



MPLS Simulation

This chapter describes how WAE Design simulates MPLS routing. All LSPs other than SR (segment routed) LSPs are routed like RSVP LSPs. For MPLS simulation information specific to these types of LSPs, see [RSVP-TE Simulation](#) and [Segment Routing Simulation](#).

- 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.

This section contains the following topics:

- [LSP Types, on page 1](#)
- [Creating and Visualizing LSPs and LSP Paths, on page 2](#)
- [Path Latency Calculations, on page 5](#)
- [Explicit LSP Paths Initializer, on page 5](#)
- [Routing Demands, on page 6](#)
- [Routing Inter-Area LSPs, on page 17](#)
- [Simulation Modes, on page 20](#)
- [LSP Establishment Order, on page 20](#)
- [Reports and Diagnostics, on page 21](#)

LSP Types

- SR LSPs—Segment routing LSPs, which do not use RSVP for routing. You can create SR LSPs using the WAE Design GUI. They are identified with a Type property of SR. For more information, see [Segment Routing Simulation](#).
- RSVP LSPs—LSPs that are established through RSVP. These are commonly known as MPLS TE tunnels. WAE discovers RSVP LSPs, and you can also create them using the WAE Design GUI. They are identified with a Type property of RSVP. For more information, see [RSVP-TE Simulation](#).

WAE Design does not model LDP tunnels as LSPs.

Creating and Visualizing LSPs and LSP Paths

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 choosing **Insert > LSPs > LSP**, or by right-clicking in the plot and choosing **New > LSPs > LSP**. 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 choosing **Insert > LSPs > LSP Mesh**, or by right-clicking in the plot and choosing **New > LSPs > LSP Mesh**. This opens a dialog box for adding a mesh of LSPs between the selected nodes.



When selected, LSPs appear in the plot as a brown arrow. To view the details of LSP routes, as shown in [Figure 1: LSP Plot View of Three LSPs with the Same Source, on page 3](#), follow these steps.

-
- Step 1** In the LSPs table, select one or more LSPs.
- When you select LSPs in the LSPs table, their associated L1 circuit paths are highlighted in the network plot.
- Step 2** Right-click and choose **Plot LSPs**.
- Step 3** To move the plot horizontally or vertically, use the left/right arrows and up/down arrows.
- Step 4** To select all LSPs within the LSP plot view, click **Select LSPs**. To deselect all of these LSPs, click in an empty plot area.
- Step 5** To visualize multilayer L1 paths when plotting LSPs, use the following drop-down options:
- **L3**—Displays the path in terms of interfaces.
 - **L1**—Displays the path in terms of L1 links.

To better align the L1 links, WAE Design checks for L3-L1 links. If they do not exist, WAE Design checks for collocated L3/L1 nodes within sites.

- **L3+L1**—Displays the L3 plot above and the L1 plot below.


WAE Design does not perform alignment between L3 and L1 nodes. When you click an L3 interface, the associated L1 links are highlighted automatically.

What to do next

You can plot one or more LSP paths in the same manner by selecting them from the LSP Paths table.

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

Figure 1: LSP Plot View of Three LSPs with the Same Source



LSP Paths

LSPs can be assigned one or more LSP paths. Like LSPs, LSP paths have properties that vary depending on whether the path is for an RSVP LSP or SR LSP. 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. For information on these properties, see [RSVP-TE Simulation](#) and [Segment Routing Simulation](#).

After an LSP path is created, you can use an initializer to create explicitly routed paths. For information, see [Creating LSP Paths, on page 3](#) and [Explicit LSP Paths Initializer, on page 5](#).

Path Options and Active Path

Each LSP path has a Path Option property. 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, where path option 1 is established first.

LSP Name:	cr1.lon_cr2.fra
Path Option:	2

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

LSP Property	
Active Path:	<div>Not set</div> <div>Not set</div> <div>1</div> <div>2</div> <div>3</div>

Creating LSP Paths

Step 1 In the LSPs table, choose 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

3.8107.8

- Step 2** Right-click in the plot choose **New > LSPs > LSP Paths**, or choose **Insert > LSPs > LSP Paths**.
- 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** If this is an RSVP LSP path, optionally set these properties. For information on these properties, see [RSVP-TE Simulation](#).
- If this is a standby LSP path, check **Standby**, which ensures the paths are always active.
 - To set the bandwidth, specify the **Setup Bandwidth** for the LSP path, or specify that it should inherit the bandwidth from the LSP.
 - To create associated named paths, check this option and complete the name using \$1 for the LSP name and \$2 for the path option.
 - To associate the LSP path with affinities, click **Edit**. 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. Click **OK**.
- Step 5** Click **OK**.


