



Cisco Crosswork Network Controller 6.0 Network Bandwidth Management

First Published: 2023-11-20

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



CHAPTER 1

Overview

- [Overview, on page 1](#)

Overview

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.
- Circuit-Style Segment Routing (CS-SR) 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.



Note

- 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](#)".
 - CS-SR and BWoD feature packs are mutually exclusive. Only one can be enabled at a time.
-



CHAPTER 2

SR Circuit Style Manager (CSM)

The SR Circuit Style Manager (CSM) feature pack provides a bandwidth-aware Path Computation Element (PCE) to compute Circuit Style SR-TE policy paths that you can visualize in your network. Circuit Style enables segment routed transport tailored for circuit-oriented services over a packet based network through the use of bi-directional, co-routed, path protected SR-TE policies. Circuit Style SR-TE policies are typically used for high priority services, such as crucial monetary transactions or important live video feed, which *require committed bandwidth with fast and fail-safe connections*. The CSM feature pack ensures dynamic Circuit Style SR-TE policies are provisioned along paths that meet strict bandwidth requirements while at the same time respecting any additional user configured constraints such as latency minimization and disjointness.

Centralized bandwidth accounting in the CSM feature pack allows the user to monitor resource reservation levels and quickly identify hot spots where available bandwidth in the circuit style bandwidth pool is low. The ability to visualize Circuit Style SR-TE policies in your network topology enables easy verification of Circuit Style SR-TE policy configurations, details, and path states. With a few clicks you can view Active and Protect paths, operational status, reserved bandwidth pool size and monitor path failover behavior for individual Circuit Style SR-TE policies.



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

This section contains the following topics:

- [Circuit Style SR-TE Important Notes, on page 4](#)
- [Workflow for Setting Up CS SR-TE Policy Visualization, on page 9](#)
- [Enable SR Circuit Style Manager, on page 10](#)
- [Configure Circuit Style SR Policies, on page 11](#)
- [Review Circuit Style SR-TE Policy Bandwidth Utilization , on page 13](#)
- [View Circuit Style SR-TE Policies, on page 14](#)
- [Trigger CSM to Recalculate a Circuit Style SR-TE Policy, on page 18](#)
- [What Happens When Bandwidth Reservation Settings are Exceeded?, on page 18](#)
- [How Does CSM Handle Path Failures?, on page 22](#)

Circuit Style SR-TE Important Notes

This topic outlines the scope of Crosswork's support for Circuit Style SR-TE policies, including requirements and constraints on the policy attribute values set in each Circuit Style SR-TE policy, and the processing logic followed during path reversions.



Note Role-based Access Control (RBAC) and task permissions have been introduced in this release. 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 RBAC and user roles, see the "[Cisco Crosswork Network Controller Administration Guide](#)".

Policy Attribute Constraints

You set policy attribute values when you create a Circuit Style SR-TE policy, using either the device's command line interface or Cisco Crosswork Network Controller. You can also change them later.

The table below describes the requirements for each attribute, and how changes affect them. It is important to understand that all the attributes that are described in the table below act as constraints. Each of them corresponds to elements of the configuration that Cisco Crosswork uses to govern how Circuit-Style path hops are computed. Each value is effectively a path computation or optimization constraint, since they either specify a required property of a path or exclude possible choices for that path.

Table 1: Circuit Style SR-TE Policy Attribute Values and 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	<p>The bandwidth constraint is required and must be the same on both sides of a Circuit Style SR-TE policy. Bandwidth changes can be made to existing policies, with these effects:</p> <ul style="list-style-type: none"> • Once you configure the new bandwidth on both sides, Crosswork will evaluate the path. This <u>will not</u> result in a recomputed path • If the new bandwidth is higher, Crosswork checks the existing path to ensure sufficient resources. If all currently delegated paths can accommodate the new bandwidth, Crosswork returns the same path with the new bandwidth value, indicating to the path computation client (PCC) that it was successful. If any of the current paths cannot accommodate the new bandwidth, it returns the old bandwidth value indicating that it was unsuccessful. This evaluation will not be retried unless the bandwidth is changed again. • If the bandwidth is lower, Crosswork returns the same path with the new bandwidth value to indicate to the PCC that it was successful. <p>The user interface shows both the requested and reserved bandwidth under each candidate path when you view the policy details. These values can differ if the requested bandwidth is increased but there is insufficient available CS pool bandwidth along one or more of the paths.</p>
Candidate Paths and Roles	<p>The Working path is defined as the highest preference Candidate Path (CP).</p> <p>The Protect path is defined as the CP of the second highest preference.</p> <p>The Restore path is defined as the lowest preference CP. The headend must have <code>backup-ineligible</code> configured.</p> <p>CPs of the same role in each direction must have the same CP preference.</p>
Bi-Directional	<p>All paths must be configured as co-routed.</p> <p>Paths of the same role on both sides must have the same globally unique bi-directional association ID.</p>

Attribute	Description
Disjointness	<p>Working and Protect paths on the same PCC must be configured with a disjointness constraint using the same disjoint association ID and disjointness type.</p> <p>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.</p> <p>Only the Node and Link disjoint types are supported. The disjoint type used must be the same in both directions of the same policy.</p> <p>The Restore path must not have a disjointness constraint set.</p> <p>Crosswork 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, Crosswork makes no automatic attempt to compute a less restrictive link type disjoint path.</p>
Metric Type	<p>Only the TE, IGP, Hop count, and Latency metric types are supported. The metric type used must match across Working, Protect and Restore paths in both directions.</p>
Segment Constraints	<p>All Working, Protect and Restore paths must have the following segment constraints:</p> <ul style="list-style-type: none"> • protection unprotected-only • adjacency-sid-only <p>To ensure persistency through link failures, configure static adjacency SIDs on all interfaces that might be used by Circuit Style policies.</p>
Unsupported Configurations	<p>The following configurations are not supported:</p> <ul style="list-style-type: none"> • Metric-bounds • SID-Algo constraints • Partial recovery is not supported 7.8.x. • State-sync configuration between PCEs of a high-availability pair. These are not required with Circuit Style SR-TE policies. Use of this feature may result in degraded performance. • Multiple Circuit Style SR-TE policies between the same nodes with the same color but different endpoint IPs.

Attribute	Description
Supported Policy Changes	<p>The following constraints may be changed for an operationally "up" Circuit Style SR-TE policy that has been previously delegated:</p> <ul style="list-style-type: none"> • Metric type • Disjoint type • MSD • Affinities <p>Once configuration changes are made in a consistent manner across all CPs and both PCCs (for example: the new metric type is the same for all CPs and both sides), Crosswork will initiate a recompute, which can result in new Working, Protect and Restore paths.</p> <p>During any transitory period in which configurations are not in sync between paths on the same PCC or between PCCs, no path updates are sent to the PCCs.</p>
Unsupported Policy Changes	<p>The following configuration changes to a previously delegated and operationally "up" Circuit Style SR-TE policy are not supported:</p> <ul style="list-style-type: none"> • CP preference • Disjoint Association ID • Bi-directional Association ID <p>To change these configurations for an existing policy, you must first shut down the policy on both sides, make the change (complying with restrictions as detailed above in terms of consistency) and then "no shut" the policy.</p>
Path Computation	<p>Crosswork 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. In cases where there is insufficient bandwidth and a path cannot be found, SR Circuit Style Manager will continue to retry after 30 minutes until a solution is found, or if Circuit Style SR-TE is disabled.</p> <p>Crosswork computes the Restore path only after the Working and Protect paths are down. The SR Circuit Style Manager feature pack configuration interface provides a configurable delay timer to control how long after Restore paths are delegated from both sides to wait before the path is computed. This delay allows topology and SR policy state changes to fully propagate to Crosswork, in cases where these changes triggered the Restore path delegation.</p> <p>Automatic re-optimization is not supported for any paths based on changes in topology, LSP state, or any periodic event. Path computation is supported for Intra/Inter area/level and Intra/Inter IGP Domain (same AS) Not supported path computation Inter-AS.</p>

