MATE Design User Guide

Release 6.1
February 2015

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide.
Addresses, phone numbers, and fax numbers
are listed on the Cisco website at
www.cisco.com/go/offices.

Text Part Number:

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

CONTENTS

CHAPTER 1

Overview

MATE Design is standalone software that can run on multiple platforms, or on standard laptops or desktops. It can be accessed in three ways: through a GUI, web browser (MATE Design Archive application only), and CLI. This guide describes how to use the MATE GUI.

Note This guide references `$CARIDEN_HOME`, which is the directory containing the MATE software. On Linux, the default is `/opt/cariden/software/mate/current`.

Applications

MATE Design enables you to model, simulate, and analyze failures, design changes, and impact of traffic growth, as well as optimize your network for maximum efficiency. For instance, with MATE Design models and tools, you can answer a wide variety of questions about traffic management.

Planning

- Where is traffic going in the network?
- When and where will my network run out of capacity?
- How do I convert my growth forecasts into upgrade plans?

Engineering

- How can I better distribute traffic throughout the network?
- How can I design for differentiated quality of service?
- Where is my network most vulnerable to failure, and how can I mitigate it?
- How will a Layer 1 design affect the resilience of Layer 3?

Operations

- How can I visualize the current health of the network?
- What caused the outage experienced earlier today?
- How would a current failure affect the network at peak time tomorrow?
- Which VPN customers will be affected by a planned outage or by an unplanned failure?

MATE Design Fundamentals

Initially, network topology and routing information is captured by Collector discovery tools and stored in plan files within a series of tables. These plan files are the basis of information displayed in and used by MATE Design. Following is a high-level overview of what they contain.

- Network configurations
- Visual layouts
- IP/MPLS routes, including multicast and LSPs
- Measured traffic
- Estimated end-to-end traffic matrix
- Operational state of network objects
- Results of analyses, such as worst-case failure analysis results

This information is displayed in the GUI through tables and a network plot where you can both visualize and manipulate it further. For example, you can change the visual layouts, simulated traffic, configurations, and operational states.

MATE Design Tools

MATE Design contains numerous tools with user-friendly features, such as easy navigation of object dependencies, bulk editing, and tight interaction between plots and tables. Following is an overview of the categories under which these tools fall.

- **Visualization**—Provides graphical layouts of the network, showing views of different regions and layers, as well as utilizations and routings.

  Related chapters: User Interface, Plan Objects, Plot Legend for Design Layouts

  Related guide: *MATE GUI Visualization Guide*

- **Simulation**—Enables you to perform interactive what-if simulations by, for example, failing objects, modifying network topology, creating and changing the traffic demands (which simulate traffic flows), and modifying routing configurations.
  - Demand Deduction—Estimates an end-to-end traffic flow matrix based on available measured traffic, such as interface utilizations and LSP traffic statistics.
  - Failures and Simulation Analysis—Allows evaluation of the effect of failures of network objects, including circuits, nodes, and Layer 1 equipment. You can evaluate overall risk using worst-case analyses over these failures.
  - Forecasting—Enables you to create plans for future expected traffic levels.

  Related chapters: Traffic Demand Modeling, Simulation, Simulation Analysis, Forecasting Traffic, IGP Simulation, MPLS Simulation, Quality of Service Simulation, BGP Simulation, Advanced Routing with External Endpoints, VPN Simulation, Multicast Simulation, Layer 1 Simulation

- **Offline Traffic Engineering**—Enables you to optimize the network's reliability and performance.
  - Metric Optimization—Suggests changes to the OSPF or IS-IS metrics that enable you to optimize the network, for example, by reducing congestion under circuit failures.

  - LSP Optimization—Explicit LSP Optimization and Tactical Explicit LSP Optimization tools optimize the placement of primary and secondary paths for selected MPLS RSVP LSPs. The LSP Loadshare Optimization tool optimizes LSP load balancing percentages between parallel LSPs.
  - Changeover—Recommends a step-by-step series of configuration changes to maneuver a network safely from one configuration state to another.

  Related chapters: Metric Optimization, LSP Optimization, Changeover

- **Reporting**—Enables you to generate reports that compare IP/MPLS topology and configuration information, demand routings, demand groupings, and LSP routings. You can also generate reports that compare traffic, compare plan files, and estimate network infrastructure costs. Other reports, such as Simulation Analysis and Demand Deduction, are automatically generated.

  Related chapters: Reports, Cost Modeling

Extensible Tools

The plan file format is open and documented. Add-on capabilities are available for extending GUI functionality. Additionally, command-line tool equivalents of the GUI tools are provided for scripting and automation. Refer to the *MATE Integration and Development Guide* for information on CLI tools, creating add-ons and reports, details on how plan files, and other reference material.

For More Information

For More Information ...	See ...
Tutorials that help you get started using and exploring MATE Design.	*MATE Design Tutorials*
Detailed descriptions of how to layout and visualize MATE plan files for use in the applications. It includes template information, as well as how to access and store plan files and templates remotely.	*MATE GUI Visualization Guide*
Advanced information, such as importing objects and traffic, creating customized add-on applications, using the CLI tools, and working with plan files and tables. This guide includes instructions for importing MATE Live traffic into MATE Design for use in forecasting.	*MATE Integration and Development Guide*
Details about tables and CLI tools	*Table Schema and CLI Reference*
Sample plan files	`$CARIDEN_HOME/samples`
CLI tools	`$CARIDEN_HOME/bin`

User Interface

This chapter is a basic framework for understanding how to use the MATE GUI.

- Plan Files—Defines what a plan file is and how to access one
- Main Window—Overall structure of the GUI
- View Plots—Network, site, interface, demands, LSPs
- Objects—Describes what a MATE object and property are
- Tables—How to use and manipulate tables
- Log and Report Windows—How to access log windows and reports

For More Information ...	**See ...**
MATE objects in the network plot and tables	Plan Objects chapter
Shapes, symbols, and colors in the network plot	Plot Legend for Design Layouts chapter
• Visualize and customize a MATE network • Create templates for use in web applications	*MATE GUI Visualization Guide*
Advanced plan file information	*MATE Integration and Development Guide*

Plan Files

The MATE Live, MATE Design Archive, and MATE Design applications all have interfaces for viewing the network through a *plan file*. A plan file is comprised of a series of tables that store information about a network, including topology, configuration information, traffic, failure state, and visual layout Additionally, MATE Design uses plan file information to perform simulations.

Plans files are typically created by the Collector discovery process for use in the applications. MATE Live stores this in an *archive* that is internal to it, and MATE Design Archive stores these in an external archive. Plan files can, however, be created using MATE CLI tools or the MATE GUI, and stored locally.

Plan files are comprised of tables, objects such as nodes and sites, properties (shown as column headings), and property values (Figure 2-1). For more information, see the Objects and Tables sections.

Plan File Formats

Each plan file is a uniquely named database of tables that is in one of two formats, both of which are editable by SQL commands.

- The .pln format is compact and thus, can be quickly loaded to and saved in the MATE GUI.
- The .txt format contains tab-delimited columns. You can create and edit these directly in a text editor or spreadsheet, such as Excel, thus allowing for quickly applying bulk edits. In these plan files, each table is labeled with angle brackets, such as <Nodes> and <Sites>.

Figure 2-1 Opened Plan Filed

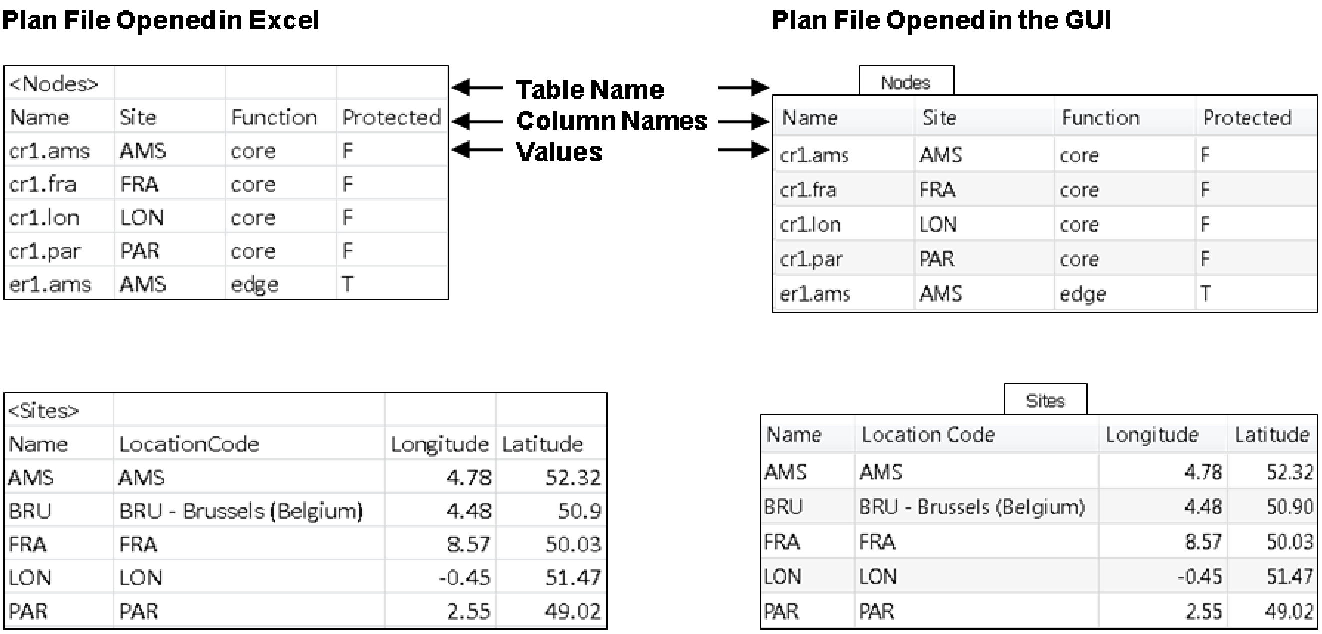

<Nodes>			
Name	Site	Function	Protected
cr1.ams	AMS	core	F
cr1.fra	FRA	core	F
cr1.lon	LON	core	F
cr1.par	PAR	core	F
er1.ams	AMS	edge	T

Nodes

Name	Site	Function	Protected
cr1.ams	AMS	core	F
cr1.fra	FRA	core	F
cr1.lon	LON	core	F
cr1.par	PAR	core	F
er1.ams	AMS	edge	T

<Sites>			
Name	LocationCode	Longitude	Latitude
AMS	AMS	4.78	52.32
BRU	BRU - Brussels (Belgium)	4.48	50.9
FRA	FRA	8.57	50.03
LON	LON	-0.45	51.47
PAR	PAR	2.55	49.02

Sites

Name	Location Code	Longitude	Latitude
AMS	AMS	4.78	52.32
BRU	BRU - Brussels (Belgium)	4.48	50.90
FRA	FRA	8.57	50.03
LON	LON	-0.45	51.47
PAR	PAR	2.55	49.02

Access Plan Files

To open a local plan file from the MATE GUI, select the File->Open menu or File->Open Recent menu. To save a plan file locally in the .pln format, select the File->Save menu. To save a plan file in a simple .txt or complete .txt format, select the File->Save As menu.

To add a description to the plan file, select the Edit->Plan Description menu.

For instructions on opening and saving plan files on remote servers, see the *MATE GUI Visualization Guide*.

Main Window

The MATE GUI is divided into the following primary sections. Figure 2-2 shows the GUI with a sample `us_wan.txt` plan in the network plot area.

- Menus—Menu selections for network modeling, simulation, and optimization.

- Visualization toolbar—Drop-down lists and icons that control how the network is displayed. For more information, see the Visualization Toolbar section.
- Plan file tabs—Identify which plan is currently opened. There can be multiple plans open; the tab with a white background is the active plan. To bring a plan to the forefront and make it the active plan, click on its associated tab.
- Network plot—Displays a visual representation of the network topology. The color on each interface represents the percentage of traffic utilization on that interface. Rows that are selected in the tables, such as demands or interfaces, are highlighted in the plot with additional graphical attributes. For more information, see the View Network Plots section.
- Tables—Display network properties, including both measured and simulated information, for the network plan. For a list of default tables and information on how to interact with tables within the GUI, see the Tables section.
- Status Bar—Identifies the network simulation mode.
- Context menus—Enable you to right-click anywhere to view a list of actions that you can take. For more information, see the View Related Information section.

Figure 2-2 MATE Design User Interface

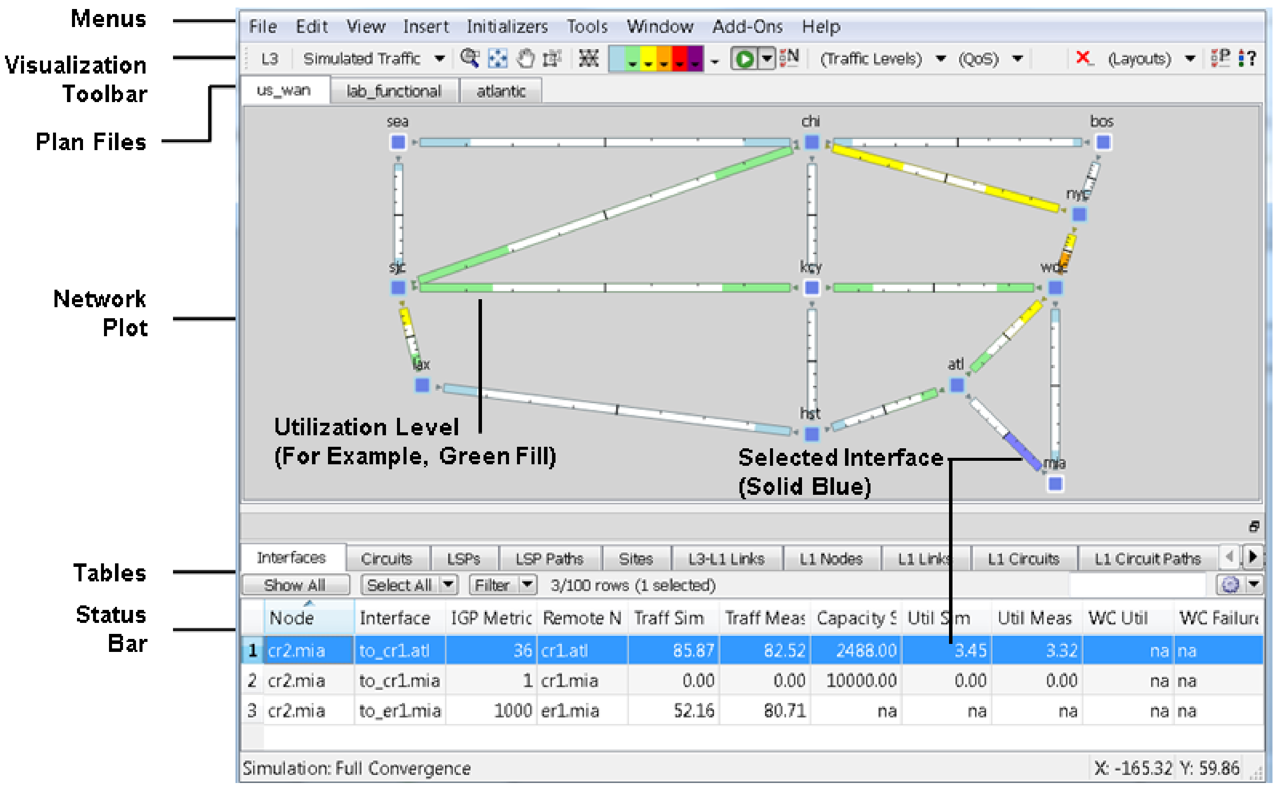

3/100 rows (1 selected)

	Node	Interface	IGP Metric	Remote N	Traff Sim	Traff Meas	Capacity S	Util Sim	Util Meas	WC Util	WC Failure
1	cr2.mia	to_cr1.atl	36	cr1.atl	85.87	82.52	2488.00	3.45	3.32	na	na
2	cr2.mia	to_cr1.mia	1	cr1.mia	0.00	0.00	10000.00	0.00	0.00	na	na
3	cr2.mia	to_er1.mia	1000	er1.mia	52.16	80.71	na	na	na	na	na

Simulation: Full Convergence X: -165.32 Y: 59.86

Visualization Toolbar

The Visualization toolbar (Figure 2-3) provides a convenient way to change the visual aspects of the network plot.

Figure 2-3 Visualization Toolbar

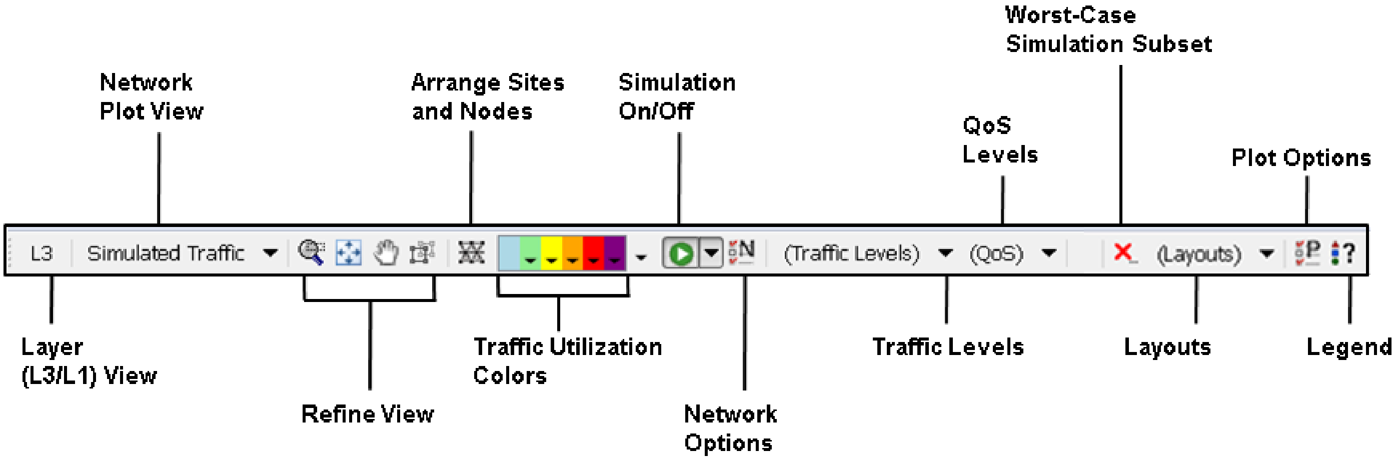

- Layout—A toggle button to change the view between Layer 3 and Layer 1.
- Network Plot—A drop-down menu from which you choose the type of traffic displayed in the plot. For more information, see the View Network Plots section.
- Refine View icons—These icons enable you to refine and enhance your view in both in the network plot and in the site plots. For information, see the Refine Viewable Objects section.
- Arrange icon—Opens an Arrange palette that enables you to move selected nodes and sites into a selected schematic alignment.
- Utilization Color—Traffic utilizations are displayed by colors as a percent of available capacity used on that interface. This enhances the other visual cues, which is the length of the fill. You can change the colors and the utilizations they represent. For more information, see *MATE GUI Visualization Guide*.

Note All traffic is displayed in Mbps. In the network plot, the traffic utilization colors represent outbound traffic.

- Simulation On/Off (white arrow in green circle) icon—By default, any change that voids the current simulation automatically triggers a resimulation. For instance, changes in topology trigger a resimulation. To turn off this resimulation, select the drop-down arrow next to the icon and deselect the “Auto resimulate” option. To turn it back on, reselect the this option.
- Network Options icon—Enables you to set global simulation modes, such as Full Convergence or Fast Reroute for MPLS, and whether to include multicast hops. You can also set IGP, BGP, and Multicast protocol options.
- Traffic Level—Drop-down menu of user-defined traffic levels in the plan. For more information, see the Traffic Demand Modeling chapter.
- QoS Level—Drop-down menu allowing the display of service classes, interface queues, and undifferentiated traffic. From the Manage QoS option, you can add/edit service classes, assign service classes to queues, and add/edit service policies. For more information, see the Quality of Service Simulation chapter.

- Worst-Case Simulation Subset (red X) icon—After a simulation analysis is performed over a specified set of failure scenarios (for example, circuits, nodes, and ports), you can click on this icon to select which subset of these failure scenarios (for instance, circuits only), as well as service classes you want included in the worst-case analysis. This avoids having to run a second simulation analysis.
- Layouts—Drop-down menu of user-defined plots. For example, you might have schematic layouts, geographic layouts, street map layouts, separate layouts for different continents, simplified layouts with hidden nodes or sites, or other representations of the network. For more information, see the *MATE GUI Visualization Guide*.
- Plot Options icon—Enables you to customize circuit width, text on interfaces, layouts (such as background images and geographic views), and other plot options. For more information, see the *MATE GUI Visualization Guide*.
- Legend (question mark) icon—Displays a legend of icons and line types in the plot and their meanings. For more information, see the Plot Legend for Design Layouts chapter.

View Plots

MATE Design opens to the main network plot. Although MATE Design shows only the sites in the network plot and tables, you can view those sites in more depth. You can also view more depth of interfaces and demands in their respective plots. In addition to these plot views, you can create and view layouts that show only those portions of the network plot that you choose. For more information, see the *MATE GUI Visualization Guide*.

View Network Plots

The network plot, also referred to simply as the *plot*, is the main area showing the network topology. It can contain both sites and nodes.

There are multiple network plot views available (top left of Figure 2-3). The default view is Simulated Traffic. To change the view, select it from the Network Plot View drop-down list in the far, left Visualization toolbar.

- Simulated Traffic—View of simulated traffic that can be used for capacity planning, what-if analysis, and failure planning. Circuits are sized proportional to their capacity. Interfaces are colored according to the percentage of simulated traffic (in comparison to available capacity as defined by the size of the interface or the QoS bound).
- Measured Traffic—View of traffic as measured from live views of the network. Circuits are sized proportional to their capacity. Interfaces are colored according to the percentage of measured traffic (in comparison to available capacity as defined by the size of the interface or the QoS bound).
- Worst-Case Traffic—Interfaces are colored to show the maximum utilization of the interface over all failure scenarios as defined by the most recently run simulation analysis. For more information, see the Simulation Analysis chapter.
- Failure Impact—Interfaces and nodes are colored to indicate how a failure of that interface or node would impact other interfaces and nodes. For more information, see the Simulation Analysis chapter.
- LSP Reservation—Circuits are sized proportional to their LSP reservable bandwidth. Interfaces are colored according to the percentage of bandwidth (in comparison to what is reserved for the LSP). For information about using LSPs in MATE Design, see the MPLS Simulation chapter.

- By editing the plan file directly, you can add a user-defined static view to display interfaces with customized color fills and fill percentages. The view then appears as an option in the Network Plot drop-down list. For information, see the *MATE Integration and Development Guide*.

Although MATE Design focuses on IP/MPLS Layer 3 (L3) network views and traffic, you can also create and view Layer 1 topologies. To toggle between views of these two layers, use the View->Layer 1 View or View->Layer 3 View menus, or use the shortcuts Ctrl-1 for Layer 1 and Ctrl-3 for Layer 3. For information, see the Layer 1 Simulation chapter.

View Site Plots

The site plot can contain other sites and/or it can contain nodes and their connecting interfaces, including external interfaces that are labeled with their destination node (Figure 2-4).

To open a site plot, follow one of these steps.

- Double-click on a single site from the plot.
- Select multiple sites from the plot and then double-click on one of those sites.
- Select one or more sites from the plot or the Sites table.
  - Select View->Selected Sites.
  - Right-click and select Open Site View from the context menu.

Figure 2-4 Site Plots

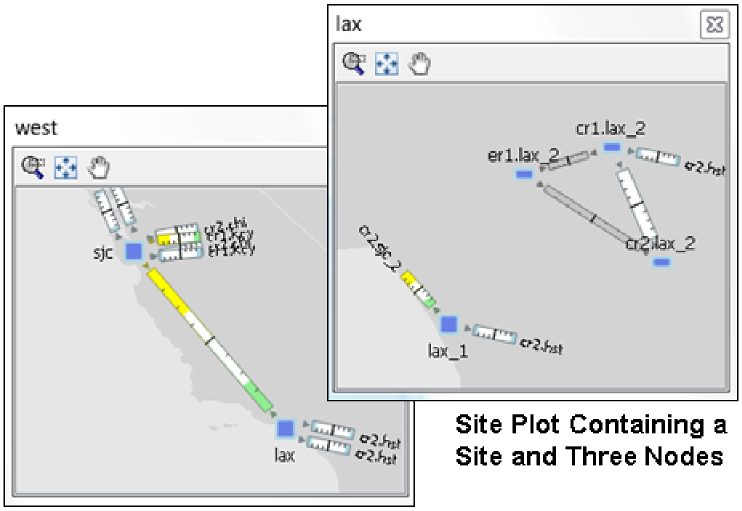

Site Plot Containing a Site and Three Nodes

Site Plot Containing Two Sites

View Interface Plots

The interface plot is a window that displays selected interfaces in a stacked format, making it easier to compare one with another (Figure 2-5).

Step 1 Select one or more interfaces from either the Interfaces table or the plot.

Step 2 Open the interface plot in one of two ways.

- Select View->Interface Plot.
- Right-click and then select Plot Interfaces from the context menu.

Following are further options for interface plots.

- To vertically move the interface list within the plot, click the up and down arrows.
- To view a table of the selected interfaces, click the Adjust button. A new Interfaces table appears.
- To sort the interfaces by ascending or descending order, follow these steps. The interface plot displays the interfaces according to the sort.
  1. Click on the Adjust button.
  2. Click on a column heading to toggle its sorting between ascending or descending order.
  3. Click OK.

Figure 2-5 Interface Plot

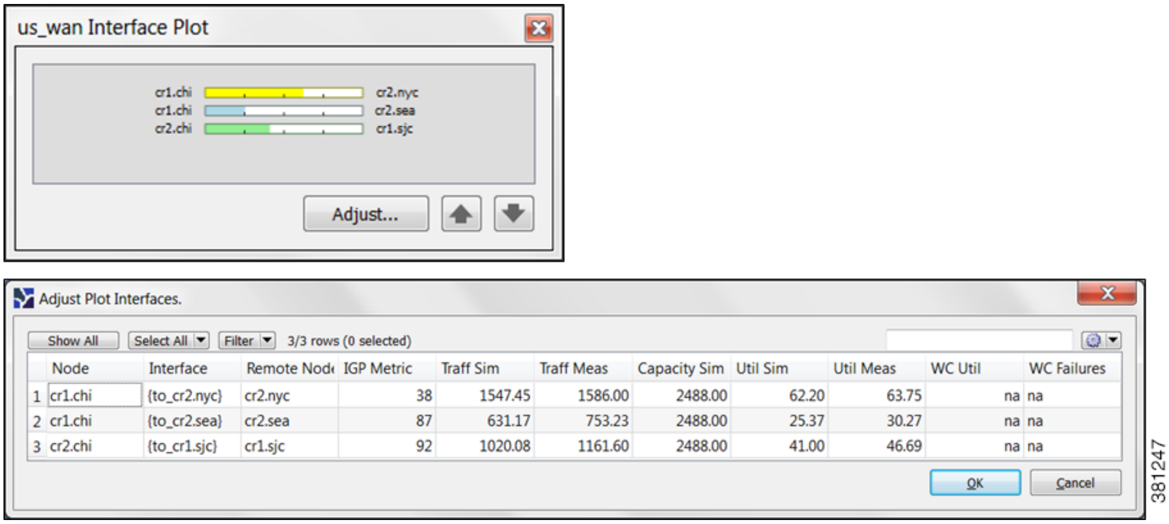

	Node	Interface	Remote Node	IGP Metric	Traff Sim	Traff Meas	Capacity Sim	Util Sim	Util Meas	WC Util	WC Failures
1	cr1.chi	{to_cr2.nyc}	cr2.nyc	38	1547.45	1586.00	2488.00	62.20	63.75	na	na
2	cr1.chi	{to_cr2.sea}	cr2.sea	87	631.17	753.23	2488.00	25.37	30.27	na	na
3	cr2.chi	{to_cr1.sjc}	cr1.sjc	92	1020.08	1161.60	2488.00	41.00	46.69	na	na

Objects

Other than a site, an *object* is a representation of elements found in live networks, such as nodes (which represent routers), circuits, interfaces, LSPs, and more. A site is also an object, and is a MATE construct for simplifying the visualization of a network by grouping nodes within a site, or even by grouping sites within a site.

Objects have *properties* that identify and define them, many of which are discovered, though they can also be manually added and changed. For example, all circuits have a discovered Capacity property that can be edited. Other properties are derived. For example, Capacity Sim is derived from the Capacity property. Another example is that interfaces have a Util Sim property that identifies the percentage of Capacity Sim the simulated traffic is using. Properties are viewable and editable through an object's Properties dialog box (right-click on the object and select Properties). They are also viewable as table columns.

For more detailed information on object creating and editing, see the Plan Objects chapter.

Object Selections

Objects are selectable through tables and plots. However, not all objects are selectable from both, though most are. Some object types are only selectable through their tables, such as LSPs.

All selections made in a plot appear in the associated table and vice versa (Table 2-1). Likewise, changes made in the plot are reflected in the tables, and vice versa. Hidden objects are excluded from the selection.

Table 2-1assumes that all objects are in the foreground. For information on selecting objects that are in the foreground versus background, see the Refine Viewable Objects section.

Table 2-1 Select Objects

To ...	Do This ...
Site Selections • If the object you select is a site, then none of the child objects (child sites, nodes, and circuits) are selected. • To select a site and all objects contained within it, press and hold the Alt key during the selections identified in this table. **Examples** • To select only a site, click it. To select a site and all of its nested child objects, press and hold Alt while clicking the site. • To select only two sites, click on one object and then Ctrl-click on each site that you want to select. To select two sites and all the objects they contain, press and hold both Alt and Ctrl while selecting both sites. • To select all sites in the plan file, press Ctrl-A in either the network plot or the Sites table. To select all sites and all objects they contain, press and hold Alt and Ctrl-A together in either the network plot or the Sites table. .	
Select an object. If other objects were selected prior to this new selection, those objects become deselected.	Click on one or more objects in any plot or table.
Select multiple objects.	Do one of these steps. • In a plot or table, click and hold while you drag the cursor across any objects or rows you want included in the selection. • In a plot or table, click on one object and then Ctrl-click on each object that you want to select.
Select all objects.	Do one of these steps. If you are selecting from a plot, all objects in that plot are selected. If you are selecting from a table, all objects within that table are selected. • Press Ctrl-A in either the network plot or a table. (Cmd-A on Mac.) • Click and hold while you drag the cursor over an entire plot or through the entire list in a table.

To ...	Do This ...
Clear the entire selection.	Do one of these steps. • Select the Select All -> Clear Selection in the table control menu. This deselects all objects within that table. If objects in other tables are selected, they remain selected. • Click anywhere in any plot. This deselects all objects in that plot and in the associated table.
Invert the current selection (everything selected becomes deselected and vice versa).	1. Show the table whose selections you want to invert. 2. Select the Select All -> Invert Selection from the table control menu.

Filter to Selected Objects

To filter only to the selected objects, select one or more objects from any plot or table, then select Filter->Filter to Selection. Alternatively, you can press and hold Ctrl-Shift (Cmd-Shift on Mac), and then double-click.

As with other object selections, to include the nested objects for site selections, add Alt to the selection process. For example, press and hold Alt while selecting one or more sites, and then select Filter->Filter to Selection.

Filter to Nested or Parent Objects

If a site contains multiple levels of child sites, you can filter to those nested objects. Right-click on a site, and select Filter to Contained <Object> and then select one of the following, where <Object> equals sites, nodes, or L1 nodes.

- Direct—To see only the objects immediately contained within the site.
- All—To see all contained objects, no matter how deeply nested they are within child sites.

If a site contains only one level, the context menu changes to reflect as such.

Conversely, you can find a parent site for any given node. Select it, and then select Filter to Sites.

Refine Viewable Objects

View icons are available both in the network plot and in the site plots. Each of these icons make it easier to find what you need.

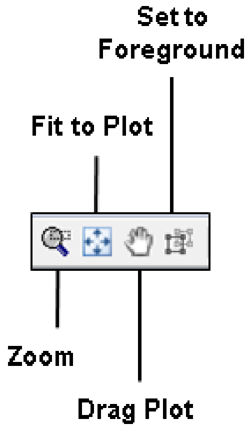

Except for the Foreground mode, when accessed through the Visualization toolbar, the view changes are applicable only to the network plot. When accessed through a site plot, the Zoom, Fit-to-Plot, and Drag modes apply only to that site plot.