Visualizing 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** submenus, you can choose to see the shortest actual paths, shortest latency paths, or shortest TE paths each time that you choose an LSP.
- To view all shortest IGP paths, shortest TE paths, or shortest latency paths for selected LSPs, use the context menu. Choose one or more LSPs, right-click, and choose **Filter To > Shortest Paths** ([Figure 2: LSP Shortest IGP, Shortest TE, and Shortest Latency Paths, on page 5](#)).

Figure 2: 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



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.
- If two paths have the same latency, paths with the smaller number of hops are preferred.

Explicit LSP Paths Initializer

The Explicit LSP Paths Initializer automates the process of explicitly routing existing LSP paths to follow their currently simulated routes.



Note


Explicit LSP Paths Initializer supports RSVP only (in terms of ABR hop creation).

- For RSVP LSPs, the initializer converts a set of dynamically routed LSPs to explicitly routed LSPs along their current paths.

The initializer creates named paths and named path hops for both primary and secondary paths. The explicit paths are named *<LSP Name> _<Path Option>*. Orphaned named paths are deleted. (Named paths are *orphans* if they are no longer used by any of the LSP paths.) For more information on RSVP LSPs, see [RSVP-TE Simulation](#).

- For SR LSPs, the only use case is if you want the initializer to create a final segment list hop on the destination node. For more information on SR LSPs, see [Segment Routing Simulation](#).

WAE Design writes a report identifying the number of selected LSP paths. To access this information later, choose **Window > Reports**



Step 1 Choose **Initializers > Explicit LSP Paths**, or right-click an LSP and choose **Explicit Path Initializer**.

Step 2 Choose whether to create explicit paths for all LSPs, selected LSPs, or tagged LSPs, and then click **OK**.

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, choose 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


Because 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 ASes 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 [Private LSPs, on page 14](#), for information on tags, see [Plan Objects](#), and for information on class-based forwarding, see [Class-Based Forwarding, on page 8](#).

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 3: Example of Two Parallel LSPs with Equal Loadsharing, on page 8](#) 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%.



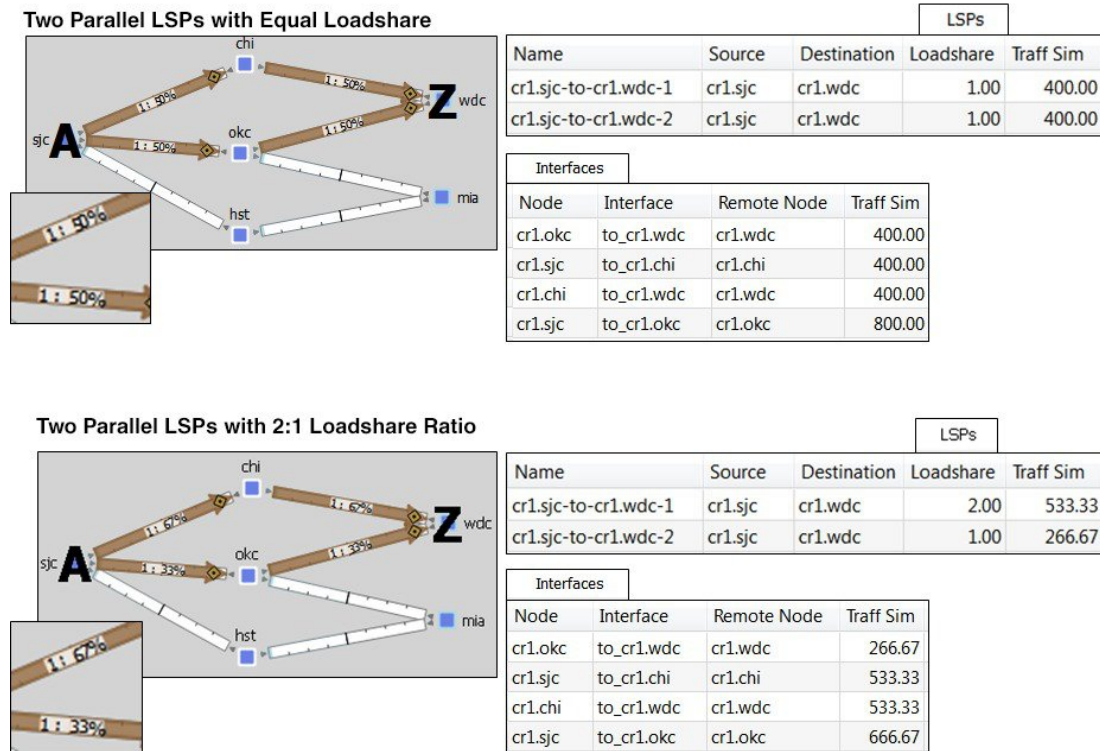
Each LSP has a path showing the percentage of traffic that crosses the LSP as a result of the Loadshare value for that LSP. The format is *<path option> :<loadshare percentage> %*. Although loadshare applies only to path option 1, the LSP itself might be routed along a different path.



