



RSVP-TE Simulation

This chapter describes how WAE Design simulates RSVP-TE routing. Unless otherwise noted, *LSPs* refer to RSVP-TE LSPs. Note that WAE Design does not distinguish between LDP and RSVP-TE LSPs.

This section contains the following topics:

- [Dynamic LSP Routing and CSPF, on page 1](#)
- [RSVP LSP Paths , on page 4](#)
- [Named Paths and Explicit LSP Routing , on page 6](#)
- [Actual Paths, on page 9](#)
- [Affinities , on page 10](#)
- [Simulation Modes, on page 14](#)
- [Advanced RSVP-TE LSP Simulations, on page 29](#)

Dynamic LSP Routing and CSPF

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

If there are equal-cost routes, the following selection criteria apply in this order:

- Bandwidth available—The route with the greatest reservable bandwidth is chosen.
- Hop count—The route with the fewest hops is chosen.
- Random—If neither of the above criteria can be used, the route is randomly chosen.

CSPF properties are set in the LSP Properties dialog box and are viewable from the LSPs table.

| | |
|------------------|--|
| CSPF | |
| Setup bandwidth: | <input type="text"/> <input checked="" type="checkbox"/> Autobandwidth |
| Setup priority: | <input type="text" value="7"/> Hold priority: <input type="text" value="7"/> |
| Hop limit: | <input type="text" value="2"/> <input type="checkbox"/> TE Metric Disabled |

- Setup Bandwidth—The amount of traffic the source node requests for this LSP in Mbps. The requested bandwidth is available from the reservable bandwidth of each interface in the path.

- **Autobandwidth**—Dynamically update the Setup BW Sim value when using the Autobandwidth Convergence mode.
- **Setup Priority**—A priority for allocating reservable bandwidth. This is the order in which the LSPs are signaled. The lower the number, the higher the priority.
- **Hold Priority**—Priority that can preempt LSPs that are 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**—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**—If checked, the LSP is routed using IGP metrics. If unchecked (default), the LSP is routed using TE metrics.

Interface MPLS Properties

Related MPLS properties are set in the MPLS tab of the interface or circuit Properties dialog box, and are viewable in the Interfaces table.

The screenshot shows the 'MPLS' tab of the 'Interface Properties' dialog box. The 'Selected Interfaces (1/32)' section contains the following settings:

- Reservable BW: [Empty text box]
- Reservable BW (%): 25
- TE Metric: 10
- Affinities: 0
- Buttons: Edit...
- Checkboxes: TE Enabled, FRR Enabled

A vertical ID '381119' is located on the right side of the dialog box.

- **TE Metric**—Determines which path an LSP takes.
- **TE Metric Sim**—Derived column that shows the effective TE metric. If the TE Metric is empty, this is set to the IGP metric.
- **TE Enabled**—Identifies whether an LSP can be routed over the interface. If checked, the value is set to T (True), and TE metrics are used in routing the LSP provided the LSP's TE Metric Disabled is unchecked (False).
- **Resv BW**—Shows the amount of bandwidth that can be reserved by LSPs routed over the interface.
- **Resv BW Percent**—Shows the percentage of bandwidth that can be reserved by LSPs routed over the interface.
- **Resv BW Sim**—Derived column.
 - If a Resv BW value is entered, this is copied to the Resv BW Sim column.

- If Resv BW and Resv BW Percent are “na,” the Resv BW Sim is copied from the Capacity Sim column.
- If Resv BW is “na,” but Resv BW Percent has a value, the Resv BW Sim value is derived by this formula:

$$\text{Capacity Sim} * (\text{Resv BW Percent} / 100)$$

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, see [Ignoring Reservable Bandwidths for Capacity Planning, on page 29](#).

- Affinities—Lists which affinities are set. For more information, see [Affinities , on page 10](#).
- Loadshare—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).
- 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, see [Fast Reroute Simulations, on page 15](#).

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.

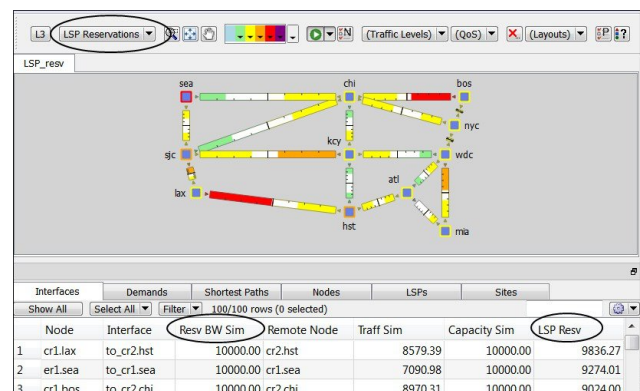
- Traff Sim LSP—The total amount of LSP traffic on the interface.
- Traff Sim Non LSP—The total amount of non-LSP traffic on the interface.

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, right-click one and choose **Filter to Demands Through Interfaces**. Then choose either **In LSP** or **Not in LSP**

Viewing LSP Reservations

To view the LSP reservations on the interfaces, choose **LSP Reservations** from the Network Plot View in the Visualization toolbar ([Figure 1: LSP Reservations View, on page 3](#)). The width of the circuit displayed in the network plot is set according to the Resv BW Sim column in the Interfaces table. The utilization colors are based on the LSP Util column in the Interfaces table. The LSP Util is the sum of all setup bandwidths for all LSPs through the interface, as a percentage of the Resv BW Sim.

Figure 1: LSP Reservations View



LSP Setup BW Initializer

The LSP Setup BW 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. For information, see [Advanced RSVP-TE LSP Simulations, on page 29](#).

-
- Step 1** Choose one or more LSPs from the LSPs table.
 - Step 2** Choose **Initializers > LSP Setup BW**, or right-click one of the selected LSPs and choose **LSP Setup BW Initializer**.
 - Step 3** Choose which traffic levels to use for demand traffic.
 - Step 4** Choose whether to set the load share values to be equal to the setup bandwidth.
 - Step 5** Click **OK**.
-

RSVP LSP Paths

Like LSPs, LSP paths have CSPF properties, such as Setup Priority and Hop Limits. For a description of LSP properties, see [Dynamic LSP Routing and CSPF, on page 1](#).

If these properties are omitted, then they are inherited from the LSP. If these properties are set in the LSP path, they override the LSP settings. Note, however, if operating in the Autobandwidth Convergence simulation mode and if the LSP has Autobandwidth set to T (true), 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.

Hot Standby

An LSP path can be defined as a “hot” standby path. Standby paths are always established even if the associated LSP is routed using a path with a different path option. For RSVP LSPs, setup bandwidth, if any, is reserved for them. These standby LSP paths are then immediately available should the currently routed path become unavailable.



Example Response to Failures

Figure 2: Example RSVP LSP and Associated LSP Paths, on page 5 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 Named Path column. The other two paths do not have associated named paths, and so are dynamically routed.

Figure 2: Example RSVP LSP and Associated LSP Paths

| Interfaces | Demands | Shortest Paths | Nodes | LSPs | LSP Paths | Named Paths |
|--|---------|----------------|----------|-------------|---------------------|-------------|
| Show All Select All Filter 1/1 rows (0 selected) | | | | | | |
| Name | Source | Destination | Setup BW | Active Path | Named Paths | |
| 1 cr1.lon_cr2.fra | cr1.lon | cr2.fra | 125.00 | 100 | cr1.lon_cr2.fra_100 | |

| Interfaces | Demands | Nodes | LSPs | LSP Paths | Sites | SRLGs | AS |
|--|---------------------|-------|----------|-------------|---------|-------|----|
| Show All Select All Filter 4/4 rows (0 selected) | | | | | | | |
| LSP | Path Name | | Setup BW | Path Option | Standby | | |
| 1 cr1.lon_cr2.fra | cr1.lon_cr2.fra_100 | | 0.00 | 100 | T | | |
| 2 cr1.lon_cr2.fra | cr1.lon_cr2.fra_200 | | 0.00 | 200 | T | | |
| 3 cr1.lon_cr2.fra | | | na | 300 | F | | |
| 4 cr1.lon_cr2.fra | | | 0.00 | 400 | F | | |

This example LSP responds to failures as follows:


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.

- 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), WAE Design 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 through the network using an ordered list of adjacencies. The route is defined by named path hops, which can be nodes or interfaces, and each hop type specifies whether the route should be strict, loose, or excluded. Named path hops can be nodes or interfaces. The hop type is identified visually on the plot, as well as listed in the Type column of the Named Path Hops table.

- **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.



Strict



Loose



Exclude

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 WAE Design simulation.

