



Traffic Demand Modeling

WAE Design uses *demands* to describe the source and destination of a potential traffic flow across a network. A route simulation determines the routes that this traffic takes from the source to the destination, which are determined by the topology, the routing protocols, and the failure state of the network. To model IGP routing, these sources and destinations are nodes or interfaces within the topology. To model basic inter-AS routing, the sources and destinations are neighboring external ASes, peering nodes in these ASes, or interfaces to these peering nodes.

Each demand has a specified amount of traffic. There are several methods for putting this traffic into the demands, including the Demand Deduction tool, which calculates a realistic amount of per-demand traffic based on measured traffic.

The demand traffic is the basis for many of the WAE Design simulation and traffic engineering tools. An accurate set of demands and demand traffic is essential for effective planning, designing, engineering, and operating of a network. Accurate knowledge of the demands is essential for accurate traffic trending and traffic growth predictions.

This chapter describes demands, including how they are created and how their traffic can be estimated.

Use Demands To...	Suggested Steps to Take...
Model discovered networks	<ol style="list-style-type: none"> 1. Create a demand mesh based on where the traffic originates. For example, if all traffic is between edge routers, create a demand mesh between those edge routers. 2. Set the demand traffic by importing it or using Demand Deduction.
Model future usage in the network	<ol style="list-style-type: none"> 1. Create a demand mesh. 2. After setting the traffic, use WAE Design tools for growing the traffic and then analyze the effects on the network. You can import demand growth, you can modify selected demand traffic to emulate growth, or you can use demand groupings and other forecasting tools (described in Traffic Forecasting).
Design networks	<ol style="list-style-type: none"> 1. Create a demand mesh. 2. Set the demand traffic using methods described in Traffic Forecasting.

Use Demands To...	Suggested Steps to Take...
Analyze existing plans	<p>Use a variety of WAE Design tools that rely on demand traffic.</p> <ul style="list-style-type: none"> • Simulation Analysis lets you run what-if analysis. • Manually edit the network topology and configuration to determine the effects on traffic utilization. • Traffic Engineering tools let you control routings and optimize utilizations.

This section contains the following topics:


- [Demands, on page 2](#)
- [Demand Traffic, on page 8](#)
- [Creating Demands and Demand Meshes, on page 15](#)
- [Adding and Modifying Demand Traffic, on page 21](#)

Demands

Because demands determine how traffic is routed through the simulated WAE Design model, creating realistic demands and demand meshes is imperative to the accuracy of other information that can be derived from WAE Design. As such, all defaults are set to create demands and demand meshes that best suit most network models.

Each demand is comprised of unique properties (keys) that define it, other properties, and traffic. The following list summarizes these, though for a complete list of properties, refer to the available columns in the Demands table.

Selected demand paths are blue. An “A” labels the source, and a “Z” labels the destination.



Unique Properties (Keys)	Each demand is defined by a unique combination of these four properties.
	Name—By default, this is blank.
	Source—Nodes, interfaces, external ASes, or external endpoints.
	Destination—Nodes, interfaces, external ASes, external endpoints, or multicast destinations.
	Service class—User-defined classification of traffic, such as for voice or video.
Commonly Used Properties	Latency bound—Policy that sets the maximum permissible latency on a demand under normal operation. This property is used by WAE Design traffic engineering tools.
	Topology—Demands can be assigned to a specific IGP, and only route through interfaces that belong to that IGP.
	Reroutable—Enable/disable the routing of demands around failures. Turning off reroutes around failures might be useful, for example, when simulating Layer 2 traffic.
	Active—Only active demands are routed during simulations.
	Tags—User-defined label that lets you group demands for subsequent processing. Most tools, for example, allow you to choose a set of tags on which to operate.
Traffic	Private LSP—If a demand is associated with a private LSP, the demand can only route through that LSP, and the only demand that is permitted to cross that LSP is this demand. You can associate an existing demand to an existing private LSP. The Private LSP drop-down list shows the private LSP that is currently associated with the selected demand. You can choose a different private LSP, or you can choose None to remove an associated LSP.
	By default, demands have zero traffic, so you must add the simulated traffic to them.
	Demand traffic belongs to the service class of the demand.
	Demand traffic can be set per traffic level.

Demand Sources and Destinations

When creating sources and destinations, follow these recommendations:

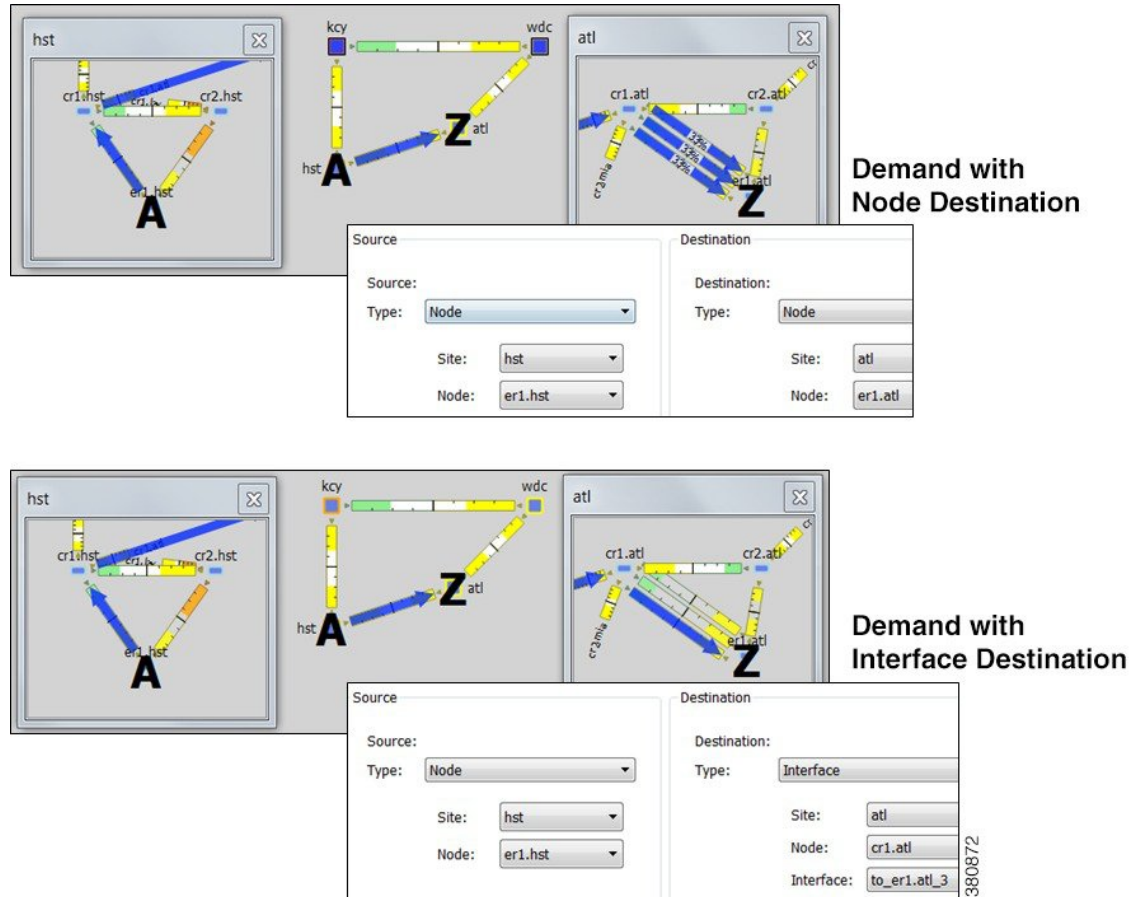
- For internal routing, use nodes.
- For external ASes, use a combination of ASes, nodes, and interfaces. Using interfaces lets you specify the exact interface on which the demand traffic is going into or out of a node.
- For more complex routing where multiple sources or destinations (and multiple failover scenarios) are required, use external endpoints.
- For multicast routing, use multicast destinations.

If multiple interfaces are attached to a node and if a demand is sourced to or destined for that node, the traffic splits across one or more of those interfaces, depending on other properties, such as IGP metrics or BGP policies (on a peering circuit). You can, however, specify just one of those interfaces. [Figure 1: Example of Demands Destined for a Node Versus an Interface, on page 4](#) shows an example of this difference between

using nodes and interfaces. In this example, the interface destination is one of three interfaces going into the er1.atl node.

Note if using an interface as a source of a demand, the source is the inbound interface. If using an interface as the destination of a demand, the destination is the outbound interface.

Figure 1: Example of Demands Destined for a Node Versus an Interface



Demand Meshes

Demand meshes are a time-efficient way of creating numerous demands for all or part of the network. By default, WAE Design creates a source-destination mesh among nodes, interfaces, external ASes, and external endpoints. There are also advanced options, such as the ability to use a different set of destinations to create the demand meshes.

Alternatively, you can import an existing <DemandMesh> table that contains estimated traffic. WAE Design uses these estimations to calculate a full mesh of demands.

Demand Latency Bounds

Each demand can have a latency bound, which is a policy that sets the maximum permissible latency on a demand under normal operation. These can then be used to guide the route selection of the traffic engineering

tools. The Simulation Analysis tool can use these values to determine if latency bounds are violated when worse-case failures occur.

The Demands table has several Latency columns. Key ones are as follows:


- Average Latency—Average latency over all ECMP subroutes.
- Minimum Latency—Minimum latency over all ECMP subroutes.
- Maximum Latency—Maximum latency over all ECMP subroutes.
- Latency Bound—Maximum permissible latency on a demand.
- Min Possible Latency—Total latency of the shortest path that the demand could take.
- % Diff Min Possible Latency—Maximum latency minus the minimum possible latency expressed as a percentage of Min Possible Latency.

Visualizing Demands

To view demand paths in the network plot, select them in the Demands table. Their path highlight is blue. An “A” labels the source, and a “Z” labels the destination. If sites are nested, these “A” and “Z” labels appear in all relevant child sites.

When you select demands in the Demands table, their associated L1 circuit paths are highlighted in the network plot.

Demands are most commonly used to show how traffic reroutes around failures. A dashed line shows the rerouted demand.



353484

Viewing Demand Plots

Demand routes can be complex because they can go through multiple sites that could hide the details of the path. For example, the network plot might not show the individual nodes traversed by the demand, as in [Figure 2: Network View of Selected Demand, on page 6](#). The demand plot offers an easier way to view end-to-end demands, letting you view the hops that the demands take without having to individually select sites to view them. For example, [Figure 3: Demand Plot View, on page 7](#) shows the demand plot view of the same selected demand as in [Figure 2: Network View of Selected Demand, on page 6](#).

-
- Step 1** In the Demand table, select one or more demands.
- Step 2** Right-click and choose **Plot Demands**.
- Step 3** To move the plot horizontally or vertically, use the left/right arrows and up/down arrows.
- Step 4** To select all demands within the demand plot view, click **Select Demands**. To deselect all of these demands, click in an empty plot area.
- Step 5** To visualize multilayer L1 paths when plotting demands, 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

Figure 2: Network View of Selected Demand




Figure 3: Demand Plot View



Viewing Shortest Paths and Shortest Latency Paths

For a description of shortest IGP paths and shortest latency paths, see [IGP Simulation](#).

To View...	Choose...	Example
Shortest IGP path	View > Demands > Shortest IGP Paths	
Shortest latency path	View > Demands > Shortest Latency Paths	

Demand Traffic

Simulated traffic is the amount of traffic a demand is attempting to propagate through the network. For example, demand traffic is used to calculate interface utilizations during simulations. By default, demands have no traffic, and thus there is no simulated traffic. The most complex and powerful method of adding demand traffic is the Demand Deduction tool, which estimates demand traffic from measured traffic values.

