Configure RSVP-TE Routing

Dynamic LSP routing and CSPF

If an RSVP LSP contains no LSP paths (also called *MPLS TE tunnel paths*), it is routed dynamically using Constrained Shortest Path First (CSPF). The weights used for the CSPF calculation are the TE metrics per interface. If the TE metric is not configured for an interface, then IGP metric is used.

If there are equal-cost routes, CSPF uses these selection criteria in order:

- Bandwidth available: The route with the greatest reservable bandwidth is chosen.
- Hop count: If bandwidth is equal, the route with the fewest hops is chosen.
- Random: If neither bandwidth nor hop count differentiate the routes, a route is chosen at random.

CSPF properties are set in the Edit LSP page and are viewable from the LSPs table. Note that these properties are available only for RSVP LSPs.

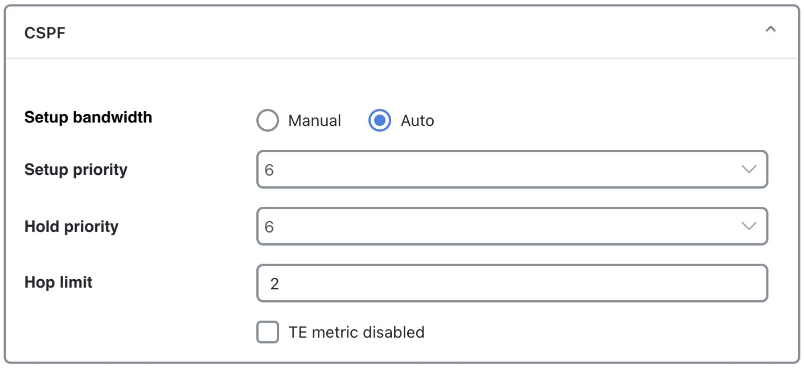

- **Setup bandwidth** (Manual): The amount of traffic (in Mbps) the source node requests for this LSP. The requested bandwidth is available from the reservable bandwidth of each interface in the path.
- **Setup bandwidth** (Auto): Automatically updates the Setup BW Sim value when using the Auto bandwidth convergence mode.
- **Setup priority**: Sets the priority for allocating reservable bandwidth. This is the order in which LSPs are signaled. The lower the number, the higher the priority.
- **Hold priority**: Determines which LSPs can preempt others on the shortest path. This is typically set for a particular service or traffic type. The lower the number, the higher the priority.
- **Hop limit**: Specifies the maximum number of hops permitted in the LSP route. If no path is available with this number of hops or fewer, the LSP is not routed.
- **TE metric disabled**: Uses IGP metrics for the LSP route if checked. If unchecked (default), the LSP uses TE metrics.

Configure interface MPLS properties

Use this task to update the MPLS properties on one or more interfaces or circuits to control reserved bandwidth, traffic engineering metrics, and LSP routing behavior.

MPLS properties are set in the MPLS tab of the Edit interface or circuit page, and are viewable in the Interfaces or Circuits table.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 Select one or more interfaces or circuits from their respective tables.

Step 3 Click .

Note

If you are editing a single interface or circuit, you can also use the ··· > **Edit** option.

Figure 1: Interface MPLS properties

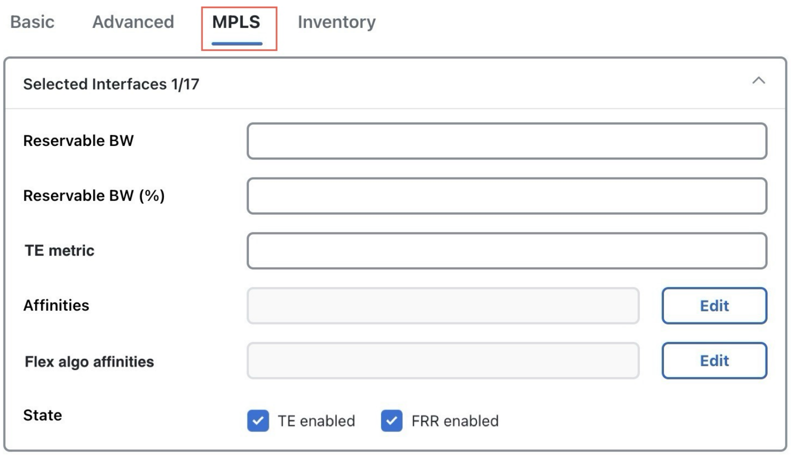

Step 4 Set these properties:

- **Reservable BW**: Shows the amount of bandwidth that can be reserved by LSPs routed over the interface.
- **Reservable BW (%)**: Shows the percentage of bandwidth that can be reserved by LSPs routed over the interface.
- **TE metric**: Determines which path an LSP takes.
- **Affinities**: Lists which affinities are set. For more information, refer to Affinities, on page 13.
- **Flex algo affinities**: Lists the FlexAlgo affinities assigned to the interface (MPLS).
- **State**:
  - **TE enabled**: Identifies whether an LSP can be routed over the interface. If checked, the value is set to True, and TE metrics are used in routing the LSP, provided the LSP's **TE metric disabled** field is unchecked.
  - **FRR enabled**: Designated interface to be avoided by the FRR LSP. Only interfaces with this property set are used by the FRR LSPs initializer when creating FRR LSPs. This property does not affect FRR LSPs that are manually created. For more information, refer to Fast Reroute (FRR) simulations, on page 19.

Step 5 Save your changes.

These columns are updated in the Interfaces or Circuits tables:

- **TE metric sim**: Derived column that shows the effective TE metric. If the TE Metric is empty, this is set to the IGP metric.
- **Resv BW sim**: Derived column.
  - If a Reservable BW value is entered, this is copied to the Resv BW sim column.
  - If Reservable BW and Reservable BW (%) are "na", the Resv BW sim is copied from the Capacity sim column.
  - If Reservable BW is "na", but Reservable BW (%) has a value, the Resv BW sim value is derived by this formula:

```
Capacity sim * (Reservable BW (%) / 100)
```

If needed, you can change this behavior so that reserved bandwidth constraints are not used during simulations, which could be useful for planning purposes. For information, refer to Ignore reservable bandwidths for capacity planning, on page 44.

- To determine how much of the interface traffic is contained in LSPs, show these columns in the Interfaces table. The total of the two equals **Traff sim**.
  - **Traff sim LSP**: The total amount of LSP traffic on the interface.
  - **Traff sim non LSP**: The total amount of non-LSP traffic on the interface.

  To determine which interfaces carry LSP traffic, sort the columns. To determine the demands associated with one or more selected interfaces based on whether or not they have LSP traffic, select them and choose > **Filter to demands** > **Through all interfaces**. Then, choose > **Filter to LSPs**.
- **Loadshare**: Show this column in the LSPs table. It indicates the value for redistributing either traffic across LSPs based on the ratio of this value with other LSPs that are parallel (have the same source and destination).

View LSP reservations

Use this task to view LSP reservations on interfaces and compare reservation utilization with simulated traffic in the network plot.

In the LSP reservations mode, interfaces are colored according to the percentage of bandwidth (in comparison to what is reserved for the LSP).

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the **Plot view** drop-down list in the toolbar, choose **LSP reservations**.

The network plot updates, displaying interfaces color-coded by the percentage of bandwidth reserved for LSPs. The utilization colors are based on the LSP util column in the Interfaces table. LSP util is the sum of all setup bandwidths for all LSPs through the interface, as a percentage of the Resv BW sim.

Figure 2: Simulated traffic versus LSP reservations, on page 5 shows how the network plot changes based on the Plot view options you select, such as Simulated traffic or LSP reservations.

Figure 2: Simulated traffic versus LSP reservations

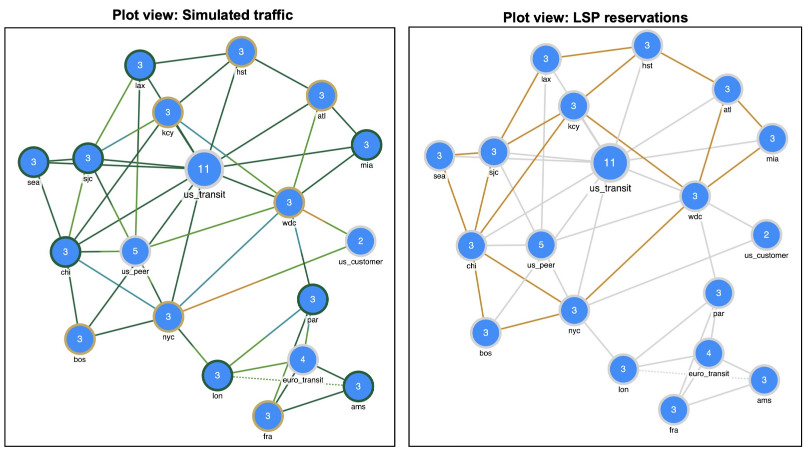

Calculate LSP setup bandwidth

Use this task to calculate setup bandwidth values for selected RSVP-TE LSPs by running the LSP setup bandwidth initializer.

The LSP setup bandwidth initializer calculates the setup bandwidth for existing RSVP-TE LSPs. The setup bandwidth is set to the maximum amount of demand traffic passing through the LSP for the traffic levels selected.

If using only one traffic level, the best practice is to use the **Autobandwidth convergence** mode ( ) > **Simulation** > **Simulation convergence mode**). For more information, refer to Autobandwidth-enabled LSPs, on page 26.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Initializers** > **LSP setup bandwidth**.

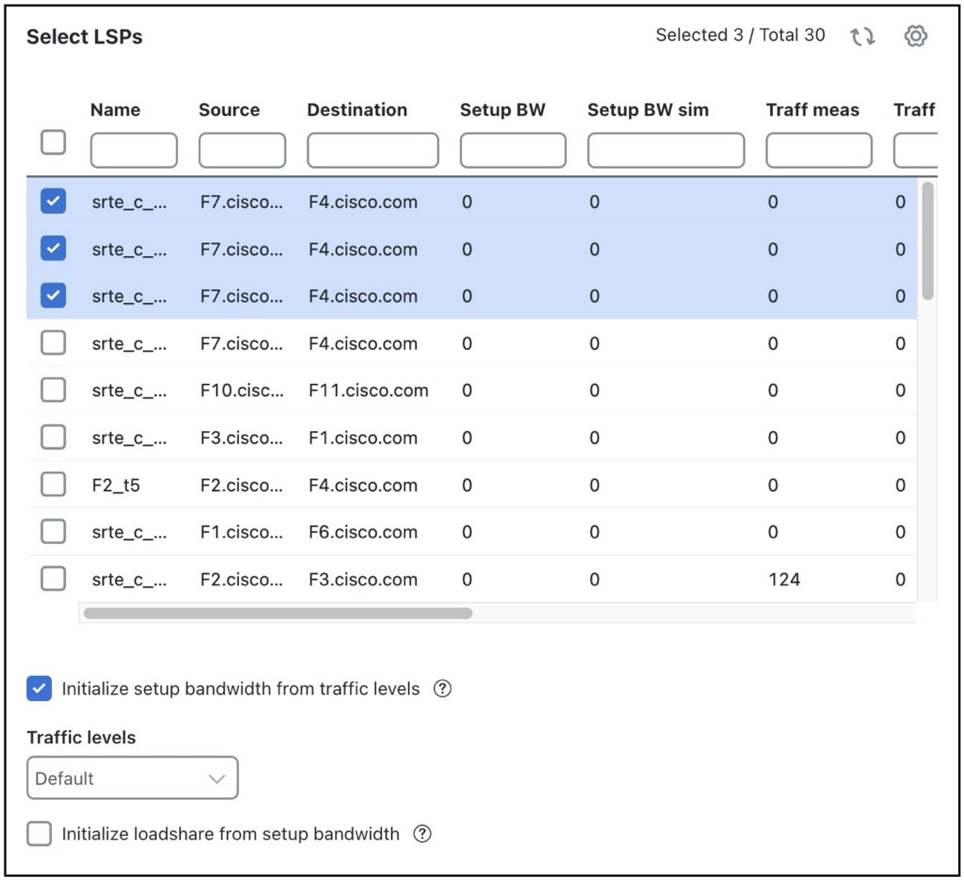

Step 3 Select the LSPs you want to optimize.

Step 4 Choose whether to set the setup bandwidth value to be equal to the maximum amount of traffic passing through each tunnel across selected traffic levels. Use the **Initialize setup bandwidth from traffic levels** check box for this purpose.

Step 5 Choose which traffic levels to use for demand traffic.

Step 6 Choose whether to set the load share values to be equal to the setup bandwidth. To set, check the **Initialize loadshare from setup bandwidth** check box.

Step 7 Submit your changes.

RSVP LSP path attributes and behavior

Like LSPs, LSP paths have CSPF properties, such as **Setup priority** and **Hop limit**. For a description of LSP properties, refer to Dynamic LSP routing and CSPF, on page 1.

Key points

- If the LSP path does not define properties like setup priority or hop limit, it inherits them from the LSP. If these properties are set in the LSP path, they override the LSP settings.

- If operating in the **Autobandwidth convergence** simulation mode and if the LSP's **Setup bandwidth** is set to **Auto**, the LSP path Setup BW property (like the LSP Setup BW) is ignored, and the LSP's Setup BW sim value is calculated and used.
- If an LSP path does not have an associated named path, it is routed dynamically. However, you can use named paths to fully or partially describe a path from the LSP source to destination.

Note This chapter refers to LSPs and LSP paths as *dynamic*, *explicit*, or *loose explicit*, depending on whether their associated named paths have no defined hops, all hops defined, or only some hops defined.

Standby paths

A standby path is an LSP path

- establishes alternative LSP paths for traffic in RSVP-TE networks, and
- reserves setup bandwidth for these alternate paths to ensure immediate availability during failure.

Standby paths are always established even if the associated LSP is routed using a path with a different path option.

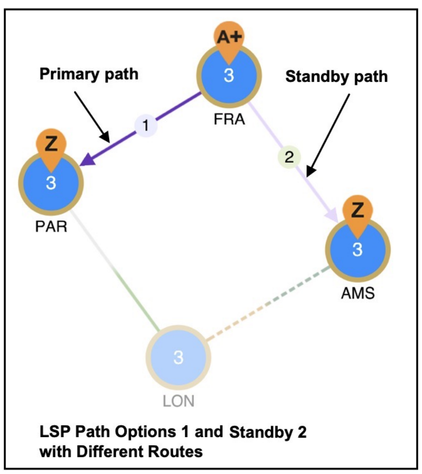

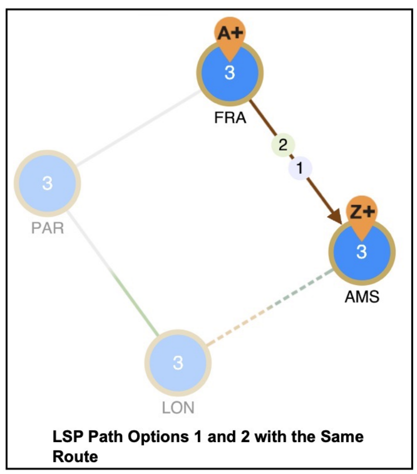

How LSP paths respond to failures

Summary

Figure 3: Example RSVP LSP and associated LSP paths, on page 8 shows an RSVP LSP (cr1.lon_cr2.fra) that uses four LSP paths. The cr1.lon_cr2.fra LSP has a primary and standby secondary path options (100 and 200), each with a setup bandwidth of 0. They are both established, using their defined named paths. Both primary and secondary LSP paths have associated named paths, which appear in the **Path name** column. The other two paths do not have associated named paths, and so are dynamically routed.