Named Path Hops Example

Figure 3: Example Named Paths for the `cr1.lon_cr2.fra` LSP, on page 7 extends the earlier example (Figure 2: Example RSVP LSP and Associated LSP Paths, on page 5), 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 3: Example Named Paths the cr1.lon_cr2.fra LSP

| Interfaces | | LSPs | | LSP Paths | | Named Paths | |
|------------|---------------------|--------------|---------|-----------|--------|------------------|--|
| Show All | | Select All ▼ | | Filter ▼ | | 2/2 rows (0 sele | |
| | Name | | Source | | Active | | |
| 1 | cr1.lon_cr2.fra_100 | | cr1.lon | | T | | |
| 2 | cr1.lon_cr2.fra_200 | | cr1.lon | | T | | |

The hops for each named path are defined in a separate table called Named Path Hops. [Figure 4: Example Named Path Hops for cr1.lon_cr2.fra_100, on page 7](#) shows the named path hops for the cr1.lon_cr2.fra_100 named path. The Steps 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 4: Example Named Path Hops for cr1.lon_cr2.fra_100

| Interfaces | | LSPs | | LSP Paths | | Named Paths | | Named Path Hops | | Site |
|------------|---------------------|--------------|---------|-----------|---------|-----------------------|--------|-----------------|--|------|
| Show All | | Select All ▼ | | Filter ▼ | | 3/3 rows (0 selected) | | | | |
| | Name | | Source | Step | Node | Interface | Type | | | |
| 1 | cr1.lon_cr2.fra_100 | | cr1.lon | 1 | cr1.par | {to_cr2.par} | strict | | | |
| 2 | cr1.lon_cr2.fra_100 | | cr1.lon | 2 | cr2.par | | strict | | | |
| 3 | cr1.lon_cr2.fra_100 | | cr1.lon | 3 | cr1.fra | {to_cr2.fra} | strict | | | |

Creating Named Paths and Their Hops

Named paths can be created in the WAE Design GUI in several ways.

Create associated named paths: \$1_\$2
 (\$1 is LSP Name, \$2 is Path Option)

- Check **Create associated named paths** when creating LSP paths. Using this option does not create hops.
- Use the Explicit LSP Paths Initializer. Using this initializer creates named paths with strict hops.
- Optimize LSP paths using the Explicit RSVP-TE LSP Optimization tool or RSVP-TE Optimization tool.

Editing Named Paths and Named Path 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 on resolving named paths, see [Unresolved LSP Destinations and Hops, on page 42](#).

The recommended method of editing the named path hop type is to access it directly from the Named Path table as described in the following steps. This is the most efficient way of viewing the entire path. However, if needed, you can double-click a named path hop in the Named Path Hops table, which takes you directly to the named path hops Properties dialog box, which is described in Step 4.

Step 1 In the Named Paths table, double-click a named path or right-click one and choose **Properties**.

Step 2 Named path options:

- To change the name, enter it in the Name field.
- To change the source site or source nodes, choose them from the drop-down lists.
- To activate or deactivate, check **Active**.

Step 3 Hop options:

- To add a new hop, choose an existing hop (if applicable) from the Hops list, and then click **Insert Before** or **Insert After** to determine where the new hop goes sequentially. Go to Step 4.
- To edit an existing hop, choose it from the Hops list and then click **Edit**. The Properties dialog box opens for named path hops. Go to Step 4.
- To delete an existing hop, choose it from the Hops list, click **Delete**, and then **OK**. You need not continue.
- To resolve nodes and interfaces that do not exist in the plan file, click **Resolve**. For information, see [Resolving Plan File to Objects, on page 43](#).

Step 4 To continue creating or editing the named path hop, use the following options and then click **OK**

- As needed, choose the site, node, and interface. For node hops, do not choose an interface.
- Choose the named path hop type as Loose, Strict, or Exclude. See [Named Paths and Explicit LSP Routing , on page 6](#) for a description of each.

Actual Paths

In WAE Design, actual paths are the actual route that routed LSP paths are taking. They are read from the live network, per LSP path. 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. WAE Design 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, WAE Design reverts to the standard CSPF routing algorithm.

Example: WAE Design 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.

It might be desirable to simulate a network without taking actual paths into account, for example, for long-term planning. To achieve this, actual paths should simply be deleted. To delete an actual path, right-click an LSP or LSP path and choose **Delete Actual Paths**.

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.

Actual paths are represented as a series of interface hops in the Actual Path Hops table. A hop is shown in the network plot as a brown circle on the interface. To view an entire actual path of an LSP, choose **View > LSPs > Actual Paths**; all hops are displayed when the LSP is selected.




Two columns in the LSP table are useful to see the result of actual paths on LSP routing.

| Column | Description |
|--------|--|
| Actual | <p>The simulation can only use actual paths if they are resolved. If the network discovery was incomplete, this might not be possible.</p> <ul style="list-style-type: none"> • T = true; the simulation converted the actual path hops into a complete path from source to destination of the LSP • F = false; the simulation did not convert the actual path hops into a complete path • na = not applicable; no actual path was available. |

| Column | Description |
|--------|---|
| Route | <ul style="list-style-type: none"> • Actual = The LSP follows an actual path. • Simulated = The LSP does not follow an actual path. • Unrouted = Routing was not possible. |

Deactivating Actual Paths for Simulations

WAE Design uses network state in its MPLS simulation, routing LSPs on actual paths where possible and using LSP active path settings. For planning purposes, when state is not relevant, you can change this behavior to disregard actual paths.



-
- Step 1** Choose **Edit > Network Options** or click the **Network Options** icon in the Visualization toolbar.
- Step 2** Click the **Simulation** tab.
- Step 3** To use or disregard actual paths and active paths in MPLS simulation, check or uncheck **Use LSP actual paths, active paths**.
- Step 4** Click **OK**.
-

Affinities

Affinities provide a mechanism for implementing LSP path diversity. By assigning affinities to physical circuits and associating LSPs with affinities, you can 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.

WAE Design 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.

Workflow:

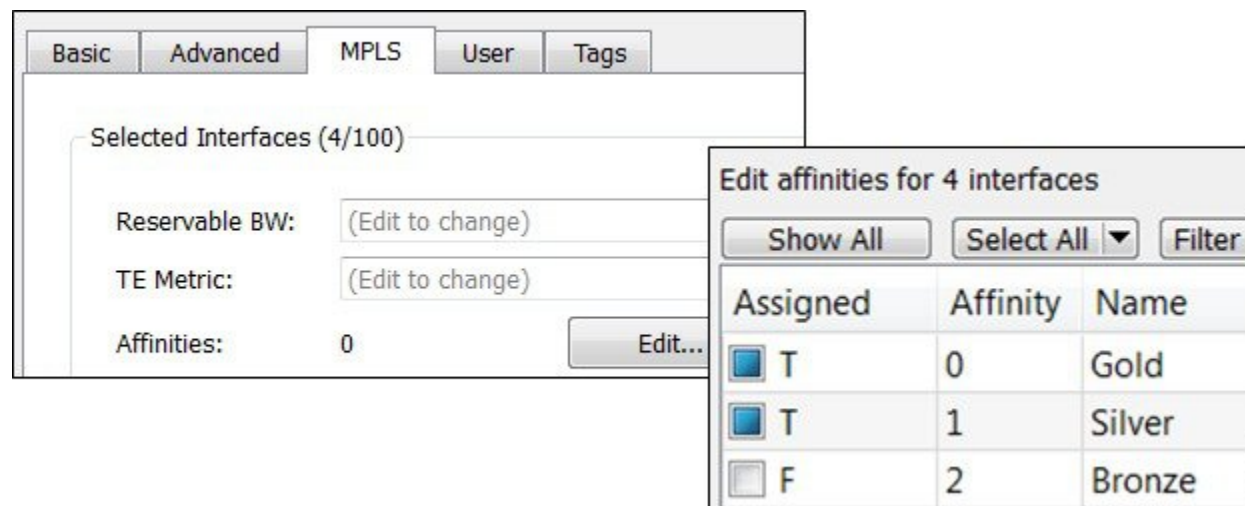
1. [Creating and Editing Affinities, on page 11.](#)
2. [Assigning Affinities to Interfaces, on page 11.](#)
3. [Associating LSPs with Affinities, on page 12.](#)

Creating and Editing Affinities

- Step 1** Choose **Edit > Admin Groups**.
- Step 2** Modify an existing affinity number and name, or click **New** to add the number and name.
Affinity numbers must be unique. Affinity names are optional and must also be unique.
- Step 3** Click **OK**.

Assigning Affinities to Interfaces

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



- Step 1** From the Interfaces table, choose one or more interfaces to which you are assigning the affinities.
- Step 2** Double-click the selection or choose **Edit > Properties**.
- Step 3** Click the **MPLS** tab.
- Click the **Affinities Edit** button.
 - In the Include column, check the boxes for the affinities that you want to associate with the selected circuits. T = true; associate the affinity with the circuit. F = false; do not associate the affinity.
 - Click **OK**.
- Step 4** Click **OK**.

You can follow this procedure to assign affinities to interfaces through a circuit Properties dialog box. However, you must choose and assign the affinity twice: once to each interface.

Associating LSPs with Affinities


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. (See [Named Paths and Explicit LSP Routing](#) , on page 6.)

Inclusion and Exclusion Rules

- 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.

