

Local Congestion Mitigation (LCM)

- Local Congestion Mitigation, on page 1
- Requirements for LCM congestion evaluation , on page 2
- Requirements for LCM congestion mitigation , on page 3
- Important considerations when using LCM, on page 4
- LCM calculation workflow, on page 6
- Example: Mitigate congestion on local interfaces, on page 8
- Configure LCM, on page 14
- 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

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

- Monitor congestion as defined by the interface thresholds you specify.
- Visually preview LCM recommendations on your network before you decide whether to commit the Tactical Traffic Engineering (TTE) SR policy deployment.
- Enable LCM to delete any down, failed, or uncommitted LCM TTE policies when there is an imminent risk of network failures based on LCM solution configurations. For more information, see the advanced configuration options (**Auto Repair Solution** and **Adjacency Hop Type**) in Configure LCM, on page 14.

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

Note Take a look at the Example: Mitigate congestion on local interfaces, on page 8 to see how to use LCM in your network.

Requirements for LCM congestion evaluation

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

To ensure LCM is receiving these traffic statistics,

- enable SNMP or gNMI on the devices whose traffic you want to monitor, including the headend device. For more information on how to configure these protocols, see the specific device platform configuration guide, for example, Configuring SNMP Support
- confirm that all the devices are reachable from the Crosswork Data Gateway, and
- enable strict SID labels on all devices within an LCM domain. See Enable strict SID for LCM usage, on page 2.

Enable strict SID for LCM usage

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

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

Procedure

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

Example:

Cisco IOS XR with ISIS:

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

Example:

Cisco IOS XR with OSPF:

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

Example:

Cisco IOS XE:

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

Step 2 Enable segment routing on the headend device and confirm that ***all*** devices

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

Example:

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

Requirements for LCM congestion mitigation

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

Autoroute steering

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

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

For example:

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

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

The *ID* parameter in this command identifies the PCC profile associated with the SR-TE policy that PCE has provisioned. The ID value can be any integer from 1 to 65535, but it must match the profile ID that PCE uses to instantiate the policy. If not, the policy will not be activated. For example, if PCE provisions a policy with profile ID `10`, you must configure `segment-routing traffic-eng pcc profile 10 autoroute force-sr-include` on the headend router to enable autoroute announcement for that policy. See the specific device platform configuration guide, for example, *Segment Routing Configuration Guide, Cisco IOS XE 17 (Cisco ASR 920 Series)*.

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

Equal Cost Multi-Path (ECMP)

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

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

Important considerations when using LCM

Consider the following information when using LCM:

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

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

  You can configure LCM to detect excessive uneven ECMP splitting among parallel TTE SR policies and issue an event to notify. To mitigate the effects of uneven ECMP, the overprovisioning factor is used in LCM. For more information, see Configure LCM.
- LCM assumes traffic in an *existing* SR-TE policy is ineligible for optimization and should not be steered into LCM TTE SR policies. To enforce this assumption, any existing non-LCM SR-TE policies should not use regular Algo-0 prefix SIDs. Any combination of Algo-1 Strict, Flexible Algorithm, or adjacency SIDs is recommended to prevent this traffic from being steered into LCM TTE SR policies.
- When domain interfaces and links are removed (intentionally or unintentionally), the following occurs:
  - If all links in a domain go down (LINK_DOWN state), LCM configuration and the Domain UI card (see Configure LCM, on page 14) will remain available until the links are aged out (after 4 hours). This behavior is intentional as it gives you time to recover domain interfaces and links if this was done by mistake.
  - If you want to force domain removal before links age out, then you can remove links manually from the UI. The domain will remain in a "ready for deletion" status until the last link is removed.
- After an HA switchover, you can manually add missing interfaces that were previously monitored or update domain configuration options after the system is stable. Missing interfaces and other configuration options occur if they were added after the last cluster data synchronization.
- When an SR-PCE goes down, **Local Congestion Mitigation** (LCM) enters a dormant stage. To exit this state, all SR-PCEs must be connected, and their associated topologies fully synchronized with the topology service. LCM will remain dormant until these conditions are met. It is important to note that LCM does not have visibility into the state of the SR-PCE redundancy set.