Both the Zoom and Foreground modes enable you to better focus and work in specific plot areas.

- Zoom mode—Enables you to drag in the area that you want enlarged to fit the entire plot.

  Shortcut = Press the letter “z,” which toggles this functionality on and off.

  To escape from the Zoom mode, press Esc.
- Fit-to-Plot mode—Resizes the diagram to fit within and fill the plot area.

  Shortcut = Press the letter “f,” which toggles this functionality on and off.
- Drag Plot mode—Enables you to move the cursor over any portion of the plot to drag it to a different location.

  Shortcut = Press the letter “d” or press the icon again to toggle this functionality on and off. The feature remains enabled until you turn it off.

  To escape from the Drag mode, either press Esc or click the icon to toggle it off.
- Foreground mode—Brings selected objects to the foreground for improved viewing while sending all other objects to the background where they appear as if they are transparent. Selected objects must be L3 nodes, L3 circuits, sites, L1 nodes, or L1 links. This feature does not apply to other objects, such as demands, LSPs, L1 circuits, and L1 circuit paths.

  To escape the Foreground mode, click the icon to toggle it off.

Set Foreground Preferences

To set foreground preferences, select the View->Preferences menu.

- Enable Selection of Background Objects—If this is enabled (selected), then you can select objects in the background. Otherwise, you cannot do so.
- Fit to Screen Zooms to Foreground Objects—If this is enabled (selected), then clicking the Fit-to-Plot icon zooms into and centers the selected objects. Otherwise, the plot itself is centered.
- Save Preferences—Save all view and tooltip preferences for all layouts when you save the plan file.

View Related Information

Tooltips

One method of seeing related information about an object is to show the tooltips and then hover the mouse over a site, node, L1 node, interface, or L1 link. For instance, hovering over a site you can see the number of contained L3 nodes, L1 nodes, and sites, as well as the total number of interfaces connected to the contained L3 nodes. To toggle this tooltip on or off, select the View->Preferences->Show Tooltips.

Context Menus

MATE Design has a hidden menu called a *context menu* (Figure 2-6). Anywhere that you right-click, whether it is in an empty area in the plot, on an object in the plot, on a table tab, or on a column or row in a table, a context menu appears giving you available options based on where you clicked. The only areas this context menu does not apply is in the main menus and in the Visualization toolbar.

The most common use case for context menus is for filtering to related objects. For example, to find the L1 circuits mapped to an L3 circuit, right-click on the L3 circuit and select Filter to L1 Circuits from the context menu. You can select multiple objects and right-click on just one to find the related information for all selected objects. If needing to see only selected objects in a table, right-click on one of the selected objects and select Filter to Selections.

To access a context menu, right-click any empty area of a plot or right-click on one or more objects from the network plot or table. If in a site, an interface, or a demand plot, the actions taken affect the network plot.

Figure 2-6 Examples of Context Menus

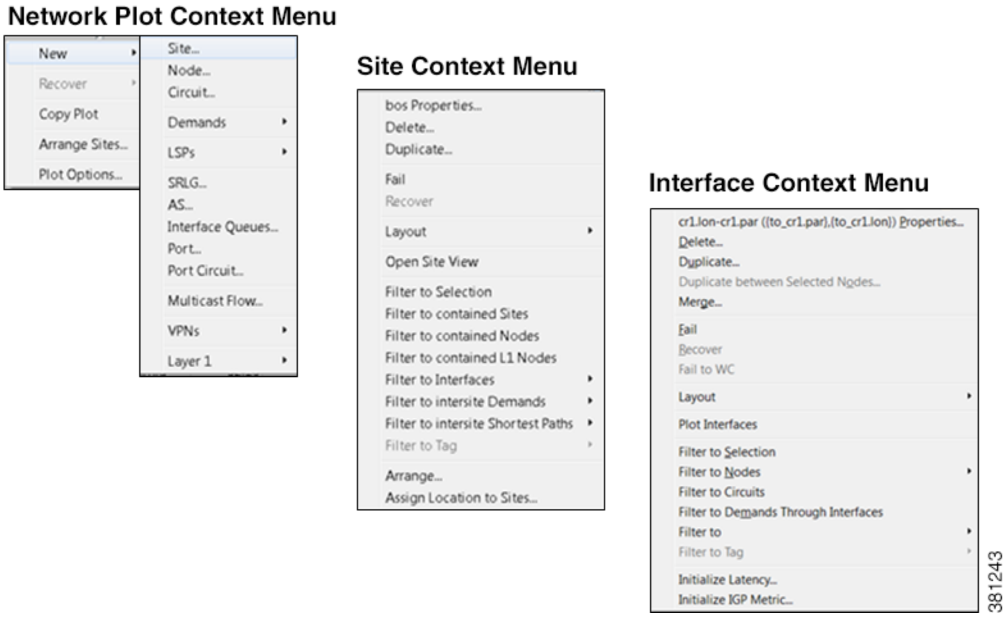

Tables

The tables that are viewable from the MATE GUI display a series of rows and columns, wherein the rows are objects and the columns are properties (Figure 2-7). Plan files also contain other tables that are not viewable from the GUI. These tables contain more complex information, such as showing complex relationships between objects. If you open a plan file using a .txt editor, each table is labeled with angle brackets, such as <Nodes> and <Sites> (Figure 2-1). For detailed information about plan files and tables, see the *MATE Integration and Development Guide*.

Common Tables

There are numerous tables available for specialized purposes, such as Multicast, P2MP LSPs, and Ports default. The tables listed below are the most commonly used and are the defaults. For information on showing and hiding tables, see Show or Hide Tables or Columns.

To bring a table to the forefront and make it the active table, click on its associated tab. For example, to open the LSP table, click on the LSP tab.

- Interfaces—A list of the interfaces in the network.
- Demands—A list of the demands in the network. Each demand specifies how much traffic is routed from a source (node, external AS, or external endpoint) to a destination (node, external AS, external endpoint, or multicast flow destination).
- Shortest Path—A list of the routes between selected sites.
- Nodes—A list of network routers, which typically have names that suggest their location and function. For example, node cr1.atl is core router 1 in the Atlanta site.
- LSPs—A list of MPLS LSPs in the network; each LSP contains both a source and destination. Note that sub-LSPs within point to multipoint (P2MP) LSPs are listed in the LSP table.
- Sites—A list of network sites.
- SRLGs—A list of the shared risk link groups (SRLGs); an SRLG is a group of objects that might all fail due to a common cause.
- AS—A list of both internal and external autonomous systems (AS's). An AS is a collection of connected IP routing prefixes that are controlled by one operator.

Figure 2-7 Common Table Functions

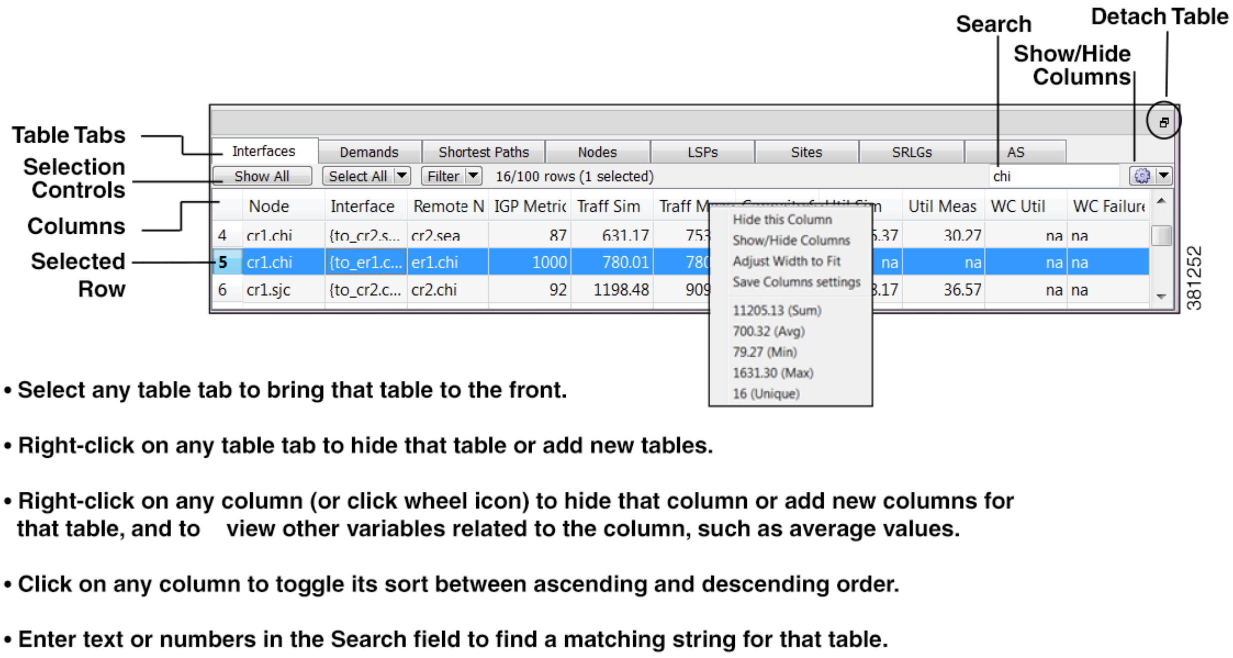

16/100 rows (1 selected)

chi

	Node	Interface	Remote N	IGP Metric	Traff Sim	Traff M	Util Sim	Util Meas	WC Util	WC Failure
4	cr1.chi	{to_cr2.s...	cr2.sea	87	631.17	753	5.37	30.27	na	na
5	cr1.chi	{to_er1.c...	er1.chi	1000	780.01	780	na	na	na	na
6	cr1.sjc	{to_cr2.c...	cr2.chi	92	1198.48	909	3.17	36.57	na	na

11205.13 (Sum)
700.32 (Avg)
79.27 (Min)
1631.30 (Max)
16 (Unique)

- Select any table tab to bring that table to the front.
- Right-click on any table tab to hide that table or add new tables.
- Right-click on any column (or click wheel icon) to hide that column or add new columns for that table, and to view other variables related to the column, such as average values.
- Click on any column to toggle its sort between ascending and descending order.
- Enter text or numbers in the Search field to find a matching string for that table.

Show or Hide Tables or Columns

Table 2-2 Show or Hide Tables or Columns

To ...	Do This ...
Show or hide one or more tables.	**1.** Do one of these two steps. – Select View->Tables->Show/Hide Tables. – Right-click on any table tab and select Show/Hide Tables. **2.** Toggle the tables you want to show to T (true) and the ones you want to hide to F (false). **3.** Click OK.
Hide a single table.	Right-click on any table tab and select Hide this Table.
Show or hide one or more columns.	**1.** Do one of these steps. – Select View->Tables->Show/Hide Columns. – Right-click on any column heading and select Show/Hide Columns. – Click the wheel icon in the table area. **2.** Toggle the columns you want to show to T (true) and the ones you want to hide to F (false). **3.** Click OK.
Hide a single column.	Right-click on any column heading and select Hide this Column.

Add Table Columns

Once a column is added, it appears in the table with a prefix::name format.

Step 1 Show the table in which you want to add the column.

Step 2 Do one of these steps.

- Select Edit->User Columns.
- Double-click a row in the table to open a Properties dialog box, and then select the User tab.

Step 3 In the dialog box that appears, click New; a New User Property dialog box appears.

a. Enter the prefix, name, and optionally, a value.

b. Click OK.

Step 4 Click OK.

Search and Filter in Tables

Most MATE Design tables, whether below the network plot or with dialog boxes, have an empty Search field in the top right as shown in Figure 2-7. Likewise, they usually have a Filter control menu that opens a Filter dialog box specific to that table. Figure 2-8 is an example of an Interfaces Filter dialog box.

Table 2-3 Use Table Search and Filter Features

To ...	Do This ...
Search for objects that match a specific string.	1. If you are not in a dialog box, show the table whose objects you want to search. 2. In the Search field, enter a search string and press Enter.
Filter the table to a list of specific search values, and either replace, add to, or narrow the list of objects currently showing in a table.	1. If you are not in a dialog box, show the table whose objects you want to filter. 2. Click the Filter control menu button (Figure 2-8). 3. Choose from the drop-down lists and enter search criteria. 4. If you are filtering by more than one criterion, you can optionally use these controls. – Select Match All to filter to only those rows that match all the specified criteria (default). – Select Match Any to filter to rows that match any one of the criteria. 5. Select whether to replace, add to, or narrow the list of objects currently showing in a table. – Replace—Replace objects showing in the table with those found in the search (default). – Union—Add the objects found to those showing in the table. – Intersection—Filter to (search for) objects only within those objects that are currently showing in the table.
Filter the selection to associated objects. For example, you might want to filter a circuit to associated interfaces.	Select the Filter -> Filter to X control menu, where X is a name of a viable object.
Toggle whether to change the filter selection with or without first asking.	Select the Filter -> Interaction with current filter -> X, where X is either Always Ask or Replace.

Figure 2-8 Interfaces Filter

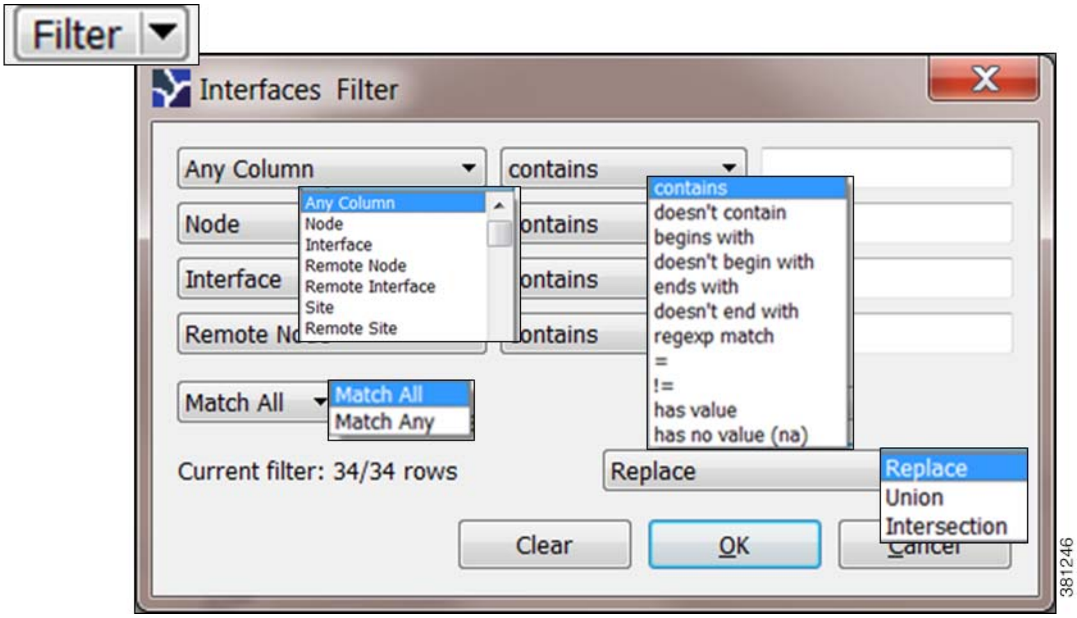

Save Table Settings or Return to Defaults

Table 2-4 Save Table Settings or Return to Defaults

To ...	Do This ...
Save table and column settings.	Select View->Tables->Save All Settings
Save table settings.	Do one of these steps. • Select View->Tables->Show/Hide Tables. • Right-click on any table tab and select Save Table Settings.
Save column settings.	**1.** Show the table whose column settings you want to save. **2.** Do one of these steps. – Select View->Tables->Show/Hide Columns. – Right-click on any table tab and select Save Column Settings.
Restore all table and column settings to system defaults.	Select View->Tables->Reset All.

To ...	Do This ...
Resize column widths.	1. Show the table whose column settings you want to save. 2. To adjust the column width to that of the longest variable, right-click on the column heading, and select Adjust Width to Fit. To adjust the column width to any desired size, click and drag the column heading edge to any size.
Restore column widths to their original size.	1. Show the table whose column widths you want to restore. 2. Select View->Tables->Reset Column Widths.

Other Table Actions

Table 2-5 Miscellaneous Table Tools

To ...	Do This ...
Show all objects for a table within the plan.	1. Show the table whose objects you want to see. 2. Click on the Show All control button.
Toggle the columns to sort in ascending or descending order.	Click on a column heading.
View the sum of the values in a column, the average, the minimum, and the maximum.	Right-click on a column heading.
View how many entries in a column are unique.	Right-click on a column heading.
Resize column widths.	• To adjust the column width to that of the longest variable, right-click on the column heading, and select Adjust Width to Fit. • To adjust the column width to any desired size, click and drag the column heading edge to any size.
Restore column widths to their original size.	1. Show the table whose column widths you want to restore. 2. Select View->Tables->Reset Column Widths.
Detach the table from the plot.	Do one of these steps. • Click and hold the empty gray area above the table while you drag it to a different location. • Click the double-page icon in the right of that same gray area.
Re-attach the table to the plot.	Click and hold the empty gray area above the table while you drag it to the bottom of the plot. When an empty blue area appears, release the mouse button.

Log and Report Windows

The following additional windows are displayed using menu items or after running tools

- Log window—This window appears automatically when warnings or errors are logged. You can view these by selecting the Window->Log Window menu.
- Reports—Reports are generated through the Tools->Reports menu, as well as automatically when running certain tools such as Simulation Analysis. The report window appears at that time, and is accessible later by selecting the Window->Reports menu. For more information, see the Reports chapter.

Modify Plan File Tables

Although you can edit plan files in the MATE GUI, we recommend that you edit them in a text editor or spreadsheet so that you can easily make global changes. A spreadsheet, such as Excel, is useful since it clearly shows table columns in a tabular format.

Any time you make a change in the MATE GUI and save the plan file, you are, in effect, changing the tables in the plan. Additionally, you can select the Edit->Plan Tables as Database menu in the MATE GUI to view and edit tables directly in the plan file. In this dialog box, you can select one or more rows to modify attribute values in the columns. You can specify the desired selection from the Advanced Filter tab, by entering SQLite queries, and specify value changes in the Set Value tab, also using SQLite syntax.

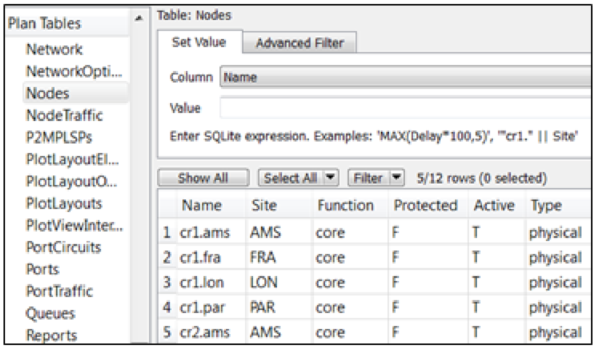

Table: Nodes

Column Name

Value

Enter SQLite expression. Examples: 'MAX(Delay*100,5)', '"cr1." || Site'

5/12 rows (0 selected)

	Name	Site	Function	Protected	Active	Type
1	cr1.ams	AMS	core	F	T	physical
2	cr1.fra	FRA	core	F	T	physical
3	cr1.lon	LON	core	F	T	physical
4	cr1.par	PAR	core	F	T	physical
5	cr2.ams	AMS	core	F	T	physical

Values defined in one table might be referenced in other tables. Thus, care should be taken when changing a value that occurs multiple places. For example, the Nodes table identifies the sites in which nodes reside. So if you change a site name in the Sites table, you need to change it in the Nodes table and all other instances, as well.

CHAPTER 3

Plan Objects

MATE Design networks consist of objects, such as nodes, sites, interfaces, circuits, SRLGs, ports, and port circuits. This chapter describes these basic objects and their relationships, as well as how to create, edit, and delete them.

Most objects are represented in the network plot, and all of them are represented in tables. Objects in the plot and in the associated tables have a set of properties that you can manage using a Properties dialog box. These object properties are represented by columns in the object's table, or by entries in tables of related objects. For example, all circuits have a Capacity property.

Note MATE Design has both a Layer 3 (L3) view and a Layer 1 (L1) view, each containing objects.

- Throughout this guide, the terms *node* and *circuit* refer to objects in the L3 view. When describing or contrasting both layers, the qualifier *L3* is used.
- The terms *L1 node* and *L1 circuit* refer to objects in the Layer 1 view.

For More Information ...	See ...
View objects in the plot	User Interface chapter
Object notation in tables	*MATE Integration and Development Guide*
Visual object notation (such as colors and symbols)	Plot Legend for Design Layouts chapter
Visualize a MATE Design network, including multiple layouts of the same network	*MATE GUI Visualization Guide*
Import objects into a plan file	*MATE Integration and Development Guide*
Set an object's state to active/inactive or failed/not failed	Simulation Analysis chapter
Additional MATE Design objects • LSPs and P2MPs • Layer 1 nodes, Layer 1 links, and Layer 1 circuits • Demands • External endpoints • AS's	• MPLS Simulation chapter • Layer 1 Simulation chapter • Traffic Demand Modeling chapter • Advanced Routing with External Endpoints chapter • BGP Simulation chapter

Editing Objects Quick Reference

This table is a simplification of basic object editing capabilities. Not all tasks apply to all objects.

Task	Select Object First?	Menu	Context Menu
Add		Insert	New
Duplicate	One object	Edit->Duplicate	Duplicate
Delete	One or more objects	Edit->Delete	Delete
Merge	One or more objects	Edit->Merge	Merge
Edit properties	One or more objects		Properties

Nodes and Sites

Both *node* and *site* are MATE Design terminology.

- **Node**—A device in the network, which can be one of three types: physical, PSN (pseudonode), or virtual. The Type property distinguishes whether the node represents a real device or an abstraction, such as a single node representing a number of edge nodes connected in the same way to the network. A physical node is a Layer 3 device, or router. Physical and virtual nodes behave in the same way within MATE. A pseudo-node (PSN) is typically used to represent a Layer 2 device or a LAN.

  Nodes can reside both inside and outside of sites. The external arrangement could be useful for small networks where routers are not geographically dispersed.

- **Site**—A collection of nodes and/or other sites that potentially form a hierarchy of sites. Any site that contains other sites is called the *parent* site.

Both nodes and sites have simulated traffic, while nodes also have measured traffic.

- Nodes are shown as rectangles. The border color is an indication of traffic sourced from and destined to the node.
- Sites are shown as squares. The border color indicates the traffic utilization of all the nodes and circuits inside the site, including all nested nodes and circuits.

  Note that sites can contain both L3 nodes and L1 nodes. Empty sites appear in both views, as do sites containing both L3 and L1 nodes. Only sites containing L3 nodes appear in the L3 view, and only sites containing L1 nodes appear in the L1 view. To change these defaults, see the *MATE GUI Visualization Guide*.

Parent Sites and Contained Objects

Unless a site is empty, there is a hierarchy of the sites and nodes that a site contains. A site can be both a parent and a child site. If a site contains another site, it is a *parent* site. Any nodes, sites, or circuits within a parent site are *contained* (*nested*) objects. The contained nodes and sites are also called *children*. Often the child nodes and sites are geographically co-located. For instance, a site might be a PoP where the routers reside.

- The site's Parent Site property defines whether it is nested within another site. If it is empty, the site is not nested.

- In the network plot, the parent site shows all egress inter-site interfaces of all nodes contained within it, no matter how deeply the nodes are nested. Similarly, this is true for each child site plot within it.
- Selecting a site from the network plot does not select the sites or nodes under it.
- To filter to the sites or nodes contained within a site, right-click to access its "Filter to contained" menu. If you select "Direct," MATE filters to only those objects contained immediately within it. If you select "All," MATE filters to all objects contained within it, no matter how deeply they are nested.

Figure 3-1 shows an example of nodes outside of sites (atl nodes), sites containing other sites (the East site contains three NYC sites and one BOS site), and nodes within a site (NYC_2 site contains three nyc nodes). While this figure only shows one of the child site plots, if you were to expand all nested sites with the East site, you would find more nested child sites and nodes. As you can see in the network plot, the East site has 10 egress interfaces extending from it, some of which are within nested sites.

Figure 3-1 Example of Node and Site Relationships

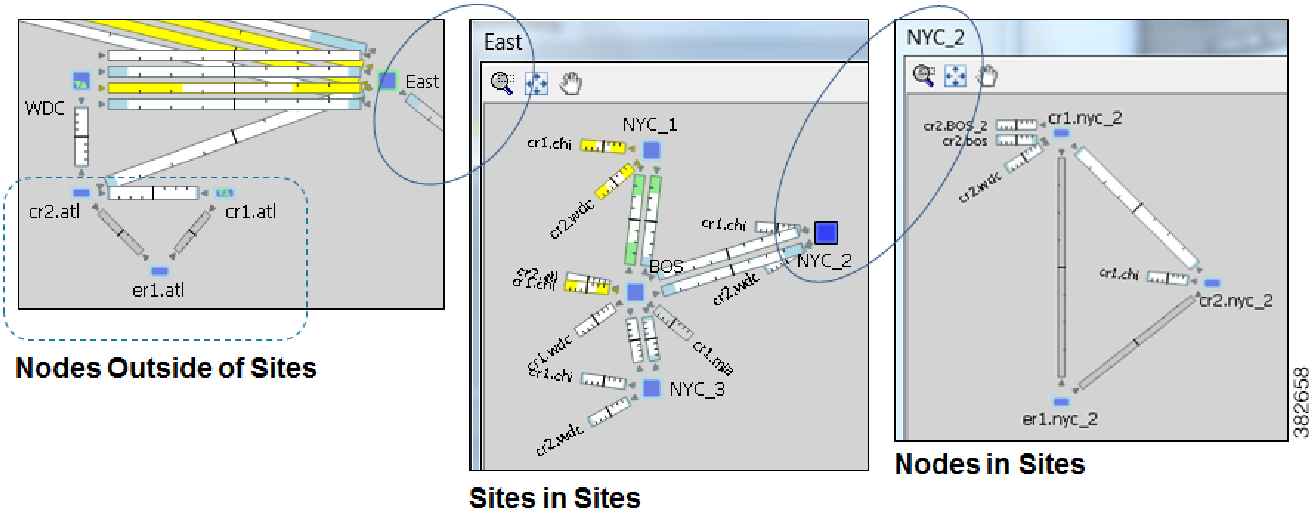

Deleting Sites

Selecting a site from the network plot does not select the sites or nodes under it except when deleting the site. In this case, all objects within a site are selected for deletion. However, in the confirmation that appears, you have the option to keep the contained sites and nodes. If you do, then the objects that are contained directly within it are moved to be on the same level as the site that is removed. The other, more deeply nested objects maintain their parent relationships.

Figure 3-2 shows a West site plot directly containing two sites (SJC and LAX) and three nodes. This parent West site resides in the network plot. Both SJC and LAX contain two sites each. By deleting the West site plot, but retaining its children, both SJC and LAX move to the network plot level, as do the three sea nodes. The sites within SJC and LAX remain as children sites to their respective parent sites.

Figure 3-2 Example Site Deletion

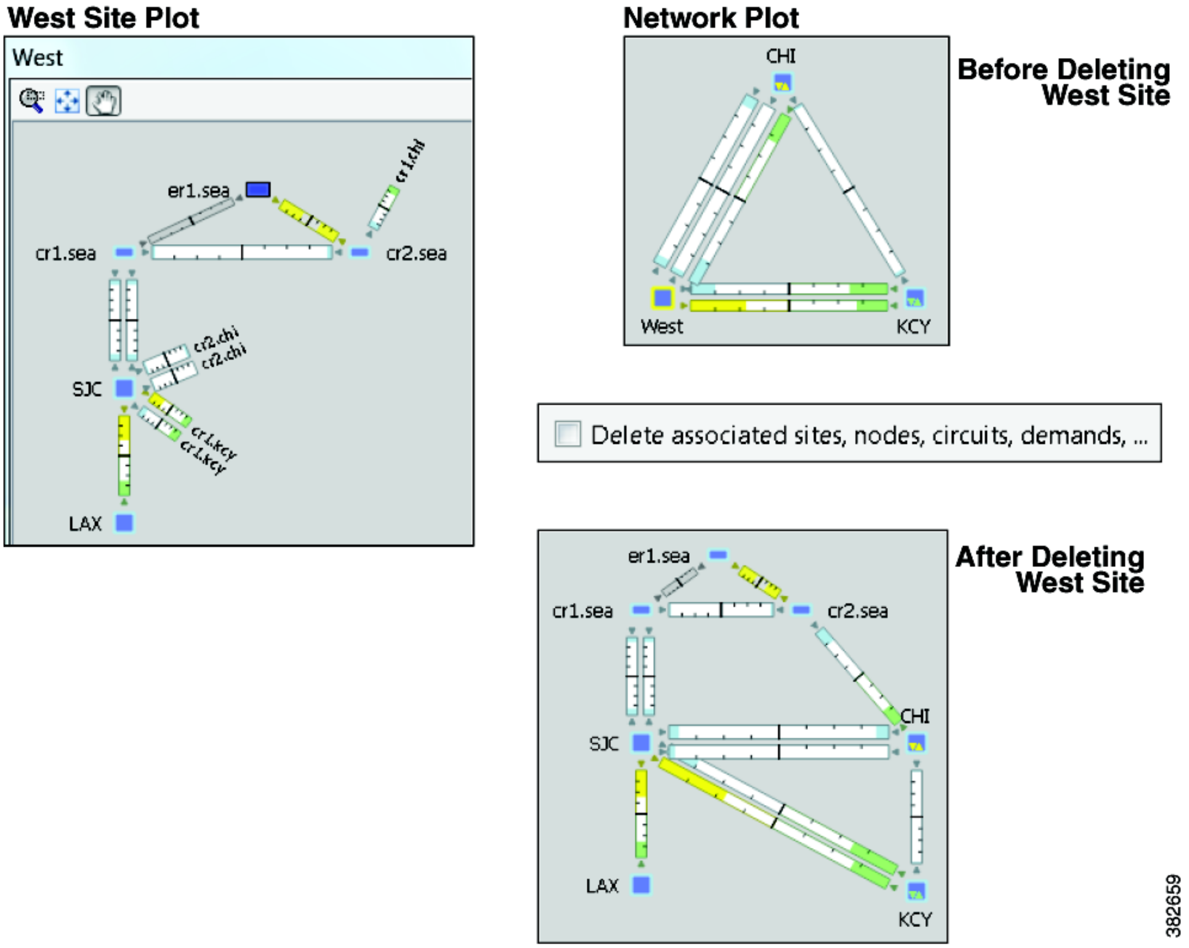

☐ Delete associated sites, nodes, circuits, demands, ...

PSN Nodes

MATE plan files can contain nodes of with a Type property of "psn." These nodes represent pseudonodes (PSNs), which are used to model LANs or switches that connect together more than two routers. They are used in two situations: for IGP modeling and BGP peer modeling.

In an IGP network, a LAN interconnecting multiple routers is represented by a PSN node, with circuits connected to each of the nodes representing the interconnected routers. Both OSPF and IS-IS have a built-in system whereby one of the routers on that LAN is the designated router (DR) for OSPF or the designated intermediate system (DIS) for IS-IS. The PSN node is named after this designated router. Collector creates nodes with a property Type of psn automatically during IGP discovery.

When BGP peers are discovered, Collector might find that a router is connected to multiple peers using a single interface. This is typical at switched Internet Exchange Points (IXPs). Collector then creates a node with a property Type of psn, and connects all the peers to it, each on a different interface.

Following are a few points to consider when working with nodes that have psn as a Type property.

- Two PSNs cannot be connected by a circuit.
- If a PSN node is created by Collector, “psn” is prepended to the designated router’s node name.

- When creating demand meshes, MATE does not create demands with nodes of Type psn as sources or destinations. This is possible in manual demand creation, but not recommended. MATE sets the IGP metric for all egress interfaces from a node of Type psn to zero. This ensures that the presence of a PSN in a route does not add to the IGP length of the path.