Figure 5: LSP Path Affinities




Example

This example ([Figure 6: Example Affinities, on page 13](#)) 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 0 (Gold) and 1 (Silver). No other interfaces are assigned affinities.
- LSP A is configured to exclude all interfaces assigned to affinity 0. 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 to exclude interfaces assigned to affinity 2 (Bronze). Because the to_cr1.nyc interface is assigned to both of these affinities, LSP B cannot route.

Figure 6: Example Affinities



LSP A

| Include | Include Any | Exclude | Affinity | Name |
|----------------------------|----------------------------|---------------------------------------|----------|--------|
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 0 | Gold |
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 1 | Silver |
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input checked="" type="checkbox"/> T | 2 | Bronze |

| Interfaces | | |
|------------|------------|---------------|
| Node | Interface | Affinities |
| cr2.wdc | to_cr1.nyc | Silver;Bronze |

LSP B

| Include | Include Any | Exclude | Affinity | Name |
|----------------------------|----------------------------|----------------------------|----------|--------|
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 0 | Gold |
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 1 | Silver |
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 2 | Bronze |

LSP Path for LSP B

| Include | Include Any | Exclude | Affinity | Name |
|----------------------------|---------------------------------------|---------------------------------------|----------|--------|
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input type="checkbox"/> F | 0 | Gold |
| <input type="checkbox"/> F | <input checked="" type="checkbox"/> T | <input type="checkbox"/> F | 1 | Silver |
| <input type="checkbox"/> F | <input type="checkbox"/> F | <input checked="" type="checkbox"/> T | 2 | Bronze |

408295

Assigning LSPs to Affinities

- Step 1** From the LSPs table, choose one or more LSPs to which you are associating the affinities.
- Step 2** Double-click the selection or choose **Edit > Properties**.
- Step 3** Click the **Affinities** tab.
- Step 4** Choose the inclusion or exclusion rule for each affinity you are associating with the selected LSPs.
- Step 5** Click **OK**.

Assigning LSP Paths to Affinities

- Step 1** From the LSP Paths table, choose one or more LSP paths to which you are associating the affinities.
- Step 2** Double-click the selection or choose **Edit > Properties**.
 - a) Click the **Affinities Edit** button. For each rule (Include, Include Any, Exclude), choose whether the LSP path should inherit the LSP affinity rule or whether it should be based on rules defined in the table below these options.
 - b) If you chose to use the table, choose the rule for each affinity you are associating with the LSP path.

c) Click **OK**.

Step 3 Click **OK** to save and edit.

Assigning Affinities When Creating LSP Meshes

You can associate affinities with LSPs when creating a new LSP mesh.


Step 1 Choose **Insert > LSPs > LSP Mesh**, or right-click in the plot and choose **New > LSPs > LSP Mesh**.

Step 2 Click the **Affinities Edit** button. Choose the inclusion or exclusion rule for each affinity you are associating with all LSPs in the mesh, and click **OK**.

Step 3 Click **OK** to save and exit.

Simulation Modes

WAE Design lets you set several global parameters that affect how LSPs are routed or rerouted. To access these options, either choose **Edit > Network Options** or click the **Network Options** icon in the Visualization toolbar. Then click the **Simulation** tab.



WAE Design supports four simulation modes for RSVP-TE LSPs:

- A *Fast Reroute* (FRR) mode using FRR LSPs, which is a common failure restoration mechanism used by RSVP-TE LSPs. In a live network, FRR restoration typically occurs within milliseconds. See [Fast Reroute Simulations, on page 15](#).
- By default, WAE Design 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. This is called the *IGP and LSP Reconvergence* simulation mode. The Optimization tools only work in IGP and LSP Reconvergence mode.
- In the Autobandwidth Convergence 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. See [Autobandwidth Simulations, on page 25](#).
- In the Autobandwidth Convergence Including Failures mode, the network is simulated after the Setup BW Sim values have been reset. See [Autobandwidth Simulations, on page 25](#).

The status bar in the lower, left of the GUI identifies which simulation modes you are currently using. Unless specified otherwise, the documentation describes behavior in the IGP and LSP Reconvergence mode.

Fast Reroute Simulations

WAE Design 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 differ in FRR simulations compared to full convergence simulations, as follows:

- 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 following routing occurs:
 - 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.
- Because 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. More details about these are further explained throughout the following sections. Note that FRR LSPs appear in the same table as do regular LSPs and are identified in the Metric Type column as FRR.

Objects and Properties

| Object | Type | Property | 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 WAE Design, this requires the Fast Reroute simulation mode. To turn this mode on, access the Network Options dialog box. For more information, see [Running Fast Reroute Simulation, on page 24](#).

- Protected SRLG—An SRLG that contains a protected interface.

Visualization

An unused FRR LSP path highlights in magenta when it is selected from the LSPs table. When routed, FRR LSPs appear as a dashed magenta line.

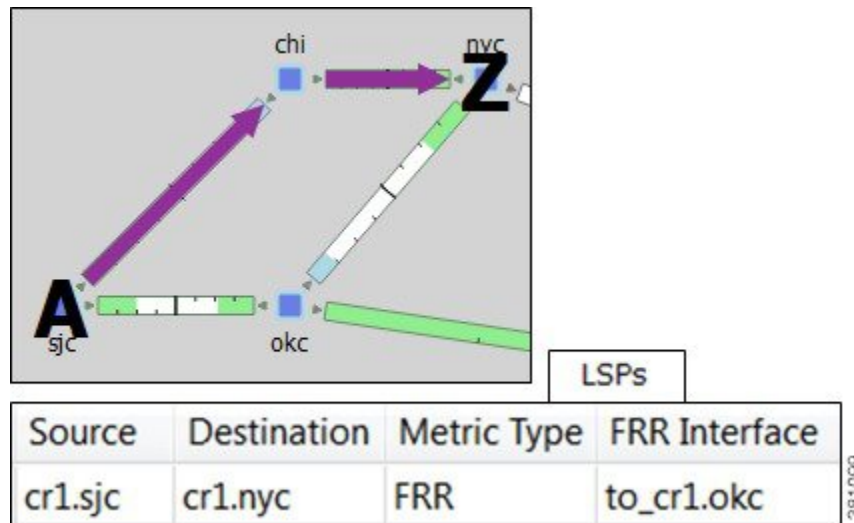


FRR LSP Routing

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 LSP Properties dialog box when manually creating the FRR LSP, or it is automatically created when running the FRR LSPs initializer. The initializer derives it from interface FRR Enabled property.

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



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.

Link and Node Protection

WAE Design supports both link (circuit) and node protection FRR LSPs. The type of protection shows in the FRR Type column of the LSPs table.

| LSPs | |
|-------------|----------|
| Metric Type | FRR Type |
| FRR | Link |
| FRR | Node |

- 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. [Figure 7: Example Full Convergence Versus Fast Reroute with Link Protection, on page 18](#) shows an example of an LSP from sjc to nyc with a failure on the circuit between slc and hst. The top portion shows the Full Convergence mode where the source node of the LSP is signaled to reroute the traffic. The lower portion shows that a link-protection FRR LSP was set up for the slc-hst interface, and that the simulation is operating in Fast Reroute mode. Here, slc reroutes the LSP traffic around 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. [Figure 8: Example Full Convergence Versus Fast Reroute with Node Protection, on page 18](#) shows an example of an LSP from sjc to nyc with a failure on the hst node. The top portion shows the Full Convergence mode where the source node of the LSP is signaled to reroute the traffic. The lower portion shows that a node-protection FRR LSP was set up to protect the hst node, and that the simulation is operating in Fast Reroute mode. Here, the slc node reroutes the LSP traffic around the node failure.