Creating and Editing Traffic Levels

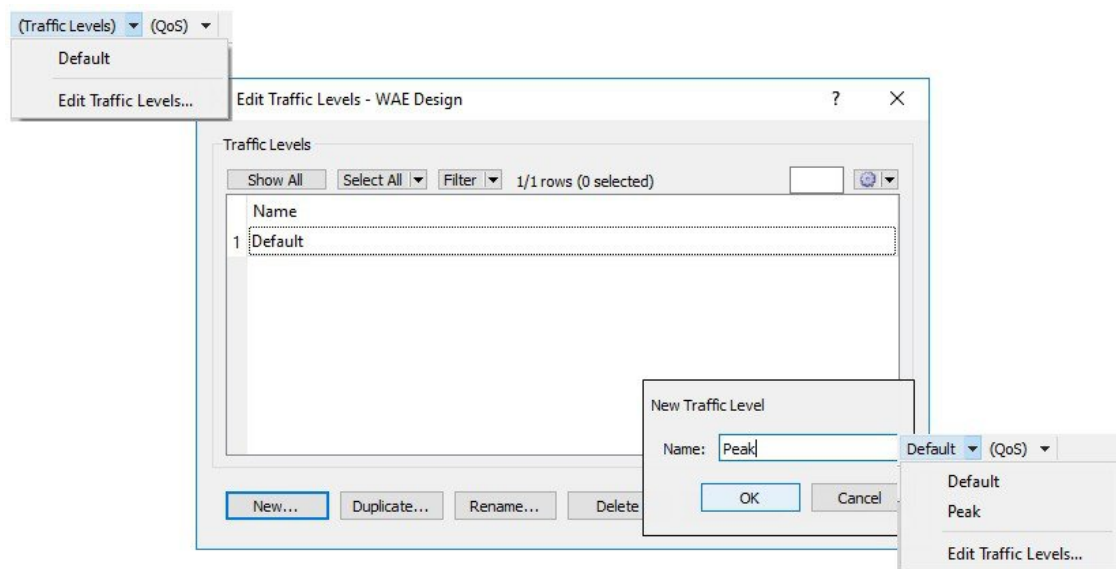
You can associate demands with specific traffic levels, which are identified in the Traffic Levels drop-down list in the Visualization toolbar. The traffic level displayed determines what is displayed in the plot and table calculations. For example, consider a totally static network that has no changes in topology or routing protocols. Throughout a day, the only thing that changes from one plan to the next is the traffic level. Thus, you can store traffic levels for different times of day all in a single plan file. Then, you can view the different states of the network from a single plan file by switching the view between different traffic levels.

You can create and edit traffic levels using the Edit Traffic Levels option ([Figure 4: Create and Edit Traffic Levels, on page 8](#)).

-
- Step 1** From the **Traffic Levels** drop-down list, choose **Edit Traffic Levels**.
 - Step 2** Click **New**, or choose an existing traffic level and click **Duplicate**.
 - Step 3** In the New Traffic Level dialog box, enter the name and click **OK**.
 - Step 4** Click **Done**. The traffic level now appears in the Traffic Level drop-down list.
-

What to do next

Figure 4: Create and Edit Traffic Levels



Demand Deduction

Plan files contain traffic measurements on the discovered network. Traffic can be measured on interfaces, interface queues, and RSVP LSPs, as well as on general traffic flows, such as from LDP LSPs. You can use Demand Deduction to estimate demand traffic based on any of these measurements.

The accuracy and usefulness of the results depend on many factors, including how much measured traffic is available, and of what type. For example, interface measurements are most often available, but LSP measurements might provide more information. The results also depend on the accuracy of the demand mesh and the routing model.

Typically, you only have interface traffic measurements. In this case, the individual demands estimated by Demand Deduction are not necessarily accurate. However, aggregates of demands can be highly accurate. For example, predicting the overall utilization after a failure, a topology change, or a metric change, can be very accurate even if the underlying demands individually are not reliable.

For more accuracy of individual demands, include point-to-point measurements, such as for RSVP LSPs or LDP flows measurements. Also, it is useful to combine different types of measurements together for use in Demand Deduction. Interface measurements are generally the most accurate measurements available, and if included in a Demand Deduction, can correct for missing or inaccurate LSP or flow measurements.


Note that you can also use Demand Deduction to set Traffic Balance (%) values for external endpoint members that are set to a Deduce Traffic type. See [Advanced Routing with External Endpoints](#).

Example Demand Deduction

This example demonstrates results when using the Demand Deduction tool on a simple network. [Figure 5: Network Containing Two Demands and No Demand Traffic, on page 10](#) shows the routes of two demands in a network. These demands split between the two parallel core circuits due to an ECMP, and they have a common routing until the last hop. The Traffic column in the Demands table shows 0 because these demands do not yet contain traffic. [Figure 6: Measured Traffic View and Interfaces Associated with the Demands, on page 11](#) shows the Measured Traffic view and the five interfaces associated with the two demands, three of which have measured traffic.