Attribute	Description
Reversion Behavior	<p>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):</p> <ul style="list-style-type: none"> • 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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 1. Working path recovers and goes to up/active state 2. Restore path is removed 3. Protect path recovers and goes to up/standby

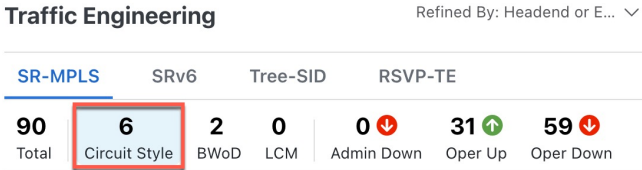
Initial State	Events	Behavior
Working path is down, Protect path is down, Restore path is up/active	Protect path comes back up, then Working path comes back up	<p>On side A: The Working path failure is local (the first Adj SID in the SegList is invalid):</p> <ol style="list-style-type: none"> 1. Protect 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. <p>On side Z: Working path failure is remote (first Adj SID in SegList is valid):</p> <ol style="list-style-type: none"> 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.

Workflow for Setting Up CS SR-TE Policy Visualization

The following tasks are necessary to start visualizing Circuit Style SR-TE policies in the topology map:

Table 3: Tasks to Complete to Start Visualizing Circuit Style SR-TE Policies

Step	Action
1. Enable the SR Circuit Style Manager (CSM) feature pack.	<p>From the main menu, choose Services & Traffic Engineering > Traffic Engineering > Circuit Style SR-TE > Configuration.</p> <p>Follow the steps in Enable SR Circuit Style Manager, on page 10.</p>
<p>2. Configure CS SR policies on the devices.</p> <p>Note If you do this step before enabling the Circuit Style SR-TE feature pack, then the CS SR policies will appear operationally down.</p>	<p>You can configure CS SR policies using one of the following methods:</p> <ul style="list-style-type: none"> • Configure CS SR policies manually on the device using the CLI. For more information, see Configure Circuit Style SR Policies, on page 11. • Configure CS SR policies using the Crosswork Network Controller UI. For more information, see the Cisco Crosswork Network Controller Solution Workflow Guide.

Step	Action
3. Verify that the CS SR policies appear in the SR Policy table.	<p>From the main menu, select Services & Traffic Engineering > Traffic Engineering > SR-MPLS > Circuit Style.</p>  <p>The SR Policy table now shows a filtered list containing only CS SR policies.</p>
4. Verify that the reserved bandwidth pool settings you defined in Step 1 are configured properly.	<p>Click on a CS SR node or policy and navigate to the Link Details > Traffic Engineering page (see Review Circuit Style SR-TE Policy Bandwidth Utilization, on page 13). From the Circuit Style section, view the reserved bandwidth pool size. You can also view current Circuit Style SR-TE bandwidth utilization and how much is still available for use.</p>

Enable SR Circuit Style Manager

In order to manage and visualize Circuit Style SR-TE policies on the topology map, you must first enable SR Circuit Style Manager (CSM) and set bandwidth reservation settings.

When CSM is enabled, it computes the best failover bidirectional paths with the requested bandwidth and other constraints defined in the Circuit Style SR policy configuration between two nodes.

-
- Step 1** From the main menu, choose **Traffic Engineering > Circuit Style SR-TE > Configuration**.
- Step 2** Toggle the **Enable** switch to **True**.
- Step 3** Enter the required bandwidth pool size and threshold information. The following list describes additional field information. See also [What Happens When Bandwidth Reservation Settings are Exceeded?](#), on page 18.

Field	Description
Basic	
Link CS BW Pool Size	The percentage of each link's bandwidth reservable for Circuit Style SR-TE policies.
Link CS BW Min Threshold	The Link CS BW Pool utilization percentage beyond which a threshold crossing event notification will be generated.
Advanced	
Validation Interval	This is the interval that 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 CSM will wait for the delegation request, to generate a notification.

Field	Description
Restore Delegation Delay	The duration until which CSM will pause before processing a restore path delegation.

- Step 4** Click **Commit Changes** to save the configuration. After enabling CSM, you must create Circuit Style SR policy configurations either manually on the device (see [Configure Circuit Style SR Policies, on page 11](#)) 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 [Circuit Style SR-TE Important Notes, 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 Performance Measurement Liveness (PM) profile on a device.

