



MPLS Simulation

This chapter describes how WAE Design simulates MPLS routing. An MPLS simulation, for a particular failure scenario and a particular traffic level, performs the following steps.

- LSPs are established under normal operation: that is, with failures not taken into account.
- Any LSPs that are affected by the failures are rerouted. Depending on the LSP path settings, reroutes might involve moving to a secondary path, dynamically rerouting the LSPs, or rerouting based on a segment list.
- Demands are routed using the established LSPs using the specified IGP protocols given the specified failure scenarios.
- LSP utilizations are calculated from the demand traffic using the specified traffic level.

Creating and Visualizing LSPs

To simulate MPLS LSPs in the network, you must first set up the LSPs to be routed. In the WAE Design GUI, enter LSPs into the plan file in the following ways.

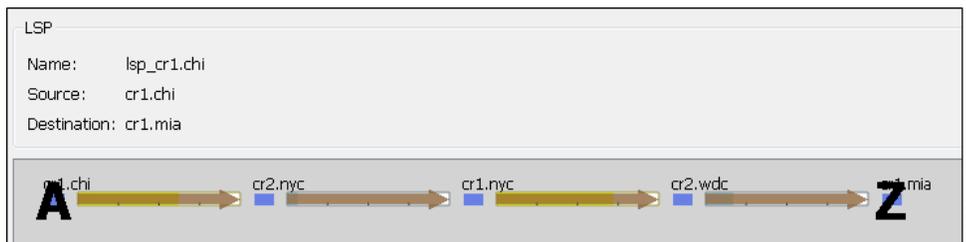


- By selecting the Insert->LSPs->LSP menu, or by right-clicking in the plot and then selecting the New->LSPs->LSP from the context menu. This opens a dialog box for entering properties for a single LSP. One such property is the Type, which determines whether this is an RSVP LSP or an SR LSP.
- By selecting the Insert->LSPs->LSP Mesh menu or by double-clicking in the plot and then selecting the New->LSPs->LSP Mesh menu. This opens a dialog box for adding a mesh of LSPs between the selected nodes. Again, the Type property determines whether the LSP is an RSVP or an SR LSP.

When selected, LSPs appear in the plot as a brown arrow. To view the details of an LSP route, open the LSP plot ([Figure 9-1](#)), select an LSP from the LSPs table. Then right-click and select Plot LSP.

To filter to related information, such as related interfaces, source and destination nodes, or demands, right-click one or more LSPs and select the option from the context menu.

Figure 9-1 LSP Plot



Routing RSVP LSPs



In this section, unless otherwise noted, *LSPs* refer to RSVP LSPs.

Dynamic LSP Routing and CSPF

If an LSP contains no LSP paths¹, 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. These and other properties affecting LSP routes are defined in the LSP and interface Properties dialog boxes.

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.



- Setup Bandwidth—The amount of traffic the source node is requesting 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. For information, see the [Autobandwidth Simulations](#) section.
- 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.

1. These are often called *MPLS TE tunnel paths*.

- 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 selected, the LSP is routed using IGP metrics. If unselected (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.

- 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 selected, the value is set to T (True), and TE metrics are used in routing the LSP provided the LSP's TE Metric Disabled is unselected (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 the [Ignore Reservable Bandwidths for Capacity Planning](#) section.

- Affinities—Lists which affinities are set. For more information, see the [Affinities](#) section.
- 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). For more information, see the [Loadsharing](#) section. For information on optimizing Loadshare values, see the [LSP Optimization](#) chapter.

- 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 the [Fast Reroute Simulations](#) section.

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.

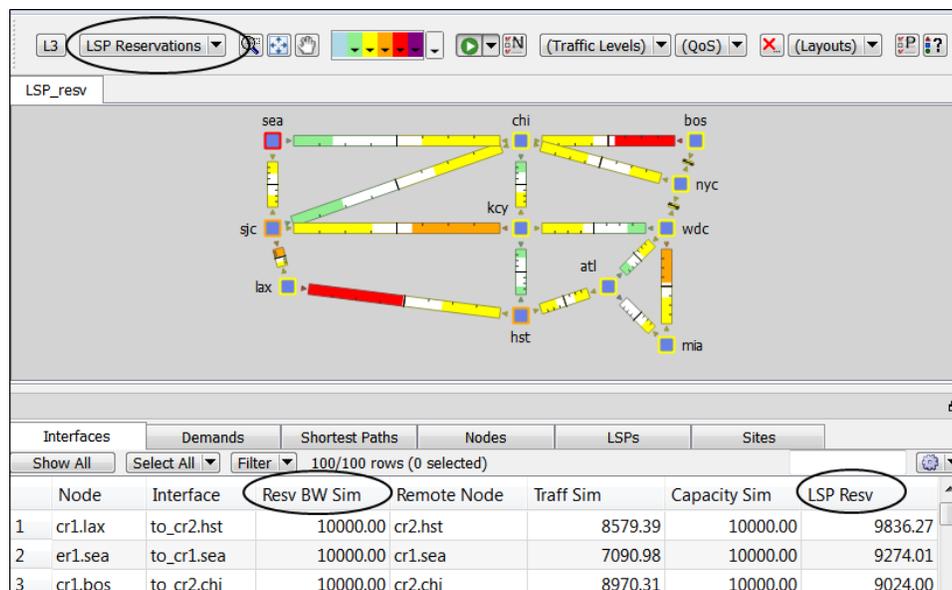
- Traff Sim LSP—The total amount of LSP traffic on the interface.
- Traff Sim Non LSP—The total amount of non-LSP traffic on the interface.

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 select Filter to Demands Through Interfaces. Then select either In LSP or Not in LSP.

View LSP Reservations

To view the LSP reservations on the interfaces, select LSP Reservations from the Network Plot View menu in the Visualization toolbar (Figure 9-2). The width of the circuit displayed in the network plot is set according to the Resv BW Sim column in the Interfaces table, rather than the Capacity column as in other views. 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 9-2 LSP Reservations View



LSP Setup BW Initializer

The LSP Setup BW Initializer calculates the setup bandwidth for existing RSVP 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 the [Autobandwidth Simulations](#) section.

Initialize Setup BW from Traffic Levels

Traffic Levels

Default

Peak

Setup BW is set to the maximum amount of traffic passing through each tunnel across selected Traffic Levels.

Initialize Load Share from Setup BW
LSP Load Share values are set equal to Setup BW

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

LSP Paths

LSPs can be assigned one or more LSP paths, each with a specified Path Option property. In this case, the LSP is routed using the first LSP path that can successfully be established. LSP paths are established in increasing order of their path option.

LSP Name: cr1.lon_cr2.fra

Path Option: 2

Alternatively, you can select which LSP path to use from the LSP Active Path drop-down list in the LSP Property dialog box.

LSP Property

Active Path: Not set

Not set

1

2

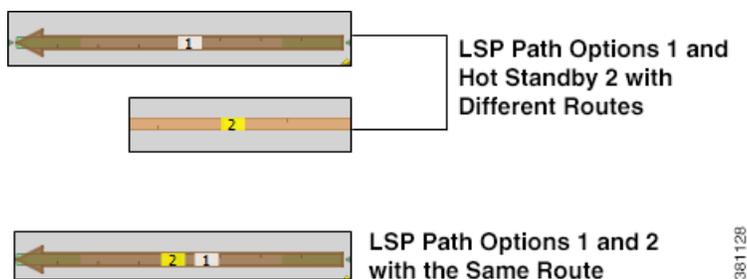
3

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

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. For more information, see the [Autobandwidth Simulations](#) section.

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. For information on named paths, see the [Named Paths and Explicit LSP Routing](#) section.

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



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, respectively.

LSP Paths Example

Figure 9-3 shows an LSP (cr1.lon_cr2.fra) that uses four LSP paths.

Figure 9-3 Example 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 | | | |

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

Following is how this example LSP responds to failures.

- Step 1** If the primary path (cr1.lon_cr2.fra_100) fails, the secondary path (cr1.lon_cr2.fra_200) is used.
- Step 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 not a standby.
- Step 3** If all these 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.

Create LSP Paths

- Step 1** In the LSPs table, select the LSP to which you are adding LSP paths.

Create LSP Paths for the selected (1/1056) LSPs

Path Option: 2

Setup Bandwidth: 0 Inherit from LSP

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

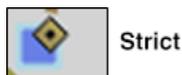
Standby

- Step 2** Either right-click in the plot and select New->LSP Paths or select the Insert->New LSP Paths menu. The Insert LSP Paths dialog box appears.
- Step 3** In the Path Option field, enter the order in which the LSP path is activated. Lower numbers have preference over higher numbers.
- Step 4** To set the bandwidth, specify the Setup Bandwidth for the LSP path, or select that it should inherit the bandwidth from the LSP. As an alternative, for information on dynamically setting bandwidth for LSPs, see the [Autobandwidth Simulations](#) section.
- Step 5** To create associated named paths, click this option and complete the name using \$1 for the LSP name and \$2 for the path option. For information on named paths, see the [Named Paths and Explicit LSP Routing](#) section.
- Step 6** If this is a standby LSP path, click Standby, which ensures the paths are always active.
- Step 7** To associate the LSP path with affinities, click Edit in the Affinities section. For example, you might want to associate a standby path with an affinity that forces it on a different path than the primary path. LSP paths can inherit the affinity from the LSP or can be assigned a specific affinity. For more information about affinities, see the [Affinities](#) section.
- Step 8** Click OK.
-

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.



The named paths and their hops are listed and selectable from the Named Paths and Named Path Hops tables, respectively. The advantage of having separate tables for these 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 9-4 extends the earlier example (Figure 9-3), 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 9-4 Example Named Paths the cr1.lon_cr2.fra LSP

| Interfaces | | LSPs | | LSP Paths | | Named Paths | |
|------------|---------------------|------------|--------|-----------|--|-----------------------|--|
| Show All | | Select All | | Filter | | 2/2 rows (0 selected) | |
| | 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 9-5 shows the named path hops for the cr1.lon_cr2.fra_100 named path. The Steps column shows the order of the hops.

- 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 9-5 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 | | | | |

Create 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)

- Select the “Create associated named paths” option when creating LSP paths. Using this option does not create hops. See the [Create LSP Paths](#) section.
- Use the Explicit LSP Path Initializer. Using this initializer creates named paths with strict hops. See the [Explicit LSP Path Initializer](#) section.

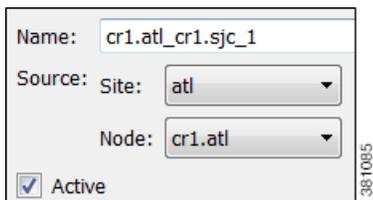
- Use the Explicit LSP Optimization and Tactical Explicit LSP Optimization tools, which minimize congestion by optimizing the placement of primary and secondary LSP paths. See the [LSP Optimization](#) chapter.

Edit Named Paths and Named Path Hops

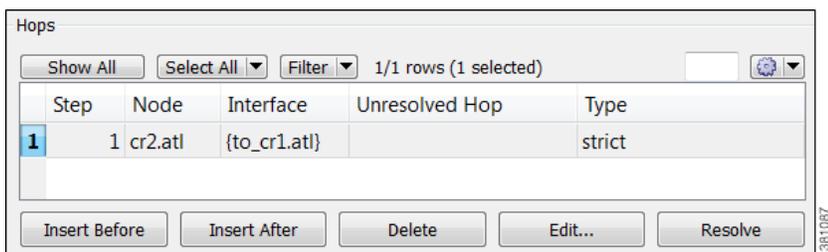
Once the named paths are created, you can create, edit, or delete their named path hops. Note that discovered named paths can contain *unresolved* hops, which are nodes and interfaces that are not in the plan file. For information on resolving named paths, see the [Unresolved LSP Destinations and Hops](#) section.

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, either double-click a named path or right-click one and select Properties from the context menu. The Properties dialog box appears with a list of all named path hops.

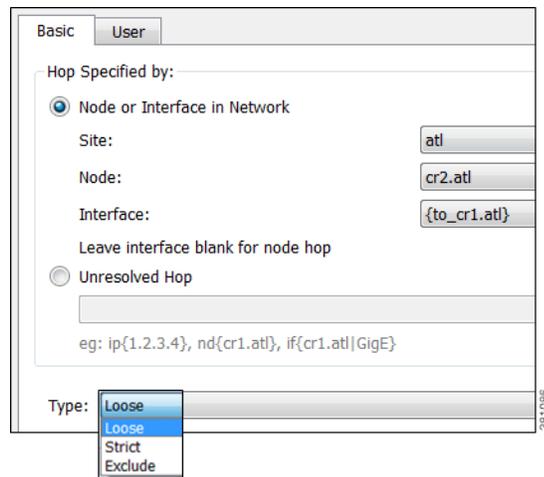


- Step 2** Named path options
- To change the name, enter it in the Name field.
 - To change the source site or source nodes, select them from the drop-down lists.
 - To activate or de-activate, select the Active option. For information on active and inactive states, see the [Simulation](#) chapter.



- Step 3** Hop options
- To add a new hop, select an existing hop (if applicable) from the Hops list, and then click either the Insert Before or Insert After button to determine where the new hop goes sequentially. Go to step 4.
 - To edit an existing hop, select it from the Hops list and then select the Edit button. This opens the Properties dialog box for named path hops. Go to step 4.
 - To delete an existing hop, select it from the Hops list, select the Delete button, and then press OK. You need not continue.

- To resolve nodes and interfaces that do not exist in the plan file, click the Resolve button. For information, see the [Resolve Plan File to Objects](#) section.
- Step 4** To continue creating or editing the named path hop, use the following options and then click OK in the named path hops Properties dialog box.
- As needed, select the site, node, and interface. For node hops, do not select an interface.
 - Select the named path hop type as Loose, Strict, or Exclude. See the [Named Paths and Explicit LSP Routing](#) section for a description of each.

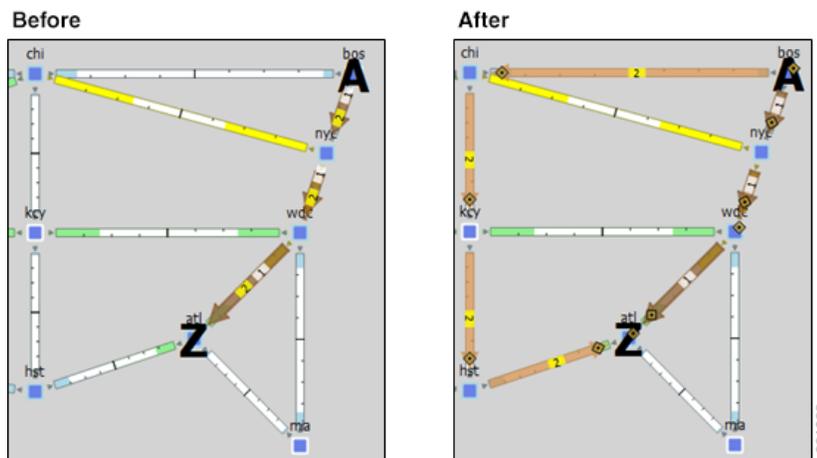


Explicit LSP Path Initializer

The Explicit LSP Paths Initializer automates the process of creating explicitly routed RSVP LSP paths, as follows.

- Convert a set of dynamically routed LSPs to explicitly routed LSPs along their current paths.
- Create disjoint paths. Two 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 primary paths for a group of LSPs, which creates a *disjoint group*.
 - Create disjoint primary and secondary paths for selected LSPs. For example, [Figure 9-6](#) shows that the original primary and secondary LSP paths use the same route. Using this option, these paths are explicitly routed to use different circuits.

Figure 9-6 Example Disjoint Primary and Secondary Explicit Paths



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

- If you choose a disjoint option that creates secondary paths, the secondary named path follows the shortest alternate LSP path that is disjoint from the primary path, according to rules that you specify.
- Orphaned named paths are deleted. Named paths are *orphans* if they are no longer used by any of the LSP paths.
- The explicit paths are named <LSP Name>_<Path Option>.

The resulting report summarizes the number of LSPs and their paths. Depending on the disjoint option selected, the report might also list the number of disjoint groups, LSP disjointness, and path disjointness, as well as the common objects amongst them. You can filter to related information from the Disjointness and Path Disjointness sections.

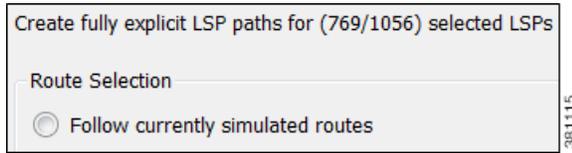
Prerequisites

- The plan file must already contain the primary and secondary LSP paths. See the [Create LSP Paths](#) section.
- If creating disjoint paths between LSPs in the same disjoint group, the LSPs must first be added to the disjoint groups in the LSP Properties dialog box. Here, you can also assign priorities to LSPs within these groups. Higher priority LSPs are assigned shorter routes based on the TE metric. The higher the number, the lower the priority.

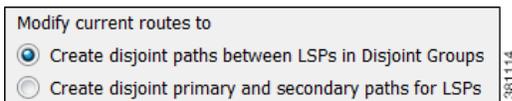
| | |
|--------------------|-----|
| Disjoint Group: | GBD |
| Disjoint Priority: | 1 |

-
- Step 1** From the LSPs table, select the LSPs to initialize with explicit named paths. If you do not select any LSPs, all LSP paths are used for creating the named paths.
 - Step 2** Select the Initializers->LSP Explicit Paths menu, or right-click a selected LSP and select Explicit Path Initializer from the context menu. The LSP Explicit Path Initializer dialog box appears.
 - Step 3** Select how to create the explicitly routed named paths.

- To create named paths for all currently simulated primary or standby LSP paths, select the “Follow currently simulated routes” option.