Figure 7: Example Full Convergence Versus Fast Reroute with Link Protection

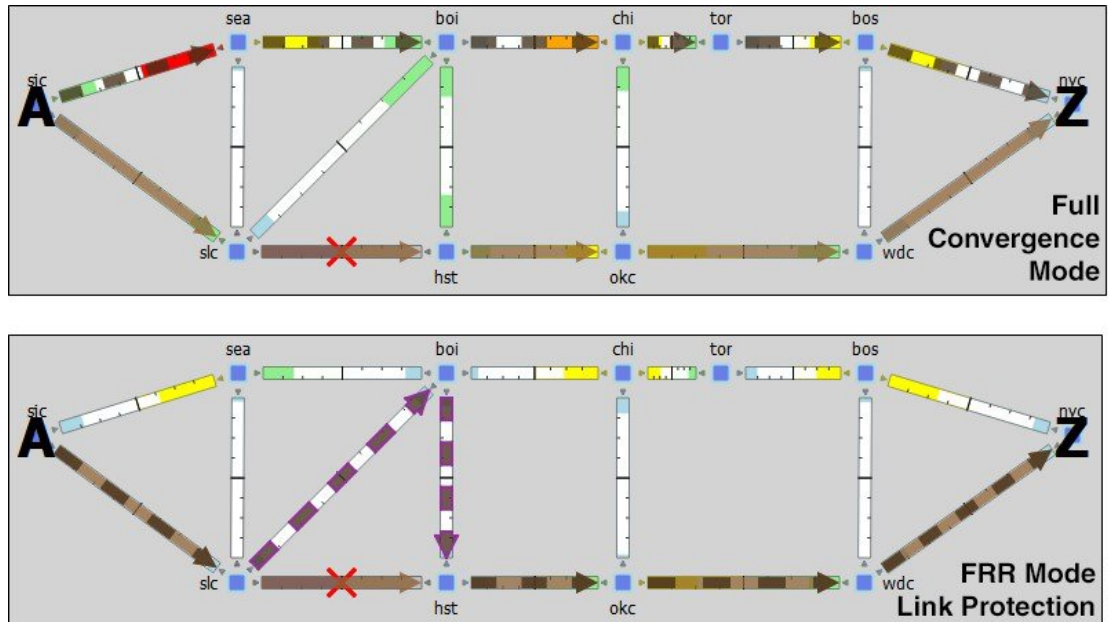



Figure 8: Example Full Convergence Versus Fast Reroute with Node Protection



SRLG Protection


Figure 9: Example Fast Reroute without and with SRLG Protection, on page 19 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 WAE Design, FRR LSPs can be configured to avoid all objects defined in the SRLG, not just interfaces on the source node of the FRR LSP.

Example: shows an example that uses the following setup:

- A protected LSP from cr2.hst to cr2.wdc.
- The protected interface is to_cr2.atl, and its source node is cr2.hst.
- An SRLG containing three circuits: cr2.hst_cr2.chi, cr2.hst_cr2.atl, and cr2.hst_cr2.mia.
- This SRLG has been configured on the cr2.hst node.

When the cr2.hst_cr2.atl circuit fails and there is no SRLG protection, circuits within the SRGL can be used to reroute the LSP. When there is SRLG protection in place, the LSP is rerouted to avoid all circuits within the SRLG.

Figure 9: Example Fast Reroute without and with SRLG Protection



382648

Routing Around Failures

Assuming one or more failures occurred, the following decisions determine which FRR LSP to take.

Prerequisites:

- LSPs must be configured to be protected by FRR LSPs. For information on setting up FRR LSPs, see [Setting Up Fast Reroute Simulations, on page 20](#).
- Any interfaces along the FRR LSP path must be configured to be protected by FRR LSPs. For information on protecting interfaces, see [Identifying LSPs for Protection, on page 20](#) or the [FRR LSPs_INITIALIZER, on page 21](#) (which automates the process of setting these).

Step 1

For each such interface on the LSP path, WAE Design 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, WAE Design 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.

Setting Up Fast Reroute Simulations

- Step 1** Set appropriate properties on LSPs to mark them for protection. See [Identifying LSPs for Protection, on page 20](#).
- Step 2** If protecting SRLGs, associate a node with the SRLG in order to establish the protected interface within it. See [Identifying SRLGs for Protection, on page 20](#).
- Step 3** If you are using the initializer to create FRR LSPs, identify which interfaces to protect. See [FRR LSPs Initializer, on page 21](#).
- Step 4** Create the FRR LSPs using the FRR LSPs initializer or manually. See [FRR LSPs Initializer, on page 21](#) or [Manually Creating FRR LSPs, on page 23](#).

Identifying LSPs for Protection

To be marked for Fast Reroute protection, LSPs must have two properties set in the LSP Properties dialog box.

- Routing type must be Autoroute or Forwarding Adjacency. The type appears as “autoroute” or “FA” in the Metric Type column of the LSPs table.
- FRR Enabled must be checked. This appears as T (true) in the FRR Enabled column of the LSPs table. If this is F (false), the LSP is not protected.

Identifying 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.

| Included | SRLG |
|---------------------------------------|--------------|
| <input checked="" type="checkbox"/> T | Midwest_SRLG |
| <input type="checkbox"/> F | South_SRLG |
| <input type="checkbox"/> F | North_SRLG |

- Step 1** Right-click the source node of a circuit within an SRLG that you want to protect and choose **Properties**.

Step 2 Click the **SRLGs** tab.

Step 3 Choose one or more SRLGs with which you want to associate this node, and click **OK**.

FRR LSPs Initializer

The FRR LSPs initializer automatically creates link-protection or node-protection FRR LSPs provided two conditions are met:

- One or more LSPs have their FRR Enabled property set. (See [Identifying LSPs for Protection](#), on page 20.)
- 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 interface Properties dialog box. The Interfaces table > FRR Enabled column shows T (true) if the interface is included or F (false) if it is not.
 - After FRR LSPs are created, the interface is listed in the LSPs table > FRR Interface column.

Interface Properties

Basic | **Advanced** | MPLS | User

Selected Interfaces (2/32)

Reservable BW:

Reservable BW (%):

TE Metric:

Affinities:

TE Enabled FRR Enabled

Interfaces

| Interface | FRR Enabled |
|------------|-------------|
| to_cr1.wdc | T |
| to_cr1.chi | T |

FRR LSPs Initializer Uses this Interface
FRR Simulations Route Away from this Interface

38118


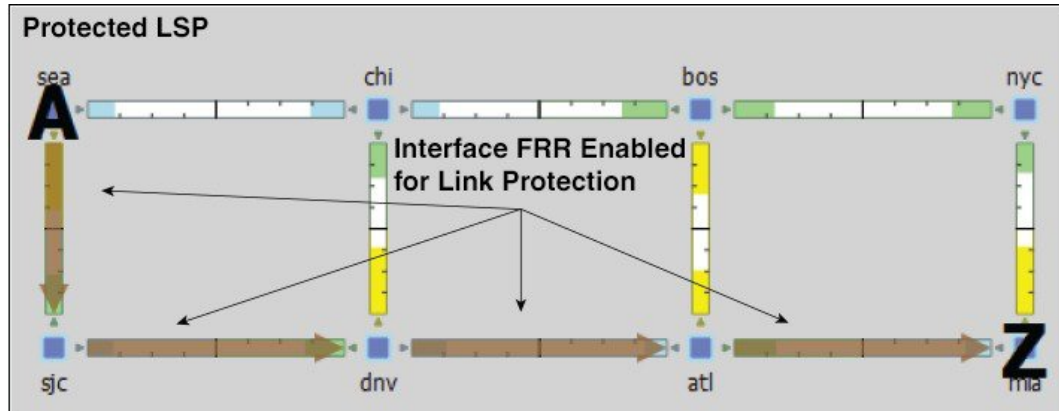
The initializer creates the FRR LSPs based on these FRR Enabled properties and based on whether you choose to create link or node protection.

- Link protection—WAE Design 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 ([Figure 10: Example Link-Protection FRR LSPs Automatically Created](#), on page 22).
- Node protection—WAE Design creates FRR LSPs between each set of next next-hop nodes (three connected nodes) in the protected LSP path where the egress interface by of the first hop has its FRR

Enabled property set (Figure 11: Example Node-Protection FRR LSPs Automatically Created, on page 23).


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

Figure 10: Example Link-Protection FRR LSPs Automatically Created



381097

Figure 11: Example Node-Protection FRR LSPs Automatically Created



Running the FRR LSPs Initializer

Step 1 Choose one or more nodes. If you do not choose any nodes, the initializer uses all nodes in the plan file.

Link Protection
 Node Protection
 Name (FRR_source_destination_postfix)
 Postfix:
 Delete current FRR LSPs

Step 2 Choose **Initializers > FRR LSPs**.

Step 3 Choose whether to create FRR LSPs that protect links (circuits), nodes, or both.

Step 4 (Optional) Enter a postfix to add to the end of the names of the newly created FRR LSPs.

Step 5 Click **OK**.

Manually Creating FRR LSPs

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

Step 1 Choose **Insert > LSPs > LSP**, or right-click in an empty area and choose **New > LSPs > LSP**. Alternatively, you can choose an existing LSP to modify to be an FRR LSP. If so, right-click the LSP and choose **Properties**.

Source
Site: chi
Node: cr1.chi

Destination
 Site: okc
Node: cr1.okc
NetInt Destination:

Step 2 Choose the source site and node of the FRR LSP.

Step 3 Choose the destination site and node, or enter a NetInt destination IP address. For information on NetInt destination IP addresses, see [Unresolved LSP Destinations and Hops](#), on page 42.

Routing
Fast Reroute
Metric value: Load Share:
Active Path: Not set
FRR Interface: to_cr1.okc
NetInt FRRInterface:

to_cr1.bos
to_cr1.okc
to_cr1.sea

Interface to Route Away from in FRR Routing


Step 4 In the Routing area, choose **Fast Reroute** as the LSP type.

Step 5 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 FRR Interface field.

Step 6 Click **OK**.

Running Fast Reroute Simulation

FRR LSPs are routed only when simulating failures in the Fast Reroute simulation mode. To enable the Fast Reroute simulation mode, follow these steps.