Workflow

Figure 3: Example RSVP LSP and associated LSP paths

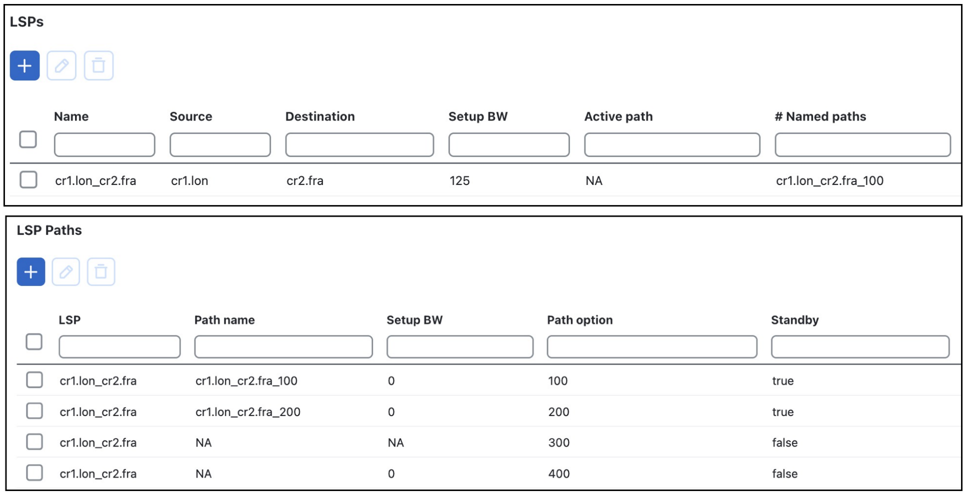

LSPs

Name	Source	Destination	Setup BW	Active path	# Named paths
cr1.lon_cr2.fra	cr1.lon	cr2.fra	125	NA	cr1.lon_cr2.fra_100

LSP Paths

LSP	Path name	Setup BW	Path option	Standby
cr1.lon_cr2.fra	cr1.lon_cr2.fra_100	0	100	true
cr1.lon_cr2.fra	cr1.lon_cr2.fra_200	0	200	true
cr1.lon_cr2.fra	NA	NA	300	false
cr1.lon_cr2.fra	NA	0	400	false

This example shows how LSP responds to failures:

1. If the primary path (cr1.lon_cr2.fra_100) fails, the secondary path (cr1.lon_cr2.fra_200) is used.
2. If both paths with associated name paths fail, a dynamic path is established with a reserved bandwidth of 125 Mbps, which it inherits from the LSP. This path is not set up initially because it is not a standby.
3. If all paths fail, the path of last resort is a dynamic path with zero bandwidth. This ensures that there is always an LSP up so that services relying on the LSP are always available.
4. The simulation changes if the LSP Active path has been set. For example, if cr1.lon_cr2.fra_200 is the active path (that is, 200 appears in the LSP's Active path column), Cisco Crosswork Planning uses cr1.lon_cr2.fra_200 first, and then goes through the previous sequence.

Named paths and explicit LSP routing

Named paths specify a route though the network using an ordered list of adjacencies. The route is defined by named path hops, which can be nodes or interfaces. Each hop type specifies whether the route should be strict, loose, or excluded.

- Strict: The LSP must reach the named path hop directly from the previous one, with no intermediate interfaces.
- Loose: The LSP is routed to this hop from the previous hop using CSPF. Intermediate interfaces can be used.
- Exclude: The node or interface is excluded from the LSP path. This hop type cannot be combined with strict or loose hops in the same named path.

The named paths and their hops are listed and selectable from the Named paths and Named path hops tables. The advantage of having separate tables is that you can have named paths that have missing or unresolved hops. Also, the path name can be reserved even if it is not part of the Cisco Crosswork Planning simulation.

The hop type is identified visually on the plot, as well as listed in the **Type** column of the Named path hops table.

Named path hops example

Figure 4: Example named paths the cr1.lon_cr2.fra LSP, on page 9 extends the earlier example (Figure 3: Example RSVP LSP and associated LSP paths, on page 8), showing two named paths for the cr1.lon_cr2.fra LSP. This shows that the first two LSP paths each have a named path. The naming convention is *<LSP_Name>_<Path_Option>*, which in this case is cr1.lon_cr2.fra_100 and cr1.lon_cr2.fra_200.

Figure 4: Example named paths the cr1.lon_cr2.fra LSP

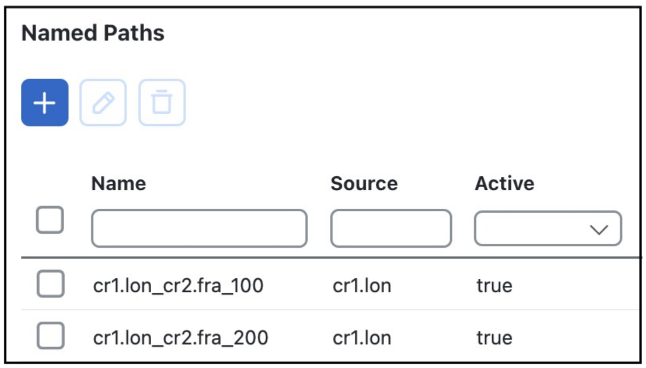
Named Paths

Name	Source	Active
cr1.lon_cr2.fra_100	cr1.lon	true
cr1.lon_cr2.fra_200	cr1.lon	true

The hops for each named path are defined in a separate table called **Named Path Hops**. Figure 5: Example named path hops for cr1.lon_cr2.fra_100, on page 9 shows the named path hops for the cr1.lon_cr2.fra_100 named path. The **Step** column shows the hop order.

- The first hop is an interface hop: a strict hop on the interface from cr1.par to cr2.par.
- The second is a node hop, also strict, to cr2.par.
- The third hop is an interface hop: a strict hop on the interface from cr1.fra to cr2.fra.
- Hops for cr1.lon_cr2.fra_200 are defined similarly.

Figure 5: Example named path hops for cr1.lon_cr2.fra_100

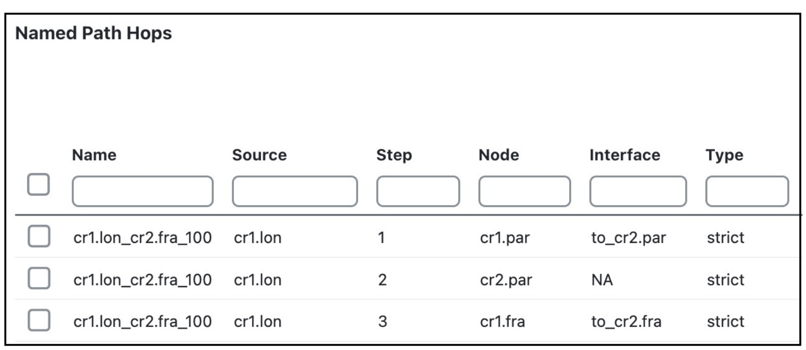
Named Path Hops

Name	Source	Step	Node	Interface	Type
cr1.lon_cr2.fra_100	cr1.lon	1	cr1.par	to_cr2.par	strict
cr1.lon_cr2.fra_100	cr1.lon	2	cr2.par	NA	strict
cr1.lon_cr2.fra_100	cr1.lon	3	cr1.fra	to_cr2.fra	strict

Create named paths and their hops

Use this task to create named paths and define ordered hops for explicit RSVP LSP routing through the network.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 Create named paths using any of these options:

- From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Named path**.
- In the Network Summary panel on the right side, click in the **Named paths** tab.

  The Named paths tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **Named paths** check box.
- When creating LSP paths, check the **Create associated named paths** check box. Using this option does not create hops.

Step 3 Enter these information:

- **Name**: Enter the name of the named path.
- **Source site** and **Source node**: Use these options to change the source site or source node.
- **Active** check box: Use this option to make the named path active.

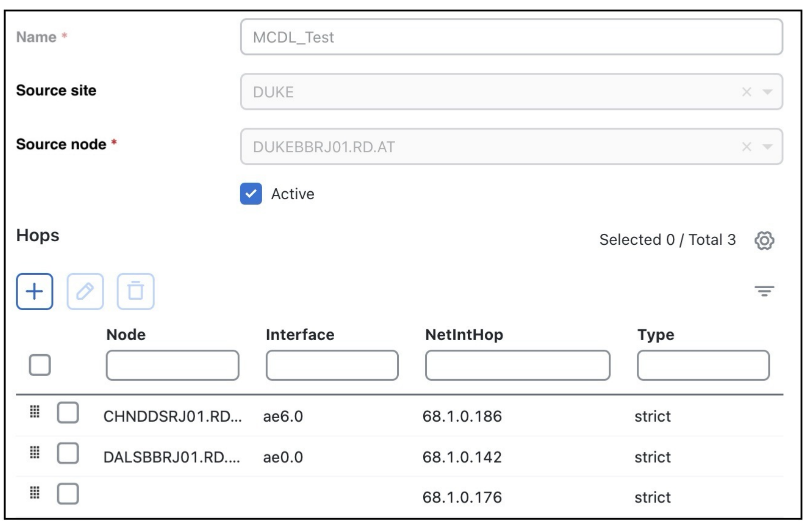

Step 4 Configure the hop options:

- To add a new hop, click .
- To edit an existing hop, select it from the Hops list and then click . The Edit page opens for named path hops.

- To delete an existing hop, choose it from the Hops list, click .

Step 5 To continue creating or editing the named path hop, use these options.

- Choose the site, node, and interface. For node hops, do not choose an interface.
- Choose the **Type** as Loose, Strict, or Exclude. For more information on these options, refer to Named paths and explicit LSP routing, on page 8.

Step 6 Click **Add**.

Edit named paths and their hops

Once the named paths are created, you can create, edit, or delete their named path hops. Note that discovered named paths can contain *unresolved* hops, which are nodes and interfaces that are not in the plan file. For information about resolving named paths, refer to Unresolved LSP destinations and hops, on page 55.

The recommended method of editing the named path hop type is to access it directly from the Named paths table. This is the most efficient way of viewing the entire path.

Procedure

Step 1 Select one or more named paths from the Named paths table.

Step 2 Click .

Note
If you are editing a single named path, you can also use the ··· > **Edit** option.

Step 3 In the **Hop** section, edit the Named path hop details, as required. For more information, refer to Create named paths and their hops, on page 10.

Step 4 Save your changes.

Actual paths

Actual path is a path that

- represents the real routes taken by routed LSP paths
- is read from the live network, per LSP path, and
- reflects current LSP path states, including successful and failed routings.

RSVP-TE routing is not always completely predictable because it depends on the order in which LSPs are established and reserve their bandwidth on the network. Cisco Crosswork Planning uses actual paths in LSP routing simulations to match the current LSP routing state of a network as accurately as possible.

When an LSP path is routed, the actual route is attempted first. If this routing fails due to CSPF constraints, such as not enough reservable bandwidth or affinity restrictions, Cisco Crosswork Planning reverts to the standard CSPF routing algorithm.

Example

Cisco Crosswork Planning routes all LSPs in its simulation according to the actual paths read from the network, and the routings completely match the network. If a circuit fails, only the LSPs through this circuit are rerouted. Because they cannot be established along their actual paths, standard CSPF is used. The result is a prediction of incremental routing changes that would occur on the current network if such a failure were to occur.

An LSP or an LSP path might have a corresponding actual path. An LSP path with no actual path inherits its LSP actual path, if available.

Two columns in the LSPs and LSP Paths tables are useful to refer to the result of actual paths on LSP routing.

Column	Description
Actual column in the **LSP Paths** table	The simulation can only use actual paths if they are resolved. If the network discovery was incomplete, this might not be possible. • true = the simulation converted the actual path hops into a complete path from source to destination of the LSP • false = the simulation did not convert the actual path hops into a complete path • NA = not applicable; no actual path was available.
Routed column in the **LSPs** table	• Actual = The LSP follows an actual path. • Simulated = The LSP does not follow an actual path. • Unrouted = Routing was not possible.

Deactivate actual paths for simulations

Normally, the simulation routes LSPs using actual paths based on current network state and LSP active path settings. For planning purposes, when state is not relevant, you can change this behavior to disregard actual paths.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 3 Click the **Simulation** tab.

Step 4 To use or disregard actual paths and active paths in MPLS simulation, check or uncheck **Use LSP actual paths, active paths** check box.

Step 5 Save your changes.

Affinities

An affinity is a mechanism that

- enables path diversity for LSPs, and
- allows assignment of attributes to physical network circuits and LSPs to implement a variety of routing policies.

For example, you can use affinities to restrict certain traffic to specific topological regions, or you can force primary and backup LSP paths onto different routes so they do not simultaneously fail.

Cisco Crosswork Planning supports an unlimited number of 64-bit affinities, each of which is defined by a number and an optional name. Named affinities are sometimes called *administrative groups* or *link coloring*.

The default plan files have 0-31 unnamed and unassigned affinities.

Create and edit affinities

Use this task to create new affinities or edit existing affinities used in RSVP LSP routing policies.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **Admin groups**.

The Admin groups page opens.

Step 3 To create a new affinity:

a) Click **Add**.
b) In the **Affinity** and **Name** fields, enter the affinity number and affinity name.

   Affinity numbers must be unique. Affinity names are optional and must also be unique.

c) Save your changes.

Step 4 To modify an existing affinity:

a) Select the affinity you want to edit.
b) Click **Edit**.
c) In the **Affinity** and **Name** fields, enter the affinity number and affinity name.

   Affinity numbers must be unique. Affinity names are optional and must also be unique.

d) Save your changes.

Assign affinities to interfaces

Use this task to assign affinities to interfaces so RSVP LSP routing can include or avoid selected network resources during simulations.

To use affinities, you must assign them to interfaces in a way that encourages favorable routes and discourages others. For example, continental paths can have one affinity and international paths can have another.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the Interfaces table, select one or more interfaces to which you are assigning the affinities.

Step 3 Click .

Note
If you are editing a single interface, you can also use the ··· > **Edit** option.

Step 4 Click the **MPLS** tab.

a) Click the **Edit** button next to the **Affinities** field.
b) In the **Include** column, check the boxes for the affinities that you want to associate with the selected circuits.
c) Click **Save**.

Step 5 Save your changes.

You can also use these steps to assign affinities to interfaces through the Edit Circuit page. However, you must choose and assign the affinity twice: once to each interface.

LSP affinity associations

For affinities to influence LSP routing, you must associate LSPs with the affinities assigned to interfaces that you want to include or exclude in the routing process. If no affinity is specified, the LSP can be routed through any path. Explicit routes can also have affinities for hops with a Loose relationship (refer to Named paths and explicit LSP routing, on page 8).

Inclusion and exclusion rules

You can choose any of these rules while associating LSPs with affinities:

- **Include**: Only use interfaces that include all affinities.
- **Include any**: Use interfaces that include at least one affinity.
- **Exclude**: Do not use interfaces with this affinity.

LSP paths inherit the LSP affinities. However, you can edit LSP path affinities to have priority over the LSP affinities.

Example of affinity-based LSP routing

This example demonstrates how affinities affect the routing of two LSPs with the same source and destination nodes.

- The to_cr1.nyc interface between wdc and nyc is assigned to both affinities 1 (Silver) and 2 (Bronze). No other interfaces are assigned affinities.
- LSP A is configured to exclude all interfaces assigned to affinity 2. While it cannot take the shortest path using the to_cr1.nyc interface, it can route around it.
- LSP B is not associated with any affinities. Its LSP path is configured to include any interfaces assigned to affinity 1 (Silver) and exclude those assigned to affinity 2 (Bronze). Since the to_cr1.nyc interface is assigned to both affinities, LSP B cannot route.