- 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 [Circuit Style SR-TE Important Notes, 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
      !
      bidirectional
        co-routed
        association-id 1050
      !

```


Link Details 

Summary **Traffic Engineering**

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A Side	Z Side
Node	NCS-3	NCS1
IF Name	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/0
FA Affinities		
FA TE Metric		
FA Delay Metric		
FA Topologies	128, 129, 130, 131, 132	128, 129, 130, 131, 132...


Circuit Style Bandwidth Pool

	A Side	Z Side
Pool Size	800 Mbps	800 Mbps
Used	4 Mbps	4 Mbps
Available	796 Mbps	796 Mbps




View Circuit Style SR-TE Policies

View Circuit Style SR-TE policy details such as the endpoints, bandwidth constraints, IGP metrics, and candidate (Working and Protect) paths.

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

Traffic Engineering Refined By: Headend or E... 

SR-MPLS SRv6 Tree-SID RSVP-TE

90	6	2	0	0 	31 	59 
Total	Circuit Style	BWoD	LCM	Admin Down	Oper Up	Oper Down

The **SR Policy** table lists all Circuit Style SR-TE policies.

Step 2 From the **Actions** column, click  > **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.

The screenshot shows the SR Circuit Style Manager interface. On the left is a topology map of the San Francisco Bay Area with nodes labeled xrv9k-14 through xrv9k-18. A path is highlighted in purple. On the right is the 'Traffic Engineering' panel. It shows a summary of 90 SR Policies, with 6 active, 2 BWD, 0 LCM, 0 Admin Down, 31 Oper Up, and 59 Oper Down. Below this is a table of SR Policies:

Head...	Endp...	Color	Admin...	Oper St...	Actions	
<input checked="" type="checkbox"/>	xrv9k-16	xrv9k-15	11056	+	+	...
<input type="checkbox"/>	xrv9k-15	xrv9k-16	11056	+	+	View Details
<input type="checkbox"/>	xrv9k-16	xrv9k-15	4294...	+	+	Edit / Delete
<input checked="" type="checkbox"/>	xrv9k-15	xrv9k-16	4294...	+	+	...
<input checked="" type="checkbox"/>	xrv9k-...	xrv9k-12	5600	+	+	...
<input checked="" type="checkbox"/>	xrv9k-12	xrv9k-...	5600	+	+	...

The **Circuit Style Policy Details** window is displayed in the side panel. By default, the Active path is displayed in the topology map and shows the bidirectional paths (Bi-Dir Path checkbox is checked) on the topology map. The Candidate Path list displays the Active (path that currently takes traffic) and Protected paths.

The screenshot shows the SR Circuit Style Manager interface. On the left is a topology map of the San Francisco Bay Area with nodes labeled xrv9k-23 and xrv9k-25. A path is highlighted in purple. On the right is the 'Circuit Style Policy Details' panel. It shows the current and history of the policy. The current policy is 'Headend xrv9k-25 | TE RID: 192.168.0.0/16 | PCC IP: 192.168.0.0 | Source IP: 192.168.0.0'. The endpoint is 'xrv9k-23 | TE RID: 192.168.0.0/16 | Dest IP: 192.168.0.0'. The color is 6905. The performance metrics section shows a traffic rate of 0 Mbps avg. The candidate path section shows a table of candidate paths:

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.16...	100	+	+
<input type="checkbox"/> cfg_srte_c_6905_ep_192.16...	50	+	+

Note The Bandwidth Constraint value can differ from the bandwidth you requested if the value was increased and insufficient resources existed to satisfy demand on all Active and Protected candidate paths.

Step 3 View Candidate path configuration details.

- 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 operational (Oper State Up) candidate path with the highest preference will always be the Active path (see [How Does CSM Handle Path Failures?](#), on page 22). In this example, the Protected path (with preference 50) is currently the Active path and is displayed on the topology map. Notice that it is designated with a green "A" icon under State to clearly indicate it is currently the operational Active path. Click **Expand All** to view more information about both paths.

- 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:


Circuit Style Policy Details


... | X

Current History

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.168....	100		↑ A ^
Path Name cfg_srte_c_6905_ep_192.168.0.25_disc			
Oper State ↑ Up A Active			
Metric Type IGP			
Bandwidth Requested: 9.006 Mbps Reserved: 0 Mbps			
Bi-Dir Association ID 5906			
Config ID CS-CS-SR-WP-601-head-end-internal			
Disjoint Group ID: 567 Association Source: 0.0.0.0 Type: Node-disjoint			
PCE Initiated false			
Affinity Exclude-Any: - Include-Any: - Include-All: -			
Segment Type Unprotected			
SID Algorithm -			

Here is a sample of a Circuit Style policy configuration. For more information, see [Configure Circuit Style SR Policies](#), on page 11.

- 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.

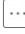



Trigger CSM to Recalculate a Circuit Style SR-TE Policy

Circuit Style SR-TE policies are static in nature, meaning once the paths are computed, they will not be automatically re-optimized based on topology or operational status changes that may affect their paths. You can manually trigger CSM to recalculate a CS-SR policy after the policy's operational status went from down to up or if bandwidth size and requirement changes have been configured.



Note You can only reoptimize an Active and Protect path. It will not work for a Restore path.

- Step 1** From the main menu, choose **Services & Traffic Engineering > Traffic Engineering > SR-MPLS** and click **Circuit Style**. The **SR Policy** table lists all Circuit Style SR-TE policies.
- Step 2** From the **Actions** column, click  > **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, click  > **Reoptimize**.

What Happens When Bandwidth Reservation Settings are Exceeded?

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 either the bandwidth pool size or bandwidth alarm threshold that were set in the CSM Configuration page.

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 for 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.

Link Details 

Summary		
Traffic Engineering		
General	SR-MPLS	SRv6
Tree-SID	RSVP-TE	
	A Side	Z Side
Node	xrv9k-15	xrv9k-16
IF Name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA Affinities		
FA TE Metric		
FA Delay Metric		
FA Topologies	128, 129, 130, 131, 132	128, 129, 130, 131, 132...
Circuit Style Bandwidth Pool		
	A Side	Z Side
Pool Size	100.00 Mbps	100.00 Mbps
Used	0 Mbps	0 Mbps
Available	100.00 Mbps	100.00 Mbps

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.

Link Details 🗑️ | ✕

Summary **Traffic Engineering**

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A Side	Z Side
Node	5501-02	5501-01
IF Name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA Affinities		
FA TE Metric		
FA Delay Metric		
FA Topologies	128, 129, 130, 131, 132	128, 129, 130, 131, 132...

Circuit Style Bandwidth Pool

	A Side	Z Side
Pool Size	1000 Mbps	1000 Mbps
Used	500 Mbps	500 Mbps
Available	500 Mbps	500 Mbps

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

Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/21
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/20
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/2
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/0
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/1/0/1
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/0/0

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 1Gbps and the alarm threshold is set for 900 Mbps.

- An existing Circuit Style SR-TE policy from node 5501-02 to node 5501-01 (*r02 - r01*) uses a bandwidth of 500 Mbps.
- 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 2 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 therefore remains operationally down. CSM will try again, every 30 minutes, to find a path that meets the bandwidth requirements.

Summary

Admin State ↑ Up

Oper State ↓ Down

Binding SID 0

Policy Type Circuit-Style

Profile ID -

Description -

Traffic Rate 0 Mbps

Unused True ?

[See more](#) ▼

Candidate Path

[Expand All](#)

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_100	100		↑ A ▼
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_50	50		↑ ▼

- Alerts are triggered.

Source	Severity	Description
Optima CSM App	⚠ Warning	Unable to compute path for 10.10.10.1 <-> 10.255.255.2 color 2000 due to CsmUpdateStatus.NO_PATH
SR Policy [10.255.255.2#10.255.255.2]	⚠ Warning	Policy 'srte_c_2000_ep_10.10.10.1' has operational status as DOWN.
SR Policy [10.255.255.1#10.255.255.2]	⚠ Warning	Policy 'srte_c_2000_ep_10.10.10.2' has operational status as DOWN.

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 polices (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.

Summary

Admin State ↑ Up

Oper State ↓ Down

Binding SID 0

Policy Type Circuit-Style

Profile ID -

Description -

Traffic Rate 0 Mbps

Unused True ?

[See more](#) ▼

Candidate Path

[Expand All](#)

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_100	100		↑ A ▼
<input checked="" type="checkbox"/> cfg_r1-r2-2_discr_50	50		↑ ▼

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

Source	Severity	Description
SR Policy [10.1#10.255...	Clear	Policy 'srte_c_2000_ep_10.2' has operational status back to UP
SR Policy [10.2#10.255...	Clear	Policy 'srte_c_2000_ep_10.1' has operational status back to UP

How Does CSM Handle Path Failures?

Cisco Crosswork computes paths for Circuit Style SR-TE policies only after a complete bidirectional, path-protected set of candidate paths has been delegated. There are three types of candidate paths that are used during path failures:

- **Working**—This is the path with the highest preference candidate path.
- **Protect**—This path is defined as the second highest preference candidate path. If the Working path goes down, the Protected path (with the lower preference value) is activated. After the Working path recovers, the Protected path remains active until the default lock duration expires.
- **Restore**—This path is defined as the lowest preference candidate path. 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 10](#)). This delay allows topology and policy state changes to fully propagate to Crosswork, in cases where these changes triggered the Restore path delegation.

To address path failures effectively and switchover from working path to protect path, you can configure Performance Measurement (PM). For more information, see [Configure Circuit Style SR Policies, on page 11](#).

Examples



Note Illustrations are for demonstration purposes only and may not always reflect the exact UI or data described within 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 Protected paths of the Circuit Style SR-TE policy are operational. The *active* path is indicated by the "A" icon.

Endpoint 5501-01 | TE RID: 10.255.255.1
Color 1000

Summary

- Admin State Up
- Oper State Up
- Binding SID 24016
- Policy Type Circuit-Style
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True [?](#)
- [See more](#)

Candidate Path Expand All

Path Name	Pref	RoleState
<input checked="" type="checkbox"/> > cfg_r1-r2_discr_100	100	Up Up
<input checked="" type="checkbox"/> > cfg_r1-r2_discr_50	50	Up

When the Active path goes down, the Protected path immediately becomes "active". When the Active path goes back up, then the Protected path takes the role of "protected" again and the Active path (with preference 100) becomes active.

Endpoint 5501-01 | TE RID: 10.255.255.1
Color 1000

Summary

- Admin State Up
- Oper State Up
- Binding SID 24016
- Policy Type Circuit-Style
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True [?](#)
- [See more](#)

Candidate Path Expand All

Path Name	Pref	RoleState
<input checked="" type="checkbox"/> > cfg_r1-r2_discr_50	50	Up Down
<input checked="" type="checkbox"/> > cfg_r1-r2_discr_100	100	Down

In the case where both Working and Protected paths go down, CSM calculates a Restore path and it becomes the active path. Note that the Restore path has the lowest preference value of 10. The Restore path only appears in this particular case. If either the Working or Protected paths become operational again, the Restore path disappears from the topology map and from the Candidate Path list.

How Does CSM Handle Path Failures?

Traffic Engineering / Traffic Engineering Last Refresh: 13-Oct-2022 03:33:54 PM GMT+11

Show Traffic Engineering Device Groups Unassigne...

All Locations / Unassigned Devices

Show: Participating Only KSP Path Bi-Dir Path

Show Groups Color: 1000

Summary

- Admin State Up
- Oper State Up
- Binding SID 24007
- Policy Type Circuit-Style
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True (i)
- [See more](#)

Candidate Path Expand All

Path Name	Pref	RoleState
<input checked="" type="checkbox"/> > cfg_r1-r2_discr_10	10	Up
<input type="checkbox"/> > cfg_r1-r2_discr_50	50	Down
<input type="checkbox"/> > cfg_r1-r2_discr_100	100	Down

Auto-Focus



CHAPTER 3

Local Congestion Mitigation (LCM)

- [Local Congestion Mitigation Overview, on page 25](#)
- [LCM Important Notes, on page 26](#)
- [LCM Calculation Workflow, on page 30](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 31](#)
- [Configure LCM, on page 38](#)
- [Add Individual Interface Thresholds, on page 41](#)
- [Monitor LCM Operations, on page 42](#)
- [Temporarily Exclude an Interface from LCM, on page 45](#)

Local Congestion Mitigation Overview

Local Congestion Mitigation (LCM) searches for congestion on a configurable cadence (as opposed to a triggered event) and 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 the minimal amount of 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 are able to 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 allows for a wider applicability of the solution in various network topologies such as that involving multiple IGP areas due to its simpler path computation and limitation to specific network elements. Focusing on the problem locally within a domain eliminates the need for simulating edge-to-edge traffic flows in the network through a full traffic matrix and allows for better scaling of large networks. Also, LCM performs the collection of TTE SR policy and interface counters via SNMP and does not require the use of SR-TM.



Note Take a look at the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 31](#) to see how to use LCM in your network.

LCM Important Notes

Consider the following information when using LCM:

- You must have the Advanced RTM license package to use LCM.
- Role-based Access Control (RBAC) and task permissions have been introduced in this release. LCM-specific Device Access Groups are automatically added to each domain that is identified as the network is discovered. Confirm that appropriate users are assigned to these Device Access Groups so that they will be able to view LCM recommendations and take the appropriate action. Admin users that have access to LCM Device Access Groups for one or more domains may modify LCM configurations. However, removing LCM Device Access Groups have to be done manually. For more information on RBAC and user roles, see the "[Cisco Crosswork Network Controller Administration Guide](#)".
- 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 2000 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.

LCM Platform Requirements

The following is a non-exhaustive list of high-level requirements for proper LCM operation:

Congestion Evaluation Requirements

LCM requires traffic statistics from the following:

- Interface traffic measurements
- Headend SR-TE policy traffic measurements

To ensure LCM is receiving these traffic statistics:

- Enable SNMP on the devices whose traffic you want to monitor, including the headend device. For more on this task, see [Configuring SNMP Support](#). Note that gNMI is also an option for collecting traffic measurements.
- Ensure that the SNMP-enabled devices are all reachable from the Crosswork Data Gateway. For more on this task, see [Check Connectivity to the Destination](#).
- Configure the headend device to use strict SID labels for SR policies. To perform this task:
 1. Enable segment routing on the headend device and configure the segment routing global block (SRGB) and the segment routing local block (SRLB) ranges. For example:

```
segment-routing
mpls
  global-block 16000 23999
  node-msd 16
!
srlb 15000 15999
```

2. Configure the SR policy candidate paths to use strict SID labels. You can use either explicit paths or dynamic paths with constraints. For example:

```
segment-routing
traffic-eng
  policy COLOR-100-TO-10.0.0.1
  color 100 end-point ipv4 10.0.0.1
  candidate-paths
  preference 100
  explicit segment-list SL1
!
```

```

        preference 200
        dynamic
        constraints
            affinity include-any RED BLUE
            sid-algorithm strict-spf
        !
        !
        !
        !
        !
        !
        segment-list SL1
        index 10 mpls label 16001 node 10.0.0.2 strict
        index 20 mpls label 16002 node 10.0.0.3 strict
        index 30 mpls label 16003 node 10.0.0.4 strict
        !

```

3. Configure the SR policy headend behavior using the binding SID and the autoroute announce option. For example:

```

!segment-routing
traffic-eng
pcc
profile 1
autoroute
include ipv4 all
force-sr-include
!
!
!
!

```

Congestion Mitigation Requirements

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 `force-sr-include` to enable traffic steering into SR-TE policies with autoroute. For example:

```
segment-routing traffic-eng pcc profile ID autoroute 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. For more information, see the [Segment Routing Configuration Guide, Cisco IOS XE 17 \(Cisco ASR 920 Series\)](#), [COE-PCE Initiated SR Policy with OSPF and IS-IS SR-TE Autoroute Announce](#).



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

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.

If all these conditions are met, then the device can support an SR-TE policy using ECMP.

BGP-LS Speaker Placement for Multiple AS Networks with a Dedicated IGP Instance Between ASBRs

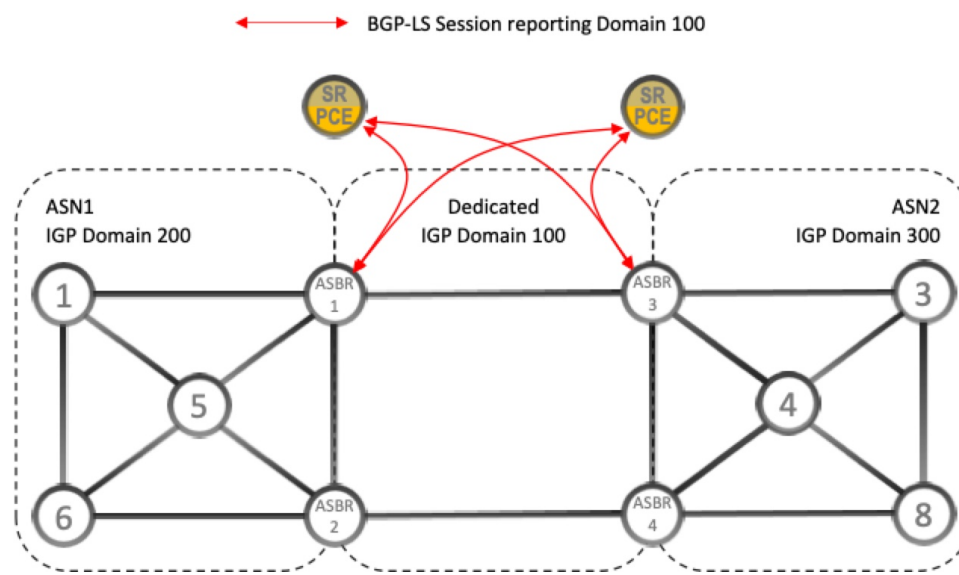
To support interdomain latency-optimized SR policy path computation by an SR-PCE (or other use cases where egress peer engineering (EPE) is not supported), a dedicated IGP instance may be configured between autonomous system border routers (ASBRs) in different ASNs. In these cases, it is important to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