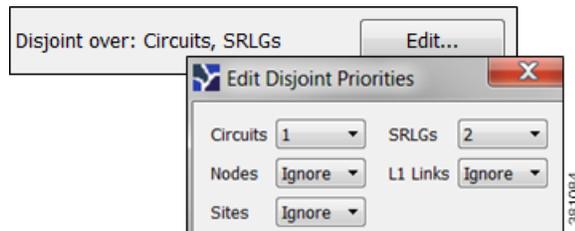


- To modify existing routes to use disjoint paths, select whether to create the disjoint paths between LSPs in the same disjoint group or between primary and secondary LSP paths.



- To modify the disjoint priorities that determine which objects are considered in the paths, click Edit. A dialog box appears enabling you to specify priorities as 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, then when the disjoint paths are created, WAE Design first creates paths with disjoint circuits. Then, if there are common SRLGs in those paths, it further creates paths such that the SRLGs are disjoint. The remaining objects are not considered in the disjoint path creating.



Step 4 Click OK.

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, using WAE Collector network discovery tools. RSVP 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.

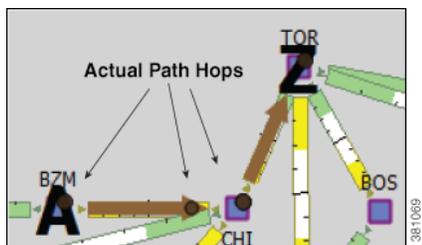
When an LSP path is routed, the actual route is attempted first. If this routing fails due to CSPF constraints, such as not enough reservable bandwidth or affinity restrictions, WAE Design reverts to the standard CSPF routing algorithm.

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. Since 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 select Delete Actual Paths from the context menu.

Note that 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, select the View->LSP->Actual Paths menu; all hops are displayed when the LSP is selected.

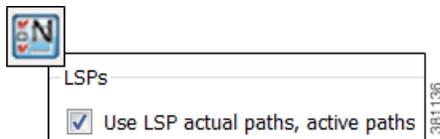


Two columns in the LSP table are useful to see the result of actual paths on LSP routing.

| 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. |
| 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. |

De-Activate 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** Either select the Edit->Network Options menu or click the Network Options icon in the Visualization toolbar.
- Step 2** Select the Simulation tab.
- Step 3** To use or disregard actual paths and active paths in MPLS simulation, select or deselect the “Use LSP actual paths, active paths” option accordingly. Then 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 32 distinct affinities, each of which is defined by a number and an optional name. These affinities always exist even if none are defined in the plan file. Named affinities are sometimes called *administrative groups* or *link coloring*.

To use affinities, perform the following tasks.

-
- Step 1** Optional: assign names to affinities.
- Step 2** Assign affinities to circuits.
- Step 3** Associate LSPs with affinities.
-

Assign Names to Affinities

Often affinities have descriptive names that suggest their purpose.

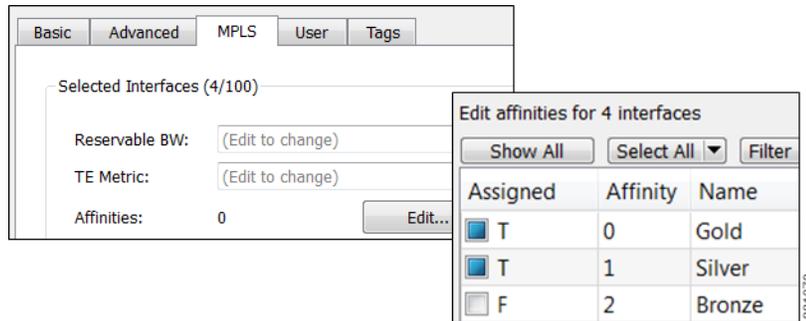


-
- Step 1** Select the Edit->Admin Groups menu; the Edit Admin Groups dialog box appears.
- Step 2** In the Name field, enter a unique, alphabetic name for each affinity.
- Step 3** Click OK.
-

Assign 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.

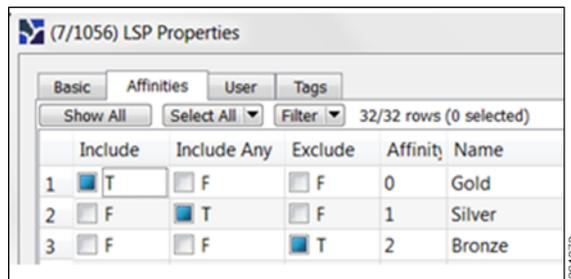
You assign affinities to interfaces. If you assign an affinity to a circuit, you are assigning it to all of its associated interfaces.



-
- Step 1** From either the Interfaces or Circuits table, select one or more interfaces or circuits to which you are assigning the affinities.
- Step 2** Either double-click the selections or select the Edit->Properties menu; the Properties dialog box appears.
- Step 3** Select the MPLS tab.
- Click the Affinities Edit button; the Edit Affinities dialog box is displayed.
 - In the Assigned column, select 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 to exit the Properties dialog box. Note, if you close without clicking OK again, the changes are not saved.
-

Associate 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. The include option has precedence over the exclude option, so if both are specified the include option prevails. 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](#).)



-
- Step 1** From either the LSPs table, select one or more LSPs to which you are associating the affinities.
- Step 2** Either double-click the selections or select the Edit->Properties menu; the Properties dialog box appears.
- Step 3** Click the Affinities tab.
- Step 4** Check one of the three boxes for each affinity you are associating with the selected LSPs.
- 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.
- Step 5** Click OK.
-

Associate LSP Mesh Groups with Affinities

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

-
- Step 1** Select the Insert->LSP Mesh menu. The LSP Mesh Creator dialog box appears.
- Step 2** Click the Affinities Edit button. An Edit Affinities dialog box appears.
- a. Check one of the three boxes for each affinity you are associating with the LSPs in the mesh group.
 - 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.
 - b. Click OK.
- Step 3** Click OK to exit the LSP Mesh Creator dialog box. Note, if you close without clicking OK again, the changes are not saved.
-

Routing Inter-Area LSPs

In WAE Design, all nodes in a single AS are assumed to belong to a single IGP. If the plan file contains more than one AS, all IGPs defined in these AS's are of the same type. For information on how nodes are assigned to areas or levels, see the [IGP Simulation](#) chapter.

An inter-area LSP is an LSP whose source node and destination node have no areas in common. If available, inter-area LSPs follow actual paths, regardless of whether doing so violates the required order of routing through areas. For example, if following an actual path, an inter-area can enter and leave an OSPF area 0 more than once.

Other factors that determine how the inter-area LSPs are routed include whether the LSP type is RSVP-TE or SR, and whether you have selected to require explicit hops at ABRs². (See the [Explicit vs. Dynamic Inter-Area LSP Routing](#) section.)


Note

Inter-area Fast Reroute LSPs and inter-area IGP shortcut LSPs are not supported. If the source and destination nodes of a Fast Route LSP are in different areas, the LSP is not routed.

Order of Routing through Areas

Regardless of the LSP type and whether explicit hops are required, inter-area LSPs route through the backbone areas, as follows, where “backbone” means area 0 for OSPF or the Level 2 area for IS-IS.

- If there are three or more areas, backbone areas must be between the source and destination nodes. Typically, there are no more than three areas and therefore, the backbone area must be in the middle. For example, an OSPF inter-area LSP would route from area 1 to area 0 (backbone) and then to area 2 ([Figure 9-7](#)).
- If there are only two areas, there can only be one backbone area, and either the source or the destination node must be in the backbone area.

Explicit vs. Dynamic Inter-Area LSP Routing

There are two modes of routing inter-area LSPs. One requires that explicit hops be set on ABR nodes. This mode correctly simulates actual router behavior where ABR explicit hops are required. The other mode does not require explicit hops at ABR nodes and can route an LSP fully dynamically across multiple areas. While this mode does not simulate actual router behavior, it is useful for planning inter-area LSP routes. For instance, you can use a fully dynamic simulation to find an appropriate path, which can then be converted into a (routable) fully explicit path using other tools such as the Explicit LSP Paths initializer. These modes are specified using the network option labeled “LSP routing requires ABR explicit hops.”

If this option is selected, inter-area LSPs are routed based on explicit hops set on the ABR nodes.

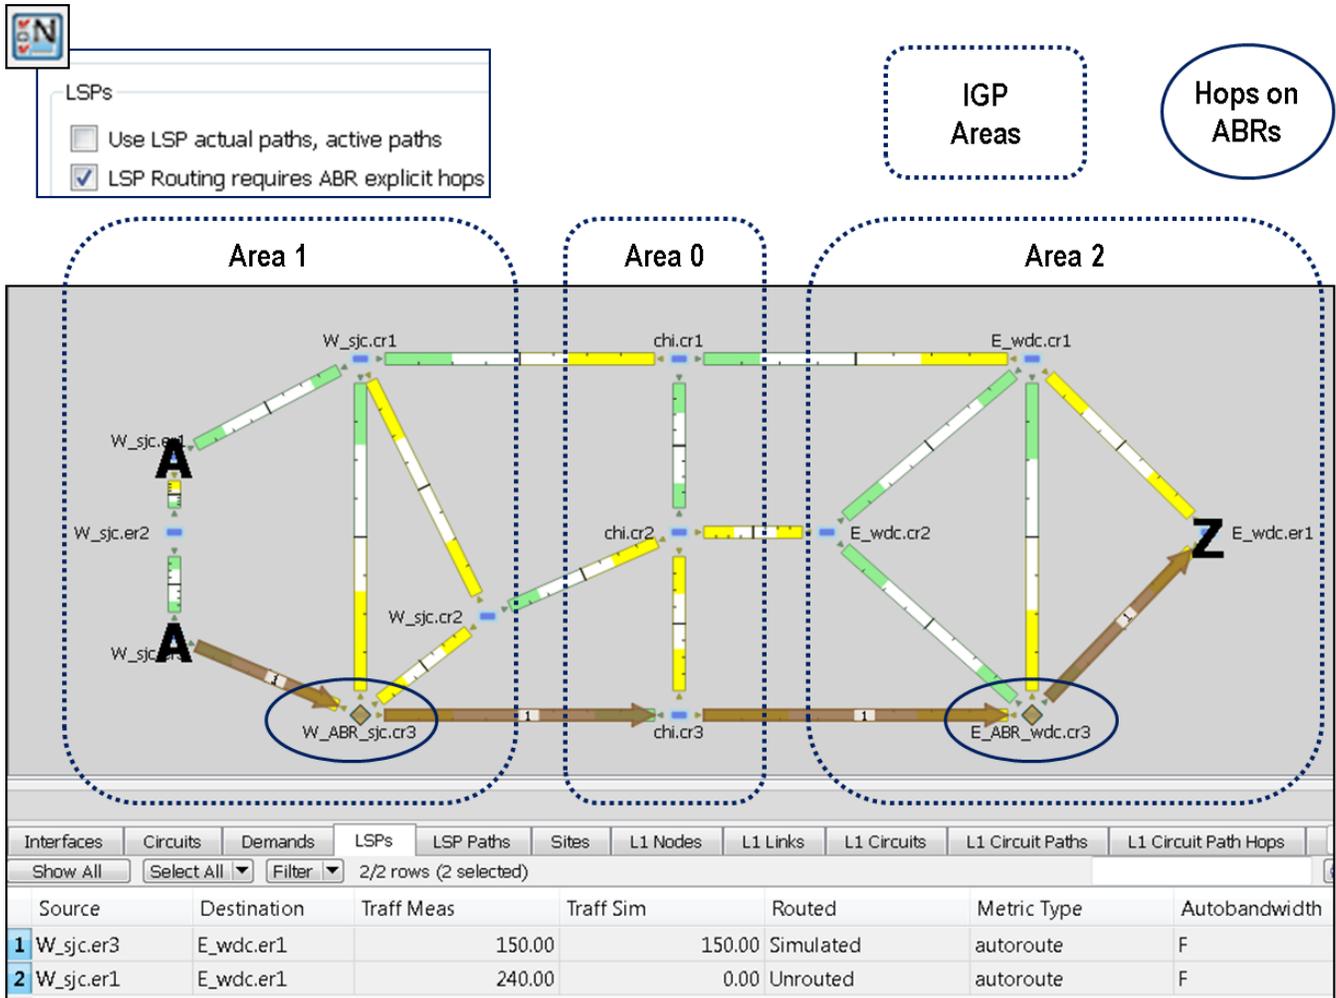
- An inter-area RSVP-TE LSP must contain a named path, and the named path must contain explicit hops at ABRs for each required area crossing.
- An inter-area SR LSP must contain a segment list, and the segment list must contain explicit node hops at ABRs for each required area crossing.

If this option is not selected, inter-area LSPs are routed dynamically and explicit hops at ABRs are not required. To leave one area and enter another, the inter-area LSP routes to the closest ABR in the current area that also borders the area it is entering.

Example: [Figure 9-7](#) shows two inter-area LSPs selected, although only the LSP with explicit hops set on ABR nodes routes. The reason is because the network option is set to require ABR hops when routing inter-area LSPs. [Figure 9-7](#) shows that by deselecting this network option, the inter-area LSP that does not have hops set at ABRs can dynamically route.

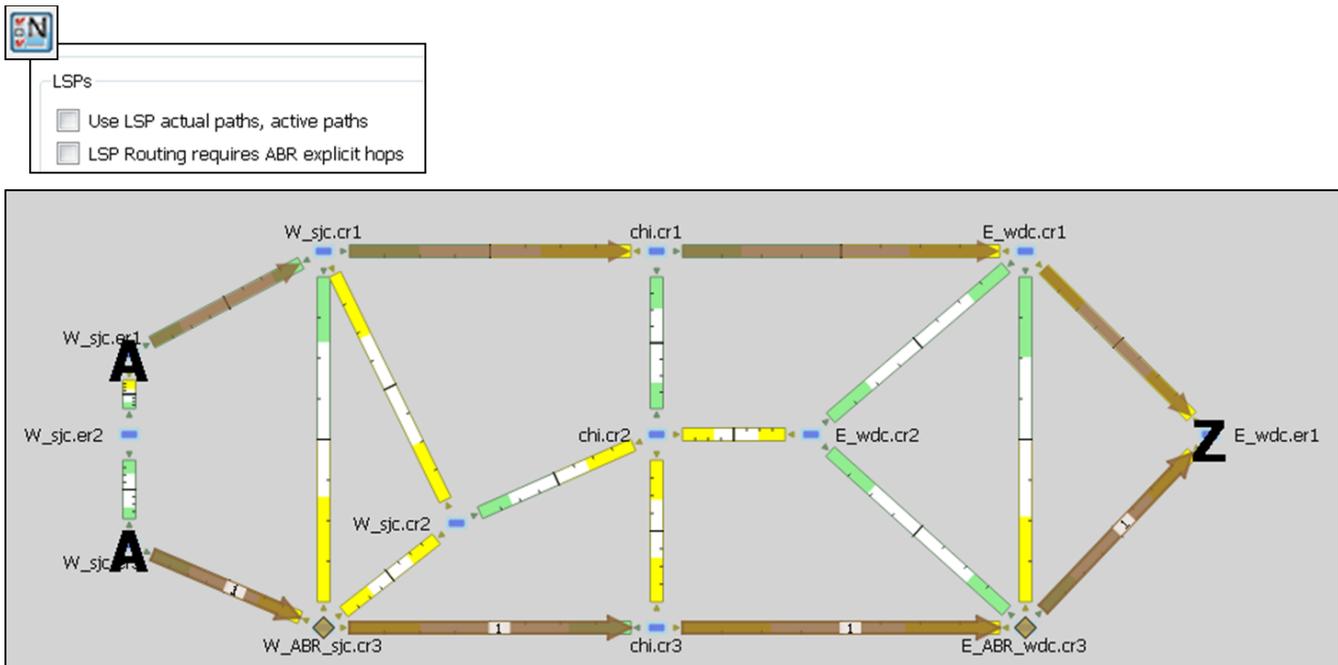
2. An OSPF Area Border Router (ABR) is a node belonging to both area 0 and other OSPF areas. An IS-IS ABR is a node belonging to both the Level 2 area and another IS-IS area.

Figure 9-7 Example Inter-Area Routing Requiring Hops at ABRs



404548

Figure 9-8 Example Inter-Area Routing without Requiring ABRs



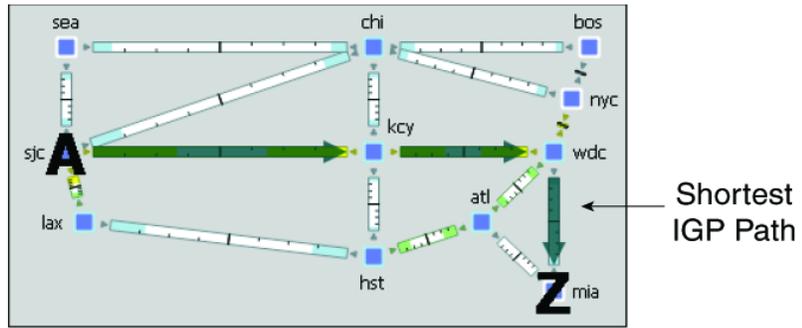
LSP Shortest Paths

To view the LSP shortest paths, use one of these methods. The source is marked with an “A” and the destination is marked with a “Z”.

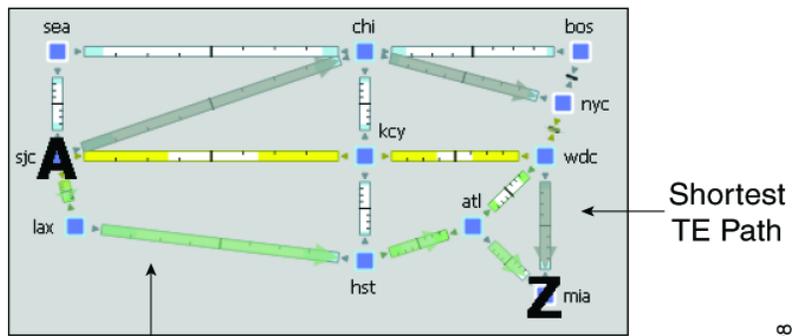
- Using the View->LSPs sub-menus, you can choose to see the shortest actual paths, shortest latency paths, or shortest TE paths each time that you select an LSP.
- To view all shortest IGP paths, shortest TE paths, and/or shortest latency paths for selected LSPs, use the context menu. Select one or more LSPs, right-click, and select Filter To->Shortest Paths (Figure 9-9).

