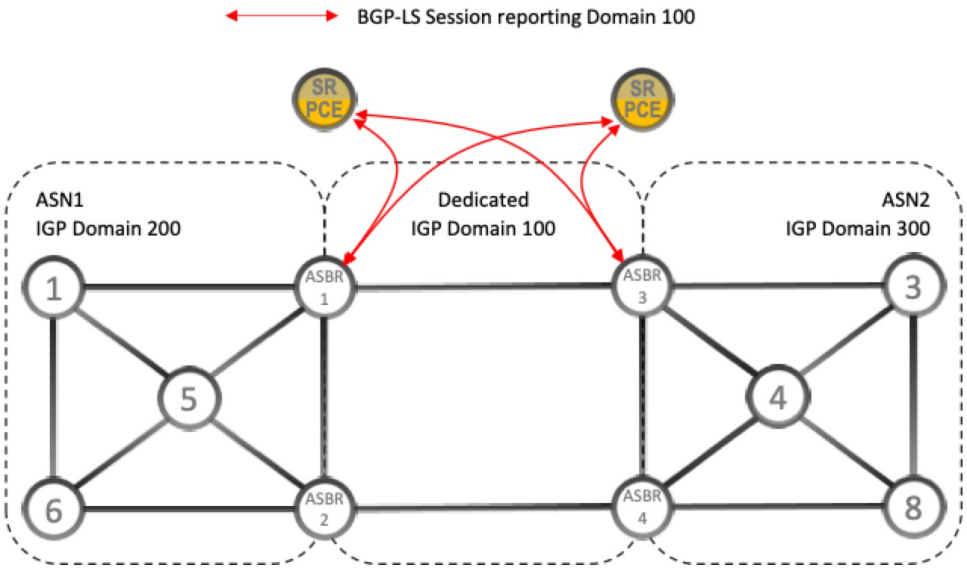
BGP-LS speaker placement for multiple AS networks with a dedicated IGP instance between ASBRs

To support interdomain latency-optimized SR policy path computation by an SR-PCE (or other use cases where egress peer engineering (EPE) is not supported), a dedicated IGP instance may be configured between autonomous system border routers (ASBRs) in different ASNs. In these cases, it is important to identify which ASBRs report the topology via BGP-LS for proper topology discovery.

In the following example, at least one ASBR in each AS participating in the dedicated inter-AS IGP (Domain 100) must have BGP-LS enabled to report the IGP between each ASBR. Each ASBR must report the domain with the same BGP-LS identifier.

Note More than one ASBR per AS reporting the BGP-LS topology is also supported.

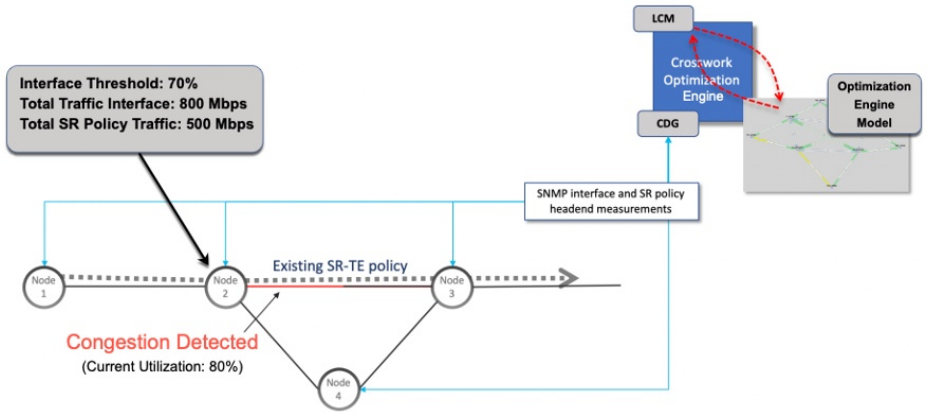
Figure 18: BGP-LS session reporting domain 100



LCM calculation workflow

This example walks you from congestion detection to the calculations LCM performs before recommending tactical tunnel deployment. These calculations are done on a per-domain basis, allowing better scalability and faster calculation for larger networks.

Figure 19: LCM Configuration Workflow Example



Procedure

Step 1 LCM first analyzes the Optimization Engine Model (a realtime topology and traffic representation of the physical network) on a regular cadence.

Step 2 In this example, LCM detects congestion after a congestion check interval when Node 2 utilization goes above the 70% utilization threshold.

Step 3 LCM calculates how much traffic is eligible to divert.

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 will not be included in the LCM calculation and will continue to travel over the original programmed path.

Eligible traffic is computed by taking the interface traffic statistics that account for all traffic on the interface and subtracting the sum of traffic statistics for all SR-TE policies that flow over the interface.

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.

The total traffic that LCM can divert in this example is 300 Mbps: 800 Mbps – 500 Mbps = 300 Mbps

Step 4 LCM calculates the amount that must be sent over alternate paths by subtracting the threshold equivalent traffic from the total traffic on the interface. In this example, the amount to be diverted is 100Mbps:

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

LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see Configure LCM, on page 38.

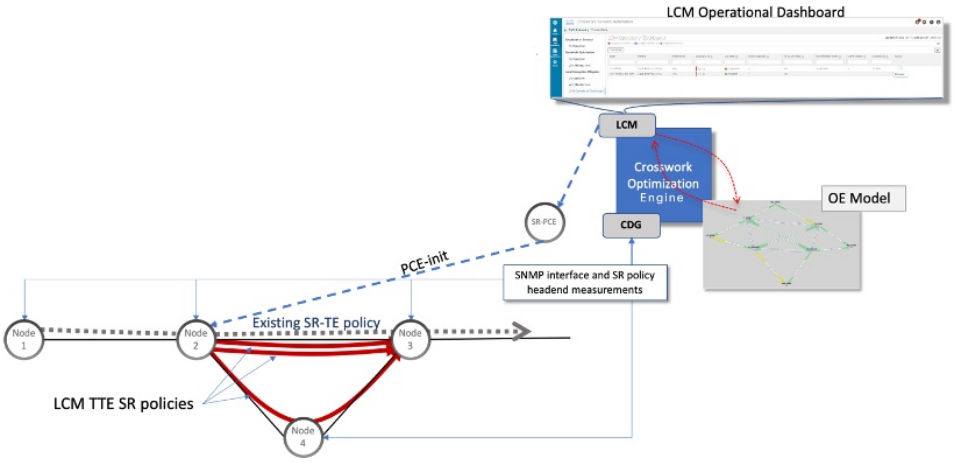
Step 5 LCM determines how many TTE SR policies are needed and their paths. The ratio of how much LCM-eligible traffic can stay on the shortest path to the amount that must be detoured will determine the number of TTE SR policies needed on the shortest versus alternate paths, respectively.