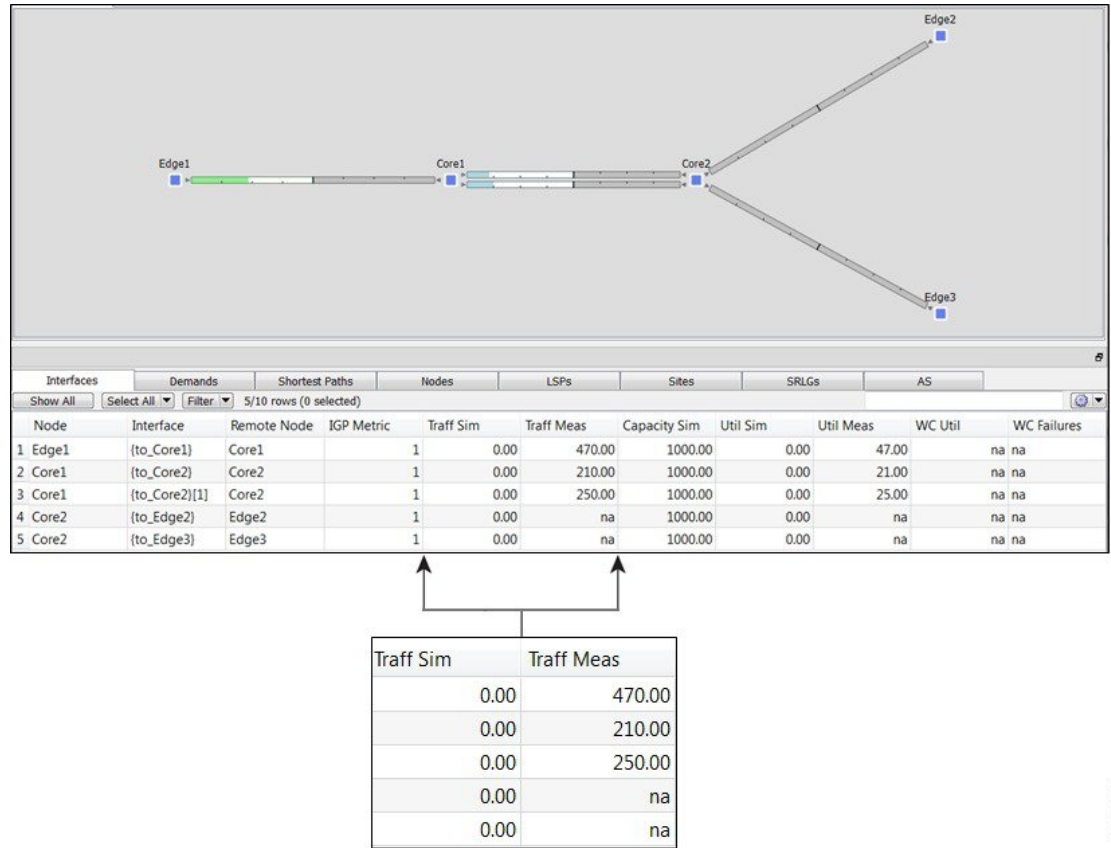
- Edge1 to Core1 has 470 Mbps of measured traffic.
- One Core1 to Core2 interface has 210 Mbps, while the other has 240 Mbps, for a total of 460 Mbps. This unequal split is due to imperfect load balancing of the ECMP.
- There is no traffic from Core2 to Edge2 or from Core2 to Edge3.

Figure 5: Network Containing Two Demands and No Demand Traffic



699096

Figure 6: Measured Traffic View and Interfaces Associated with the Demands



380870

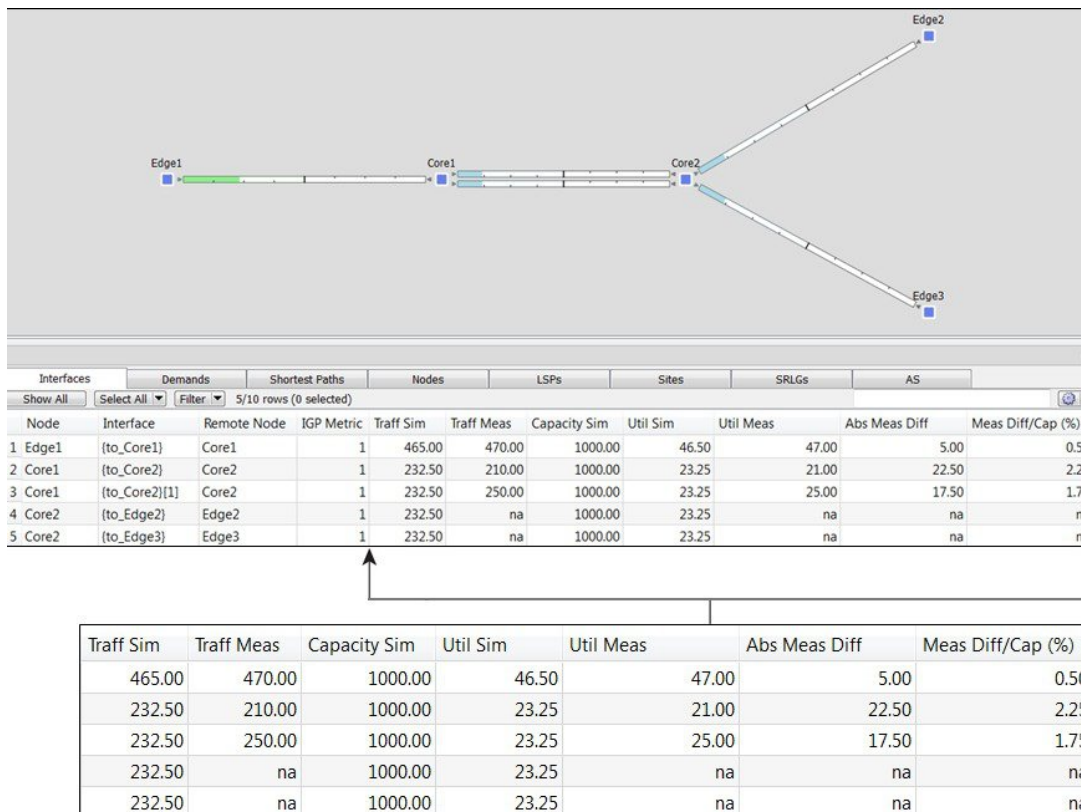
Upon running Demand Deduction with its default options, the Simulated Traffic view appears. Other than the measured interface traffic, there is no other information about the demand traffic. So Demand Deduction first splits the difference between the measured 470 Mbps of traffic (Edge1 to Core1) and the measured traffic of 460 Mbps (Core1 to Core2) to get an estimated total demand traffic of 465 Mbps. In the absence of any other information, it divides this 465 equally to give 232.5 Mbps of traffic to each demand (Figure 7: Simulated View Showing Demand Traffic, on page 12). In the Interfaces table, the Traff Sim column now has values and the network plot shows simulated traffic percentages on all five interfaces associated with the demands.

- Edge1 to Core1 has 465 Mbps of simulated traffic.
- Both Core1 to Core2 interfaces have 232.50 Mbps.
- Core2 to Edge2 and Core2 to Edge3 both have 232.50 Mbps.