Create Nodes and Sites

Create Nodes

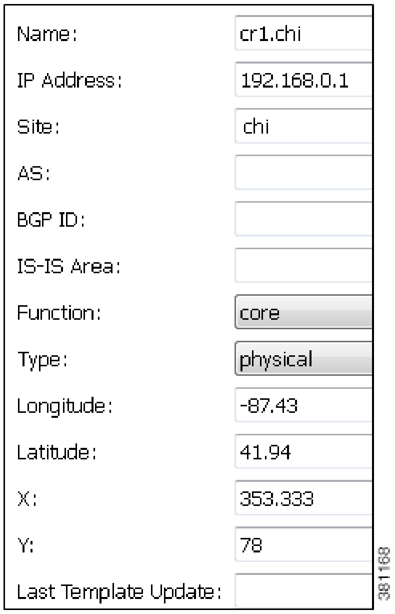

Name:	cr1.chi
IP Address:	192.168.0.1
Site:	chi
AS:	
BGP ID:	
IS-IS Area:	
Function:	core
Type:	physical
Longitude:	-87.43
Latitude:	41.94
X:	353.333
Y:	78
Last Template Update:	

To create a node, either select the Insert->Node menu or right-click in an empty plot area and select New->Node from the context menu. Following are some of the frequently used fields.

- Name—Required unique name for the node.
- IP Address—Often the loopback address used for the router ID.
- Site—Name of the site in which the node exists. If left empty, the node resides in the network plot. This offers a convenient way to create a site while creating the node, move nodes from one site to another, or remove nodes from a site so it stands alone in the network plot.
- AS—Name of the AS in which this node resides, which identifies its routing policy. This can be left empty if no BGP is being simulated.
- BGP ID—IP address that is used for BGP.
- Function—Identifies whether this is a core or edge node.
- Type—The node type, which is physical, PSN, or virtual.
  Because a PSN node represents a Layer 2 device or a LAN, interfaces on a PSN must all have their IGP metrics set to zero, and two PSN nodes cannot be directly connected to one another. If you change a node type to PSN, MATE Design automatically changes the IGP metrics on its associated interfaces to zero.
- Longitude and Latitude—Geographic location of the node within the network plot. These values are relevant when using geographic backgrounds.
- X and Y—Location of the node within the network plot. These values are relevant when using schematic backgrounds.

Create Sites

To create a site, either select the Insert->Site menu or right-click in an empty plot area and select New->Site from the context menu. Following are some of the frequently used fields.

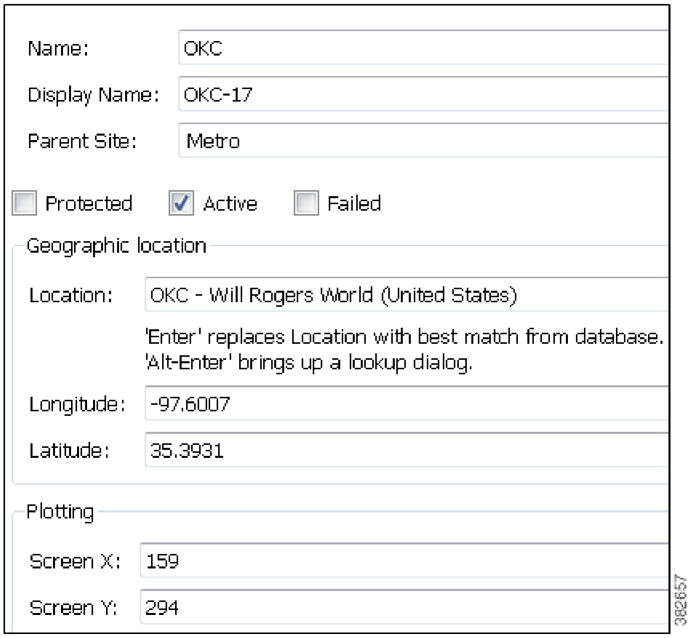

Name: OKC

Display Name: OKC-17

Parent Site: Metro

□ Protected ☑ Active □ Failed

Geographic location

Location: OKC - Will Rogers World (United States)

'Enter' replaces Location with best match from database.
'Alt-Enter' brings up a lookup dialog.

Longitude: -97.6007

Latitude: 35.3931

Plotting

Screen X: 159

Screen Y: 294

- Name—Required unique name for the site.
- Display Name—Site name that appears in the plot. If this field is empty, the Name entry is used.
- Parent Site—The site that immediately contains this site. If empty, the site is not contained within another one.
- Location—To automatically place a site in its correct geographic location and update the Longitude and Latitude fields, enter the airport code and press Enter. To select from a list of cities worldwide, press Alt-Enter.
- Longitude and Latitude—Geographic location of the site within the network plot. These values are relevant when using geographic backgrounds.
- Screen X and Y—Location of the site within the network plot. These values are relevant when using schematic backgrounds.

Duplicate Nodes

Duplicating a node creates an identical node and duplicates the connectivity of that node by creating new, identical circuits containing the same amount of traffic. By default, each demand and each LSP that uses the node as a source or destination is also duplicated. The newly created node is sequentially named the same. For example, if you duplicated a node named “chi,” MATE Design would create a chi[1] node. If you were to duplicate it again, a chi[2] node would be created.

Step 1 Use one of these methods to access the Duplicate Node dialog box.

- Right-click a node, and select Duplicate from the context menu.
- Select a node, and select the Edit->Duplicate menu.

Step 2 Select whether to duplicate demands and LSPs, and then click OK.

Duplicate Sites

Duplicating a site creates an identical site, including the nodes within it and its connectivity. The default site name and display name are the same as the copied site plus a sequential number. For example, duplicating OKC would create OKC[1]. However, you have the option to change both the site name and the display name. The nodes that are created are likewise sequentially renamed. To change a duplicated node name, open the node’s Properties dialog box.

Note You cannot duplicate a site that contains other sites.

Step 1 Use one of these methods to access the Duplicate Site dialog box.

- Right-click a site, and select Duplicate from the context menu.
- Select a site, and select the Edit->Duplicate menu.

Step 2 If desired, rename the new site and/or display name, and click OK.

Merge Nodes

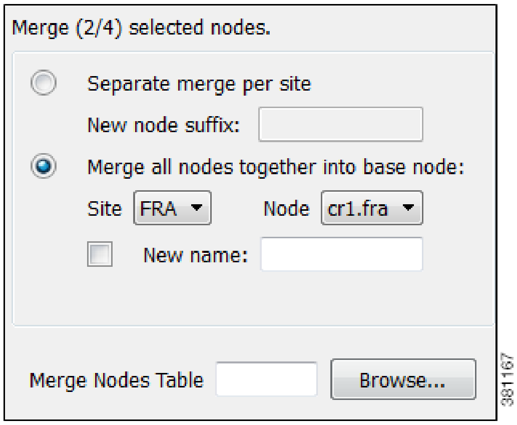

Merge (2/4) selected nodes.

○ Separate merge per site

New node suffix:

● Merge all nodes together into base node:

Site FRA ▾ Node cr1.fra ▾

□ New name:

Merge Nodes Table [Browse...]

Real network topologies often have a number of nodes, typically edge nodes, connected to the network in the same way. For example, they might all be connected to the same core node or pair of core nodes. For planning and design of the network core, it is often desirable to merge these physical nodes into a single virtual node, which simplifies the plan and accelerates the calculations and simulations performed. Note that merging nodes changes the plan itself, not just the visual representation.

The name of a newly merged node can be based on the site name, selected node name (base node), or have a new user-specified name. The node merge effects are as follows.

- Re-attach circuits from other nodes to the base node.
- Move demands to or from other nodes to the base node.
- Move LSPs to or from other nodes to the base node.
- Set base node traffic measurements to the sum of measurements of the selected nodes.
- Delete other nodes.

Step 1 Select multiple nodes. If you do not select nodes, MATE Design merges all nodes.

Step 2 Either select the Initializers->Merge Nodes menu, or right-click on one of the nodes and select Merge from the context menu.

Step 3 Select whether to merge the nodes per site or merge them into one node.

- Separate merge per site—Merges nodes on a per-site basis. For example, if you selected all nodes in the plan, the result would be one merged node per site. If you do not specify a new suffix, the default name is the same as the site.
- Merge all nodes together into base node—Merges all nodes selected into one node. For example, if you selected two nodes in one site and three nodes in another, the result would be a single node in the site and node combination selected as the base node.

  If you do not specify a new name, the default is to use the name of the base node.

Step 4 Click OK.

Trim Nodes and Move Their Traffic

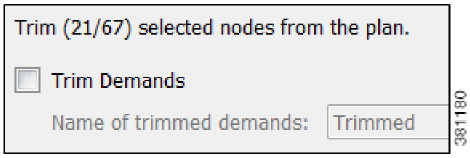
Trim (21/67) selected nodes from the plan.

□ Trim Demands

Name of trimmed demands: Trimmed

A plan file created through Collector discovery often has more edge nodes in the plot than you are interested in viewing or simulating. These might be nodes that forward very little traffic or are singly connected to the core. However, you can reduce such a network model in MATE Design for simplicity and performance improvements by trimming (removing) these nodes. This process also removes all associated circuits, demands, and external endpoints.

Typically, plan files are trimmed before demands are created. However, you can optionally trim demands in a plan file. This moves a demand that is sourced on a trimmed node to the first hop on the demand path in the remaining network. Likewise, a destination to a trimmed node is moved to the last hop of the demand path in the remaining network. If two or more demands in the same service class are trimmed so that their resulting source and destination nodes are equal, these demands are aggregated into one demand with their traffic summed. Demands split by ECMPs are converted into multiple demands, each with traffic divided proportionally to the ECMP split.

Trimming a node is not the same as deleting a node. When you trim a node, the measured traffic carried through that node is transferred to the nodes that remain so that the overall traffic in the network remains unchanged. If the node is deleted, this traffic is discarded. Only measured traffic is preserved in this way.

Note that multicast demands are removed, and cannot be trimmed.

Example: Edge node E is singly connected to node C in the core. When E is trimmed, the measured traffic on the interface from E to C is added to the total source traffic on node C. The measured traffic on the interface from C to E is added to the total destination traffic on node C. By default, demands to and from node E are removed.

Example: A demand source is node A, and destination is node E, and is routed through nodes A-B-C-D-E. If nodes A and E are trimmed and if you select the option to trim demands, then the new route for the demand is B-C-D.

Step 1 Optional: If you want to trim the traffic of nodes under a specific utilization, sort one of the Traffic columns to find the lowest to highest utilization. Alternatively, use the Nodes table filter to search for traffic of a certain level using, for example, the less than operator (<).

Step 2 Select one or more nodes.

Step 3 Right-click and select Trim from the context menu.

Step 4 Optional: Select the option to trim demands.

By default, the moved demands are named “Trimmed,” though you can enter a different name.

Step 5 Click Yes to confirm the action.

Circuits and Interfaces

In MATE Design, an interface is a either an individual logical interface or a LAG logical interface. If there is a one-to-one mapping between a logical and physical interface, then the interface contains both Layer 3 properties (for example, Metric) and physical properties (for example, Capacity). If there is a one-to-many mapping between logical and physical interfaces, then the interface is the logical LAG and the ports are included in the plan file as the physical ports in the LAG. For more information on ports and port circuits, see the Ports, Port Circuits, and LAGs section.

Each circuit connects a pair of interfaces on two different nodes. Therefore, an interface always has an associated circuit. Both the interface Properties dialog box and the circuit Properties dialog box enable you to simultaneously edit properties for the pair of interfaces and the circuit.

Interfaces have both measured and simulated traffic. The traffic that appears in the Circuits table is the higher of the traffic in the two interfaces.

The Latency and Distance Initializer uses mapped L1 circuit delays to calculate L3 circuit delays. If L1 circuits are not mapped to the L3 circuits, node Longitude and Latitude properties are used in the calculation. For information, see the Simulation chapter.

Create Circuits and Interfaces

To create a circuit, either select the Insert->Circuit menu or right-click in an empty plot area and select New->Circuit from the context menu. Following are some of the frequently used fields, with the critical ones being the definition of the interfaces and the capacity.

As a result of creating a circuit, two interfaces are also created.

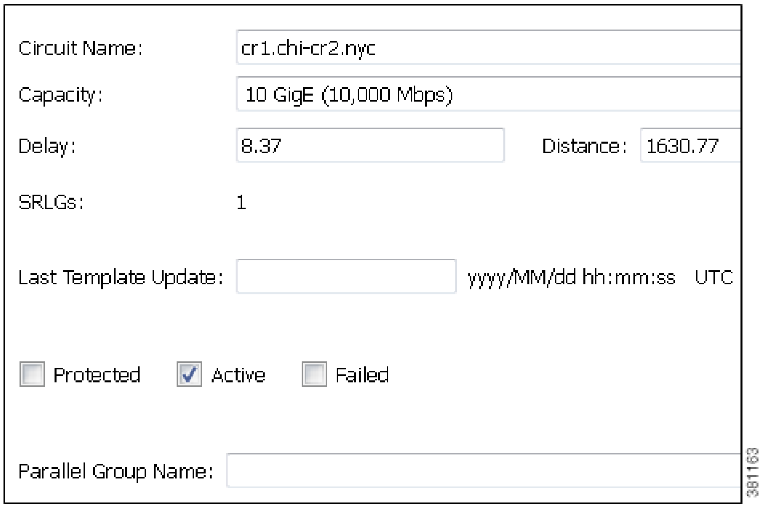

Circuit Name: cr1.chi-cr2.nyc
Capacity: 10 GigE (10,000 Mbps)
Delay: 8.37 Distance: 1630.77
SRLGs: 1
Last Template Update: yyyy/MM/dd hh:mm:ss UTC
□ Protected ☑ Active □ Failed
Parallel Group Name:

- Capacity—The amount of total traffic this circuit can carry. The drop-down list has a selection of the most widely used capacities.
- SRLGs—If you want this circuit to belong to an SRLG, select it from this list or create a new one by clicking the Edit option. See the Create SRLGs for Circuits Only section for more information.
- Parallel Group Name—To include this circuit in a new or existing parallel grouping, enter its name. To visually combine these parallel groups in the network plot, see the *MATE GUI Visualization Guide*.
- Interface A and B—You must specify two interfaces that are connected by the circuit.

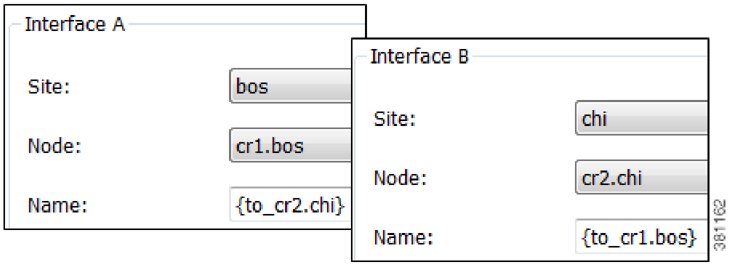

Interface A
Site: bos
Node: cr1.bos
Name: {to_cr2.chi}

Interface B
Site: chi
Node: cr2.chi
Name: {to_cr1.bos}

Duplicate Circuits

Duplicating a circuit has the following effects on the new circuit and its interfaces.

- Duplicates the circuit itself and all its properties, including the circuit name.
- As a result, it duplicates both associated interfaces and renames them by adding an incremental numeric suffix to the original names, such as [1]. All other interface properties are copied to the new interfaces.

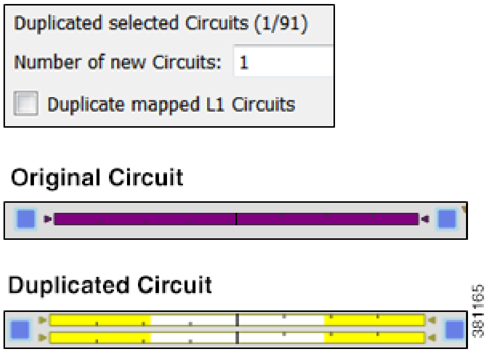

Duplicated selected Circuits (1/91)

Number of new Circuits: 1

□ Duplicate mapped L1 Circuits

Step 1 Select one or more circuits to duplicate.

Step 2 Right-click on one of these circuits, and select Duplicate.

Step 3 Enter the number of new circuits you want to create for each selected circuit. For example, if you selected 2 circuits and you entered a 3, each of those selected circuits would have 3 new parallel circuits for a total of 6 new circuits.

Step 4 To duplicate L1 circuits that are mapped to selected L3 circuits, select “Duplicate mapped L1 Circuits.”

Step 5 Click OK.

Create a Circuit Mesh Between Nodes

To create a mesh of circuits between multiple nodes, follow these steps. Figure 3-3 shows an example circuit mesh created between all the nodes within a site.

Step 1 Select multiple nodes that you want to interconnect.

Step 2 Right-click on the interface you want to duplicate, and select Duplicate Between Selected Nodes from the context menu.

Step 3 Click OK to confirm.

Figure 3-3 Example Circuit Mesh

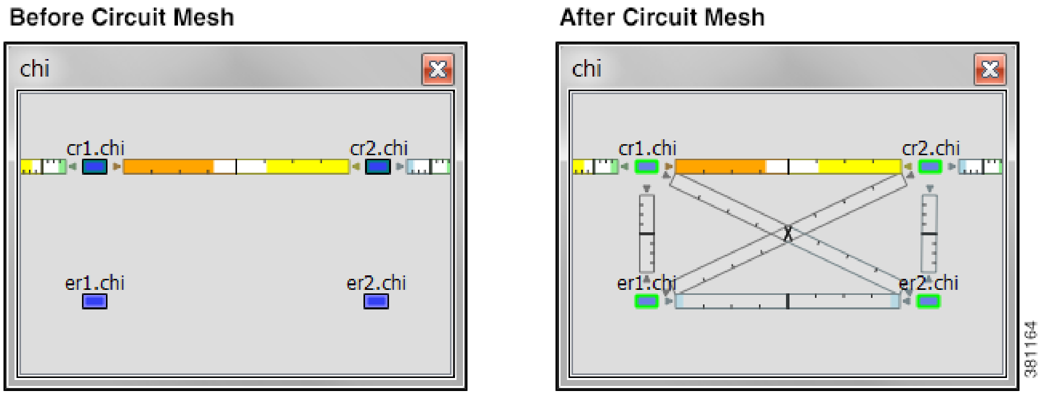

Merge Circuits

You can simplify the plan by merging circuits that have the same source and destination endpoints (nodes). This capability is useful, for example, in long-term capacity planning where multiple parallel circuits can be ignored and only the site-to-site connections are of interest.

Merging circuits changes the plan itself, not just the visual representation. (To change only the visual representation, see the *MATE GUI Visualization Guide*.).

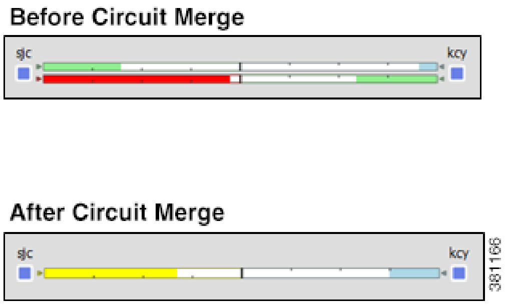

The circuit merge effects are as follows.

- Set base circuit capacity to the sum of capacities.
- Set base circuit metric to minimum of metrics.
- Set base traffic measurements to the sum of measurements.
- Delete other circuits.

Step 1 Select multiple circuits that have the same source and destination endpoints.

Step 2 Either select the Initializers->Merge Circuits menu, or right-click on one of the circuits and select Merge from the context menu.

Step 3 The name of a newly merged circuit is based on the selected circuit. If you would like to base the name on a different circuit, select it from the Base Circuit list.

Step 4 Click OK.

SRLGs

An SRLG is a group of objects that might all fail due to a common cause. For example, an SRLG could contain all the circuits whose interfaces belong to a common line card.

For More Information ...	See ...
Include SRLGs in Fast Reroute simulations	MPLS Simulation chapter
Protect circuits from being included in SRLG failures and SRLG worst-case analysis	Simulation Analysis chapter

Create SRLGs

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

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

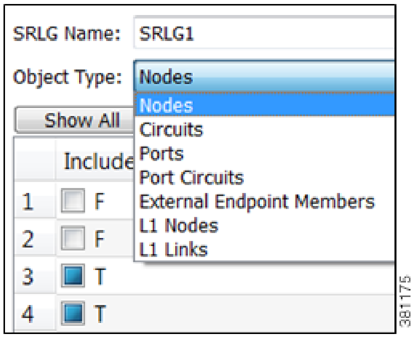

SRLG Name: SRLG1

Object Type: Nodes

Step 2 Enter the SRLG name, or keep the sequentially named default names (SRLG1, SRLG2, etc.).

Step 3 In the Object Type list, select the type of object you want included in the SRLG, which can be nodes, circuits, ports, port circuits, external endpoint members, L1 nodes, or L1 links.

Step 4 For each object you want to include in the SRLG, click on the F (false) check box to mark it T (true).

Step 5 Repeat steps 3 and 4 for each object included, and then click OK.

Create SRLGs for Circuits Only

Step 1 Open the Circuits Properties dialog box.

- If you are creating SRLGs for selected circuits, select one or more circuits. Then double-click one of them. This is a very quick method of adding numerous circuits to an SRLG.
- If you are creating a new circuit, right-click in an empty plot area and select New->Circuit from the context menu. For more information, see the Create Circuits and Interfaces section.

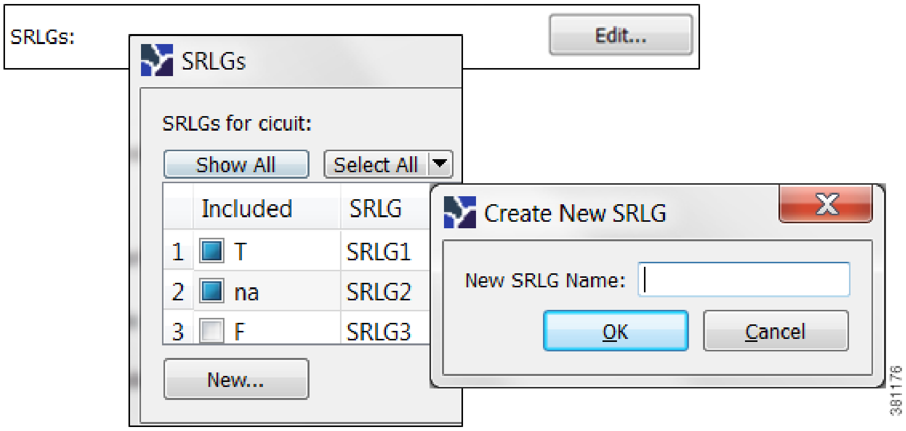

	Included	SRLG
1	■ T	SRLG1
2	■ na	SRLG2
3	□ F	SRLG3

Step 2 Click the Edit button associated with the SRLG field. An SRLG dialog box appears.

Step 3 Either associate the circuits with one or more existing SRLGs or create a new SRLG.

Associate Circuits With Existing SRLGs	Create New SRLG
For each SRLG in which you want to include the selected circuits, click the F (false) or "na" check box to make it T (true). An "na" next to an SRLG means that one or more of the selected circuits are contained in that SRLG, but not all circuits are contained in it.	**1.** Click the New button. A Create New SRLG dialog box appears. **2.** Enter the new SRLG name, and click OK.

Step 4 Click OK in the SRLG dialog box to save the changes, and click OK again in the Properties dialog box.

Ports, Port Circuits, and LAGs

In MATE Design, a port is a physical interface. You can model LAGs[1] and port channels using MATE port and port circuits (Figure 3-4).

By default, each logical interface listed in the Interfaces table corresponds to a single physical port, and these ports need not be explicitly modeled. The exception is when the logical interface is a LAG, which bundles more than one physical port. In this case, the physical ports are listed in the Ports table.

A port circuit is a connection between two ports. However, ports are not required to be connected to other ports by port circuits.

Port circuits can be mapped to L1 circuits, and ports can be connected to L1 nodes. For information, see the Layer 1 Simulation chapter.

[1] A link aggregation group (LAG) is a group of physical ports that are bundled into a single logical interface. A LAG is also known as *bundling* or *trunking*.

Figure 3-4 Ports, Port Circuits, and LAGs

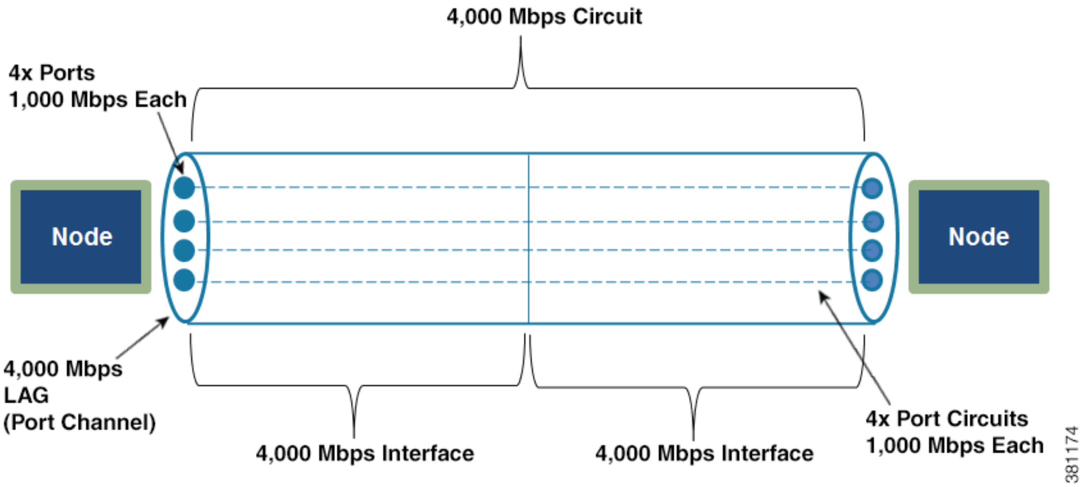

Create Ports

To create a port, either select the Insert->Port menu or right-click in an empty plot area and select New->Port from the context menu. Following are some of the frequently used fields.

- Site and Node—Site and node on which this port exists.
- Interface—Logical interface to which this port is mapped. This must be defined to create port circuits using this port.
- Capacity—The amount of total traffic this port can carry. The drop-down list has a selection of the most widely used capacities.

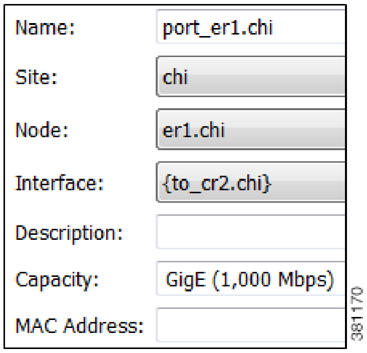

Name:	port_er1.chi
Site:	chi
Node:	er1.chi
Interface:	{to_cr2.chi}
Description:	
Capacity:	GigE (1,000 Mbps)
MAC Address:	

Each port can be connected to an L1 node. Creating these connections populates the L1 Node column in the Ports table. If this port belongs to a port circuit, the Remote L1 Node column is also updated. To connect the port to an L1 node, select the Layer 1 tab. Then select the site containing the L1 node and also the L1 node to which this port is connected. For more information on Layer 1, see the Layer 1 Simulation chapter.

Create LAGs

To make one of the existing interfaces into a LAG by assigning ports to it, follow these steps. If the interface does not contain ports, you must first create them (see Create Ports).

Step 1 From the Ports table, select all ports that belong to the LAG. You can quickly do this by filtering to their common interface.

Step 2 Double-click one of the ports to open the Properties dialog box.

a. Select the interface that contains these ports.

b. Click OK.

Create Port Circuits

A port circuit specifies a pair of connected ports.

- Both of the ports must exist and be mapped to interfaces that are connected by a circuit.
- When selecting the two ports for the port circuit, note that if one is assigned to an interface, the other must be assigned to the remote interface on the same circuit.

To create ports and map them to their interfaces, see Create Ports.

To create a port circuit, either select the Insert->Port Circuit menu or right-click in an empty plot area and select New->Port Circuit from the context menu. The site, node, and port are all required fields.

- Site and Node—Site and node on which the port exists.
- Port—Name of the port.
- Capacity—The amount of total traffic this port circuit can carry. The drop-down list has a selection of the most widely used capacities.

Each port circuit can be mapped to an L1 circuit. Creating these mappings populates the L1 Circuit Name, L1 Circuit Node A, and L1 Circuit Node B columns in the Port Circuits table. For more information on Layer 1, see the Layer 1 Simulation chapter.

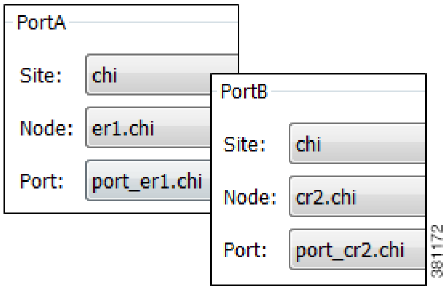

PortA
Site: chi
Node: er1.chi
Port: port_er1.chi

PortB
Site: chi
Node: cr2.chi
Port: port_cr2.chi

Step 1 Select the Layer 1 tab.

Step 2 Select the sites and the L1 nodes within them on both ends of the L1 circuit.

Step 3 Either map an existing L1 circuit or create a new one, and then click OK.

- To map an existing L1 circuit, click Select and then select the L1 circuit from the drop-down list.
- To create an L1 circuit, click Create. Either enter a new L1 circuit name or use the default one that is shown. Once you click OK, you have the option to create L3-L1 links.

Set LAG Simulation Properties

Each interface is considered to be a LAG (port channel). You can configure LAG properties so that if it loses too much capacity due to inoperational ports, the entire LAG will be taken down.

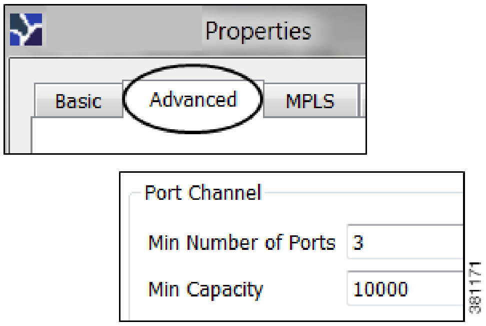

Port Channel
Min Number of Ports 3
Min Capacity 10000

Step 1 Double-click on an interface or circuit; the Properties dialog box appears.

Step 2 Select the Advanced tab and go to the Port Channel section. You can set one or both of the following parameters, and you can set them for either one or both interfaces on the circuit.

a. In the Min Number of Ports field, enter the minimum number of ports that must be active and operating for the LAG circuit to be up.

b. In the Min Capacity field, enter the minimum capacity that must be available for the LAG circuit to be up.

Step 3 Click OK.

Edit Object Properties

Most objects in the plot and in the associated tables have a set of properties that you can manage using a Properties dialog box. These are the properties that MATE Design uses to define and simulate an object.

To change the properties of individual objects, follow one of these two steps to open the Properties dialog box.

- Right-click on one or more objects from the network plot or a table, and then select Properties from the context menu.
- Double-click on one or more objects from the network plot or table. (The exception is that double-clicking a site from the plot opens a site plot.)

The Properties dialog box always lists how many objects were selected prior to opening the dialog box. For example, if the plan file includes 500 interfaces, and you have 342 of them selected, the notation is 342/500.

Tags

Many objects can be *tagged* as a way to group them for subsequent processing. Once created, tags appear as options in dialog boxes, enabling you to select and act upon all objects with a given tag name. Thus, tags provide a very flexible way to group interfaces, demands, nodes, LSPs, sites, and AS's that are not easily selected using standard attributes. For example, when using LSP paths you might want to tag selected paths to make it easier to assign or change their path affinities.

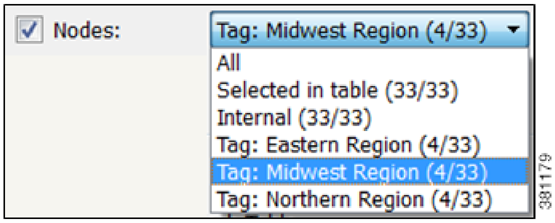
☑ Nodes:
Tag: Midwest Region (4/33)
All
Selected in table (33/33)
Internal (33/33)
Tag: Eastern Region (4/33)
Tag: Midwest Region (4/33)
Tag: Northern Region (4/33)
381179

Adding tags to a parent site does not add tags to its children sites and nodes.

Add Tags

Step 1 Right-click one or more objects and select its Properties menu.

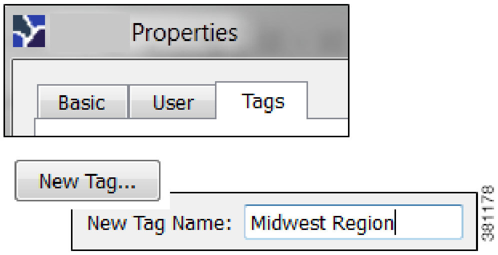

New Tag Name: Midwest Region

Step 2 Select the Tags tab.

