



Local Congestion Mitigation (LCM)

- [Local Congestion Mitigation Overview, on page 1](#)
- [LCM Important Notes, on page 2](#)
- [LCM Calculation Workflow, on page 6](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 8](#)
- [Mitigate Congestion Automatically, on page 16](#)
- [Configure LCM, on page 17](#)
- [Add Individual Interface Thresholds, on page 19](#)
- [Monitor LCM Operations, on page 21](#)
- [Temporarily Exclude an Interface from LCM, on page 23](#)

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. This is available when LCM is in Manual mode.
- Enable LCM to automatically deploy TE tunnel recommendations based on thresholds that you specified. This feature is available when LCM is in Automated mode.
- 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 17](#).

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 8](#) 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 17](#)) 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

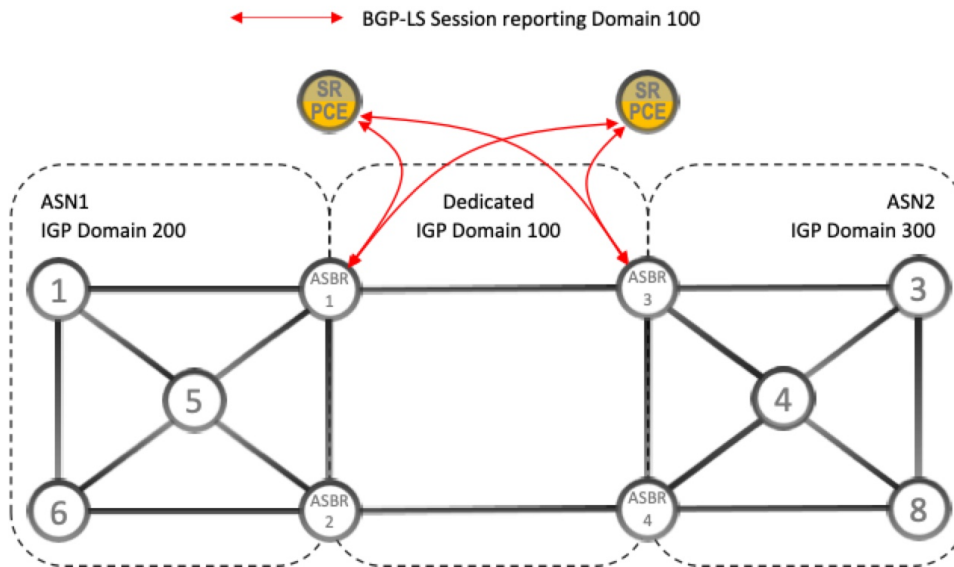
To support interdomain latency-optimized SR policy path computation by an SR-PCE (or other use cases where egress peer engineering (EPE) is not supported), a dedicated IGP instance may be configured between autonomous system border routers (ASBRs) in different ASNs. In these cases, it is important to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

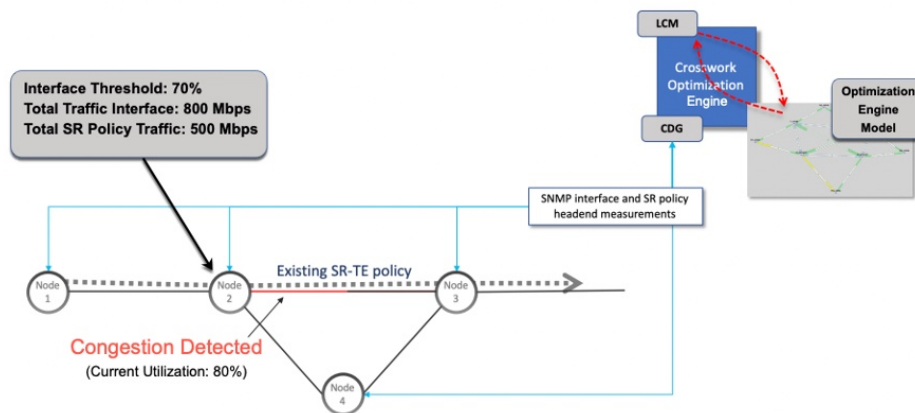
Figure 1: BGP-LS Session Reporting Domain 100