The Abs Meas Diff and Meas Diff/Cap (%) columns in the Interfaces table show mismatches between measured and simulated values.

- Edge1 to Core1 has a difference of 5 Mbps, or 0.5%.
- One Core1 to Core2 has a difference of 22.5 Mbps, or 2.25%, while the other has a difference of 17.5%, or 1.75%.
- Neither the Core2 to Edge2, nor the Core2 to Edge3 interfaces have values because they had no measured traffic.

Figure 7: Simulated View Showing Demand Traffic



In this same example, if the Core2 to Edge2 interface had 50 Mbps traffic, the results would have been different. Because this interface is used only by the one demand, the measured 50 Mbps of traffic would be used as an estimate only for that one demand. Using the same logic as before, the demands should total to 465 Mbps, so the other demand is set to the difference, which is 415 Mbps.

Differences in Measured and Simulated Traffic

Demand Deduction relies on accurate topologies, demand meshes, and traffic measurements. These can affect the results of the traffic simulated in the demands and cause simulated traffic to differ from the measured traffic, thus affecting the accuracy of WAE Design simulations. You can see how close these values are by showing the Abs Meas Diff and Meas Diff/Cap (%) columns in the Interfaces column.

Abs Meas Diff—The difference between measured traffic (Traff Meas) and simulated traffic (Traff Sim).

Meas Diff/Cap (%)—The absolute measured difference expressed as a percentage of capacity.

If these columns show large values, one of the following situations likely exists:

- Inaccurate measurements—Different measurements, for example of traffic through different interfaces, can be made at slightly different points in time. Fluctuations in traffic levels might take place between the times that measurements are being taken. This means that measurements could be inconsistent with one another. Usually, these inconsistencies are small and do not seriously affect the Demand Deduction results.

- Insufficient measurements—There are typically many more demands in a network than measurements, and many solutions will fit the observed data well. Demand Deduction chooses between possible solutions using knowledge of typical behavior of point-to-point traffic.
- Incorrect network configurations—If the network topology is incorrect in the plan file, the simulated routes would naturally be incorrect and measurements would not be adequately interpreted.
- Unbalanced ECMP—ECMP hashing can result in imperfect load balancing. Demand Deduction, however, distributes traffic evenly across ECMPs.
- Static routes—WAE Design does not model static routes. If these are present, demands routes might be simulated incorrectly, leading to deduction errors.
- Incomplete demand meshes—Demand meshes do not contain nodes even though traffic is routed between those nodes.
- Inappropriate priorities—In the Demand Deduction dialog box, you have the option to set the priority for calculations as 1 or 2. WAE Design first uses the measurements identified as Priority 1 to calculate the demands. Therefore, if the priority settings do not match the consistency of the traffic measurements in the network, the simulated traffic measurements will be less than optimal.

Additionally, Demand Deduction displays warnings for misleading or undesirable results:

- AS “(AS Name)” contains both dynamic LSPs and interface traffic. Interface traffic in AS has been ignored.

Routing of dynamic LSPs is nondeterministic. So it is not possible to make use of both measured interface traffic and measured dynamic LSP traffic for LSPs that may (or may not) traverse these interfaces. If the network contains an AS with both dynamic LSPs and interface traffic, this warning is issued and the interface traffic is not used.

- Some interface measurements exceed capacities by as much as (percent).

This warning is issued if a specified measurement exceeds the corresponding circuit capacity.

Minimizing Differences Between Measured and Simulated Traffic

Demand Deduction estimates demands that predict interface utilizations under incremental changes to the topology, for example failures, metric changes, or design changes, such as adding a new express route. If interface measurements alone are available, you might choose to fine-tune the Demand Deduction calculations to get better results, such as for site-to-site traffic. To enhance the accuracy of Demand Deduction results, consider the following suggestions:

- Include RSVP LSP or LDP measurements in the network discovery process.
- Restrict demand meshes to exclude demands that are known to be zero. For example, if you know that core nodes do not source traffic, then exclude core nodes when creating the demand mesh.
- Check the Nodes table to see if there is a node where the measured traffic going into it (Dest Traff Meas) and out of it (Source Traff Meas) are very different. Ensure these nodes are included in the demand mesh because they are either sources or destinations for traffic.
- In the Demand Deduction dialog box, always set the most consistent measurements to a Priority 1. The most reliable measurements are usually interface measurements. Likewise, LSP measurements are end-to-end, and thus also generally highly reliable. You can set multiple measurements to priority 1.

For example, if the flow measurements are inconsistent and the interface measurements are very consistent, then interfaces should be set to Priority 1 and the flow measurements to Priority 2.

- The Use Selected option in the Demand Deduction dialog box calculates demands for the selected rows in the Demand table. If you know beforehand the traffic for certain demands, you can set the traffic for those demands and then run Demand Deduction on just the unknown demands. Select the Use Selected—Fix Remaining option to fix the traffic of the remaining demands. A typical example is multicast demands, which often have a fixed, known bandwidth. (You can identify multicast demands in the Cast column of the Demands table.)
- If only a few measurements are available or if there are many inaccurate measurements, the tool sometimes estimates more traffic in a circuit than its capacity. To prevent this, in the Demand Deduction dialog box, select the option to keep the interface utilization below 100%. This forces the resulting simulated calculations to be below the given percentage of circuit capacity.

Flow Measurements in Demand Deduction

Besides node, interface, and LSP traffic measurements, Demand Deduction allows more general flow measurements to be used. These flow measurements can be flows from (or through) a specified node, to (or through) another node. Measurements can also be combinations of these node-to-node flows. This measurement format can be used to enter, for example, peer-to-peer flow measurements, or traffic measurements obtained from LDP routing or from NetFlow.

Flow measurements are entered in the plan file in the <Flows> table, and appear in the GUI in the Flows table. [Table 1: Flows Table Columns](#), on page 14 lists some of the more useful columns in the Flows table. Note that exactly which traffic is included is defined in the FromType and ToType columns.