Step 3 Click New Tag; a New Tag dialog box appears.

  a. Enter a name that identifies the objects you are tagging.

  b. Click OK.

Step 4 Click OK.

Enable and Disable Tags

Step 1 Right-click one or more objects and select its Properties menu.

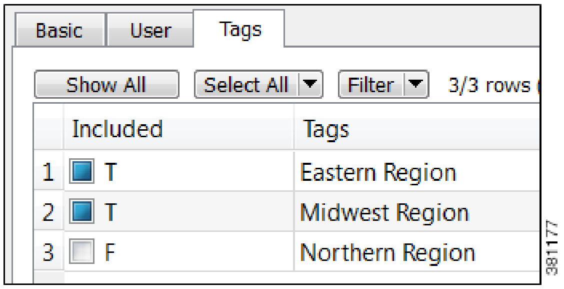

	Included	Tags
1	■ T	Eastern Region
2	■ T	Midwest Region
3	□ F	Northern Region

Step 2 Select the Tags tab; a list of available tags appears.

Step 3 Toggle the tags to enable (T = true) or disable (F = false).

Step 4 Click OK.

CHAPTER 4

Traffic Demand Modeling

MATE 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 AS's, peering nodes in these AS's, or interfaces to these peering nodes.

Each demand has a specified amount of traffic. There are a variety of 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 MATE Design simulation and traffic engineering tools. As such, an accurate set of demands and demand traffic is essential for effective planning, designing, engineering, and operating of a network. Further, 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 or specified, as well as modified. As an overview, following are a few example use cases for demands and steps you would take to deploy them.

Use Demands To...	Suggested Steps to Take...
Model discovered networks	1. Create a demand mesh based on where the traffic originates. For example, if all traffic is between edge routers, then 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	1. Create a demand mesh. 2. After setting the traffic, use MATE 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 the Forecasting Traffic chapter).

Use Demands To...	Suggested Steps to Take...
Design networks	1. Create a demand mesh. 2. Set the demand traffic using methods described in the Forecasting Traffic chapter.
Analyze existing plans	Use a variety of MATE Design tools that rely on demand traffic, such as the following. • Simulation Analysis enables you to run what-if analysis. • Manually editing the network topology and configuration to determine the effects on traffic utilization. • Traffic Engineering tools enable you to control routings and optimize utilizations.

Demands

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

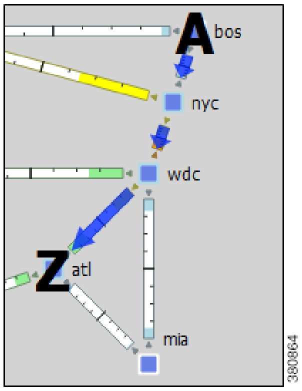

Unique Properties (Keys)	Each demand is defined by a unique combination of these four properties.
	Name—By default, this is blank.
	Source—Nodes, interfaces, external AS’s, or external endpoints.
	Destination—Nodes, interfaces, external AS’s, external endpoints, or multicast destinations.
	Service class—User-defined classification of traffic, such as for voice or video. The default service class for demands is named Default.

Commonly Used Properties	Latency bound—Policy that sets the maximum permissible latency on a demand under normal operation. This property is used by MATE Design traffic engineering tools.
	Topology—Demands can be assigned to a specific IGP, and only route through interfaces belonging 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 enables you to group demands for subsequent processing. Most tools, for example, allow you to choose a set of tags on which to operate.
	Private LSP Name and Source—If a demand is associated with a private LSP, then the demand can only route through that LSP, and the only demand that is permitted to cross that LSP is this demand.
Traffic	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, following are a few recommendations.

- For internal routing, use nodes.
- For external AS's, use a combination of AS's, nodes, and interfaces. Using interfaces enables you to 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 4-1 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 4-1 Example of Demands Destined for a Node Versus an Interface

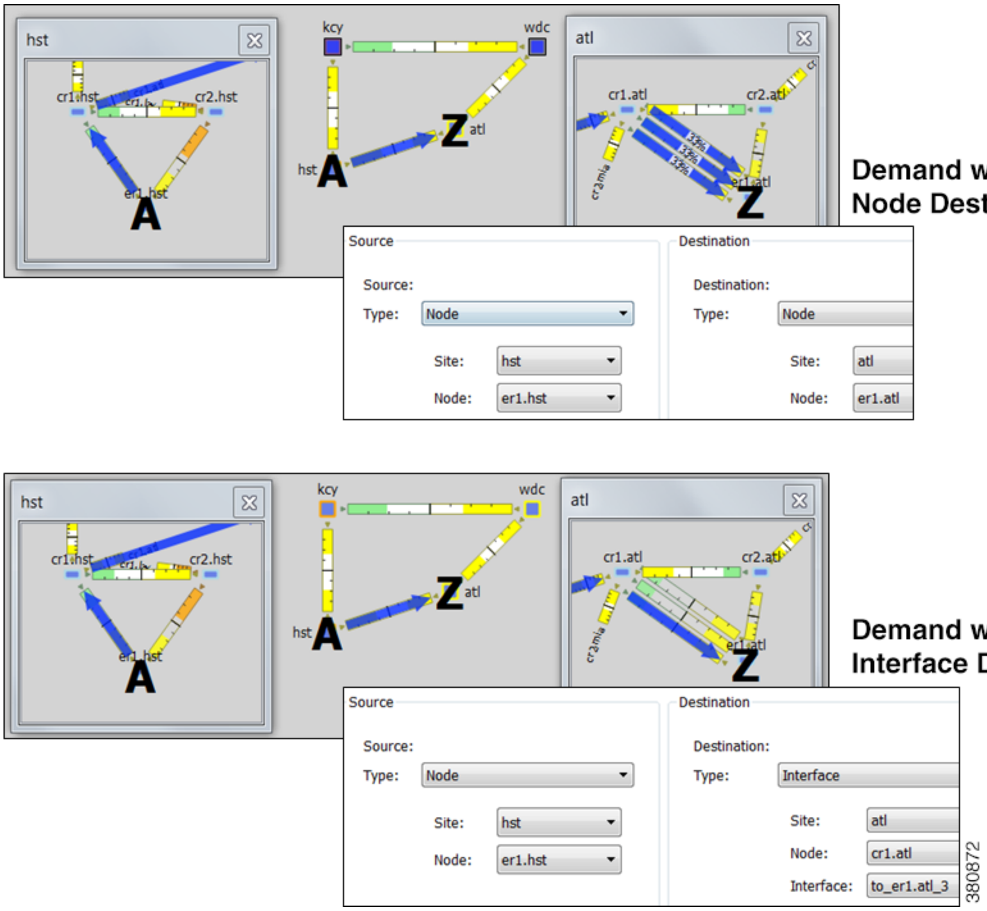

Demand with Node Destination

Source

Source:

Type: Node

Site: hst

Node: er1.hst

Destination

Destination:

Type: Node

Site: atl

Node: er1.atl

Demand with Interface Destination

Source

Source:

Type: Node

Site: hst

Node: er1.hst

Destination

Destination:

Type: Interface

Site: atl

Node: cr1.atl

Interface: to_er1.atl_3

For More Information...	See...
Create demands	Create Demands section
External endpoints	Advanced Routing with External Endpoints chapter
Multicast demands	Multicast Simulation chapter
Demand and demand grouping cost modeling	Cost Modeling chapter
Demand groupings	Forecasting Traffic chapter
Import demand groupings	*MATE Integration and Development Guide*
Put traffic into demands	• Demand Traffic • Demand Deduction
Import demand grouping traffic	*MATE Integration and Development Guide*

Demand Meshes

Demand meshes are a time-efficient way of creating numerous demands for all or part of the network. By default, MATE Design creates a source-destination mesh amongst either all or the selected nodes, interfaces, external AS's, 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 demand mesh <DemandMesh> table that contains estimated traffic. MATE Design uses these estimations to calculate a full mesh of demands.

For More Information...	See...
Create demand meshes	Create Demand Meshes
Create demand meshes for LSPs	Create Demand Meshes for LSPs

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 sub-routes.
- Minimum Latency—Minimum latency over all ECMP sub-routes.
- Maximum Latency—Maximum latency over all ECMP sub-routes.
- Latency Bound—The maximum permissible latency on a demand.
- Min Possible Latency—The 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.

For More Information...	See...
Create demand latency bounds	Create Demand Latency Bounds
Set latency and distance	Simulation chapter

Visualize 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.

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

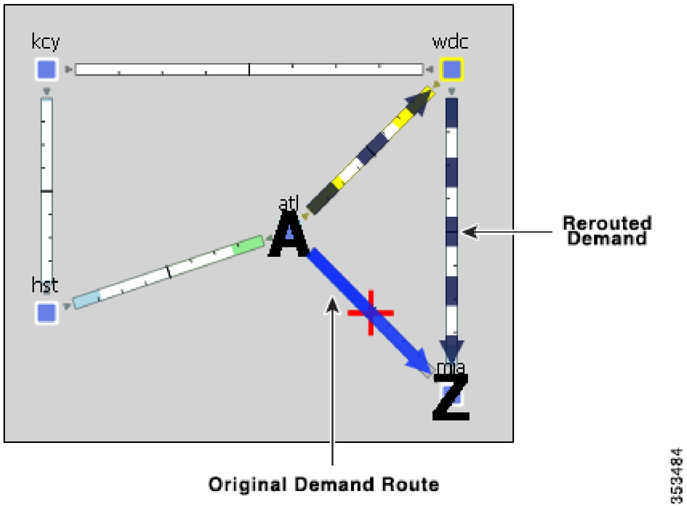

Rerouted Demand

Original Demand Route

View Demand Plots

Demand routes can be complex since they can go through multiple sites that could potentially hide the details of the path. For example, the network plot might not show the individual nodes traversed by the demand, as in Figure 4-2. The demand plot offers an easier way to view end-to-end demands, enabling you to view the hops that the demands take without having to individually select sites to view them, such as in Figure 4-3.

Although you can open a demand plot for only one demand, you can repeat these steps to open multiple demand plots.

Step 1 Select a demand from the Demand table.

Step 2 Right-click and select Plot Demand from the context menu.

To move the plot horizontally or vertically, use the left/right arrows and up/down arrows, respectively.

Figure 4-2 Network View of Selected Demand

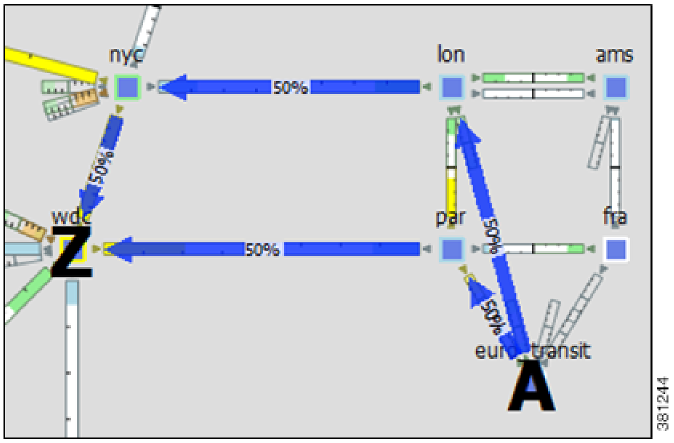

Figure 4-3 Demand Plot View

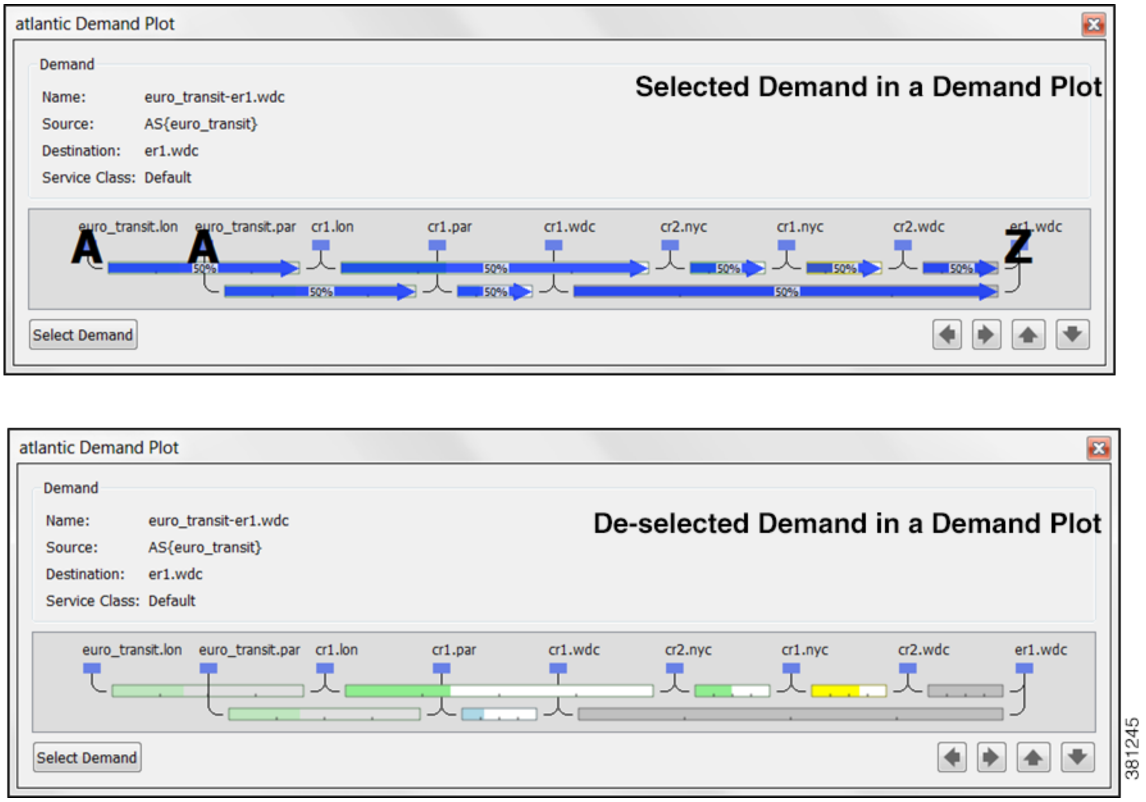

atlantic Demand Plot

Selected Demand in a Demand Plot

Demand

Name: euro_transit-er1.wdc
Source: AS{euro_transit}
Destination: er1.wdc
Service Class: Default

atlantic Demand Plot

De-selected Demand in a Demand Plot

Demand

Name: euro_transit-er1.wdc
Source: AS{euro_transit}
Destination: er1.wdc
Service Class: Default

View Shortest Paths and Shortest Latency Paths

For a description of shortest IGP paths and shortest latency paths, see the IGP Simulation chapter.

To View...	Select this View->Demands Menu...	Example
Shortest IGP path	Shortest IGP Paths	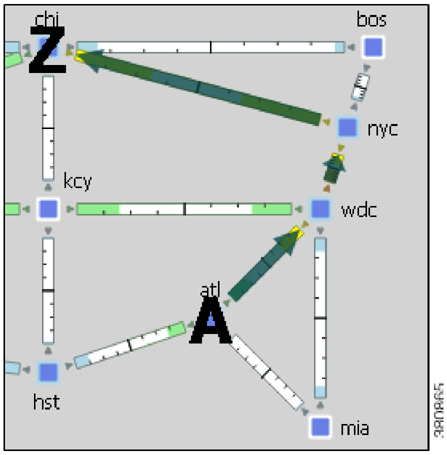 chi Z bos nyc kcy wdc atl A hst mia 380865
Shortest latency path	Shortest Latency Paths	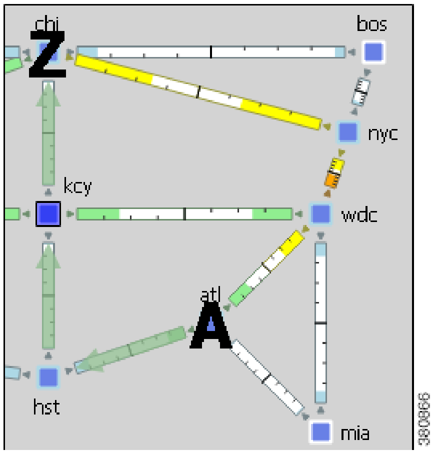 chi Z bos nyc kcy wdc atl A hst mia 380866

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.

For More Information...	See...
Demand Deduction	• Demand Deduction section • Estimate Demand Traffic Using Demand Deduction
Generate demand traffic using node traffic estimates	Generate Demand Traffic Using Node Traffic Estimates section
Modify demand traffic	Modify Demand Traffic section

For More Information...	See...
Enter traffic values directly editing the Plan Table Database Editor using the GUI	• *MATE GUI Visualization Guide* • *MATE Integration and Development Guide*
Use demand traffic growth and demand grouping traffic in growth plans	Forecasting Traffic chapter
Import demand traffic <DemandTraffic> and growth rates	*MATE Integration and Development Guide*
Create demand grouping traffic table	*MATE Integration and Development Guide*

Traffic Levels

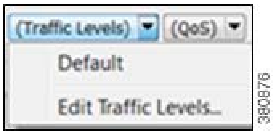

Demands can be associated with specific traffic levels, which are identified in the Traffic Levels drop-down list in the Visualization tool bar. The traffic level displayed determines what is displayed in the plot and table calculations. For example, consider a totally static network that has no changes 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 4-4).

Step 1 Select Edit Traffic Levels from the Traffic Levels drop-down list. The Edit Traffic Levels dialog box appears.

Step 2 You can either click New, or select an existing traffic level and then click Duplicate.

Step 3 In the New Traffic Level dialog box, enter the name and click OK.

Step 4 Click Done in the Edit Traffic Levels dialog box. The traffic level now appears in the Traffic Level drop-down list.

Figure 4-4 Create and Edit Traffic Levels

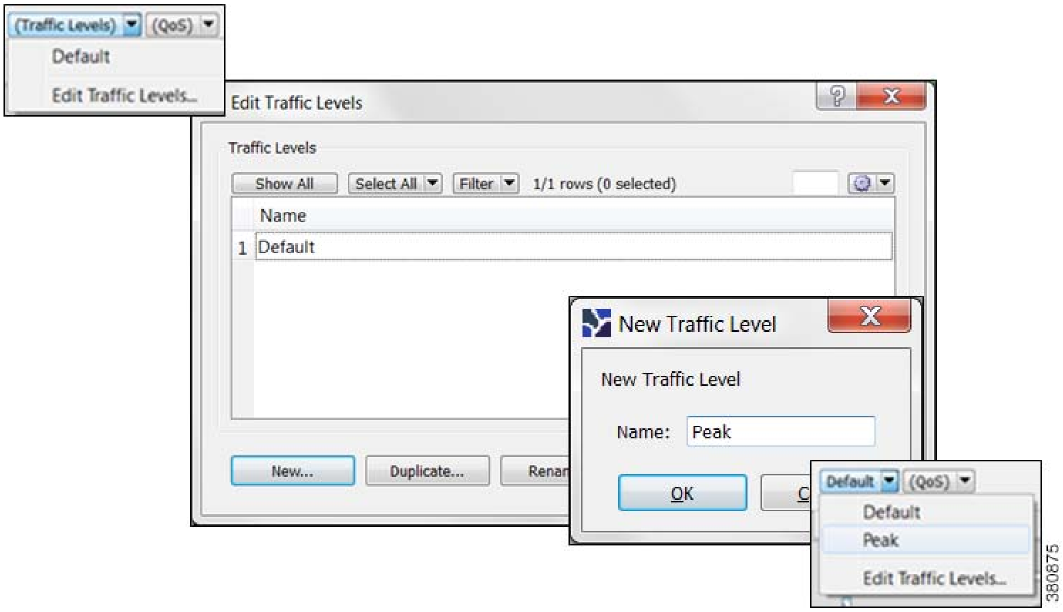

	Name
1	Default

Name: Peak

Demand Deduction

Plan files created by

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, for example 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. For information, see the Advanced Routing with External Endpoints chapter.

For More Information...	See...
Discover measured traffic	*WAE Platform Configuration Guide*
Common causes of suboptimal results	Differences in Measured and Simulated Traffic

For More Information...	See...
Improve accuracy of Demand Deduction results	Minimize Differences Between Measured and Simulated Traffic
Flow measurements used in Demand Deduction	• Flow Measurements in Demand Deduction section • *MATE Integration and Development Guide*
Set up external endpoints to use Demand Deduction estimates on members	Advanced Routing with External Endpoints chapter

Example Demand Deduction

This example demonstrates results when using the Demand Deduction tool on a very simple network. Figure 4-5 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 since these demands do not yet contain traffic. Figure 4-6 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 4-5 Network Containing Two Demands and No Demand Traffic

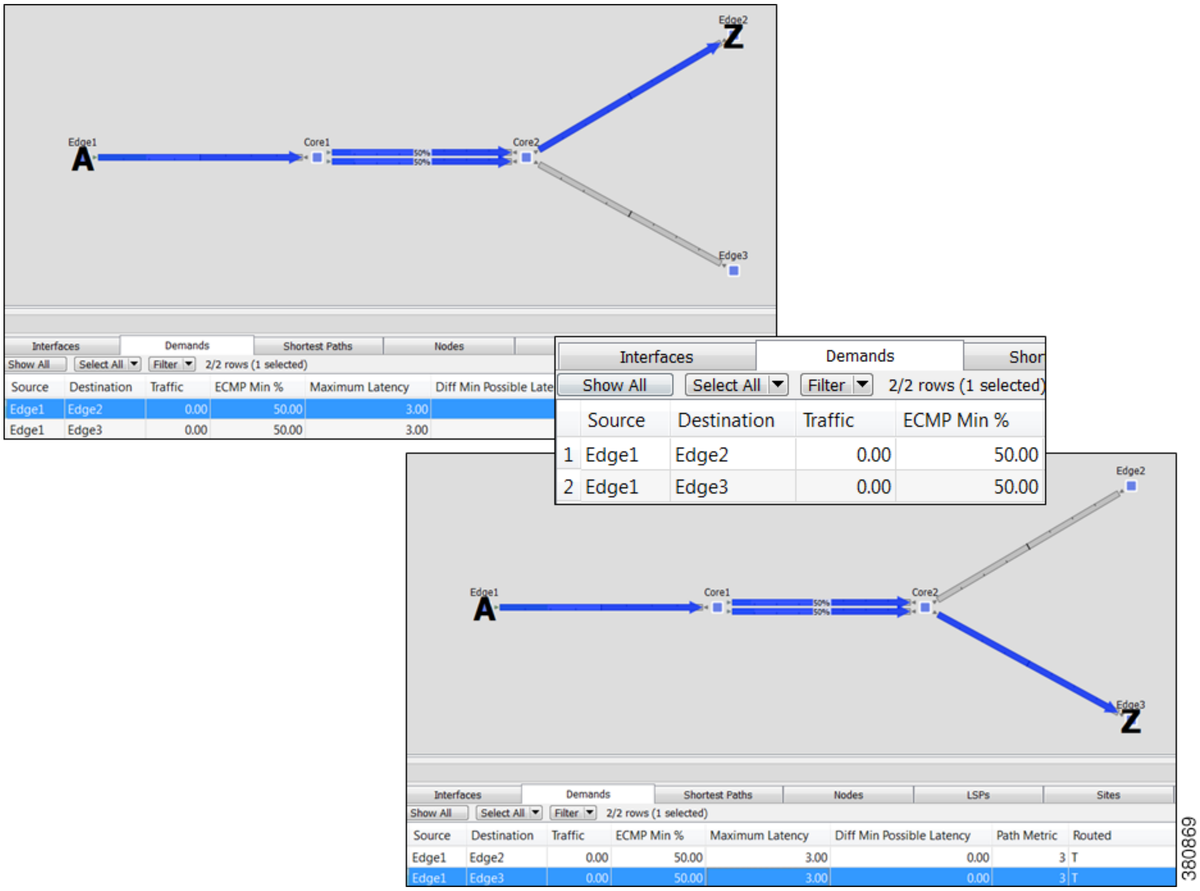

2/2 rows (1 selected)

Source	Destination	Traffic	ECMP Min %	Maximum Latency	Diff Min Possible Late
Edge1	Edge2	0.00	50.00	3.00	
Edge1	Edge3	0.00	50.00	3.00	

2/2 rows (1 selected)

	Source	Destination	Traffic	ECMP Min %
1	Edge1	Edge2	0.00	50.00
2	Edge1	Edge3	0.00	50.00

2/2 rows (1 selected)

Source	Destination	Traffic	ECMP Min %	Maximum Latency	Diff Min Possible Latency	Path Metric	Routed
Edge1	Edge2	0.00	50.00	3.00	0.00	3	T
Edge1	Edge3	0.00	50.00	3.00	0.00	3	T

Figure 4-6 ***Measured Traffic View and Interfaces Associated with the Demands***

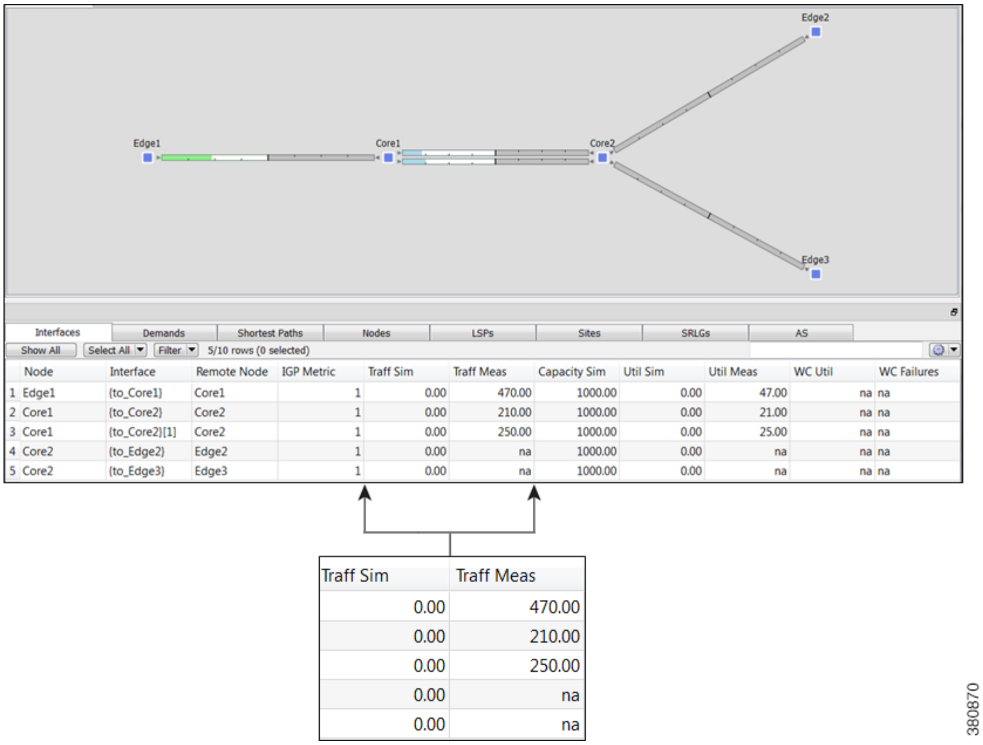

	Node	Interface	Remote Node	IGP Metric	Traff Sim	Traff Meas	Capacity Sim	Util Sim	Util Meas	WC Util	WC Failures
1	Edge1	{to_Core1}	Core1	1	0.00	470.00	1000.00	0.00	47.00	na	na
2	Core1	{to_Core2}	Core2	1	0.00	210.00	1000.00	0.00	21.00	na	na
3	Core1	{to_Core2}[1]	Core2	1	0.00	250.00	1000.00	0.00	25.00	na	na
4	Core2	{to_Edge2}	Edge2	1	0.00	na	1000.00	0.00	na	na	na
5	Core2	{to_Edge3}	Edge3	1	0.00	na	1000.00	0.00	na	na	na

Traff Sim	Traff Meas
0.00	470.00
0.00	210.00
0.00	250.00
0.00	na
0.00	na

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 4-7). 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 since they had no measured traffic.

Figure 4-7 Simulated View Showing Demand Traffic

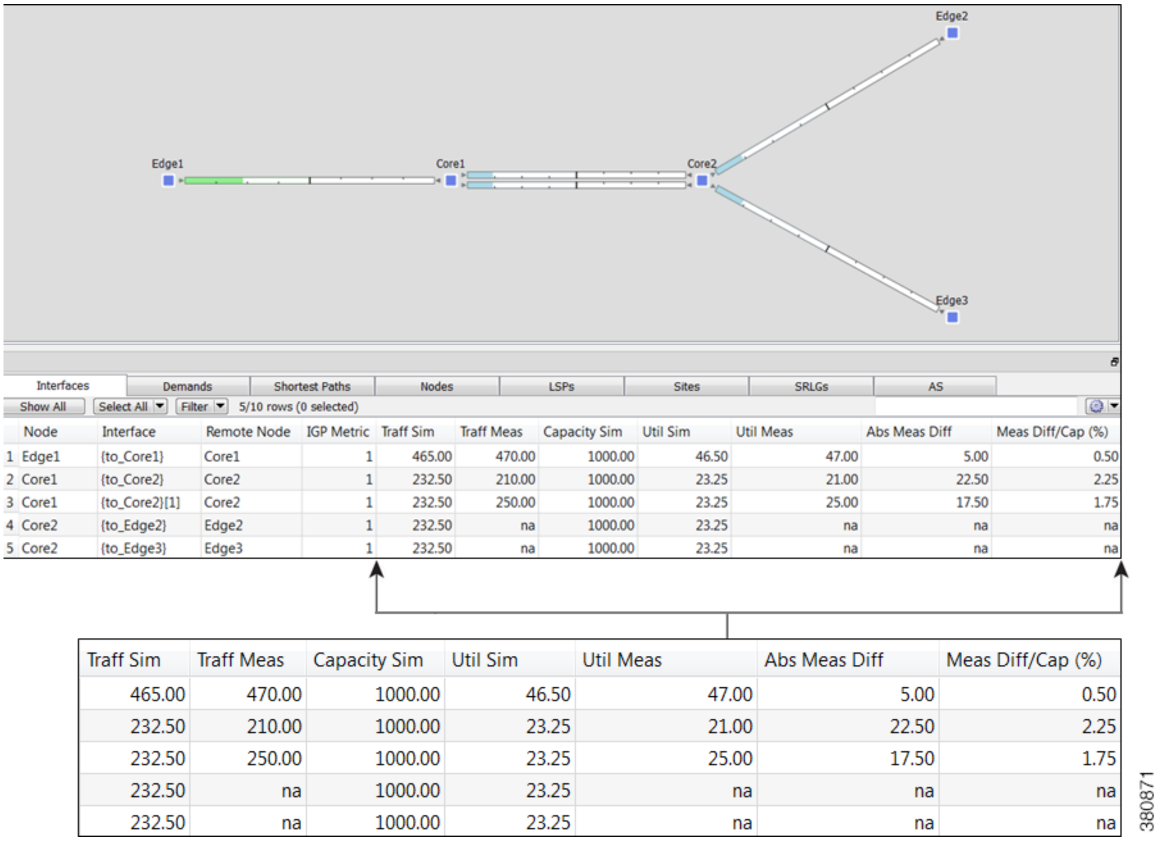

	Node	Interface	Remote Node	IGP Metric	Traff Sim	Traff Meas	Capacity Sim	Util Sim	Util Meas	Abs Meas Diff	Meas Diff/Cap (%)
1	Edge1	{to_Core1}	Core1	1	465.00	470.00	1000.00	46.50	47.00	5.00	0.50
2	Core1	{to_Core2}	Core2	1	232.50	210.00	1000.00	23.25	21.00	22.50	2.25
3	Core1	{to_Core2}[1]	Core2	1	232.50	250.00	1000.00	23.25	25.00	17.50	1.75
4	Core2	{to_Edge2}	Edge2	1	232.50	na	1000.00	23.25	na	na	na
5	Core2	{to_Edge3}	Edge3	1	232.50	na	1000.00	23.25	na	na	na

Traff Sim	Traff Meas	Capacity Sim	Util Sim	Util Meas	Abs Meas Diff	Meas Diff/Cap (%)
465.00	470.00	1000.00	46.50	47.00	5.00	0.50
232.50	210.00	1000.00	23.25	21.00	22.50	2.25
232.50	250.00	1000.00	23.25	25.00	17.50	1.75
232.50	na	1000.00	23.25	na	na	na
232.50	na	1000.00	23.25	na	na	na

In this same example, if the Core2 to Edge2 interface had 50 Mbps traffic, the results would have been different. Since 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 MATE 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 are showing large values, then it is likely that one of the following situations exist.