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

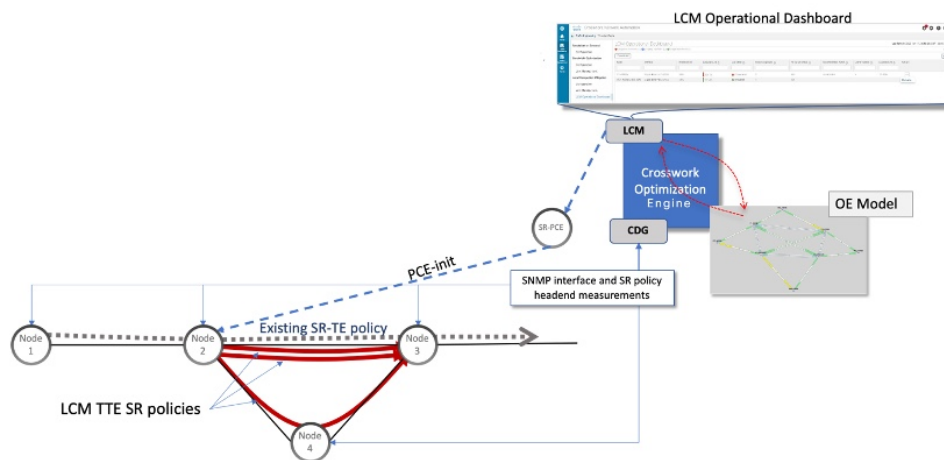
LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see [Configure LCM, on page 17](#).

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.

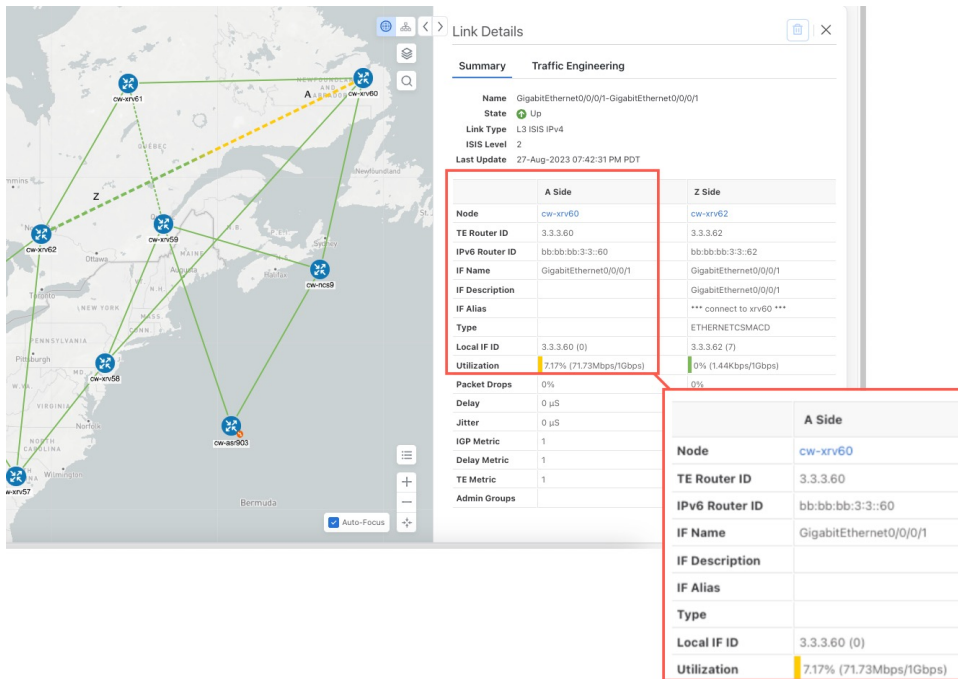
In this example, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE SR policies when utilization on a device's interface surpasses a defined utilization threshold. We will preview the recommended TTE SR policies before committing them to mitigate the congestion. The 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 polices 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 19](#).