In this example, LCM must divert one-third of the total eligible traffic (100Mbps out of 300Mbps) away from the congested link. Assuming a perfect ECMP, LCM estimates three tactical SR-TE policies are required to create 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 three TTE SR policies (each expected to route approximately 100Mbps) to be deployed 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 will be listed in the **LCM Operational Dashboard**.

Figure 20: LCM Recommendation Example



Step 6 Assuming you deploy these TTE SR policies, LCM continues to monitor the deployed TTE policies and will recommend modifications or deletions as needed in the LCM **Operational Dashboard**. LCM recommends deleting deployed TTE SR policies if the mitigated interface will remain uncongested after they are removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Example: Mitigate congestion on local interfaces

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

In this example, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE SR policies when utilization on a device's interface surpasses a defined utilization threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion. The example goes through the following steps:

1. View uncongested topology.
2. Set utilization thresholds for individual interfaces.

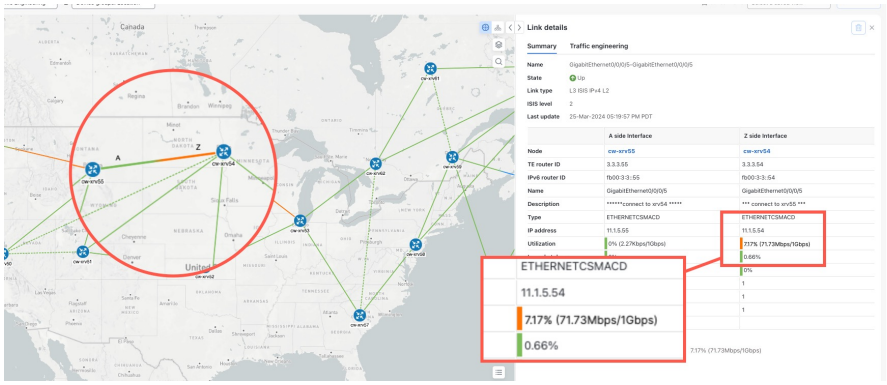
3. Enable and configure LCM in manual mode. Manual mode allows you to view recommended TTE policies prior and decide whether or not to deploy them.
4. After LCM detects congestion, view LCM recommendations on the Operational dashboard.
5. Preview the recommended LCM TTE policies to deploy visually on the topology map.
6. Commit and deploy all LCM TTE policy recommendations to mitigate the congestion.
7. Verify that the LCM TTE policies have been deployed.

Procedure

Step 1 View initial topology and utilization prior to LCM configuration.

a) In this example, note that the node cw-xrv54 has a utilization of 7.17%.

Figure 21: Initial utilization



Step 2 Define any individual interface thresholds.

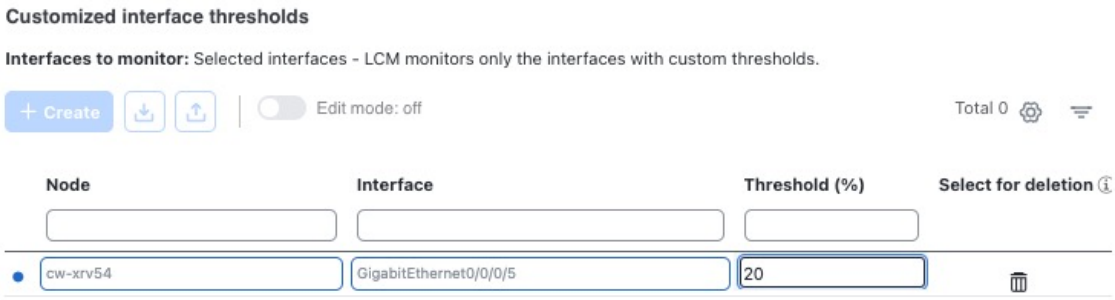
LCM allows you to configure a *global* utilization threshold that can be used for all interfaces. When traffic utilization surpasses the threshold, LCM will try to find bypass policies to remediate the congestion. You set the global utilization threshold on the **LCM Configuration** page. However, if 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 (**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 43.

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 22: Customized interface thresholds



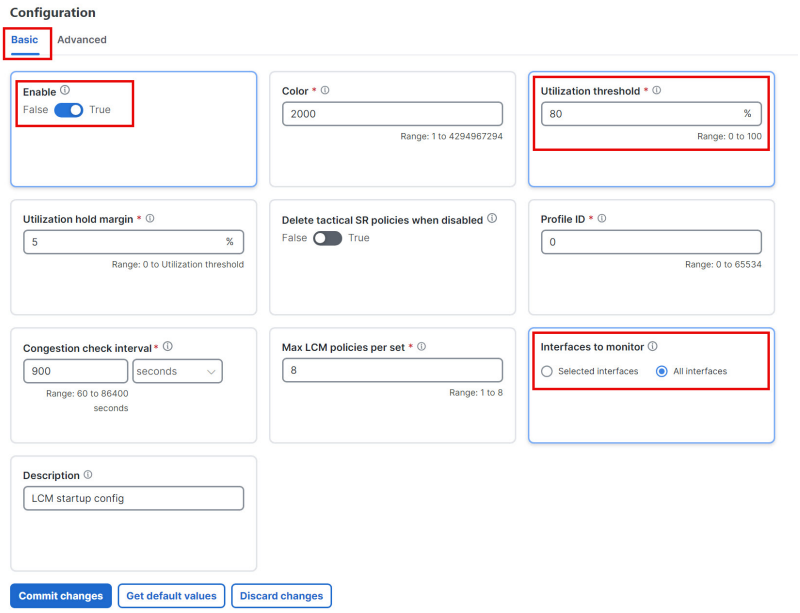
Note
By default, LCM monitors all interfaces. This includes any individual thresholds that are imported to this page. The rest of the interfaces will be monitored using the global **Utilization threshold** defined on the **LCM Configuration** page.