Figure 6: Affinity-based LSP routing example

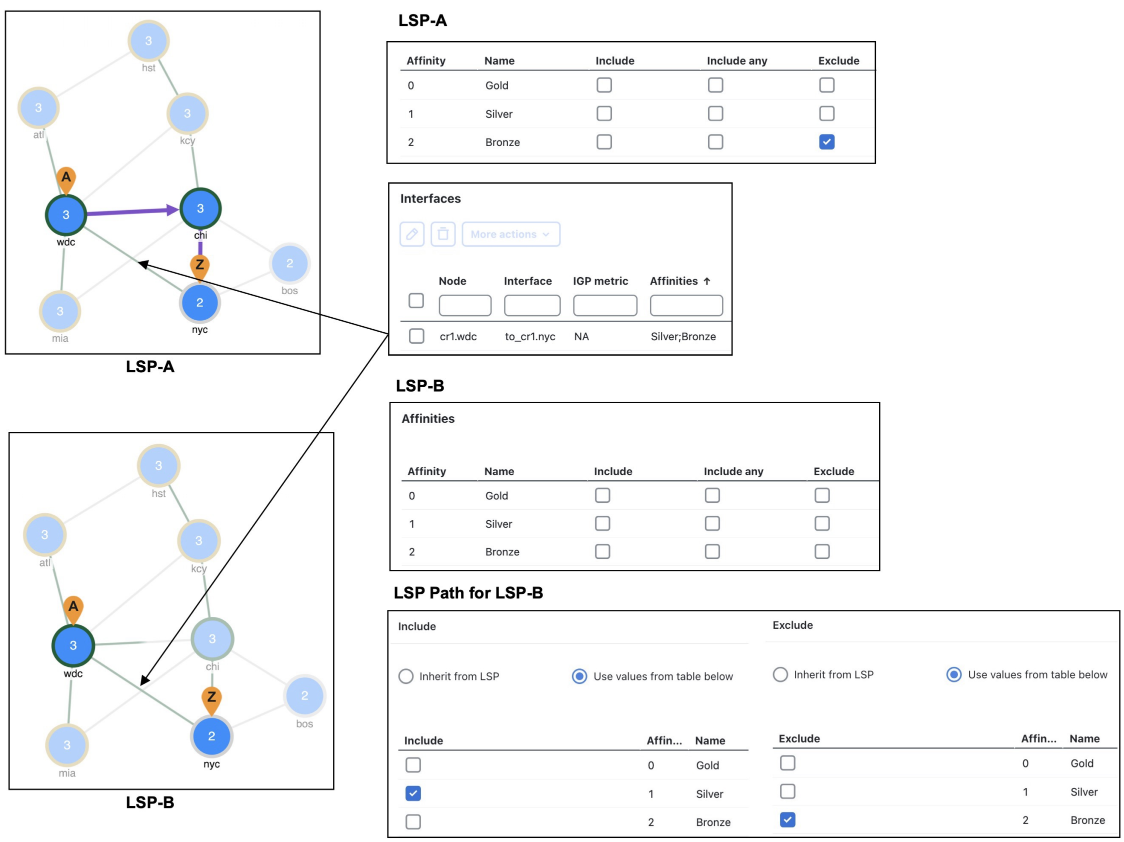

Assign LSPs to affinities

Use this task to associate LSPs with affinities so routing can include or exclude interfaces with matching affinity settings.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the **LSPs** table, choose one or more LSPs to which you are associating the affinities.

Step 3 Click .

Note
If you are editing a single LSP, you can also use the > **Edit** option under the **Actions** column.

Step 4 Click the **Affinities** tab.

Figure 7: LSP affinities

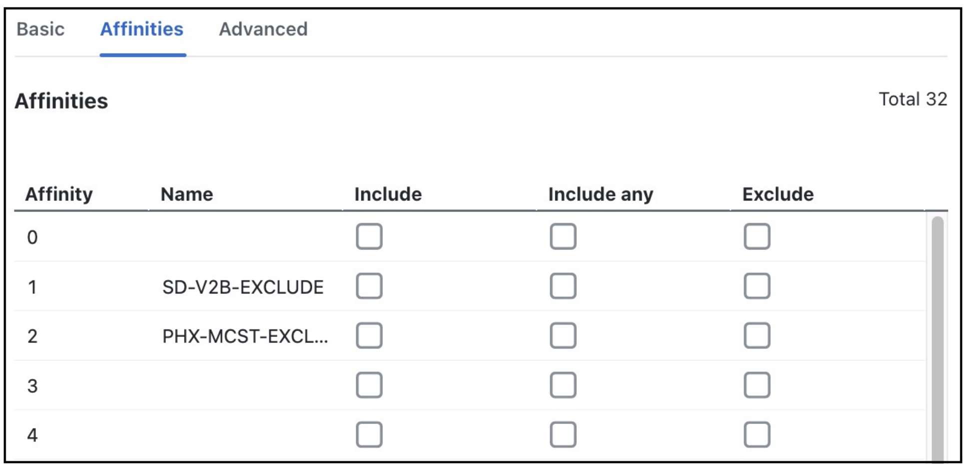

Step 5 Choose the inclusion or exclusion rule (Include, Include any, or Exclude) for the affinity you are associating with the selected LSPs.

Step 6 Save your changes.

Assign LSP paths to affinities

Use this task to associate LSP paths with affinities.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the LSP paths table, choose one or more LSP paths to which you are associating the affinities.

Step 3 Click .

Note
If you are editing a single LSP path, you can also use the ··· > **Edit** option under the **Actions** column.

Step 4 In the **Affinities** section, click the **Edit** button.

Step 5 Choose the inclusion or exclusion rule (Include, Include any, or Exclude) for the affinity you are associating with the selected LSP paths. You can choose the LSP paths to inherit the LSP affinities or choose the values from the table. After making the selections, click **Save**.

Figure 8: LSP path affinities

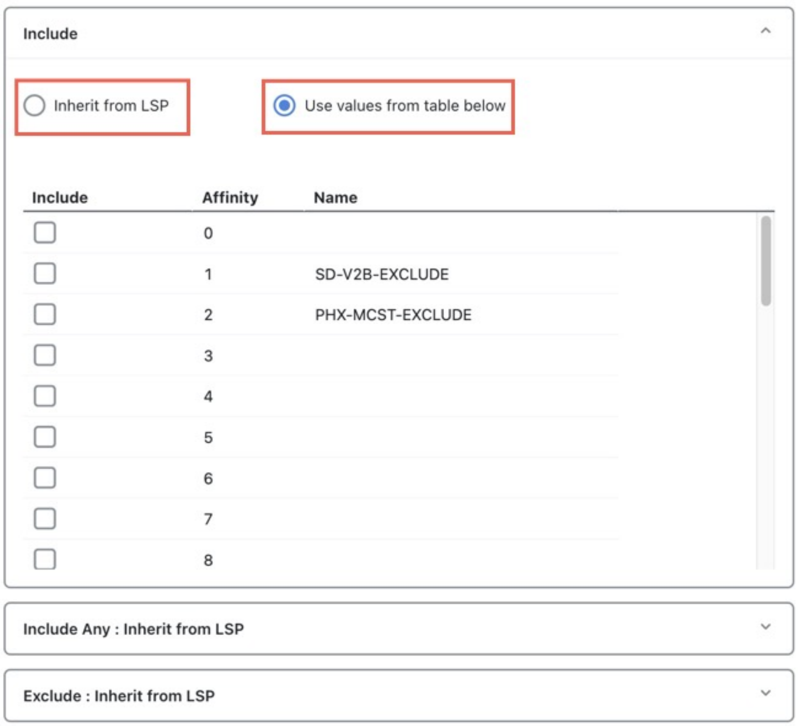

Step 6 Save your changes.

Assign affinities when creating LSP meshes

Use this task to assign affinities while creating LSP meshes so that the generated LSPs use the intended routing constraints.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP mesh**.

Alternatively, in the Network Summary panel on the right side, click > **LSP mesh** in the **LSPs** tab.

Step 3 After selecting the required nodes and options, go to the **Affinities** page.

Step 4 In the **Affinities** section, click **Choose affinities**.

Step 5 Choose the inclusion or exclusion rule for each affinity you are associating with all LSPs in the mesh, and click **Add**.

Step 6 Submit your changes.

Global simulation parameters

Cisco Crosswork Planning lets you set global parameters that affect how LSPs are routed or rerouted.

To access these options:

1. In the toolbar, click or choose **Actions** > **Edit** > **Network options**.
2. Click the **Simulation** tab.

Simulation modes

Cisco Crosswork Planning supports four simulation modes for RSVP-TE LSPs. These are listed as options in the **Layer 3** drop-down list in the **Simulation convergence mode** section.

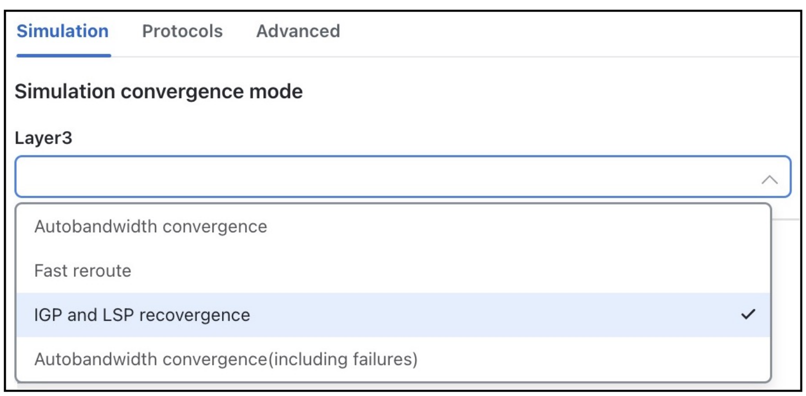

- **Autobandwidth convergence**: In this mode, the network is simulated after the traffic (Traff sim) has been rerouted to other LSPs, but before the Setup BW sim values have been reset. Refer to Autobandwidth-enabled LSPs, on page 26.
- **Fast reroute**: This mode uses FRR LSPs, which is a common failure restoration mechanism used by RSVP-TE LSPs. In a live network, FRR restoration typically occurs within milliseconds. Refer to Fast Reroute (FRR) simulations, on page 19.
- **IGP and LSP reconvergence**: By default, Cisco Crosswork Planning simulates the state of the network once it has fully responded to a failure. Specifically, this is the network state after LSPs have re-established new routes around the failure, and the IGP has fully reconverged. The optimization tools only work in IGP and LSP reconvergence mode.

- **Autobandwidth convergence (including failures)**: In this mode, the network is simulated after the Setup BW sim values have been reset. Refer to Autobandwidth-enabled LSPs, on page 26.

Fast Reroute (FRR) simulations

A Fast Reroute (FRR) simulation is a network analysis technique that

- models the behavior of Label Switched Paths (LSPs) under failure scenarios
- evaluates how locally protected LSPs reroute traffic using pre-signaled FRR LSPs or bypass tunnels, and
- measures restoration performance, often aiming for a 50 millisecond recovery.

Cisco Crosswork Planning simulates a Fast Reroute convergence mode using FRR LSPs. If a failure occurs in the route of a protected RSVP-TE LSP, a presignaled FRR LSP (or bypass tunnel) forwards traffic locally around the point of failure, typically from the node just before the failure to a node downstream. This restoration mechanism gives the source node (head-end) of the FRR LSP time to re-establish an alternate route around the failure. This period is sometimes described as the *50 millisecond* behavior of the network because FRR restoration ideally becomes effective within approximately 50 milliseconds of a failure.

Routing of LSPs and demands in FRR simulations differs from full convergence simulations:

- The source node of the protected LSP does not reroute the traffic. If a failure occurs on a node or circuit (link) within the LSP path, the routing occurs as described in these steps:
  - If an FRR LSP exists to protect a failed node or circuit, the protected LSP continues to be routed, but its new path includes the FRR LSP route that redirects the traffic around the failure. This traffic enters the FRR LSP's source node just prior to the failure and rejoins the protected LSP path after the failure at the FRR LSP's destination node.
  - If the LSP is not protected by an FRR LSP, then the LSP traffic is not routed.
- Since there is no IGP reconvergence, demands through unrouted LSPs are unrouted. Additionally, demands through failures outside of these LSPs are not routed.

FRR fundamentals

This section summarizes key properties and terms, and how FRR LSPs are visualized in the network plot.

Objects and properties

Note FRR LSPs appear in the same table as the regular LSPs and are identified in the **Metric type** column as **FRR**.

Table 1: Objects and properties

Object	Type	Column	Description
LSP	FRR LSP	FRR interface	Designated interface to be avoided by the FRR LSP.
		FRR type	Type of FRR protection: Link or Node.
	Autoroute Forwarding Adjacency	FRR enabled	Marks an LSP to be protected by FRR LSPs.
Interface	na	FRR enabled	Marks an interface to be protected by FRR LSPs. This property is required if an interface is to be used when running the FRR LSPs initializer.

Terms

- Protected interface: An interface that is avoided when using FRR LSPs.
- Protected LSP: An LSP that would be rerouted through FRR LSPs in the event of failures. In Cisco Crosswork Planning, this requires the Fast Reroute simulation mode. To turn this mode on, refer to Run fast reroute simulation, on page 25.
- Protected SRLG: An SRLG that contains a protected interface.

Visualization

An unused FRR LSP path highlights in purple when it is selected from the LSPs table. When routed under failure, FRR LSPs appear as a dotted line.

FRR LSP routing

The source and destination of an FRR LSP can be any pair of nodes, although in practice the source and destination nodes are usually one hop (link protection) or two hops (node protection) apart. The source node for the FRR LSP has a designated interface that is configured so that it is avoided during Fast Reroute simulations. This is the *protected interface*.

This is configured using the **FRR interface** property, which is set in the Edit LSP page when manually creating the FRR LSP, or it is automatically created when running the FRR LSPs initializer. The initializer derives it from interface's FRR enabled property.

Example

This image shows an FRR LSP sourced at sjc, destined for mia, and whose protected interface (FRR interface) is to_cr1.hst. Therefore, the FRR LSP path routes towards okc and away from hst. If this FRR interface property were to_cr1.okc instead, the FRR LSP path would route sjc-hst-mia, away from okc.

Figure 9: FRR LSP routing example

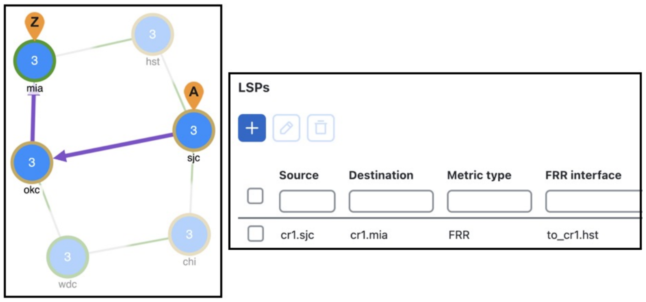

Like other LSPs, FRR LSPs can use actual and named paths. Unlike other LSPs, they do not use setup bandwidth when routing and they only use one LSP path.

Types of FRR LSP protection

Cisco Crosswork Planning supports three main types of protection for FRR LSPs: link protection, node protection, and SRLG protection.

Link and node protection

The type of link and node protection are displayed in the **FRR type** column of the LSPs table.

- Link-protection LSPs: Protect an LSP from a circuit failure by routing an FRR LSP around the failed circuit. The source node of the FRR LSP is the local node of the failed circuit, and if configured correctly, the destination of the FRR LSP is the remote node of the failed circuit.
- Node-protection LSPs: Protect an LSP from a node failure by routing an FRR LSP from the node one hop prior to the failed node in the protected LSP path to the node one hop after the failed node.

SRLG protection

In a network, if interfaces are sourced from the same node (router) and are configured as an SRLG, an FRR LSP protecting one of these interfaces can be configured to reroute around that interface. If possible, the FRR LSP avoids all other interfaces in the SRLG. In Cisco Crosswork Planning, FRR LSPs can be configured to avoid all objects defined in the SRLG, not just interfaces on the source node of the FRR LSP.

Configure routing around failures

When one or more failures occur, these decisions determine which FRR LSP to take:

Before you begin

- LSPs must be configured to be protected by FRR LSPs. For information about setting up FRR LSPs, refer to Set up fast reroute simulations, on page 22.
- Any interfaces along the FRR LSP path must be configured to be protected by FRR LSPs. For information about protecting interfaces, refer to Mark LSPs for protection, on page 22 or the FRR LSPs initializer, on page 23 (which automates the process of setting these).

Procedure

Step 1 For each such interface on the LSP path, Cisco Crosswork Planning checks for an FRR LSP protecting that interface whose destination is a node further down the path of the LSP.

Step 2 If more than one of these FRR LSPs exist, Cisco Crosswork Planning chooses the one that is farthest down that path of the protected LSP, thus simulating standard FRR behavior that prefers node protection over link protection.

If more than one eligible FRR LSP has the same destination farthest down the protected LSP path, the FRR LSP is chosen arbitrarily among them.

Step 3 If more failed circuits or nodes occur after this destination, further FRR LSPs are chosen in the same manner.

Set up fast reroute simulations

Use this task to set up FRR simulations by configuring protection properties, SRLGs, FRR LSP initialization, and simulation execution.

Procedure

Step 1 Set appropriate properties on LSPs to mark them for protection. Refer to Mark LSPs for protection, on page 22.

Step 2 If protecting SRLGs, associate a node with the SRLG in order to establish the protected interface within it. Refer to Identify SRLGs for protection, on page 23.

Step 3 If you are using the initializer to create FRR LSPs, identify which interfaces to protect. Refer to FRR LSPs initializer, on page 23.

Step 4 Create the FRR LSPs using the FRR LSPs initializer or manually. Refer to FRR LSPs initializer, on page 23 or Create FRR LSPs manually, on page 25.

Mark LSPs for protection

Use this task to mark LSPs for FRR protection by setting properties that identify protected interfaces, nodes, or SRLG resources.