Figure 9-9 LSP Shortest IGP, Shortest TE, and Shortest Latency Paths

| Shortest Paths | | | |
|----------------|---------|-------------|---------|
| | Source | Destination | Type |
| 1 | er1.sjc | er1.mia | IGP |
| 2 | er1.sjc | er1.mia | TE |
| 3 | er1.sjc | cr1.mia | Latency |



| Shortest Paths | | | |
|----------------|---------|-------------|---------|
| | Source | Destination | Type |
| 1 | er1.sjc | er1.mia | IGP |
| 2 | er1.sjc | er1.mia | TE |
| 3 | er1.sjc | cr1.mia | Latency |



Shortest Latency Route

380898

Path Latency Calculations

WAE Design calculates shortest latency paths for LSPs and demands as follows.

- By default, WAE Design uses the shortest path from the source to the destination of the LSP or demand, where the weight on each interface is the delay value for that interface.
- If the delay for any interface is zero (as it is by default), a small (0.00001 ms) delay is used instead. This means that if all the delays for the interfaces in the plan are zero, the path with the smallest number of hops is selected.
- In cases where two different paths have the same latency, paths with the smaller number of hops are preferred.

Routing Demands

Intra-Area LSPs

Intra-area LSPs are modeled as IGP shortcuts. That is, the source of the demand using the LSP can be a node other than the ingress node into the IGP, and the destination node of the LSP can be a node other than the egress node from the IGP. A demand through intra-area LSPs does not require the LSP to traverse the full demand path.

Each intra-area LSP has a metric that helps determine which traffic is routed through the LSP. By default, autoroute LSPs have a metric equal to the shortest IGP distance from the source to the destination. However, you can configure static metrics for these LSPs that override this default. Static metrics can also be defined relative the shortest IGP distance, but these relative metrics are not currently supported in WAE Design.

Forwarding adjacency LSPs always have a metric. If not specified, the default is 10. Forwarding adjacency LSP metrics are injected into the IGP so that nodes other than the source node of the LSP are aware of the path length through the forwarding adjacency LSP, and use it in their shortest path calculations.

To edit autoroute and forwarding adjacencies, select one or more LSPs from the LSP table; double-click to bring up the Properties dialog box and edit the values in the Routing area.

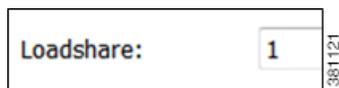
Inter-Area LSPs

Since it is not possible to define metrics distances across ISP areas, there can be no well-defined metric for inter-area LSPs. Therefore, in WAE Design, a demand only routes through an inter-area LSP if the demand's endpoints are nodes matching the source and destination of the LSP, interfaces on these nodes, or external AS's whose ingress and egress points through the IGP are these nodes. This requirement is true regardless of the network option settings.

To be routed, these demands must also match the privacy, tags, and class-based forwarding requirements. Autoroute, forwarding adjacency properties, and LSP metrics are ignored. For information on privacy, see the [Private LSPs](#) section, for information on tags, see the [Plan Objects](#) chapter, and for information on class-based forwarding, see the [Class-Based Forwarding](#) section.

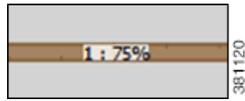
Loadsharing

Two or more LSPs with the same source and destination (and metrics, if these are defined) loadshare traffic between them. How the load is shared between the LSPs is determined by the LSP Loadshare property. By default, LSPs have a Loadshare property of 1, and thus route traffic between them in equal proportions. Changing the Loadshare value changes the distribution of LSP traffic and interface traffic in proportion to these values.



Example: If two LSPs are parallel, and one has a Loadshare property of 2 and one has a Loadshare property of 1, there will be a 2-to-1 ratio of traffic shared between them. The top half of [Figure 9-10](#) shows an example of two parallel LSPs that are routed using strict explicit paths. Each has a Loadshare value of 1, which means the traffic is routed using a 1:1 loadshare ratio so that each LSP carries 50% of the traffic. In contrast, the lower half shows the same parallel LSPs with a 2:1 ratio. That is, one LSP has a Loadshare value of 2, and one has a Loadshare property value of 1. The LSP with a Loadshare value of 2 carries 67% of the traffic, while the other carries 33%.

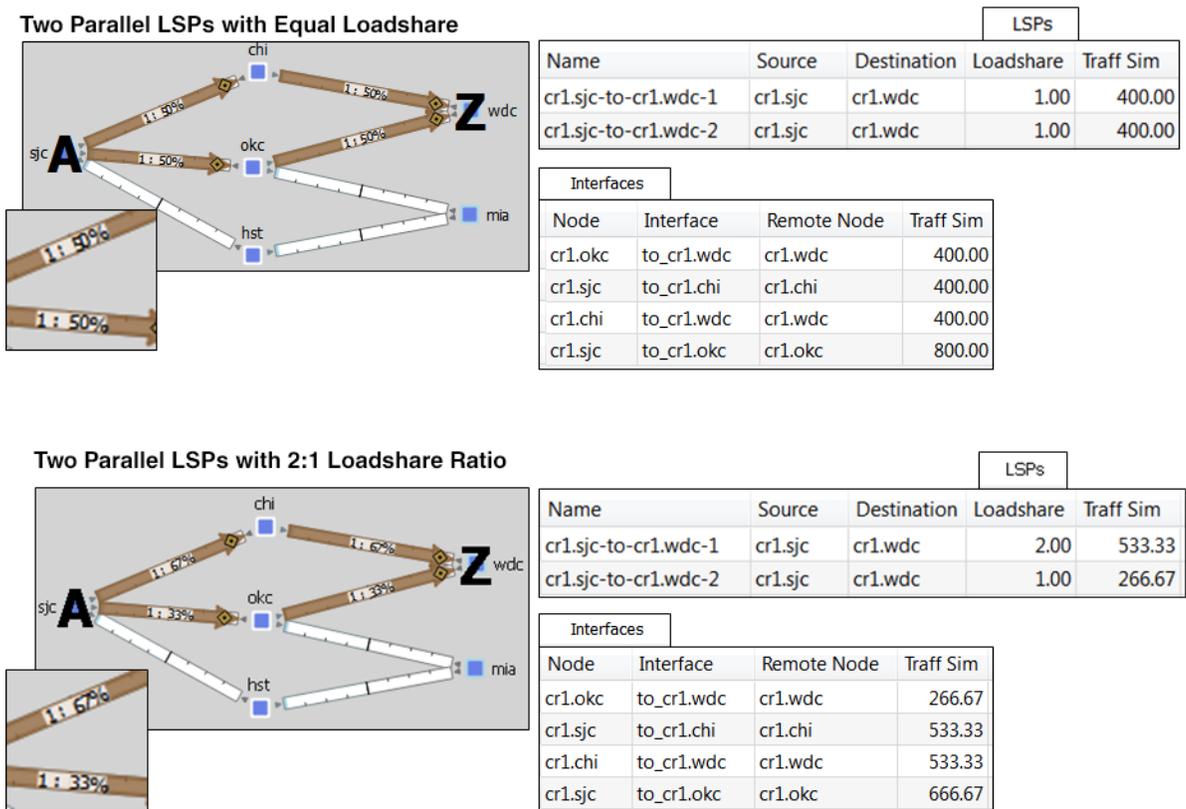
Each LSP has a path showing the percentage of traffic that crosses the LSP as a result of the Loadshare value for that LSP. The format is <path option>:<loadshare percentage>%. Although loadshare applies only to path option 1, the LSP itself might be routed along a different path.



Note that this percentage may or may not be the same as the Loadshare value in the LSPs table. This is because the Loadshare value is relative to the other Loadshare values for parallel LSPs. For this reason, if manually editing the Loadshare property, it is recommended that you consider the entire set of parallel LSPs. Note also that since an LSP can be in multiple sets of parallel LSPs, the percentage on its path can differ.

To optimize these Loadshare values, use the LSP Loadshare Optimization tool. For information, see the [LSP Optimization](#) chapter.

Figure 9-10 Example of Two Parallel LSPs with Equal Loadsharing



381100

Class-Based Forwarding

Class-based forwarding (CBF) enables you to manage differentiated service requirements by restricting or allowing demands into LSPs according to a set of user-defined mappings between demand service classes and LSPs. For instance, a demand with an Assured Forwarding service class could be mapped to an LSP saved for carrying low-latency voice traffic, while a different LSP using that same demand could be used for only best-effort services, such as file transfers.

Class-Based Forwarding Simulations

Demands can be assigned a Service Class property. LSPs can be assigned a Class property. CBF is enabled by creating a table that maps the demand service class to the LSP classes. The resulting table defines a set of LSPs to which a demand with a specific service class has access.

Figure 9-11 Example Service Class LSP Class Mapping Table

| | Service Class | Priority | LSP Class |
|---|---------------|----------|-------------|
| 1 | AF | 1 | LSPClass_AF |
| 2 | AF | 2 | LSPClass_BE |
| 3 | AF | 3 | Drop |
| 4 | BE | 1 | LSPClass_BE |

- Each demand service class in the table is mapped to one or more LSP classes, or with a Drop value.
- If a demand’s service class is in the mapping table, the demand can only use LSPs with LSP classes that match that service class.

If a demand service class is not in the mapping table, it can use any LSP.

Example: There are three demand service classes: AF, BE, and EX. The mapping table contains only AF and BE (Figure 9-11). This means that the demands with EX service class can use any LSP in the plan file.

- If an LSP class does not appear in this mapping table, then any demand can use any LSP of that class.
- If multiple LSPs are sourced from one node, demands routed through that node can use any LSP to which the demand’s service class is mapped.
 - If all available LSPs have LSP classes with equal priority, the traffic is load balanced between these LSPs.
 - If the LSPs have prioritized LSP classes (1, 2, 3, etc.), those with the highest priority (lowest number) are routed first.

Example: A demand with an AF service class must first use LSPs with an LSPClass_AF because within the AF service class, this is the only Priority 1 LSP class. If no LSPs with LSPClass_AF are available, the demand will attempt to route across LSPs with an LSPClass_BE class (Figure 9-11).

- If there are no available LSPs with matching LSP classes, the demand follows these rules.
 - If the demand’s service class is not mapped to a Drop value, the demand forgoes using LSPs and takes the shortest IGP path.
 - If the demand’s service class is mapped to a Drop value, the demand is unrouted.

Example: If there are no available LSPs for the AF demand to route across, it will not route due to the Drop value associated with it (Figure 9-11). If the BE demand cannot route across an LSP that has a class of LSPClass_BE, it will route across the shortest IGP path available.

Example Simulations

In these examples, we want to route the assured-forwarding (AF) Europe-to-Japan traffic directly over the more expensive, but lower latency, circuits through Russia. We want to route the best-effort (BE) traffic along the cheaper, but higher latency, circuits through the US.

| Demands | | | |
|------------|---------|-------------|---------------|
| Name | Source | Destination | Service Class |
| Lon-Tok-D1 | cr1.lon | cr1.tok | AF |
| Lon-Tok-D2 | cr1.lon | cr1.tok | BE |

| LSPs | | | |
|-------|---------|-------------|-------------|
| Name | Source | Destination | Class |
| LSP_A | cr1.lon | cr1.tok | LSPClass_AF |
| LSP_B | cr1.lon | cr1.tok | LSPClass_BE |

381094

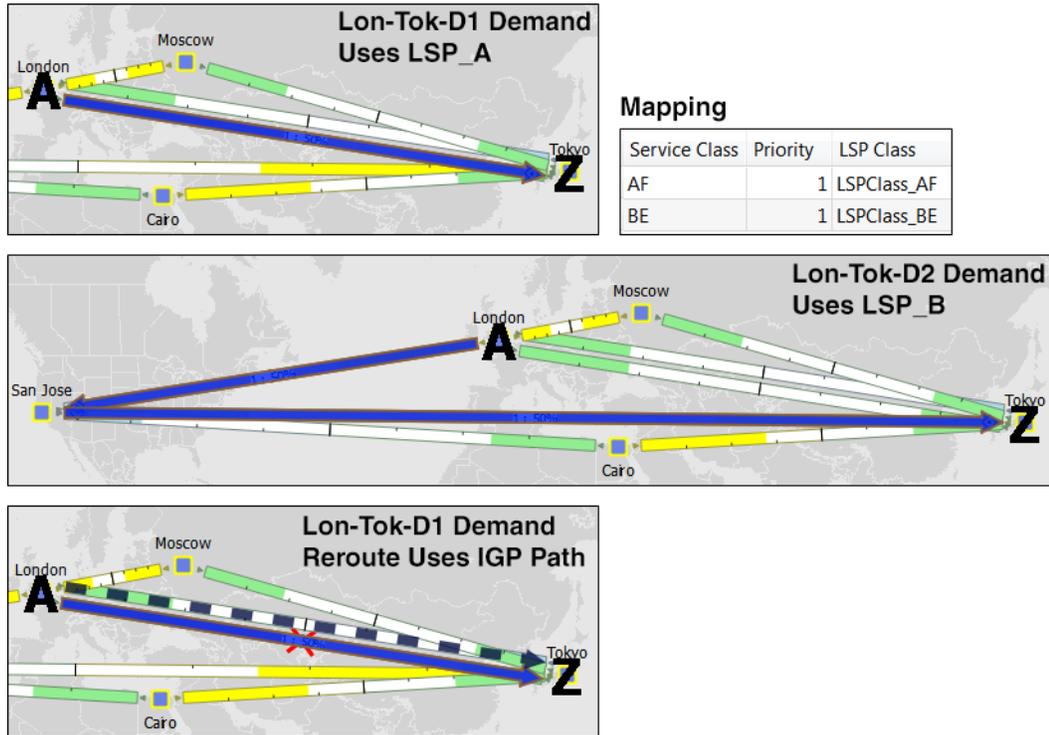
In these example simulations, there are only two demands and two LSPs. Both demands and both LSPs go from London to Tokyo.

- Lon-Tok-D1 demand uses an AF service class.
- Lon-Tok-D2 demand uses an BE service class.
- The LSPs are named LSP_A and LSP_B. Each LSP is explicitly routed to reach Tokyo via a different path.
 - LSP_A is assigned to LSPClass_AF.
 - LSP_B is assigned to LSPClass_BE.

Figure 9-12 shows a simple one-to-one mapping between a service class and an LSP class.

- The AF service class is mapped to LSPClass_AF.
- The BE service class is mapped to LSPClass_BE.
- If the circuit between London and Tokyo fails, the Lon-Tok-D1 demand reroutes using the shortest IGP path. The reason is because there are no alternative LSPs available.

Figure 9-12 Example Reroute Around Failure Using IGP Path



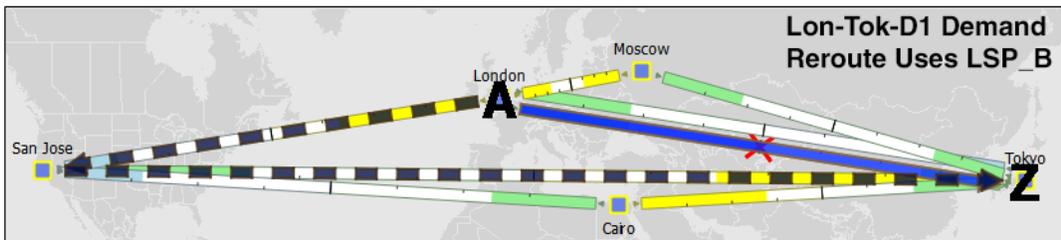
381088

Figure 9-13 shows that the AF service class is mapped to both LSPClass_AF and LSPClass_BE. The Lon-Tok-D1 demand now reroutes around a failure between London and Tokyo by using LSP_B. This is because the AF service class now has a Priority 2 mapping to the LSPClass_BE class, and the fact that LSP_B is assigned to this LSPClass_BE class.

Figure 9-13 Example Reroute Around Failure Using Alternative LSP Class

Mapping

| Service Class | Priority | LSP Class |
|---------------|----------|-------------|
| AF | 1 | LSPClass_AF |
| AF | 2 | LSPClass_BE |
| BE | 1 | LSPClass_BE |



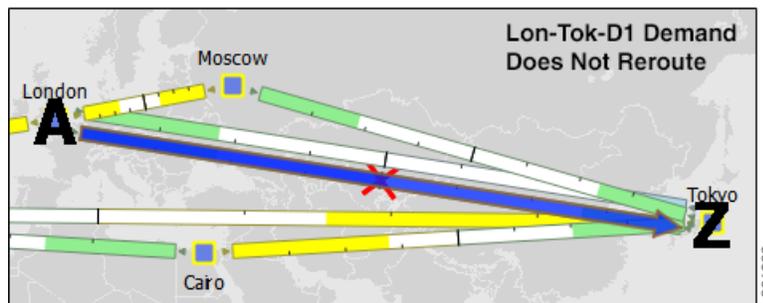
381089

Figure 9-14 shows that the AF service class is mapped to both LSPClass_AF and a Drop value. The Lon-Tok-D1 demand now does not reroute around a failure between London and Tokyo. This is because demands using AF service class are mapped such that if an LSP with an LSPClass_AF class is not available, then the traffic is dropped.