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 [Upload] [Download] | 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. In the **Advanced** tab, Operation mode is set to **Manual**. For more information on all the available options, see [Configure LCM, on page 17](#).

Figure 5: LCM Configuration Page

Configuration

Basic Advanced

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

Configuration

Basic

Advanced

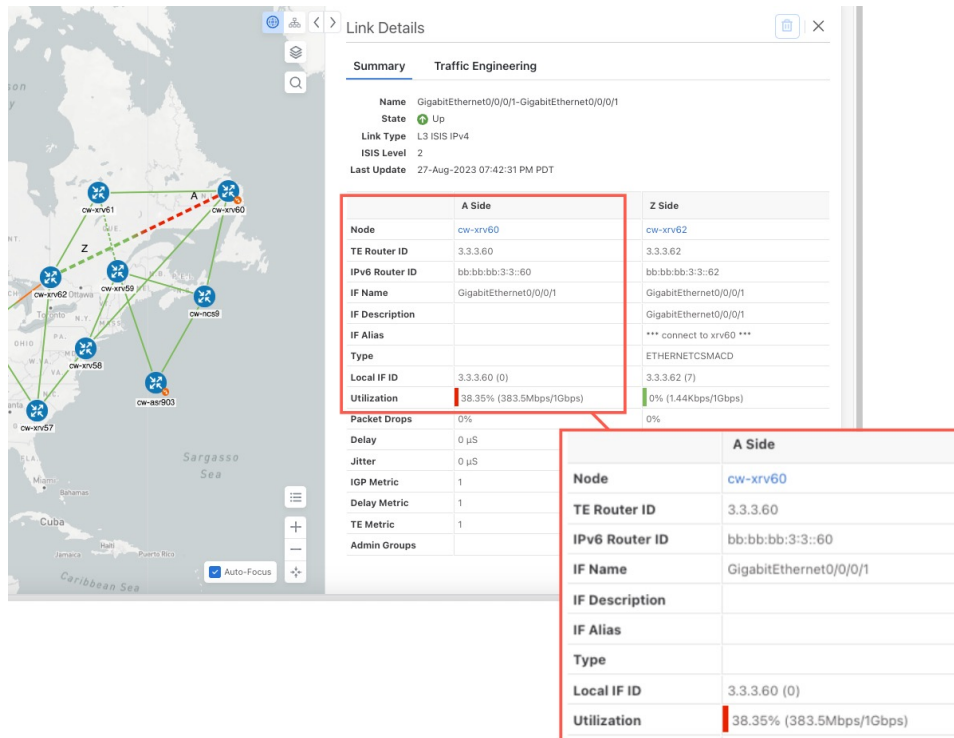
Auto Repair Solution ⓘ False <input type="radio"/> True <input checked="" type="radio"/>	Stay in Area ⓘ False <input type="radio"/> True <input checked="" type="radio"/>	Adjacency Hop Type ⓘ Unprotected <input type="text"/>
Operation Mode ⓘ Manual <input checked="" type="radio"/> Automated <input type="radio"/>	Optimization Objective ⓘ Minimize the IGP metric <input type="text"/>	Deployment Timeout * ⓘ 180 <input type="text"/> Range: 10 to 300
Congestion Check Suspension Interval * ⓘ 600 <input type="text"/> Range: 600 to 3600	Over-Provisioning Factor * ⓘ <input type="text"/> Range: 0.0 to 100.0	Uneven ECMP Traffic Threshold * ⓘ <input type="text"/> Range: 0.0 to 100.0
Throttle Mode Threshold * ⓘ 5 <input type="text"/> Range: 0 to 10		

- 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. LCM will *not* commit or deploy new TTE policies automatically since Manual mode is enabled. Later, you will be able to preview the recommended TTE policies and decide whether or not to commit and deploy them onto your network.

Step 4

After some time, congestion occurs surpassing the custom LCM threshold defined at 16% for node cw-xrv60 with interface GigabitEthernet0/0/0/1.