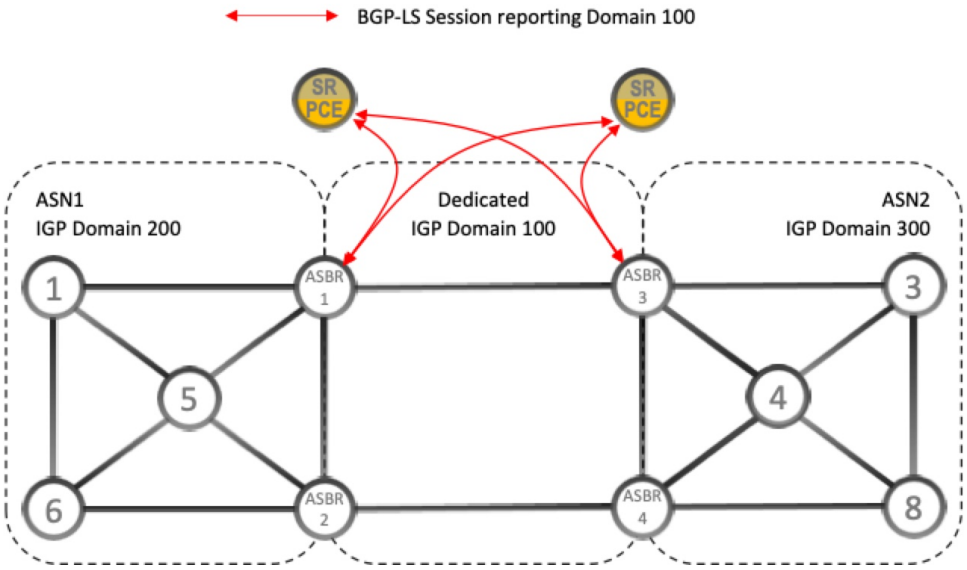
BGP-LS speaker placement for multiple AS networks with a dedicated IGP instance between ASBRs

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

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

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

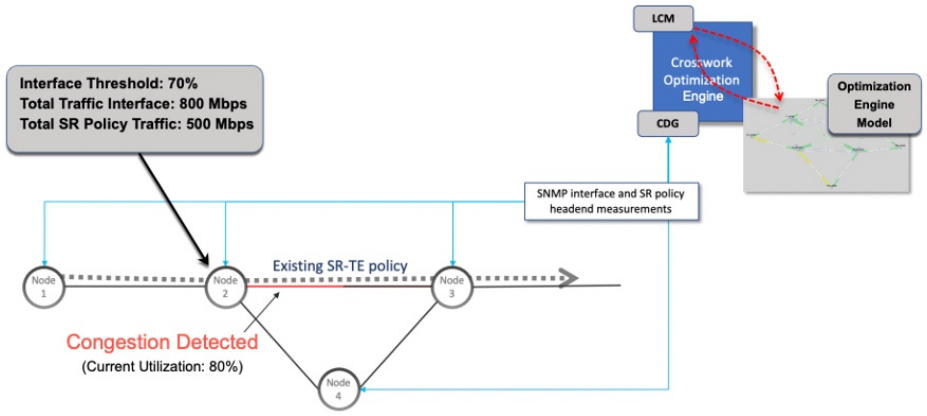
Figure 1: BGP-LS session reporting domain 100



LCM calculation workflow

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

Figure 2: LCM Configuration Workflow Example



Procedure

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

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

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

LCM only diverts traffic that is not already routed on an existing SR policy or RSVP-TE tunnel (for example, unlabeled, IGP routed, or carried via FlexAlgo-0 SIDs). Traffic within an SR-TE policy will not be included in the LCM calculation and will continue to travel over the original programmed path.

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

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

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

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

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

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

LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path. Note that if the Over-provisioning Factor (OPF) percentage is set to 10, then LCM must route 110 (100 Mbps x 1.10) of the eligible traffic. The OPF can be set in the Advanced tab within the LCM Configuration window. For more information, see Configure LCM, on page 14.