Table 1: Flows Table Columns

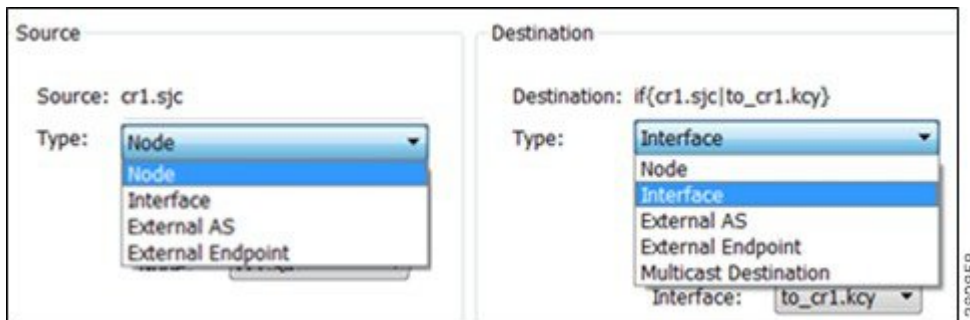
Column	Description
From	Source node.
FromType	<ul style="list-style-type: none"> • Source—Traffic originating at the From node is included in the flow. • Interior—Traffic is included that passes through the From node, entering that node from another node in the same AS • Border—Traffic is included that passes through the From node, entering that node from another node in a different AS.
To	Destination node.
ToType	<ul style="list-style-type: none"> • Dest—Traffic that is destined for the To node. • Interior—Traffic that passes through the To node to another node in the same AS. • Border—Traffic that passes through the To node to another node in a different AS.
Traff Meas	Measured traffic used by Demand Deduction in its calculations. If more than one node is included in either the From or To columns, this measurement is the sum of the traffic over all flows between individual pairs of From and To nodes.

Creating Demands and Demand Meshes

Creating Demands

All selections and entries are optional except for identifying the source and destination.

Step 1 Choose **Insert > Demands > Demand**, or right-click in an empty plot area and choose **New > Demands > Demand**.



Step 2 Enter a demand name. The default is to create the demand without a name.

Step 3 Specify the source as a node, interface, external AS, or external endpoint.

Step 4 Specify the destination as a node, interface, external AS, external endpoint, or multicast group.

Step 5 Choose the service class. If there are no service classes, the demand operates on a service class named *Default*.

Service Class: Default

Latency Bound:

Topology: (No change)

Private LSP: (None)

Active Reroutable Require LSP

Step 6 Enter a value for the latency bound.

Step 7 Choose a topology to restrict the demand routes only to interfaces or LSPs belonging to that topology. The default is unrestricted routing.

Step 8 Retain the **Active** default to include the demand in WAE Design simulations, or uncheck **Active** to exclude this demand from simulations.

	Traffic Level	Traffic	Growth %
1	Default	na	
2	Morning	1300.00	
3	Peak	2020.00	

380857

- Step 9** For each traffic level, specify the amount of traffic or leave it empty for Demand Deduction to complete. You can also create a traffic level if one does not exist (click **New Traffic Level**).
- Step 10** For each traffic level, specify the amount of growth rate you want to use for forecasting purposes. For more information, see [Traffic Forecasting](#).
- Step 11** Click **OK**.

Duplicating Demands

- Step 1** Choose one or more demands.

Duplicate selected demands (3/96)

Number of new demands: 1

Share current traffic evenly


380867

- Step 2** Right-click one of them and choose **Duplicate**.
- Step 3** Enter the number of new demands you want to create per selected demand. For example, if you have three demands selected and you enter 2, each of the three demands is duplicated twice for a total of six new demands.
- Step 4** To evenly distribute the traffic of each demand being duplicated, select “Share current traffic evenly.” If this is unselected, the amount of traffic on the newly created demand is replicated. For example, if the original demand has 1000 Mbps of traffic and you share it with one newly created demand, the resulting traffic is 500 Mbps per demand. If you do not share it, both the original and the new demand have 1000 Mbps of traffic.
- Step 5** Click **OK**. Duplicated demands are named sequentially.

Creating Demand Meshes

Step 1

To restrict the demand mesh source to a specified set of nodes, external ASes, or external endpoints, select them before opening the Demand Mesh dialog box.




Step 2

Choose **Insert > Demands > Demand Mesh**, or right-click in the plot and choose **New > Demands > Demand Mesh**.

Step 3

Create a demand mesh based on the current plan file or import a Demand Mesh table.

- If creating a demand mesh based on the current plan file, select one or more sources from the Source drop-down lists. Options include using the selected objects (in Step 1) as the source, which is the default. Another option is to use all nodes, external ASes, or external endpoints as the source, which is the default if none were selected. The remaining option is to select the source based on tags. (Tags are set in an object's Properties dialog box.)
- If importing demands from a table, select the Import from Demand Mesh Table option. Then browse to the desired table. Note that this table specifies a mesh of sources and destinations, rather than creating demands. For information on how to create this table, see the *Cisco WAE Design Integration and Development Guide*.



Step 4

Enter a demand name. The default is to have no name to prevent large numbers of demands using the same name from being created. The names are useful if needing to identify a specific area of the network, such as a VPN. However, not using demand names helps ensure you do not create a large number of demands that all have the same name.

Step 5

Choose a service class. The default is the **Default** service class.

Step 6 Choose a topology. The default is that demands are applied to all topologies.

Step 7 If you want to create demand traffic immediately upon creating the demands, check **Edit demand traffic after insertion**. For information on entering traffic in the dialog box that appears, see [Generating Demand Traffic Using Node Traffic Estimates, on page 22](#). Note, however, if you have measured traffic in the plan file, a more accurate means of populating the demands with traffic is to use the Demand Deduction tool. For information, see [Demand Deduction, on page 9](#).

Step 8 If you want to delete all existing demands before new ones are created, check **Delete existing demands with same name**. The default (unselected) is to keep the existing demands and add only new ones.

Step 9 For any of the following options, click **Advanced Options**, make changes, and click **OK**.

- Destination—Choose this option if you want to create demands to destinations other than what has been selected as the source.