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 **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.

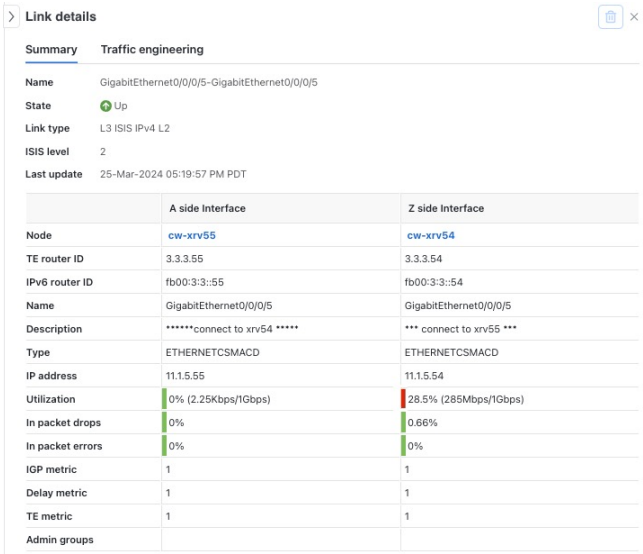
Figure 23: LCM Configuration page



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

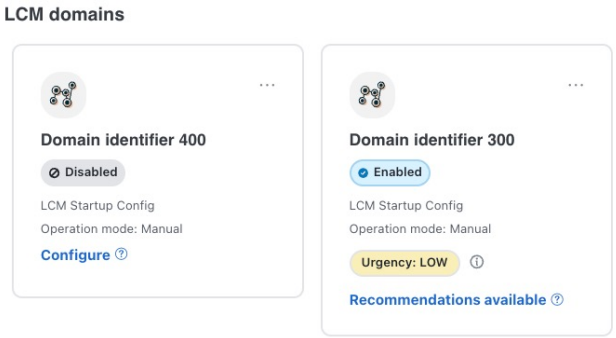
Figure 24: Observed congestion



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

a) Navigate to **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 25: Congested detected and LCM recommendations



b) (Optional) View LCM events.

From the top-right corner of the Crosswork 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** (**Traffic Engineering > Local Congestion Mitigation >** ***Domain-Identifier*** **>** > **Operational dashboard**).

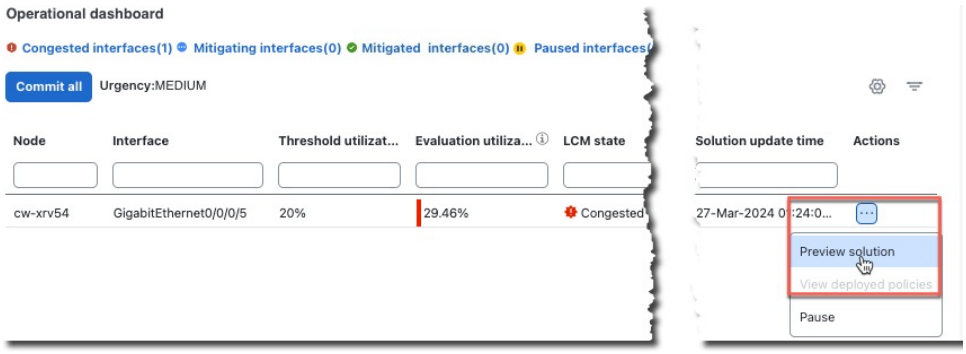
The dashboard shows that cw-xrv54 utilization has surpassed 20% and is now at 29.46%. In the **Recommended action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended action - Create set**) to address the congestion on the interface. For more information, see Monitor LCM operations, on page 45.

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 38).

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**.

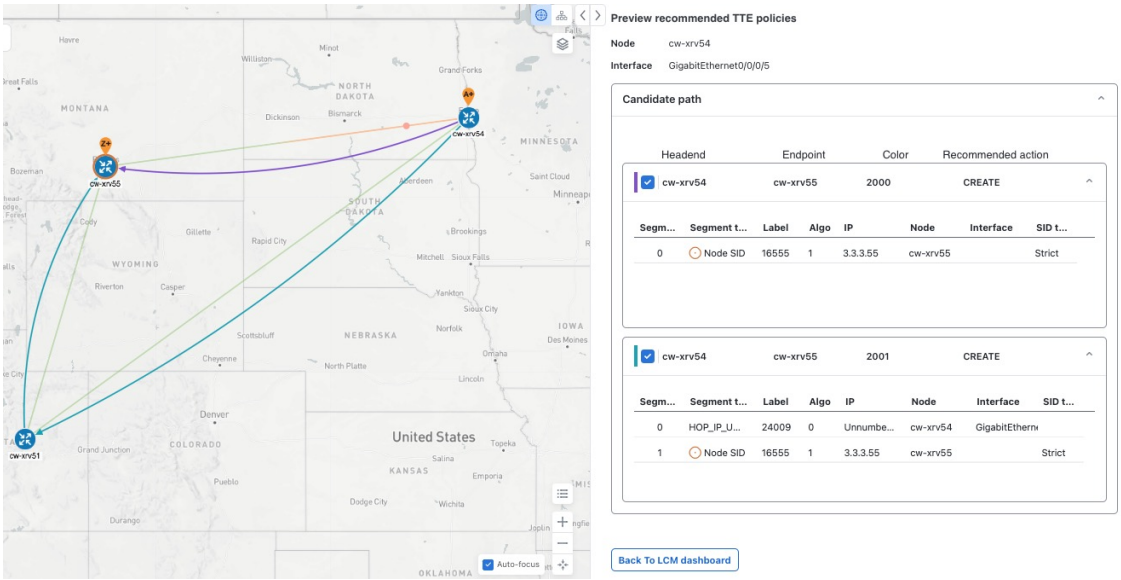
Figure 26: Preview solution



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.