Figure 9-14 Example of Not Rerouting Based on Drop LSP Class

Mapping

| Service Class | Priority | LSP Class |
|---------------|----------|-------------|
| AF | 1 | LSPClass_AF |
| AF | 2 | Drop |
| BE | 1 | LSPClass_BE |

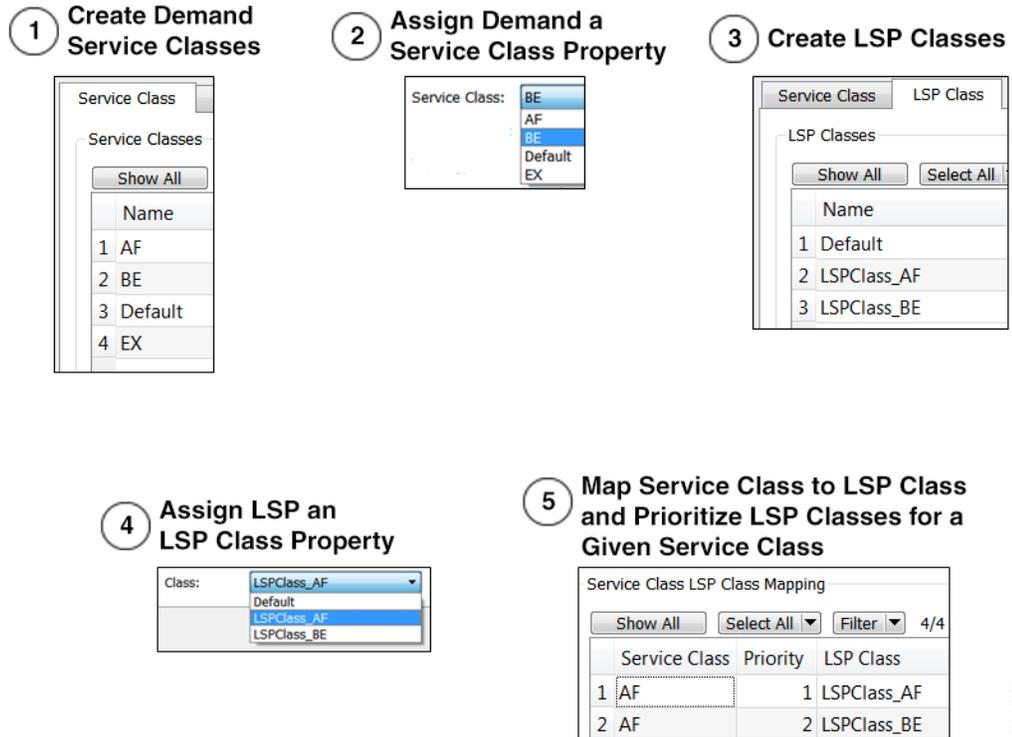


Workflow for Creating Service Class and LSP Class Mappings

The workflow for creating mappings between service classes and LSP classes is as follows (Figure 9-15).

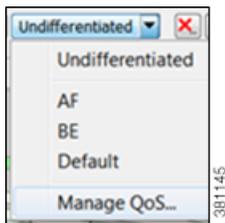
- Step 1** Create the service classes using the Manage QoS dialog box. For information, see the [Quality of Service Simulation](#) chapter.
- Step 2** Assign demands to the service classes. To assign this property, use the Service Class menu in the demand's Properties dialog box. (Right-click one or more selected demands, and select Properties.)
- Step 3** Create the LSP classes using the Manage QoS dialog box. For information, see the [Create LSP Classes](#) section.
- Step 4** Assign LSPs to the LSP classes. For information, see the [Assign LSP Classes to LSPs](#) section.
- Step 5** As needed, assign a service class to one or more LSP classes using the Manage QoS dialog box. This creates a mapping table between the demands with a given service class and LSPs with a given LSP class. Optionally, prioritize multiple LSP classes within a service class. For information, see the [Create Mappings Between Service Classes and LSP Classes](#) section.

Figure 9-15 Workflow Process for Creating Service Class to LSP Class Mappings

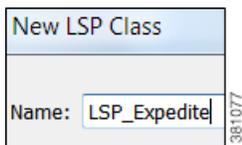


Create LSP Classes

- Step 1** Access the Manage QoS dialog box in one of two ways.
- Select Edit->Manage QoS.
 - Select Manage QoS from the QoS drop-down menu in the Visualization toolbar.



- Step 2** Select the LSP Class tab.
- Step 3** In the LSP Class area, click New.



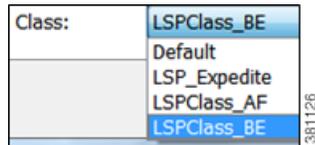
- Step 4** Enter the name of the LSP class, and click OK.

381148

Step 5 Click OK in the Manage QoS dialog box.

Assign LSP Classes to LSPs

Step 1 Select one or more LSPs.



Step 2 Right-click one of these LSPs and select Properties from the context menu.

Step 3 Select the LSP class from the Class drop-down list, and click OK.

Create Mappings Between Service Classes and LSP Classes

Step 1 Access the Manage QoS dialog box in one of two ways.

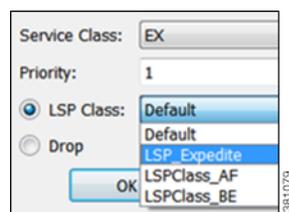
- Select Edit->Manage QoS.
- Select Manage QoS from the QoS drop-down menu in the Visualization toolbar.



Step 2 Select the LSP Class tab.

Step 3 In the Service Class LSP Class Mapping section, click New.

Step 4 In the Service Class drop-down list, select the service class to which you are mapping LSP classes.



Step 5 Enter a priority (number) for the LSP class. This determines the order in which the demand attempts to route available LSPs. The lower the number, the higher the priority.

Step 6 To map the selected service class to an LSP class, select the LSP Class radio button. Then select the class from the LSP Class drop-down list.

To set the mapping such that the demand should not reroute if LSPs are unavailable, select the Drop radio button. Note that Drop should only be used on the lowest priorities.

Step 7 Click OK, and then click OK again in the Manage QoS dialog box.

Private LSPs

WAE Design provides two ways to route selected traffic demands through specific LSPs. One way is to dedicate the traffic for specific demands to a private LSP, which is a special LSP that carries those demands only. This type of LSP models an MPLS Layer 2 VPN, providing an exclusive route for the associated demands. If the LSP goes down, all traffic associated with the LSP is interrupted.

The other way to dedicate traffic to an LSP is to filter the demands that the LSP can carry. You can do this by using the WAE Design tag feature to identify the desired demands, and then configuring the LSP with a list of tags values. This type of LSP makes it possible to route specific types of traffic over the LSP, and the type of traffic is any arbitrary group that you define.

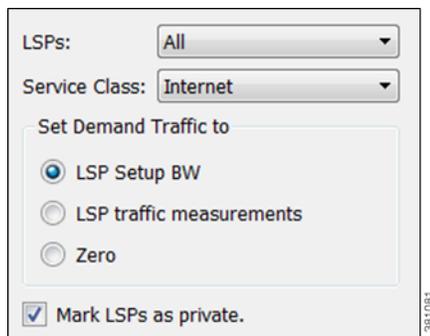
To configure LSPs that simulate Layer 2 VPNs, use one of these two tools.

- One tool creates dedicated demands for existing LSPs. The created demands will match the LSPs in source and destination.
- One tool creates private LSPs from existing demands. The created LSPs will match the existing demands in source and destination.

The LSP Private property is set to T (true) in the LSPs table, and the Private LSP Node and Private LSP Name properties in the Demands table are set.

Create Private Demands for Existing LSPs

LSPs must currently exist in the plan file.



Step 1 If you are creating demands for only selected LSPs, then select those LSPs from the LSP table. If they are tagged, you can skip this step and specify the tag value in step #3 instead. If you are creating demands for all LSPs, you can skip this step.

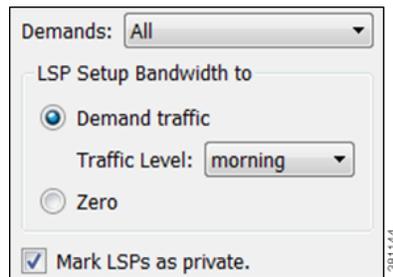
Step 2 Either select the Insert->LSPs->Demands for LSPs menu or right-click in the plot and then select New->LSPs-> Demands for LSPs menu. The Create Demands for LSPs dialog box appears.

Step 3 In the LSPs list, select All, Selected in Table (step #1), or tag name.

- Step 4** Select the service class to which these LSPs belong.
- Step 5** Set the demand traffic to equal the LSP setup bandwidth, the LSP traffic measurements, or zero.
- Step 6** Select this option: Mark LSP as private.
- Step 7** Click OK. The newly created demands are highlighted in the Demands table.

Create Private LSPs for Demands

Demands must currently exist in the plan file. See the [Traffic Demand Modeling](#) chapter.



- Step 1** If you are creating LSPs for only selected demands, then select those demands from the Demands table. If they are tagged, you can skip this step and specify the tag value in step #3 instead.
- Step 2** Either select the Insert->LSPs-> LSPs for Demands menu or right-click in the plot and then select New->LSPs->LSPs for Demands menu. The Create LSPs for Demands dialog box appears.
- Step 3** In the Demands list, select All, Selected in Table (step #1), or tag name.
- Step 4** Set the bandwidth traffic to a specific traffic level or to zero.
- Step 5** Select this option: Mark LSP as private.
- Step 6** Click OK. The newly created LSPs are highlighted in the LSPs table.

Configure LSPs to Transport Specific Traffic Types

This section describes how to configure LSPs that only route specific types of traffic. For example, you might want to route traffic for a voice service class through an LSP with a low latency, and route traffic for an Internet service class through another LSP with best-effort only. Such LSPs only carry the specified demands, but those demands might also be routed elsewhere.

- Step 1** In the Demands table, select and double-click the demands that you want to group. The Properties dialog box appears.
 - a.** Select the Tags tab.
 1. If a tag exists that you want to use, toggle its F (false) value to T (true) to enable the tag for the selected demands. If a tag does not exist that you want to use, click the New Tag button.
 2. Enter the name of the tag.
 3. Click OK.

- b. Click OK to exit the Properties dialog box. Note, if you close without clicking OK again, the changes are not saved.

Step 2 In the LSP table, select and double-click the LSPs that are to transport demands that were tagged (step #1). The Properties dialog box appears.

- a. In the Include Any Demands Tags field, enter a list of demand tags and separate each entry with a semicolon.
- b. Click OK.



Simulation Modes

WAE Design supports four simulation modes for RSVP LSPs.

- A *Fast Reroute* (FRR) mode using FRR LSPs, which is a common failure restoration mechanism used by RSVP LSPs. In a live network, FRR restoration typically occurs within milliseconds. See the [Fast Reroute Simulations](#) section.
- 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.
- 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 the [Autobandwidth Simulations](#) section.
- In the Autobandwidth Convergence Including Failures mode, the network is simulated after the Setup BW Sim values have been reset. See the [Autobandwidth Simulations](#) section.



Note

The Optimization tools only work in IGP and LSP Reconvergence mode. If you attempt to run one while in a different convergence mode, you are prompted whether to continue, and if you do, the simulation changes to 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 LSP, a pre-sigaled 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 since FRR restoration ideally becomes effective within approximately 50 milliseconds of a failure occurring.

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, then 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.
- Since there is no IGP reconvergence, demands through unrouted LSPs are unrouted. Additionally, demands through failures outside of these LSPs will not be 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

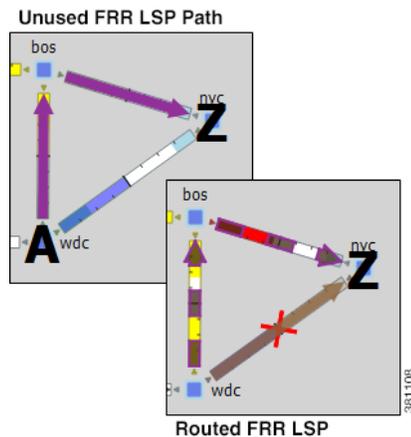
| Object | Type | Property | Description |
|-----------|--------------------------------|---------------|--|
| LSP | FRR LSP | FRR Interface | Designated interface to be avoided by the FRR LSP. |
| | | FRR Type | Type of FRR protection: Link or Node. |
| | Autoroute Forwarding Adjacency | FRR Enabled | 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 the [Run Fast Reroute Simulation](#) section.
- 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.

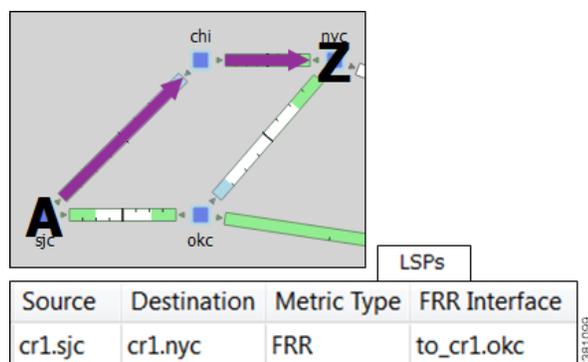


FRR LSP Routing

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

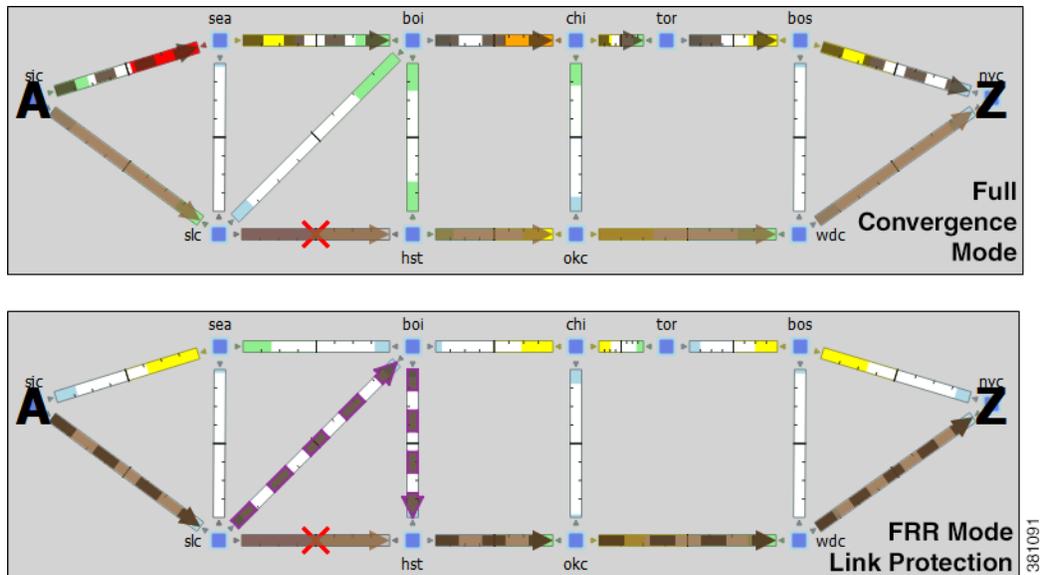
WAE Design supports both link (circuit) and node protection FRR LSPs. The type of protection shows in the FRR Type column of the LSPs table.

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

381109

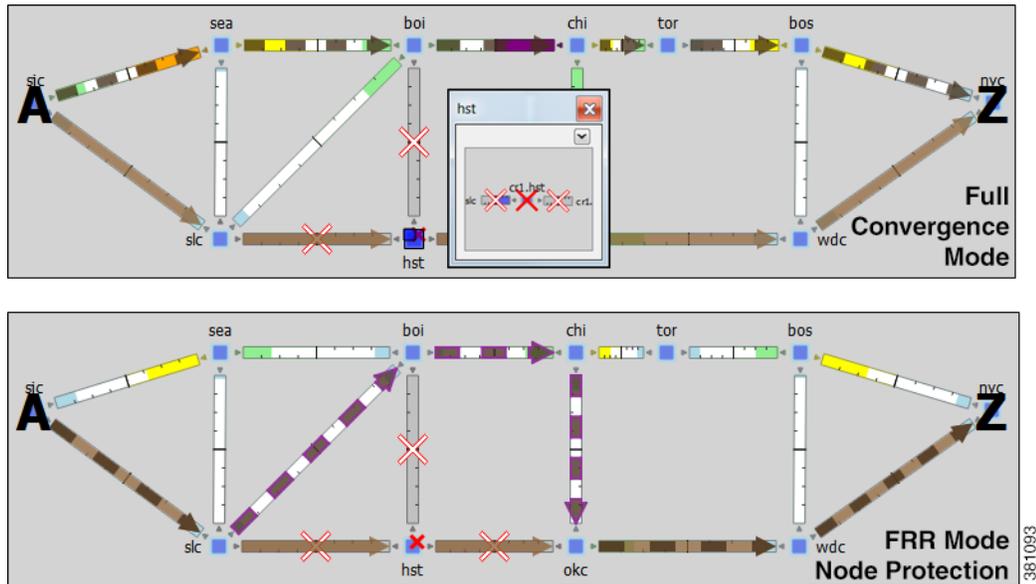
- Link-protection LSPs—Protect an LSP from a circuit failure by routing an FRR LSP around the failed circuit. The source node of the FRR LSP is the local node of the failed circuit, and if configured correctly, the destination of the FRR LSP is the remote node of the failed circuit. Figure 9-16 shows 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 9-16 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 9-16 Example Full Convergence vs. Fast Reroute with Link Protection



381091

Figure 9-17 Example Full Convergence vs. Fast Reroute with Node Protection



SRLG Protection

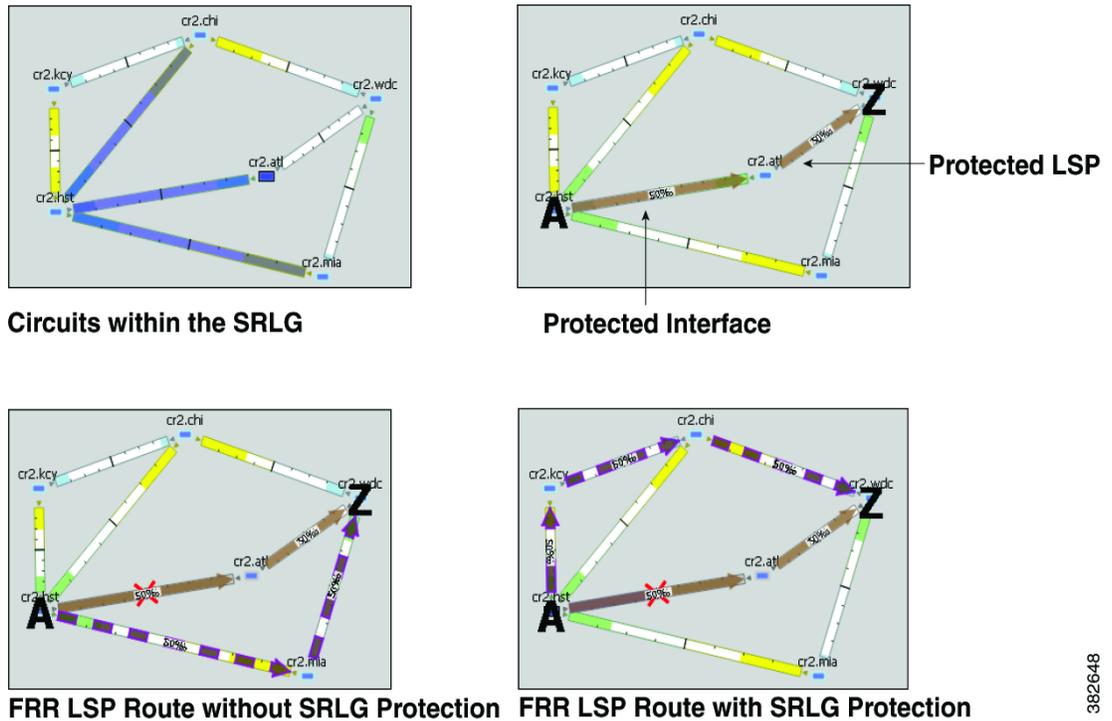
In a network, if a group of interfaces are sourced from the same node (router) and are configured as an SRLG, an FRR LSP protecting one of these interfaces can be configured to reroute around that interface. If possible, the FRR LSP avoids all other interfaces in the SRLG. In WAE Design, FRR LSPs can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

Example: Figure 9-18 shows an example of how this works using 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 9-18 Example Fast Reroute without and with SRLG Protection



Routing Around Failures

Assuming one or more failures have occurred, the following series of decisions determine which FRR LSP is taken.