Procedure

Step 1 Select the required LSPs from the LSPs table and click .

Step 2 In the **LSPs details** section, check the **FRR enabled** check box.

This appears as "true" in the **FRR enabled** column of the **LSPs** table. If this is "false", the LSP is not protected.

Step 3 In the **Routing** section, set the **Routing type** as **Autoroute** or **FA** (Forwarding Adjacency).

The type appears as "autoroute" or "FA" in the **Metric type** column of the LSPs table.

Identify SRLGs for protection

To simulate a router's "exclude protect" configuration, associate the source node of the protected circuit with the SRLG. If that circuit fails, the FRR LSP reroutes around all circuits within that SRLG.

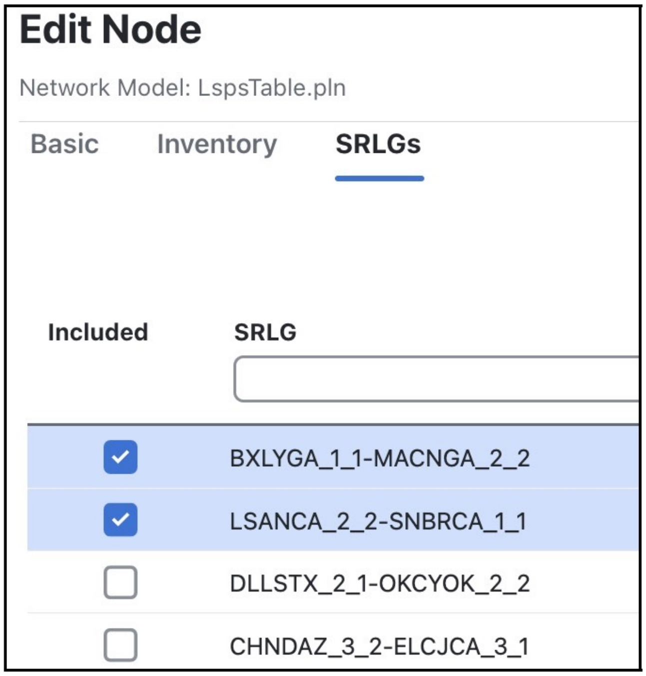

Procedure

Step 1 Select the source node of a circuit within an SRLG that you want to protect. Click or choose > **Edit**.

Step 2 Click the **SRLGs** tab.

Step 3 Choose one or more SRLGs with which you want to associate this node.

Step 4 Save your changes.

FRR LSPs initializer

The **FRR LSPs** initializer automatically creates link-protection or node-protection FRR LSPs when these conditions are met:

- One or more LSPs have their **FRR enabled** property set (refer to Mark LSPs for protection, on page 22).

- One or more interfaces on the path of these LSPs have their **FRR enabled** property set. This tells the FRR LSPs initializer how to create the FRR LSPs so that in the event of a failure, the interface would be avoided when the LSP is being rerouted.
  - This property is set in the **MPLS** tab of the Edit Interface window. The **FRR enabled** column in the Interfaces table shows "true" if the interface is included or "false" if it is not.
  - After FRR LSPs are created, the interface is listed in the **FRR interface** column in the LSPs table.

The initializer creates the FRR LSPs based on these FRR enabled properties and based on whether you choose to create link or node protection.

- Link protection: Cisco Crosswork Planning creates FRR LSPs for each pair of next-hop nodes (two connected nodes) in the protected LSP path where the egress interface of the first hop has its FRR enabled property set.
- Node protection: Cisco Crosswork Planning creates FRR LSPs between each set of next next-hop nodes (three connected nodes) in the protected LSP path where the egress interface of the first hop has its FRR enabled property set.

These new FRR LSPs are named FRR_*<source>* _*<destination>* _*<postfix>*, where *postfix* is optional and set in the initializer's page.

Run the FRR LSPs initializer

Use this task to run the FRR LSPs initializer and create protection LSPs for eligible protected LSPs.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Initializers** > **FRR LSPs**.

Step 3 Select the nodes you want to optimize. If you do not choose any nodes, the initializer uses all nodes in the plan file. Click **Next**.

Step 4 Choose whether to create FRR LSPs that protect links (circuits), nodes, or both.

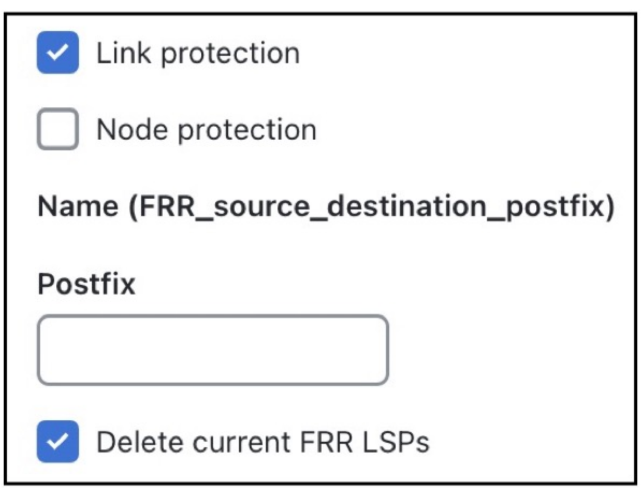

Step 5 (Optional) Enter a postfix to add to the end of the names of the newly created FRR LSPs.

Step 6 (Optional) Uncheck the **Delete current FRR LSPs** check box to retain the existing FRR LSPs.

Step 7 Submit your changes.

Create FRR LSPs manually

Use this task to create FRR LSPs manually and set the properties required for link or node protection.

This section describes the required properties for an FRR LSP. It does not describe all the options available when creating an LSP.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 Create an LSP using one of these options:

- In the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP**.
- In the Network Summary panel, click + > **LSPs** in the LSPs page.
- You can also choose an existing LSP to modify to be an FRR LSP. If doing so, click or choose ··· > **Edit**.

Step 3 Choose the source site and node of the FRR LSP.

Step 4 Choose the destination site and node, or enter a NetInt destination IP address. For information about NetInt destination IP addresses, refer to Unresolved LSP destinations and hops, on page 55.

Step 5 In the **Routing** area, choose **FRR** as the Routing type.

Step 6 Specify the interface that is to be avoided during routing by choosing it from the **FRR interface** drop-down list or by entering its IP address in the **NetInt FRRInterface** field.

Step 7 Save your changes.

Run fast reroute simulation

FRR LSPs are routed only when simulating failures in the Fast Reroute simulation mode. Complete these steps to enable the Fast Reroute simulation mode.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the Network Design page.

Step 2 In the toolbar, click or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 3 Click the **Simulation** tab.

Step 4 In the **Simulation convergence mode** section, choose **Fast reroute** from the Layer 3 drop-down list.

Step 5 Save your changes.

Autobandwidth-enabled LSPs

In a network, autobandwidth-enabled LSPs reset their setup bandwidth periodically based on the configured autobandwidth timers, and establish new routes if necessary. In networks with constrained bandwidth, changes in LSP routes can make other LSPs unroutable, which can then affect the amount of traffic in each LSP, the setup bandwidth settings, and the routes taken, without ever converging.

While Cisco Crosswork Planning does not simulate this instability, it does simulate autobandwidth-enabled LSPs. In this simulation mode, internally Cisco Crosswork Planning first routes the autobandwidth-enabled LSPs with a setup bandwidth of zero. Then, demands are routed to determine the simulated traffic (Traff sim) through each LSP. The result in the plan file and UI is that the Traff sim value is copied to the Setup BW sim for each autobandwidth-enabled LSP, and these LSPs are then routed using their new Setup BW sim value. Note that if some LSPs cannot be routed, the Traff sim and Setup BW sim values resulting from this process might not match one another for all autobandwidth-enabled LSPs.

Autobandwidth simulation modes

There are two autobandwidth simulation modes available. Before simulating failures or running Simulation analysis on autobandwidth-enabled LSPs, choose the appropriate mode depending on which state you need to simulate.

Note These simulation modes are only available in plan files with a single traffic level.

- **Autobandwidth convergence**: In this mode, the network is simulated after the traffic (Traff sim) has been rerouted to other LSPs, but before the Setup BW sim values have been reset. This simulates the immediate, less optimal response to failures.
- **Autobandwidth convergence (including failures)**: In this mode, the network is simulated after the Setup BW sim values have been reset. This simulates the longer, more optimal response to failures.

These autobandwidth simulations require that the LSPs have their Autobandwidth property set to true, and that one of the two autobandwidth modes be selected from the Network options page (in the toolbar, click or choose **Actions** > **Edit** > **Network options**).

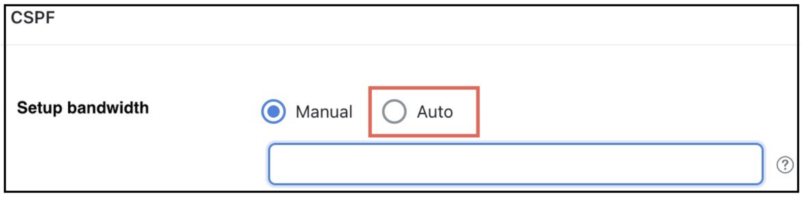

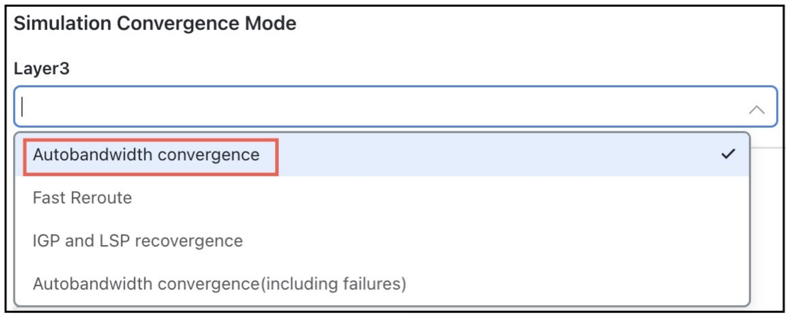

Autobandwidth convergence example with and without failures

This example describes how autobandwidth convergence behaves in a sample MPLS network under normal operation and when failures occur. The scenario highlights key network parameters, routing decisions, and bandwidth allocation across two LSPs.

Network parameters

The network in this example has these parameters:

- Demand: One demand originates from er1.chi and terminates at cr1.wdc.
- Two LSPs: LSP-A from cr1.chi to cr1.wdc and LSP-B from cr2.chi to cr1.wdc.
- Setup bandwidth: Neither LSPs has a Setup BW.
- Circuit capacities: All circuits have a 10,000 Mbps capacity except for one (between cr2.chi and cr1.wdc), which is only 1000 Mbps.

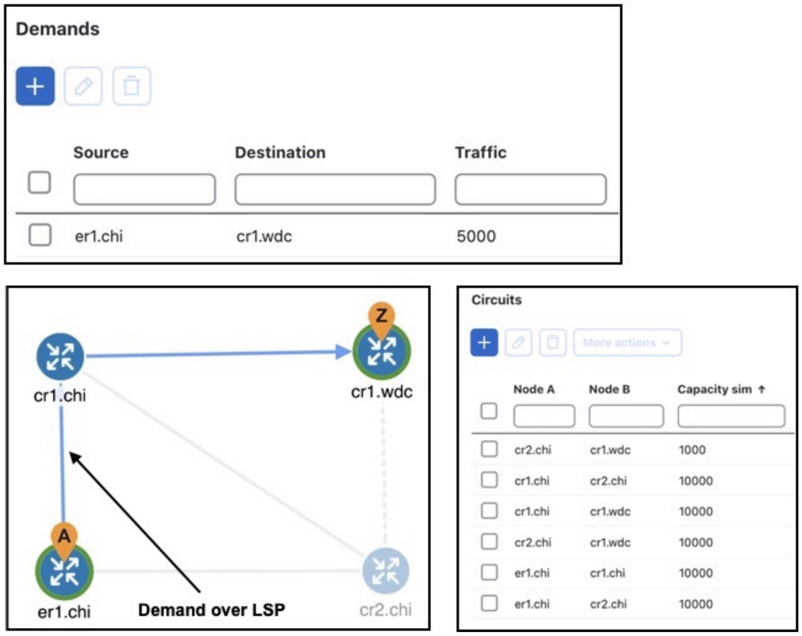

Scenario comparisons

- **Normal operation (no failures)**: Figure 10: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence, on page 29 shows that in the IGP and LSP Reconvergence mode, the Setup BW sim value for both LSPs is 0.
  - In the Autobandwidth convergence mode, the Setup BW sim value for LSP-A is copied from its Traff sim value, which is 5000.
  - There are no failures, and each LSP follows the same route in both convergence modes.

Figure 10: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence

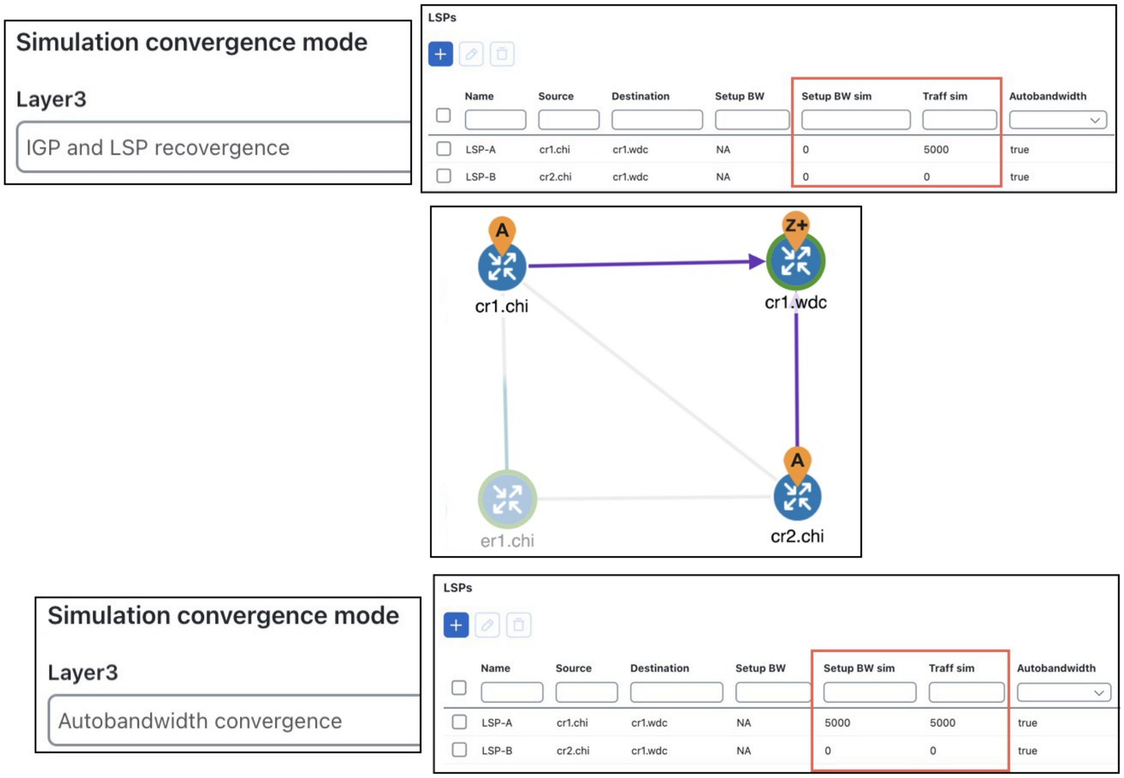

- **Failure event (autobandwidth convergence mode)**: Figure 11: Example autobandwidth convergence with failures, on page 30 shows the Autobandwidth convergence mode when there is a failure.
  - Failing the circuit between cr1.chi and cr1.wdc causes LSP-A to reroute through cr2.chi.
  - The shortest path from er1.chi through cr1.chi is now longer than the shortest path from er1.chi through cr2.chi, and so the demand moves over to LSP-B. This is shown in the LSPs table where the traffic (Traff Sim) of LSP-A is moved to LSP-B.
  - Because the smaller of the two circuits is congested, LSP-A uses the larger circuit, which has enough capacity to carry the traffic.
  - LSP-B continues to route on the lower capacity circuit, thus causing congestion.
  - The Setup BW sim value is not updated for either LSP because the Autobandwidth convergence mode does not update setup bandwidth as a result of traffic shifts due to failure.