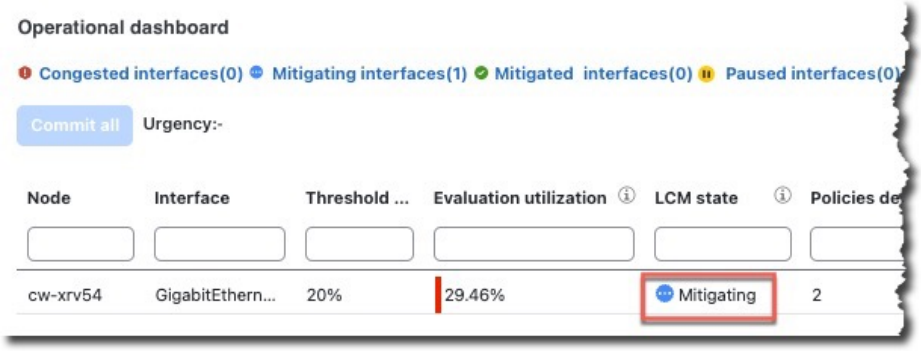
The following figure shows the recommended TTE policies for node cw-xrv54.

Figure 27: 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 28: 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 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 to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands, but at the same time, gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations and see how the potential deployment will take effect in your network before you deploy them. As traffic changes,

LCM tracks the deployed TTE SR-TE policies and decides whether or not they are still needed. If not, LCM recommends deleting them.

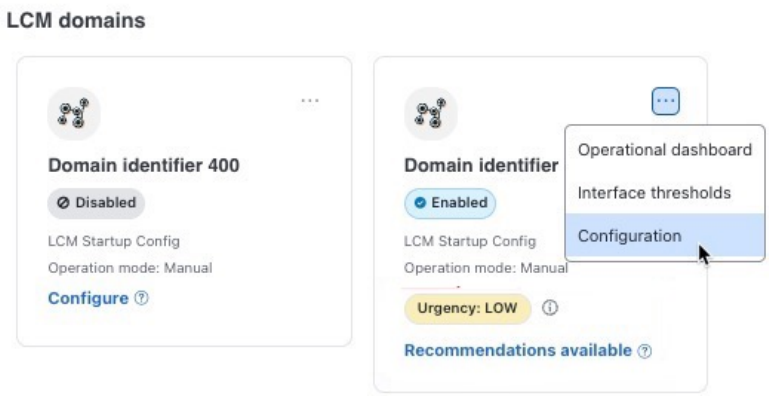
Configure LCM

To enable and configure LCM:

Procedure

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

Figure 29: LCM Configuration



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

Step 3 Enter the required information. Hover the mouse pointer over ⓘ to view a description of each field. See Additional information on LCM configuration options, on page 38.

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

Additional information on LCM configuration options

These tables provide additional information that are not described in the UI.

- Basic LCM configuration

- Advanced LCM configuration

Basic LCM options

Table 3: Basic configuration options

Option	Description
Utilization threshold	Set 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.
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)	This value 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.
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 (**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.

Advanced LCM options

Table 4: Advanced configuration options

Option	Description
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.
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.
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.
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.
Over-provisioning factor (OPF)	This option helps address unequal ECMP traffic distribution (elephant flows). This value determines the percentage of extra traffic that should be accounted for when computing a path for a by-pass policy. If LCM needs to divert *i* amount of traffic due to congestion, then it will search for a path that can support *x* * (1 + OPF) traffic. For more information, see LCM calculation workflow, on page 30. The default value is 0.

Option	Description
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 42 and Configure link affinities, on page 41.

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 42.

Step 2 Add affinities in Crosswork Network Controller, on page 42.

Step 3 Configure LCM, on page 38 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 in Crosswork Network Controller, the affinity name is displayed as "UNKNOWN" in the UI. Follow these steps to add affinities in Crosswork Network Controller:

Before you begin

Configure and note down the affinities on your devices.

Procedure

Step 1 From the main menu, choose **Administration** > **System 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 bit it will be assigned. For example:

Figure 30: Affinity

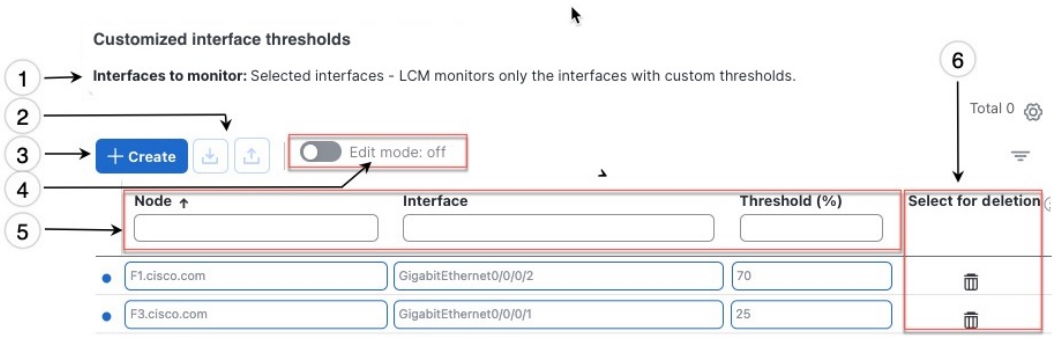


Step 4 Click **Save**. To create another mapping, you must click **+ Create** and save the entry.