Prerequisites

- LSPs must be configured to be protected by FRR LSPs. For information on setting up FRR LSPs, see the [Set Up Fast Reroute Simulations](#) section.
- Any interfaces along the FRR LSP path must be configured to be protected by FRR LSPs. For information on protecting interfaces, see [Identify LSPs for Protection](#) or the [FRR LSPs Initializer](#) (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 selects the one that is farthest down that path of the protected LSP, thus simulating standard FRR behavior where node protection is preferred over link protection.
- If more than one eligible FRR LSP has the same destination farthest down the protected LSP path, then the FRR LSP is chosen arbitrarily amongst them.
- Step 3** If more failed circuits or nodes occur after this destination, further FRR LSPs are selected in the same manner.
-

Set Up Fast Reroute Simulations

WAE Collector discovers FRR LSPs, as well as LSPs, interfaces, and SRLGs that are marked for Fast Reroute protection. Setting up Fast Reroute simulations when they are not includes the following steps.

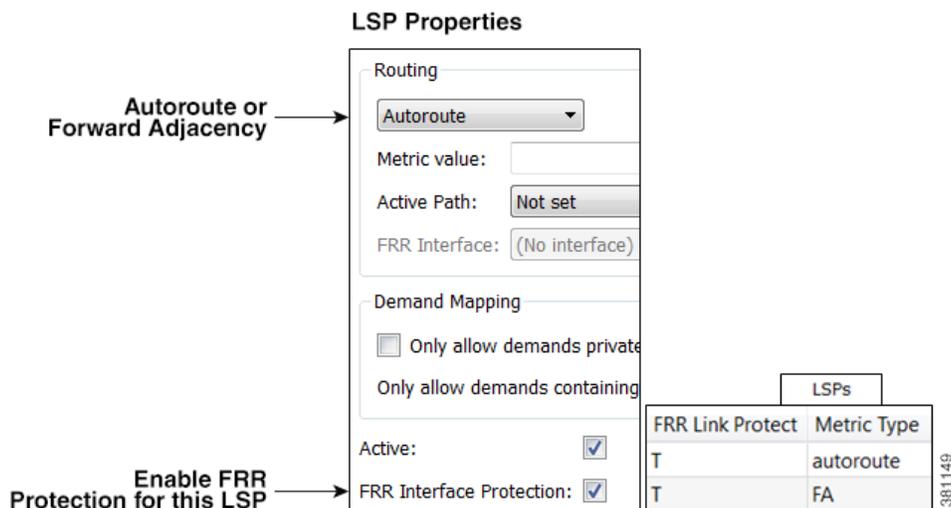
-
- Step 1** Set appropriate properties on LSPs to mark them for protection. See the [Identify LSPs for Protection](#) section.
 - Step 2** If protecting SRLGs, associate a node with the SRLG in order to establish the protected interface within it. See the [Identify SRLGs for Protection](#) section.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

Identify LSPs for Protection

To be marked for Fast Reroute protection, LSPs must have two properties set in the LSP Properties dialog box ([Figure 9-19](#)).

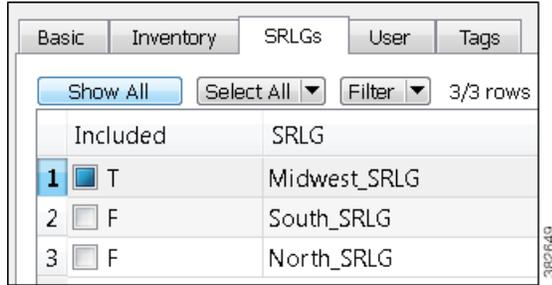
- Routing type must be Autoroute or Forwarding Adjacency. The type appears as “autoroute” or “FA” in the Metric Type column of the LSPs table.
- FRR Enabled must be selected. This appears as T (true) in the FRR Enabled column of the LSPs table. If this is F (false), the LSP is not protected.

Figure 9-19 Protected LSP Properties



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. For an example, see [Figure 9-18](#).

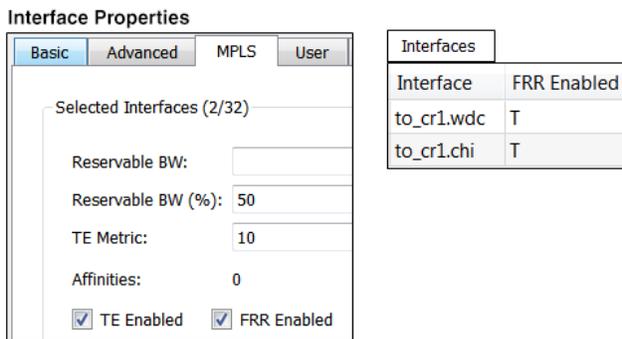


- Step 1** Right-click the source node of a circuit within an SRLG that you want to protect, and select Properties.
- Step 2** Select the SRLG tab.
- Step 3** Select one or more SRLGs with which you want to associate this node, and click OK.

FRR LSPs Initializer

The FRR LSPs initializer automatically creates link-protection and/or node-protection FRR LSPs provided two conditions are met, as follows.

- One or more LSPs have their FRR Enabled property set. (See the [Identify LSPs for Protection](#) section.)
- 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 FRR Enabled column in the Interfaces table shows T (true) if the interface is included or F (false) if it is not.
 - Once the FRR LSPs are created, this interface is listed in the FRR Interface column of the LSPs table.



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

The initializer creates the FRR LSPs based on these FRR Enabled properties and based on whether you select 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 9-19).
- 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 9-20).

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

Figure 9-20 Example Link-Protection FRR LSPs Automatically Created

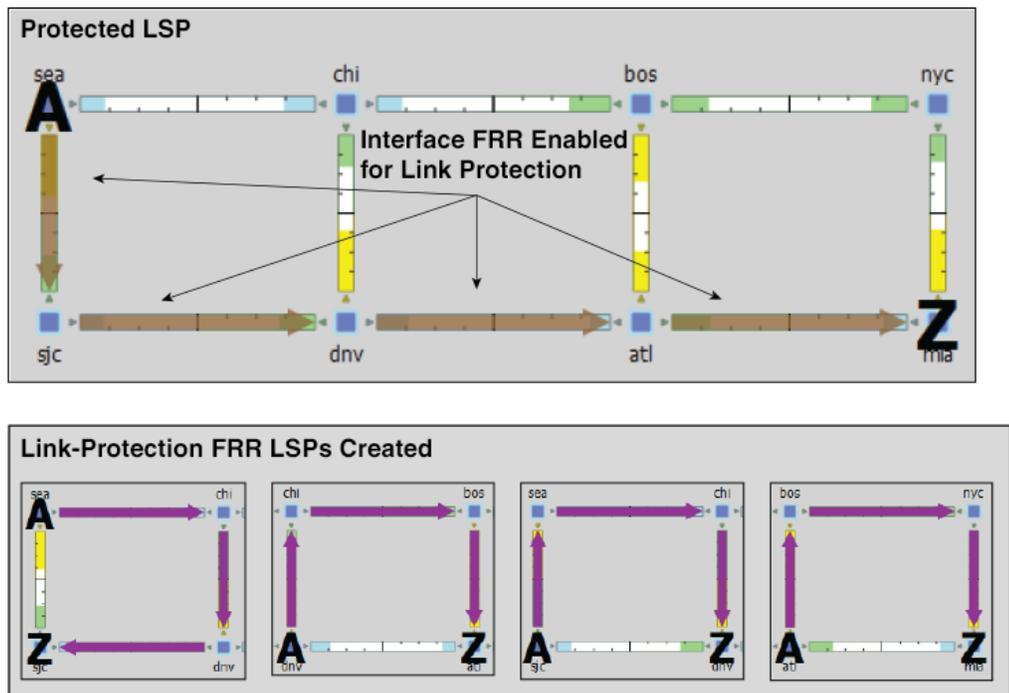
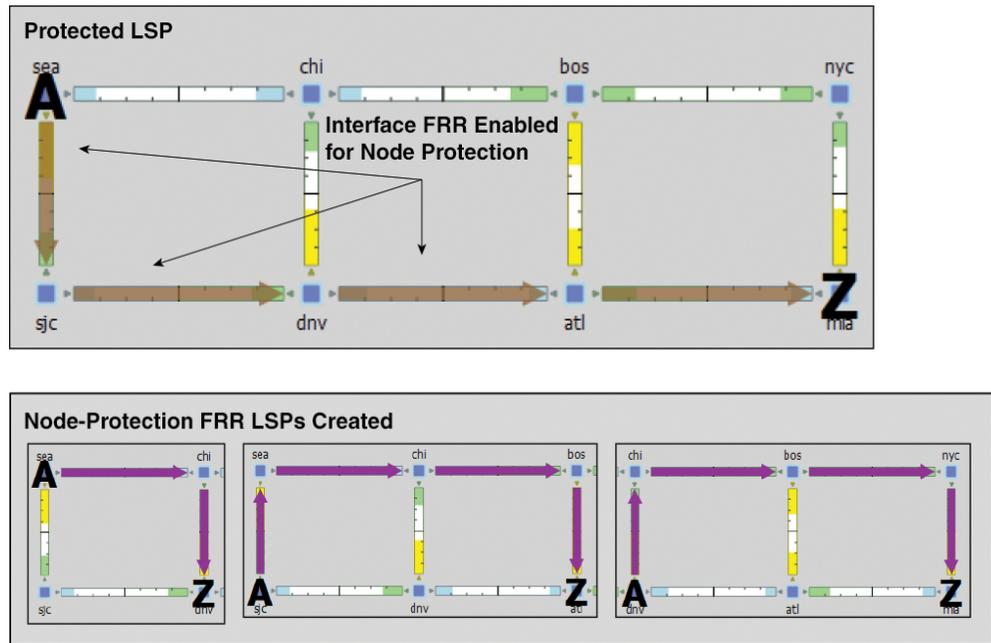
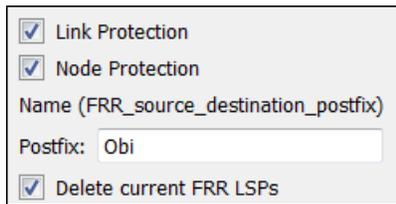


Figure 9-21 Example Node-Protection FRR LSPs Automatically Created



Run FRR LSPs Initializer

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



Step 2 Select the Initializers->FRR LSPs menu.

Step 3 Select 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 Create 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 Either select the Insert->LSPs->LSP menu, or right-click in an empty area and select New->LSPs->LSP from the context menu. Alternatively, you can select an existing LSP to modify to be an FRR LSP. If so, right-click the LSP and select Properties.

- Step 2** Select the Source site and node of the FRR LSP.
- Step 3** Either select the Destination site and node, or enter a NetInt Destination IP address. For information on NetInt Destination IP addresses, see the [Unresolved LSP Destinations and Hops](#) section.

- Step 4** In the Routing section, select Fast Reroute as the LSP type
- Step 5** Specify the interface that is to be avoided during routing by selecting it from the FRR Interface drop-down list or by entering its IP address in the NetInt FRR Interface field.
- Step 6** Press OK.

Run 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** Either select the Edit->Network Options menu or click the Network Options icon in the Visualization toolbar.
- Step 2** Select the Simulation tab.

Step 3 Select the Fast Reroute option, 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. Prior to simulating failures or running Simulation Analysis on autobandwidth-enabled LSPs, select the appropriate mode depending on which state you need to simulate.

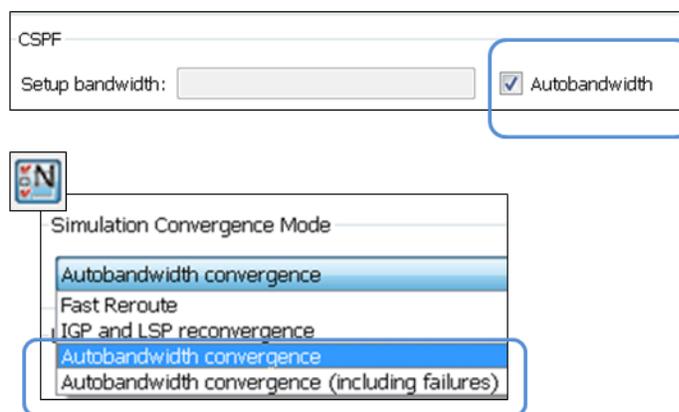
- In the Autobandwidth Convergence mode, the network is simulated after the traffic (Traff Sim) has been rerouted to other LSPs, but before the Setup BW Sim values have been reset. This simulates the immediate, less optimal response to failures.
- In the Autobandwidth Convergence Including Failures mode, 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.

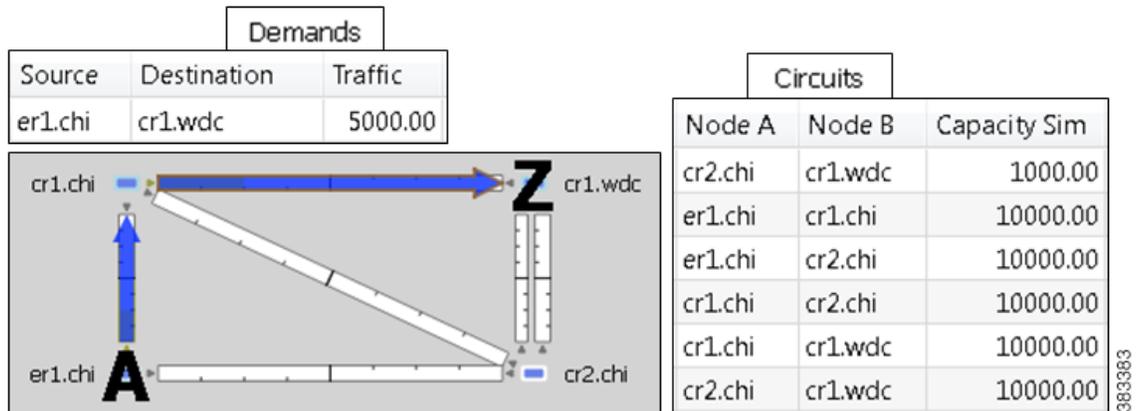
These autobandwidth simulations require that the LSPs have their Autobandwidth property set to T (True), and that one of the two autobandwidth modes be selected from the Network Options dialog box.



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 1,000 Mbps.



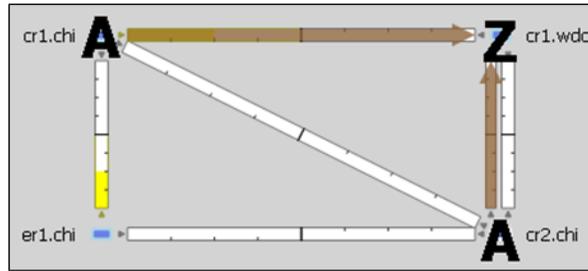
- **Figure 9-22** shows that in the IGP and LSP Reconvergence mode, the Setup BW Sim value for both LSPs is 0.
 - In the Autobandwidth Convergence mode, the Setup BW Sim value for LSP-A is copied from its Traff Sim value, which is 5,000.
 - There are no failures, and each LSP follows the same route in both convergence modes.
- **Figure 9-22** 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 ("sloshes") over to LSP-B. This is shown in the LSPs table where the traffic (Traff Sim) of LSP-A is moved to LSP-B.
 - Since 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 9-23** 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 since the LSP has found a path with sufficient capacity, there is no more congestion.