Figure 11: Example autobandwidth convergence with failures

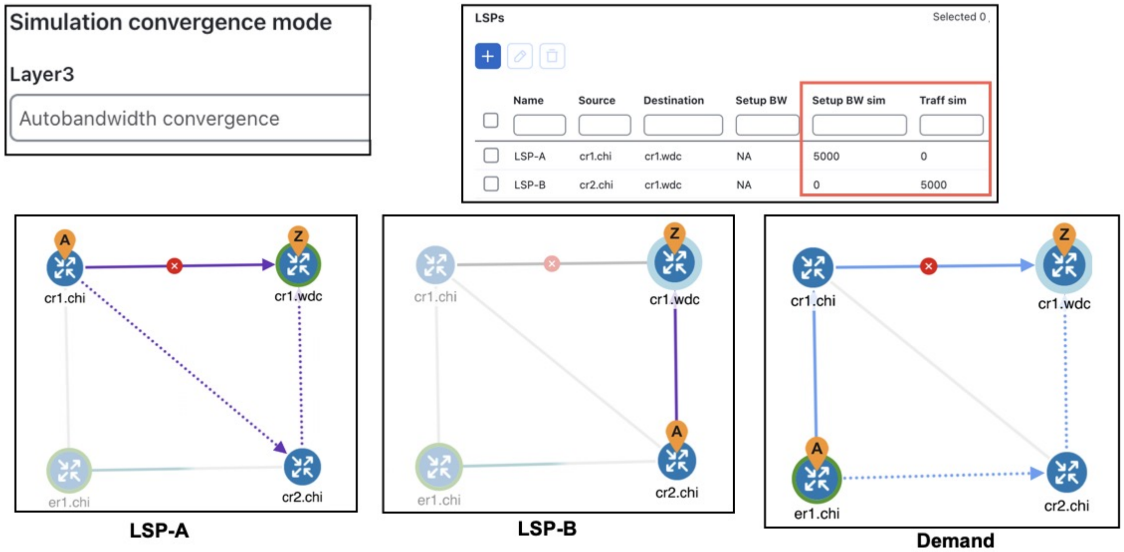

- **Failure event (autobandwidth convergence including failures mode)**: Figure 12: Example autobandwidth convergence including failures, on page 30 shows the same failure in Autobandwidth convergence (including failures) mode.
  - The LSP-B Setup BW sim is updated to the Traff sim value, forcing it to take the larger circuit to reach its destination. Note that the plot shows both the normal (solid) and failure (dotted) route for LSP-B, indicating that it has rerouted under the failure, even though the failure did not interrupt its path.
  - The demand continues to route on LSP-B, but because the LSP has found a path with sufficient capacity, there is no more congestion.

Figure 12: Example autobandwidth convergence including failures

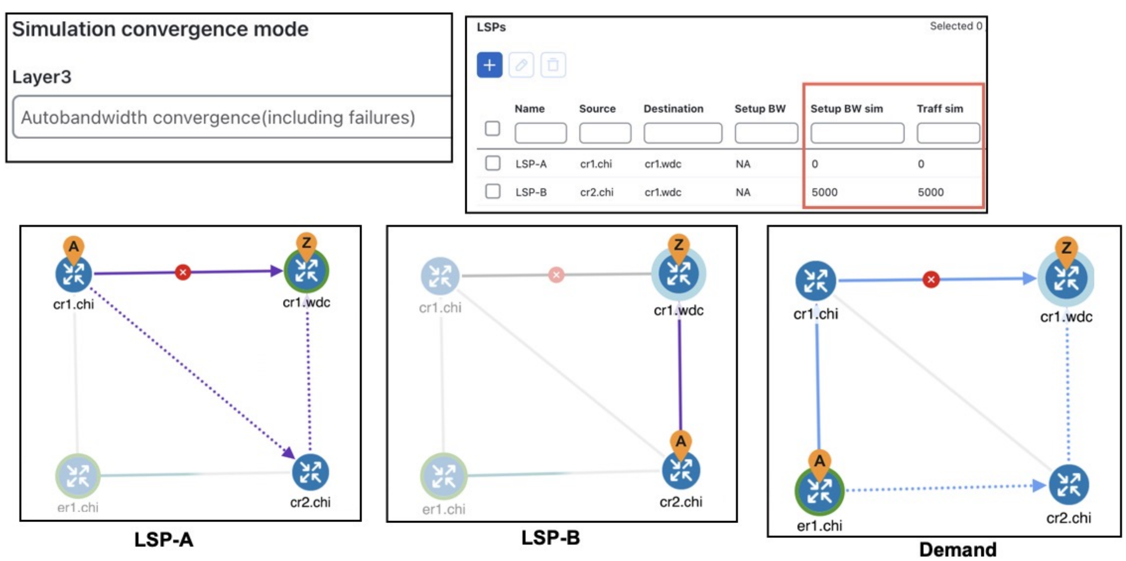

Circuit Style RSVP LSPs

Circuit Style RSVP (CS-RSVP) LSPs are logical entities that bundle two uni-directional RSVP LSPs with the same endpoints to form bi-directional RSVP LSPs. This allows traffic to consistently travel in both directions between the endpoints. Cisco Crosswork Planning supports creating and visualizing CS-RSVP LSPs.

CS-RSVP LSP key features

These are the key features of CS-RSVP LSPs:

- **Persistent end-to-end traffic engineered paths**: provides predictable and identical latency in both directions.
- **Strict bandwidth commitment**: ensures no impact on the SLA due to changing network load from other services.
- **End-to-end protection and restoration mechanism**: implements robust protection and restoration mechanisms.
- **Monitoring and maintenance of path integrity**: continuously monitors and maintains the integrity of the paths.
- **Data plane resilience**: ensures the data plane remains operational even if the control plane is down.

CS-RSVP LSP paths

To improve redundancy, a CS-RSVP LSP has four paths: Working (W), Protect (P), and the corresponding restore paths, named WR and PR respectively.

The paths are categorized as W, P, and restore paths according to the "Path option" field in the LSP paths table. From the lowest to the highest path option, they are classified as W, P, WR, and PR, respectively.

- Working path (W): This is the path with the lowest path option value.
- Protect path (P): This is the path with the second lowest path option value.
- Restore paths (WR and PR): These paths are calculated dynamically when the corresponding W or P is down, considering disjointness constraints. By default, node-disjoint is considered. If node-disjointness is not possible, it automatically falls back to link-disjointness. If an SRLG-disjoint is specified, it takes precedence over node and link disjointness. If the system cannot find a disjointed restore path, it keeps the corresponding CS-RSVP LSP operational without backup path protection.

  Restore paths are calculated based on the "Protected by" field of the LSP Path. For more information, refer to Restore path calculation, on page 31.

Cisco Crosswork Planning ensures that both W and P paths have a fully explicit path defined in the named path.

Restore path calculation

A restore path calculation is a method in CS-RSVP LSP configuration that

- determines how traffic is rerouted when a primary LSP path fails
- identifies which LSPs serve as working, protect, and restore paths under different failure scenarios, and
- uses attributes like **Protected by** and **Routed** to control path behavior.

The idea of a restore path is to protect the path that carries traffic. This is defined in the **Protected by** field of the LSP path. This field specifies which alternative LSPs can take over if the primary (working) path fails. The "Routed" field shows whether a specific LSP is actively routing traffic.

Example

Figure 13: Restore path calculation example, on page 33 shows the path configuration of the four LSP paths associated with a CS-RSVP LSP. Consider LSP path with path option 1 as LSP-1, Path option 2 as LSP-2, Path option 3 as LSP-3, and Path option 4 as LSP-4. Observe that LSP-1 and LSP-2 have the **Routed** option as "true", and LSP-3 and LSP-4 have this option as "false".

These steps describe how the restore path is calculated:

1. If LSP-1 fails, the traffic shifts to LSP-2 and LSP-3 becomes the Protect path. This happens because the **Protected by** field for LSP-1 is set to **2, 4**, indicating that when LSP-1 goes down, LSP-2 becomes the Working path and LSP-4 becomes the restore path. Since LSP-2 has the **Protected by** field as **1, 3** and since LSP-1 is already down, the LSP-3 becomes the Protect path.

   Notice that the **Routed** column value for LSP-1 changes to "false", and LSP-2 and LSP-3 have this value as "true".

2. Now, if LSP-2 goes down, traffic shifts to LSP-3 and LSP-4 becomes the Protect path. This happens because the LSP-2 has **1, 3** as the **Protected by** option. Since LSP-1 is down as indicated in Step 1, LSP-3 becomes the Working path. LSP-4 becomes the Protect path because LSP-3 has the **Protected by** field as 4.

   Observe that the **Routed** column value for LSP-2 changes to "false", and LSP-3 and LSP-4 have this value as "true".

Figure 13: Restore path calculation example

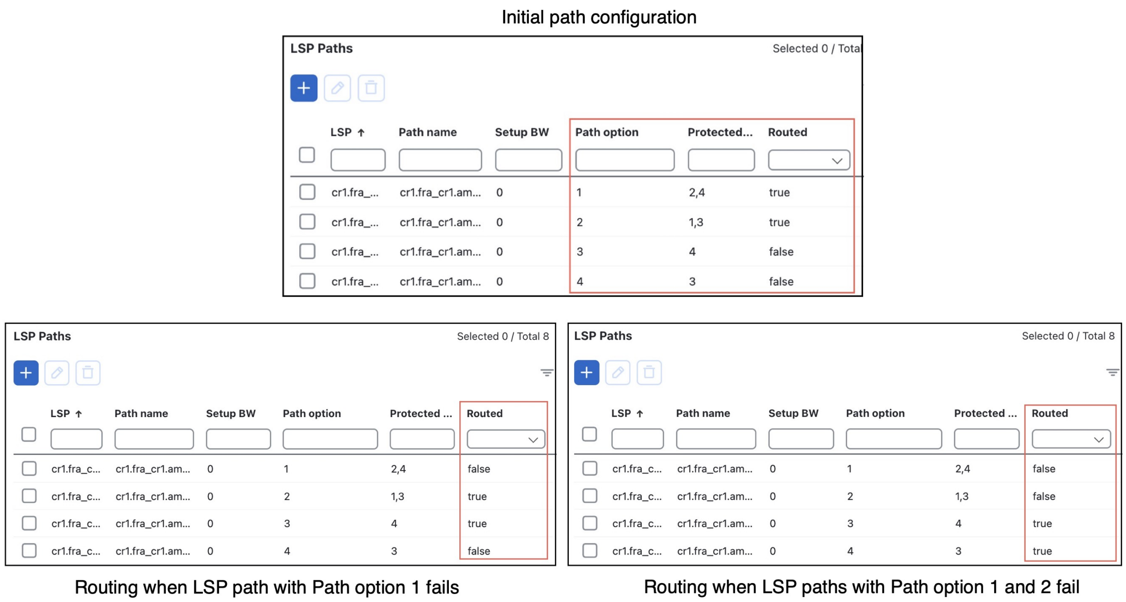

How CS-RSVP LSP path simulation works

Summary

Cisco Crosswork Planning performs these checks to determine if the path is routable:

Workflow

1. The explicit paths go through the same hops in both directions. This check applies only to restore paths routed dynamically.

2. The explicit path can reach the endpoints in both directions. Note that the configuration check is not performed.

3. The path has enough bandwidth to satisfy the setup bandwidth requirement. Note that asymmetric configuration is not checked.

4. If W or P is down, Cisco Crosswork Planning calculates the corresponding restore path. If the **Biased routing** option in the Network Options page is enabled, the restore path is calculated by reusing the existing reserved resources. If **Biased routing** is disabled, the restore paths are calculated regularly. For more information, refer to Configure biased routing, on page 34.

5. When routing an LSP path, the route follows the hops defined in the named path. W and P have fully explicit path listed in the named path. If the named path is not explicit or unroutable, the LSP path becomes unrouted. Paths are found dynamically only for restore paths.

6. When routing the restore path, node disjointness and link disjointness from the existing path carrying traffic are considered by default. If node disjointness is not possible, then link disjointness is considered. SRLG disjointness is considered if SRLG disjointness is configured at the LSP level.

Configure biased routing

Use this task to set biased routing options that influence CS-RSVP LSP restore path calculation during simulations.

In normal routing, restore path is not required to follow the corresponding W or P path as much as possible. The restore path is routed regularly with the constraint of disjointness.

In Biased Routing, the restore paths (WR and PR) reuse the corresponding W and P paths as much as possible, while ensuring that the total metric of the restore path does not exceed the maximum metric constraint. If biased routing is disabled, the restore paths are calculated regularly.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the **Network Design** page.

Step 2 In the toolbar, click or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 3 Under the **Simulation** tab, in the **Metric constraint** section, enable biased routing.

a) Check the **Biased routing** check box.

b) In the **Maximum metric constraint** field, enter the maximum allowable metric for the restore path. It can be set to an extremely large value if this constraint is not required.

Step 4 Save your changes.

Create CS-RSVP LSPs

Use this task to create CS-RSVP LSPs by configuring participating RSVP LSPs, co-routing settings, association IDs, and path roles.

You can create CS-RSVP LSPs in the Cisco Crosswork Planning Design application either manually or with the **CS-RSVP LSP optimizer** tool.

This task describes the manual method. For the optimizer method, refer to Create CS-RSVP LSPs using the optimizer tool, on page 35.

Before you begin

In the Add/Edit LSP page of the two RSVP LSPs participating in the CS-RSVP LSP, ensure that:

- The **Co-routed** option is enabled. This option is available in the **CSPF** section of the Add/Edit LSP page.
- They have the same Association ID.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **CS-RSVP LSP**.

Alternatively, in the Network Summary panel on the right side, click + in the **CS-RSVP LSPs** tab.

The CS-RSVP LSPs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **CS-RSVP LSPs** check box.

Step 3 In the **Node A** and **Node B** sections, choose the endpoints of the CS-RSVP LSP.

Step 4 Enter the details in these fields.

- **Association ID**: This field is used to associate two RSVP LSPs. This is a required field.
- **Source address**: This field represents the source address for the bidirectional association. This is an optional field. If you do not enter a value, it defaults to an empty value.
- **Global ID**: This field specifies the global ID for the association's global source. This is an optional field. If you do not enter a value, it defaults to an empty value.
- **Routed**: This field indicates whether the CS-RSVP LSP is routable. It derives its value based on specific criteria and cannot be manually set.

Step 5 Click **Add**.

CS-RSVP LSP optimizer

CS-RSVP LSP optimizer is a Cisco Crosswork Planning tool that

- helps to create, find, and optimize CS-RSVP LSPs based on specified endpoint pairs, association ID, source address, and global ID
- optimizes both existing and newly created CS-RSVP LSPs together, and
- allows you to choose to optimize only the existing LSPs or only the newly created ones

Create CS-RSVP LSPs using the optimizer tool

If CS-RSVP LSPs do not exist, the optimizer creates the corresponding LSPs based on the given configurations. The configuration process involves specifying several key parameters as described in this procedure.

The tool creates two LSP paths as W and P for each LSP. It creates the hops of each path as the named path of the corresponding LSP path. In addition to W and P paths, it also creates two additional paths, WR and PR. These additional paths do not have named paths and are associated with their respective W and P paths.

Procedure

Step 1 Open the plan file (refer to Open plan files). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.

Step 3 (Optional) If you want to optimize the existing CS-RSVP LSPs, select them and click **Next**.

Step 4 Click **Create** in the Create new CS-RSVP LSPs page.

Step 5 Enter the details in the required fields. In the **LSP path configurations** section, define the configurations for the four LSP paths: W, P, WR, and PR. For field descriptions, refer to Table 2: CS-RSVP LSP configuration parameters, on page 38.

