



Local Congestion Mitigation (LCM)

- [Local Congestion Mitigation Overview, on page 1](#)
- [LCM Important Notes, on page 2](#)
- [LCM Calculation Workflow, on page 4](#)
- [Workflow Example: Mitigate Congestion on Local Interfaces, on page 6](#)
- [Configure LCM, on page 15](#)
- [Add Individual Interface Thresholds, on page 17](#)
- [Monitor LCM Operations, on page 19](#)

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. If the user approves, LCM performs the mitigation through the deployment of Tactical Traffic Engineering (TTE) SR policies. LCM will not modify paths of existing deployments of SR policies to mitigate congestion. 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 deploy changes in the network automatically to address congestion and 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 15](#).

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 6](#) 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.
- 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 assigned to an IGP process. Domains are learned from the network. The domain ID is taken from 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 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 over-provisioning 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.
- LCM will not operate in portions of the network carrying Tree-SID LSPs.
- When domain interfaces and links are removed (intentionally or unintentionally), the following occurs:
 - As links go down (LINK_DOWN state), LCM configuration and the Domain UI card (see [Configure LCM, on page 15](#)) 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.

LCM Platform Requirements

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

Congestion Evaluation:

- LCM requires traffic statistics from the following:
 - SNMP interface traffic measurements
 - SNMP headend SR-TE policy traffic measurements
- Strict SID labels should be configured for SR.

Congestion Mitigation:

- The headend device should support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies
- The headend device must support PCE-initiated SR-TE policies with autoroute steering