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, we recommend that you consider the entire set of parallel LSPs. Note also that because 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 [LSP Loadshare Optimization](#).

Figure 3: Example of Two Parallel LSPs with Equal Loadsharing



381100

Class-Based Forwarding

Class-based forwarding (CBF) lets you 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 example, 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 4: Example Service Class LSP Class Mapping Table

Service Class LSP Class Mapping			
<div> <div>Show All</div> <div>Select All ▼</div> <div>Filter ▼</div> <div>4/4 rows</div> </div>			
	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 4: Example Service Class LSP Class Mapping Table, on page 9). 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 tries to route across LSPs with an LSPClass_BE class (Figure 4: Example Service Class LSP Class Mapping Table, on page 9).

- 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 4: Example Service Class LSP Class Mapping Table, on page 9). 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 5: Example Reroute Around Failure Using IGP Path, on page 11 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 5: Example Reroute Around Failure Using IGP Path

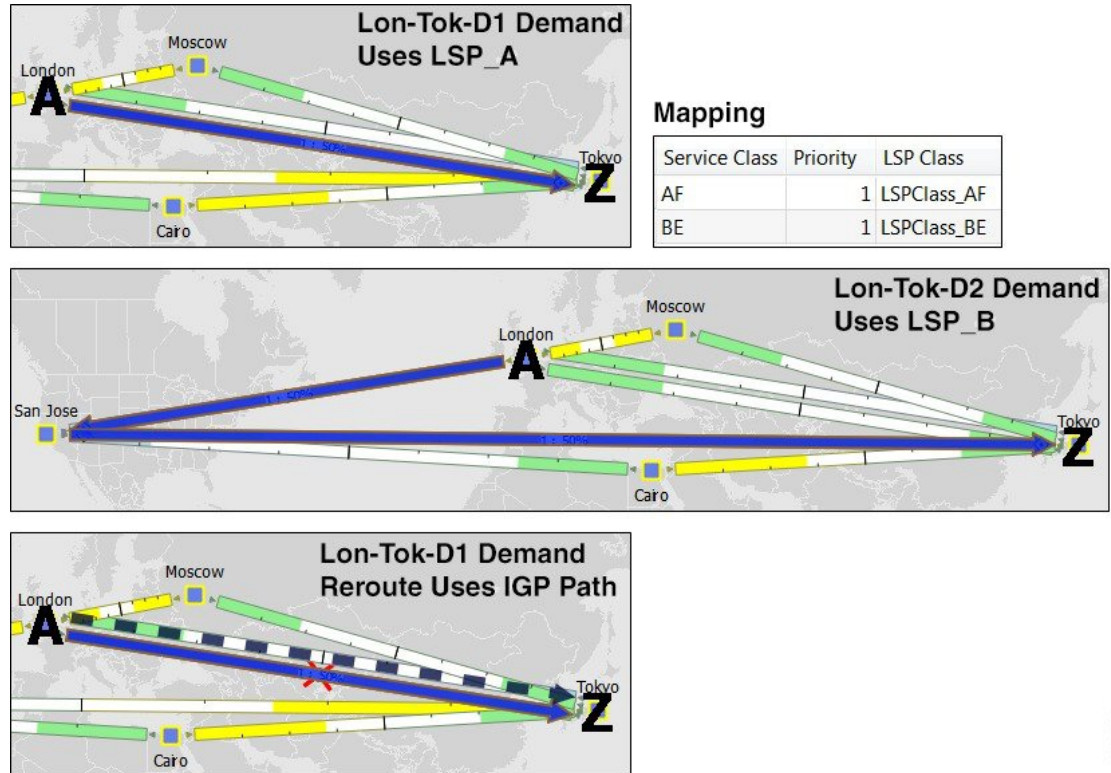


Figure 6: Example Reroute Around Failure Using Alternative LSP Class, on page 11 shows that the AF service class is mapped to both LSPClass_AF and LSPClass_BE. The Lon-Tok-D1 demand now reroutes around a failure between London and Tokyo by using LSP_B. This is because the AF service class now has a Priority 2 mapping to the LSPClass_BE class, and the fact that LSP_B is assigned to this LSPClass_BE class.

Figure 6: 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




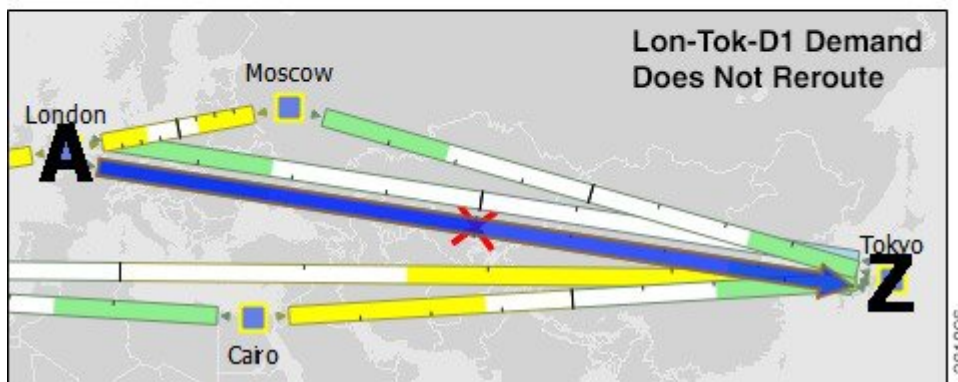
Figure 7: Example of Not Rerouting Based on Drop LSP Class, on page 12 shows that the AF service class is mapped to both LSPClass_AF and a Drop value. The Lon-Tok-D1 demand now does not reroute around a

failure between London and Tokyo. This is because demands using AF service class are mapped so that if an LSP with an LSPClass_AF class is not available, the traffic is dropped.