Figure 14: Create new CS-RSVP LSPs page

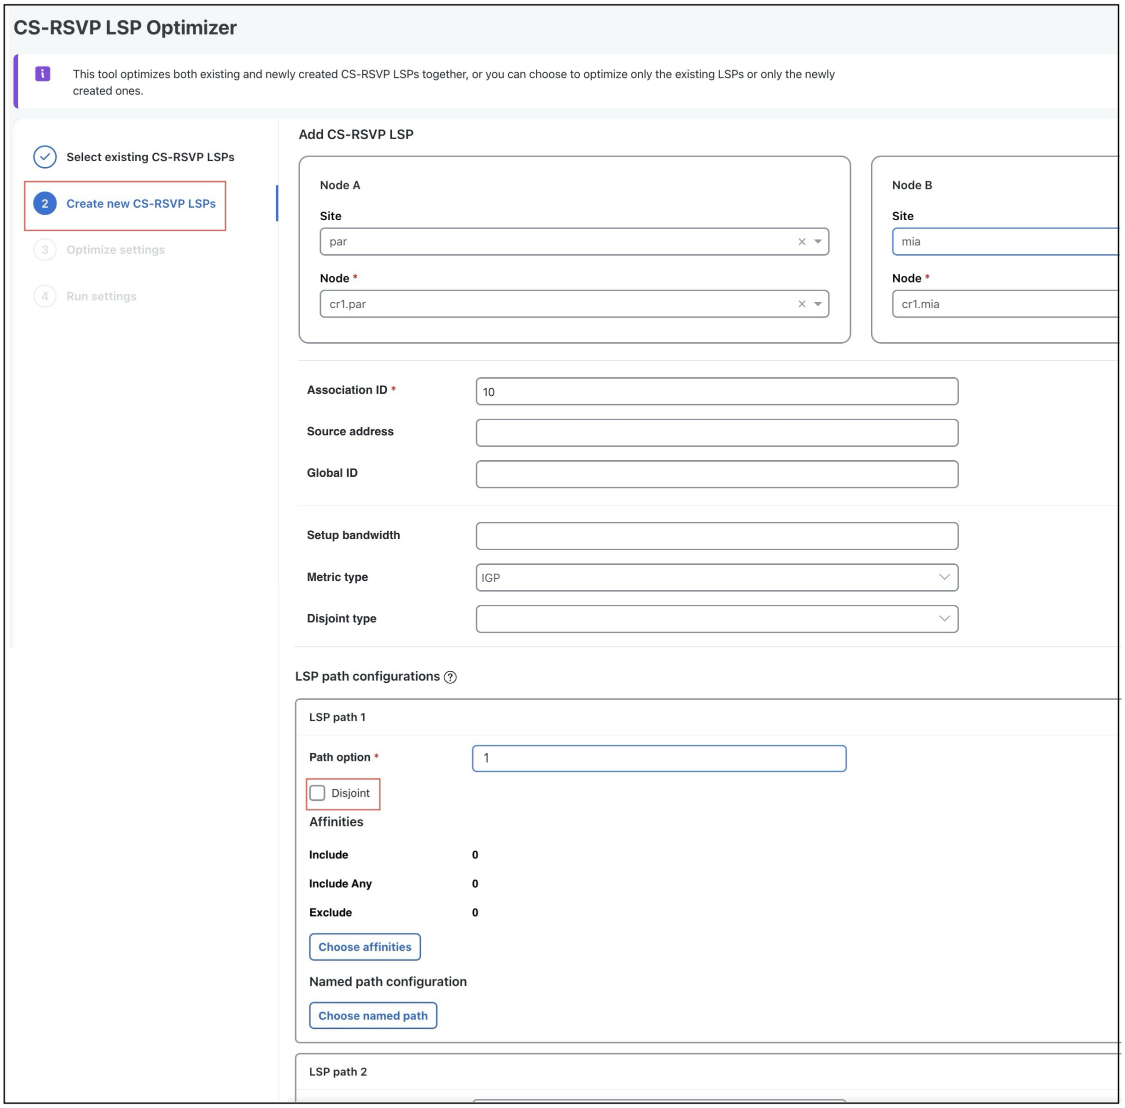

Step 6 Click **Save**. Then, select the newly created CS-RSVP LSP and click **Next**.

To create additional CS-RSVP LSPs, click [+].

Step 7 Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

- Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.

- As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
  - Priority: Select the priority of the task.
  - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
  - Schedule: Set the time at which you want to run the tool.

  The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager.

  Note
  Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

Step 9 (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

In the previous step:

- If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
- If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

Observe these changes:

- The newly created CS-RSVP LSPs are available on the CS-RSVP LSPs table with *NewCSRSVPLSPs* as tags.
- The corresponding LSPs are available on the LSPs table.
- The corresponding LSP paths are available on the LSP paths table. Note that the LSP paths are created based on the Path option values you entered and the **Protected by** field is automatically calculated.

Optimize existing CS-RSVP LSPs with the optimizer tool

If CS-RSVP LSPs already exist, the optimizer can optimize them based on the current configurations and requirements. Use this task to optimize the existing CS-RSVP LSPs using the CS-RSVP LSP optimizer tool.

Procedure

Step 1 Open the plan file (refer to Open plan files) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.

Step 3 Select the existing CS-RSVP LSPs you want to optimize and click **Next**.

Step 4 (Optional) If you want to create new CS-RSVP LSPs, click **Create** and configure the required parameters. For more information, refer to Create CS-RSVP LSPs using the optimizer tool, on page 35.

Step 5 Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.

Step 6 Click **Next**.

Step 7 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

- Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
- As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
  - Priority: Select the priority of the task.
  - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
  - Schedule: Set the time at which you want to run the tool.

The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager.

Note
Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

Step 8 (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

In the previous step:

- If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
- If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

Observe that the selected CS-RSVP LSPs are updated and have the tags*UpdatedCSRSVPLSPs*.

CS-RSVP LSP configuration parameters

This table lists the parameters to configure when creating CS-RSVP LSPs using the optimizer tool.

Table 2: CS-RSVP LSP configuration parameters

Field	Description
Node A and Node B	These fields indicate the endpoints of the LSPs corresponding to the CS-RSVP LSP.
Association ID	This field is used to associate two RSVP LSPs. This is a required field.

Field	Description
Source address	This field represents the source address for the bidirectional association. This is an optional field. If you do not enter a value, it defaults to an empty value.
Global ID	This field specifies the global ID for the association's global source. This is an optional field. If you do not enter a value, it defaults to an empty value.
Setup bandwidth	This field indicates the amount of traffic the source nodes requests for this CS-RSVP LSP in Mbps. This is an optional field.
Metric type	This field indicates the metric type for route calculation. Choose either IGP or TE as per your requirement.
Disjoint type	This field indicates the disjoint type to consider. Choose **SRLGs** to enable SRLG disjointness in addition to the default node and link disjointness.
LSP path configurations	
Path option	This field specifies the path option of the LSP paths corresponding to the CS-RSVP LSP. The path with the lowest Path option value is assigned as the W path, the next lowest as P, the third lowest as WR, and the fourth lowest as the PR. Ensure that each LSP path has a unique Path option value. For example, if the path options are 1, 2, 3, and 4, then the LSP path with Path option 1 is designated as W, path option 2 as P, path option 3 as WR, and path option 4 as PR.
Disjoint	This check box indicates whether to consider this LSP path for disjointness with any other LSP path that has this check box enabled. Ensure that exactly two LSP paths are configured as disjoint.
Affinities	This is an optional field. Click **Choose affinities** to assign affinities to the LSP paths.
Named path configuration	This is an optional field. Click **Choose named path** to assign named path configurations to the LSP paths.

CS-RSVP LSP simulation diagnostics

The **CS-RSVP LSP simulation diagnostics** tool in Cisco Crosswork Planning performs configuration and disjointness checks on the CS-RSVP LSPs.

Configuration check

The configuration check ensures that these properties of a given CS-RSVP LSP are identical on the associated RSVP LSPs:

- Association ID, Source address, and Global ID
- Setup bandwidth
- LSP type and LSP path count

- Protected by and Dynamic field configurations
- Hold priority and setup priority
- Named path, ensuring that Working and Protect paths have named paths
- Source and Destination
- Path options on LSP paths

Disjointness check

The disjointness check ensures that all CS-RSVP LSPs in the network model are disjoint, meaning they do not share certain resources. This helps to maintain redundancy and avoid single points of failure.

The tool summarizes these violations:

- SRLG disjointness violation (only for CS-RSVP LSP with SRLG disjointness configuration)
- Node disjointness violation
- Link disjointness violation
- Node to link disjointness fallback

Run the CS-RSVP LSP simulation diagnostics tool

Use this task to run the CS-RSVP LSP simulation diagnostics tool and identify configuration, association, and path issues.

Procedure

Step 1 Open the plan file (refer to Open plan files) with CS-RSVP LSPs. It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Tools** > **Diagnostics** > **CS-RSVP LSP simulation**.

Step 3 Click **Proceed** in the information dialog box.

A report is generated indicating the details of the failures.

CS-RSVP LSP simulation diagnostics report

Each time the CS-RSVP LSP simulation diagnostics tool is run, a report is automatically generated. You can access this information at any time by choosing **Actions** > **Reports** > **Generated reports** and then clicking the **CS-RSVP LSP Simulation** link in the right panel. Note that new reports replace the previous ones.

The **Summary** tab provides information on the number of failed CS-RSVP LSPs.

The **CS-RSVL LSP simulation** tab provides failure details. Expand the **FAILUREMESSAGE** column to view the failure messages.

CS-RSVP LSP simulation diagnostics example

The plan file has a CS-RSVP LSP with configuration issues, such as associated LSPs having different Setup priority, Hold priority, and Setup bandwidth. After running the CS-RSVP LSP simulation diagnostics tool, the generated report shows this information in the **FAILUREMESSAGE** column.

By filtering from the CS-RSVP LSPs table to the LSPs table, you can refer to that the values for these fields do not match.

Figure 15: CS-RSVP LSP simulation diagnostics report and mismatching LSPs

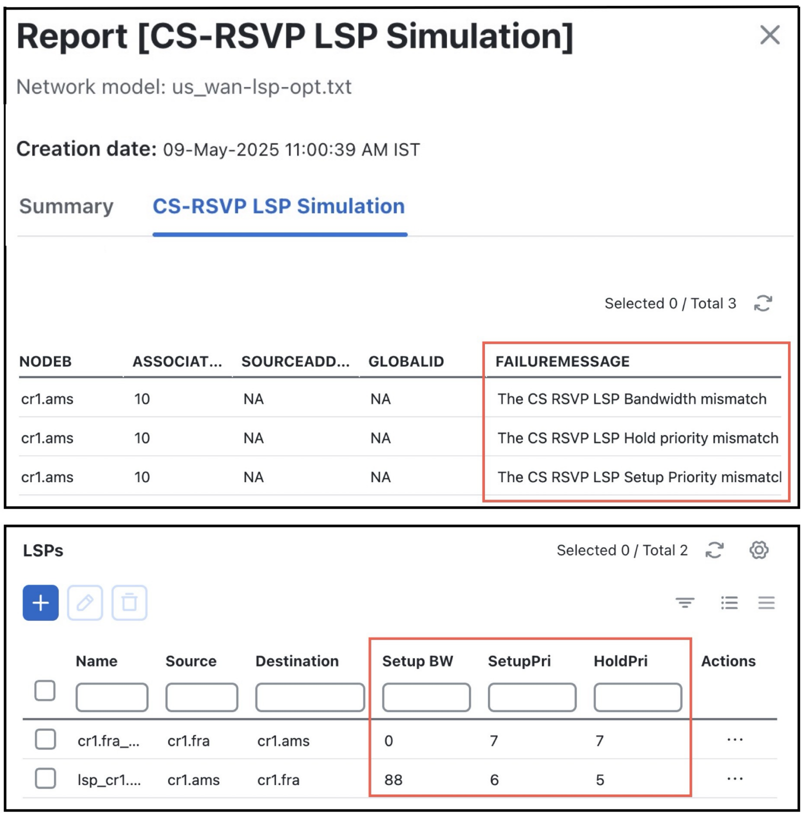

CS-RSVP LSP visualization

CS-RSVP LSPs are associated with LSPs and LSP paths. This section describes how to visualize:

- CS-RSVP LSPs
- LSPs associated with CS-RSVP LSPs
- LSP paths associated with CS-RSVP LSPs

Visualize CS-RSVP LSPs

Use this task to visualize selected CS-RSVP LSPs in the network plot.

Procedure

Step 1 Open the plan file (refer to Open plan files) with CS-RSVP LSPs. It opens in the Network Design page.

Step 2 In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab.

The CS-RSVP LSPs tab is available under the **More** tab. If it is not visible, then click the Show/hide tables icon ( ) and check the **CS-RSVP LSPs** check box.

Note
You can select only one CS-RSVP LSP at a time to view on the network plot.

Step 3 Select the required CS-RSVP LSP from the list.

Step 4 Observe these in the network plot on the left side:

- Both LSPs associated with the CS-RSVP LSP are plotted. The purple arrow indicates the LSP path is active (carrying traffic). The number on the path (here 1) is the path option. For example, refer to the first image in Figure 16: CS-RSVP LSP path visualization, on page 43.

  A dim purple arrow indicates the standby path, which is an LSP path that is not a primary path and is set as a Standby path. To view the standby paths, check the **All LSP paths** check box in the network plot. For example, refer to the second image in Figure 16: CS-RSVP LSP path visualization, on page 43.

- The bi-directional LSPs are plotted when you check the **Bi-dir** check box in the network plot. For example, refer to the third image in Figure 16: CS-RSVP LSP path visualization, on page 43.

Figure 16: CS-RSVP LSP path visualization

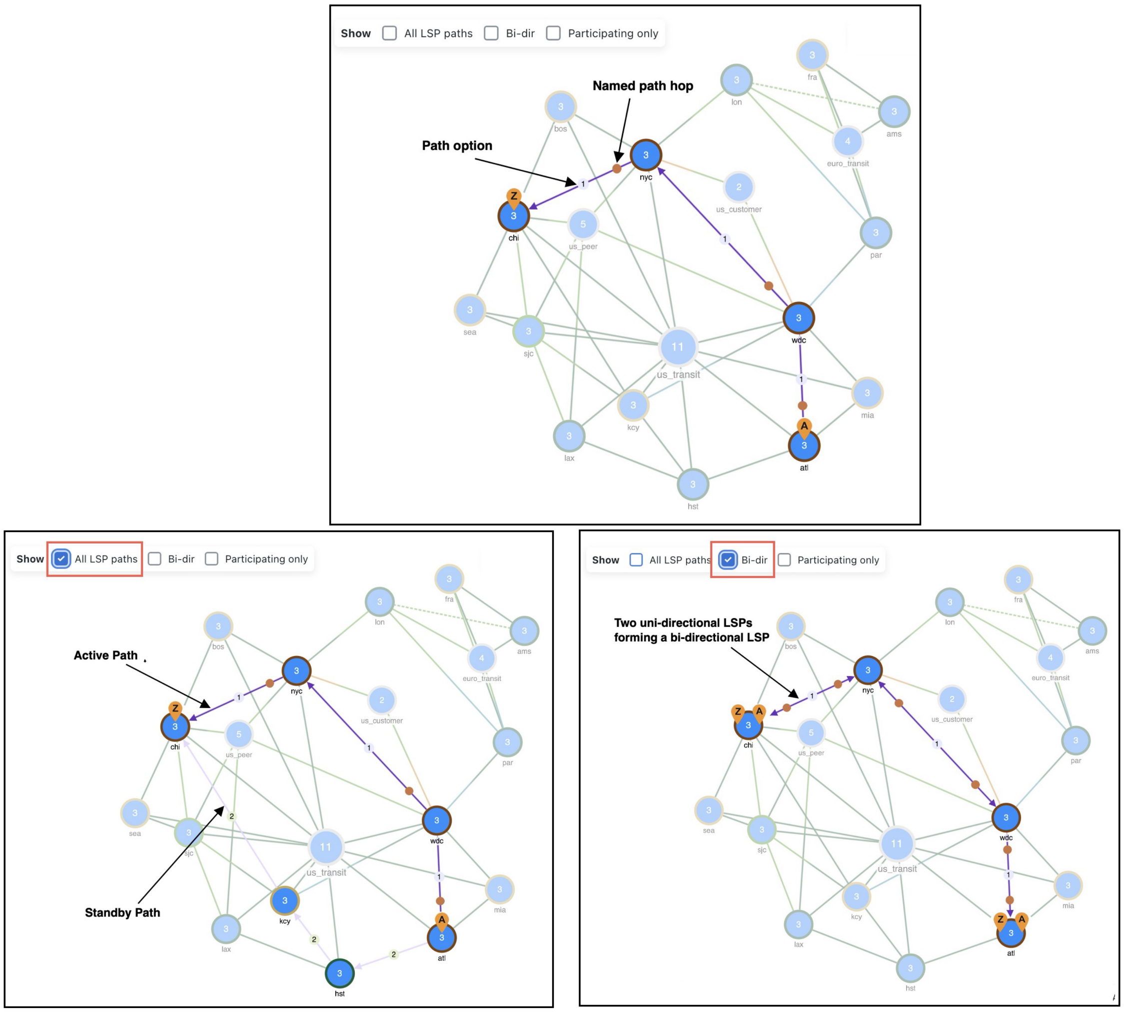

Visualize LSPs associated with CS-RSVP LSPs

Use this task to visualize RSVP LSPs associated with selected CS-RSVP LSPs.