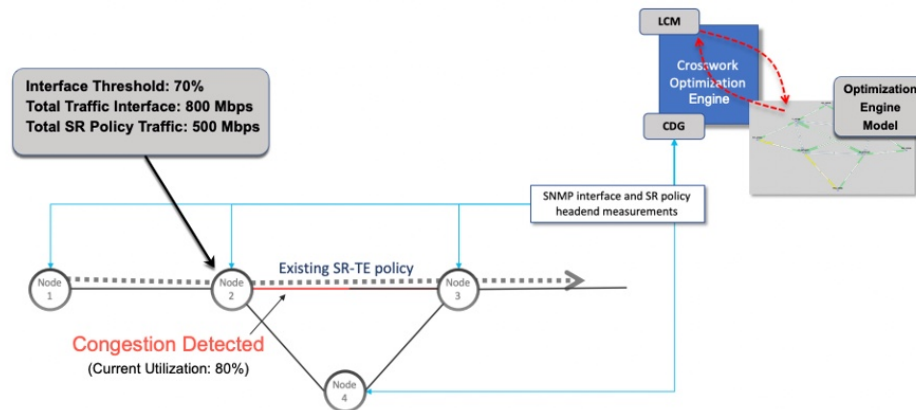
Figure 1: BGP-LS Session Reporting Domain 100



LCM Calculation Workflow

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

Figure 2: LCM Configuration Workflow Example



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, after a congestion check interval, LCM detects congestion 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). The traffic within an SR-TE policy will not be included in 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. 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

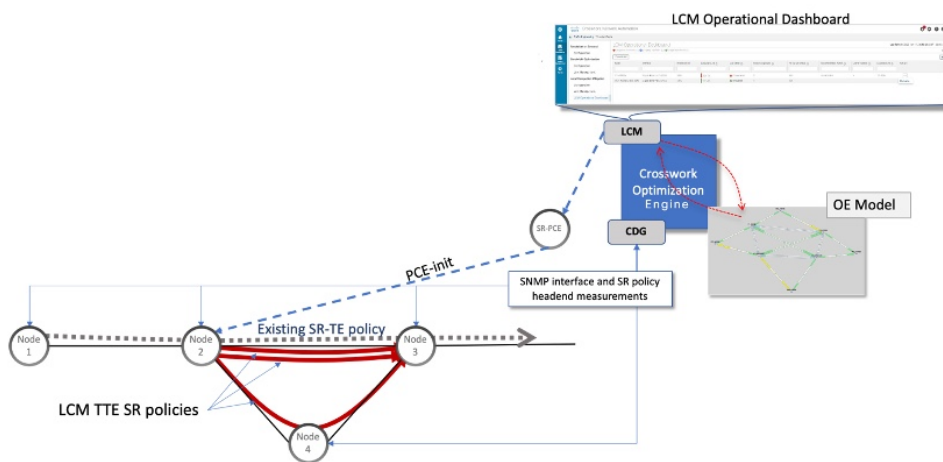
LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information,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



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**. TTE SR policy removal recommendations will occur if the mitigated interface would not be congested if these policies were removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Workflow 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:

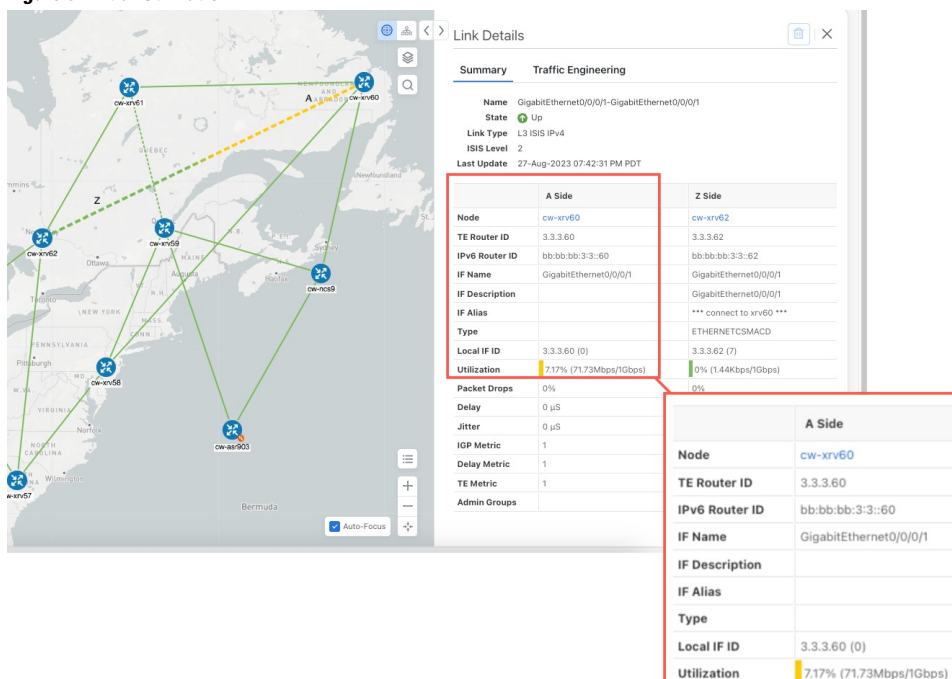
1. View uncongested topology.
2. Set utilization thresholds for individual interfaces.

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.

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

- a) In this example, note that the node cw-xrv60 has a utilization of 7.17%.

Figure 3: 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 in the **LCM Configuration** page. However, if you want to define different thresholds for individual interfaces, we recommend that you define them in the **Customized Interface Threshold** page *prior* to 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-ID > ... > 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 41](#).

See the following example and note the defined threshold for cw-xrv60 with interface GigabitEthernet0/0/0/1 is 16%.

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

Figure 4: Customized Interface Thresholds

Customized Interface Thresholds

Interfaces to Monitor: Selected Interfaces - LCM monitors only the interfaces with custom thresholds.

+ Create | Edit Mode: OFF

Node	Interface	Threshold (%)	Select for
cw-xrv60	GigabitEthernet0/0/0/1	16	<input type="checkbox"/>

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 in the **LCM Configuration** page (see **Step 3**).

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-ID** and click **Configuration**. Toggle the **Enable** switch to **True** and configure other LCM options. In this example, the global threshold is set at 80%, the **Interfaces to Monitor > All Interfaces** option is selected.