Figure 7: 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



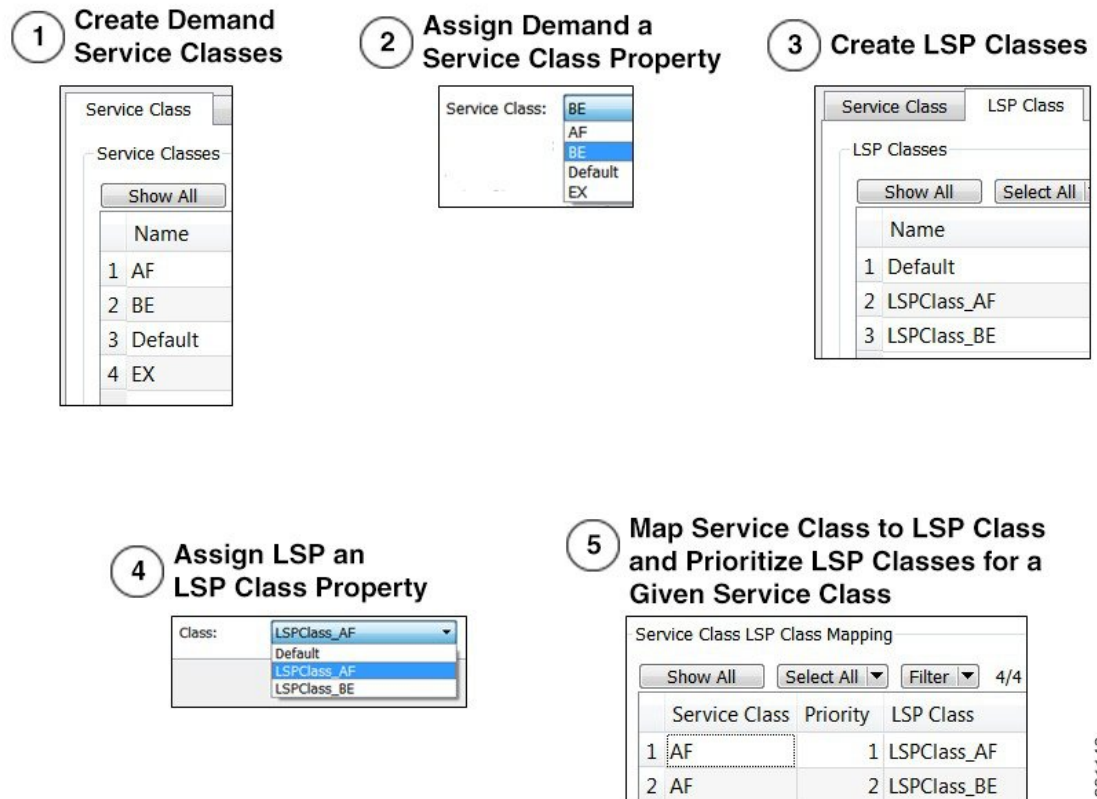
Workflow for Creating Service Class and LSP Class Mappings

The workflow for creating mappings between service classes and LSP classes is as follows ([Figure 8: Workflow Process for Creating Service Class to LSP Class Mappings, on page 13](#)).

-
- Step 1** Create the service classes using the Manage QoS dialog box. For information, see [Quality of Service Simulation](#).
 - 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 choose **Properties**.)
 - Step 3** Create the LSP classes using the Manage QoS dialog box. For information, see [Creating LSP Classes, on page 13](#).
 - Step 4** Assign LSPs to the LSP classes. For information, see [Assigning LSP Classes to LSPs, on page 13](#).
 - 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 [Creating Mappings Between Service Classes and LSP Classes, on page 14](#).
-

What to do next

Figure 8: Workflow Process for Creating Service Class to LSP Class Mappings



381148

Creating LSP Classes

- Step 1** Open the Manage QoS dialog box in one of two ways:
- Choose **Edit > Manage QoS**.
 - Choose **Manage QoS** from the **QoS** drop-down menu in the Visualization toolbar.
- Step 2** Click the **LSP Class** tab.
- Step 3** In the LSP Classes area, click **New**.
- Step 4** Enter the name of the LSP class, and click **OK**.
- Step 5** Click **OK** in the Manage QoS dialog box.

Assigning LSP Classes to LSPs

- Step 1** Choose one or more LSPs.
- Step 2** Right-click one of the LSPs and choose **Properties**.

Step 3 Choose the LSP class from the Class drop-down list and click **OK**.

Creating Mappings Between Service Classes and LSP Classes


Step 1 Open the Manage QoS dialog box in one of two ways:

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

Step 2 Click the **LSP Class** tab.

Step 3 In the Service Class LSP Class Mapping area, 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 routes available LSPs. The lower the number, the higher the priority.

Step 6 Click one of the following radio buttons:

- To map the selected service class to an LSP class, click **LSP Class** and select the class from the drop-down list.
- To set the mapping so the demand doesn't reroute if LSPs are unavailable, click **Drop**. 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 tagged 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.

Creating Private Demands for Existing LSPs


LSPs must currently exist in the plan file.



-
- Step 1** If you are creating demands for only selected LSPs, 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** Choose **Insert > LSPs > Demands for LSPs**, or right-click in the plot and choose **New > LSPs > Demands for LSPs**.
- Step 3** In the LSPs list, choose **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** Check **Mark LSPs as private**.
- Step 7** Click **OK**. The newly created demands are highlighted in the Demands table.
-

Creating Private LSPs for Demands

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



-
- Step 1** If you are creating LSPs for only selected demands, 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** Choose **Insert > LSPs > LSPs for Demands**, or right-click in the plot and choose **New > LSPs > LSPs for Demands**.
- Step 3** In the Demands list, choose **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** Check **Mark LSPs as private**.
- Step 6** Click **OK**. The newly created LSPs are highlighted in the LSPs table.
-

Deleting Demands when Private LSPs are Deleted

You can choose to delete a demand when deleting a Private LSP. By default, when a private LSP is deleted, the corresponding demand is not deleted.