- 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—MATE 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 prioritities—In the Demand Deduction dialog box, you have the option to set the priority for calculations as 1 or 2. MATE 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, upon running Demand Deduction warnings are displayed if you might be receiving 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 any of the specified measurements exceed the corresponding circuit capacity.

Minimize 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. Following is a list of suggestions to consider for enhancing the accuracy of Demand Deduction results.

- Include RSVP LSP or LDP measurements in the network discovery process for inclusion in the snapshot.
- 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 since 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 4-1 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 4-1 Flows Table Columns

Column	Description
From	Source node
FromType	• 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	• 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.

Create Demands and Demand Meshes

Create Demands

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

Step 1 Either select the Insert->Demands->Demand menu or right-click in an empty plot area and then select New->Demands->Demand from the context menu. The New Demand dialog box appears.

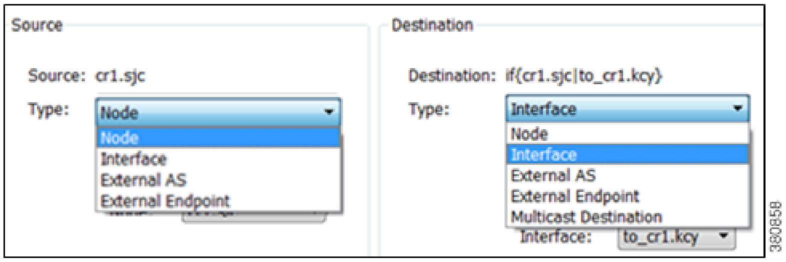

Source

Source: cr1.sjc

Type: Node

Destination

Destination: if{cr1.sjc|to_cr1.kcy}

Type: Interface

Interface: to_cr1.kcy

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

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

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

Step 5 Select the service class. If there are on service classes, the demand operates on a service class named Default.

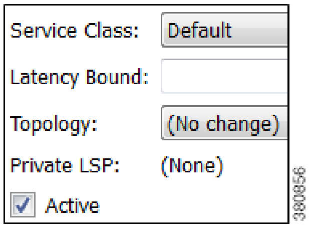

Service Class: Default

Latency Bound:

Topology: (No change)

Private LSP: (None)

☑ Active

Step 6 Enter a value for the latency bound.

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

Step 8 Keep the Active default to include the demand in MATE Design simulations, or deselect Active to exclude this demand from simulations.

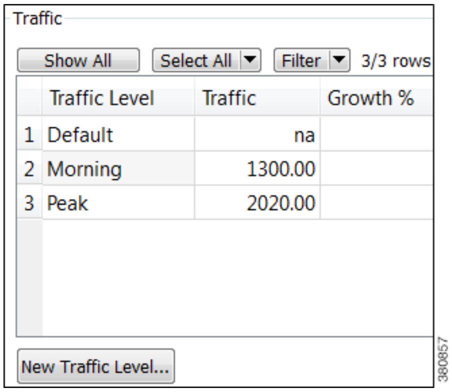

	Traffic Level	Traffic	Growth %
1	Default	na	
2	Morning	1300.00	
3	Peak	2020.00	

Step 9 For each traffic level, either specify the amount of traffic or leave it empty for Demand Deduction to complete. You also have the option to create a traffic level if one does not exist (New Traffic Level button).

Step 10 For each traffic level, specify the amount of growth rate you want to use for forecasting purposes. For more information, see the Forecasting Traffic chapter.

Step 11 Click OK.

Duplicate Demands

Step 1 Select one or more demands.

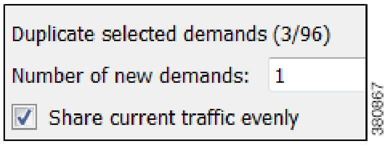

Duplicate selected demands (3/96)

Number of new demands: 1

☑ Share current traffic evenly

Step 2 Right-click on one of them, and select Duplicate from the context menu.

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 1,000 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 1,000 Mbps of traffic.

Step 5 Click OK. Duplicated demands are named sequentially.

Create Demand Meshes

Step 1 To restrict the demand mesh source to a specified set of nodes, external AS’s, or external endpoints, select them before opening the Demand Mesh dialog box.

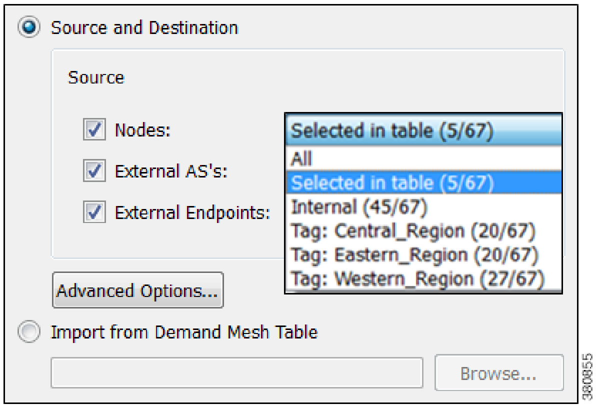

● Source and Destination

Source

☑ Nodes: Selected in table (5/67)

☑ External AS's:

☑ External Endpoints:

○ Import from Demand Mesh Table

All
Selected in table (5/67)
Internal (45/67)
Tag: Central_Region (20/67)
Tag: Eastern_Region (20/67)
Tag: Western_Region (27/67)
Advanced Options...
Browse...
380855

Step 2 Either select the Insert->Demands->Demand Mesh menu, or right-click in the plot and select New->Demands->Demand Mesh from the context menu. The Insert Demand Mesh dialog box appears.

Step 3 Either 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 AS's, 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 *MATE Integration and Development Guide*.

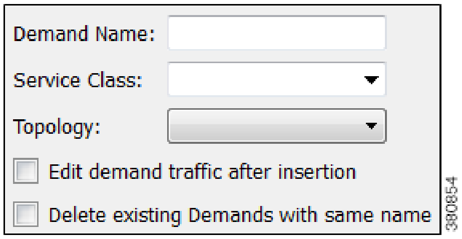

Demand Name:

Service Class:

Topology:

☐ Edit demand traffic after insertion

☐ Delete existing Demands with same name

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 Select a service class. The default is the Default service class.

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

Step 7 If you would like to create demand traffic immediately upon creating the demands, select the "Edit the demand traffic after insertion" option. For information on entering traffic in the dialog box that appears, see Generate Demand Traffic Using Node Traffic Estimates. 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 the Demand Deduction section.

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

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

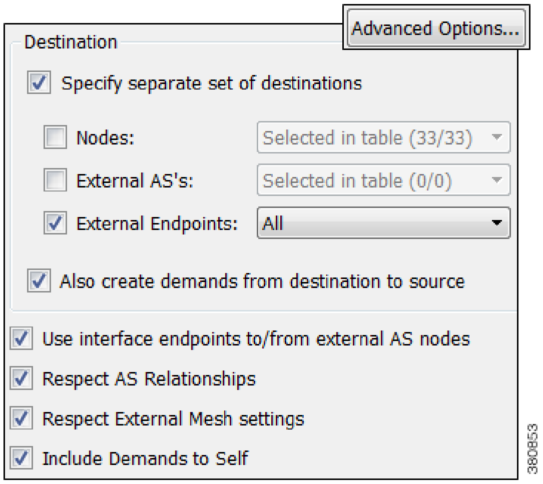

Destination

☑ Specify separate set of destinations

□ Nodes: Selected in table (33/33)

□ External AS's: Selected in table (0/0)

☑ External Endpoints: All

☑ Also create demands from destination to source

☑ Use interface endpoints to/from external AS nodes

☑ Respect AS Relationships

☑ Respect External Mesh settings

☑ Include Demands to Self

- 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—Deselect 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 AS's, 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 the BGP Simulation chapter.
- Respect AS Relationships—If selected, keep the existing AS relationships defined by the Routing Policy (default). If unselected, 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 the BGP Simulation chapter.
- Respect External Mesh Settings—If selected, keep the existing External Mesh settings defined for external AS meshes (default). If unselected, 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.

Create Demand Meshes for LSPs

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

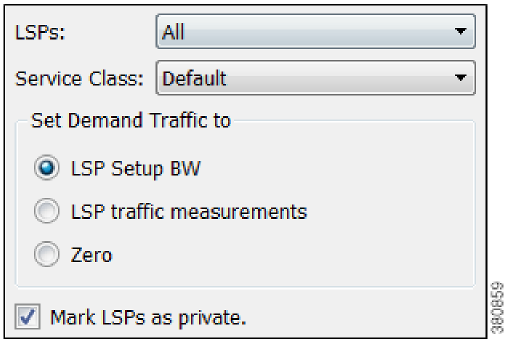

LSPs: All

Service Class: Default

Set Demand Traffic to

● LSP Setup BW

○ LSP traffic measurements

○ Zero

☑ Mark LSPs as private.

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

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

Step 4 Select the service class for the resulting demands.

Step 5 Select 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, deselect the “Mark LSPs as private” option. Otherwise, the default is to restrict these LSPs so that they use only the resulting demands, and vise versa, that is, restrict the demands to run only over these LSPs.

Step 7 Click OK.

Create Demand Latency Bounds

You can set demand latency bounds to fixed values or to the following relative values (Table 4-2). 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.

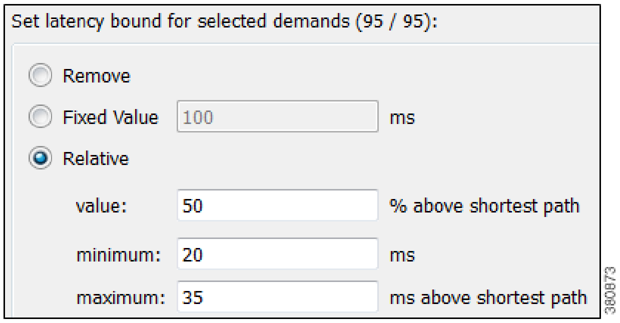

Set latency bound for selected demands (95 / 95):

○ Remove

○ Fixed Value 100 ms

● Relative

value: 50 % above shortest path

minimum: 20 ms

maximum: 35 ms above shortest path

  - 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.

    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.
  - 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 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 4-2 Editing Demand Latency Bounds

Using Demands Properties Dialog Box	Using the Initializer Tool
1. Select one or more demands in the Demands table. 2. Right-click on one of these demands, and select 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.	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. Select the Initializers->Demand Latency Bounds menu. 3. You have several options. – To delete latency bounds, select Remove. – To set a fixed value, enter this number in the Fixed Value field. – To set a relative value, select Relative and enter desired values in value, minimum, and maximum fields. 4. Click OK.

Add and Modify Demand Traffic

Estimate 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 the Minimize Differences Between Measured and Simulated Traffic section. For information on setting up external endpoint members to be included in Demand Deduction calculations, see the Advanced Routing with External Endpoints chapter.

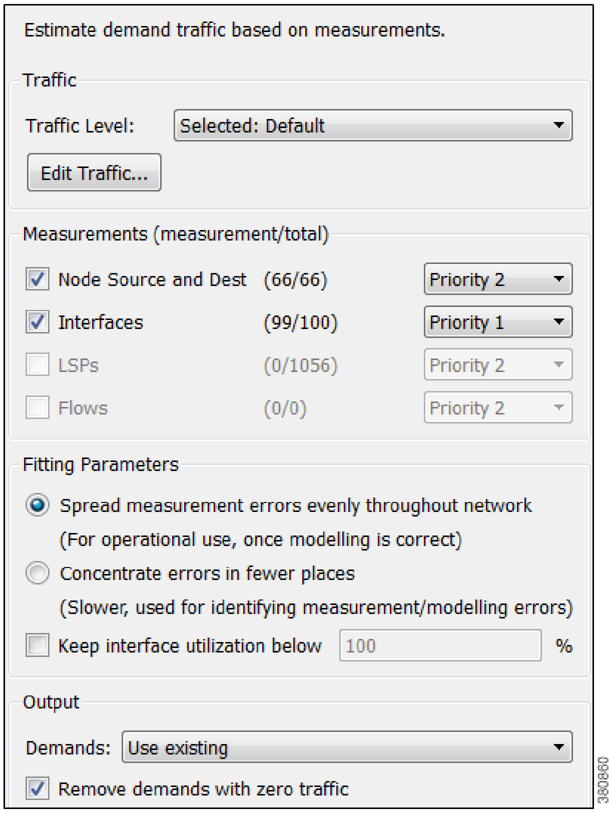
Estimate demand traffic based on measurements.

Traffic

Traffic Level: Selected: Default

Measurements (measurement/total)

☑ Node Source and Dest (66/66) Priority 2
☑ Interfaces (99/100) Priority 1
□ LSPs (0/1056) Priority 2
□ 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
☑ Remove demands with zero traffic

Step 1 If applicable, select one or more demands.

Step 2 Select the Tools->Demand Deduction menu. The Demand Deduction dialog box appears.

Step 3 Select 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 the Forecasting Traffic chapter.

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 Select the Fitting Parameters.

Step 8 If you need to keep the traffic utilization below a different percentage than 100% (default), select the “Keep utilization below” option and enter a value in its field.

Step 9 Select 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 remove demands with zero traffic. The default is to remove them since 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 11 Verify your parameters, and click OK. Demand Deduction calculates the simulated traffic and lists the results in a Demand Deduction report.

Generate 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 the Forecasting Traffic chapter.

- 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.

- Select the Edit->Traffic-> Demand Mesh menu.
- When creating demand meshes, select the “Edit traffic after insertion” option 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, select either Gravity or Proportional.
- To change the demand traffic to a fixed amount or on range of random values, follow these steps.

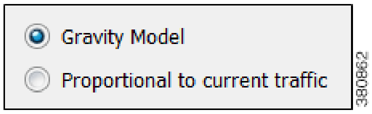
● Gravity Model
○ Proportional to current traffic

1. Select one or more demands.
2. Right-click in a SrcTraffic or DestTraffic field, and click Set Value.
3. In the dialog box that appears, select either Fixed or Random.

   If you choose Random, enter the minimum and maximum traffic values.
4. Click OK.

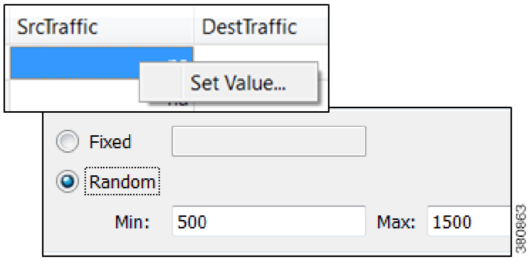

● Random
Min: 500 Max: 1500

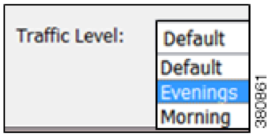

Traffic Level: Default

Step 4 Select the traffic level to which you want to apply the calculations.

Step 5 Click OK.

Modify 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 4-3).
- To set a relative value, use the Modify Demand Traffic dialog box (Figure 4-8 and Table 4-3). 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 1,000 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 1,000 Mbps of traffic and the other had 2,000 Mbps, and if you added 50 Mbps proportionally, one would have 1,016.67 Mbps and the other would have 2,033.33 Mbps.
  - Add a set amount of traffic uniformly to all the demands. For example, if one demand had 1,000 Mbps of traffic and the other had 2,000 Mbps, and if you added 50 Mbps uniformly, one would have 1,025 Mbps and the other would have 2,025 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 1,000 Mbps of traffic and the other had 2,000 Mbps, and if you set them to 4,000 Mbps proportionally, one would have 1,333.33 Mbps and the other would have 2,666.67 Mbps.
  - Set a specified amount of traffic, in Mbps, uniformly to all the demands. For example, if one demand had 1,000 Mbps of traffic and the other had 2,000 Mbps, and if you set them to 4,000 Mbps uniformly, they would both be 2,000 Mbps.

Table 4-3 Modifying Demand Traffic

Modify Fixed Demand Traffic	Modify Fixed or Relative Demand Traffic
1. Select one or more demands in the Demands table. 2. Right-click on one of these demands, and select Properties from the context menu. 3. For each applicable traffic level, enter the desired amount of simulated traffic in the Traffic field. If needed, click the New Traffic Level button to create a new traffic level. 4. Click OK.	1. Select the traffic level you want the changes to apply to from the Traffic Level menu in the Visualization tool bar. 2. If applicable, select one or more demands in the Demands table. Or, if you do not select demands, the result applies to all demands in the plan file for the selected traffic level. 3. Open the Modify Demand Traffic dialog box in one of two ways. – Right-click on a demand and select Modify Demand Traffic from the context menu. – Select the Initializer->Modify Demand Traffic menu. 4. Select one of the options identified in the above list and enter a value accordingly. 5. Click OK.

Figure 4-8 Modify Demand Traffic

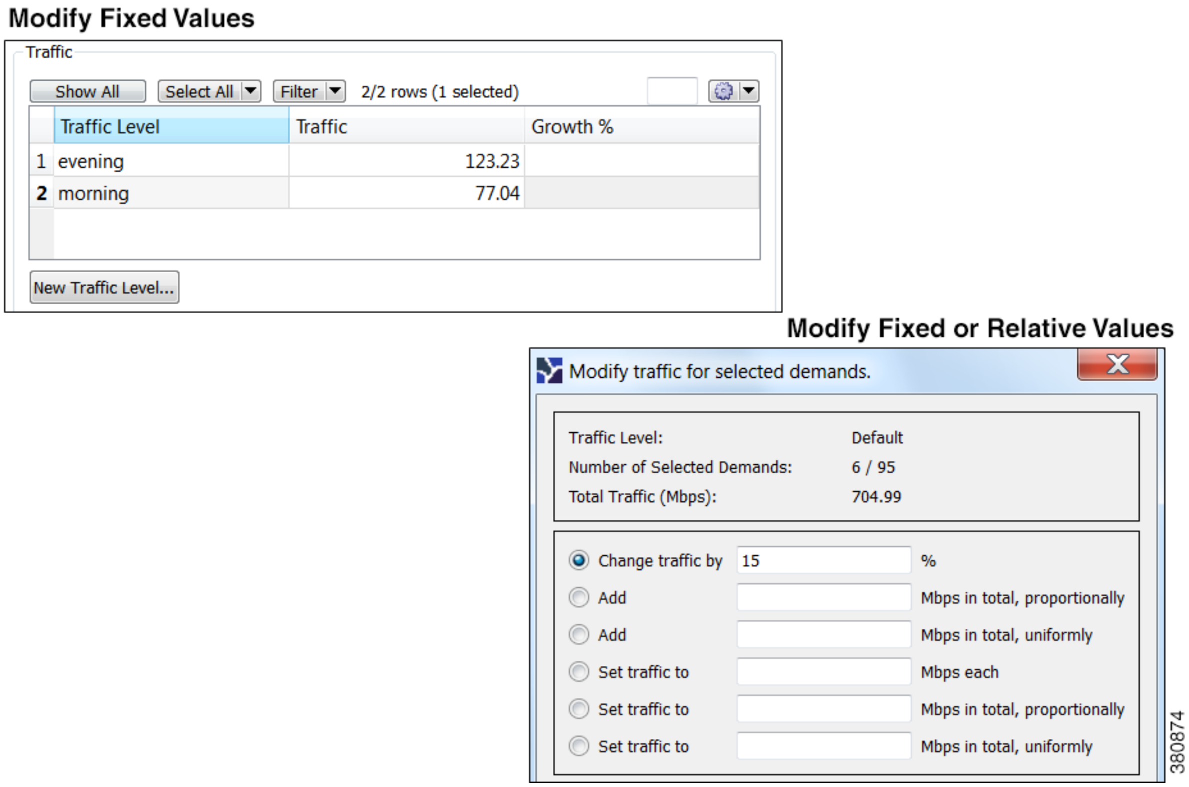

Modify Fixed Values

Traffic

2/2 rows (1 selected)

	Traffic Level	Traffic	Growth %
1	evening	123.23	
2	morning	77.04	

Modify Fixed or Relative Values

Modify traffic for selected demands.

Traffic Level: Default
Number of Selected Demands: 6 / 95
Total Traffic (Mbps): 704.99

● Change traffic by 15 %
○ Add Mbps in total, proportionally
○ Add Mbps in total, uniformly
○ Set traffic to Mbps each
○ Set traffic to Mbps in total, proportionally
○ Set traffic to Mbps in total, uniformly

CHAPTER 5

Simulation

MATE Design network simulations calculate demand routings and traffic distributions through the network based on the given traffic demand, network topology, configuration, and state. Simulation is the fundamental capability of MATE Design on which most of the other tools are built, including those for planning, traffic engineering, and worst-case failure analysis. A number of protocols and models are supported, including IGP, MPLS RSVP-TE, BGP, QoS, VPNs, Multicast, and Layer 1 (L1).

This chapter focuses on the general features of MATE Design simulations. The individual protocols and models are described in their respective chapters.

Simulations can be performed under different simulation convergence modes (Fast Reroute, IGP and LSP Reconvergence, and Autobandwidth Convergence), depending on which stage of the network recovery after failure is being investigated. If in Autobandwidth simulation mode, the results for simulations with failures are identical to those in the IGP and LSP Reconvergence simulation mode. The default simulation mode is IGP and LSP Reconvergence, and except where identified, the documentation describes this simulation mode. For more information, see the MPLS Simulation chapter.

For More Information ...	See ...
Demands and Demand Deduction	Traffic Demand Modeling chapter
Advanced routing with external endpoints	Advanced Routing with External Endpoints chapter
Fast Reroute, Autobandwidth Convergence, and Autobandwidth Convergence (Including Failures) simulation modes	MPLS Simulation chapter
IGP simulation	IGP Simulation chapter
MPLS simulation	MPLS Simulation chapter
Quality of service (QoS) simulation	Quality of Service Simulation chapter
BGP simulation	BGP Simulation chapter
VPN simulation	VPN Simulation chapter
Multicast simulation	Multicast Simulation chapter
L1 simulation	Layer 1 Simulation chapter
Simulate worst-case failures and analyze failure impact	Simulation Analysis chapter
Forecast traffic growth and create growth plans	Forecasting Traffic chapter
IGP metric optimization	Metric Optimization chapter
LSP optimization	LSP Optimization chapter

Simulation

MATE Design uses demands to simulate traffic between source and destination endpoints within a network. Each demand has a specified amount of traffic.

By default, an up-to-date simulation of an open plan is always available. Any change in the plan that affects or invalidates the current simulation automatically triggers a resimulation. Types of changes that result in re-simulations are those that would typically affect routing in a network.

- Changes in topology, such as adding and deleting objects or changing explicit paths.
- Changes to an object's state, such as failing an object or making it inactive.
- Changes in numerous properties, such as metrics, capacities, and delay.

To toggle automatic resimulation on and off, follow one of these two steps.

- Select the drop-down arrow next to the Simulation On/Off icon (white arrow in green circle) in the Visualization toolbar. Click the "Auto resimulate" option to toggle it on and off.
- Select the Tools->Simulator->On/Off menu.

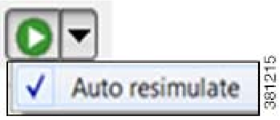

State

The state of an object affects simulation and determines whether an object is operational.

- **Failed State**—Identifies whether the object is failed. See the Failed State section.
- **Active State**—Identifies whether the object is available for use in the simulated network. For example, an object might be unavailable because it has been set administratively down. See the Active State section.
- **Operational State**—Identifies whether an object is operational. For instance, an object might be non-operational because it is failed, is inactive, or because other objects on which it depends are not operational. See the Operational State section.

The Failed and Active columns in the tables show a visual representation of their status. Likewise, you can show the Operational column to show the calculated operational state. In each column, T means the object is in that state and F means it is not. Note that the T or F for Active state in the Interfaces table is reflective of the associated circuit. The plot shows graphical representations of these states with crosses of various colors and fills.

Note For detailed information about visual representations of failed and inactive states in the L1 view, see the Layer 1 Simulation chapter.

Failed State

The quickest way to see the effects of failures in a Simulated Traffic view is to have demands in place and then fail an object. The plot immediately displays where the traffic increases as a result (Figure 5-1), and the Util Sim column in the Interfaces table reflects the traffic changes. Demands are then rerouted around the failure (Figure 5-2).

To specify that a demand should not reroute around failures, deselect the Reroutable option in the demand's Properties dialog box. This can be used as a way of including L2 traffic on an interface. For example, a one-hop, non-reroutable demand can be constructed over the interface to represent the L2 traffic. Other, reroutable, demands, can be constructed through the interface as usual. If the interface fails, the L2 traffic is removed and the L3 traffic reroutes.

You can simultaneously fail one or more objects of the same type. If you select an interface to fail, you are actually failing its associated circuit.

Objects That Can Be Failed	
• Circuits • Nodes • Sites • Ports • Port circuits	• SRLGs • L1 links • L1 nodes • External endpoint members

Figure 5-1 Failed Circuit

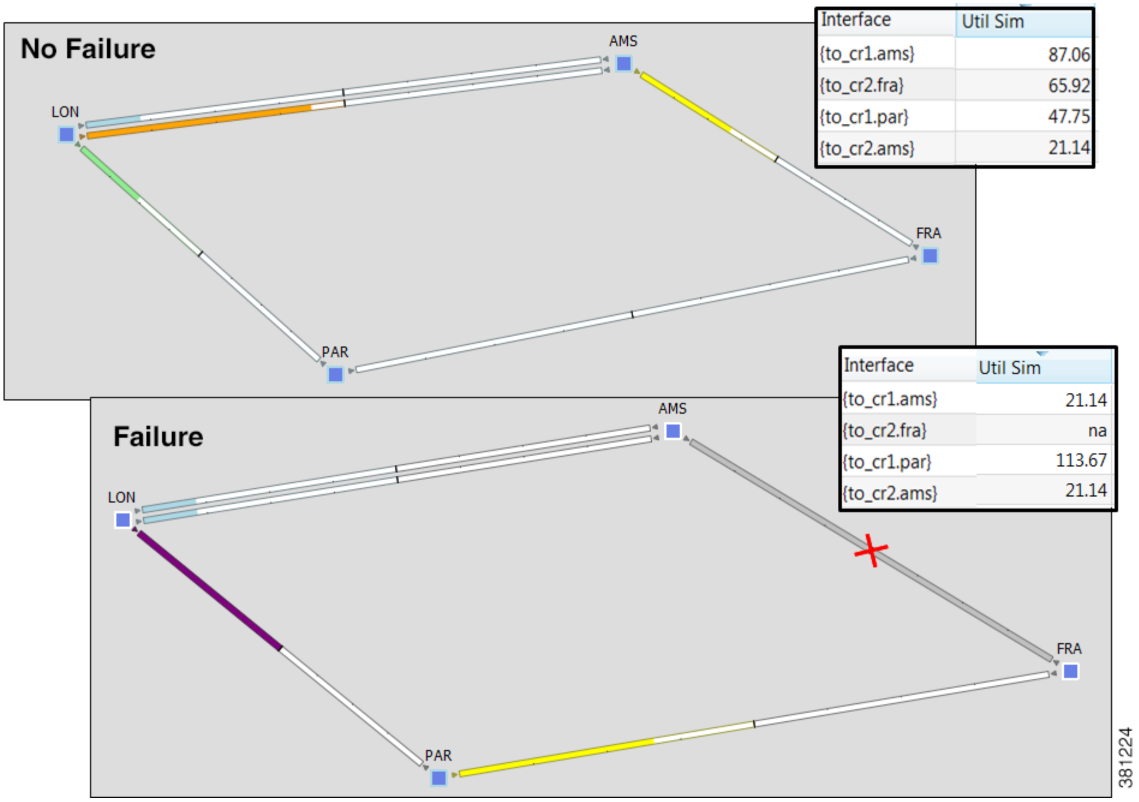

Interface	Util Sim
{to_cr1.ams}	87.06
{to_cr2.fra}	65.92
{to_cr1.par}	47.75
{to_cr2.ams}	21.14

Interface	Util Sim
{to_cr1.ams}	21.14
{to_cr2.fra}	na
{to_cr1.par}	113.67
{to_cr2.ams}	21.14

Figure 5-2 Reroute of Demand Around Failed Circuit

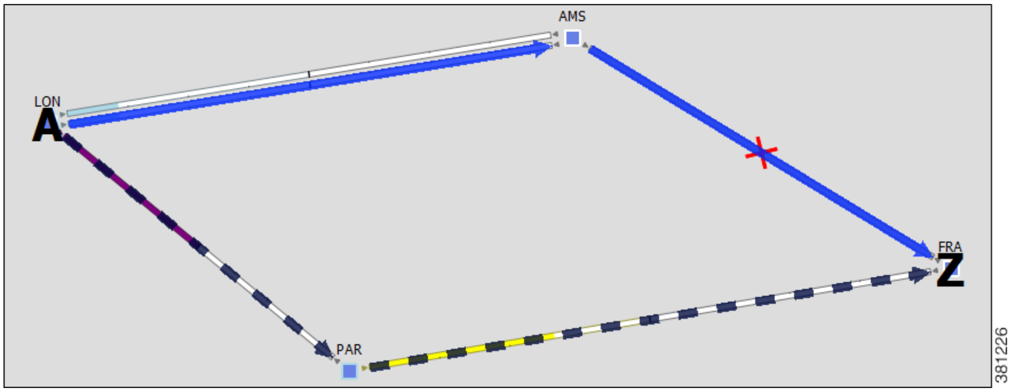

Failures in the Network Plot

In the network plot, a solid red X appears to show failure.

Failure Notations	Description
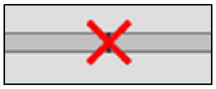	Failed circuit.
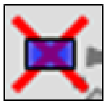	Failed node.
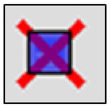	Failed site.
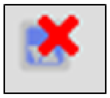	One or more sites, nodes, or circuits within the site failed.

For failed ports or port circuits that do not take down the entire circuit, a circuit with a red strikethrough line in the middle shows the reduced capacity.

Reduced Capacity Notation	Description
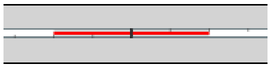	The circuit is a LAG, and one or more ports or port circuits in the LAG failed.The red strikeout line is the percent of unavailable capacity due to the failure.

Fail and Recover Objects

	Fail Objects	Recover Objects
From Context Menu	**1.** Select one or more objects. **2.** Right-click on one, and select Fail from the context menu.	**1.** Right-click in an empty plot area. **2.** Either select All to recover all failed objects or individually select an object to recover it.
From Properties Dialog Box	**1.** Select one or more objects. **2.** Right-click on one, and select Properties from the context menu. **3.** Select the Failed option, and click OK. 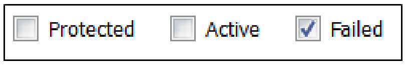 □ Protected □ Active ☑ Failed	**1.** Select one or more objects. **2.** Right-click on one, and select Properties from the context menu. **3.** Deselect the Failed option, and click OK.

Protect Circuits within SRLGs

You can protect circuits from being included in SRLG failures and SRLG worst-case analysis, though there are differences in behavior, as follows.

- These circuits do not fail if you individually fail an SRLG as described in the above steps. However, if you fail the circuit itself, it will fail.
- These circuits are protected from being included in Simulation Analysis regardless of whether they are in an SRLG or not. For more information, see the Simulation Analysis chapter.

Note that this setting has no effect on how FRR SRLGs are routed. For information on FRR SRLGs, see the MPLS Simulation chapter.

Step 1 Set the Protected property for circuits.

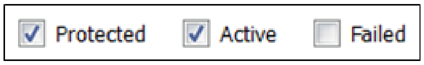

☑ Protected ☑ Active □ Failed

a. Select one or more circuits.

b. Right-click a selected circuit, and select Properties from the context menu.

c. Select the Protected option, and click OK.

Step 2 Set the Network Option property for protecting circuits included in SRLGs.

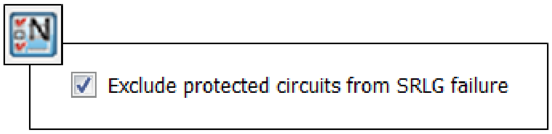

☑ Exclude protected circuits from SRLG failure

a. Select the Edit->Network Options menu or the Network Options icon from the Visualization toolbar.

b. Select the Simulation tab.

c. Select the option to exclude protected circuits from SRLG failures, and click OK.

Active State

An active/inactive state identifies whether an object is available for measured or simulated traffic calculations. An object can be inactive for one of three reasons.

- It is administratively down.
- It is a placeholder. For instance, you might be planning to install an object and want its representation in the network plot.
- It exists in a copied template plan, but was not in the original plan that was discovered.