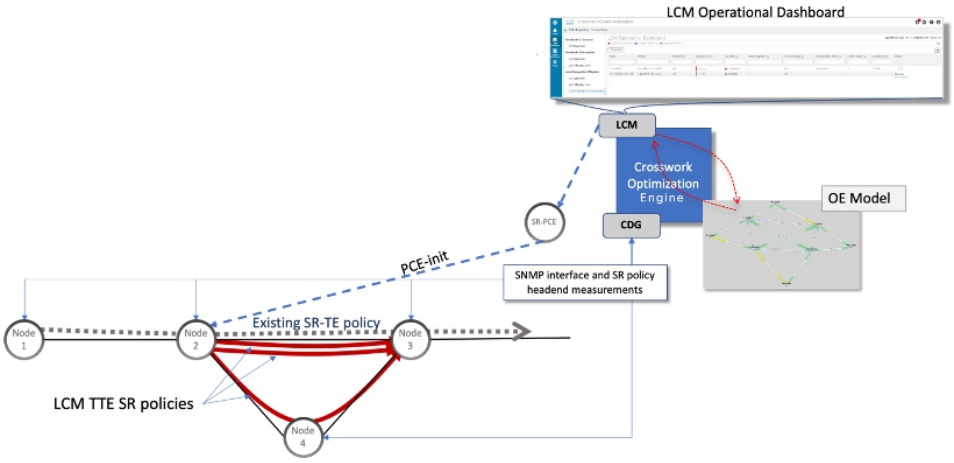
Step 5 LCM determines how many TTE SR policies are needed and their paths. The ratio of how much LCM-eligible traffic can stay on the shortest path to the amount that must be detoured will determine the number of TTE SR policies needed on the shortest versus alternate paths, respectively.

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

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

These recommendations will be listed in the **LCM Operational Dashboard**.

Figure 3: LCM Recommendation Example



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

Example: Mitigate congestion on local interfaces

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

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

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

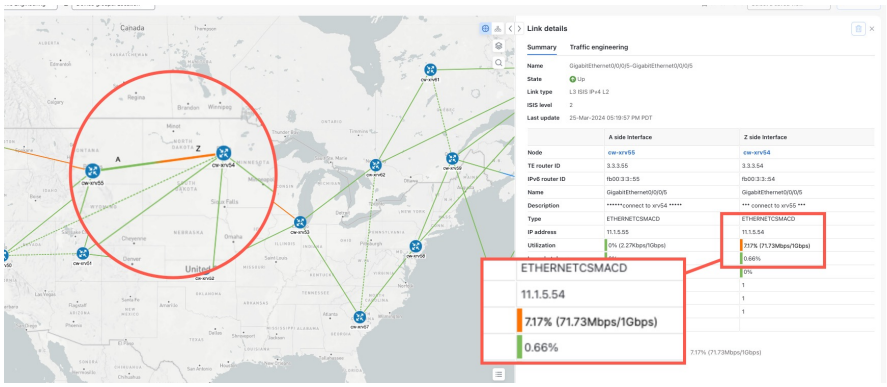
3. Enable and configure LCM in manual mode. Manual mode allows you to view recommended TTE policies prior and decide whether or not to deploy them.
4. After LCM detects congestion, view LCM recommendations on the Operational dashboard.
5. Preview the recommended LCM TTE policies to deploy visually on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

Procedure

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

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

Figure 4: Initial utilization



Step 2 Define any individual interface thresholds.

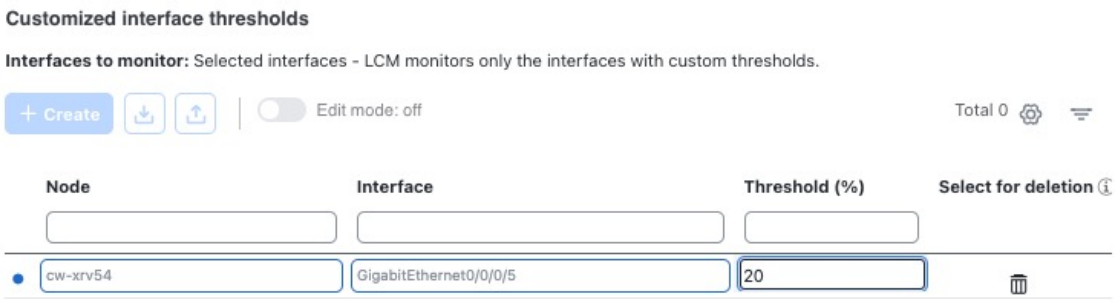
LCM allows you to configure a *global* utilization threshold that can be used for all interfaces. When traffic utilization surpasses the threshold, LCM will try to find bypass policies to remediate the congestion. You set the global utilization threshold on the **LCM Configuration** page. However, if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