Figure 9-22 Example LSP Routing Using IGP and LSP Reconvergence and Using Autobandwidth Convergence

Simulation Convergence Mode

IGP and LSP reconvergence

| LSPs | | | | | | | |
|-------|---------|-------------|----------|--------------|-----------|---------------|--|
| Name | Source | Destination | Setup BW | Setup BW Sim | Traff Sim | Autobandwidth | |
| LSP-A | cr1.chi | cr1.wdc | na | 0.00 | 5000.00 | T | |
| LSP-B | cr2.chi | cr1.wdc | na | 0.00 | 0.00 | T | |



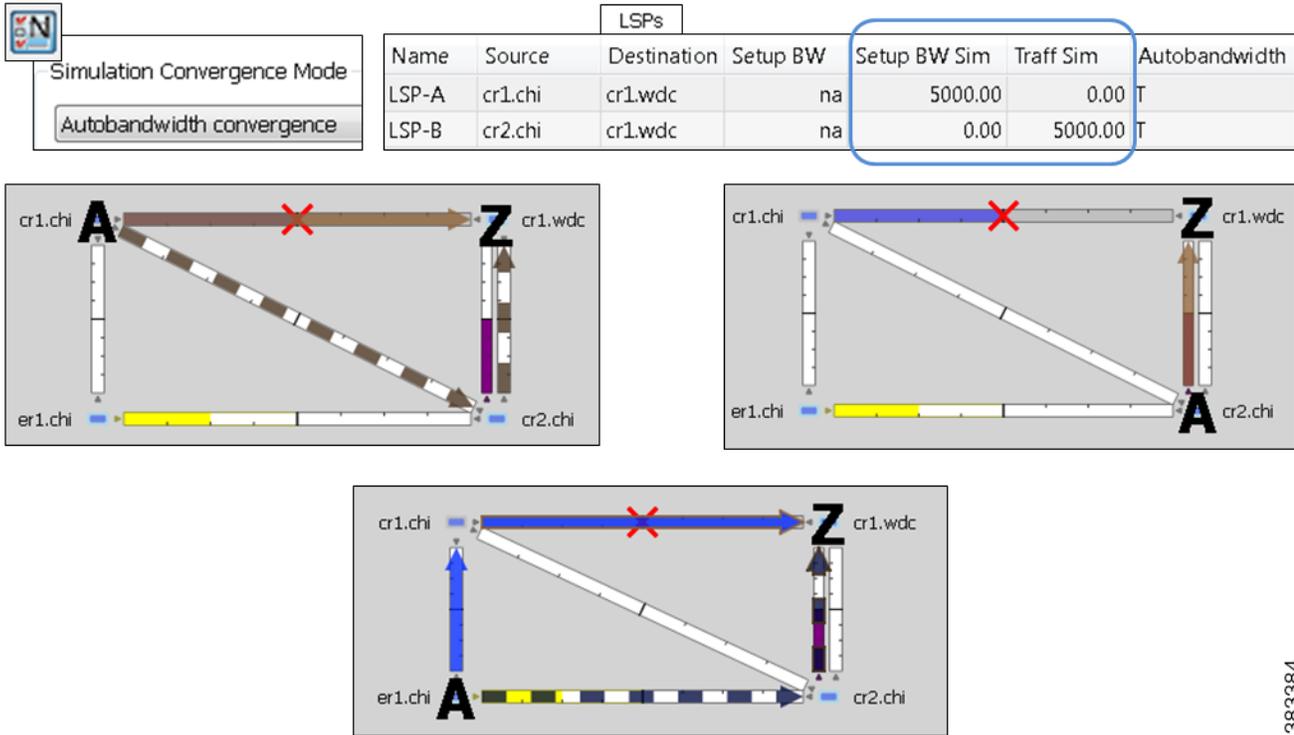
Simulation Convergence Mode

Autobandwidth convergence

| LSPs | | | | | | | |
|-------|---------|-------------|----------|--------------|-----------|---------------|--|
| Name | Source | Destination | Setup BW | Setup BW Sim | Traff Sim | Autobandwidth | |
| LSP-A | cr1.chi | cr1.wdc | na | 5000.00 | 5000.00 | T | |
| LSP-B | cr2.chi | cr1.wdc | na | 0.00 | 0.00 | T | |

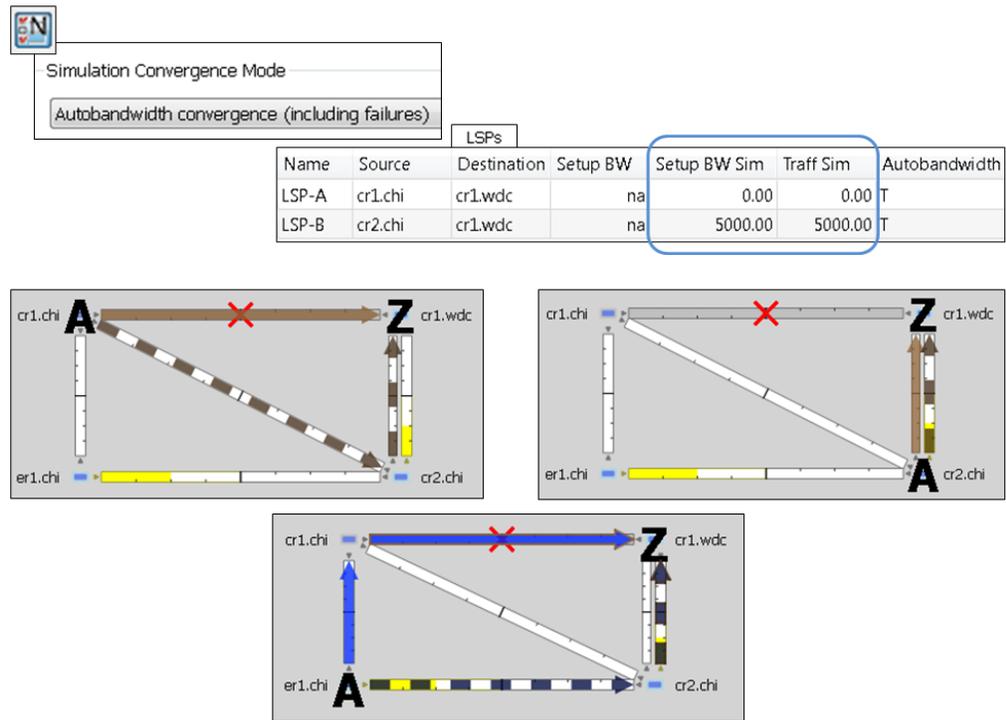
383385

Figure 9-23 Example Autobandwidth Convergence with Failures



383384

Figure 9-24 Example Autobandwidth Convergence Including Failures



Advanced RSVP LSP Simulations

Ignore 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 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 *WAE Design Integration and Development Guide*.

This operates only on LSPs that do not have autobandwidth enabled (the Autobandwidth property is set to False).

Step 1 Route LSPs as usual. (For information, see the [Dynamic LSP Routing and CSPF](#) section.)

- Step 2** If an LSP is not routable (call it “LSP-A”), attempt to route it with zero setup bandwidth.
- If LSP-A cannot be routed, leave it unrouted and attempt to route the next LSP in the LSPs table. If this LSP cannot be routed, attempt 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 1,000 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 9-24 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 purple circuits indicate they are being over utilized and likely need their capacity increased. (For information on the LSP Reservations view, see the [View LSP Reservations](#) section.)

Figure 9-25 Comparison of Reservable Bandwidth Being Used vs Not Being Used

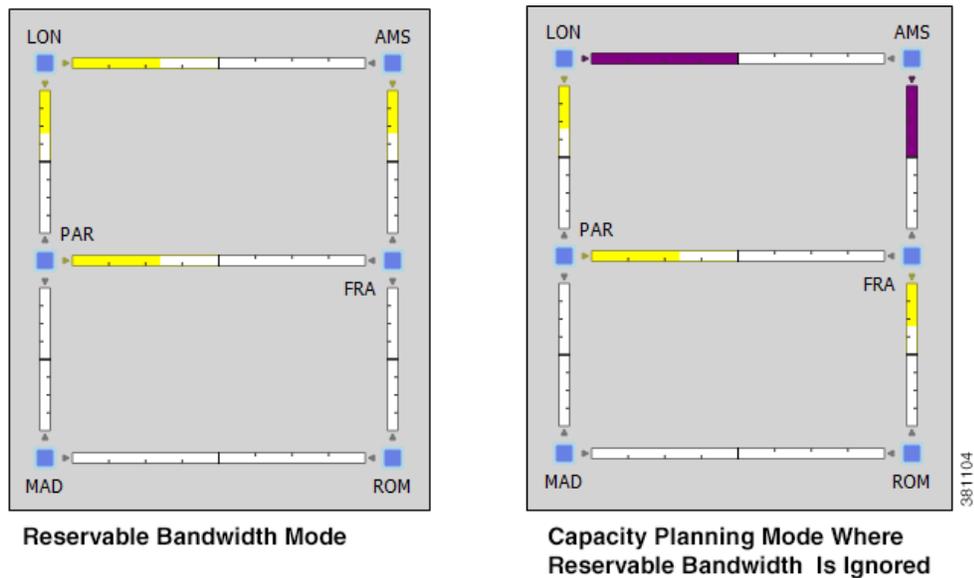
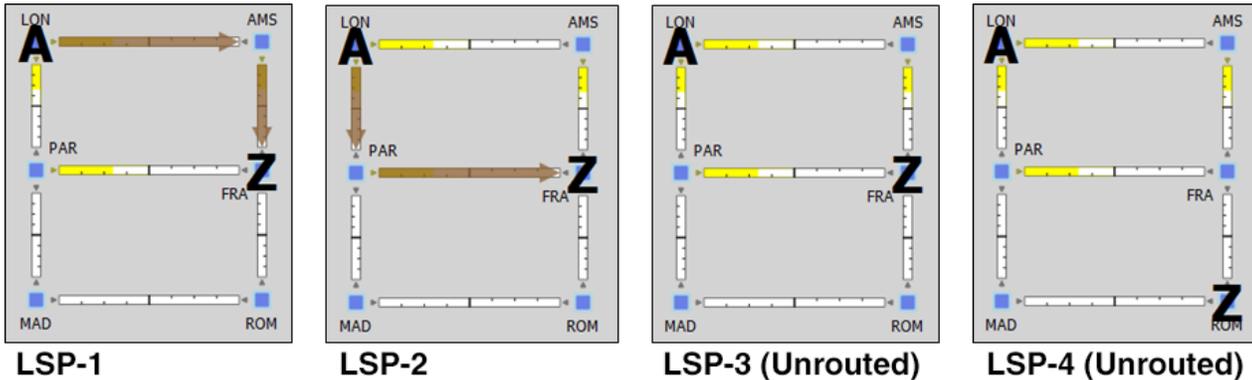


Figure 9-26 demonstrates routing simulations where reservable bandwidth constraints are applied, which is the default behavior.

- 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 9-26 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 |

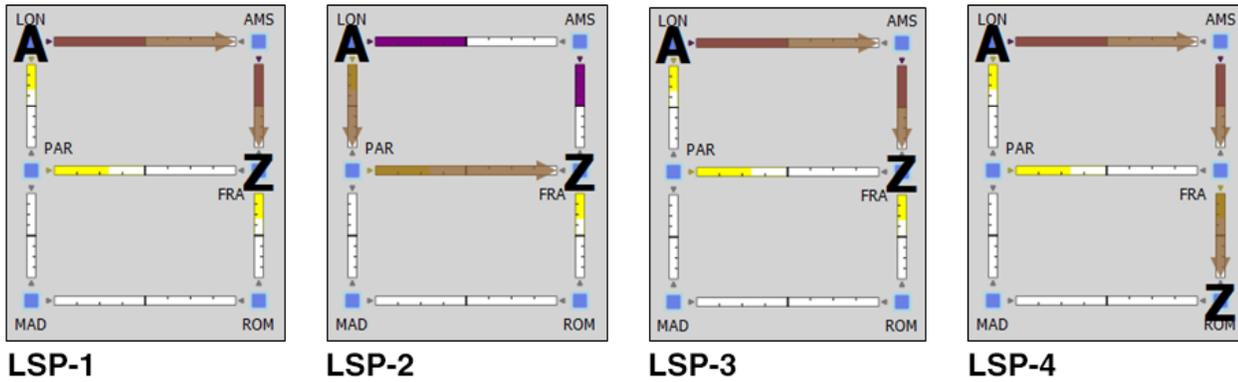
| LSPs | | | | | |
|-------|---------|-------------|----------|--------------|-----------|
| Name | Source | Destination | Setup BW | Setup BW Sim | Routed |
| LSP-1 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-2 | cr1.lon | cr1.fra | 600.00 | 600.00 | Simulated |
| LSP-3 | cr1.lon | cr1.fra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

381107

In Figure 9-27, 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 9-27 Simulated Routing Without Using Reservable Bandwidth Constraints

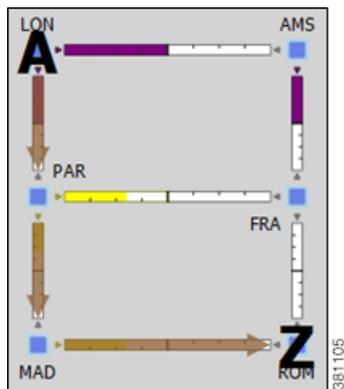


| Interfaces | | | | |
|------------|------------|---------|-------------|-------------|
| Node | Interface | Resv BW | Resv BW Sim | Remote Node |
| cr1.lon | to_cr1.par | 1000.00 | 1000.00 | cr1.par |
| cr1.ams | to_cr1.fra | 1000.00 | 1000.00 | cr1.fra |
| cr1.par | to_cr1.fra | 1000.00 | 1000.00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

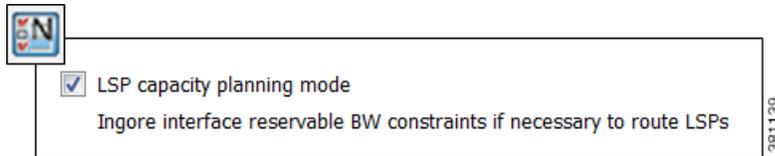


LSP-4

381105

Enable Capacity Planning Mode for LSPs

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



-
- Step 1** Either select the Edit->Network Options menu or click the Network Options icon in the Visualization toolbar.
- Step 2** Select the Simulation tab.
- Step 3** Select the option to use the LSP capacity planning mode, thus ignoring reservable bandwidth constraints, and click OK.
-

P2MP LSPs

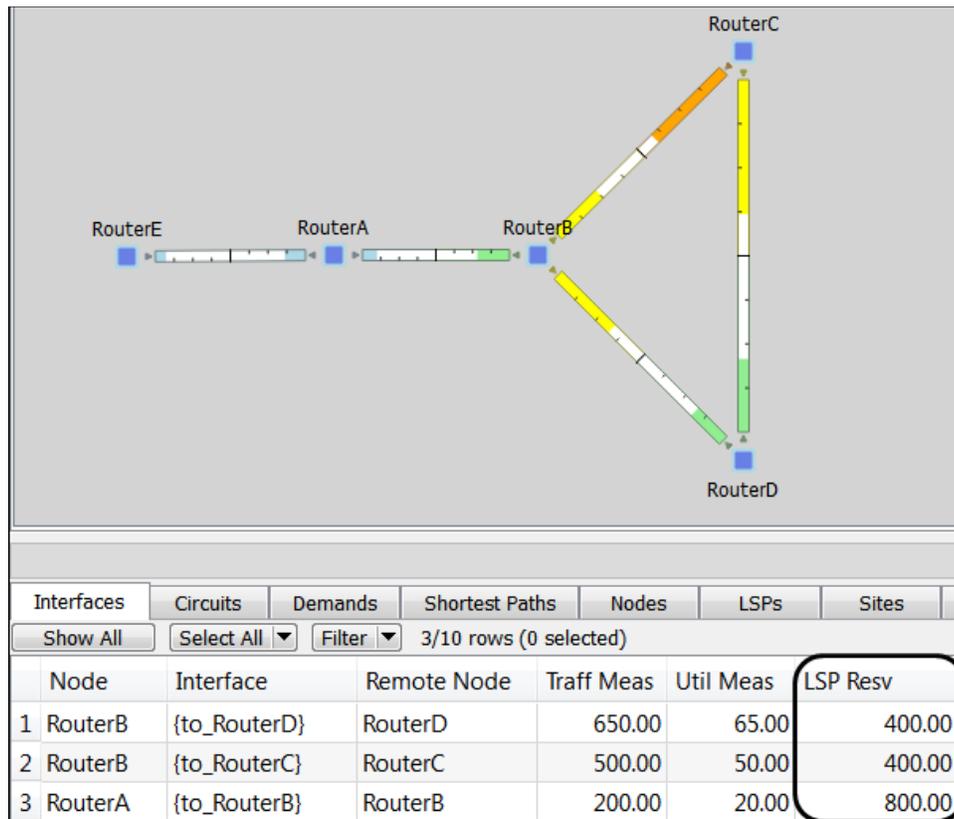
Point-to-multipoint (P2MP) LSPs offer an MPLS alternative to IP multicast for setting up multiple paths to multiple locations from a single source. In live networks, the signal is sent once across the network and packets are replicated only at the relevant or designated branching MPLS nodes.

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, select them from the P2MP LSPs and LSPs tables.

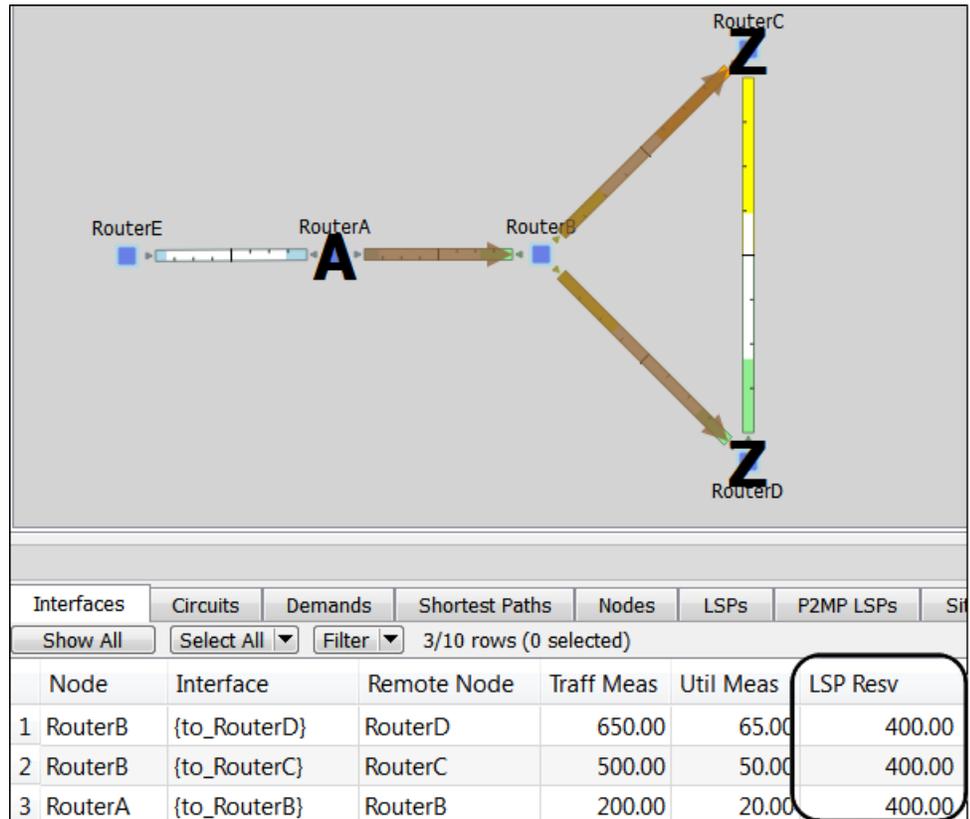
- To determine if an LSP belongs to a P2MP LSP, you can show the P2MP LSP column in the LSPs table. You can also right-click an LSP, and select Filter to P2MP LSPs.
- To determine which LSPs belong to a P2MP LSP, you can visually see them by selecting the P2MP LSP from the P2MP LSPs table. You can also right-click a P2MP LSP, and select Filter to sub LSPs.