Procedure

Step 1 Open the plan file (refer to Open plan files) with CS-RSVP LSPs. It opens in the Network Design page.

Step 2 In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab and select the required CS-RSVP LSP from the list.

Note
You can select only one row at a time.

Step 3 Choose > **Filter to LSPs**. The LSPs table opens, displaying the LSPs associated with the selected CS-RSVP LSP.

Step 4 Select the LSPs.

Note
You can select up to three LSPs at a time.

Step 5 Observe these in the network plot on the left side:

- All LSP paths which belong to the LSPs and are routed are plotted.
- The active path is distinguishable from the other alternate or standby paths.
- If the Working (W) path fails and traffic is rerouted, both the no-fail path and fail (rerouted) path are plotted.

Visualize LSP paths associated with CS-RSVP LSPs

Use this task to visualize LSP paths associated with selected CS-RSVP LSPs.

Procedure

Step 1 Open the plan file (refer to Open plan files) with CS-RSVP LSPs. It opens in the Network Design page.

Step 2 In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab and select the required CS-RSVP LSP from the list.

Note
You can select only one row at a time.

Step 3 Choose > **Filter to LSP paths**. The LSP Paths table opens, displaying the LSP paths associated with the selected CS-RSVP LSP.

Step 4 Select the LSP paths to visualize in the network plot.

Note
You can select up to three LSP paths at a time.

Advanced RSVP-TE LSP simulations

Ignore reservable bandwidths for capacity planning

When using Cisco Crosswork Planning for capacity planning, it is common to increase the demand traffic to a level expected at a future date. Various simulations are then performed, such as failures and worst-case analysis, to determine which circuits will likely experience over-utilization.

In MPLS networks using RSVP-TE LSPs, a setup bandwidth is configured for each LSP. If the demand traffic is increased, LSP setup bandwidths are usually increased correspondingly (for example, by using the LSP

Setup Bandwidth initializer). When these LSP setup bandwidths become too large, some LSPs might not be routable, thus rendering an improbable view of the future network. When this occurs, you have the option to temporarily remove the reservable bandwidth constraints for a more realistic view into future capacity requirements. While this is not a routing method that actually exists on routers, it is a convenient Cisco Crosswork Planning simulation that can assist you when planning networks.

When you enable this routing mode, Cisco Crosswork Planning routes LSPs in the order they appear in the plan file LSPs table. Note that static order of the LSPs table in the plan file likely differs from the order of LSPs in the GUI, which changes frequently due to user interaction.

This operates only on LSPs that do not have autobandwidth enabled.

Procedure

Step 1 Route LSPs as usual. For information, refer to Dynamic LSP routing and CSPF, on page 1.

Step 2 If an LSP is not routable (call it “LSP-A”), try to route it with zero setup bandwidth.

- If LSP-A cannot be routed, leave it unrouted and try to route the next LSP in the LSPs table. If this LSP cannot be routed, try to route the next LSP listed in the LSPs table in the plan file.
- If LSP-A can be routed, then there will be one or more interfaces on its route whose reservable bandwidth is exceeded. These interfaces are marked as having unlimited reservable bandwidth during the simulation of all subsequent LSPs, thus allowing for the routing of LSP-A and other LSPs that would otherwise exceed the reservable bandwidth on these interfaces. Note that actual Resv BW and Resv BW sim values in the Interfaces table do not change.

On completion of the simulation, a number of interfaces will have their reserved bandwidth greater than their configured reservable bandwidth. The circuits containing these interfaces are candidates for capacity expansion.

Ignore reservable bandwidth constraints

Use this task to configure routing simulations to ignore reservable bandwidth constraints during routing simulations.

Procedure

Step 1 From the toolbar, click or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 2 Click the **Simulation** tab.

Step 3 In the **Label switched paths** section, check the **LSP capacity planning mode** check box. This ignores the reservable bandwidth constraints.

Step 4 Save your changes.

Reservable bandwidth capacity planning example

This example demonstrates RSVP LSP routing in a simple network and examines how reservable bandwidth influences routing outcomes during simulation analysis.

Network configuration

In this example, there are four LSPs, each with a setup bandwidth (Setup BW) of 600 Mbps. All interfaces have a reservable bandwidth (Resv BW) of 1000 Mbps.

- LSP details
  - LSP-1, LSP-2, and LSP-3 go from LON to FRA.
  - LSP-4 goes from LON to ROM.
- IGP metrics
  - Interfaces between PAR and FRA, and interfaces between FRA and ROM have an IGP metric of 2.
  - All other interfaces have an IGP metric of 1.

Figure 17: Comparison of reservable bandwidth being used vs not being used, on page 46 shows the LSP reservations view of this simple network in both simulations modes: the usual where reservable bandwidths are observed and the capacity planning version wherein they are ignored. In this latter mode, the two red circuits indicate they are being over utilized and likely need their capacity increased. (For information about the LSP Reservations view, refer to View LSP reservations, on page 4.)

Figure 17: Comparison of reservable bandwidth being used vs not being used

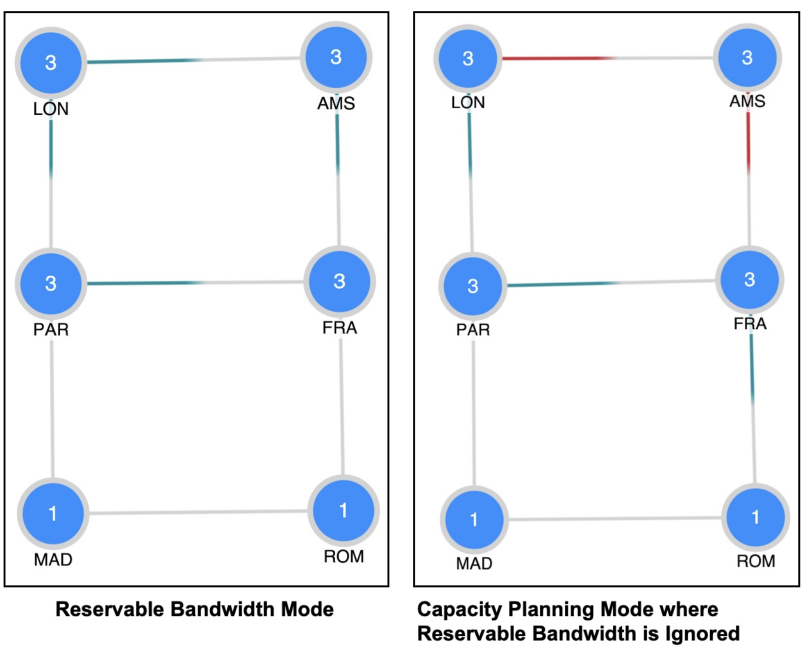

Figure 18: Simulated routing using reservable bandwidth constraints, on page 47 demonstrates routing simulations where reservable bandwidth constraints are applied (**LSP capacity planning mode** unchecked), which is the default behavior.

- LSP-1 routes across LON-AMS-FRA, leaving circuits LON-AMS and AMS-FRA each with only 400 Mbps of available bandwidth.
- LSP-2 routes across LON-PAR-FRA, leaving circuits LON-PAR and PAR-FRA each with only 400 Mbps of available bandwidth.
- LSP-3 cannot be routed because there is not enough available bandwidth on LON-AMS or LON-PAR. Its Setup BW sim value in the LSPs table is set to “NA” and the Routed column is set to “Unrouted.”
- LSP-4 cannot be routed, again because there is not enough available bandwidth. Its Setup BW Sim value in the LSPs table is set to “na” and the Routed column is set to “Unrouted.”

Figure 18: Simulated routing using reservable bandwidth constraints

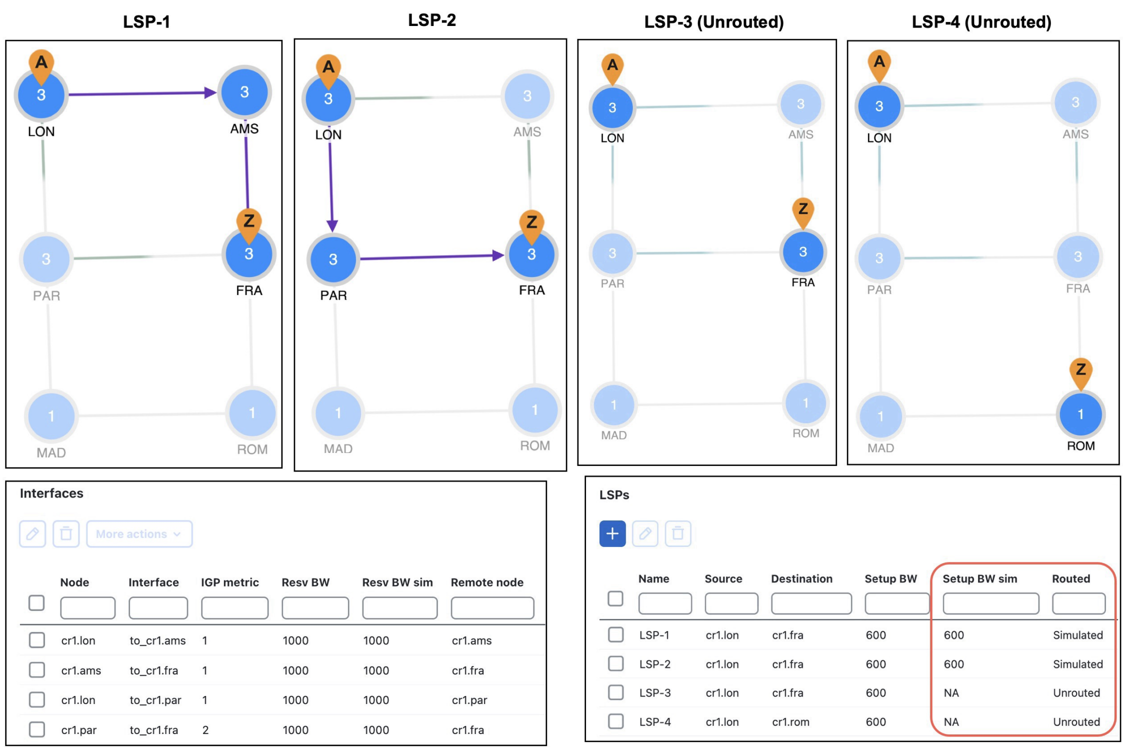

	Node	Interface	IGP metric	Resv BW	Resv BW sim	Remote node
	cr1.lon	to_cr1.ams	1	1000	1000	cr1.ams
	cr1.ams	to_cr1.fra	1	1000	1000	cr1.fra
	cr1.lon	to_cr1.par	1	1000	1000	cr1.par
	cr1.par	to_cr1.fra	2	1000	1000	cr1.fra

	Name	Source	Destination	Setup BW	Setup BW sim	Routed
	LSP-1	cr1.lon	cr1.fra	600	600	Simulated
	LSP-2	cr1.lon	cr1.fra	600	600	Simulated
	LSP-3	cr1.lon	cr1.fra	600	NA	Unrouted
	LSP-4	cr1.lon	cr1.rom	600	NA	Unrouted

In Figure 19: Simulated routing without using reservable bandwidth constraints, on page 48, the use of reservable bandwidth constraints has been disabled (**LSP capacity planning mode** checked). The LSP routes are:

- As with the default behavior, LSP-1 routes across LON-AMS-FRA and LSP-2 routes across LON-PAR-FRA.
- Once it is determined that there is not enough reservable bandwidth for LSP-3 to route, it takes the shortest path, LON-AMS-FRA. The reservable bandwidth on interfaces LON-to-AMS and AMS-to-FRA is ignored for the remainder of this simulation. LSP-3’s Setup BW sim value in the LSPs table is set to 600.

- LSP-4 routes across LON-AMS-FRA-ROM because these reservable bandwidth constraints have been removed and because it is the only route available. Even though LON-PAR-MAD-ROM is the shortest path, there is not enough reservable bandwidth on LON-to-PAR or PAR-to-FRA. LSP-4's Setup BW sim value in the LSPs table is set to 600.

Figure 19: Simulated routing without using reservable bandwidth constraints

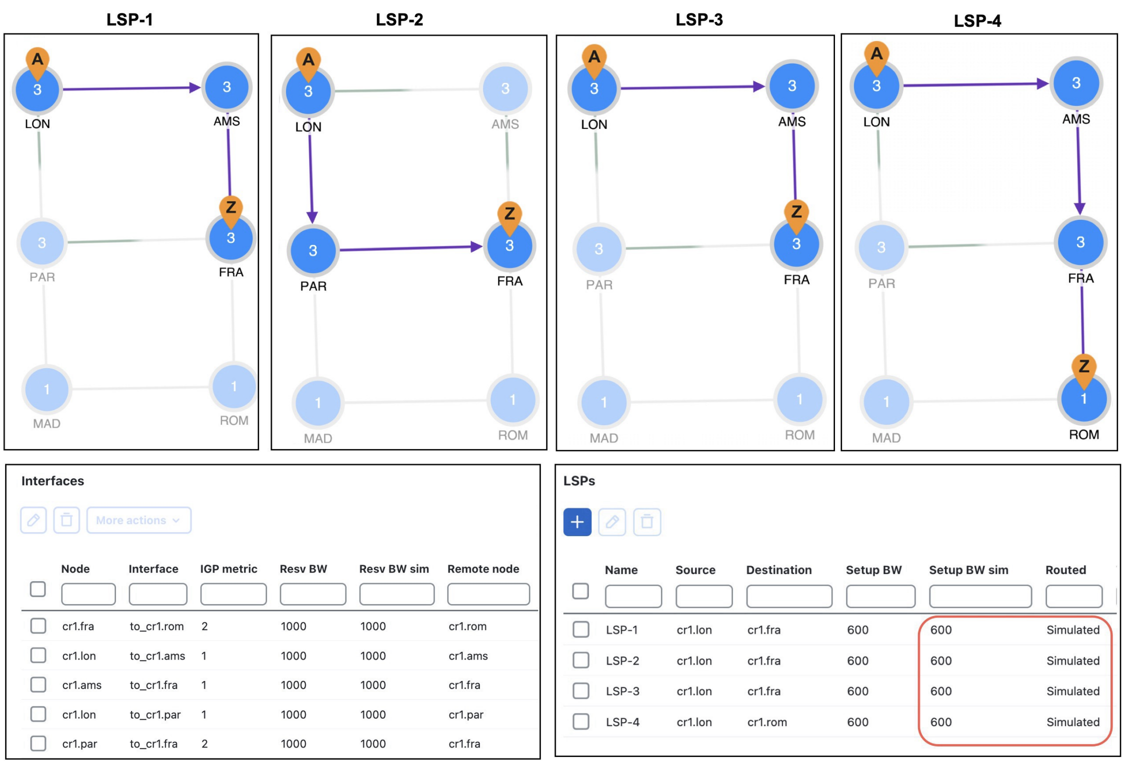

If, however, LSP-4 were routed before LSP-3, then LSP-4 would route across its shortest path of LON-PAR-MAD-ROM, and the LON-to-PAR interface would have its reservable bandwidth constraint ignored. With this routing, LSP-3 would still route LON-AMS-FRA, and the interfaces LON-to-AMS and AMS-to-FRA would still be set to remove their reservable bandwidth constraint.

P2MP LSPs

Point-to-multipoint (P2MP) LSPs offer an MPLS alternative to IP multicast for setting up multiple paths to multiple locations from a single source. In live networks, the signal is sent once across the network and packets are replicated only at the relevant or designated branching MPLS nodes.

A P2MP LSP consists of two or more sub-LSPs that have the same source node. Most of Cisco Crosswork Planning refers to sub-LSPs as LSPs, and there is no distinction between the two.

To view the P2MP LSPs or their associated sub-LSPs, choose them from the P2MP LSPs and LSPs tables. To determine if an LSP belongs to a P2MP LSP, you can show the **P2MP LSP** column in the LSPs table.