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 31: Customized interface thresholds



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

Callout No.	Description
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 when using LCM, do the following:

Procedure

Step 1 From the main menu, choose **Services & Traffic Engineering > Local Congestion Mitigation> *Domain-identifier*** > ··· > **Interface thresholds** and click one of the following:

- **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 32: 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.

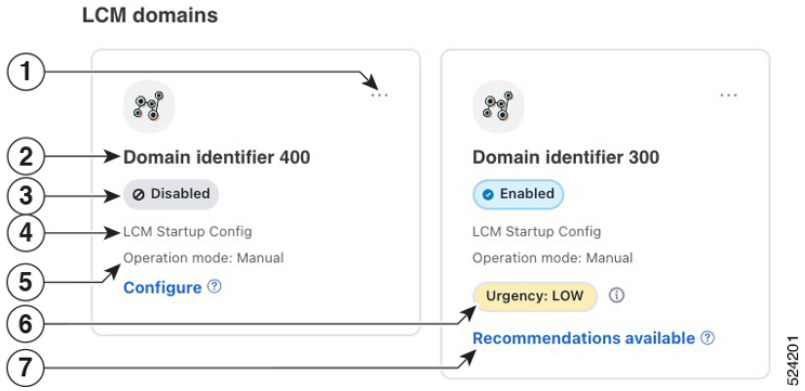
Monitor LCM operations

Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the Example: Mitigate congestion on local interfaces, on page 32 topic.

LCM Domains Dashboard

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

Figure 33: LCM Domains Dashboard



Callout No.	Description
1	**Main Menu**: Allows you to navigate to the following pages: • Operational Dashboard • Interface Thresholds • Configuration
2	**Domain identifier**: The domain ID is taken from the router configuration (`link-state instance-id`) that you use to advertise IGP with BGP-LS.
3	**LCM status**: Indicates whether LCM has been enabled for the domain or 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: Manual**—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.

Callout No.	Description
6	**Urgency**: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following: • **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 38.
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 (**Traffic Engineering** > **Local Congestion Mitigation** > ***Domain-ID*** > ··· > **Operational Dashboard**) shows congested interfaces as defined by the configured utilization threshold.

Each interface lists details such as current utilization, recommended action, status, expected utilization after committing recommendations, and so on. Hover the mouse pointer over ⓘ to view a description of the type of information each column provides.

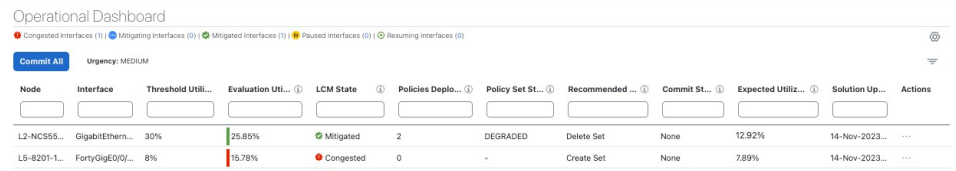
From the Actions column, you can do the following:

- Preview TTE policies prior to deployment (··· > **Preview Solution**)
- Verify deployment (··· > **View Deployed Policies**)
- Pause or resume an interface (··· > **Resume / Pause**)

To gain a better understanding of what information the LCM Operational Dashboard provides, see the following example:

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

Figure 34: LCM Operational Dashboard



In this example, the following information is conveyed:

- The first row is an interface that is currently in a Mitigated state. It shows that two policies have been deployed (**Policies Deployed - 2**) to mitigate a previous congestion. However, the current recommendation (**Recommended Action - Delete Set**) is to delete the policies since they are no longer needed (congestion should not occur even if the previously deployed policies are removed). Since the current recommendation has not been committed, the current Commit Status is None.
- The second row is an interface that is currently in a Congested state. LCM detects congestion and suggests to deploy policies to remediate the congestion (**Recommended Action - Create Set**). You can choose to preview the solution (⋯ > **Preview Solution**).

Note If LCM cannot find a solution (**Recommended Action - No Solution**), it may be due to constraints enabled on the **LCM Configuration** page. For more information, see Configure LCM, on page 38.

Recommendations are listed as part of a set, and if deployed, all changes are committed. You must click **Commit All**.

Temporarily exclude an interface from LCM

You can temporarily pause LCM from including an interface for mitigations. 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 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:

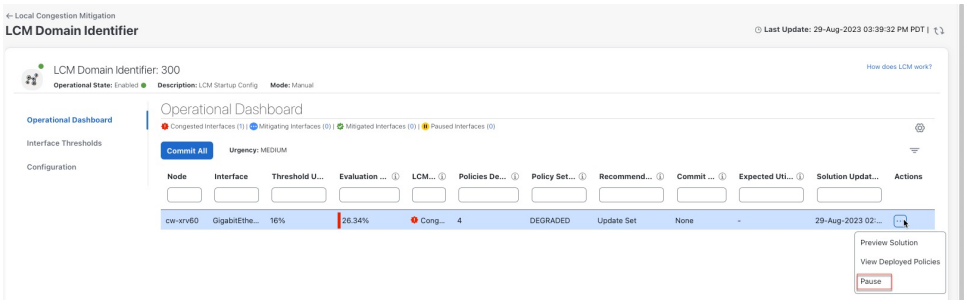
- No LCM SR policy traffic
- Excessive imbalance in LCM policy traffic
- Excessive LCM oscillations or removals per hour

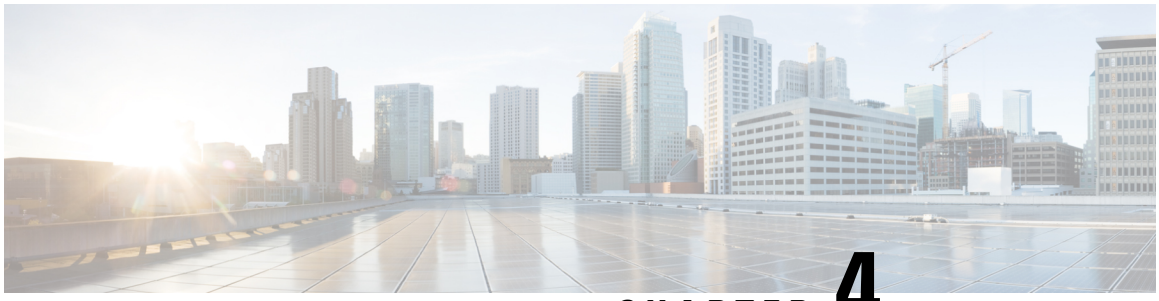
In these circumstances, the user may perform a corrective action, and manually resume the interface.