a) In this example, we will define an individual interface threshold. Go to the **Customized interface thresholds** page (**Services & Traffic Engineering > Local Congestion Mitigation > *Domain-Identifier* > ⋯ > Interface thresholds**). You can add interfaces individually or upload a CSV file with a list of nodes and interfaces with custom utilization thresholds. For more information, see Add individual interface thresholds, on page 19.

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

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

Figure 5: Customized interface thresholds



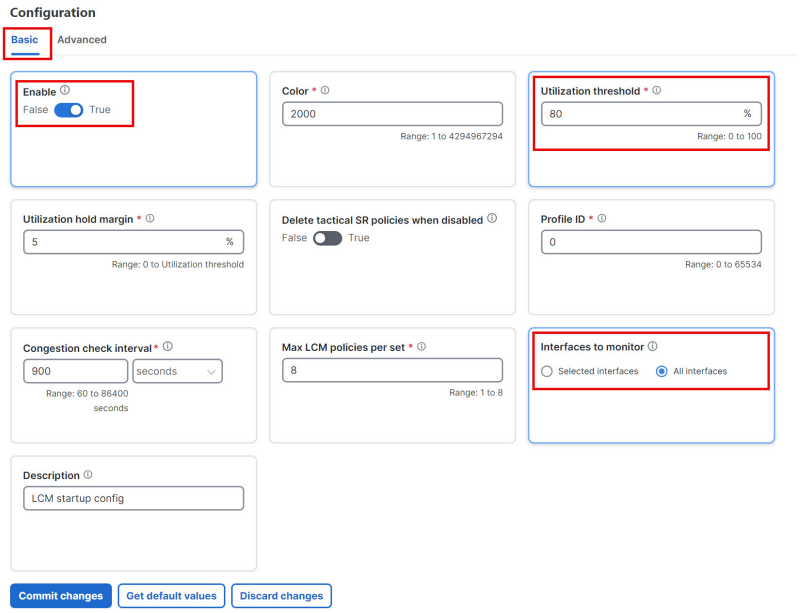
Note
By default, LCM monitors all interfaces. This includes any individual thresholds that are imported to this page. The rest of the interfaces will be monitored using the global **Utilization threshold** defined on the **LCM Configuration** page.

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

Step 3 Enable LCM and configure the global utilization thresholds.

a) From the main menu, choose **Services & Traffic Engineering > Local Congestion Mitigation >** ***Domain-Identifier*** and click **Configuration**. Toggle the **Enable** switch to **True** and configure other LCM options. In this example, the global threshold is set at 80%, and the **Interfaces to monitor > All interfaces** option is selected.