You can simultaneously change the active state of one or more objects. If you change the active state of an interface, you are actually changing its associated circuit.

Objects That Can Be Set to Inactive	
• Circuits	• L1 links
• Nodes	• L1 nodes
• Sites	• External endpoint members
• Ports	• Demands
• Port circuits	• LSPs
• SRLGs	• LSP paths

Like failures, changing an object from active to inactive immediately affects how demands are routed, and affects the Util Sim column in the Interfaces tables (Figure 5-3).

Figure 5-3 Inactive Circuit

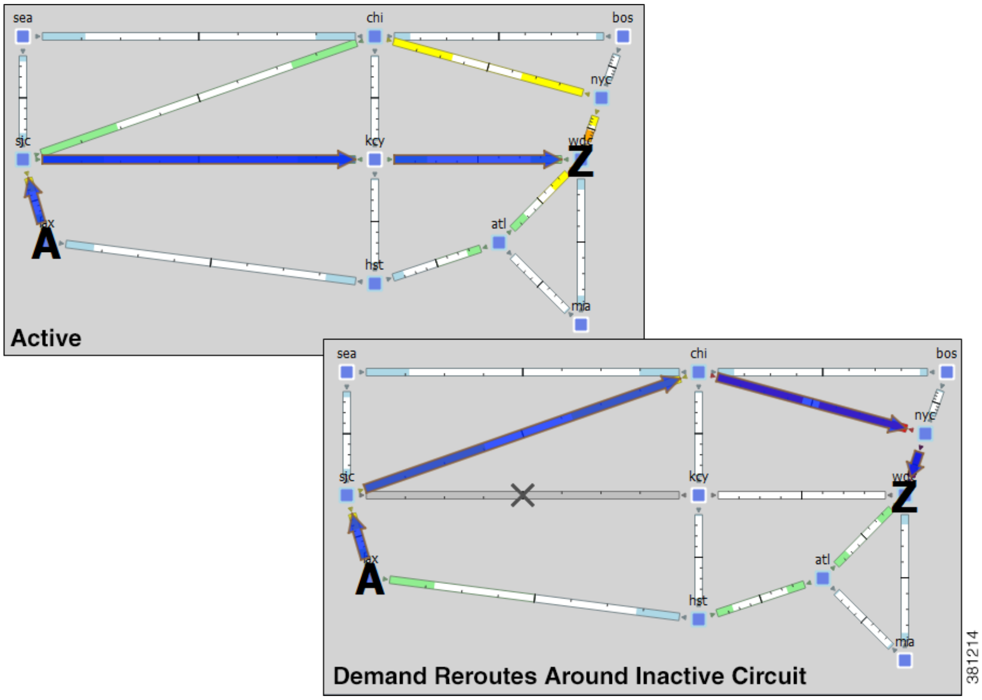

Inactive States in the Network Plot

In the network plot, a solid black X appears to show an inactive state.

If a demand is inactive, the plot shows the ingress and egress sites marked with A and Z, respectively, but the demand (blue) path between them does not appear.

The inactive state takes precedence over a failure. For example, if a circuit is not active and then you fail it, the circuit continues to show a black X on it.

Inactive Notations	Description
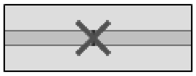	Circuit is inactive.
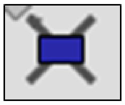	Node is inactive.
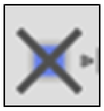	Site is inactive.
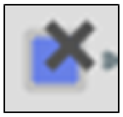	One or more sites, nodes, or circuits within the site are inactive.

If a port or port circuit is inactive, a circuit with a black strikethrough line in the middle shows the reduced capacity.

Reduced Capacity	Reasons
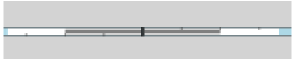	The circuit is a LAG and one or more ports or port circuits in the LAG is inactive. The gray strikeout line is the percent of unavailable capacity due to the inactive ports or port circuits.

Set Active or Inactive States

Step 1 Right-click on one or more like objects from their respective tables. The Properties dialog box appears.

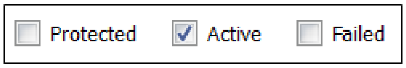

□ Protected ☑ Active □ Failed

Step 2 Click the Active check box to toggle it on or off. A check mark means the object is active.

Step 3 Click OK.

Operational State

The operational state identifies whether the object is functioning. You cannot set an operational state; rather, it is automatically calculated based on the failed and active states.

- Any object that is failed or inactive is operationally down.
- If the object relies on other objects that comprise it to be functioning, then its operational state mirrors the state of those objects.

If This Object Fails or Is Inactive	Then These Objects Will Be Operationally Down
Node	Circuits connected to the failed node
Site	Sites, nodes, and circuits within the failed site
SRLG	Objects within the failed SRLG
Port	Port circuits that contain the failed port

Non-Operational Due to Failures

A red X with a white interior shows that an object is operationally down because of a failure of other objects. For example, if a node fails, a red X appears on the node itself and a red X with a white interior appears on all circuits to which it is attached.

Non-Operational Due to Failures Notations	Description
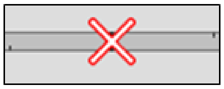	Circuit is not operational for one of the following reasons. • A node connected to the circuit failed. • A site failed that contains a node connected to the circuit. • The circuit is a LAG that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to failure of some of its ports or port circuits. • Layer 1 link or L1 node used by the circuit failed. • An SRLG failed that contains the circuit or contains one of the above objects that are required for circuit operation.
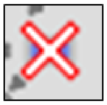	Node is not operational because it is contained in a failed site or SRLG.
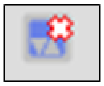	One or more nodes or circuits within the site are not operational because they are contained in a failed SRLG.

Non-Operational Due to Inactive States

A black X with a white interior shows that an object is operationally down because of an inactive state of other objects. For example, if a node is inactive, a black X appears on the node itself and a black X with a white interior appears on both circuits to which it is attached.

Non-Operational Due to Inactive State Notations	Description
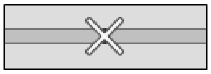	Circuit is not operational for one of the following reasons. • A node connected to the circuit is inactive. • An inactive site contains a node connected to the circuit. • The circuit is a LAG that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to some of its ports or port circuits being inactive. • Mapped Layer 1 circuit cannot be routed. • An SRLG is inactive, and the SRLG contains the circuit or contains one of the above objects that are required for circuit operation.
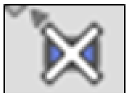	Node is not operational because it is contained in a site or an SRLG that is not active.
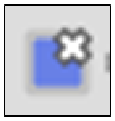	One or more nodes or circuits within the site are not operational because they are contained in an SRLG that is not active.

Simulated Capacity

The Capacity column displays the configured physical capacity of interfaces, circuits, ports, and port circuits. Each circuit, port, and port circuit has a physical capacity that you can set in the Capacity field of the Properties dialog box. The interfaces have a configurable capacity that you can set in the Configured Capacity field. From these properties, a simulated capacity is derived for each object.

The Capacity Sim column is the calculated capacity of the object given the state of the network, which could include failures that reduce capacity. All utilization figures in the Interfaces, Circuits, and Interface Queues tables are calculated based on this Capacity Sim value. When referencing the Capacity Sim value, there are a few rules to note regarding its calculation.

- If a circuit’s capacity is specified, this becomes the simulated capacity of the circuit, and all other capacities (interface and constituent port capacities) are ignored. Specifying circuit capacities (rather than interface capacities) is the simplest way to modify existing capacities, which is useful, for example, for build-out planning.
- A circuit Capacity Sim is the same as the Capacity Sim for each of its constituent interfaces.
- If the circuit capacity is not specified, then the associated interface Capacity Sim values are the minimum of the configured interface capacity and that of its connected remote interface. Thus, the circuit capacity is effectively negotiated down to the minimum of the two interface capacities.
- If two ports are connected by a port circuit, then the Capacity Sim of the two ports and of the port circuit is set to the minimum capacity of the three, which again, effectively negotiates down the capacity of each side of the connection.

- In a LAG interface, if any of the constituent ports are operationally down, then the interface Capacity Sim column displays a value that is reduced by the aggregate capacity of all the ports that are down. For example, if a 1,000 Mbps port of a four-port 4,000-Mbps LAG is operationally down, the simulated capacity for that LAG interface becomes 3,000 Mbps.

Simulated Delay and Distance

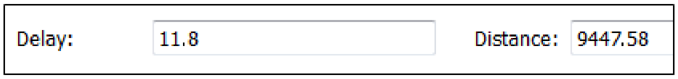

Delay: 11.8 Distance: 9447.58

Both Delay and Distance are properties that can be set in the circuit Properties dialog box. To automate the process of setting these, use the Latency and Distance Initializer. For information, see the Latency and Distance Initializer section.

All MATE Design delay calculations that use the L3 circuit delay, such as Metric Optimization, use the Delay Sim value.

Circuits Table Column	Description
Delay	One-way transmission latency over the circuit in milliseconds (ms).
Distance	Distance of the circuit in kilometers (km).
Delay Sim	• (Derived) If the circuit Delay value is entered, it is copied to the Delay Sim column. • If no Delay is explicitly set, the Delay Sim column in the Circuits table is copied from the L1 Delay Sim column of the mapped L1 circuit (if one exists), which is based on the L1 circuit path that is currently the active path. • If the Delay is not specified (na) and if there are no L1 circuits mapped to the circuit, the Delay Sim column is 0.
Distance Sim	• (Derived) If the circuit Distance value is entered, it is copied to the Distance Sim column. • If no Distance is explicitly set for the circuit, the Distance Sim column in the Circuits table is copied from the L1 Distance Sim column of the mapped L1 circuit (if one exists), which is based on the L1 circuit path that is currently the active path. • If the Distance is not specified (na) and if there are no L1 circuits mapped to the circuit, the Distance Sim column is 0.

For more information on the delay and distance of L1 objects, see the Layer 1 Simulation chapter.

Latency and Distance Initializer

Following is the default behavior for the Latency and Distance Initializer. Note that if the L1 nodes, nodes, or sites used by the initializer do not have Longitude and Latitude properties, both the Delay Sim and Distance Sim properties are set to zero.

- **L1 Defaults**

- The current L1 circuits Delay and Distance properties are removed unless the "Keep existing distances if available" option is selected.
- The L1 node's Longitude and Latitude properties are used to estimate the L1 distances. If an L1 node does not have these properties set, the shortest distance calculation is based these same properties of its immediate parent site, if applicable. (L1 nodes are not required to have parent sites.) The L1 link delays are calculated from these distances.
- Each L1 circuit path's Distance Sim and Delay Sim columns are calculated as the sum of the distances and delays of L1 links over which the L1 circuit path is routed.

- **L3 Defaults**
  - L3 circuit delays and distances are inherited from L1 circuits if there are L1 circuits mapped to the L3 circuits. Otherwise, the L3 delay and distance are estimated using the shortest distance between nodes at the endpoints of the circuits. If a node does not have Longitude and Latitude properties, the calculation is based on these same properties of its immediate parent site, if applicable. (Nodes are not required to have parent sites.) Thus, if you have an L1 model, the delays and distances are more precise because the physical paths that the L3 circuits take are known.
  - ITU G.826 adjustments are used to convert L3 circuit distance (d) to an adjusted circuit distance (r) as follows.
    - If d < 1000, set r = 1.5 * d
    - If 1000 <= d <= 1200, set r = 1500
    - If d > 1200, set r = 1.25 * d

    If this option is unchecked, set r equal to d.
- **Both the L3 and L1** delays are calculated to be the speed of light along a distance r, multiplied by the speed of light in medium correction factor (between 0 and 1). The default for this correction factor is 0.67.

Step 1 This initializer can act on two different selected object types (L1 links and L3 circuits), so the defaults for the dialog box differ, depending both on selections and on how you access it. Regardless of the method, you can modify your selection options once the dialog box is opened. Following are guidelines for having the dialog box open with the most desirable defaults.

To Calculate ...	Select These Objects	Select This Menu
L1 link distance and delay, and L3 circuit delay	• No selections or • One or more L1 links and one or more L3 circuits	Initializers->Latency and Distance menu.

To Calculate ...	Select These Objects	Select This Menu
L1 link distance and delay	• No selections or • One or more L1 links	Right-click on an L1 link, and select Initialize Latency from the context menu.
L3 circuit delay	• No selections or • One or more L3 circuits	Use one of these options. • Right-click on an L3 circuit, and select Initialize Latency from the context menu. • Right-click on a circuit or interface, and select Properties. Click the Estimate button next to the Delay field. This option is specific to L3 circuits.

Step 2 In each section, Layer 1 and Layer 3, verify or change your selection to act on the desired objects.

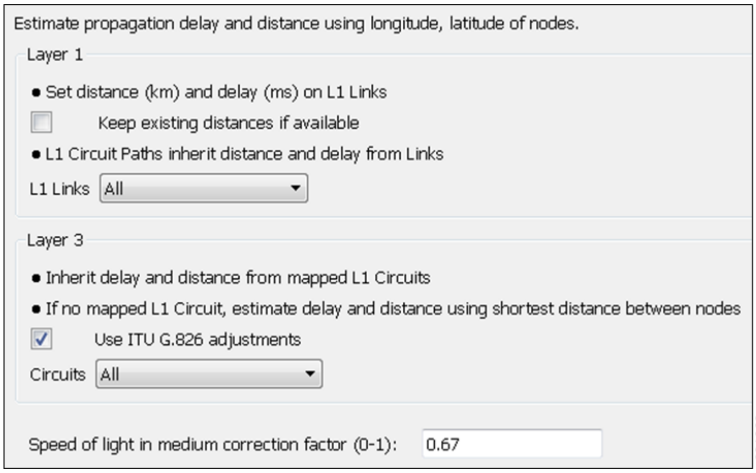

Estimate propagation delay and distance using longitude, latitude of nodes.

Layer 1

● Set distance (km) and delay (ms) on L1 Links

□ Keep existing distances if available

● L1 Circuit Paths inherit distance and delay from Links

L1 Links: All

Layer 3

● Inherit delay and distance from mapped L1 Circuits

● If no mapped L1 Circuit, estimate delay and distance using shortest distance between nodes

☑ Use ITU G.826 adjustments

Circuits: All

Speed of light in medium correction factor (0-1): 0.67

- For L1 links, the default is to delete distances. To keep them, select the option to keep existing distances.
- For L1 circuit paths, the default is to inherit distance and delay from L1 links. If desired, use the L1 Links drop-down list to base these calculations on all, selected, or tagged L1 links. If you do not want the distance and delay set on L1 circuit paths, select None.
- For L3 circuits, if desired, use the Circuits drop-down list to calculate L3 circuit delay and distance for all, selected, or tagged L3 circuits. If you do not want the delay or distance set on L3 circuits, select None.

  Verify the inclusion ITU G.826 adjustments in the calculation, which is the default and which is usually left unchanged.

Step 3 For both L1 links and L3 circuits, verify or change default speed of light factor of 0.67. This is usually left unchanged.

Step 4 Click OK.

Simulation Analysis

The Simulation Analysis tool combines the simulation results of a large set of failure scenarios. These results are useful for determining how vulnerable a network is to congestion and high latencies under failures, thus allowing you to plan sufficient capacity for any given failure scenario.

Simulation Analysis is run across a set of failure scenarios that include selected objects, such as circuits, and traffic levels. MATE Design calculates these failure scenarios across all service classes. Each scenario is simulated and results in the following available analyses, which vary depending on whether the network has QoS parameters and depending on which options are selected when running the simulation.

- Worst-Case Traffic Utilization on interfaces per service class
- Optional: VPN worst-case utilization and latency
- Optional: Worst-Case Demand Latency
- Failure Impact, which analyzes the impact that each failed object has on interface utilizations throughout the network

Upon completion, a report window opens with a summary of each analysis, along with the list of simulations performed. Each time you run a simulation, this information is updated (replaced).

After performing a Simulation Analysis across multiple failure sets, you can fine-tune the analysis for a subset of failure scenarios and a subset of service classes without running a new analysis. For example, if you run a simulation for nodes, circuits, and ports, you can later go back and view the results for any one of those three objects. See View Simulation Analysis on a Subset of Failure Scenarios or Service Classes.

Simulation Analysis can be performed under different simulation convergence modes (Fast Reroute, IGP and LSP Reconvergence, Autobandwidth Convergence, and Autobandwidth Convergence Including Failures), depending on which stage of the network recovery after failure is being investigated. The default simulation mode is IGP and LSP Reconvergence, and except where identified, the documentation describes this simulation mode. For more information, see the MPLS Simulation chapter.

For More Information ...	See ...
Visualization of Layer 1 (L1) failures in both the L1 and Layer 3 (L3) views	Layer 1 Simulation chapter
VPN worst-case (WC) utilization and latency columns in the VPNs table	VPN Simulation chapter
Fast Reroute, Autobandwidth Convergence, and Autobandwidth Convergence (Including Failures) simulation modes	MPLS Simulation chapter

Worst-Case Traffic Utilization

The default analysis is to identify up to 10 failures for the worst-case utilization on each interface in the network.

Worst case is the highest utilization that a particular interface experiences over all the failure sets and traffic levels that you selected. MATE Design also determines which combination of failures would cause this worst-case utilization.

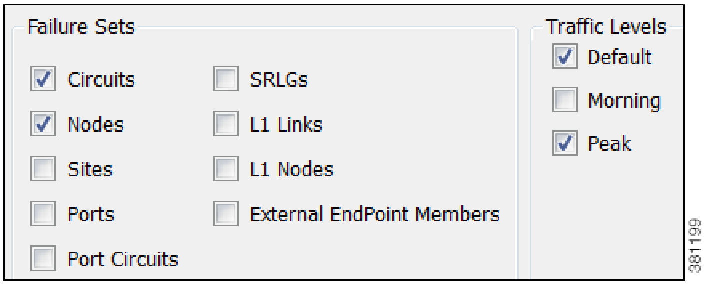
Failure Sets
☑ Circuits
☑ Nodes
☐ Sites
☐ Ports
☐ Port Circuits
☐ SRLGs
☐ L1 Links
☐ L1 Nodes
☐ External EndPoint Members

Traffic Levels
☑ Default
☐ Morning
☑ Peak

Alternatively, you can record failures causing utilizations within a specified percent of the worst-case utilization.

Upon finishing the analysis, MATE Design switches to the Worst-Case Traffic view and updates the plot to simultaneously display all worst-case failures (Figure 6-1). It also updates the following columns in the Interfaces and Circuits tables.

- WC Util—The worst-case utilization for that interface. The worst-case for a circuit is defined to be whichever of the worst cases of the two constituent interfaces results in the larger utilization. Thus, for circuits, this value is the larger of the WC Util values for the two interfaces in the circuit.
- WC Traffic—The actual traffic (Mbps) through the interface under the worst-case scenario.
- WC Traff Level—The traffic level under which this worst-case scenario occurs.
- WC Failure—List of one or more failures that cause the worst-case failure of the circuit. An easier way to read this list is to right-click on an interface and select Fail to WC from the context menu.

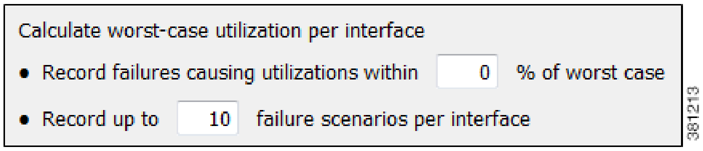
Calculate worst-case utilization per interface
- Record failures causing utilizations within 0 % of worst case
- Record up to 10 failure scenarios per interface

Example: The circuit between cr1.lon to cr1.par, the cr1.lon node, and the L1 link between lon and par would all cause the worst-case utilization failure. ct{cr1.lon|to_cr1.par}|cr1.par|{to_cr1.lon}}; nd{cr1.lon};L1lnk{lon|par|lon-par}

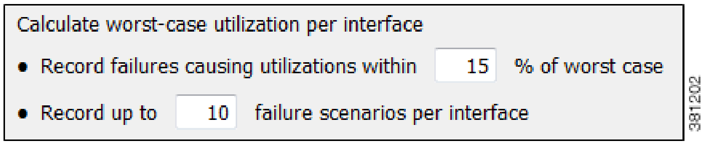
Calculate worst-case utilization per interface
- Record failures causing utilizations within 15 % of worst case
- Record up to 10 failure scenarios per interface

If you record failures causing utilizations within a given percent of worst case, this column shows QoS violations as a percent. (See Worst-Case QoS Violations.) If the number is positive, then the allotted capacity has been surpassed. If negative, the capacity has not been surpassed. For example, if a circuit has 10,000 Mbps capacity, and if the amount of traffic on it as a result of three different failures is 11,000, 8,000, and 4,000 Mbps, the utilizations are 10%, -20%, and -60%, respectively and in descending order.

Example: The circuit between cr1.ams and cr2.lon is the worst-case failure and would cause this interface to exceed its capacity by 24.75%. The failure of the cr2.lon node would cause the interface traffic to go to its second highest utilization, which is 2.8% less than its capacity.

ct{cr1.ams|{to_cr2.lon}|cr2.lon|{to_cr1.ams}} (-24.75%);nd{cr2.lon} (2.8%)

For information on reading notation of objects within tables, see the *MATE Integration and Development Guide*.

- WC Service Class—The service class for which this worst-case scenario occurs. For information on running a Simulation Analysis with QoS, see Worst-Case QoS Violations.

Note For information on worst-case calculations for VPNs, see the VPN Simulation chapter.

Figure 6-1 Worst-Case Traffic Utilization for All Interfaces

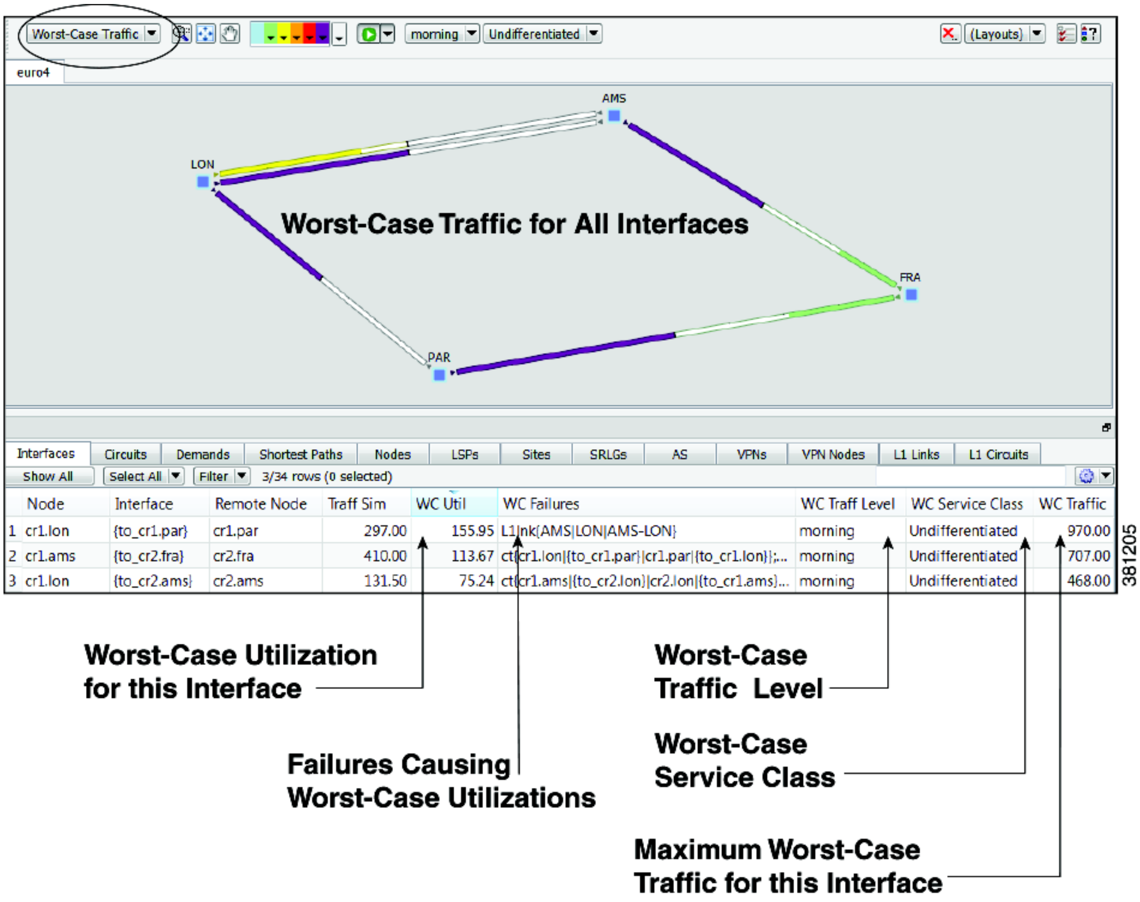

	Node	Interface	Remote Node	Traff Sim	WC Util	WC Failures	WC Traff Level	WC Service Class	WC Traffic
1	cr1.lon	{to_cr1.par}	cr1.par	297.00	155.95	L1Link{AMS|LON|AMS-LON}	morning	Undifferentiated	970.00
2	cr1.ams	{to_cr2.fra}	cr2.fra	410.00	113.67	ct{cr1.lon|{to_cr1.par}|cr1.par|{to_cr1.lon}};...	morning	Undifferentiated	707.00
3	cr1.lon	{to_cr2.ams}	cr2.ams	131.50	75.24	ct{cr1.ams|{to_cr2.lon}|cr2.lon|{to_cr1.ams}...	morning	Undifferentiated	468.00

Worst-Case QoS Violations

MATE Design includes QoS bound (maximum available capacity) as part of the worst-case calculations. If there are no QoS parameters set, then the QoS bound is 100% and violations occur if utilization goes over that 100%. However, if a worst-case policy has been set on a service class or if interface queue parameters have been set, then worst-case QoS violations are calculated. In these instances, MATE Design identifies the interface with the highest percentage of QoS violation as the worst-case possibility (Figure 6-2). The following columns are updated accordingly.

- WC QoS Bound—The worst-case interface capacity available without violating these QoS requirements. This value is based on available capacity, traffic utilization, worst-case policies set on service classes, and interface queue parameters. The WC QoS Bound (%) column identifies this same value as a percentage of the total capacity.
- WC QoS Violation—The worst-case traffic minus the worst-case capacity permitted (worst-case QoS bound). A violation occurs if the QoS capacity allotted through worst-case policies for service classes is exceeded or if QoS capacity allotted through interface queue parameters was exceeded. If the number appearing in the WC QoS Violation column is positive, then the allotted capacity has been surpassed. If negative, the capacity has not been surpassed. The WC QoS Violation (%) column identifies this same value as a percentage of total capacity.

  To see the cause of worst-case QoS violations, right-click on a circuit and select Fail to WC. The table that appears lists all causes of this interface's worst-case utilization and its worst-case QoS violations. Select the worst-case failure you want to view, and click OK.
- WC Service Class—The service contributing to the worse-case QoS violation.

For More Information ...	See ...
QoS parameters and QoS calculations	Quality of Service Simulation chapter
Set worst-case policies on service classes	Quality of Service Simulation chapter
Set interface queue parameters	Quality of Service Simulation chapter
Worst-case QoS calculations for VPNs	VPN Simulation chapter

Figure 6-2 Worst-Case QoS Violations

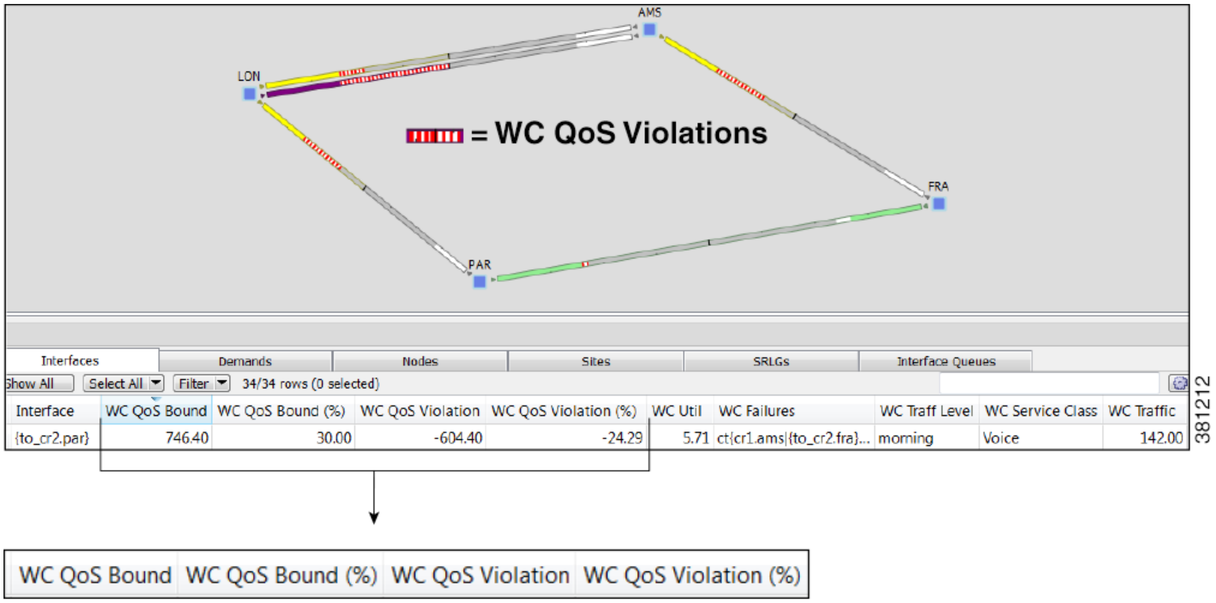

Interface	WC QoS Bound	WC QoS Bound (%)	WC QoS Violation	WC QoS Violation (%)	WC Util	WC Failures	WC Traff Level	WC Service Class	WC Traffic
{to_cr2.par}	746.40	30.00	-604.40	-24.29	5.71	ct{cr1.ams|{to_cr2.fra}...	morning	Voice	142.00

Fail Circuit to Worst-Case Utilization

You have the option to selectively view each failure scenario that causes the worst-case utilization or worst-case QoS violation for a single interface.

- Right-click on an interface or a circuit, and select Fail to WC. If there is only one worst-case failure (listed in the WC Failure column), that object causing the worst-case utilization is listed. You can select it to fail it immediately.

  If there are multiple possibilities for a worst-case failure, or if there is a range of failures within a percentage of the worst-case failure, a Worst-Case Failures dialog box appears listing each failure, its worst-case utilization percent, and its QoS violation percent (Figure 6-3). Select the worst-case failure you want to view, and click OK. The network plot changes to show this particular failure scenario. If you select to fail an L1 link or L1 node, switch to the L1 view to see the failure.

Note If you select an interface, you are actually failing its associated circuit to its worst-case.

- Alternatively, to filter to these worst-case failures for an interface without invoking a failure, right-click on an interface or a circuit, select Filter to -> Filter to WC, and then select the worst-case failure scenario of interest. If you fail this object from here, it achieves the same as if you had selected it from the Worst-Cae Failures dialog box.