-
- Step 1** Choose **Edit > Network Options**.
- Step 2** Click **Advanced** tab.
- Step 3** Under the Demands section, choose:
- Unrouted if the Private LSPs are unrouted
 - Removed if the Private LSPs are removed
- Step 4** Click **OK**.
-

Configuring 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.

- a) Click the **Tags** tab.

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 **New Tag**.

Enter the name of the tag.

Click **OK**.


- b) Click **OK** to save and exit the Properties dialog box.

Step 2

In the LSP table, select and double-click the LSPs that are to transport demands that were tagged (Step 1).

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

- b) Click **OK**.



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 ASes are of the same type. For information on how nodes are assigned to areas or levels, see [Traffic Demand Modeling](#).

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 or SR, and whether you have selected to require explicit hops at ABRs. (See [Explicit Versus Dynamic Inter-Area LSP Routing](#), on page 18.)

**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: Example Inter-Area Routing Requiring Hops at ABRs, on page 19](#)).
- 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 Versus Dynamic Inter-Area LSP Routing

An OSPF ABR is a node that belongs to both area 0 and other OSPF areas. An IS-IS ABR is a node that belongs to both the Level 2 area and another IS-IS level.

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 example, 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 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: Example Inter-Area Routing Requiring Hops at ABRs, on page 19](#) 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 10: Example Inter-Area Routing Without Requiring ABRs, on page 19](#) shows that by deselecting this network option, the inter-area LSP that does not have hops set at ABRs can dynamically route.