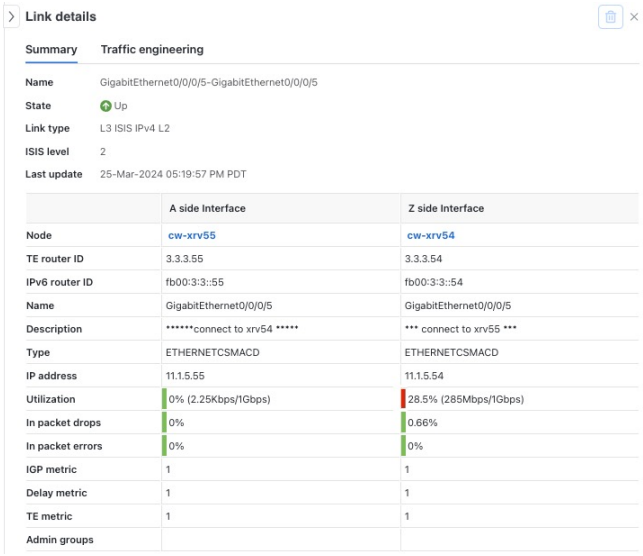
Figure 6: LCM Configuration page



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

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

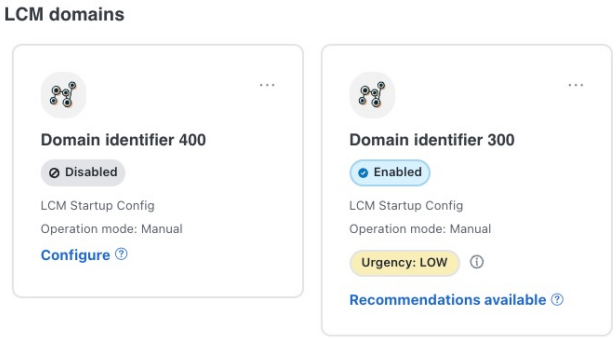
Figure 7: Observed congestion



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

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

Figure 8: Congested detected and LCM recommendations



b) (Optional) View LCM events.

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

c) Open the **Operational dashboard** (**Services & Traffic Engineering > Local Congestion Mitigation >** ***Domain-Identifier*** **>** > **Operational dashboard**).

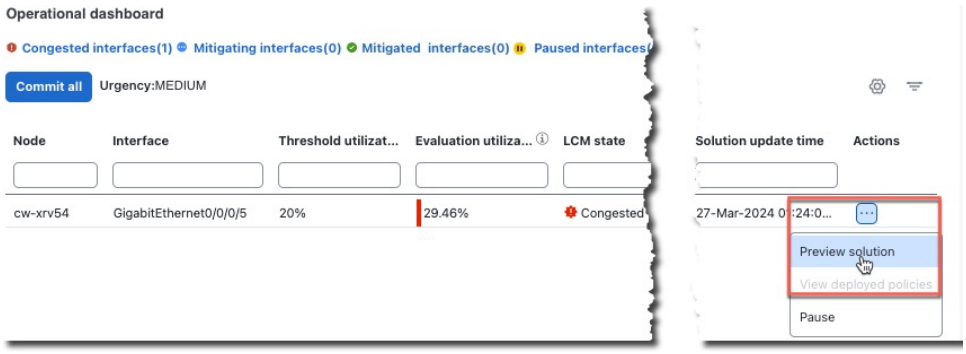
The dashboard shows that cw-xrv54 utilization has surpassed 20% and is now at 29.46%. In the **Recommended action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended action - Create set**) to address the congestion on the interface. For more information, see Monitor LCM operations, on page 21.

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 14).