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 include ipv4 all
segment-routing traffic-eng pcc profile <id> autoroute force-sr-include
```

where `<id>` is the user configured ID (any number as allowed per router).

See SR configuration documentation for your specific device to view descriptions and supported configuration commands (for example: [Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers](#))

Contact your Cisco sales representative for an exhaustive list of platform requirements.

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.

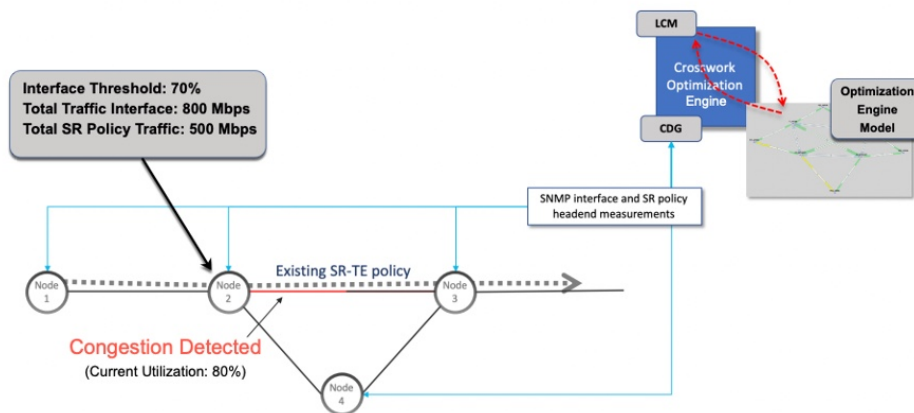
Figure 1: BGP-LS Session Reporting Domain 100



LCM Calculation Workflow

This example walks you from congestion detection to the calculations LCM performs prior to recommending tactical tunnel deployment. With the release of Crosswork Optimization Engine 3.0, 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 (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 = 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: $800 \text{ Mbps} - 500 \text{ Mbps} = 300 \text{ Mbps}$

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:

$800 \text{ Mbps} - 700 \text{ Mbps (70\% threshold)} = 100 \text{ Mbps}$


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 ($100 \text{ Mbps} \times 1.10$) 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 15](#).

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 needs to 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.

This example demonstrates the following workflow:

1. View uncongested topology.
2. Set utilization thresholds for individual interfaces.
3. Enable and configure LCM.
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.

The following image shows the topology that will be used for this example.

Figure 3: Initial Topology



Step 1

View initial topology and utilization prior to LCM configuration.

- a) Click on the link between F3.cisco.com and F5.cisco.com to view link details. Note that utilization on F3.cisco.com is 9%.

Figure 4: Initial Utilization

	A Side	Z Side
Node	F3.cisco.com	F5.cisco.com
TE Router ID	192.168.100.3	192.168.100.5
IF Name	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/0
IF Description	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/0
Type	ETHERNETCSMACD	ETHERNETCSMACD
IP Address	100.100.1.17	100.100.1.18
Utilization	9% (159Mbps/1Gbps)	0% (0Bps/1Gbps)
IGP Metric	1	
Delay Metric	1	
TE Metric	1	
OSPF Router ID	192.168.100.3	
OSPF Area	0	
FA Affinities		
Admin Groups		

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 some individual interface thresholds. Go to the **Customized Interface Thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Interface Thresholds**). Add interfaces 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 17](#).


See the following example and note the defined thresholds for F3.cisco.com with interface GigabitEthernet0/0/0/1 (13%) and F5.cisco.com with interface GigabitEthernet0/0/0/1 (11%).


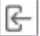