Figure 9-29 Bandwidth Behavior without a P2MP LSP



361122

Figure 9-30 Bandwidth Behavior with a P2MP LSP



P2MP LSP Demands

Demands must be private to the P2MP LSP. (See the [Private LSPs](#) section.) 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.

Explicit P2MP LSP Path Initializer

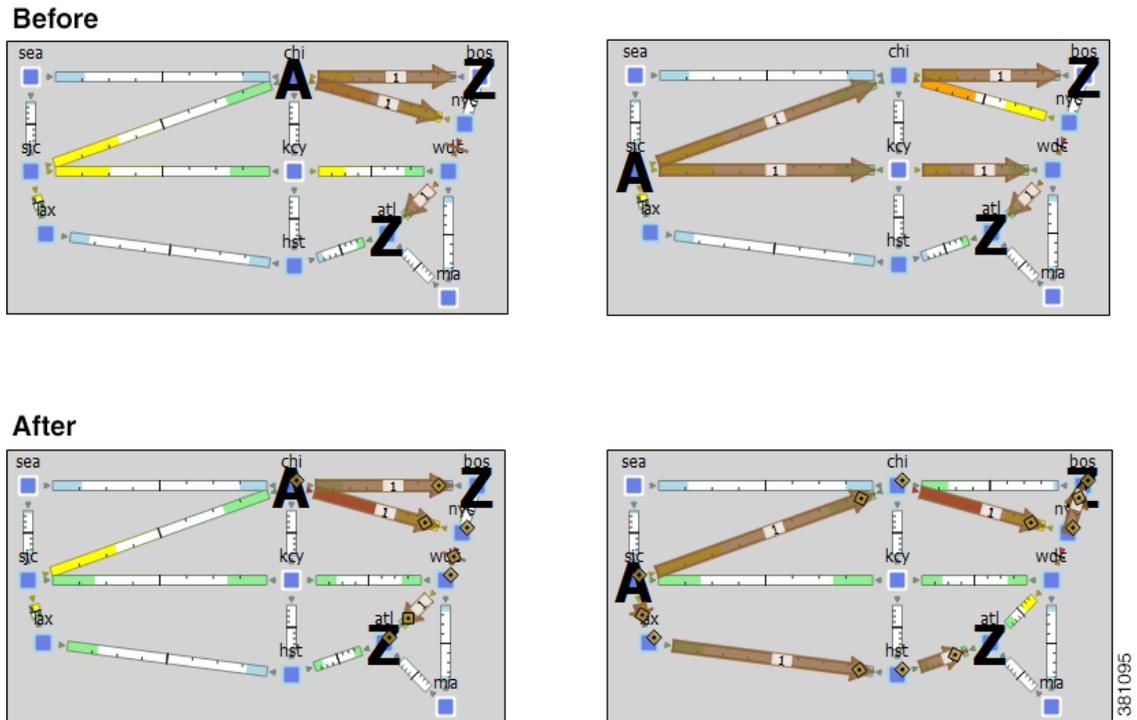
The Explicit P2MP LSP Paths Initializer automates the process of creating explicitly routed sub-LSP paths, as follows.

- Convert a set of dynamically routed sub-LSPs to explicitly routed sub-LSPs along their current paths.
- Create disjoint paths. Two P2MP LSP paths are *disjoint* if they do not route over common objects. These objects are configurable and may 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 9-31 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 such that they arrive at the common destinations using different explicit paths.

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

Figure 9-31 Example Disjoint Paths for P2MP LSPs



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 table, respectively.

- 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.
- The explicit paths are named <Sub-LSP Name>_<Path Option>.

Prerequisites

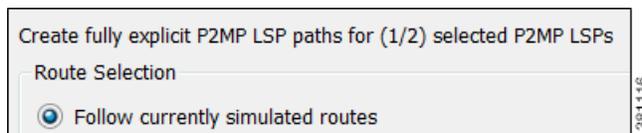
- The plan file must already contain the LSP paths. See the [Create LSP Paths](#) section.

- 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.



Disjoint Group: GBD
Disjoint Priority: 1

-
- Step 1** From the P2MP LSPs table, select the P2MP LSPs to initialize with explicit named paths. If you do not select any P2MP LSPs, all sub-LSP paths are used for creating the named paths.
- Step 2** Select the Initializers->P2MP LSP Explicit Paths menu, or right-click a selected P2MP LSP and select Explicit Path Initializer from the context menu. The P2MP LSP Explicit Path Initializer dialog box appears.
- Step 3** Select 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, select the “Follow currently simulated routes” option.

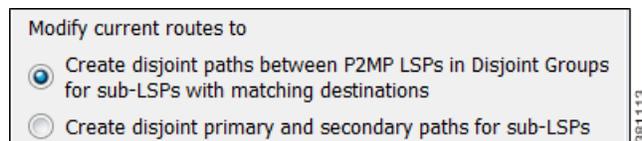


Create fully explicit P2MP LSP paths for (1/2) selected P2MP LSPs

Route Selection

Follow currently simulated routes

- To create disjoint paths, select 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.



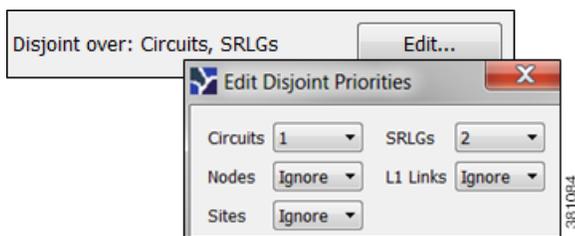
Modify current routes to

Create disjoint paths between P2MP LSPs in Disjoint Groups for sub-LSPs with matching destinations

Create disjoint primary and secondary paths for sub-LSPs

- To modify the disjoint priorities that determine which objects are considered in the paths, click Edit. A dialog box appears enabling you to specify priorities as 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, then when the disjoint paths are created, WAE Design first creates paths with disjoint circuits. Then, if there are common SRLGs in those paths, it further creates paths such that the SRLGs are disjoint. The remaining objects are not considered in the disjoint path creating.



Create P2MP LSPs and Sub-LSPs

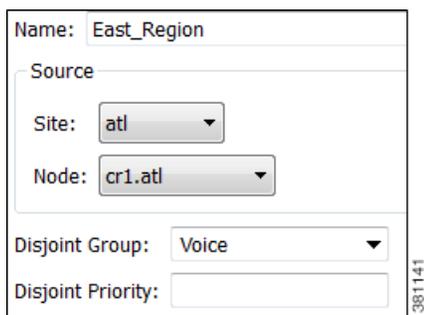
Rules and Guidelines

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

Create P2MP LSPs

Step 1 Create an empty P2MP LSP.

- Either select the Insert->LSPs-> P2MP LSP menu, or right-click in the plot and select New->LSPs->P2MP LSP from the context menu. A P2MP LSP dialog box appears.
- Enter a unique P2MP LSP name. If you do not enter a name, WAE Design creates one using the format p2mp_lsp_<source_node>.
- Select both the source site and node. This selection is the root for all sub-LSPs contained within this P2MP LSP.
- Optional: Enter or select a disjoint group and disjoint priority for WAE Design to consider when using the Explicit P2MP LSP Path Initializer. For more information, see the [Explicit P2MP LSP Path Initializer](#) section.



Step 2 If you have not done so, create sub-LSP for the P2MP LSP. See the [Create Sub-LSPs](#) section.

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

- Double-click the sub-LSPs from the LSP table. The LSP Properties dialog box appears.
- In the Sub-LSP of P2MP LSP list (at the bottom), select the P2MP LSP to which you want these LSPs to belong.

- c. Click OK.
-

Create Sub-LSPs

- Step 1** Either select the Insert->LSPs->LSP menu, or right-click in the plot and select New->LSPs->LSP from the context menu. An LSP dialog box appears.
 - Step 2** In the Source fields, select the site and node that will be the source of the P2MP LSP. All sub-LSPs within a P2MP LSP must have the same source site and source node.
 - Step 3** In the Destination fields, select the site and node that are the final destination for the sub-LSP.
 - Step 4** Complete the remaining optional fields and tabs, as needed. Since 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, select that P2MP LSP from the Sub-LSP of P2MP LSP list (at the bottom).
 - Step 6** Click OK.
-

Create 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, select and then double-click all nodes that will be included in the P2MP LSP. The Properties dialog box appears.
 - b. Select the Tags tab.
 1. Click New Tag.
 2. Enter a name that collectively identifies the sub-LSPs you are creating.
 3. Click OK.
 - c. Click OK in the Properties dialog box. Note if you close without clicking OK again, the above changes are not saved.
- Step 2** In the Nodes table, select the node you will be designating as the source for the P2MP LSP.
- Step 3** Either select the Insert->LSPs->LSP Mesh menu, or right-click in the plot and select New->LSPs->LSP Mesh from the context menu. An LSP Mesh Creator dialog box appears.
- Step 4** Deselect these two fields.
 - Same Destinations as Sources
 - Also create LSPs from destination to source
- Step 5** In the Source Nodes menu, select the Selected in table (##/#) option. All sub-LSPs within a P2MP LSP must have the same source node. A P2MP LSP cannot have multiple source nodes.
- Step 6** In the Destination Nodes menu, select 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, select the Autobandwidth option.
- Step 8** Complete the remaining optional fields as needed.

Step 9 Click OK.

Route Traffic through P2MP LSPs

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

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

Step 2 Either select the Insert->LSPs->Demands for P2MP LSPs menu, or right-click in the plot and select New->LSPs->Demands for P2MP LSPs from the context menu. A Create Demands for P2MP LSPs dialog box appears.

- a. Select the P2MP LSPs through which you are routing traffic.
 - b. Select the desired bandwidth: Minimum sub-LSP Setup BW or Zero.
 - c. Click OK.
-

Delete P2MP LSPs and Sub-LSPs



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

Disassociate Sub-LSPs from their P2MP LSPs

Step 1 In the LSPs table, select 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, either

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

Step 4 Click OK.

Delete Sub-LSPs or P2MPs

| Delete Sub-LSPs but Keep P2MP LSP | Delete P2MP LSP and Associated Sub-LSPs |
|--|---|
| 1. Right-click one or more sub-LSPs from the LSPs table. | 1. Right-click one or more P2MP LSPs from the P2MP LSP table. |
| 2. Select Delete. | 2. Select Delete. |
| 3. Select Yes to continue. | 3. Select Yes to continue. |

Unresolved LSP Destinations and Hops

RSVP 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 Collector discovery tools have not been able to read all the nodes, or not able to 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 Collector discovery tools attempt to resolve as many of these references as possible. Following is how the related 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 or not.

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.

Resolve Plan File to Objects

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

- Resolve LSP destinations, hops in named paths, or hops in actual paths using the Resolve Plan Initializer.
- Resolve named path hops in the named paths Properties dialog box.
- Resolve named path hops in the named path hops Properties dialog box.

Segment Routing

By default, WAE Design creates and routes RSVP LSPs. However, you can create Segment Routing (SR) LSPs³ by changing the LSP Type property to SR. These SR LSPs use autoroute, forwarding adjacencies, and class-based forwarding, just like RSVP LSPs. It is once the packet enters the tunnel that the different

3. WAE Design uses the terms *SR LSPs* that use *segment lists*, which contain *segment list hops*. In Cisco routing terminology, these terms are *SR tunnels* that use *segment ID (SID) lists*, which contain *SIDs*.

routing mechanisms determine the path. An RSVP LSP is first established along a path, with each node on the path accepting the reservation request and maintaining the reservation state. In contrast, SR LSPs rely on a segment list that is created at the LSP's source node (head-end). This segment list, which is stored in the packet (demand), directs the traffic over a series of configured node segments and interface segments. SR LSPs use IP routing (and thus, potentially ECMP) between segments.

Segment Types

The segment list of an SR LSP contains one or more node segments, interface segments, and/or anycast groups.

- Node segment—The SR LSP takes the shortest path to the specified node. [Figure 9-32](#) shows an example SR LSP from cr2.kcy node to cr2.mia node using one segment node, which is cr2.okc. Since the IGP metrics are equal, the traffic is routed using ECMP to reach okc.



Note

The process of creating and optimizing segment lists and their segment list hops can be automated using the SR-TE Optimization tool. For information, see the [SR-TE Optimization](#) chapter.

- Interface segment—The SR LSP routes traffic over this interface. On routers, interface segments can only be used in segment lists if the previous segment is the node containing the interface, or if the interface is the egress of the source node. In WAE Design, there is no restriction that an interface segment is local. Meaning, there is no requirement that the segment list contain a preceding segment to the node containing the interface.

[Figure 9-32](#) shows an example of an interface segment from cr2.chi to cr2.wdc. To ensure the path reaches cr2.chi (the node the packet must reach to use this interface segment), a node segment is configured first to reach chi. This figure also shows how a demand through an SR LSP reroutes when a node fails. The demand uses IGP to route around the failure and back to the segment, if possible.

- Anycast group segment—The SR LSP routes through the node in the anycast group that has the shortest path to it. In case of a tie between multiple nodes in an anycast group, ECMP is applied. This mechanism enables you to impose routing restrictions in terms of potential intermediate segment destinations. It also enables the SR LSP to choose among the possible next hops (next segment list hops), potentially reducing latency and improving load balancing.

[Figure 9-33](#) shows an example of anycast group nodes cr2.chi and cr1.chi. In one instance, the IGP metrics to those nodes have equal IGP metrics, so the SR LSP uses ECMP to route through both of them. When the IGP metric is increased from cr2.lax to cr1.chi, the SR LSP routes only through cr2.chi since it has the shortest path. However, as [Figure 9-34](#) shows, if a failure prevents this shortest path from being used, the SR LSP can still route through the anycast group using the next-highest cost path.

To view a segment list, right-click the SR LSP's source node to open the Properties dialog box, and then open the Segment tab. Alternatively, right-click an SR LSP and select Filter to Segment Lists. You can then right-click the segment list and select Filter to Segment List Hops, or double-click the segment list.