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

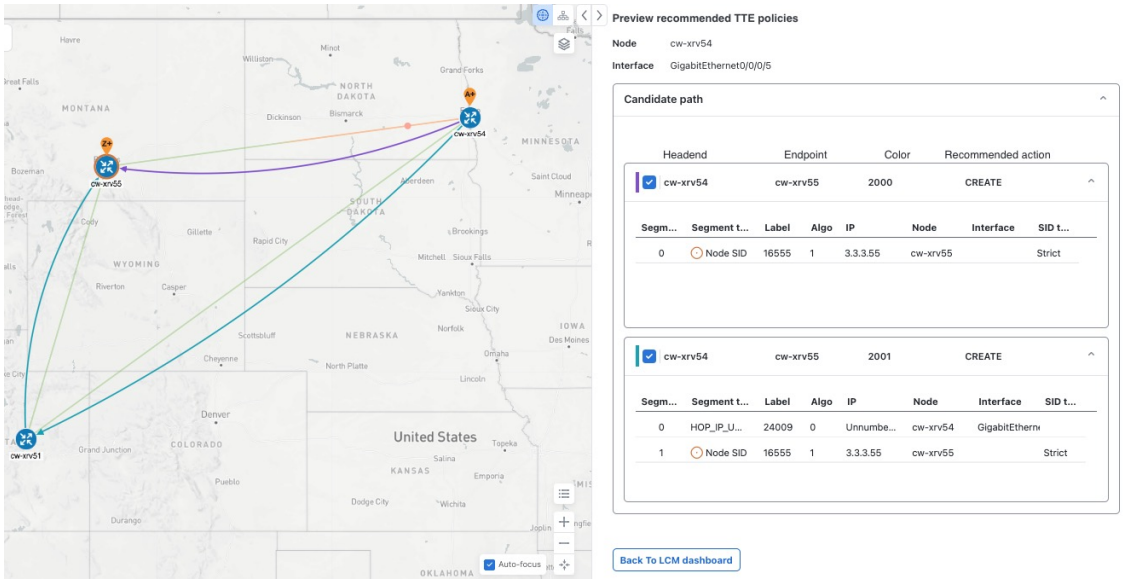
Figure 9: Preview solution



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

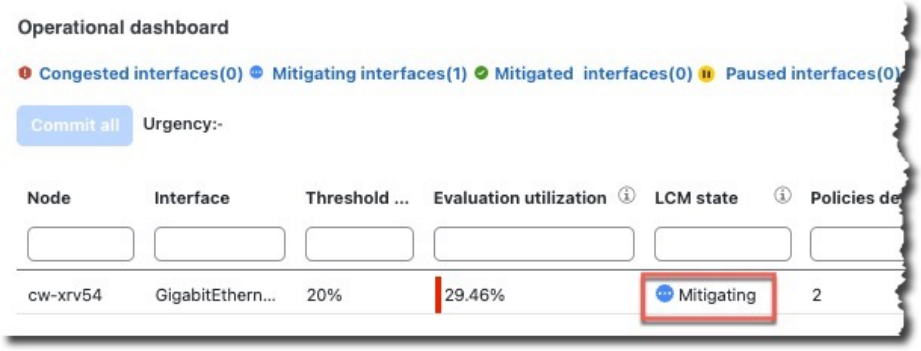
The following figure shows the recommended TTE policies for node cw-xrv54.

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

Step 6 Validate TTE SR policy deployments.

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

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

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

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

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

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

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

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

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

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

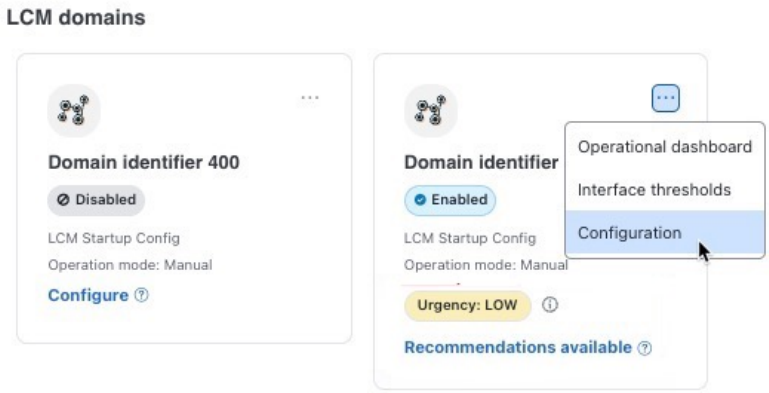
Configure LCM

To enable and configure LCM:

Procedure

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

Figure 12: LCM Configuration



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

Step 3 Enter the required information. Hover the mouse pointer over ⓘ to view a description of each field. See Additional information on LCM configuration options, on page 14.