Figure 6-3 Failure of Single Circuit to Its Worst-Case

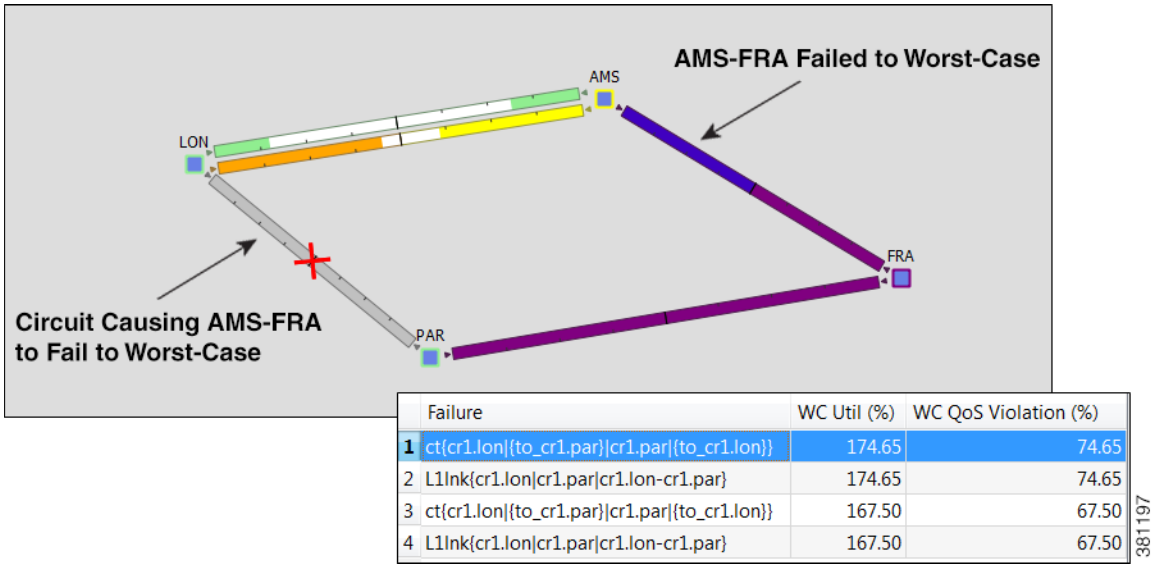
AMS-FRA Failed to Worst-Case
AMS
LON
FRA
PAR
Circuit Causing AMS-FRA to Fail to Worst-Case
381197

	Failure	WC Util (%)	WC QoS Violation (%)
1	ct{cr1.lon|{to_cr1.par}|cr1.par|{to_cr1.lon}}	174.65	74.65
2	L1lnk{cr1.lon|cr1.par|cr1.lon-cr1.par}	174.65	74.65
3	ct{cr1.lon|{to_cr1.par}|cr1.par|{to_cr1.lon}}	167.50	67.50
4	L1lnk{cr1.lon|cr1.par|cr1.lon-cr1.par}	167.50	67.50

Worst-Case Demand Latency

When running a Simulation Analysis, you have the option to simulate worst-case latency for each demand in the plan. MATE Design calculates the maximum latency of each demand under the failure scenarios selected. The result does not depend on service classes or traffic levels since demand routing is independent of these plans. The simulation also records the failures that cause this maximum latency.

The following columns in the Demands table are updated when the "Calculate demand worst-case latency" option is selected for Simulation Analysis.tool

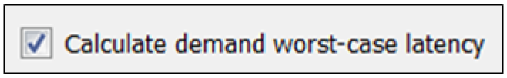
☑ Calculate demand worst-case latency

- WC Latency—The highest demand latency over all failure scenarios in the analysis.

- WC Latency Failures—The failures that caused this worst-case latency. Up to 10 failures are identified.

Note For information on worst-case latency calculations for VPNs, see the VPN Simulation chapter.

Fail Demand to Worst-Case Latency

After running Simulation Analysis to calculate demand worst-case latency, you have the option to fail a single demand to its worst-case latency. Simply right-click on the demand and select Fail to WC Latency (Figure 6-4).

Figure 6-4 Example of Worst-Case Demand Latency

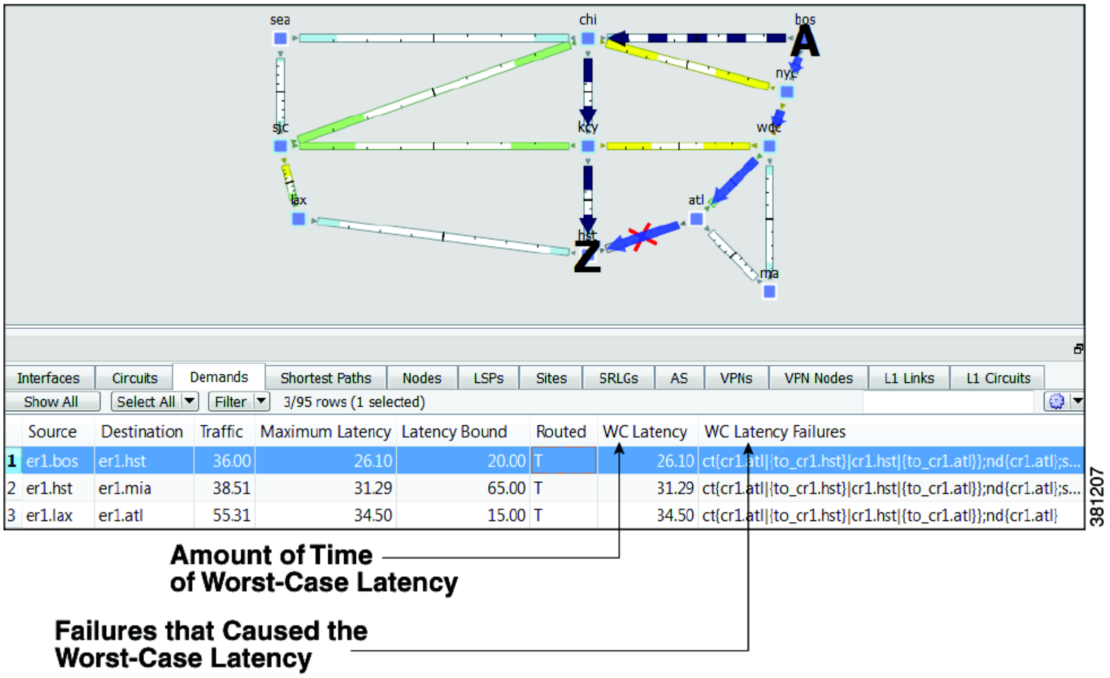

	Source	Destination	Traffic	Maximum Latency	Latency Bound	Routed	WC Latency	WC Latency Failures
1	er1.bos	er1.hst	36.00	26.10	20.00	T	26.10	ct{cr1.atl|{to_cr1.hst}|cr1.hst|{to_cr1.atl}};nd{cr1.atl};s...
2	er1.hst	er1.mia	38.51	31.29	65.00	T	31.29	ct{cr1.atl|{to_cr1.hst}|cr1.hst|{to_cr1.atl}};nd{cr1.atl};s...
3	er1.lax	er1.atl	55.31	34.50	15.00	T	34.50	ct{cr1.atl|{to_cr1.hst}|cr1.hst|{to_cr1.atl}};nd{cr1.atl}

Amount of Time of Worst-Case Latency

Failures that Caused the Worst-Case Latency

Failure Impact

The Failure Impact view is available upon running a Simulation Analysis (Figure 6-5). The plot in this view colors the nodes and circuits according to the maximum utilization level that would be caused elsewhere in the network should the node or circuit fail. The color indicates the resulting utilization and severity of the congestion.

Example: In the Failure Impact view, a PAR-LON has a utilization of 90-100% and its color representation is red. This means that if If PAR-LON were to fail, one or more interfaces would react by exceeding a 90% utilization level and correspondingly, would turn red in the plot.

The Node, Interface, and Circuit tables contain the Failure Impact and Failure Impact Interface columns. In the Interfaces table, the information describes the failure impact of the circuit containing the interface.

- Failure Impact—The failure impact of each node or circuit. For instance, if the value is 80%, it means that if this node or circuit failed, the resulting traffic utilization on one or more interfaces would exceed 80%.
- Failure Impact Interface—The interface that will experience the highest utilization as a result of the node or circuit going down.

  Format = if{Node|Interface}

  Example: if{cr2.lon|{to_cr1.ams}} means if the circuit goes down, it will have the greatest traffic impact on the cr1.lon to cr1.ams interface.

Note The Failure Impact view only shows the impact of circuit and node failures. It does not show failures of other objects.

Figure 6-5 Example Failure Impact

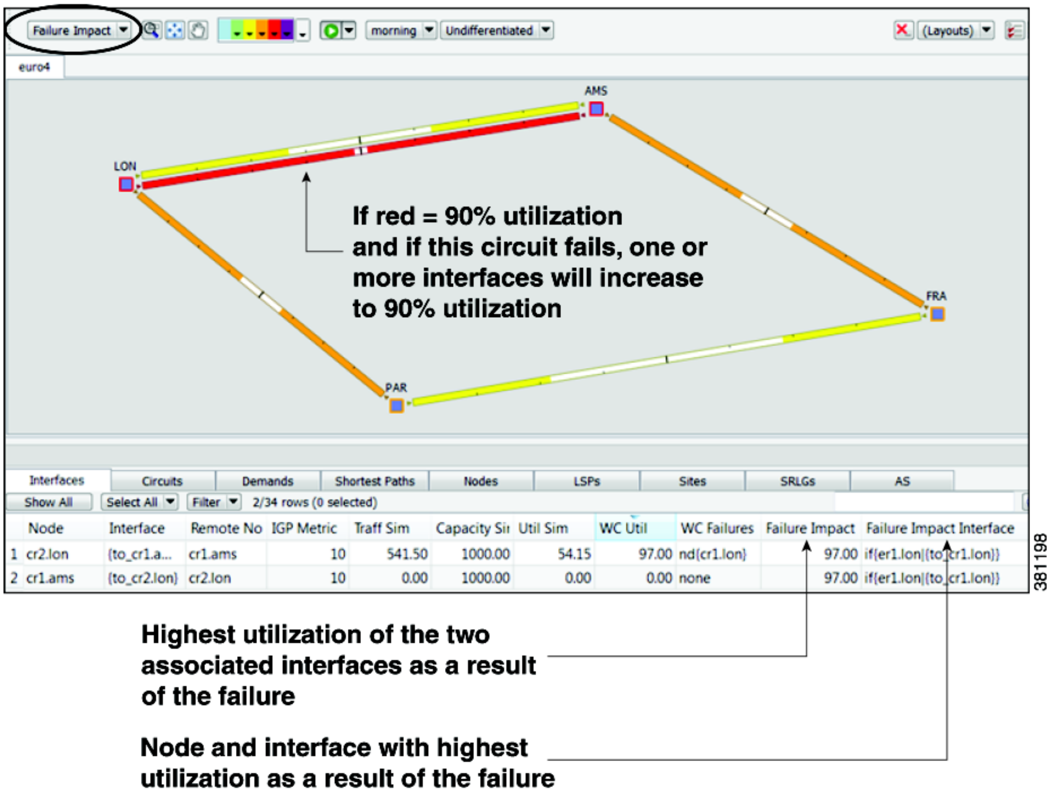

If red = 90% utilization and if this circuit fails, one or more interfaces will increase to 90% utilization

	Node	Interface	Remote No	IGP Metric	Traff Sim	Capacity Si	Util Sim	WC Util	WC Failures	Failure Impact	Failure Impact Interface
1	cr2.lon	{to_cr1.a...	cr1.ams	10	541.50	1000.00	54.15	97.00	nd{cr1.lon}	97.00	if{er1.lon|{to_cr1.lon}}
2	cr1.ams	{to_cr2.lon}	cr2.lon	10	0.00	1000.00	0.00	0.00	none	97.00	if{er1.lon|{to_cr1.lon}}

Highest utilization of the two associated interfaces as a result of the failure

Node and interface with highest utilization as a result of the failure

Simulation Analysis Reports

Each time Simulation Analysis is run, a report is automatically generated. You can access this information again by selecting the Window->Reports menu. Note that new reports replace previous ones.

The Summary information details the options used in the analysis and summarizes the most important problems identified, such as QoS violations and latency bound violations.

The Simulations table lists each simulation that was performed in the Simulation Analysis. (Figure 6-5).

Table 6-1 Simulations Table in Simulation Analysis Report

Simulation Data Point	Description
Failure	Failure scenario used in the analysis.
Service Class	Service class used in the analysis.
Traffic Level	Traffic level used in the analysis.
Network Breakpoint	Identifies whether there are network breaks resulting from the failures in this simulation. If multiple network breakpoints occur, the most serious one is listed. • Yes (Total) = A break exists that completely partitions the network into two or more disconnected sections. • Yes (AS) = A break exists that completely partitions an AS into two or more sections. However, routes exist between the sections of the AS through other AS's. • Yes (OSPF Area 0) = A break exists that completely partitions Area 0 of an AS running OSPF. Under OSPF, traffic cannot route between the partitions even if a path is available through non-zero areas in the AS. • No = No break in the network.
Num Unrouted Demands	Number of demands that cannot be routed under this failure for any of the reasons identified by the network breakpoint.
Unrouted Traffic	Total amount of demand traffic that cannot be routed under this failure for any of the reasons identified by the network breakpoint.
Max Util	Maximum utilization over all interfaces in this simulation. Utilization is the traffic through the interface as a percentage of the capacity of the interface.
Max QoS Bound Percent	Worst-case capacity available without violating QoS bounds, expressed as a percentage of the total capacity.
Num QoS Violations	Number of times the QoS bound is violated. QoS bounds are set through service class policies and interface queue parameters.
Latency Bound Violations	Number of demands with maximum latency in excess of the latency bound specified for the demand.
Num Unrouted LSPs	Number of unrouted LSPs in the analysis.

Protect Objects

To exclude an object from the list of those objects failed when performing a Simulation Analysis, you can mark it as *Protected* in its Properties dialog box. For example, if you want to run a Simulation Analysis only on core nodes, you could first protect all edge nodes. Note that you cannot protect an SRLG itself, though you can protect the objects within it.

You can protect nodes, sites, circuits ports, port circuits, external endpoint members, L1 links, and L1 nodes. You cannot protect SRLGs.

For information on protecting circuits from being failed as part of simple SRLG failures and as part of worst-case SRLG failure analysis, see the Simulation Analysis chapter.

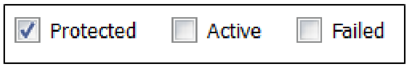

☑ Protected □ Active □ Failed

Note If you select an interface, you are actually protecting its associated circuit.

Step 1 Right-click on one or more like objects from their respective tables. The Properties dialog box appears.

Step 2 Click the Protected check box to toggle it on or off. A check mark means the object is protected.

Step 3 Click OK.

Run Simulation Analysis

The Simulation Analysis tool is the basis of four failure analysis options: worst-case utilization on interfaces, worst-case VPN utilization and latency, worst-case demand latency, and failure impact.

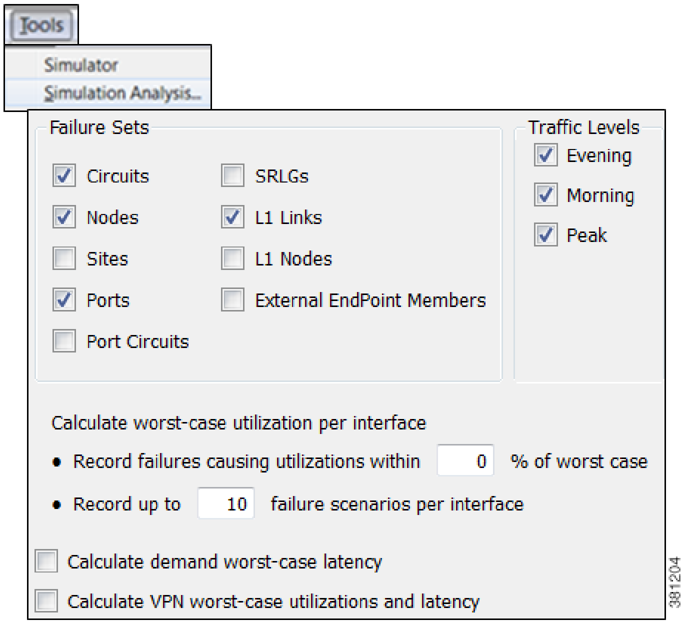

Failure Sets

☑ Circuits □ SRLGs
☑ Nodes ☑ L1 Links
□ Sites □ L1 Nodes
☑ Ports □ External EndPoint Members
□ Port Circuits

Traffic Levels

☑ Evening
☑ Morning
☑ Peak

Calculate worst-case utilization per interface

- Record failures causing utilizations within 0 % of worst case
- Record up to 10 failure scenarios per interface

□ Calculate demand worst-case latency

□ Calculate VPN worst-case utilizations and latency

Note Recording worst-case latencies or VPN worst-case utilizations increases the time it takes to perform a worst-case analysis.

Step 1 Select the Tools->Simulation Analysis menu.

Step 2 Select one or more failure sets.

Step 3 Select one or more traffic levels.

Step 4 In the “Record failures causing utilizations within __% of worst case” field, either enter 0 to record only worst-case failures, or enter a number to find all failures causing utilizations within that percentage range of the worst-case failure.

Step 5 Define the maximum number of failure scenarios to record per interface. The default is 10.

Example: If you record failures causing utilizations within 10% of the worst case, and if the worst-case utilization for an interface is 90%, then MATE Design records failures on this interface resulting in utilization of 81% or higher (90 -(90/10)). In this same scenario, if you record 10 failure scenarios per interface, and if there are failures that could cause utilizations of 90%, 85%, 82%, and 76% for an interface, MATE Design does not record the failure causing 76% utilization.

Step 6 Select whether or not to record demand worst-case latency calculations.

Step 7 Select whether or not to record VPN worst-case utilizations and latencies. For more information, see the VPN Simulation chapter.

Step 8 Click OK.

You can now use the Worst-Case Traffic view to analyze the worst-case traffic utilization, worst-case QoS violation, and worst-case latency information. Here you can also fail interfaces and nodes to their worst case, and fail demands to their worst-case latency. You can also use the Failure Impact view to identify the circuits that are responsible for worst-case traffic congestion.

View Simulation Analysis on a Subset of Failure Scenarios or Service Classes

Once you have executed a Simulation Analysis, you can select a subset of failure scenarios or service classes, and view the results for just these subsets. This practice saves time in that you do not have to rerun the analysis.

After running a Simulation Analysis, follow these steps.

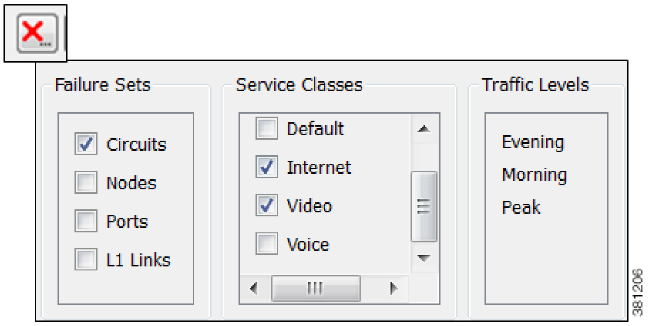

Failure Sets
☑ Circuits
□ Nodes
□ Ports
□ L1 Links

Service Classes
□ Default
☑ Internet
☑ Video
□ Voice

Traffic Levels
Evening
Morning
Peak

Step 1 Click on the red cross in the top, right toolbar. A list of failure scenarios, service classes, and traffic levels appears. These are the options on which you last ran a Simulation Analysis.

Step 2 Click on one or more failure sets and service classes.

Step 3 Click OK.

Note If you need to change the traffic level selection, you must rerun a Simulation Analysis.

CHAPTER 7

Forecasting Traffic

MATE Design traffic forecasting tools enable you to create plans containing estimates of future traffic based on projected growth of current measured or simulated traffic. From these new plans you can determine the impact of the new traffic on the network and plan for upgrades. The traffic growth estimates used in this process can be manually entered, for example, based on knowledge of new services being introduced. The traffic can also be imported from a MATE Live traffic (trending) report or be generated in an external table. After creating these growth plans, you can analyze their capacity requirements using Simulation Analysis, and use the GUI to upgrade the network to meet these requirements.

One example use of these growth plan tools would be to predict the effects of a certain percentage of increased traffic for each quarter of the coming year. You could generate a separate plan for each of these traffic predictions to forecast the impact of each quarter's anticipated growth. Another example application would be to group the demands into demand groupings to predict the effects of different traffic growth at each site, for instance 10% for Paris, 20% for London, and 30% for Amsterdam.

For More Information...	See...
Demands	Traffic Demand Modeling
Import demand groupings	*MATE Integration and Development Guide*
Import traffic and growth rates, including demand grouping traffic and MATE Live traffic reports for use in forecasting	*MATE Integration and Development Guide*
Create MATE Live traffic reports for use in MATE Design forecasting	*MATE Live User Guide* *MATE Integration and Development Guide*

Demand Groupings

Demand Groupings define a group of demands. They provide a convenient way of specifying aggregated traffic in a plan, which can be used as a basis for traffic reports and growth plans. For example, you can use demand groupings to represent the following.

- Total traffic sourced from one specific site
- Total VPN traffic, or any defined service class, sourced from one specific site to another
- Total traffic destined for a particular AS

To generate a report on demand groupings, including costs associated with them, select the Tools->Reports->Demand Groupings menu. For detailed information, see the Demand Groupings Reports section. For a description of the cost modeling columns in the report output, see the Cost Modeling chapter.

Demand Groupings and Their Filters

The term *demand grouping* refers to a set of filters (rules) that determine a group of demands. For example, you could create a demand grouping that filters all demands sourced from Paris, and uniquely name it ParisSrc. When you apply the filter, if the result were such that 52 of the existing demands contained Paris as a source, then those 52 demands combined would comprise a single demand grouping called ParisSrc.

Each row in the Demand Groupings table identifies a uniquely named demand grouping and the associated filters that define which demands are in it.

- Each column other than Name is a criterion that a demand must meet to be in that demand grouping.
- A demand must meet all criteria defined in these columns in order for it to be a member of the demand grouping.
- If the field is empty, it is not used as a criterion.
- The column headings with "Matches" in the name use regular expressions and the filters are applied accordingly.

  Example: If Destination Node Matches is .*\.atl$, a demand can be a member of this demand grouping only if it has destination nodes ending in .atl and if it meets all other criteria (listed in the other columns).
- The column heading that use "Equals" in the name must be exact matches.

  Example: If Source Site Equals is chi, a demand can be a member of this demand grouping if its source site is named exactly chi and if it meets all other criteria.
- If the Tags Include column has an entry, a demand must contain this tag to be included in this demand grouping. (Demands can have multiple tags.)

  Example: If Tags Include is atlantic, and if a demand has three tags named europe, atlantic, and apac, then the demand would be included in this demand grouping provided it meets all other criteria.

  There are other object-based Tags Include columns for source and destination sites, nodes, and AS's. If one of these columns has an entry, the source or destination object must contain this tag to be included in this demand grouping.

  Example: If Source Site Tags Include is Europe and if the Dest Site Tags Include is Asia, the demand would be included in this demand grouping if its source site is tagged with Europe, if its destination site is tagged with Asia, and if it meets all other criteria.
- Demand Grouping columns can contain information that does not exist in the current plan. This flexibility allows a demand grouping to be applied to any network.

To see a list of demands that comprise a demand grouping, in the Demands Grouping table, right-click on the row containing the unique demand grouping name and select Filter to Associated Demands.

- In Figure 7-1, if you were to filter to associated demands on the first row (chi to bos), the resulting Demands table would show four demands because they are the only ones with both chi as the source and bos as the destination.
- If you were to filter to associated demands on the second row (chi to All), the resulting Demands table would show all demands with chi as the source, no matter the destination.

To see all demand groupings in which a demand exists, in the Demands table, right-click on the row containing the demand of interest and select Filter to Demand Groupings.

- In Figure 7-2, if you were to filter to associated demand groupings on the second row (cr1.chi_cr1.mia), the resulting Demand Groupings table would show the two demand groupings that contain this demand (chi to All, and All to mia).
- If you were to filter to associated demand groupings on the first row (cr1.bos_cr1.mia), the resulting Demand Groupings table would show the only demand grouping that contains it (All to mia).

Figure 7-1 Demand Groupings Contain Demands

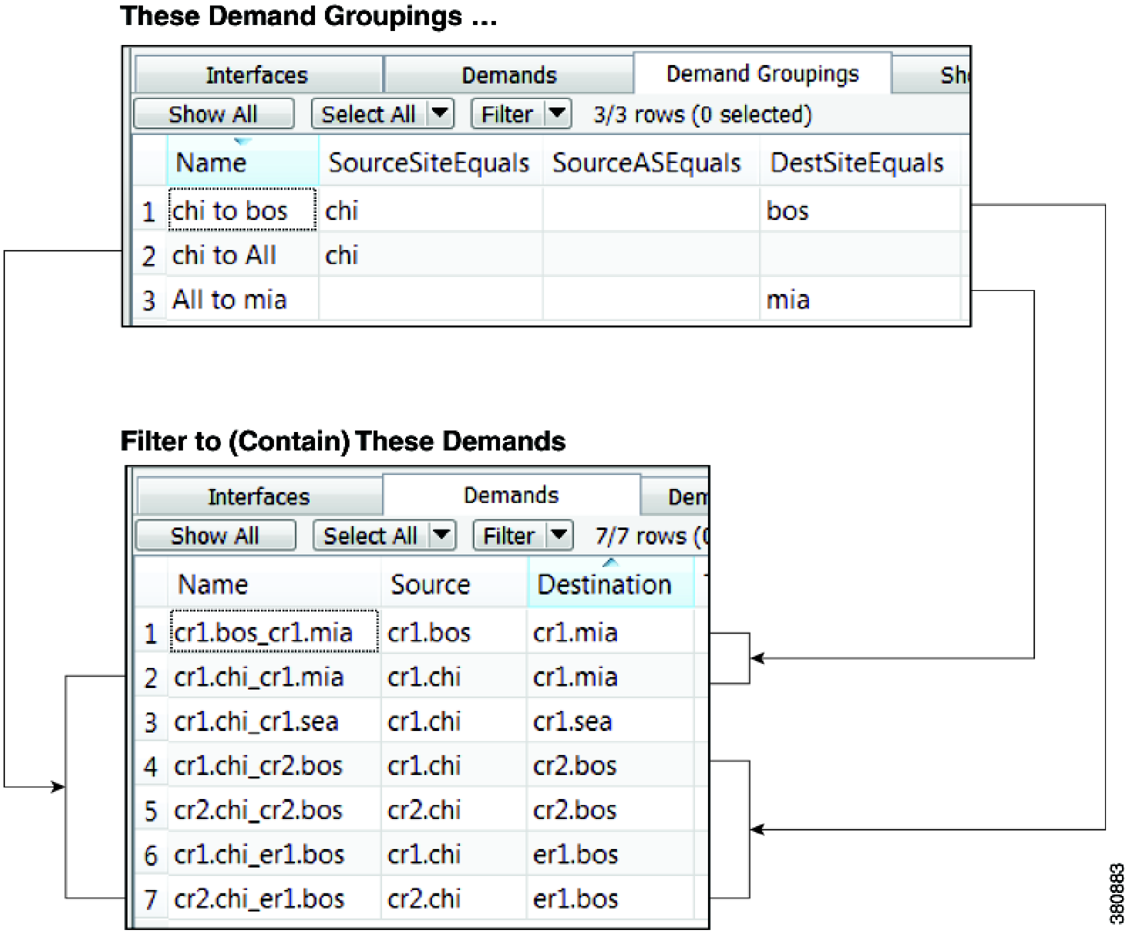

	Name	SourceSiteEquals	SourceASEquals	DestSiteEquals
1	chi to bos	chi		bos
2	chi to All	chi		
3	All to mia			mia

	Name	Source	Destination
1	cr1.bos_cr1.mia	cr1.bos	cr1.mia
2	cr1.chi_cr1.mia	cr1.chi	cr1.mia
3	cr1.chi_cr1.sea	cr1.chi	cr1.sea
4	cr1.chi_cr2.bos	cr1.chi	cr2.bos
5	cr2.chi_cr2.bos	cr2.chi	cr2.bos
6	cr1.chi_er1.bos	cr1.chi	er1.bos
7	cr2.chi_er1.bos	cr2.chi	er1.bos

Figure 7-2 Demands Contained in Demand Groupings

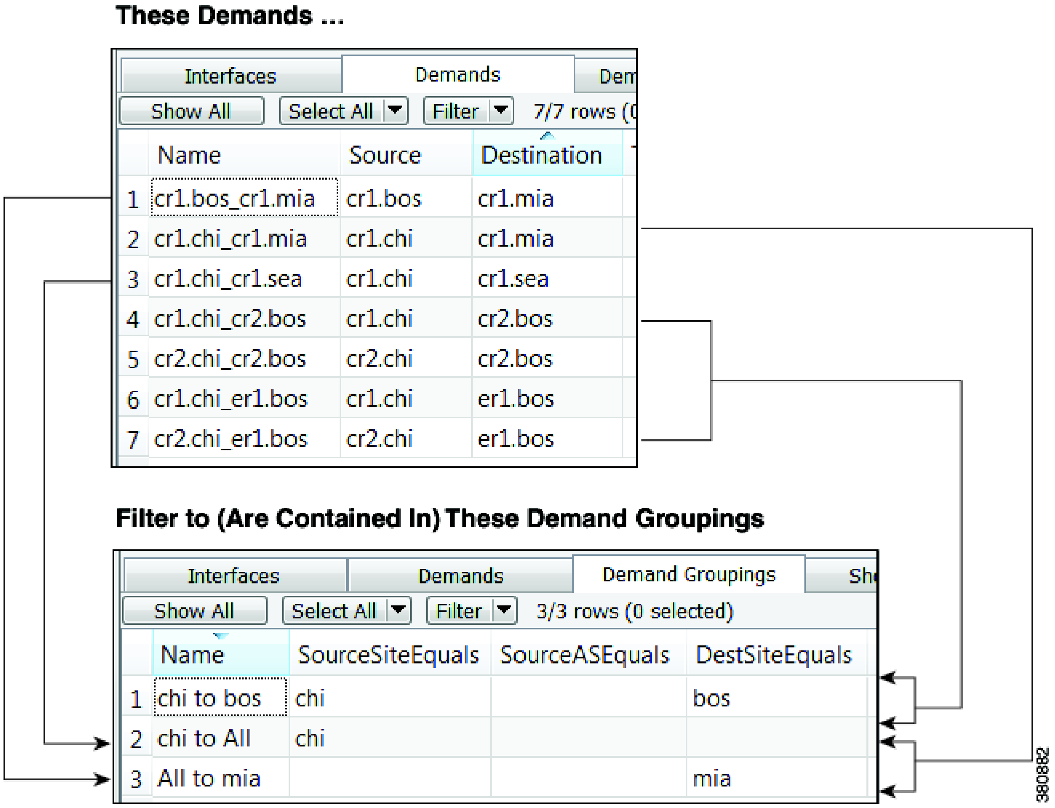

These Demands ...

	Name	Source	Destination
1	cr1.bos_cr1.mia	cr1.bos	cr1.mia
2	cr1.chi_cr1.mia	cr1.chi	cr1.mia
3	cr1.chi_cr1.sea	cr1.chi	cr1.sea
4	cr1.chi_cr2.bos	cr1.chi	cr2.bos
5	cr2.chi_cr2.bos	cr2.chi	cr2.bos
6	cr1.chi_er1.bos	cr1.chi	er1.bos
7	cr2.chi_er1.bos	cr2.chi	er1.bos

Filter to (Are Contained In) These Demand Groupings

	Name	SourceSiteEquals	SourceASEquals	DestSiteEquals
1	chi to bos	chi		bos
2	chi to All	chi		
3	All to mia			mia

Create a Single Demand Grouping

This method enables you to create a single demand grouping. All fields other than the Name field combine to create a set of rules, or filters, for the demand grouping. A demand must meet all of the specified criteria for it to be included in the demand grouping. If you specify one or more growth values, the demand grouping can be used in creating growth plans.

Note that you can also import demand groupings. For information, see the *MATE Integration and Development Guide*.

Step 1 Either select the Insert->Demands->Demand Grouping menu, or right-click in the empty plot area and select New->Demands->Demand Grouping from the context menu. The New Demand Grouping dialog box appears.

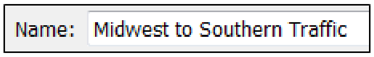

Name: Midwest to Southern Traffic

Step 2 Enter a unique name for the demand grouping. If you do not add a name, MATE Design renames the grouping to make it unique.

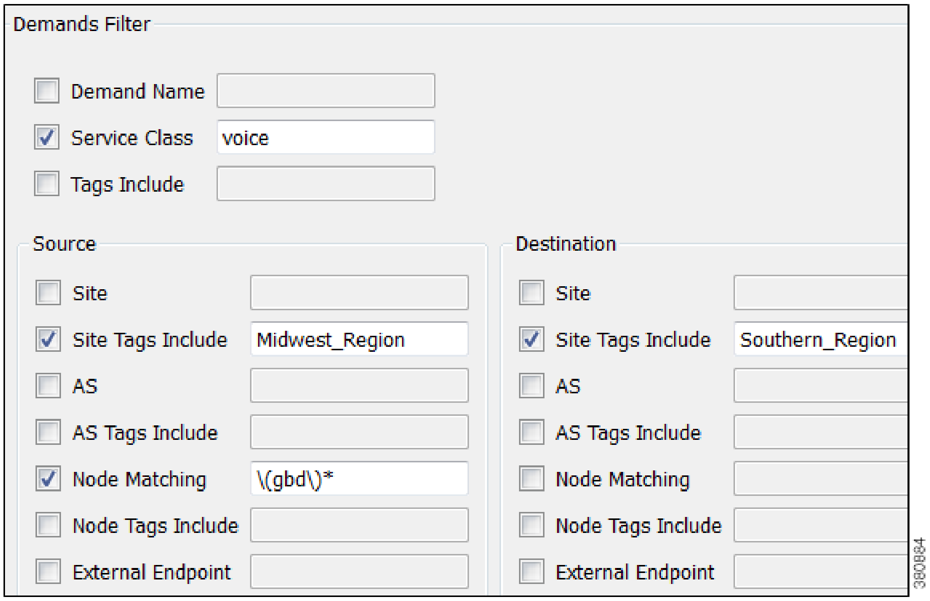

Demands Filter

□ Demand Name	
☑ Service Class	voice
□ Tags Include	

Source		Destination	
□ Site		□ Site	
☑ Site Tags Include	Midwest_Region	☑ Site Tags Include	Southern_Region
□ AS		□ AS	
□ AS Tags Include		□ AS Tags Include	
☑ Node Matching	\(gbd\)*	□ Node Matching	
□ Node Tags Include		□ Node Tags Include	
□ External Endpoint		□ External Endpoint	

Step 3 Enter information as needed in the Demands, Source, and Destination fields to create the set of rules that define the demand grouping. Each of these is a demand filter.