Figure 9-32 Example Node Segment

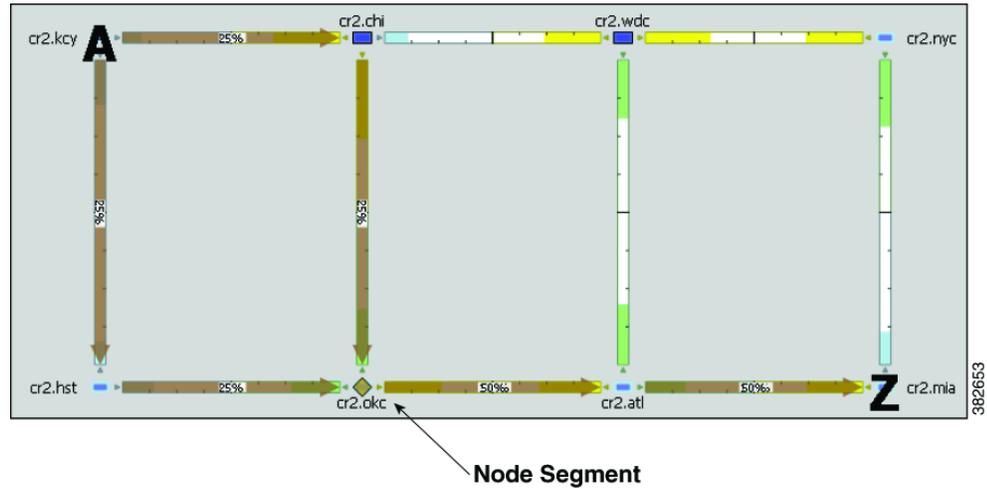


Figure 9-33 Example Interface Segment

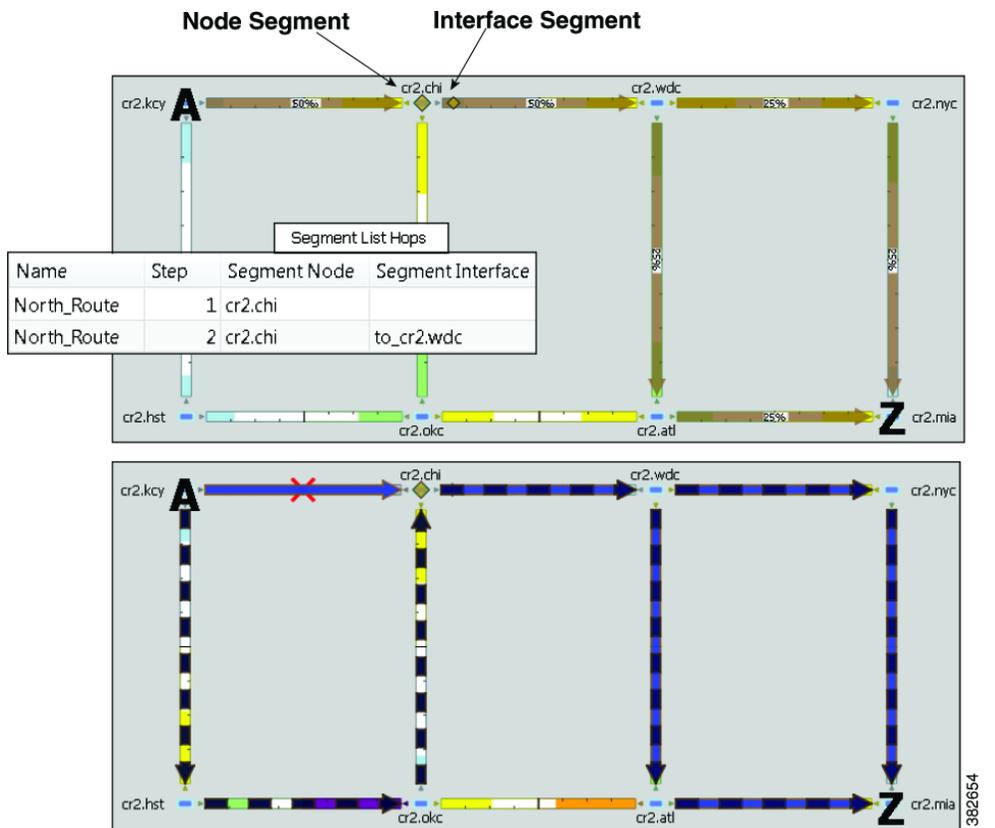
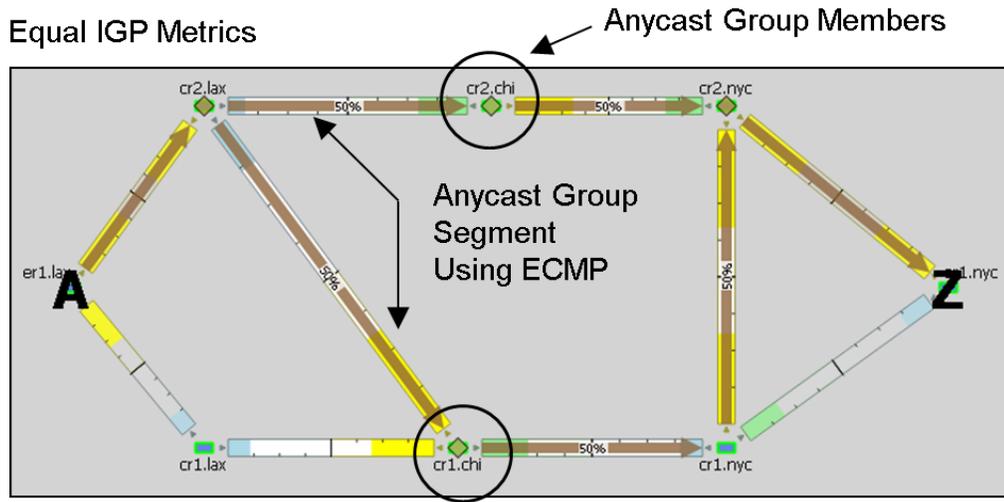
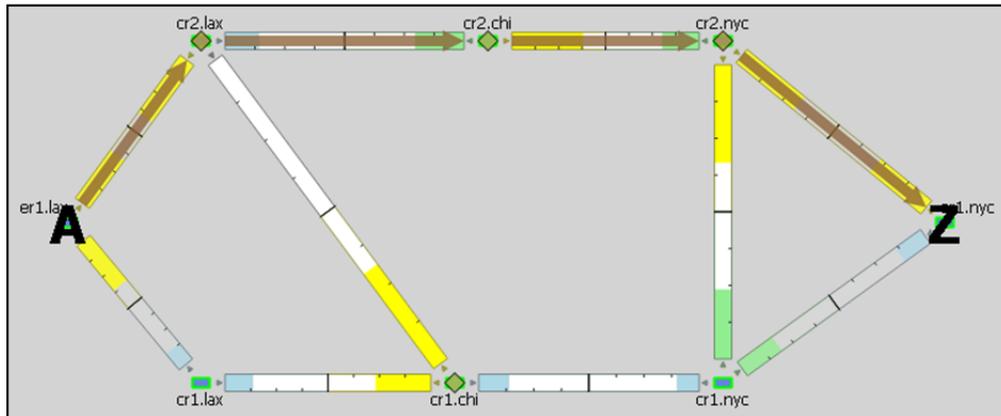


Figure 9-34 Example Anycast Group Segment

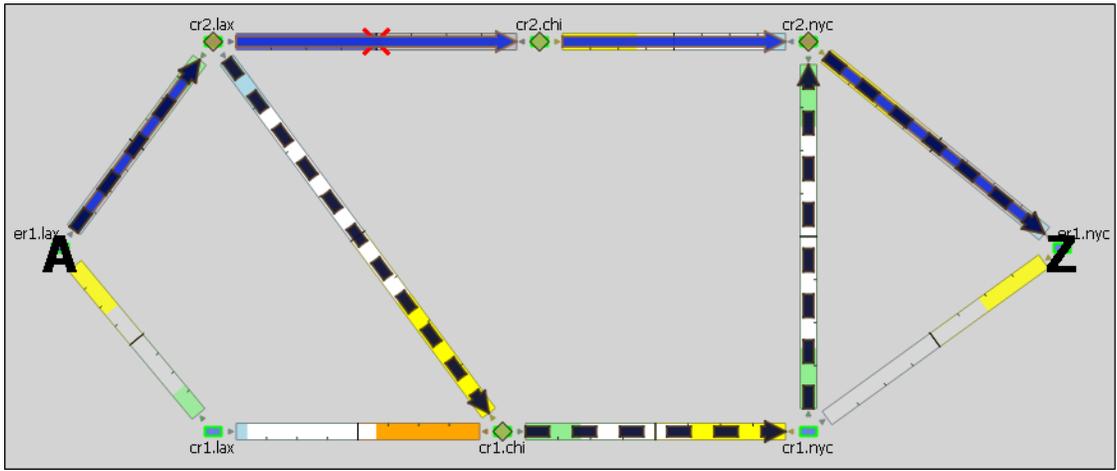


SR LSP Route Based on Anycast Group Node with Shortest Path



404549

Figure 9-35 Example Anycast Group Segment Reroute Around Failure



404650

SR LSP Paths

Like RSVP LSPs, SR LSPs can have multiple LSP paths, where each LSP path has its own segment list. There is no distinction between standby and non-standby LSP paths for SR LSPs.

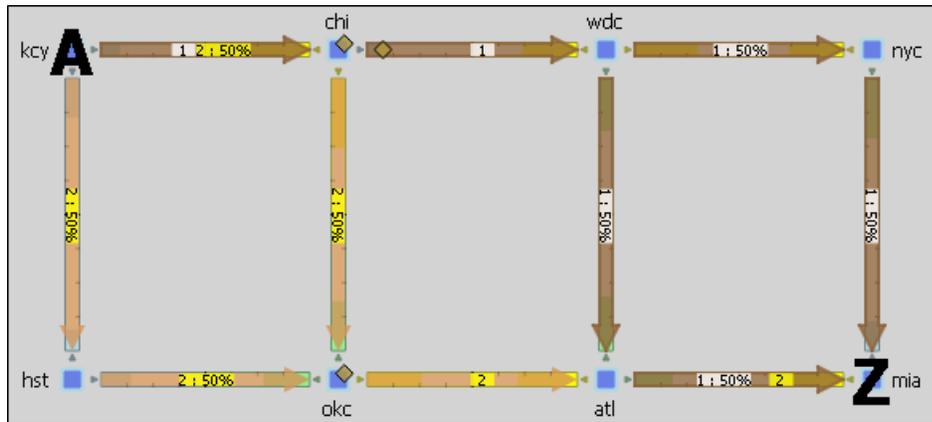


Note

If both the SR LSP and its SR LSP paths have segments, the segments on the LSP paths override those on the SR LSP.

Figure 9-35 shows an SR LSP from kcy to mia that has a primary and a secondary path. The primary path is configured with a chi node segment and chi-to-wdc interface segment, and the secondary path is configured with a node segment in okc. This can be seen by selecting the LSP paths individually from the LSP Paths table.

Figure 9-36 Example SR LSP Paths



382655

SR LSP Routing

Assuming the common LSP mechanisms, such as autoroute and forwarding adjacencies have determined that a demand will enter an SR LSP, the LSP traffic is routed as follows. Note that if the destination node or any of the segments are not reachable, the SR LSP is not established.

**Note**

These descriptions are for the default behavior where SR LSPs take the shortest IGP path. Using the SR-TE LSP Optimization tool, you can optimize SR LSPs to use paths that are based on TE metrics or delays (latencies). You can also use this tool to avoid specified nodes. For information, see the [SR-TE Optimization](#) chapter.

- If the SR LSP does not have a segment list, the demand is routed via IGP to the destination node.
If the SR LSP has a segment list, the demand is routed using the segments within the segment list in the sequential order in which they appear.
- If it is using a node segment, the demand takes the shortest path to that node.
If it is using an interface segment, the demand takes the shortest path of the source node of that interface, and then routes to that interface's destination node.
If it is using an anycast group segment, the demand routes through the node with the shortest path. In case of a tie between multiple nodes in an anycast group, the SR LSP uses ECMP to route.
- If you select both a node segment and an anycast group segment when creating the segment list hop, the node segment has precedence.
- If the SR LSP contains LSP paths, then the demand is routed on the LSP path with the lowest path option for which the demand can be routed from source to destination. Demands for SR LSP paths route in the same manner as described above for SR LSPs.
- If the demand cannot be routed on the segments defined in the SR LSP or its LSP paths (for example, because of a failure in a node segment), the demand is routed to the destination using the IGP shortest path.

Manage Segment Lists

SR LSPs and mesh SR LSPs are created the same way RSVP LSPs are created. The difference is that you must set the Type property to SR. Once created, the next step is to either create the segment list or create LSP paths and their segment lists.

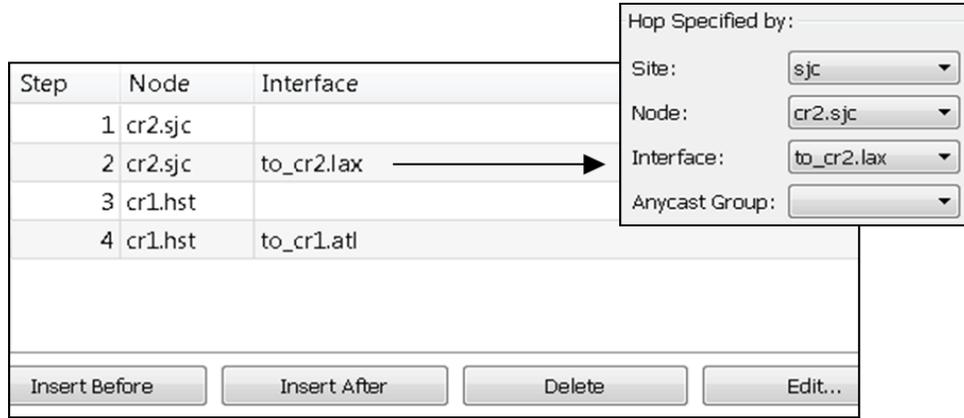
Likewise, LSP paths are created in the same way as RSVP LSP paths. Since you are creating them from SR LSPs, the only editable properties are Path Option and Active.



Step 1 In the LSP Properties dialog box or in the LSP path Properties dialog box, click the Segment tab. Add the segments in the order in which they are to be followed.

- To add a new segment, select an existing segment (if applicable), and then click either the Insert Before or Insert After button to determine where the new segment goes sequentially. Go to step 2.

- To edit an existing segment, select it from the list and then select the Edit button. This opens the Edit Segment List dialog box. Go to step 2.
- To delete an existing segment, select it from the list, select the Delete button, and then press OK. You need not continue.



Step 2 To continue creating or editing a segment, use the following options and then click OK in the Edit Segment List dialog box. As needed, select the site, node, interface, or anycast group.



Note Interface segments must be the egress interface of a node segment.

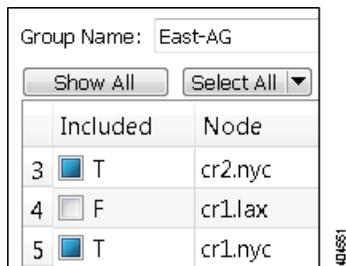
- For node segments, do not select an interface.
- For anycast group segments, do not select a site, node, or an interface.

Once the segment list is created, another way to manage segments is to click one from the Segments List Hops table. From there you can change the site, node, and interface, though you cannot change the sequence.

Create Anycast Groups

Step 1 Follow one of these two steps to open the Anycast Group dialog box.

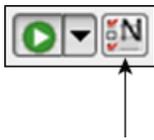
- Select the Insert->Anycast Group menu.
- Right-click in the plot and select New->Anycast Group from the context menu.



- Step 2** Enter a unique name to identify the anycast group.
- Step 3** For each node you want to include in or exclude from the anycast group, click the T (true) or F (false) check box, respectively. Then click OK.

Network Options for MPLS Simulations

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

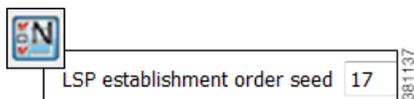


The status bar in the lower, left of the GUI identifies which simulation modes you are currently using. Note that unless specified otherwise, the documentation describes behavior in the IGP and LSP Reconvergence mode.

| For More Information... | See... |
|--|--|
| Using Fast Reroute mode | Fast Reroute Simulations section |
| Using the Autobandwidth Convergence mode | Autobandwidth Simulations section |
| Using the Autobandwidth Convergence Including Failures mode | Autobandwidth Simulations section |
| Ignoring reservable bandwidth constraints on RSVP LSPs for capacity planning | Ignore Reservable Bandwidths for Capacity Planning section |
| Ignoring actual paths and active paths for RSVP LSPs | De-Activate Actual Paths for Simulations section |
| Setting the order in which LSPs are routed | LSP Establishment Order section |

LSP Establishment Order

WAE Design establishes LSPs in the order in which they appear in the plan. Thus, the routing of a specific LSP might depend on previously established LSP routes. You can modify this order by changing a random seed. WAE Design then establishes LSPs in a random order that is completely determined by this number. Although you cannot predict the order based on the number, if you use the same number multiple times, WAE Design establishes the LSPs in the same order each time. Varying the LSP establishment order enables you to check, for example, whether certain orders result in higher utilizations.

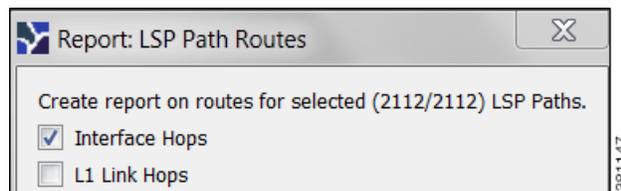


-
- Step 1** Either select the Edit->Network Options menu or click the Network Options icon in the Visualization toolbar.
 - Step 2** Select the Simulation tab.
 - Step 3** In the LSPs section, enter a number in the “LSP establishment order seed” field. The default is 0.
 - Step 4** Click OK.
-

Reports and Diagnostics

LSP Path Reports

You can create reports on the LSP path routes that the LSP is actually taking across both L3 and L1 topologies. Each report identifies the complete path route, as well as disjointness information, including whether it contains common nodes or common circuits. Additionally, the report contains all other columns in an LSP Paths table, though they are initially hidden.



-
- Step 1** Select one or more LSP paths. If you do not select one, the all LSP paths are used.
 - Step 2** Select the Tools->Reports->LSP Path Routes menu.
 - Step 3** Select one or both hop options, and click OK.

After closing the report, you can access it again by selecting the Windows->Reports menu.

To export the report, right-click the Routes report name in the left Reports navigation pane, and select Export Routes.

Simulation Diagnostics

To help troubleshoot RSVP LSP and LSP path simulations, WAE Design analyzes simulation routes and provides reasons for certain types of routing behavior. The Simulation Diagnostics tool identifies why an LSP is not routed, why an LSP is routed away from its actual path, and why an LSP is not following the shortest TE path.

These routing diagnostics assume that all other LSPs, except for the one being tested, have been routed. That is, WAE Design calculates whether an LSP can route on the actual path after all other routed LSPs have had their bandwidth reserved.

Note that running a report overwrites the previous report of the same type. For example, an LSPs diagnostic report would overwrite the previous LSPs report, but not the LSP Path diagnostic report.

Using Simulation Diagnostics to Troubleshoot

In the “Reason” column, information is provided to help you troubleshoot the LSP and LSP path issues.

| Reason | Description |
|---------------------|--|
| Affinities | The affinity settings prevent the LSP or LSP path from routing. |
| Available BW | There is insufficient bandwidth to route the LSP or LSP path. |
| Explicit Hops | The route is determined by a named path that contains one or more explicit hops. |
| Hop Limit | The hop limit is too low. |
| Invalid Actual Path | The actual path is invalid and cannot be interpreted as an LSP route. |
| No Actual Path | The LSP does not have an actual path. |
| No Attempt | The LSP path is not routed because it is not a standby path and it is not the lowest routable path option for the LSP. |
| No Destination | There is no defined destination. The collected network might not contain the destination node of the LSP. |
| Simulation Option | The simulation option for following actual paths is not enabled. |
| TE not enabled | There are interfaces that the LSP or LSP path traverses that are not TE enabled. |
| Topology | The source and destination are disconnected. |

Run LSP Simulation Diagnostics

-
- Step 1** Either select one or more LSPs or LSP paths from their respective tables.
- Step 2** Open the diagnostics information in one of two ways.
- Select the Tools->Diagnostics->LSP Simulation or Tools->Diagnostics->LSP Path Simulation.
 - Right-click the selection in the table and select Simulation Diagnostics.
-