Figure 5: LCM Configuration Page

Configuration

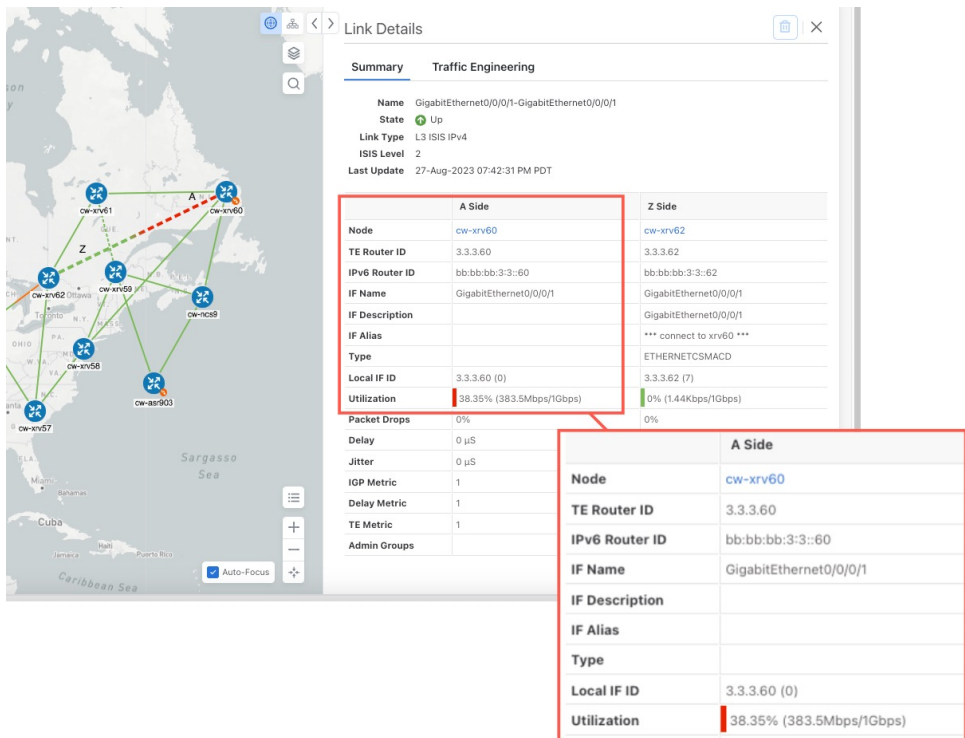
Basic Advanced

Enable <input checked="" type="checkbox"/> True False <input type="checkbox"/> True	Color <input type="text" value="2000"/> Range: 1 to 4294967295	Utilization Threshold <input type="text" value="80"/> Range: 0 to 100
Utilization Hold Margin <input type="text" value="5"/> Range: 0 to Utilization Threshold	Delete Tactical SR Policies when Disabled <input checked="" type="checkbox"/> True False <input type="checkbox"/> True	Profile ID <input type="text" value="1981"/> Range: 0 to 65535
Congestion Check Interval <input type="text" value="300"/> seconds Range: 60 to 86400 seconds	Max LCM Policies per Set <input type="text" value="8"/> Range: 1 to 8	Interfaces to Monitor <input type="radio"/> Selected Interfaces <input checked="" type="radio"/> All Interfaces
Description <input type="text" value="LCM Startup Config"/>		

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 16% for node cw-xrv60 with interface GigabitEthernet0/0/0/1.

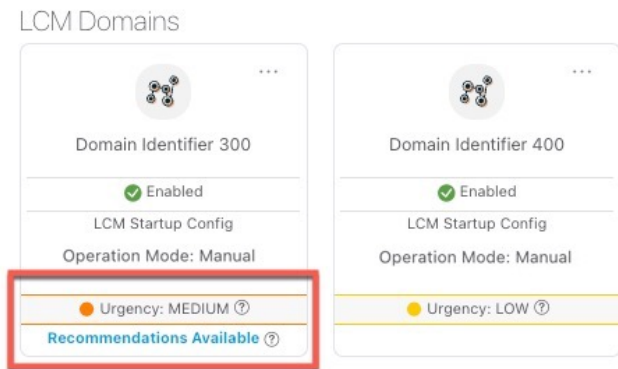
Figure 6: Observed Congestion



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 7: 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-ID > ... > Operational Dashboard)**.

The dashboard shows that cw-xrv60 utilization has surpassed 16% and is now at 38.5%. In the **Recommended Action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended Action - Update Set**) to address the congestion on the interface. For more information, see [Monitor LCM Operations, on page 42](#).

Figure 8: LCM Operational Dashboard



Node	Interface	Threshold Ut...	Evaluation ...	LCM State	Policies De...	Policy Set ...	Recommende...	Commit ...	Expected Util...	Solution Update...
cw-xrv60	GigabitEther...	16%	38.35%	Congested	4	DEGRADED	Update Set	None	14%	29-Aug-2023 04:...

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](#)).

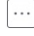
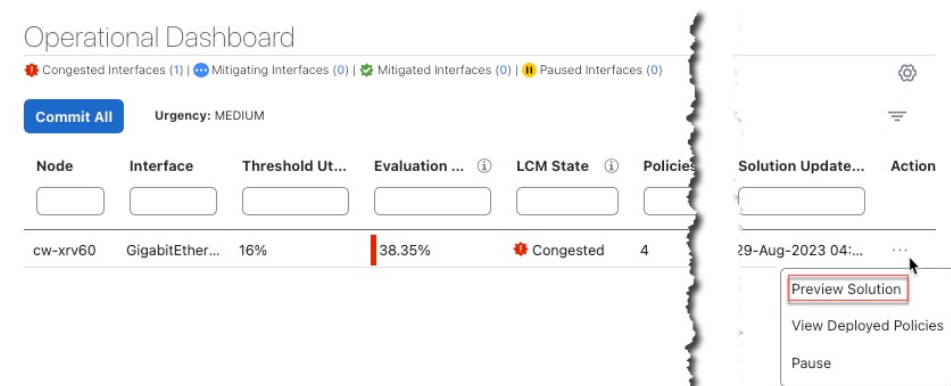
- d) Before committing TTE policies, you can preview the deployment of each TTE policy solution set. Click  in the **Actions** column and choose **Preview Solution**.

Figure 9: 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 do in 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-xrv60.

Figure 10: 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 11: Mitigating State

Note All LCM recommendations per domain must be committed in order 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 that are 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.

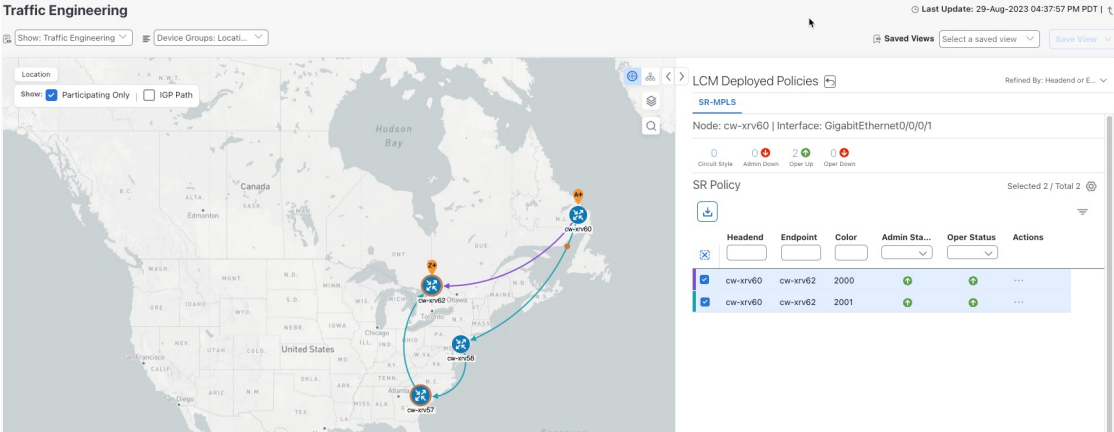
- b) Return to the **Operational Dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

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.

Figure 12: View TTE Deployment Policies on Topology Map



The screenshot displays the Cisco Traffic Engineering interface. On the left, a map of the United States shows a network topology with nodes labeled 'cw-xrv60', 'cw-xrv62', and 'cw-xrv67'. On the right, a panel titled 'LCM Deployed Policies' is open, showing details for the node 'cw-xrv60' on interface 'GigabitEthernet0/0/0/1'. The panel displays two SR Policies:

Headend	Endpoint	Color	Admin Sta...	Oper Status	Actions
cw-xrv60	cw-xrv62	2000			...
cw-xrv60	cw-xrv62	2001			...

d) View SR policy details.

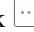
From the **Actions** column of one of the deployed policies click  and choose **View Details**. Note that the **Policy Type** is **Local Congestion Mitigation**.

Figure 13: SR Policy Details

The screenshot displays the 'SR Policy Details' window. It features two tabs: 'Current' (selected) and 'History'. Under the 'Current' tab, the following information is shown:

- TE RID: 3.3.3.60 | IPv6 RID: bb:bb:bb:3::60
- PCC IP: 3.3.3.60
- Endpoint: cw-xrv62 | Dest IP: 3.3.3.62
- TE RID: 3.3.3.62 | IPv6 RID: bb:bb:bb:3::62
- Color: 2000

Below this is a 'Summary' section with the following details:

- Admin State: Up (green up arrow)
- Oper State: Up (green up arrow)
- Binding SID: 24013
- Policy Type: Local Congestion Mitigation
- Profile ID: 2021
- Description: -
- Traffic Rate: 198.93 Mbps
- Unused: False (info icon)
- Delay: 1 (info icon)
- Accumulated Metric: 0
- Delegated PCE: 10.194.132.94
- Non-delegated PCEs: 10.194.132.93
- PCE Computed Time: 28-Aug-2023 04:59:16 PM PDT
- Last Update: 28-Aug-2023 04:59:19 PM PDT

A 'See less' link with an upward arrow is located at the bottom of the summary section.