- Also create demands from destination to source—Uncheck this option if you want the demands created in only one direction. This applies only if you have selected a different set of destinations.
- Use interface endpoints to/from external AS nodes—When creating demands for external ASes, use a source/destination type of interface, and create a demand for all interfaces connected to each node in the external AS. For information on AS relationships and routing policies, see [BGP Simulation](#).
- Respect AS Relationships—If checked, keep the existing AS relationships defined by the Routing Policy (default). If unchecked, recreate the AS relationships. The Routing Policy property is defined in the Edit AS Relationships dialog box. For information on AS relationships and routing policies, see [BGP Simulation](#).
- Respect External Mesh Settings—If checked, keep the existing External Mesh settings defined for external AS meshes (default). If unchecked, recreate the external AS mesh. The External Mesh property is set in the AS Properties dialog box.
- Include Demands to Self—Creates demands that have the same source and destination node (default).

Step 10 Click **OK**.

Creating Demand Meshes for LSPs

Step 1 If applicable, choose LSPs over which you want to run the demands.

The screenshot shows a dialog box with the following elements:

- LSPs:** A dropdown menu currently showing 'All'.
- Service Class:** A dropdown menu currently showing 'Default'.
- Set Demand Traffic to:** A section containing three radio buttons:
 - LSP Setup BW
 - LSP traffic measurements
 - Zero
- Mark LSPs as private.

Step 2 Choose **Insert > LSPs > Demands for LSPs**, or right-click in the plot and choose **New > LSPs > Demands for LSPs**.

Step 3 Choose which LSPs to use: All, those selected, or tags.

Step 4 Choose the service class for the resulting demands.

Step 5 Choose the traffic for the newly created demands.

- Traffic equal to the LSP setup bandwidth
- Traffic equal to the LSP measurements
- Zero, which is appropriate if you need to insert demand traffic in other ways, such as using Demand Deduction, importing it, or manually modifying it.

Step 6 To remove the restriction of setting these demands to only these LSPs, uncheck **Mark LSPs as private**. Otherwise, the default is to restrict these LSPs so that they use only the resulting demands.

Step 7 Click **OK**.

Creating Demand Latency Bounds

You can set demand latency bounds to fixed values or to the following relative values ([Table 2: Editing Demand Latency Bounds](#), on page 20). All values are in ms (milliseconds).

- To set a fixed value or to remove demand latency bounds, use either the Properties dialog box or the Initializer tool.
- To set a relative value, use the Initializer tool.

Set latency bound for selected demands (95 / 95):

Remove

Fixed Value ms

Relative

value: % above shortest path

minimum: ms

maximum: ms above shortest path

380873

- **Minimum**—Minimum possible latency. Any latency bounds that are less than this value are set to exactly to it. The default is 20 ms.
- **Maximum __ ms above shortest path**—Maximum possible latency. Any resulting latency bounds that are more than their minimum possible latency are set to exactly to this value above their minimum possible latency. The default is 35.
- **Value __ % above shortest path**—The latency bound is set to a percentage over the minimum possible latency for that demand. The default is 50%.

For example, if the minimum possible latency is 20 ms and this maximum is 35 ms, if a demand's latency were to go up to 70 ms, it would be reset to 55 ms.

For example, if this value is set to 50% and if the minimum is 20 ms, then two demands with shortest paths by latency of 10 ms and 16 ms would have bounds set to 20 ms and 24 ms, respectively.

Table 2: Editing Demand Latency Bounds

Using Demands Properties Dialog Box	Using the Initializer Tool
<ol style="list-style-type: none"> 1. Select one or more demands in the Demands table. 2. Right-click a demand and choose Properties. 3. To set a fixed value for the latency bound, enter a value in the Latency Bound field. To delete a latency bound, delete the text in this field. 4. Click OK. 	<ol style="list-style-type: none"> 1. If applicable, select one or more demands in the Demands table. If you do not select demands, the result applies to all demands in the plan file. 2. Choose Initializers > Demand Latency Bounds. 3. You have several options: To delete latency bounds, click Remove. To set a fixed value, enter this number in the Fixed Value field. To set a relative value, click Relative and enter values. 4. Click OK.

Adding and Modifying Demand Traffic

Estimating Demand Traffic Using Demand Deduction

Demand Deduction calculates demand traffic when traffic measurements are available.

The options available can significantly affect the calculations. For information on improving accuracy of results, see [Minimizing Differences Between Measured and Simulated Traffic, on page 13](#). For information on setting up external endpoint members to be included in Demand Deduction calculations, see [Advanced Routing with External Endpoints](#).

Demand Deduction

Estimate demand traffic based on measurements.

Traffic

Traffic Level: Selected: Default

Edit Traffic...

Measurements (measurement/total)

<input checked="" type="checkbox"/> Node Source and Dest	(66/66)	Priority 2
<input checked="" type="checkbox"/> Interfaces	(99/100)	Priority 1
<input type="checkbox"/> LSPs	(0/0)	Priority 2
<input type="checkbox"/> Flows	(0/0)	Priority 2

Fitting Parameters

Spread measurement errors evenly throughout network
(For operational use, once modelling is correct)

Concentrate errors in fewer places
(Slower, used for identifying measurement/modelling errors)

Keep interface utilization below 100 %

Output

Demands: Use existing

Fix multicast demands

Remove demands with zero traffic

OK Cancel

Step 1 If applicable, choose one or more demands.

Step 2 Choose **Tools > Demand Deduction**.

Step 3 Choose the traffic level on which you are creating simulated traffic. The measurements used come from this traffic level, and the simulated traffic for each demand is placed in this traffic level.