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

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

Additional information on LCM configuration options

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

- Basic LCM configuration

- Advanced LCM configuration

Basic LCM options

Table 1: Basic configuration options

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

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

Advanced LCM options

Table 2: Advanced configuration options

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

Option	Description
Over-provisioning factor (OPF)	This option helps address unequal ECMP traffic distribution (elephant flows). This value determines the percentage of extra traffic that should be accounted for when computing a path for a by-pass policy. If LCM needs to divert *i* amount of traffic due to congestion, then it will search for a path that can support *x* * (1 + OPF) traffic. For more information, see LCM calculation workflow, on page 6. The default value is 0.
Uneven ECMP traffic threshold	The percentage of sensitivity to detect uneven amounts of traffic across solution bypass tunnels.
Debug optimizer	Enable debug optimizer to log plan files to the Crosswork Network Controller file system. Files are saved to the maximum number of files you specify in **Debug opt max plan files**.
Maximum segment hops	Prior to using this option, you must create device tag groups to which you want to assign certain MSD values. For information on creating tags and assigning them to devices, see the *Cisco Crosswork Network Controller Administration Guide* When calculating bypass TTE policies, LCM uses the effective Maximum SID Depth (MSD) value (as entered here) for specified device tags. You can assign up to five device tags with specific MSD values. A **0** value will not result in a solution. Setting a **0** value is equivalent to LCM monitoring and indicating when there is congestion in the network without providing a recommendation. The system learns from SR-PCE the MSD for each platform advertising the hardware limit in the IGP and BGP-LS. It represents the hardware limit that can be imposed exclusive of any service/transport/special labels. Therefore, you may want to use this new option to assign less than the advertised MSD value that LCM can use for bypass TTE policy calculation. To view the MSD value for a device, navigate to the **Traffic Engineering** topology map and click on the device. From the **Device details** page, click **SR-MPLS** > > **Prefixes** > **Expand all**.
Affinity	You can configure LCM to include or exclude links by using affinities to route data based on specific criteria. For example, if an affinity is excluded, LCM will try to alleviate a congested link by diverting traffic using paths that do not have that affinity. Affinities must already be configured on devices and then mapped using the Crosswork Network Controller UI in order to see the list of affinity names. See Example: Cisco IOS-XR affinity configuration, on page 18 and Configure link affinities, on page 17.

Configure link affinities

Link affinities are attributes or tags associated with links. Link affinities help in directing traffic along preferred paths based on specific criteria, such as bandwidth availability, latency, or cost. The affinity configuration on interfaces simply turns on some bits. It is a 32-bit value, with each bit position (0–31) representing a link

attribute. Affinity mappings can be colors representing a certain type of service profile (for example, low delay, high bandwidth, and so on). Crosswork Network Controller sends bit information to the SR-PCE during provisioning.

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

Procedure

Step 1 Configure affinities on your devices. See Example: Cisco IOS-XR affinity configuration, on page 18.

Step 2 Add affinities in Crosswork Network Controller, on page 18.

Step 3 Configure LCM, on page 14 using the advanced affinity option.

Example: Cisco IOS-XR affinity configuration

There are different ways to apply affinity configurations on a device. See Segment Router configuration documentation for your specific device to view descriptions and supported configuration commands.

Cisco IOS-XR affinity configuration example:

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

Add affinities in Crosswork Network Controller

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

Before you begin

Configure and note down the affinities on your devices.

Procedure

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

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

Step 3 Enter the name and the bit it will be assigned. For example:

Figure 13: Affinity

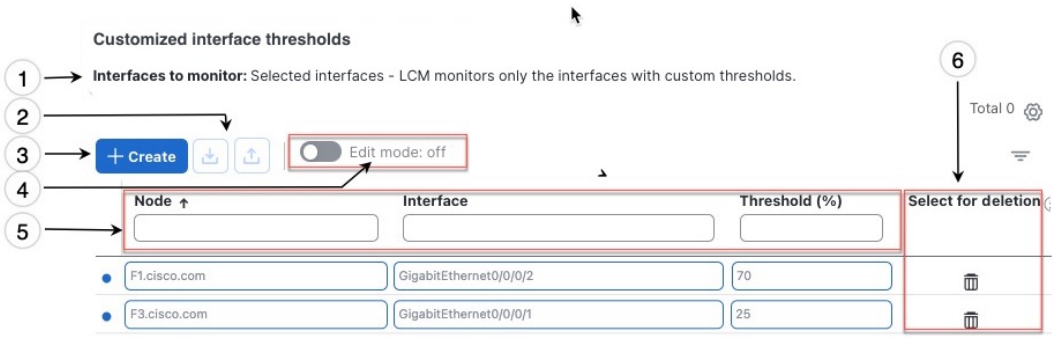


Step 4 Click **Save**. To create another mapping, you must click **+ Create** and save the entry.