From the Actions column of the LCM Operational Dashboard, click ··· > **Pause** for the interface you would like to exclude from LCM calculations. To include the interface in LCM calculations again, click ··· > **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 35: Pause interface





CHAPTER 4

Bandwidth on Demand (BWoD)

Bandwidth on Demand (BWoD) provides a bandwidth-aware Path Computation Element (PCE) in conjunction with SR-PCE for segment routing policies (SR policies). BWoD policies can be PCC-initiated (PCE-delegated) or PCE-initiated. BWoD is designed to deliver soft bandwidth guarantee services over SR policies. BWoD monitors network conditions and re-optimizes BWoD paths to prevent total BWoD traffic on any interface from exceeding the configured threshold percent.

BWoD does not track total interface utilization, and therefore, interfaces can still be congested if the combined BWoD traffic and non-BWoD traffic exceed the interface capacity. In addition, BWoD does not enforce the total amount of traffic entering BWoD SR policy. BWoD policies may traverse Equal Cost Multi-Path (ECMP) and assume even traffic distribution over these paths. However, actual ECMP distribution can be uneven, especially with large flows.

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

This section contains the following topics:

- Important considerations when using BWoD, on page 49
- PCC-initiated BWoD SR-TE policies, on page 50
- Provision an SR-TE policy to maintain intent-based bandwidth requirements example, on page 51
- Configure Bandwidth on Demand, on page 57
- **Troubleshoot BWoD**, on page 58

Important considerations when using BWoD

Consider the following information when using BWoD:

- To provision a BWoD policy, you must have write-access to the head-end device based on Device Access Groups and assigned 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 will disable itself when an unexpected error is encountered to avoid network disruption.

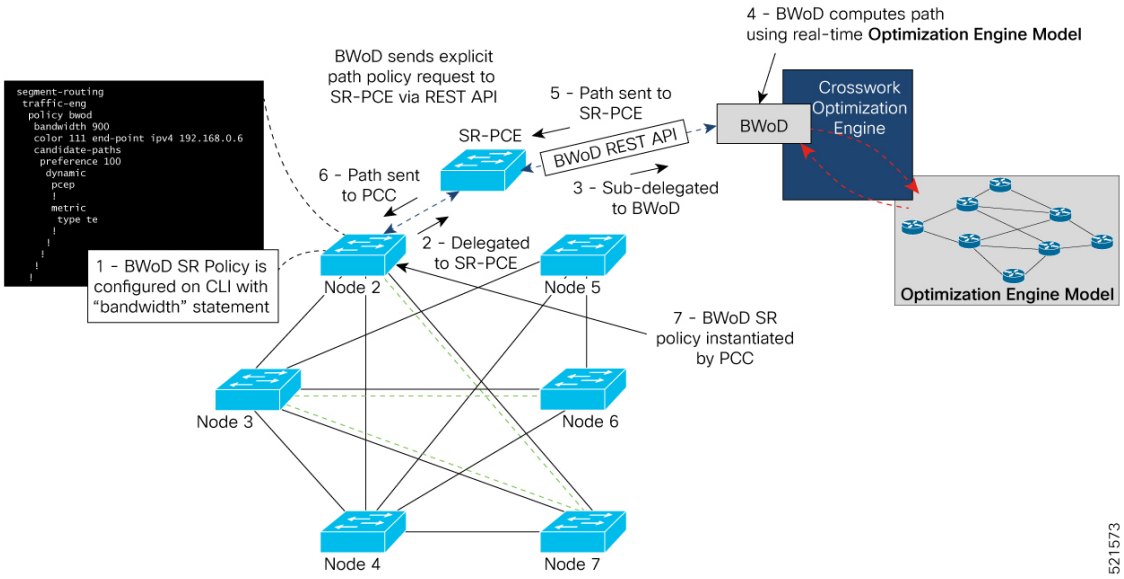
- BWoD temporarily pauses operation whenever the Optimization Engine model is unavailable due to an Optimization Engine restart or a rebuild of the topology from Topology Services. Any requests to BWoD during this time are rejected. When the model becomes available and BWoD receives two traffic updates 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 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

When enabled, BWoD automatically connects to all SR-PCE providers configured in Crosswork Network Controller. A persistent connection is made to the SR-PCE BWoD Rest API, which is registered as a PCE for bandwidth-constrained SR-TE policies.

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

Figure 36: PCC-Initiated BWoD SR-TE Policies



Callout No.	Description

1	A BWoD policy is configured on a PCC via the CLI. 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	The bandwidth statement is added to a PCE-delegated SR policy to create a BWoD policy. Once committed, the PCC delegates the path compute to SR-PCE.
3, 4	SR-PCE then sub-delegates the policy to BWoD, which attempts to compute a path that meets the bandwidth constraint.
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.