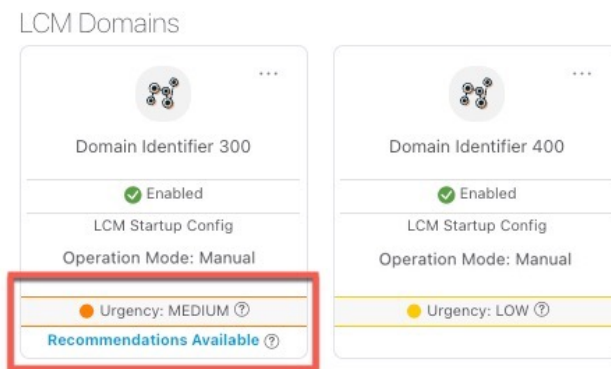
Figure 6: Observed Congestion




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

- 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



- (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 21.

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 17).

- 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**.

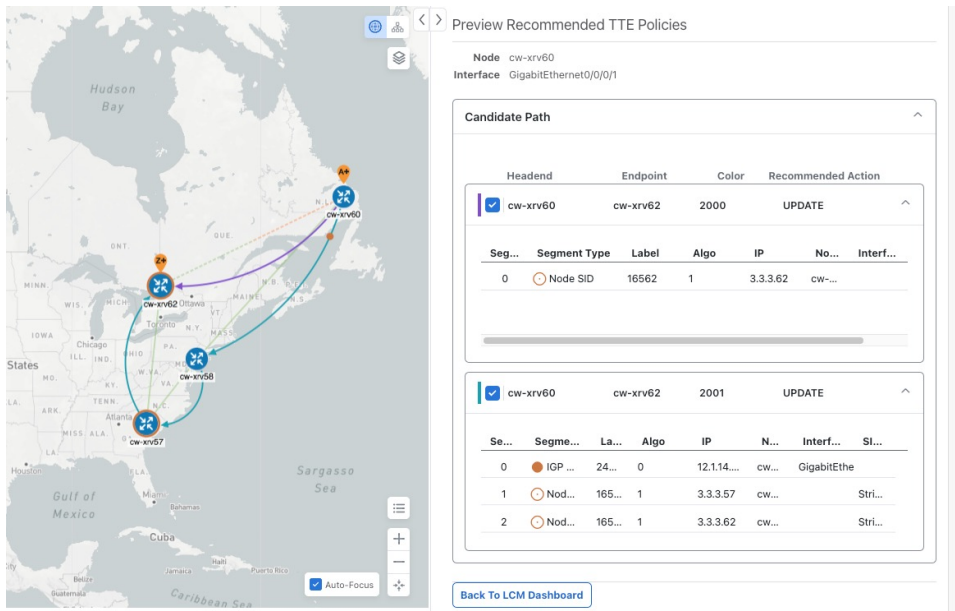
Figure 9: Preview Solution

Node	Interface	Threshold Ut...	Evaluation ...	LCM State	Policies	Solution Update...	Action
cw-xrv60	GigabitEther...	16%	38.35%	Congested	4	29-Aug-2023 04:...	<ul style="list-style-type: none"> Preview Solution View Deployed Policies Pause

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.

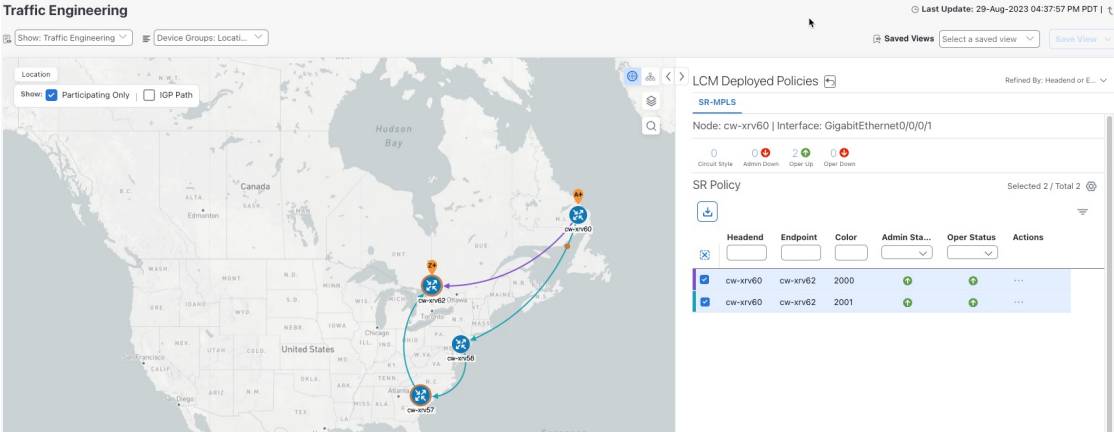
- b) Return to the **Operational Dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