- Step 4** If needed, modify traffic measurements for one or more demands by clicking the Edit Traffic button. In the dialog box that appears, you can modify measured traffic of interfaces, nodes, LSPs, flows, multicast flows, and/or ports. You can further change these measurements for a traffic level and interface queue.
- Another option in this dialog box is to enter growth percents for use with Forecasting tool. For information, see [Traffic Forecasting](#).
- Step 5** Identify one or more types of measurements used in the calculations: nodes (source and destination), interfaces, LSPs, and flows.
- Step 6** For each type, set its priority. Select Priority 1 for high priority and Priority 2 for lower priority. You can have multiple measurements of the same priority. Like priorities are calculated simultaneously with equal consideration for the measurements.
- By default, all available measurements in the selected traffic set are used, and the interface measurements have priority over node, LSP, and flow measurements.
- Step 7** Choose the **Fitting Parameters**.
- Step 8** If you need to keep the traffic utilization below a different percentage than 100% (default), check **Keep interface utilization below** and enter a value.
- Step 9** Choose the demands for use in constructing the demand calculations.
- Use existing—Calculates demands using the existing demands only. This option is useful when simulating a pattern of demands that cannot be represented as a simple mesh between nodes. If you did not select one or more demands before opening this dialog box, use this option.
 - Use selected—Calculates demands for the selected rows in the Demand table. This option is helpful when you want to recalculate some of the demands, for example, such as a VPN submesh.
- Ignore Remaining—Does not include unselected demands in its calculations (that is, it treats those unselected demands as zero).
- Fix Remaining—Includes unselected demands in its calculations without changing their values. For example, if you select 6 out of 10 demands, then those 4 unselected demands would remain as they are, but would be included in the calculations.
- Step 10** Determine whether to fix multicast demands. If selected, the multicast demands are fixed at their current traffic value.
- Step 11** Determine whether to remove demands with zero traffic. The default is to remove them because Demand Deduction typically estimates a significant percentage of the simulated traffic to be zero when a large number of point-to-point utilizations in a mesh are extremely small. Using this default can substantially improve simulation and optimization performance in large plans. Do not remove demands with zero traffic if all demand routes are of interest, irrespective of traffic.
- Step 12** Verify your parameters, and click **OK**. Demand Deduction calculates the simulated traffic and lists the results in a Demand Deduction report.

Generating Demand Traffic Using Node Traffic Estimates

When a live network cannot be measured, you can estimate the demand traffic by editing the demand mesh. This is a simple way to create traffic in a mesh of demands by specifying the total source and destination traffic per node using one of the following algorithms. For traffic mesh options used in more sophisticated growth estimates, see [Traffic Forecasting](#).

- Gravity Model—This estimation method allocates total traffic among the demands with roughly equal weights.
- Proportional to current traffic—This estimation method allocates total traffic among the demands in proportion to simulated traffic on existing demands.

After you manually enter total traffic data in the dialog box, you can export it to a file for later use.

Step 1 Open the Demand Mesh Traffic dialog box in one of two ways:

- Choose **Edit > Traffic > Demand Mesh**.
- When creating demand meshes, choose **Edit traffic after insertion** from the Insert Demand Mesh dialog box.

Step 2 Each node used in a demand is listed. The SrcDest column identifies whether it is a source, a destination, or both.

For each node, enter the traffic measurement on which you want the estimated demand traffic to be based. You can specify different values for source and destination traffic.


Step 3 You can set demand traffic based on an estimation model, a fixed value, or randomly based on a range.

- To base the demand traffic on an estimation model, click **Gravity Model** or **Proportional to current traffic**.
- To change the demand traffic to a fixed amount or on range of random values, follow these steps.



- Choose one or more demands.
- Right-click in a SrcTraffic or DestTraffic field, and choose **Set Value**.
- In the dialog box that appears, click **Fixed** or **Random**. If you choose Random, enter the minimum and maximum traffic values.
- Click **OK**.

Step 4 Choose the traffic level to which you want to apply the calculations, and then click **OK**.



Modifying Demand Traffic

You have numerous options for modifying demand traffic, all from the same dialog box. The changes that you make apply to the selected demands for the current traffic level.

You can modify demand traffic to either fixed values or to the following relative values.


- To set a fixed value, use either the Properties or the Modify Demand Traffic dialog box ([Table 3: Modifying Demand Traffic](#), on page 25).
- To set a relative value, use the Modify Demand Traffic dialog box ([Figure 8: Modify Demand Traffic](#), on page 25 and [Table 3: Modifying Demand Traffic](#), on page 25). Following are the available options. Except for the percentage option, all values are in Mbps.
 - Change traffic by a specified percentage. Positive percentages add to the traffic, and negative percentages subtract. For example, if the traffic were 1000 Mbps and you entered -10 , the traffic would be reduced to 900 Mbps.
 - Add a set amount of traffic spread over all demands in proportion to their current traffic. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you added 50 Mbps proportionally, one would have 1016.67 Mbps and the other would have 2033.33 Mbps.
 - Add a set amount of traffic uniformly to all the demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you added 50 Mbps uniformly, one would have 1025 Mbps and the other would have 2025 Mbps.
 - Set traffic to a fixed value.
 - Set traffic to a specific value that is spread proportionally over all demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you set them to 4000 Mbps proportionally, one would have 1333.33 Mbps and the other would have 2666.67 Mbps.
 - Set a specified amount of traffic, in Mbps, uniformly to all the demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you set them to 4000 Mbps uniformly, they would both be 2000 Mbps.

Table 3: Modifying Demand Traffic

Modify Fixed Demand Traffic	Modify Fixed or Relative Demand Traffic
<ol style="list-style-type: none"> Choose one or more demands in the Demands table. Right-click one of these demands and choose Properties. For each applicable traffic level, enter the desired amount of simulated traffic in the Traffic field. If needed, click New Traffic Level to create a new traffic level. Click OK. 	<ol style="list-style-type: none"> Choose the traffic level you want the changes to apply to from the Traffic Level menu in the Visualization toolbar. If applicable, choose one or more demands in the Demands table. If you do not choose demands, the result applies to all demands in the plan file for the selected traffic level. Open the Modify Demand Traffic dialog box in one of two ways: Right-click a demand and choose Modify Demand Traffic. Choose Initializers > Modify Demand Traffic. Choose one of the options identified in the preceding list and enter a value accordingly. Click OK.

Figure 8: Modify Demand Traffic

Modify Fixed Values



Modify Fixed or Relative Values