Step 7 Remove the TTE SR policies upon LCM recommendation.

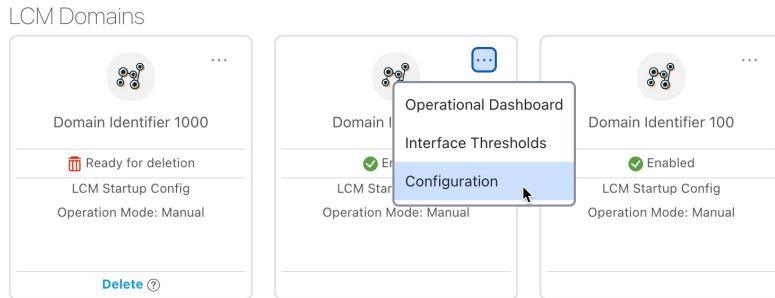
- After some time, the deployed TTE SR policies may no longer be needed. This occurs if the utilization will continue to stay under the threshold without the LCM-initiated TTE tunnels. If this is the case, LCM generates new recommended actions to delete the TTE SR policy sets.
- Click **Commit All** to remove the previously deployed TTE SR policies.
- 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:

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



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

Step 3 Enter the required information. Hover the mouse pointer over **i** to view a description of each field.

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

The following list describes additional **Basic** field information not described in hover text:

- **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 in 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 to recommendations. 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 in the **Customized Interface Thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Customized Interface Thresholds**). Only interfaces defined in the **Customized Interface Thresholds** page will be monitored. If set to **All Interfaces**, LCM will monitor the interfaces with custom thresholds that are uploaded in the **Customized Interface Thresholds** page and the rest of the interfaces using the **Utilization Threshold** value configured on this page.

The following list describes additional **Advanced** field information not described in hover text:

- **Advanced > 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.

- **Advanced > 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.

- **Advanced > Adjacency Hop Type**—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.

Note 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.

- **Advanced > Optimization Objective**—LCM calculates tactical SR policies based on the metric type chosen to minimize.
- **Advanced > Deployment Timeout**—Enter the maximum amount of seconds allowed to confirm deployment of tactical SR policies.
- **Advanced > Over-provisioning Factor (OPF)**—This option helps address unequal ECMP traffic distribution (elephant flows). This value determines the percentage of how much extra traffic should be accounted for when computing a path for a by-pass policy. If LCM needs to divert x 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.
- **Advanced > Maximum Segment Hops**—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.

Note 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.

Crosswork 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, and click **SR-MPLS > Prefixes** tab > **Expand All**.

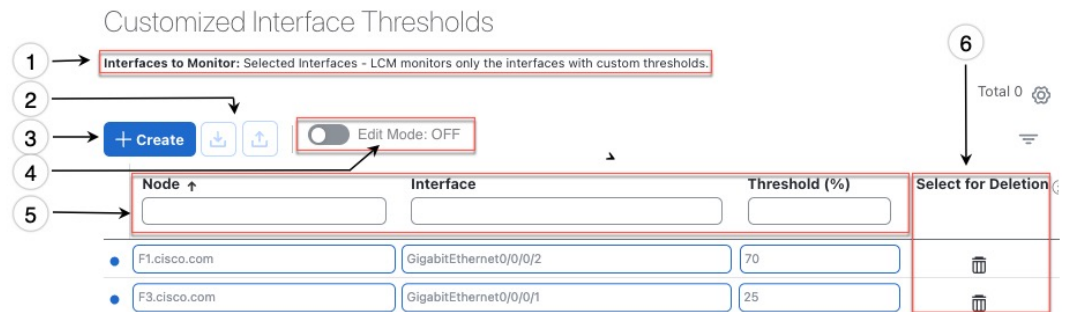
Note Prior to using this option, you must create device tag groups that you want to assign certain MSD values to. For information on creating tags and assigning them to devices, see the [Cisco Crosswork Network Controller Administration Guide](#)


- **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.

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 14: Customized Interface Thresholds



Callout No.	Description
1	Interfaces to Monitor: Displays the option that is currently configured in the Configure LCM 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.
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:

- Step 1** From the main menu, choose **Services & Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > 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 15: Add First Interface

Node	Interface	Threshold (%)	Select for Deletion
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>

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

- Step 4** Confirm that the information appears correctly in 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

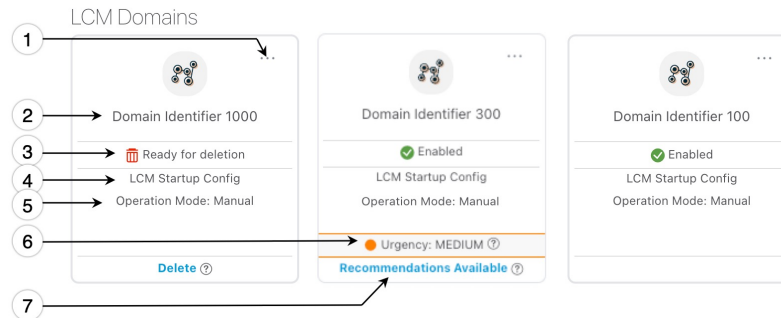


Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 31](#) topic.

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 16: LCM Domains Dashboard



Callout No.	Description
1	<p>Main Menu: Allows you to navigate to the following pages:</p> <ul style="list-style-type: none"> Operational Dashboard Add Individual Interface Thresholds Configure LCM
2	<p>Domain Identifier: The domain ID is taken from the router configuration (<code>link-state instance-id</code>) that you use to advertise IGP with BGP-LS.</p>
3	<p>LCM Status: Indicates whether LCM has been enabled for the domain or can be deleted. Also</p>
4	<p>LCM Configuration Description: The description is defined in the Configure LCM page. The default description is "LCM Startup Config".</p>
5	<p>Operation Mode: Manual—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.</p>
6	<p>Urgency: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following:</p> <ul style="list-style-type: none"> 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	<p>Configure: This link appears if LCM has not yet been configured. Click Configure to go to the Configure LCM page.</p> <p>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.</p> <p>Delete: Indicates that the domain card can be removed from LCM monitoring.</p>

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.

For each interface, it 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 what 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 17: LCM Operational Dashboard

Node	Interface	Threshold Util...	Evaluation Util...	LCM State	Policies Deplo...	Policy Set St...	Recommended ...	Commit St...	Expected Utiliz...	Solution Up...	Actions
L2-NCSS5...	GigabitEthern...	30%	25.85%	Mitigated	2	DEGRADED	Delete Set	None	12.92%	14-Nov-2023...	⋮
L5-8201-L...	FortyGigE0/0/...	8%	15.78%	Congested	0	-	Create Set	None	7.89%	14-Nov-2023...	⋮

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 in 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 are not carrying traffic
- When an interface is continuously throttling between different solutions

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

- No LCM SR policy traffic
- Excessive LCM Policy Traffic Imbalance
- Excessive LCM Oscillations/removals per hour

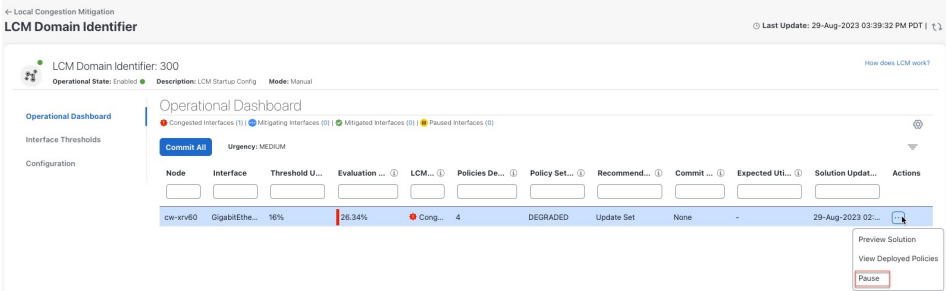
In these circumstances, the user may perform a corrective action, and manually Unpause the interface.