Note The utilization thresholds used in this example are extremely low and are best used for lab environments.


Figure 5: Customized Interface Thresholds

Customized Interface Thresholds

Interfaces to Monitor: All Interfaces - LCM monitors the interfaces with custom thresholds. All other interfaces are monitored using the Utilization Threshold defined in the Configuration page.

Total 5 




Edit Mode: OFF


Node	Interface	Threshold (%)	Select to Delete 
<input type="text"/>	<input type="text"/>	<input type="text"/>	
F4.cisco.com	GigabitEthernet0/0/0/1	14	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/1	13	<input type="checkbox"/>
F5.cisco.com	GigabitEthernet0/0/0/1	11	<input type="checkbox"/>
F1.cisco.com	GigabitEthernet0/0/0/1	20	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/2	10	<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% and the **Interfaces to Monitor > All Interfaces** option is selected. For more information on all the available options, see [Configure LCM, on page 15](#).

Figure 6: LCM Configuration Page

Configuration

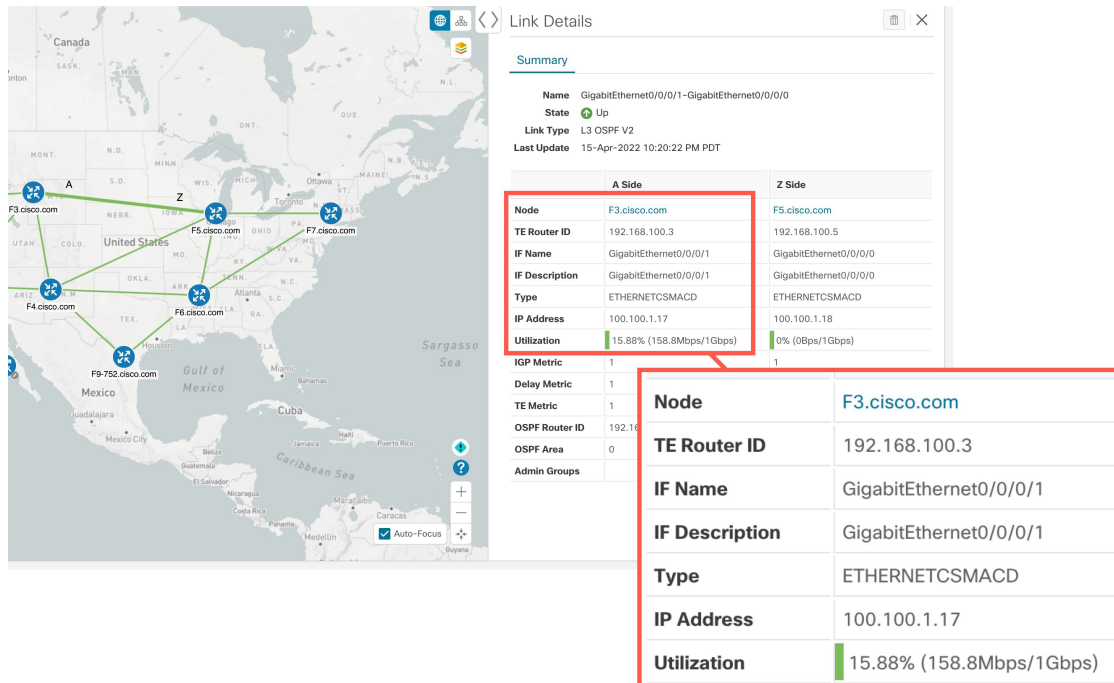
Basic Advanced

Enable ? False <input checked="" type="checkbox"/> True	Color ? 2000 Range: 1 to 4294967295	Utilization Threshold ? 80 Range: 0 to 100
Utilization Hold Margin ? 5 Range: 0 to Utilization Threshold	Delete Tactical SR Policies when Disabled ? False <input checked="" type="checkbox"/> True	Profile ID ? 1981 Range: 0 to 65535
Congestion Check Interval ? 300 seconds Range: 60 to 86400 seconds	Max LCM Policies per Set ? 8 Range: 1 to 8	Interfaces to Monitor ? <input type="radio"/> Selected Interfaces <input checked="" type="radio"/> All Interfaces
Description ? 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. LCM will *not* commit or deploy new TTE policies automatically. 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 13% for node F3.cisco.com with interface GigabitEthernet0/0/0/1.


Figure 7: 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 8: 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 F3.cisco.com utilization has surpassed 13% and is now at 16.05%. It also shows that F5.cisco.com utilization has also surpassed the 11% threshold and is now 19.26%. In the **Recommended Action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended Action - Create Set**) to address the congestion on the interface. The **Expected Utilization** column shows the expected utilization of each of the interface after the recommended action is committed. For more information, see [Monitor LCM Operations, on page 19](#).

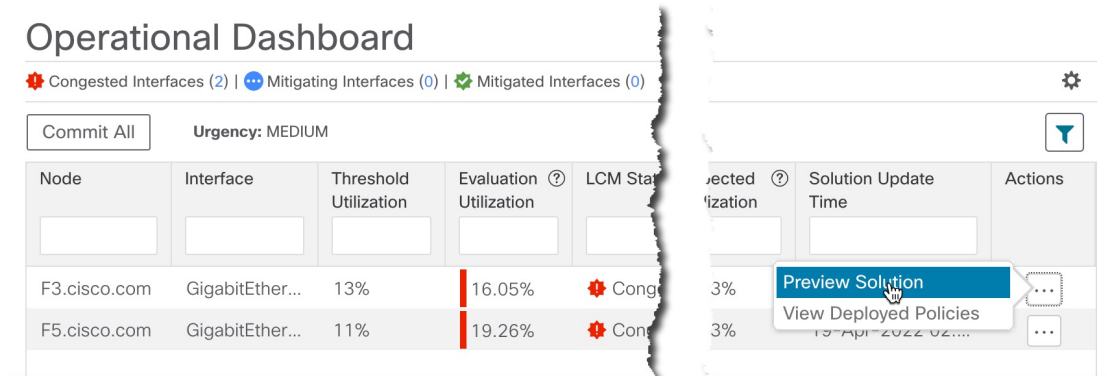
Figure 9: LCM Operational Dashboard

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F3.cisco.com	GigabitEther...	13%	16.05%	Congested	0	-	Create Set	None	8.03%	19-Apr-2022 02:...	...
F5.cisco.com	GigabitEther...	11%	19.26%	Congested	0	-	Create Set	None	9.63%	19-Apr-2022 02:...	...

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