-
- Step 1** Choose **Edit > Network Options** or click the **Network Options** icon in the Visualization toolbar.
- Step 2** Click the **Simulation** tab.
- Step 3** Choose **Layer 3 Fast Reroute** and click **OK**.
-

Autobandwidth Simulations

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 WAE Design does not simulate this instability, it does simulate autobandwidth-enabled LSPs. In this simulation mode, internally WAE Design 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 GUI 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.

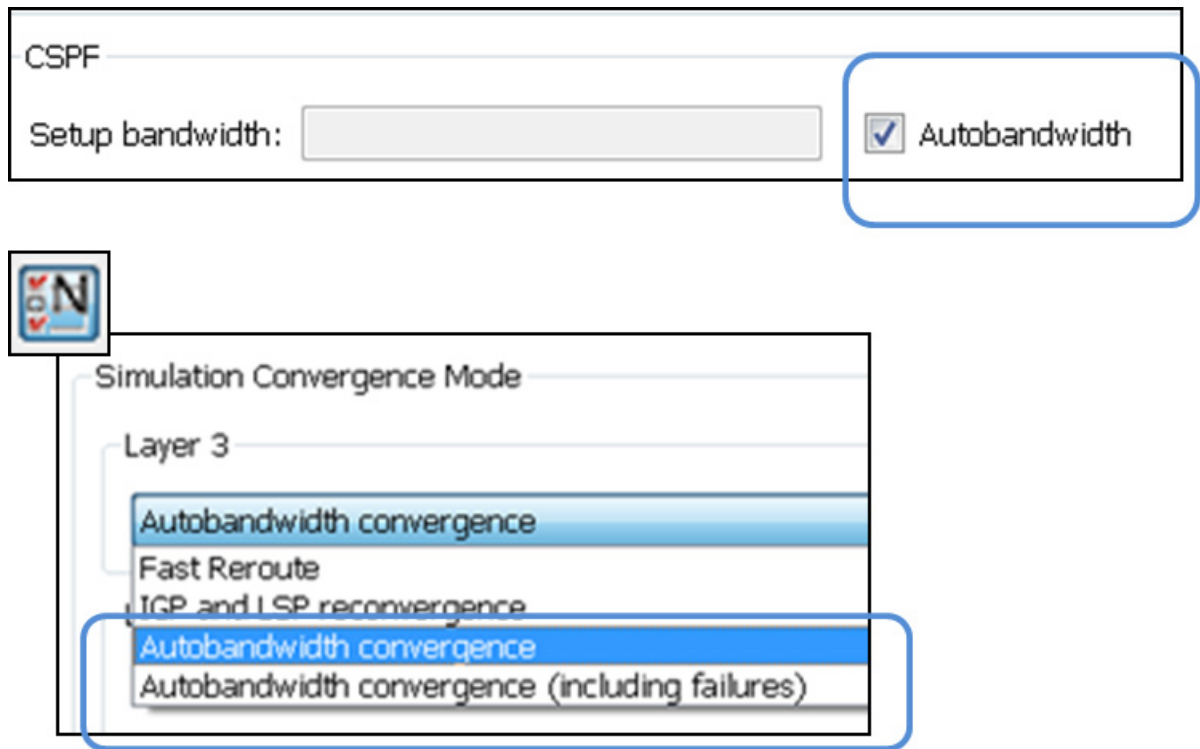
There are two autobandwidth simulation modes available. Before simulating failures or running Simulation Analysis on autobandwidth-enabled LSPs, choose the appropriate mode depending on which state you need to simulate.

- In the Autobandwidth Convergence 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.
- In the Autobandwidth Convergence Including Failures mode, the network is simulated after the Setup BW Sim values have been reset. This simulates the longer, more optimal response to failures.



Note The Autobandwidth Convergence and Autobandwidth Convergence Including Failures simulation modes are only available in plan files with a single traffic level.

These autobandwidth simulations require that the LSPs have their Autobandwidth property set to T (True), and that one of the two autobandwidth modes be selected from the Network Options dialog box.

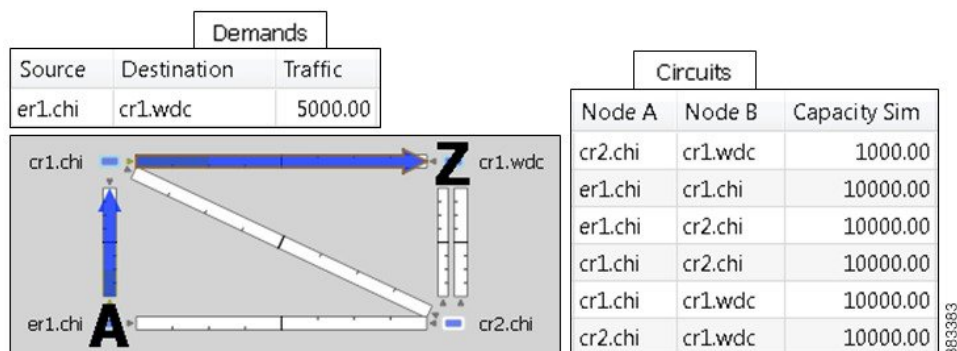


408290

Example Autobandwidth Convergence Without and with Failures

The network in this example has the following parameters:

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



- [Figure 12: Example LSP Routing Using IGP and LSP Reconvergence and Using Autobandwidth Convergence, on page 27](#) 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 13: Example Autobandwidth Convergence with Failures, on page 28](#) 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 14: Example Autobandwidth Convergence Including Failures, on page 28](#) 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 (dashed) 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 LSP Routing Using IGP and LSP Reconvergence and Using Autobandwidth Convergence

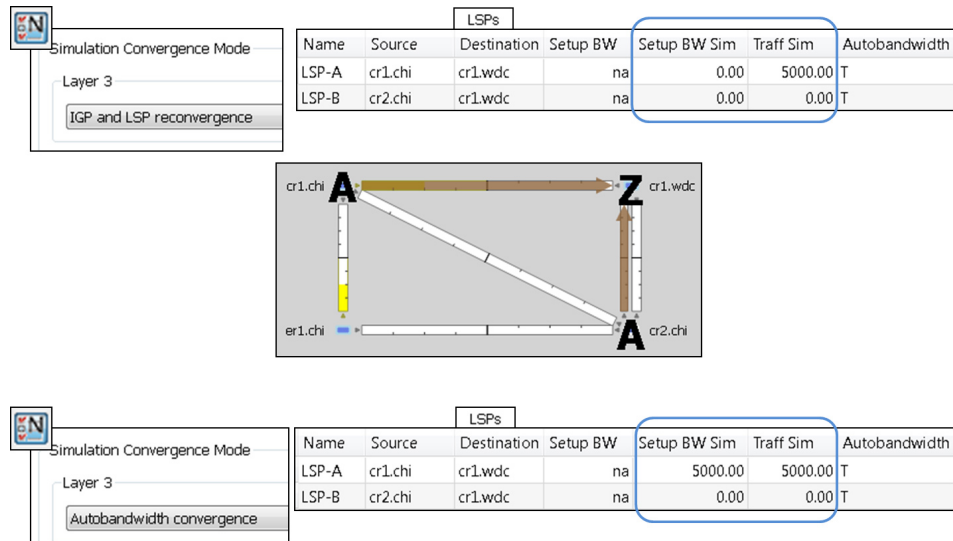


Figure 13: Example Autobandwidth Convergence with Failures





Figure 14: Example Autobandwidth Convergence Including Failures



Advanced RSVP-TE LSP Simulations

Ignoring Reservable Bandwidths for Capacity Planning

When using WAE Design 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 WAE Design simulation that can assist you when planning networks.

When you enable this routing mode, WAE Design routes LSPs in the order they appear in the plan file LSPs table, as described in the following steps. 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. For information on tables in plans files, see the *Cisco WAE Design Integration and Development Guide*.

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

Step 1 Route LSPs as usual. (For information, see [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.

Example

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.