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 18: Pause Interface



The screenshot shows the LCM Operational Dashboard for LCM Domain Identifier 300. The dashboard displays a table of interfaces with columns for Node, Interface, Threshold, Evaluation, LCM, Policies, Policy Set, Recommendation, Commit, Expected Uptime, Solution Update, and Actions. The interface 'GigabitEthernet0/29' is highlighted in blue, and a 'Pause' button is visible in the Actions column.

Node	Interface	Threshold U...	Evaluation ...	LCM...	Policies De...	Policy Set...	Recommend...	Commit ...	Expected Uti...	Solution Updat...	Actions
sw-er060	GigabitEthe...	10%	26.34%	Cong...	4	DEGRADED	Update Set	None	-	29-Aug-2023 02:...	Preview Solution View Deployed Policies Pause



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 for the delivery of 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 when there are large flows.



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

This section contains the following topics:

- [BWoD Important Notes, on page 47](#)
- [PCC-Initiated BWoD SR-TE Policies, on page 48](#)
- [Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example, on page 50](#)
- [Configure Bandwidth on Demand, on page 51](#)
- [Troubleshoot BWoD, on page 52](#)

BWoD Important Notes

Consider the following information when using BWoD:

- You must have the Advanced RTM license package to use BWoD.
- Role-based Access Control (RBAC) and task permissions have been introduced in this release. 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. For more information on RBAC and user roles, 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.

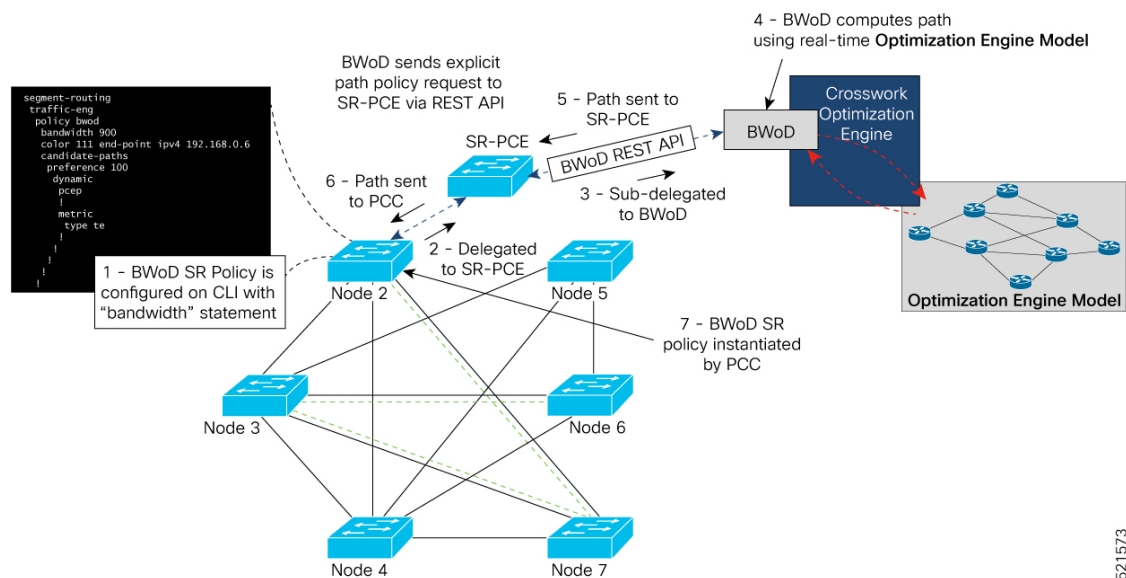
- BWoD will disable itself when an unexpected error is encountered to avoid network disruption.
- BWoD temporarily pauses operation whenever the Optimization Engine model is unavailable due to an Optimization Engine restart or a rebuild of the topology from Topology Services. Any requests to BWoD during this time are rejected. When the model becomes available and BWoD receives 2 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. The persistent connection is made to the SR-PCE BWoD Rest API, registering it as a PCE for bandwidth constrained SR-TE policies.

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

Figure 19: PCC-Initiated BWoD SR-TE Policies




521573

Callout No.	Description
1	<p>A BWoD policy is configured on a PCC via the CLI. For example:</p> <pre> 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 ! ! ! ! ! </pre>
2	<p>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.</p>
3, 4	<p>SR-PCE then sub-delegates the policy to BWoD which attempts to compute a path that meets the bandwidth constraint.</p>
5, 6	<p>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 attempt to minimize violations. This occurrence will also trigger BWoD to issue an event to the COE events UI indicating which BWoD policies are now on best effort paths.</p>
7	<p>A BWoD SR-TE policy is instantiated.</p>


Provision an SR-TE Policy to Maintain Intent-Based Bandwidth Requirements Example

Figure 20: Initial BWoD Topology Example




In this scenario we are using the above topology. The goal is to create a path from F2.cisco.com to F7.cisco.com that can accommodate 920 Mbps of traffic while keeping the utilization at 80%. The above example highlights the utilization on nodes F2.cisco.com and node F4.cisco.com to show that the link is being utilized and has a capacity of 1 Gbps. BWoD will initially try to find a single path that does not include this link since the addition of the requested bandwidth would exceed the utilization threshold.

Step 1 Enable and Configure BWoD.


- From the main menu, choose **Services & Traffic Engineering > Bandwidth on Demand > Configuration**.
- Toggle the Enable switch to **True** and enter **80** to set the utilization threshold percentage. To find descriptions of other options, simply hover the mouse over .
- Click **Commit Changes**.

Step 2 Create a PCE-initiated BWoD SR-TE Policy.

- From the main menu, choose **Traffic Engineering > SR-TE** tab and click **+Create**.
- Enter the required SR-TE policy details.
- In the **Policy Path** field, click **Bandwidth on Demand** and enter a unique name for the BWoD path. In this case, **bwOD-Policy-new**.
- From the **Optimization Objective** drop-down list, select **Traffic Engineering (TE) Metric**.
- In the **Bandwidth** field enter the requested bandwidth. In this case, we are requesting **920 Mbps**.
- Click **Preview**.



In the below example, the BWoD SR policy uses the existing ECMP paths on the network without explicitly splitting the traffic to avoid congestion. The traffic may be distributed by ECMP, but BWoD does not influence that. It is only aware of it and takes it into consideration if it occurs on the path computed.



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

Step 3


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

- From the main menu, choose **Traffic Engineering** > **SR-TE**.
- Select the new BWoD SR-TE policy and view the SR policy details (click and choose **View**). Note that the Policy Type is **Bandwidth on Demand**.

Configure Bandwidth on Demand

There are two parts to configure 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.

-
- 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, navigate to **Traffic Engineering > Traffic Engineering**.
- Step 6** From the SR Policy table, click **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** Click **Preview** to view the proposed SR policy.
- Step 9** 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 4: Errors

Error Event Message	Possible Causes and Recommended Corrective Action
OptimaModelError	<p>The network model used by BWoD from the Optimization Engine is corrupt or is missing key data that is 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.</p> <p>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 which should suffice for small to medium sized networks. However, larger networks may require additional time.</p>
NATSTimedOutError	<p>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.</p>
Traceback or other errors found in the log file	Please contact your Cisco service representative.