- 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 10: Select 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 F3.cisco.com and interface GigabitEthernet0/0/0/1. The top path shows the node SID (orange outline), headend and endpoint (A and Z) because the mouse pointer hovers over that segment.

Figure 11: LCM TTE Deployment Preview

Preview Recommended TTE Policies

Node F3.cisco.com
Interface GigabitEthernet0/0/0/1

Headend	Endpoint	Color	Recommended Action
<input checked="" type="checkbox"/> F3.cisco.com	F5.cisco.com	2000	CREATE

Se...	Segme...	L...	Algo	IP	N...	Interf...	St...
0	Nod...	16...	1	192.16...	F5...		Strict

Headend	Endpoint	Color	Recommended Action
<input checked="" type="checkbox"/> F3.cisco.com	F5.cisco.com	2001	CREATE

Se...	Segme...	L...	Algo	IP	N...	Interf...	St...
0	Nod...	16...	1	192.16...	F9-...		Strict
1	IGP ...	10...	0	100.10...	F9-...	GigabitEthe U	
2	Nod...	16...	1	192.16...	F5...		Strict

[Back To LCM Dashboard](#)

- 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**.

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.

Figure 12: Mitigating LCM State

Operational Dashboard

🔥 Congested Interfaces (0) |
 🔄 Mitigating Interfaces (2) |
 🟢 Mitigated Interfaces (0)

Commit All Urgency: LOW

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F5.cisco.com	GigabitEther...	11%	19.78%	🔄 Mitigating	2	OK	No Change	CONFIRMED	9.89%	19-Apr-2022 03:...	⋮
F3.cisco.com	GigabitEther...	13%	15.88%	🔄 Mitigating	2	OK	No Change	CONFIRMED	7.94%	19-Apr-2022 03:...	⋮

Step 6 Validate TTE SR policy deployments.

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

Note Crosswork Optimization Engine 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.


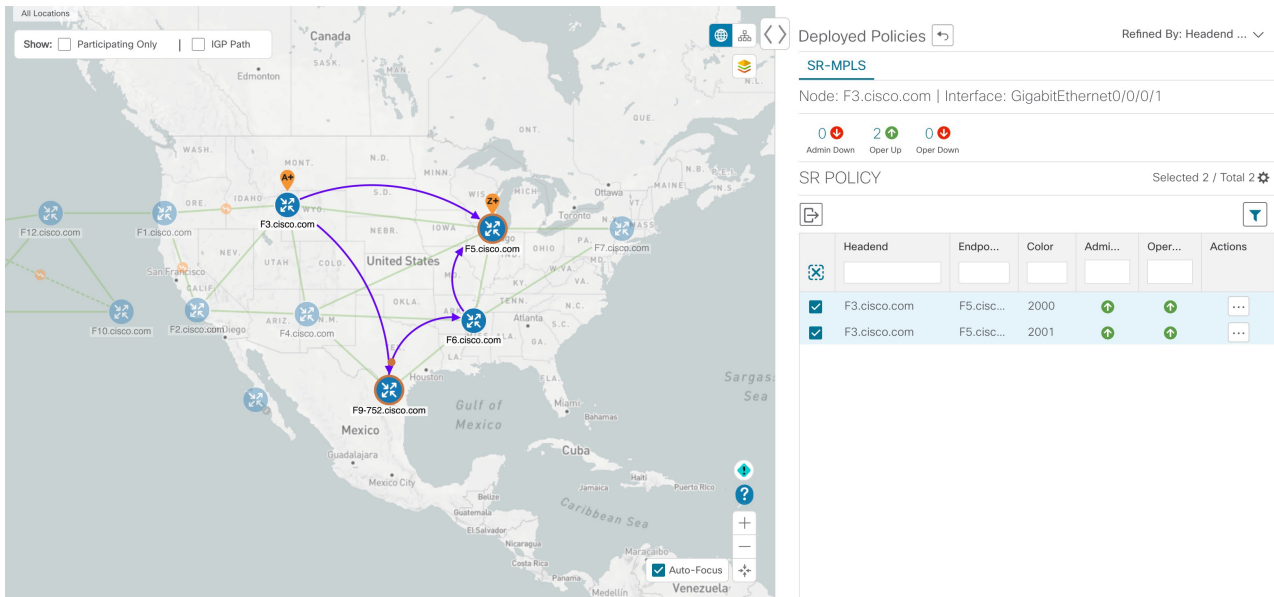


Click  in the **Actions** column and choose **View Deployed Policies**. The deployed policies are displayed in focus within the topology map. All other policies are dimmed.




Figure 13: View TTE Deployment Policies on Topology Map




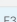
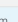

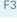
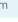
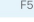
Deployed Policies  Refined By: Headend ... 

SR-MPLS

Node: F3.cisco.com | Interface: GigabitEthernet0/0/0/1

0  Admin Down 2  Oper Up 0  Oper Down

SR POLICY Selected 2 / Total 2 

	Headend	Endpo...	Color	Admi...	Oper...	Actions
<input checked="" type="checkbox"/>	F3.cisco.com	F5.cisc...	2000			
<input checked="" type="checkbox"/>	F3.cisco.com	F5.cisc...	2001			

- d) View SR policy details.

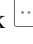
From the **Actions** column of one of the deployed policies click  and choose **View Details**. Note that the **Policy Type** is **Local Congestion Mitigation**.

Figure 14: SR Policy Details

SR Policy Details ⋮ | ✕

Details

Historical Data

Headend A F3.cisco.com | Source IP: 192.168.100.3
 TE RID: 192.168.100.3
 PCC IP: 192.168.100.3

Endpoint Z F5.cisco.com | Dest IP: 192.168.100.5
 TE RID: 192.168.100.5

Color 2000

∨ Summary

- Admin State ↑ Up
- Oper State ↑ Up
- Binding SID 1005011
- Policy Type Local Congestion Mitigation
- Profile ID 1981