Example

In Figure 20: P2MP LSPs and associated sub-LSPs, on page 49, there are two P2MP LSPs, LSP_er12 and LSP_er13. LSP_er12 and LSP_er13 have "er12" and "er13" as their source nodes, respectively. In the LSPs table, the first three LSPs have er12 as the source node and are sub-LSPs of LSP_er12 P2MP LSP (as indicated by the **P2MP LSP** column). Similarly, the LSPs with er13 as the source node are the sub-LSPs of LSP_er13 P2MP LSP. The **Sub LSP count** in the P2MP LSPs table indicates the number of sub-LSPs of each P2MP LSP.

Figure 20: P2MP LSPs and associated sub-LSPs

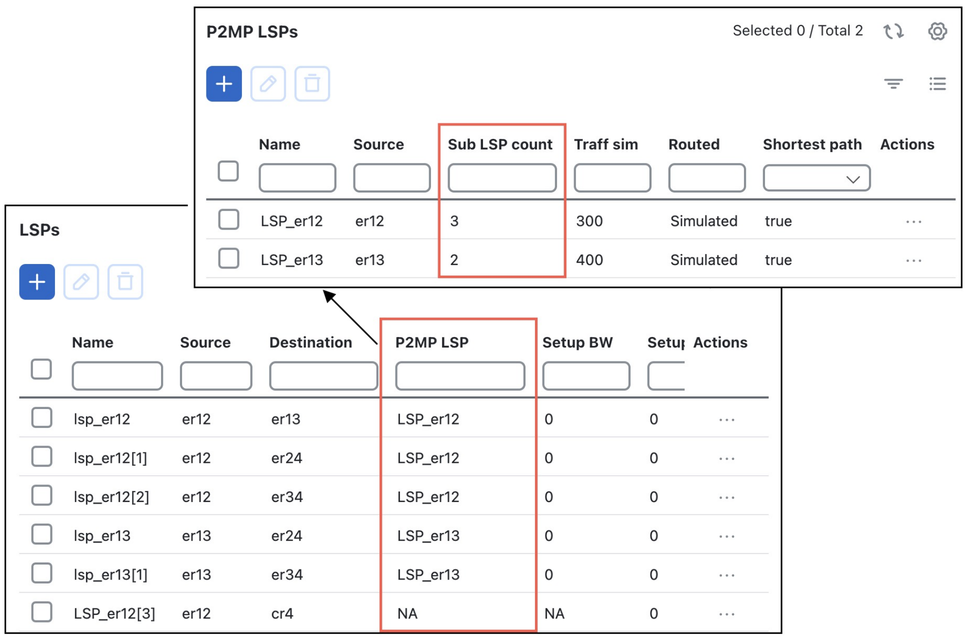

P2MP LSP bandwidth

All sub-LSPs within a given P2MP LSP share bandwidth and therefore, common circuits between these sub-LSPs have the same bandwidth, rather than an aggregate. Common practice is to set all sub-LSPs within a P2MP LSP to the same bandwidth. Otherwise, the bandwidth for the shared circuit is the highest bandwidth set.

To refer to the reserved LSP bandwidth, show the LSP Resv column in the Interfaces table. You can also view the associated traffic by choosing **LSP Reservations** in the Network Plot menu in the Visualization toolbar. For more information, refer to View LSP reservations, on page 4.

Example

Example: If A-C LSP and A-D LSP each have a bandwidth 400 Mbps, the A-B circuit has an aggregate bandwidth of 800 Mbps (Figure 21: Bandwidth behavior without a P2MP LSP, on page 50). However, if A-C and A-D are sub-LSPs within a P2MP LSP, the A-B circuit has a bandwidth of 400 Mbps (Figure 22: Bandwidth behavior with a P2MP LSP, on page 51).

Figure 21: Bandwidth behavior without a P2MP LSP

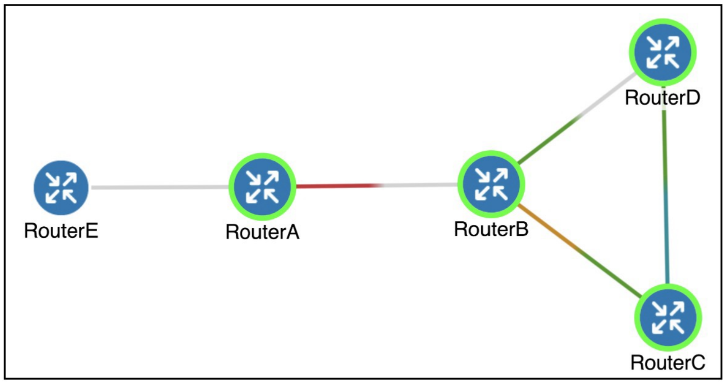

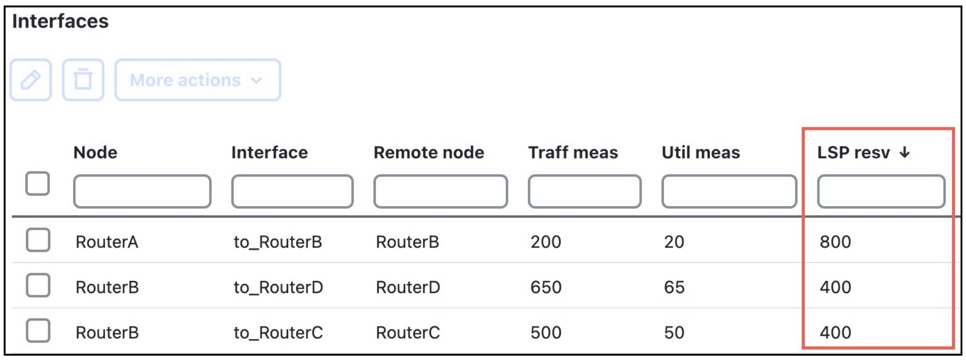

Interfaces

	Node	Interface	Remote node	Traff meas	Util meas	LSP resv ↓
	RouterA	to_RouterB	RouterB	200	20	800
	RouterB	to_RouterD	RouterD	650	65	400
	RouterB	to_RouterC	RouterC	500	50	400

Figure 22: Bandwidth behavior with a P2MP LSP

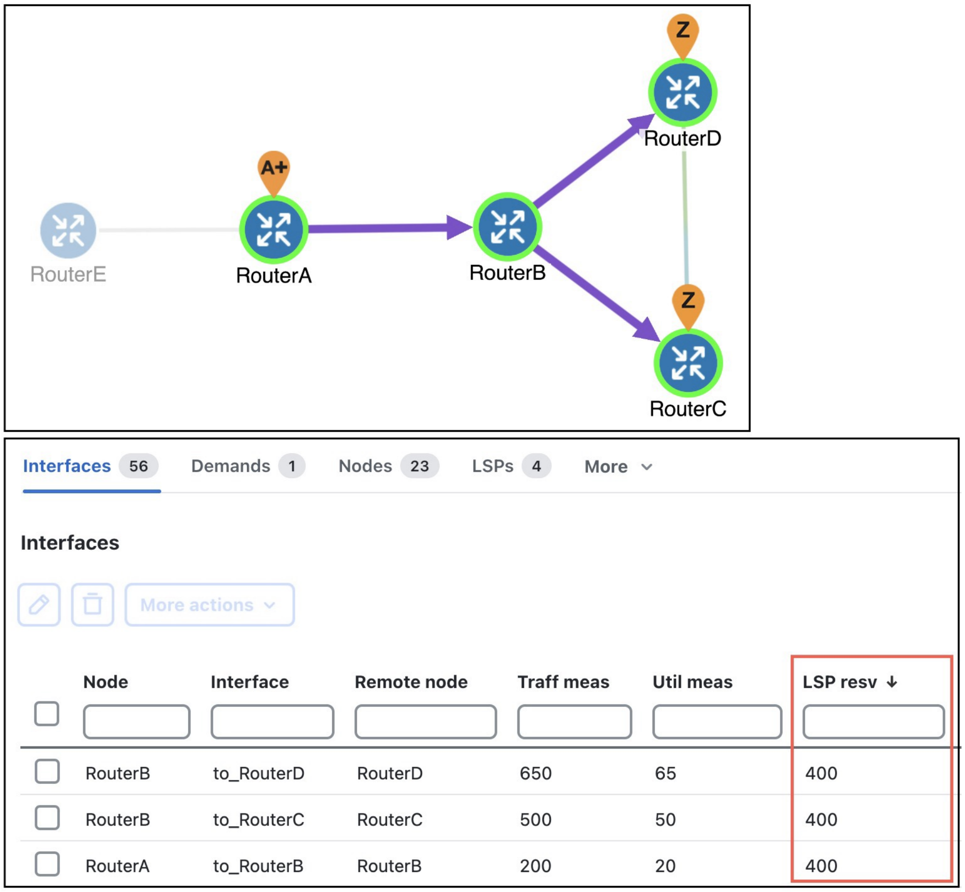

Node	Interface	Remote node	Traff meas	Util meas	LSP resv
RouterB	to_RouterD	RouterD	650	65	400
RouterB	to_RouterC	RouterC	500	50	400
RouterA	to_RouterB	RouterB	200	20	400

P2MP LSP demands

Demands routed through P2MP LSPs must have the same source as the P2MP LSP; the demand then terminates at each of the P2MP LSP destinations. No other demands can go through the P2MP LSP except for the one created for it. If the P2MP LSP fails, its demand cannot be routed. For information on creating these demands, refer to Create demands for P2MP LSPs, on page 53.

Demands must be private to the P2MP LSP. This privacy is automatically created when you insert a demand for a P2MP LSP.

Requirements for P2MP LSPs and sub-LSPs

Rules and guidelines:

- All sub-LSPs within a given P2MP LSP must have the same source site and source node.
- Because sub-LSPs within a P2MP LSP share setup bandwidth, typically bandwidth is set to the same value for each sub-LSP.

Create P2MP LSPs

Use this task to create P2MP LSPs by selecting sources, destinations, bandwidth settings, and routing properties.

Procedure

Step 1 Create an empty P2MP LSP.

a) From the toolbar, choose **Actions** > **Insert** > **LSPs** > **P2mp LSP**, or click + in the **P2MP LSPs** table.
b) Enter a unique P2MP LSP name.
c) Choose the source site and node. This selection is the root for all sub-LSPs contained in this P2MP LSP.
d) (Optional) Enter or choose a disjoint group and disjoint priority.

Step 2 If you have not done so, create sub-LSPs for the P2MP LSP. For more information, refer to Create sub-LSPs, on page 52.

Step 3 Associate the sub-LSPs to the P2MP LSP. You can use these steps individually or you can choose multiple LSPs and edit them collectively. If you are editing multiple sub-LSPs, verify that all of them have the same source site and node as the P2MP LSP.

a) Select the sub-LSPs from the LSPs table and click .
b) In the **Sub-LSP of P2MP LSP** list under the **Other** section at the bottom, choose the P2MP LSP to which you want these LSPs to belong.
c) Click **Save**.

Create sub-LSPs

Use this task to create sub-LSPs under a P2MP LSP and configure their destination and bandwidth settings.

Procedure

Step 1 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP** or click + > **LSPs** in the **LSPs** table.

Step 2 Choose the site and node that will be the source of the P2MP LSP. All sub-LSPs in a P2MP LSP must have the same source site and source node.

Step 3 Choose the site and node that are the final destination for the sub-LSP.

Step 4 Complete the remaining optional fields and tabs, as needed. Because sub-LSPs share bandwidth, typically the CSPF Setup bandwidth field is set to the same value for each sub-LSP in the P2MP LSP.

Step 5 If you have already created the P2MP LSP to which this sub-LSP will belong, choose that P2MP LSP from the Sub-LSP of P2MP LSP list (at the bottom). For creating P2MP LSPs, refer to Create P2MP LSPs, on page 52.

Step 6 Save your changes.

Create a mesh of sub-LSPs

Use this task to create a mesh of sub-LSPs from selected source and destination nodes for P2MP routing simulations.

Procedure

Step 1 In the Nodes table, choose the node you will be designating as the source for the P2MP LSP. All sub-LSPs in a P2MP LSP must have the same source node. A P2MP LSP cannot have multiple source nodes.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP mesh** or click + > **LSP mesh** in the **LSPs** table.

Step 3 In the **Destination nodes** section, uncheck these check boxes and click **Next**.

- **Same destinations as sources**
- **Create LSPs from destination to source**

Step 4 Complete the remaining optional fields as needed and click **Next**.

Step 5 To set the bandwidth to be the same for all sub-LSPs in this mesh, choose **Fixed Value** in the **Bandwidth** drop-down list. Enter that value in the **Fixed value** field (in Mbps).

To use the auto bandwidth feature so that these LSPs have their Setup BW sim value automatically calculated and used in auto bandwidth simulations, choose **Auto Bandwidth** in the **Bandwidth** drop-down list.

Step 6 Complete the remaining optional fields as needed and click **Submit**.

Create demands for P2MP LSPs

To simulate traffic routing through a P2MP LSP, create a demand for it.

Procedure

Step 1 If you want to create demands for P2MP LSPs, choose one or more P2MP LSPs from the P2MP LSP table.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Demands for P2MP LSPs** or click + > **Demands for P2MP LSPs** in the **LSPs** table.

Step 3 Choose the P2MP LSPs through which you are routing traffic.

Step 4 (Optional) Choose a service class for the newly created demand.

Step 5 Choose the desired bandwidth: **Minimum sub-LSP Setup BW** or **Zero**.

Step 6 Submit your changes.

Deletion behavior for P2MP LSPs and sub-LSPs

Deleting a P2MP LSP will invalidate the simulation and remove all demands that use the P2MP LSP as a private LSP. If you do not want all the sub-LSPs deleted, disassociate them first.

Disassociate sub-LSPs from their P2MP LSPs

Use this task to disassociate sub-LSPs from P2MP LSPs so selected sub-LSPs can be preserved during deletion.

Procedure

Step 1 In the **LSPs** table, choose one or more sub-LSPs.

Step 2 Click .

Note

If you are editing a single sub-LSP, you can also use the ··· > **Edit** option under the **Actions** column.

Step 3 In the Sub-LSP of P2MP LSP list, perform one of these actions:

- Choose the empty row to disassociate the sub-LSP and change it to an LSP.
- Choose a different P2MP LSP to associate the sub-LSP with a different P2MP LSP.

Step 4 Click **OK**.

Delete sub-LSPs or P2MP LSPs

Use this task to delete sub-LSPs or P2MP LSPs and choose the cleanup behavior for associated objects.

Procedure

Step 1 To delete sub-LSPs but keep the P2MP LSP:

a) Select one or more sub-LSPs in the **LSPs** table.

b) Click . If you are deleting a single sub-LSP, you can also use the ··· > **Delete** option under the **Actions** column.

c) Click **Delete** in the confirmation dialog box to continue.

Step 2 To delete the P2MP LSP:

a) Select one or more P2MP LSPs in the **P2MP LSPs** table.

b) Click . If you are deleting a single P2MP LSP, you can also use the ··· > **Delete** option under the **Actions** column.

c) Click **Delete** in the confirmation dialog box to continue.

Unresolved LSP destinations and hops

RSVP-TE LSP configurations and state are read using network discovery from the source node of the LSP, through SNMP, the Parse config tool, or other methods. These configurations might reference nodes or interfaces that do not exist in the plan file, so they are referred to as being unresolved. For example, the LSP destination node, read from the configuration on the source, might not be in the plan file.

There are several reasons for these differences:

- The plan file was modified to contain fewer nodes than are in the IGP.
- Cisco Crosswork Planning cannot read all the nodes, or cannot obtain the IP addresses of all the nodes.
- The LSPs themselves are not configured correctly.

References that might not be resolved include:

- Destination nodes of the LSPs.
- Hops in named paths configured on the source node.
- Hops in the actual paths read from the source node.

Cisco Crosswork Planning tries to resolve as many of these references as possible. Tables and columns are updated:

- If the LSP destinations are resolved, they are updated in the Destination column of the LSPs table. If they are not resolved, the column remains empty. The original IP address remains in the NetInt Destination column regardless of whether the destination is resolved or not.
- If the hops are resolved, they are updated in the Node and Interface columns in the Named Path Hops and Actual Path Hops tables. If they are not resolved, the columns remain empty. The original IP address remains in the NetIntHop column regardless of whether the destination is resolved.

Cisco Crosswork Planning does not make further attempts to resolve these references unless you explicitly make the request. There are special cases in which this might be useful. For example, if additional nodes are added to a network from a second network discovery procedure and if LSPs in the original network have unresolved references to these nodes, it is useful to resolve the LSPs again.