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

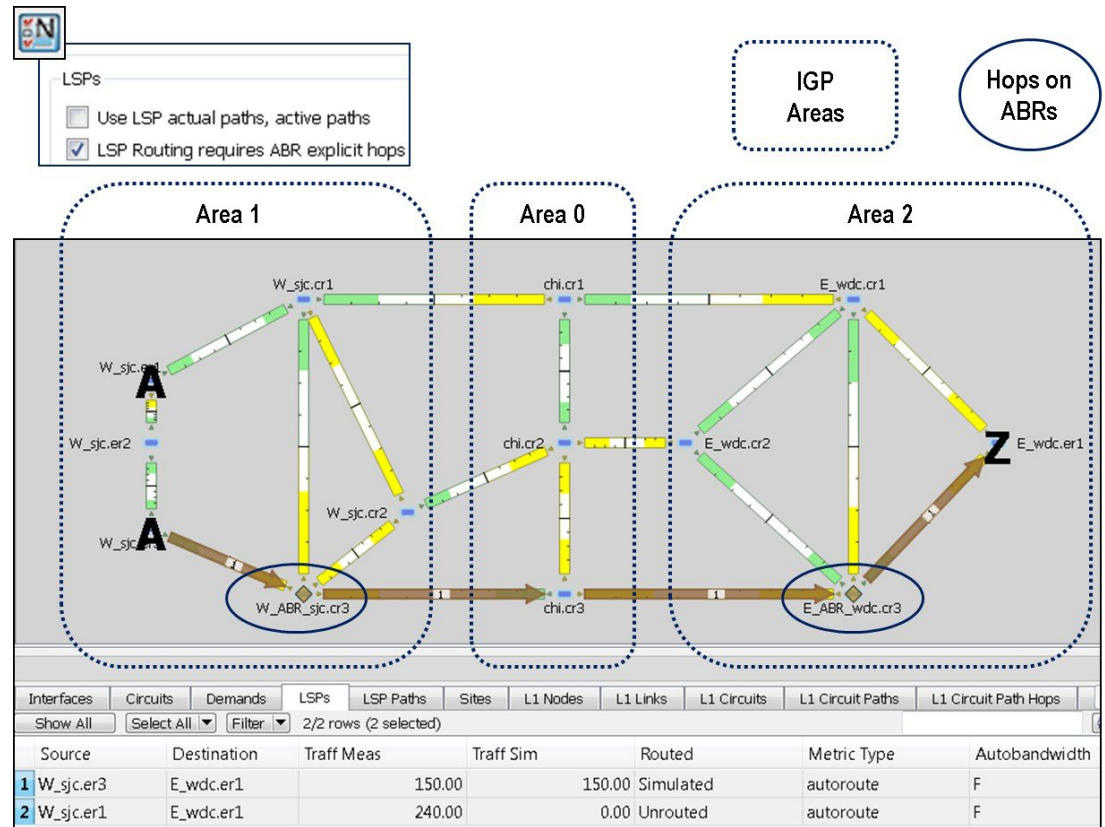
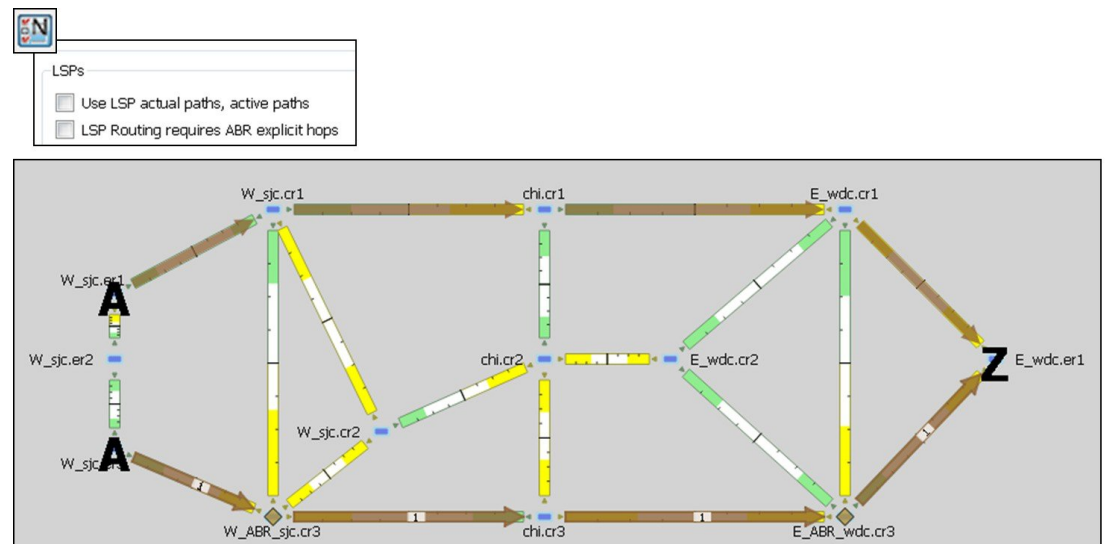



Figure 10: Example Inter-Area Routing Without Requiring ABRs



Simulation Modes

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



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.


Other simulation modes include Fast Reroute (FRR), Autobandwidth Convergence, and Autobandwidth Convergence Including Failures. For information, see [RSVP-TE Simulation](#). For information on the network option for rerouting of non-standby L1 circuit paths under failure, see [Layer 1 Simulation](#).

The Optimization tools only work in IGP and LSP Reconvergence mode. If you try 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.

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.

LSP Establishment Order

WAE Design establishes LSPs in the order in which they appear in the plan. 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 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 lets you check, for example, whether certain orders result in higher utilizations.



-
- Step 1** Choose **Edit > Network Options** or click the **Network Options** icon in the Visualization toolbar.
 - Step 2** Click the **Simulation** tab.
 - Step 3** In the LSPs area, 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, all LSP paths are used.
 - Step 2** Choose **Tools > Reports > LSP Path Routes**.
 - Step 3** Check one or both hop options, and click **OK**.
 - Step 4** To export the report, right-click the **Routes** report name in the left Reports navigation pane, and choose **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

The following information is provided to help you troubleshoot 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.

Reason	Description
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 destination defined. 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.

Running LSP Simulation Diagnostics

Step 1 Choose one or more LSPs or LSP paths from their respective tables.

Step 2 Open the diagnostics information in one of two ways:

- Choose **Tools > Diagnostics > LSP Simulation** or **LSP Path Simulation**.
 - Right-click the selection in the table and choose **Simulation Diagnostics**.
-