- Description -
- Traffic Rate 39.28 Mbps
- Unused False
- Delay 1 i
- BWOD Policy Bandwidth 0 Mbps
- Accumulated Metric 0
- Delegated PCE 10.194.60.51
- Non-delegated PCEs -
- PCE Computed Time -
- Last Update 22-Apr-2022 01:31:10 PM PDT

[See less](#) ^

∨ Candidate Path [Collapse All](#)

Path Name	Preference	Path Type	State
<input checked="" type="checkbox"/> ∨ lcm_to_F5_cisco_com_c_2000	100	Explicit	↑ A

Segment	Segment T...	Label	Algo	IP	Node	Interface	SID
0	⊙ Node SID	16505	1	192.168.1...	F5.cisco.com		Stric

Path Name lcm_to_F5_cisco_com_c_2000

Oper State ↑ Up | A Active

Metric Type UNKNOWN

Disjoint Group ID:
 Association Source: -
 Type: -

PCE Initiated true

Affinity Exclude-Any: -
 Include-Any: -
 Include-All: -

Segment Type Unprotected

SID Algorithm -

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.

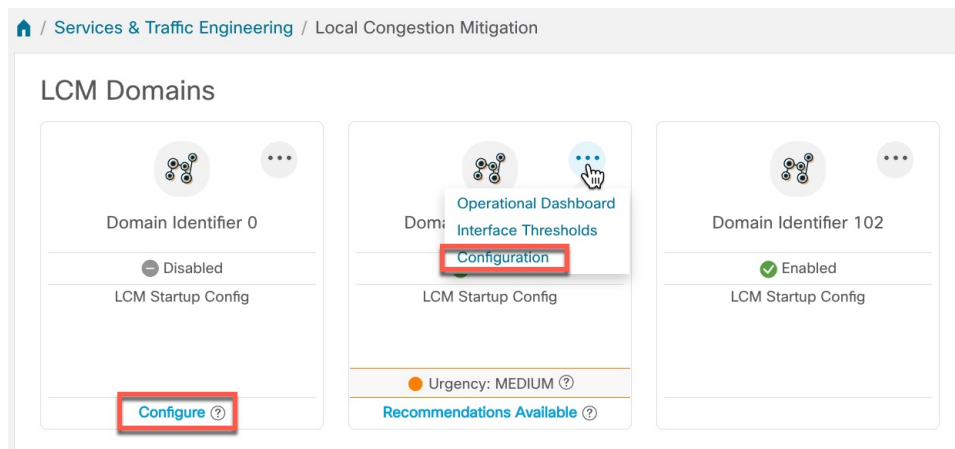
In this scenario we observed how to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands but at the same time gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations 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 **Traffic Engineering > Local Congestion Mitigation > Domain-ID-card** and click one of the following:

- **Configuration**
- **Configure**



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

Step 3 Enter the required information. Hover the mouse pointer over **?** 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 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.
- **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 4. 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

Interface	Threshold (%)	Select to Delete	
F4.cisco.com	GigabitEthernet0/0/0/1	14	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/1	13	<input type="checkbox"/>
F5.cisco.com	GigabitEthernet0/0/0/1	11	<input type="checkbox"/>
F1.cisco.com	GigabitEthernet0/0/0/1	20	<input type="checkbox"/>
F3.cisco.com	GigabitEthernet0/0/0/2	10	<input type="checkbox"/>

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

Callout No.	Description
2	Import CSV File: All interfaces currently in the table will be replaced with the data in the CSV file you import.
3	Add: Click this icon to add new interface threshold rows.
4	Export CSV File: All interfaces are exported to a CSV file. You cannot filter data for export.
5	Edit Mode: When Edit Mode is ON , you can edit multiple fields in one session, then click Save .
6	Filter: By default, this row is available for you to enter text in which to filter content. To disable or enable the filtering feature, click  .
7	Select to Delete: 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 **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 a web interface for adding interface thresholds. At the top, there are navigation icons and a toggle for 'Edit Mode: OFF'. Below this is a table with four columns: 'Node', 'Interface', 'Threshold (%)', and 'Select for Delete...'. The first row of the table is highlighted with a red border, indicating it is the active row for data entry. A blue selection bar is visible at the bottom of the table, with a small trash icon on the right side.