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 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, the 'LCM Deployed Policies' panel is open, showing details for the SR-MPLS policy on interface GigabitEthernet0/0/0/1. The panel includes a table of 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.

Mitigate Congestion Automatically

You can choose to have LCM automatically deploy TE tunnel solutions to mitigate traffic congestion based on thresholds you configure. In Automated mode, LCM does not require user intervention to make changes

in the network. To enable this mode, go to the LCM **Configuration - Advanced** page (**Local Congestion Mitigation > LCM-Domain-Card > ... > Configuration > Advanced** tab) and toggle the Operation Mode option to **Automated**.

In the LCM Operations Dashboard, you can see the current interface and mitigation status that is associated with the domain. You can also view deployed policies and paused interfaces (see [Temporarily Exclude an Interface from LCM, on page 23](#)). Unlike Manual mode, recommendation-related columns are not displayed.

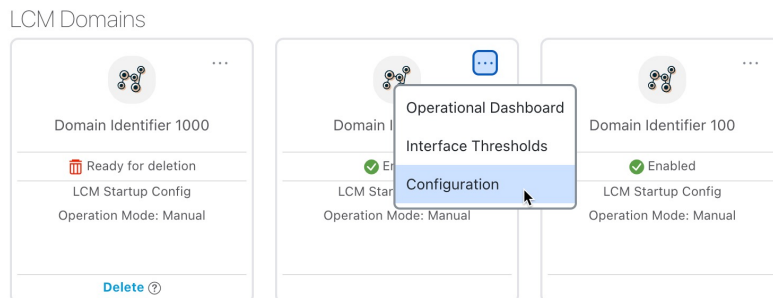
Figure 14: LCM Dashboard - Automated Mode

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Actions
cx-xrv54	GigabitEthernet0/0/0/5	20%	17.20%	Paused	0	OK	...
cx-xrv55	GigabitEthernet0/0/0/6	13%	4.12%	Mitigated	2	OK	...

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.

- **Operation Mode:**

- **Automated Mode**—This option allows LCM to automatically deploy TE tunnel recommendations based on thresholds that a user configures.
- **Manual Mode**—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
- **Throttle Mode Threshold** —In Automated mode, enter the number of times LCM throttles between solutions per hour until the interface is automatically paused.
- **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 6](#). 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 15: Customized Interface Thresholds

Customized Interface Thresholds

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

2 →


3 → + Create

4 →

5 →

6 → Total 0

Node ↑	Interface	Threshold (%)	Select for Deletion
F1.cisco.com	GigabitEthernet0/0/0/2	70	🗑️
F3.cisco.com	GigabitEthernet0/0/0/1	25	🗑️

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:

- Click the **Download sample configuration file** link.
- Click **Cancel**.
- Open and edit the configuration file (LCMLinkManagementTemplate.csv) you just downloaded. Replace the sample text with your specific node, interface, and threshold information.
- Rename and save the file.
- Navigate back to the **Customized Interface Thresholds** page.
- Click **Import .CSV File** and navigate to the CSV file you just edited.
- Click **Import**.