Add individual interface thresholds

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

Figure 14: Customized interface thresholds



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

Callout No.	Description
6	**Select for deletion**: Click to delete the row. When **Edit mode** is **ON**, you can check multiple rows to delete, then click **Save**.

To assign specific threshold values for individual interfaces when using LCM, do the following:

Procedure

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

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

Step 2 If you import a CSV file:

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

Step 3 If you manually add individual interfaces:

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

Figure 15: Add first interface



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

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

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

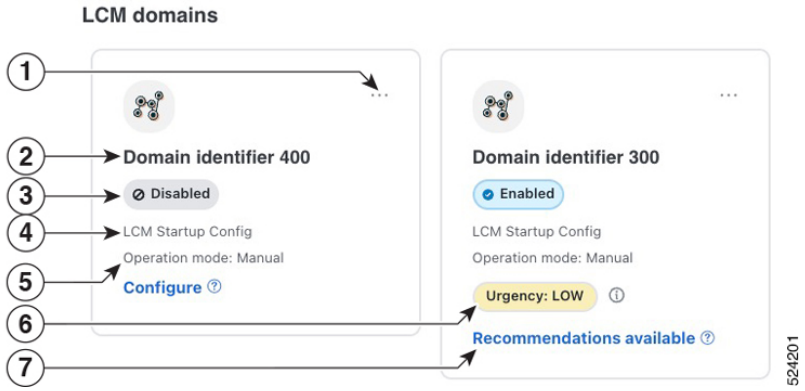
Monitor LCM operations

Note This topic describes how to use and configure the LCM Domain Dashboard and the LCM Operational Dashboard to monitor LCM operations. For information on how to use LCM in your network, see the Example: Mitigate congestion on local interfaces, on page 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 16: LCM Domains Dashboard



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

Callout No.	Description
6	**Urgency**: Indicates the importance of the recommendation deployment or action. Urgency values can be one of the following: • **Low**—Indicates that LCM instantiated policies can be removed because they are no longer needed or that no changes are required. • **Medium**—Indicates new or modified recommendations. • **High**—Indicates network failures and recommendations should be deployed. This is a candidate that can be addressed automatically if the **Auto Repair Solution** advanced option was enabled. See Configure LCM, on page 14.
7	**Configure**: This link appears if LCM has not yet been configured. Click **Configure** to go to the LCM Configuration page. **Recommendations available**: This link appears if LCM has detected congestion and has TTE policy recommendations. To view LCM recommendations, click the link to go to the **LCM Operational dashboard**. **Delete**: Indicates that the domain card can be removed from LCM monitoring.

LCM Operational Dashboard

The LCM Operational Dashboard (**Services & Traffic Engineering > Local Congestion Mitigation >** ***Domain-ID*** **>** ··· **> Operational Dashboard**) shows congested interfaces as defined by the configured utilization threshold.

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

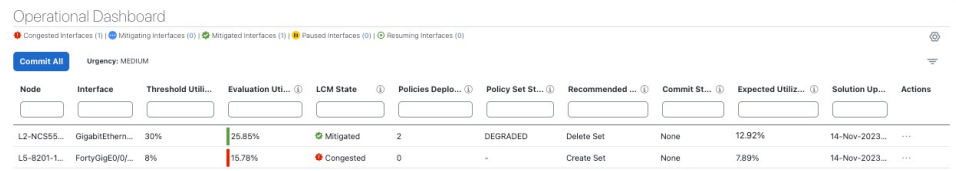
From the Actions column, you can do the following:

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

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

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

Figure 17: LCM Operational Dashboard



In this example, the following information is conveyed:

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

Note If LCM cannot find a solution (**Recommended Action - No Solution**), it may be due to constraints enabled on the **LCM Configuration** page. For more information, see Configure LCM, on page 14.

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

Temporarily exclude an interface from LCM

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

- Where deployed solutions do not result in the intended resolution
- When there is uneven ECMP traffic
- When there are policies that do not carry traffic
- When an interface is continuously throttling between different solutions

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

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

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

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

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

Figure 18: Pause interface