- Click  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. For more information, see Figure 15.

Monitor LCM Operations

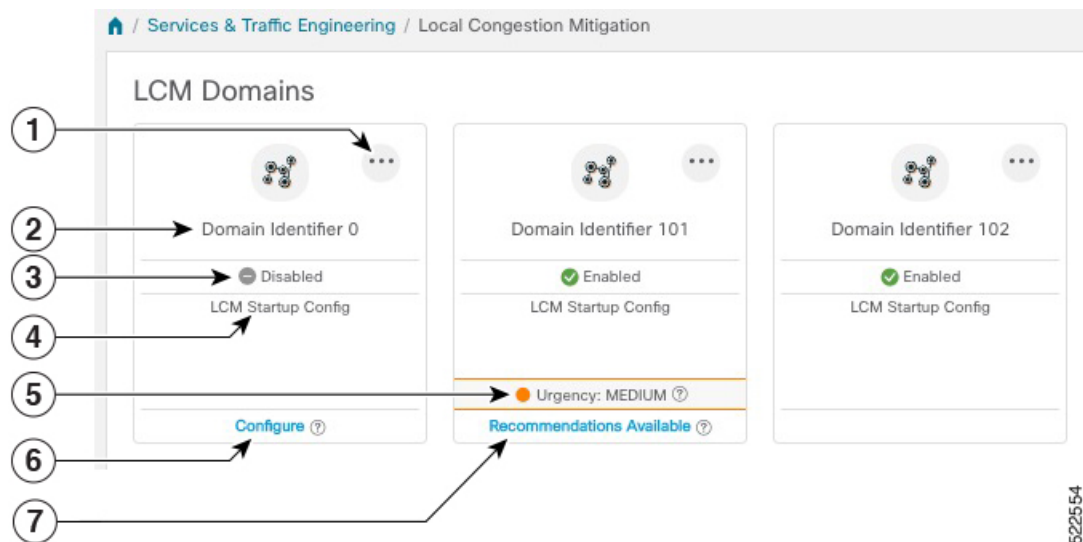


Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the [Workflow Example: Mitigate Congestion on Local Interfaces, on page 6](#) topic.

LCM Domains Dashboard

The LCM Domain Dashboard (**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 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 had been enabled for the domain.

Callout No.	Description
4	LCM Configuration Description: The description is defined in the Configure LCM page. The default description is "LCM Startup Config".
5	Urgency: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following: <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 15.
6	Configure: This link appears if LCM has not yet been configured. Click Configure to go to the Configure LCM page.
7	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 .

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. You can also preview TTE policies prior to deployment (**... > Preview Solution**) or to verify deployment (**... > View Deployed Policies**) visually on a topology map. Hover the mouse pointer over  to view a description of what type of information each column provides. 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

Operational Dashboard

● Congested Interfaces (2) | ● Mitigating Interfaces (0) | ● Mitigated Interfaces (1)

Commit All Urgency: MEDIUM

Node	Interface	Threshold Utilization	Evaluation Utilization	LCM State	Policies Deployed	Policy Set Status	Recommended Action	Commit Status	Expected Utilization	Solution Update Time	Actions
F3.cisco.com	GigabitEthernet0/0/0/1	13%	4.59%	Mitigated	2	OK	Delete Set	None	8.14%	20-Apr-2022 09:07:44 PM PDT	...
F3.cisco.com	GigabitEthernet0/0/0/2	10%	13.47%	Congested	0	-	No Solution	None	-	25-Apr-2022 04:32:10 PM PDT	...
F5.cisco.com	GigabitEthernet0/0/0/1	11%	13.56%	Congested	0	-	Create Set	None	6.78%	20-Apr-2022 08:52:43 PM PDT	...

In this example, the following information is conveyed:

- f3.cisco.com with interface GogabitEthernet0/0/1—The current LCM state is Mitigated and 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.
- f3.cisco.com with interface GogabitEthernet0/0/2—LCM detects congestion however it cannot find bypass policies to remediate the congestion (**Recommended Action - No Solution**).



Note If LCM cannot find a solution (**No Solution**), it may be due to constraints enabled in the **LCM Configuration** page. For more information, see [Configure LCM, on page 15](#).

- f5.cisco.com with interface GogabitEthernet0/0/1—LCM detects congestion and has recommended to deploy policies to remediate the congestion (**Recommended Action - Create Set**).

Recommendations are listed as part of a set, and if deployed, all changes are committed. You must click **Commit All** if you want to remediate the congestion on F5.cisco.com with interface GogabitEthernet0/0/1.