- LSPs
 - 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 15: Comparison of Reservable Bandwidth Being Used vs Not Being Used, on page 30 shows the LSP Reservations view of this simple network in both simulation modes: the usual where reservable bandwidths are observed and the capacity planning version wherein they are ignored. In this latter mode, the two purple circuits indicate they are being over utilized and likely need their capacity increased. (For information on the LSP Reservations view, see [Viewing LSP Reservations, on page 3.](#))

Figure 15: Comparison of Reservable Bandwidth Being Used vs Not Being Used

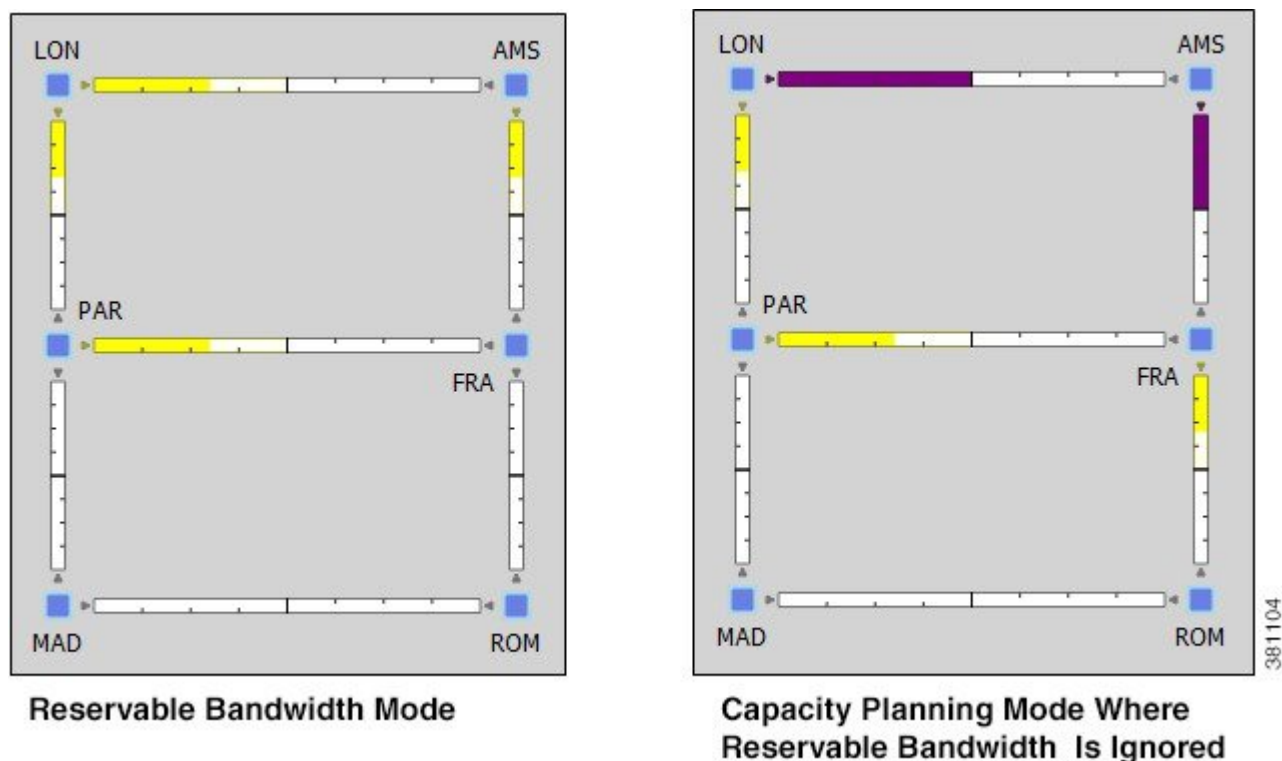



Figure 16: Simulated Routing Using Reservable Bandwidth Constraints, on page 31 demonstrates routing simulations where reservable bandwidth constraints are applied, which is the default behavior.

- 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 16: Simulated Routing Using Reservable Bandwidth Constraints




| Interfaces | | | | |
|------------|------------|---------|-------------|-------------|
| Node | Interface | Resv BW | Resv BW Sim | Remote Node |
| cr1.lon | to_cr1.par | 1000.00 | 1000.00 | cr1.par |
| cr1.ams | to_cr1.fra | 1000.00 | 1000.00 | cr1.fra |
| cr1.par | to_cr1.fra | 1000.00 | 1000.00 | cr1.fra |
| cr1.lon | to_cr1.ams | 1000.00 | 1000.00 | cr1.ams |

| LSPs | | | | | |
|-------|---------|-------------|----------|--------------|-----------|
| Name | Source | Destination | Setup BW | Setup BW Sim | Routed |
| LSP-1 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-2 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-3 | cr1.lon | cr1.fra | 600.00 | na | Unrouted |
| LSP-4 | cr1.lon | cr1.rom | 600.00 | na | Unrouted |

In Figure 17: Simulated Routing Without Using Reservable Bandwidth Constraints, on page 32, the use of reservable bandwidth constraints has been disabled. The LSP routes are now as follows.

- 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 17: Simulated Routing Without Using Reservable Bandwidth Constraints




| Interfaces | | | | |
|------------|------------|---------|-------------|-------------|
| Node | Interface | Resv BW | Resv BW Sim | Remote Node |
| cr1.lon | to_cr1.par | 1000.00 | 1000.00 | cr1.par |
| cr1.ams | to_cr1.fra | 1000.00 | 1000.00 | cr1.fra |
| cr1.par | to_cr1.fra | 1000.00 | 1000.00 | cr1.fra |
| cr1.lon | to_cr1.ams | 1000.00 | 1000.00 | cr1.ams |

| LSPs | | | | | |
|-------|---------|-------------|----------|--------------|-----------|
| Name | Source | Destination | Setup BW | Setup BW Sim | Routed |
| LSP-1 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-2 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-3 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-4 | cr1.lon | cr1.rom | 600.00 | 600.00 | Simulated |

381106


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.



381105

Enabling Capacity Planning Mode for LSPs

To ignore reservable bandwidth constraints in routing simulations, follow these steps.



- Step 1** Choose **Edit > Network Options**, or click the **Network Options** icon in the Visualization toolbar.
- Step 2** Click the **Simulation** tab.
- Step 3** Check **LSP capacity planning mode**, thus ignoring reservable bandwidth constraints, and click **OK**.

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 WAE Design 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. You can also right-click an LSP and choose **Filter to P2MP LSPs**.
- To determine which LSPs belong to a P2MP LSP, you can visually see them by choosing the P2MP LSP from the P2MP LSPs table. You can also right-click a P2MP LSP and choose **Filter to sub LSPs**.

Figure 18: P2MP LSP with chi as the Source for Three Destinations



Plotting P2MP LSPs

- Step 1** In the P2MP LSPs table, choose one or more P2MP LSPs.
- Step 2** Right-click and choose **Plot P2MP LSPs**.
- Step 3** To move the plot horizontally or vertically, use the left/right arrows and up/down arrows.
- Step 4** To choose all P2MP LSPs within the plot, click **Select P2MP LSPs**. To deselect all of these P2MP LSPs, click in an empty plot area.
- Step 5** To visualize multilayer L1 paths when plotting P2MP LSPs, use the following drop-down options:
- **L3**—Displays the path in terms of interfaces.
 - **L1**—Displays the path in terms of L1 links.
To better align the L1 links, WAE Design checks for L3-L1 links. If they do not exist, WAE Design checks for collocated L3/L1 nodes within sites.
 - **L3+L1**—Displays the L3 plot above and the L1 plot below.
WAE Design does not perform alignment between L3 and L1 nodes. When you click an L3 interface, the associated L1 links are highlighted automatically.
-