Step 3 If you manually add individual interfaces:

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

Figure 16: Add First Interface



The screenshot shows the LCM interface with a table for adding interfaces. At the top, there is a '+ Create' button, a download icon, an upload icon, and a toggle switch for 'Edit Mode: OFF'. The table has four columns: 'Node', 'Interface', 'Threshold (%)', and 'Select for Deletion'. The first row is highlighted with a red box, indicating it is the first empty row for adding a new interface.

- 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

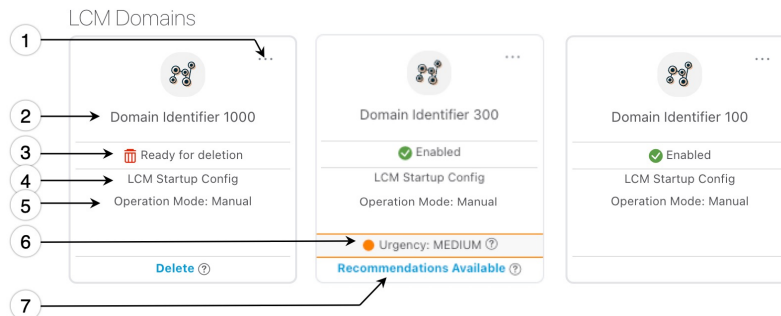


Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 8](#) topic.

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 17: LCM Domains Dashboard



Callout No.	Description
1	Main Menu: Allows you to navigate to the following pages: <ul style="list-style-type: none"> Operational Dashboard Add Individual Interface Thresholds Configure LCM
2	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.
3	LCM Status: Indicates whether LCM has been enabled for the domain or can be deleted. Also
4	LCM Configuration Description: The description is defined in the Configure LCM page. The default description is "LCM Startup Config".

Callout No.	Description
5	<p>Operation Mode: Indicates if LCM is running in Automatic or Manual mode. The default is Manual mode.</p> <ul style="list-style-type: none"> • Automated Mode—LCM automatically deploys TE tunnel recommendations based on thresholds that a user configures. • Manual Mode—This option requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
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 17.
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 18: 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 17](#).

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

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

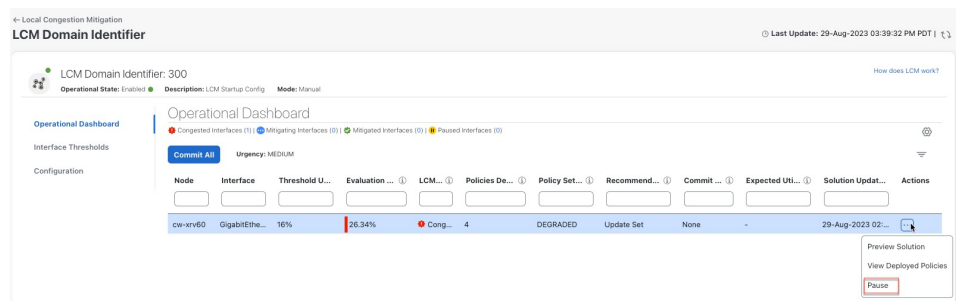
In these circumstances, the user may perform a corrective action, and manually Unpause the interface.

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 19: Pause Interface



The screenshot shows the LCM Operational Dashboard for a domain with identifier 300. The dashboard displays a table of interfaces with the following columns: Node, Interface, Threshold U..., Evaluation ..., LCM..., Policies De..., Policy Set..., Recommend..., Commit ..., Expected Util..., Solution Updat..., and Actions. The table contains one row for interface 'ce-xrv60' on node 'GigabitEthe...'. The evaluation shows a congestion level of 26.34% and a state of 'Cong...'. The Actions column for this interface has a dropdown menu with options: 'Preview Solution', 'View Deployed Policies', and 'Pause'.

Node	Interface	Threshold U...	Evaluation ...	LCM...	Policies De...	Policy Set...	Recommend...	Commit ...	Expected Util...	Solution Updat...	Actions
ce-xrv60	GigabitEthe...	16%	26.34%	Cong...	4	DEGRADED	Update Set	None	-	29-Aug-2023 02:...	<ul style="list-style-type: none"> Preview Solution View Deployed Policies Pause