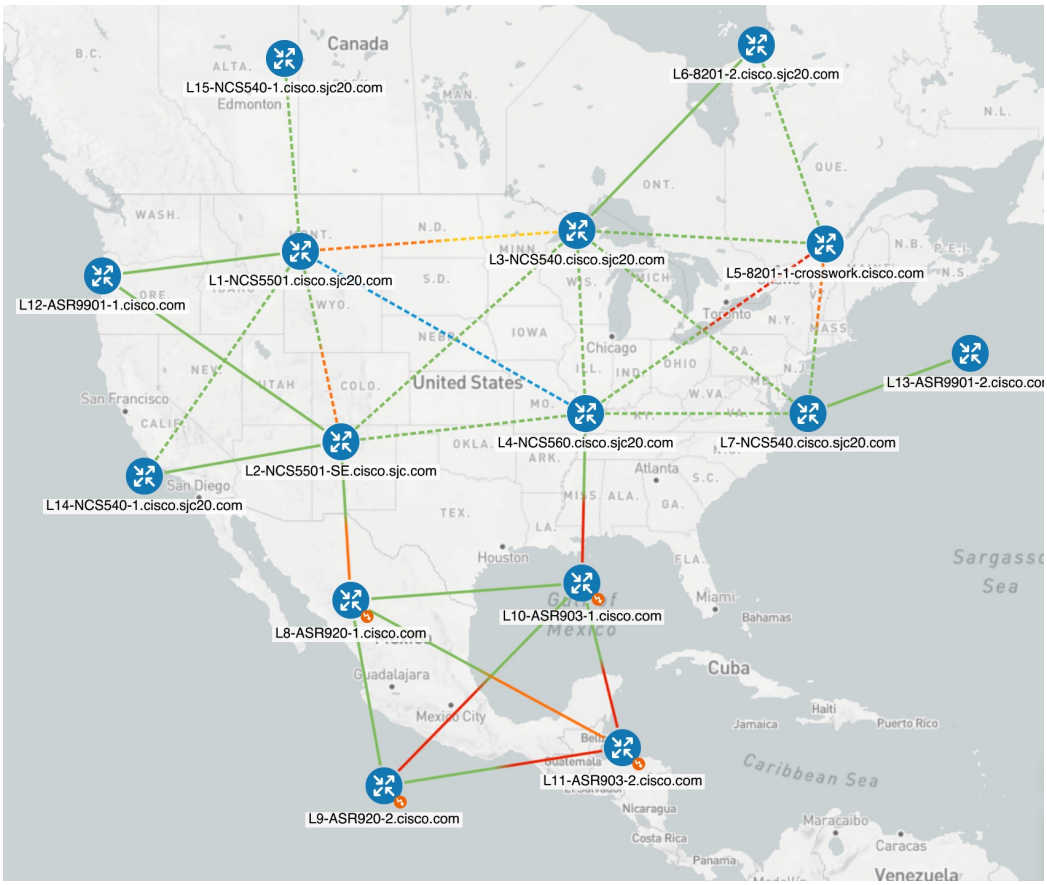
Provision an SR-TE policy to maintain intent-based bandwidth requirements example

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 particular, three BWoD policies will be created using the ***same*** headend (L1-NCS5501.cisco.sjc20.com) and endpoint (L5-8201-1-crosswork.cisco.com) with a bandwidth requirement of 700 and 1000 Mbps, while keeping the utilization at 80%. In this example, all interfaces have the capacity of **1 Gbps**.

Figure 37: Initial BWoD topology



Procedure

Step 1 Enable and configure BWoD.

Note
If you have CSM enabled, you must disable it before enabling BWoD. It is recommended to delete all CSM policies from the network before disabling the CSM function pack. If CSM policies remain, this may cause issues with new policy delegation and increase processing time for function packs.

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.

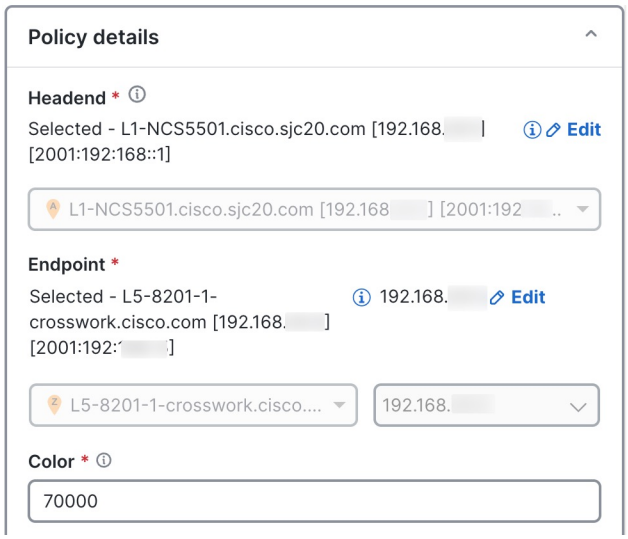
a) From the main menu, choose **Traffic Engineering** > **SR-TE** tab and click **Create > PCE init**.

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**

Example:

Figure 38: Policy details

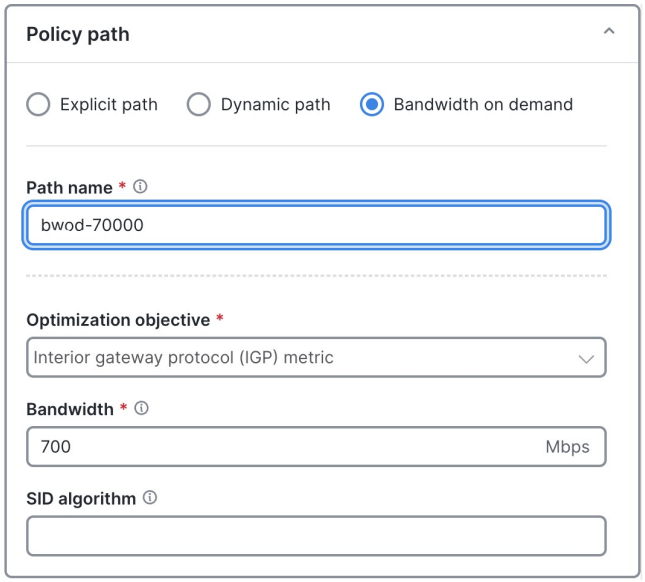


c) In the **Policy path** area, click **Bandwidth on demand**, and enter the required policy path details. In this example, we use these values:

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

Example:

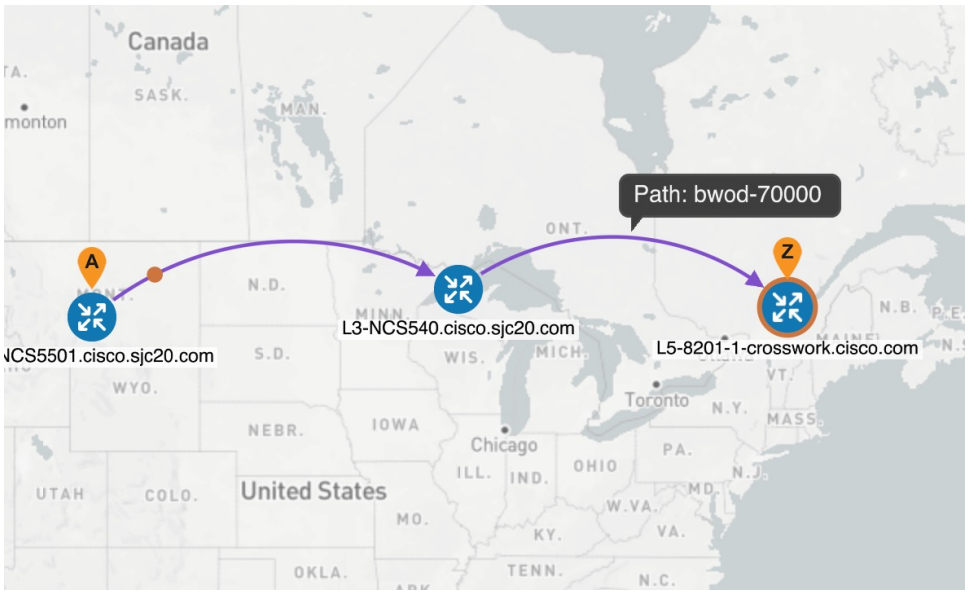
Figure 39: Policy path details



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.

Example:

Figure 40: 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 **Traffic Engineering** > **SR-TE**.

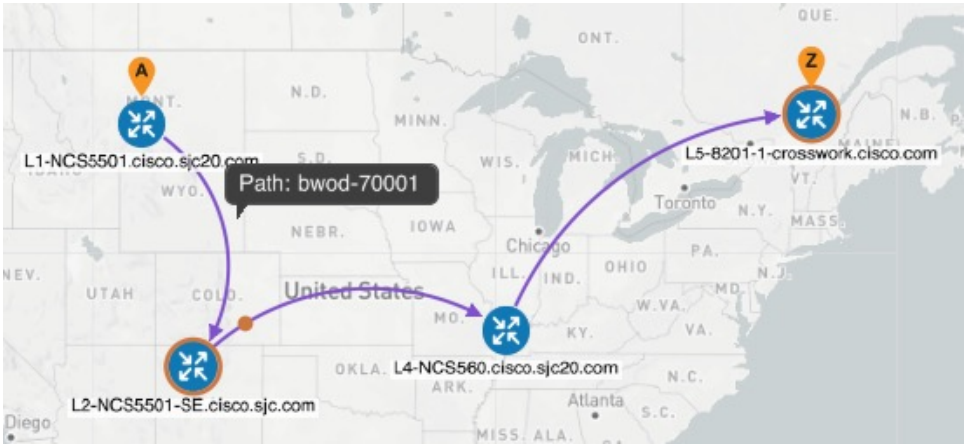
b) Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View details**).

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 41: New bwod-70001 policy

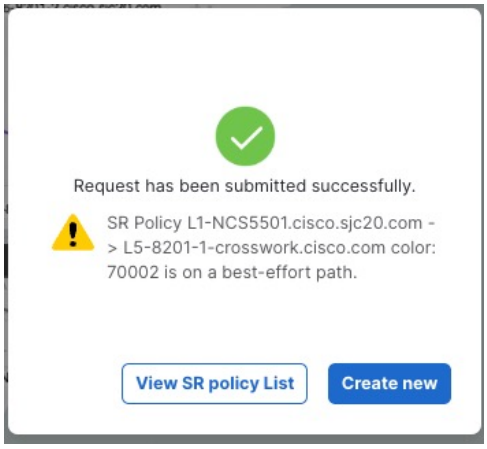


Step 5 Create a third BWoD policy. In this example, we use these values:

- 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 42: Best effort message



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

Example:

Figure 43: BWoD policies with Loose option



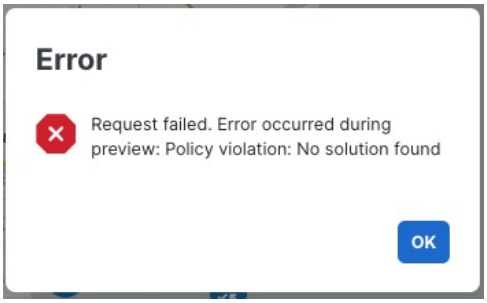
Step 6 Change the BWoD policy violation option to **Strict** (**Services & Traffic Engineering** > **Bandwidth on Demand** > **Configuration** > **Advanced**).

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 44: No solution found



Configure Bandwidth on Demand

There are two main steps in using Bandwidth on Demand (BWoD):

1. Enable and configure BWoD options.
2. Create BWoD SR policies. As long as BWoD is enabled, you can create multiple BWoD SR policies.

Before you begin

Both CSM and BWoD cannot be enabled at the same time. You must disable one of these function packs before enabling the other.

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

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 additional options. Hover the mouse pointer over to view a description of each field.

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

Step 5 To create BWoD SR policies, choose **Traffic Engineering** > **Traffic Engineering**.

Step 6 From the SR Policy table, choose **Create** > **PCE Init**.

Step 7 In addition to entering the required SR policy details, click the **Bandwidth on demand** option and enter the required bandwidth.

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

Step 9 Click **Preview** to view the proposed SR policy.

Step 10 Click **Provision** to commit the SR policy.

Troubleshoot BWoD

The following are some of the most common error conditions for BWoD and some possible corrective actions that may fix the issue.

Table 5: Errors

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	Please contact your Cisco service representative.