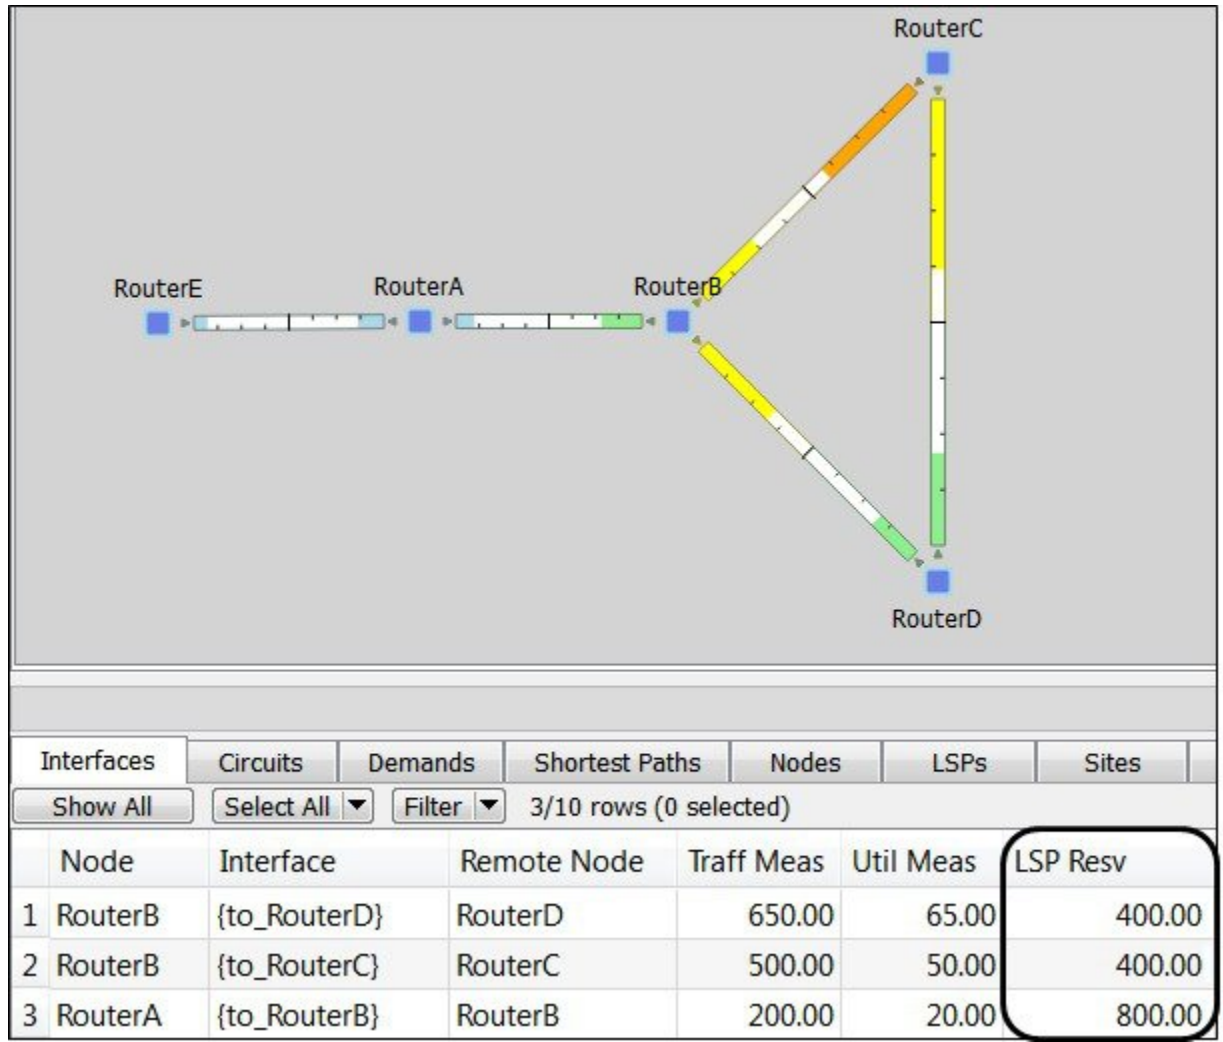
P2MP LSP Bandwidth

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 see 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, see [Viewing LSP Reservations, on page 3](#).

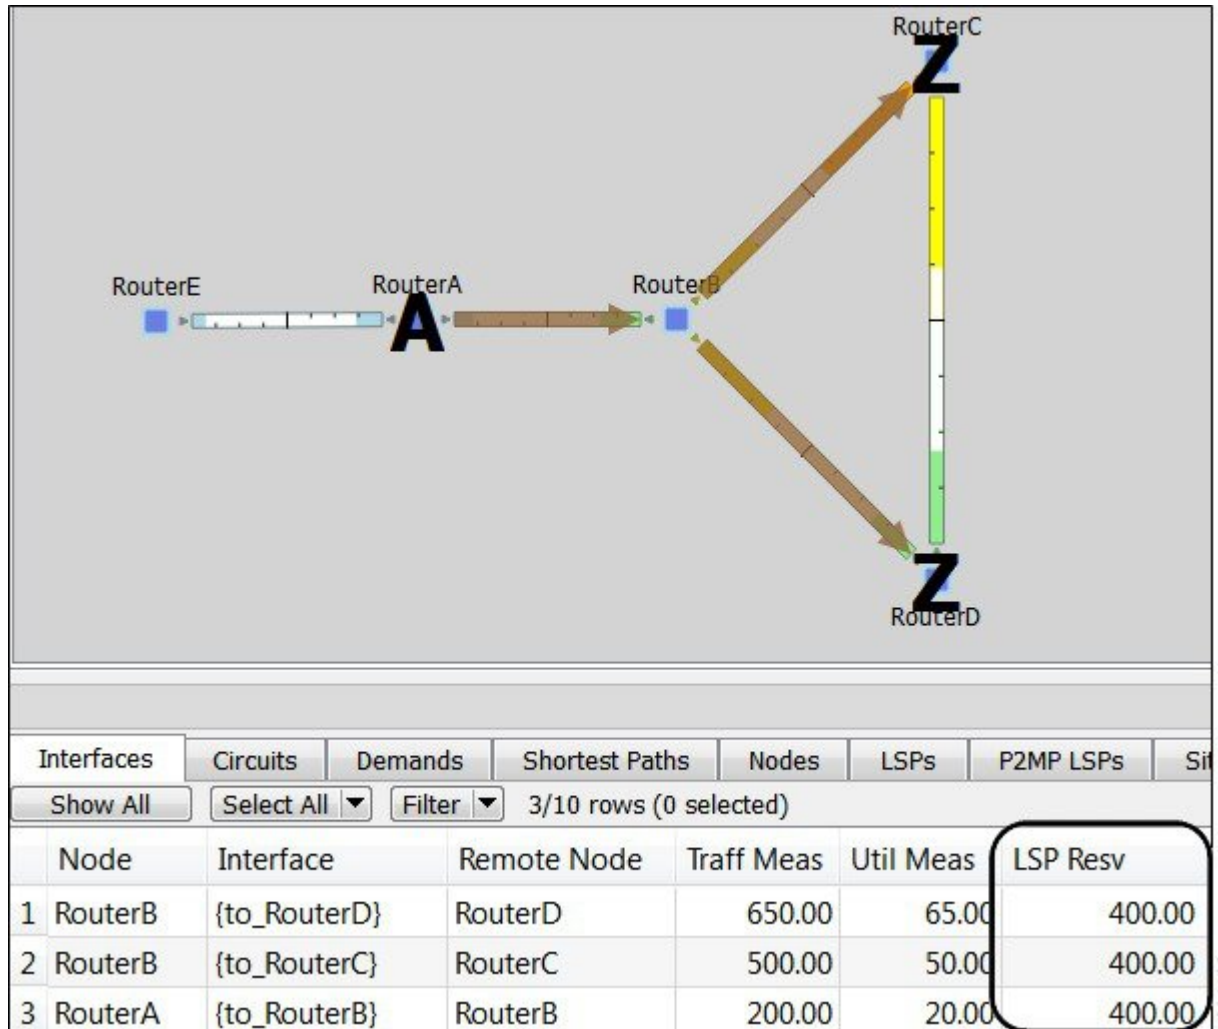
Example: If A-C LSP and A-D LSP each have a bandwidth 400 Mbps, the A-B circuit has an aggregate bandwidth of 800 Mbps ([Figure 19: Bandwidth Behavior Without a P2MP LSP, on page 35](#)). However, if A-C and A-D are sub-LSPs within a P2MP LSP, the A-B circuit has a bandwidth of 400 Mbps ([Figure 20: Bandwidth Behavior with a P2MP LSP, on page 36](#)).

Figure 19: Bandwidth Behavior Without a P2MP LSP



381122

Figure 20: Bandwidth Behavior with a P2MP LSP




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. (See [Creating Demands for P2MP LSPs, on page 41](#) for information on creating these demands.)

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

For complex P2MP LSP demands, filter to the demand, right-click the demand, and choose **Plot Demand**. [Figure 21: Example Demand Plot for P2MP LSP, on page 37](#) shows an example of a demand plot for the P2MP LSP shown in [Figure 18: P2MP LSP with chi as the Source for Three Destinations, on page 33](#).

Figure 21: Example Demand Plot for P2MP LSP



Explicit P2MP LSP Path Initializer

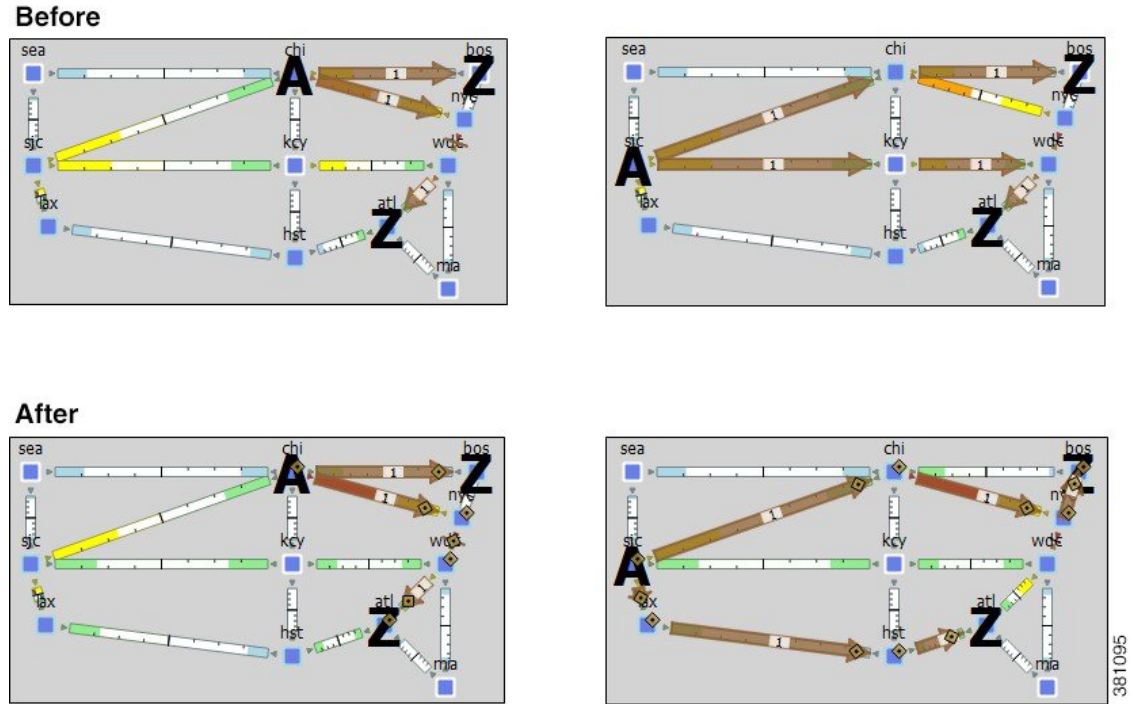
The Explicit P2MP LSP Paths Initializer automates the process of creating explicitly routed sub-LSP paths, as follows:

- Convert a set of dynamically routed sub-LSPs to explicitly routed sub-LSPs along their current paths.
- Create disjoint paths. Two PSMP LSP paths are *disjoint* if they do not route over common objects. These objects are configurable and can be circuits, nodes, sites, SRLGs, or L1 links.
 - Create disjoint paths between sub-LSPs that have the same destination and whose P2MP LSPs belong to the same disjoint P2MP LSP group. This is useful for creating an alternate path if one of the sub-LSP paths becomes unavailable.

Example: [Figure 22: Example Disjoint Paths for P2MP LSPs, on page 38](#) shows two P2MP LSPs (Region A and Region B) that are in the same disjoint group. Their sub-LSPs share two common destinations: bos and atl. Upon choosing this option, the sub-LSPs in Region B are rerouted to arrive at the common destinations using different explicit paths.

- Create disjoint primary and secondary paths for selected sub-LSPs.

Figure 22: Example Disjoint Paths for P2MP LSPs



Note Explicit P2MP LSP Paths Initializer supports Inter-Area functionality.

WAE Design creates named paths and named path hops for the sub-LSP paths. After running the initializer, you can verify the named paths and their hops in the Named Paths and Named Path Hops tables.

- If you choose a disjoint option that creates secondary paths, the secondary named path follows the shortest alternate sub-LSP path that is disjoint from the primary path, according to rules that you specify.
- Orphaned named paths are deleted.
- Explicit paths are named *<Sub-LSP Name> _<Path Option>*.

Prerequisites:


- The plan file must already contain the LSP paths.
- If creating disjoint paths between P2MP LSPs in the same disjoint group, the P2MP LSPs must first be added to the disjoint group in the P2MP LSP Properties dialog box. Here, you can also assign priorities to P2MP LSPs within these groups. Higher priority P2MP LSPs are assigned shorter routes based on the TE metric. The higher the number, the lower the priority.

Step 1 From the P2MP LSPs table, choose the P2MP LSPs to initialize with explicit named paths. If you do not choose any P2MP LSPs, all sub-LSP paths are used for creating the named paths.


Step 2 Choose **Initializers > Explicit P2MP LSP Paths**, or right-click a P2MP LSP and choose **Explicit Path Initializer**.

Step 3 Specify how to create the explicitly routed named paths, and then click **OK**.

- To create named paths for all currently simulated primary or standby sub-LSP paths, click **Follow currently simulated routes**.




- To create disjoint paths, choose whether to create them between disjoint PSMP LSPs that have sub-LSPs with the same destination, or between disjoint primary and secondary sub-LSP paths.



- To modify the disjoint priorities that determine which objects are considered in paths, click **Edit**. Choose **1**, **2**, **3**, or **Ignore**, where the lower the number, the higher the priority. Click **OK**.

Example: If you give circuits a priority of 1, SRLGs a priority of 2, and ignore the other objects, when the disjoint paths are created, WAE Design first creates paths with disjoint circuits. Then, if there are common SRLGs in those paths, it creates paths so the SRLGs are disjoint. The remaining objects are not considered in the disjoint path creating.



Creating 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.

Creating P2MP LSPs

Step 1 Create an empty P2MP LSP.

- Choose **Insert > LSPs > P2MP LSP**, or right-click in the plot and choose **New > LSPs > P2MP LSP**.
- Enter a unique P2MP LSP name. If you do not enter a name, WAE Design creates one using the format `p2mp_lsp_<source_node>`.
- Choose the source site and node. This selection is the root for all sub-LSPs contained in this P2MP LSP.
- (Optional) Enter or choose a disjoint group and disjoint priority for WAE Design to consider when using the Explicit P2MP LSP Path Initializer.

The screenshot shows a configuration dialog for a P2MP LSP. The 'Name' field contains 'East_Region'. Under the 'Source' section, the 'Site' dropdown is set to 'atl' and the 'Node' dropdown is set to 'cr1.atl'. Below this, the 'Disjoint Group' dropdown is set to 'Voice', and the 'Disjoint Priority' field is empty. A vertical number '381141' is printed on the right edge of the dialog box.

Step 2 If you have not done so, create sub-LSPs for the P2MP LSP.

Step 3 Associate the sub-LSPs to the P2MP LSP. You can follow these steps individually, or you can choose multiple sub-LSPs and edit them collectively.

- Double-click the sub-LSPs from the LSP table.
- In the Sub-LSP of P2MP LSP list (at the bottom), choose the P2MP LSP to which you want these LSPs to belong. Then click **OK**.

Creating Sub-LSPs

Step 1 Choose **Insert > LSPs > LSP**, or right-click in the plot and choose **New > LSPs > LSP**.

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). Then click **OK**.

Creating a Mesh of Sub-LSPs

Step 1 For each node that will be the egress node of the P2MP LSP, create the same tag name.

- a) In the Nodes table, choose and then double-click all nodes that will be included in the P2MP LSP.
- b) Click the **Tags** tab, then click **New Tag**.
- c) Enter a name that collectively identifies the sub-LSPs you are creating, and then click **OK**.
- d) Click **OK** to save and exit.

- Step 2** In the Nodes table, choose the node you will be designating as the source for the P2MP LSP.
- Step 3** Choose **Insert > LSPs > LSP Mesh**, or right-click in the plot and choose **New > LSPs > LSP Mesh**.
- Step 4** Uncheck these check boxes:
- **Same Destinations as Sources**
 - **Also create LSPs from destination to source**
- Step 5** In the Source Nodes drop-down list, choose **Selected in table (###)**. All sub-LSPs in a P2MP LSP must have the same source node. A P2MP LSP cannot have multiple source nodes.
- Step 6** In the Destination Nodes drop-down list, choose the newly created tag name.
- Step 7** To set the bandwidth to be the same for all sub-LSPs in this mesh, enter that value in the Fixed Value ___ Mbps field.
- To use the autobandwidth feature so that these LSPs have their Setup BW Sim value automatically calculated and used in autobandwidth simulations, click **Autobandwidth**.
- Step 8** Complete the remaining optional fields as needed, and then click **OK**.

Creating Demands for P2MP LSPs

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

- Step 1** If you want to create demands for P2MP LSPs, choose one or more P2MP LSPs from the P2MP LSP table.
- Step 2** Choose **Insert > LSPs > Demands for P2MP LSPs**, or right-click in the plot and choose **New > LSPs > Demands for P2MP LSPs**.
- a) Choose the P2MP LSPs through which you are routing traffic.
 - b) (Optional) Choose a service class for the newly created demand.
 - c) Choose the desired bandwidth: Minimum sub-LSP Setup BW or Zero. Then click **OK**.

Deleting P2MP LSPs and Sub-LSPs

Deleting a P2MP LSP also deletes all sub-LSPs within it. If you do not want all the sub-LSPs deleted, disassociate them first.

Disassociating Sub-LSPs from Their P2MP LSPs

- Step 1** In the LSPs table, choose one or more sub-LSPs.
- Step 2** Double-click to open the Properties dialog box.
- Step 3** In the Sub-LSP of P2MP LSP list, do one of the following:
- 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**.

Deleting Sub-LSPs or P2MPs

| Delete Sub-LSPs but Keep P2MP LSP | Delete P2MP LSP and Associated Sub-LSPs |
|---|--|
| Right-click one or more sub-LSPs in the LSPs table. Click Delete . Click Yes to continue. | Right-click one or more P2MP LSPs in the P2MP LSP table. Click Delete . Click Yes 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 Configs 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.
- WAE 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 are as follows:

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

WAE tries to resolve as many of these references as possible. Tables and columns are updated as follows:

- 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.

WAE Design 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.

Resolving Plan File to Objects

If plan files have configurations that are not matched or refer to objects outside the modeled network, there are several ways to resolve these differences:

- Resolve LSP destinations, hops in named paths, hops in actual paths, or segment list hops by choosing **Initializers > Resolve Plan**.
- Resolve named path hops in the named paths Properties dialog box.
- Resolve named path hops in the named path hops Properties dialog box.