- You can enter only one parameter per field.
- If a field is unchecked, the filter is not used.
- If you make a selection, you must enter a value in the empty field.
  - All fields using "Tags Include" require a single tag to be specified. The list of tags for the corresponding object must include this tag.
  - All fields using "Matching" allow you to use regular expressions.
  - All fields without "Matching" require an exact match.

Step 4 If you want to use this demand grouping when creating growth plans, enter one or more values in the Growth area. See the Single Period Demand Grouping Traffic section for information on these values.

Step 5 Click OK.

Create a Demand Grouping Mesh

A mesh of demand groupings is a set of demand groupings created between a specified set of sources and destinations. It is often more convenient to create an entire mesh at once than to create the demand groupings individually.

There are two methods for quickly creating a demand grouping mesh.

- Selecting the Insert->Demands->Demand Grouping Mesh menu or the New->Demands->Demand Grouping context menu as described the following steps.
- Importing existing demand groupings from another plan file or table file. For information on importing demand groupings, see the *MATE Integration and Development Guide*.

Since you are creating a set of filters, not all steps are required. By default, the Insert Demand Grouping Mesh uses selected sites (or all sites if none are selected) as the source and does not group by destination. To change the default and further restrict the demand grouping, you can choose from any combination of source, destination, service class, and demand name, as well as use object tags. The only required steps below are #2 and #6.

If you select multiple sites before opening the dialog box and use those sites for either the source, destination, or both, a unique demand grouping is created for each selection.

Examples: If you select both chi and mia sites for the source and nothing for destination, two demand groupings are generated: one named *chi to All* and one named *mia to All*. If you select both bos and nyc as both source and destination, four demand groupings are created: *bos to bos*, *bos to nyc*, *nyc to nyc*, and *nyc to bos*.

Step 1 Select sites, nodes, and/or AS's if you want to create the demand grouping for a specific set of selected objects, whether as source, destination, or both.

Step 2 Either select the Insert->Demands-> Demand Grouping Mesh menu or right-click in an empty area and select New->Demands->Demand Grouping Mesh from the context menu. The Insert Demand Grouping Mesh dialog box appears.

Step 3 For both source and destination, you have these options.

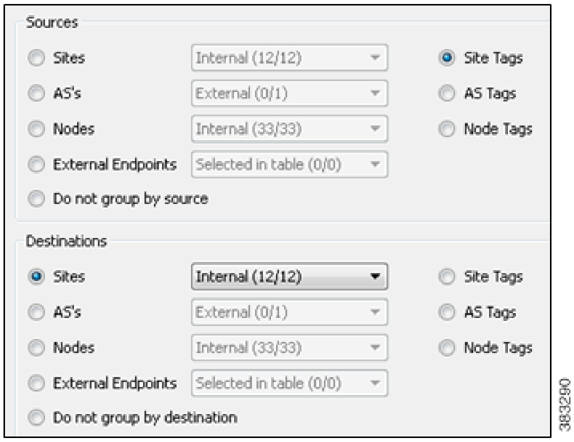

Sources

○ Sites Internal (12/12) ● Site Tags
○ AS's External (0/1) ○ AS Tags
○ Nodes Internal (33/33) ○ Node Tags
○ External Endpoints Selected in table (0/0)
○ Do not group by source

Destinations

● Sites Internal (12/12) ○ Site Tags
○ AS's External (0/1) ○ AS Tags
○ Nodes Internal (33/33) ○ Node Tags
○ External Endpoints Selected in table (0/0)
○ Do not group by destination

- Sites—All, those selected prior to opening the dialog box, internal sites only (those not containing external AS nodes), or a sites with a specific tag. For example, if you selected source sites tagged with Acme, only demand groupings whose source site was tagged with Acme would be included.
- AS's—All, those selected prior to opening the dialog box, external only, or AS's with a specific tag.
- Nodes—All, those selected prior to opening the dialog box, internal nodes only, or nodes with a specific tag.
- External Endpoints—All, those selected prior to opening the dialog box, or external endpoints with a specific tag.
- Site Tags, AS Tags, Node Tags—Objects that are tagged with any value. For instance, if you select "Node Tags" for destination nodes, demand groupings would include those demands whose destination nodes have one or more tags of any name. If a destination node was tagged with APAC and another destination node was tagged with Europe, both would be included in the demand grouping mesh.
- Do not group by—Ignore sources or destinations as filters for the demand grouping meshes.

Step 4 In addition to or instead of selecting sources or destinations, you can create the grouping by restricting it to specific demands.

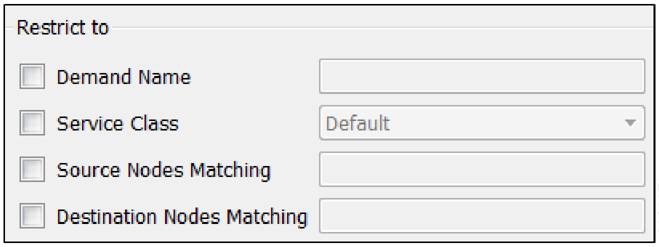

Restrict to

□ Demand Name
□ Service Class Default
□ Source Nodes Matching
□ Destination Nodes Matching

- Enter the exact demand name.
- Select a service class.
- Enter the exact source and destination node names or use regular expressions.

Step 5 The demand grouping names are created based on the name or variables used in the Name of Grouping field.

- Since each demand grouping created must have a unique name, we recommend that the string "$1 to $2" be used as the name, which is the default. The $1 variable is source site, node, or AS, or All if none of these are selected but you have chose to include sources in the grouping. Likewise, $2 is the destination site, node, AS, or All if neither is selected. Thus, a name can potentially be *All to All*.

  If a demand grouping is selected using source and destination tags, then $1 is the name of the source tag, and $2 is the name of the destination tag.

- You can this field to create uniquely named demand groupings for the same site selections by entering extensions.

  Example: You could select and use chi as the source, leave the destination empty, and add an extension of voice, thus creating a demand grouping named *chi to All voice*.

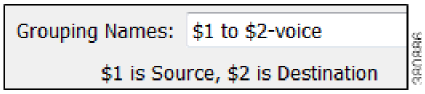

Grouping Names: $1 to $2-voice

$1 is Source, $2 is Destination

Step 6 Click OK.

Demand Groupings Reports

You can generate Demand Groupings reports on plan files that contain demand groupings. Each report has two sections.

- Demand Groupings—For each demand grouping in the plan file, this section lists the name, traffic level to which it applies, and the number of demands. All other columns are based on the total (aggregate) of all demands in the demand grouping, traffic, capacity bound, capacity violation, QoS bound, QoS violation, cost by capacity, and cost by utilization. These costs are the sum of the demands in the demand grouping.
  - Capacity Bound—The maximum amount that the sum of the demands in the demand grouping can be without causing over 100% utilization on any interface in the network. The calculation is based on the assumption that all demand traffic increases proportionally.
  - Capacity Violation—Total traffic minus the capacity permitted for the demand grouping (Capacity Bound). A violation occurs if the Capacity Bound is exceeded. If the number appearing in the Capacity Violation column is positive, then the Capacity Bound has been surpassed. If the number is negative, the Capacity Bound has not been reached.
  - For information on QoS bound and QoS violation, see the Quality of Service Simulation chapter.
  - For information on how these costs are calculated, see the Cost Modeling chapter.
- Demands—For each demand in the plan file, regardless of whether it is in a demand grouping or not, this section lists the name, source and destination, service class, and number of demand groupings to which it belongs.

To generate a Demand Groupings report, select the Tools->Reports->Demand Groupings menu, and then click OK in the confirmation dialog box.

Example of Capacity Bound and Capacity Violation

Each interface imposes a limit on the traffic that can be carried by all demands in the demand grouping.

The goal of this example is to determine which interface is limiting the demand grouping's ability to deliver traffic. To do this, MATE looks at the interfaces over which each demand is routed. MATE then calculates the maximum traffic that can be used for the demand without exceeding 100% utilization for any routed interface.

In this example, the demands and interfaces are as follows.

- All interfaces have 1,000 Mbps of traffic.
- The demand grouping is Acme_Region.
  - sjc_wdc_demand has 300 Mbps of traffic, and it traverses A_interface and B_interface.
  - sjc_hst_demand has 200 Mbps of traffic, and it traverses A_interface and C_interface.

- There are other demands outside this demand grouping.
- A_interface has 600 Mbps of traffic. Of this, 500 Mbps of it is from the Acme Region demand grouping, and the remaining 100 Mbps is from demands outside this Acme_Region demand grouping.
- B_interface has 500 Mbps of traffic, of which 300 Mbps is from the sjc_wdc_demand. The remaining 200 Mbps of traffic is from demands outside this demand grouping.
- C_interface has 400 Mbps of traffic, of which 200 Mbps is from the sjc_hst_demand. The remaining 200 Mbps of traffic is from demands outside this demand grouping.

The calculation is as follows.

- A_interface
  - A_interface can hold 900 Mbps of demand grouping traffic.

    A_interface capacity - non-demand grouping traffic

    1,000 - 100 = 900
  - The amount of sjc_wdc_demand traffic that A_interface can hold is 540 Mbps.

    Total available capacity for the demand grouping * (sjc_wdc_demand/total demand grouping traffic)

    900* (300/500) = 540
  - The amount of sjc_hst_demand traffic that A_interface can hold is 360 Mbps.

    Total available capacity for the demand grouping * (sjc_hst_demand/total demand grouping traffic)

    900* (200/500) = 360
- B_interface can hold 800 Mbps of sjc_wdc_demand traffic.

  B_interface capacity - non-demand grouping traffic

  1,000 - 200 = 800
- C_interface can hold 800 Mbps of sjc_hst_demand traffic.

  C_interface capacity - non-demand grouping traffic

  1,000 - 200 = 800
- The **limiting interface** is A_interface because it imposes the strictest limits in the demand grouping's ability to deliver traffic. It supports only 360 Mbps of sjc_hst_demand traffic, whereas B_interface supports 800 Mbps of sjc_hst_demand traffic. Also, A_interface can support only 540 Mbps of sjc_wdc_demand traffic, whereas C_interface supports 800 Mbps of sjc_wdc_demand traffic.
- The **Capacity Bound** is 900 Mbps. If the sum of the demands in the demand grouping becomes 900 Mbps, the utilization of A_interface (the limiting interface) will be 100%. This 900 Mbps is the total available capacity for the demand grouping based on the limiting interface.

  A_interface capacity - non-demand grouping traffic

  1,000 - 100 = 900
- The **Capacity Violation** is -400.

  Total A_interface traffic permitted for the demand grouping - Capacity Bound

  500 - 900 = -400

Figure 7-3 Example Demand Groupings Report

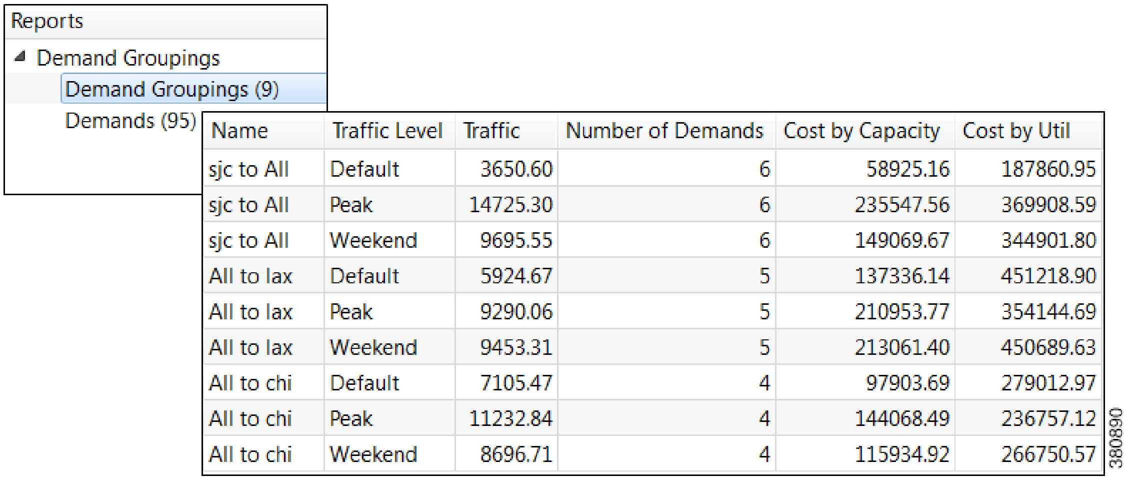

Name	Traffic Level	Traffic	Number of Demands	Cost by Capacity	Cost by Util
sjc to All	Default	3650.60	6	58925.16	187860.95
sjc to All	Peak	14725.30	6	235547.56	369908.59
sjc to All	Weekend	9695.55	6	149069.67	344901.80
All to lax	Default	5924.67	5	137336.14	451218.90
All to lax	Peak	9290.06	5	210953.77	354144.69
All to lax	Weekend	9453.31	5	213061.40	450689.63
All to chi	Default	7105.47	4	97903.69	279012.97
All to chi	Peak	11232.84	4	144068.49	236757.12
All to chi	Weekend	8696.71	4	115934.92	266750.57

Growth Plans

The Create Growth Plan tool enables you to generate new plan files containing predictions of future traffic, based on growth forecasts for the current plan. The predicted traffic growth can be generated in one of three ways.

- Per demand grouping using demand grouping traffic growth rates
- Per demand using demand traffic growth rates
- Per interface using interface traffic growth rates

Regardless of the method, the growth is compounded per period increment identified in the Create Growth Plan tool.

A series of plan files is created for each growth period, and each plan contains a report on the prediction process used to generate it. Each newly created plan shows the effects of the traffic growth visually in the plot using default traffic utilization colors, as well as in the Traffic column in the Demands table. By automatically running a simulation analysis on the new plans, you can incorporate worst-case performance into planning decisions.

For instructions on creating growth plans, see the Create Growth Plans section.

Demand Grouping Traffic Growth

Growth plans for network traffic, such as those produced either by trending analyses or by financial or sales forecasts, are typically created for traffic aggregates. You can specify these aggregate growth forecasts or trending estimates as demand grouping traffic growth information. The Create Growth Plans tool allocates traffic to demands to match the aggregate traffic growth specified for each demand grouping. If two demand groupings contain a common demand (for example, a grouping of demands sourced from one site, and a grouping of demands destined for another site), this demand traffic is balanced between the requirement to match each of the demand grouping traffic aggregates.

MATE Design supports two methods of creating growth plans based on demand groupings: one is based on a single period, while the other is for multi-period traffic forecasting.

Single Period Demand Grouping Traffic

The basic single-period method grows demands groupings based on the values identified in the Growth Traffic Total, Growth Traffic (%), and Growth Traffic Inc (increment) columns of the Demand Groupings table in the current plan. This approach is useful when all traffic in a given demand grouping is forecast to grow at the same rate or when you want to create growth plans for only one growth period.

If Growth Traffic Total is specified, the base traffic is set to this value. Otherwise, the base traffic is set to the sum of the current demand traffic in the demand grouping. The Create Growth Plans tool then applies the Growth Traffic (%), followed by the Growth Traffic Inc.

Example: If the demands in the demand grouping have a total of 15,000 Mbps of traffic, the Growth Traffic (%) is 5, and the Growth Traffic Inc is 800, the resulting total demand traffic in the first period would be 16,550 Mbps, which is (15,000 *1.05)+ 800. If you created growth plans for two period increments, the total demand traffic in second increment would be 18,177.50 Mbps, which is (16,550 *1.05)+ 800.

If the Growth Traffic Total is specified as 20,000 Mbps, then the current demand grouping traffic total is ignored, and the new demand grouping traffic total for the first period would be 21,800 Mbps, which is (20,000*1.05)+800.

You can populate these growth columns in the Demand Groupings table as follows.

- Using the Import->Demand Grouping Growth menu.
  - By importing a plan file containing the table.
  - By importing MATE Live traffic trending data.

  For information see the *MATE Integration and Development Guide*.
- Using the Edit->Traffic->Demand Groupings menu.
- By editing the Demand Groupings Properties dialog box as described in the following steps.

Step 1 If setting the same growth for multiple demand groupings, select them from the Demand Groupings table.

Step 2 Either double-click on a demand grouping, or right-click on one and select Properties from the context menu.

Step 3 Enter one or more values in the Growth area.

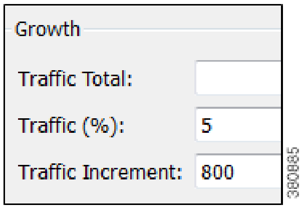

Growth

Traffic Total:

Traffic (%): 5

Traffic Increment: 800

- Traffic Total—Specifies an exact traffic total in Mbps.
- Traffic (%)—Specifies the percent by which to grow current traffic in Mbps.

  Example: If the Traffic Total is 30,000 and Traffic (%) is 5, then the growth plan grows the traffic to 31,500 Mbps.
- Traffic Increment—Specifies the amount of traffic to increment by in Mbps.

  Example: If Traffic Total is 11,000 and Traffic Increment is 800, then the growth plan grows the traffic to 11,800 Mbps

Step 4 Click OK.

Multi-Period Demand Grouping Traffic

The Demand Grouping Traffic table is a way of specifying growth based on demand groupings for one or more specified periods. Example use scenarios could be as follows.

- Traffic sourced and destined for CHI will increase by 10% next quarter
- VPN traffic between MIA and BOS will increase by 1,000 Mbps in the third quarter
- Transit traffic will reach 2 Gbps within a year

This table has the flexibility to specify the traffic total, traffic growth in specified increments (for example 10,000 Mbps), or traffic growth at a specified rate (percentage) per demand grouping.

The Demand Grouping Traffic table is created outside of MATE Design (for example, in a spreadsheet). You can then use it, including all of its periods, when creating growth plans. When used with the Create Growth Plans tool, the demand groupings in the current plan must match the names of those you are importing.

For information on creating this external table, see the *MATE Integration and Development Guide*. For information on creating growth plans, see the Create Growth Plans section.

Demand Traffic Growth

This method grows demands based on the value identified in the Growth (%) column of the Demands table in the current plan. This approach is useful when all traffic in the network, all traffic within a given traffic level, or all traffic within a service class is forecast to grow at the same rate. When generating the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

Example: If traffic on a demand is 1,000 Mbps, if the growth for that demand is defined as 10%, and if you create two growth plans, the first will show 1,100 Mbps and the second will show 1,210 Mbps on that same demand, representing 10% growth per period.

You can populate the Growth (%) column in the Demands table using either of these two methods.

- Select one or more demands, right-click on one, and select Properties from the context menu. In the Traffic section, enter the growth rate in the Growth % column.
- Import growth rates from a different plan file or table file. For instructions, see the *MATE Integration and Development Guide*.

Interface Traffic Growth

This method grows measured traffic based on the Growth (%) column in the Interfaces table in the current plan. It differs from demand-based growth methods because measured traffic, rather than demands (simulated traffic), is grown. This is useful for basic forecasting of the current network state, rather than worst-case failure states, which require simulation. When generating the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

You can also use this method in conjunction with Demand Deduction to create demands from the forecast measurements.

Example: If traffic on an interface is 525 Mbps, if the growth for that interface is defined as 5%, and you specify three growth plans, the first will show 551.25 Mbps, the second will show 578.81 Mbps, and the last will show 607.75 on the same interface, representing 5% growth per period.

You can populate the Growth (%) column in the Interfaces table using either of these two methods.

- Select the Edit->Traffic->Measured Traffic menu. In the Growth (%) column of the Interfaces table, enter the percent by which you want to grow the interface traffic for each applicable interface.

  To set growth rates for multiple interfaces to the same value or to radom values, select one or more interfaces, right-click, and select Set Value. You then have the option to set fixed values for all selected interfaces or have MATE Design randomly set them within the range that you specify.
- Import growth rates from a different plan file or table file. For instructions, see the *MATE Integration and Development Guide*.

Create Growth Plans

Step 1 Select the Tools->Create Growth Plans menu; a Create Growth Plans dialog box appears.

Step 2 Specify one of the following growth methods.

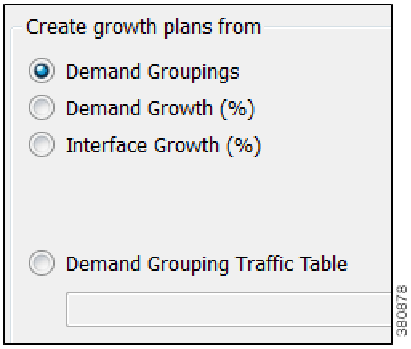

Create growth plans from

● Demand Groupings
○ Demand Growth (%)
○ Interface Growth (%)
○ Demand Grouping Traffic Table

- Demand Groupings—For information on how to set demand grouping growth based on fixed, incremental, or percents, see the Demand Grouping Traffic Growth section.
- Demand Growth (%) in current plan—For information on how to set this percent, see the Demand Traffic Growth section.
- Interface Growth (%) in current plan—For information on how to set this percent, see the Interface Traffic Growth section.
- Demand Grouping Traffic Table—Browse to or enter the name of the file containing the Demand Grouping Traffic Table. For information on how to create these tables, see the *MATE Integration and Development Guide*.

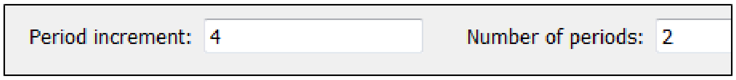

Period increment: 4 Number of periods: 2

Step 3 If creating growth plans based on demand groupings, demands, or interfaces, enter the number of plans to create and the increment by which you want to create them.

- Period increment—Number of percentage increments per period. For example, if a demand is set to grow at 10%, and the period increment is 2, then the demand will grow by 21% in the first growth plan.
- Number of periods—Number of growth plan files to create.

Plans are named using the period increment. For example, if you enter 4 as the increment and 2 as period, MATE Design generates two growth plan files: one ending with _4 and one ending with _8.

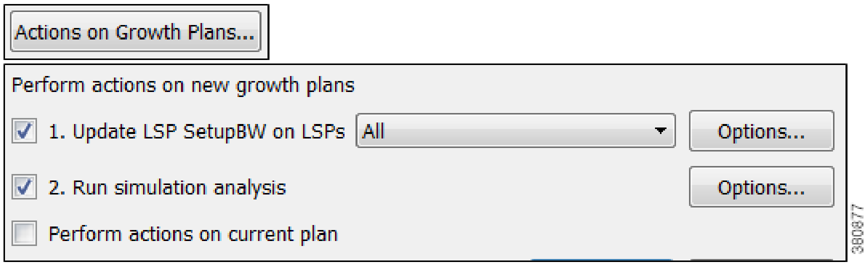

Perform actions on new growth plans

☑ 1. Update LSP SetupBW on LSPs: All

☑ 2. Run simulation analysis

□ Perform actions on current plan

Step 4 Click on the Actions on Growth Plans button to perform these two tasks on the newly created growth plan files. To further specify how to execute them, click on their associated Options button. You can also choose to execute these tools on the current plan.

- Upgrade LSP setup bandwidth on LSPs—For more information, see the MPLS Simulation chapter.
- Run simulation analysis—For more information, see the Simulation Analysis chapter.

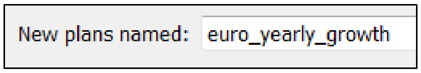

New plans named: euro_yearly_growth

Step 5 By default, MATE Design uses the current plan file name as the root name and adds the extensions generated by step #3 for each growth plan. If you want the new plan files to use a different root name, enter that name in the New plans named field.

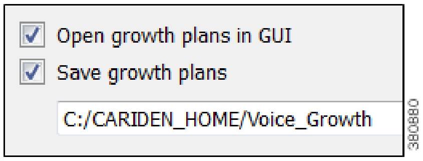

☑ Open growth plans in GUI

☑ Save growth plans

C:/CARIDEN_HOME/Voice_Growth

Step 6 To open these new plans upon exiting the dialog box, select Open growth plans in GUI.

Step 7 To automatically save these new plan files, select Save growth plans, and then enter or browse to the directory in which you want to save the files.

Step 8 Click OK.

IGP Simulation

This chapter describes options for simulating an IGP. MATE Design can simulate IS-IS, OSPF, and EIGRP, as well as supports multi-topology routing for IS-IS and OSPF IGPs.

All simulations use demands. For detailed information on demands, see the Traffic Demand Modeling chapter.

IGP Simulation Options

The IGP simulation options are OSPF, IS-IS, and EIGRP.

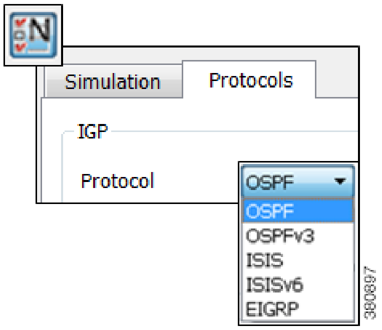

Protocol OSPF

- IS-IS and OSPF simulations are identical with the following exceptions.
  - OSPF routing uses OSPF areas, if specified. By default all interfaces are assigned area zero.
  - IS-IS routing uses IS-IS levels, if specified. By default, all interfaces are set to Level 2. Interfaces can belong to Level 1, Level 2, or both. If both, then an alternate metric for Level 1 can also be specified. You can change these properties using the Advanced tab in an interface Properties dialog box.
- EIGRP uses the following rules to calculate interface metrics for shortest-path calculations.

  metric = [K1 * bandwidth + K2 * bandwidth / (256 - load) + K3 * delay] * [K5 / (reliability + K4)]

  - MATE Design only supports bandwidth and delay parameters.

    K2 = K4 = K5 = 0, and K1 = K3 = 1
  - Variance=1 (standard equal cost load-balancing).
  - Bandwidth is in kbps, and delay is in tens of microseconds.
  - MATE Design uses Mbps and milliseconds. The MATE Design conversion of the above formula is as follows.

    metric = (10,000/(bandwidth in Mbps) + 100 * (delay in ms)) * 256
  - The last multiplication of [K5/(reliability+K4)] is only performed if K5 > 0.

To set these IGP options, follow these steps.

Step 1 Select the Edit->Network Options menu or the Network Options icon from the Visualization toolbar.

Step 2 Select the Protocols tab.

Step 3 Select the appropriate protocol from the IGP Protocol list.

Step 4 Click OK.

OSPF Area Simulation

OSPF area membership can be specified per interface. The two interfaces on each circuit must belong to the same area. Area names can be any string. Area zero, the backbone area, must be denoted by ‘‘0’’, ‘‘0.0.0.0’’, or an empty string. MATE Design simulates the OSPF area routing configuration in which the areas import LSAs from the backbone. A demand from a source node to a destination node in a different area will only be routed if it can reach the destination by passing through the source area, directly to area zero, and from there directly to the destination area.

By default, all nodes in a single AS are assumed to belong to a single OSPF area. Nodes are assigned to areas, as follows.

- If OSPF areas are not defined for interfaces, all nodes are assumed to be in the same area.
- Each interface can be assigned to only one OSPF area. Note that each node can be assigned to one or more areas.
- If a node has an interface in an OSPF area, the node is assigned to that area.
- An Area Border Router (ABR) is a node belonging to both area 0 and other OSPF areas.

IS-IS Multi-Level Simulation

By default, IS-IS interfaces are assigned to Level 2, though you can assign them to Level 1 or to both Levels 1 and 2. (For instructions, see the Set IS-IS Multi-Level Simulation section.)

- The IGP metric defines the Level 1 metric for interfaces in Level 1, in Level 2, or in both Level 1 and 2 if these have equal metrics. However, you can change the Level 1 metric.
- If an interface is in both Level 1 and Level 2 with uneven metrics, the IGP metric defines the Level 2 metric and the Level 1 metric defines the Level 1 metric.

The IS-IS level is listed in the ISIS Level column of the Interfaces table, and the Level 1 metric is listed in the Metric Level 1 ISIS column.

By default, all nodes in a single AS are assumed to belong to a single IS-IS level. Nodes are assigned to levels, as follows.

- Each node can be assigned to one or more levels.
- Two Level 1 nodes are placed in different areas if any route between them (under normal operation) passes through a Level 2 node.
- A node is assigned to a single Level 2 area if there is at least one interface that is Level 2 or both Level 1 and Level 2.
- A node is assigned to one of potentially multiple Level 1 areas if any interface is Level 1 or both Level 1 and Level 2.
- An ABR is a node belonging to both the Level 2 area and another IS-IS area.

Set IS-IS Multi-Level Simulation

Follow these steps to set IGP metrics, IS-IS levels, and Level 1 metrics.

Step 1 Select one or more interfaces. Either right-click on one and select Properties from the context menu or double-click on an interface.

IGP Metric: 36

Step 2 To set the IGP metric, enter its value in the IGP metric field or select it from the drop-down list.

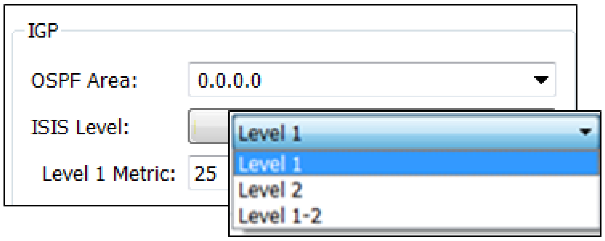

IGP

OSPF Area: 0.0.0.0

ISIS Level: Level 1

Level 1 Metric: 25

Step 3 To set IS-IS multi-level simulation, select the Advanced tab.

a. In the IGP area, select the level from the ISIS Level drop-down list.

b. In the Level 1 Metric field, enter the metric.

Step 4 Click OK.

IGP Multipath Simulation

By default, demands are split equally between all paths from a route of equal distance to the destination, and there is no ECMP limit.

You can specify a maximum number of ECMP paths. In this case, demands transiting through a router are distributed among the available paths up to this maximum ECMP value. Paths are chosen by lowest next-hop IP address, which is the interface IP address for IGP or the destination IP address for LSPs. Paths with no IP address are chosen last.

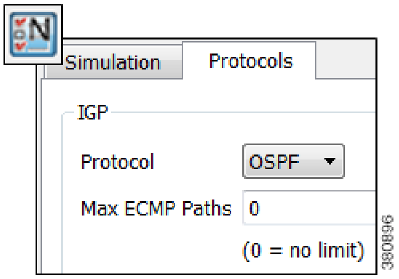

IGP

Protocol OSPF

Max ECMP Paths 0

(0 = no limit)

To configure ECMP, follow these steps.

Step 1 Select the Edit->Network Options menu or the Network Options icon from the Visualization toolbar.

Step 2 Select the Protocols tab.

Step 3 Enter the maximum number of ECMPs in the Max ECMP Paths field. A 0 (zero) means there is no limit.

Step 4 Click OK.

View Shortest IGP, Latency, and TE Paths

The Shortest Paths table is dynamically generated to show routes that traffic would take between selected nodes or between all nodes in selected sites. Columns include information about the total IGP metrics used in the routes, as well as numerous latency columns. Since shortest paths are not objects in the plan file, there is no associated Properties dialog box. You can, however, right-click on paths to filter to interfaces or to source and destination nodes to change their properties, as needed, to change the path.

- Shortest IGP paths are the shortest paths as calculated by IGP metric.
- Shortest latency paths are the shortest paths in an IGP network as calculated by the circuit latencies.
- Shortest TE paths are the shortest paths as calculated by the interface TE metric.

To view these shortest paths in the network plot, select them from the table (Figure 8-1) using the following steps.

Step 1 Select one or more nodes, or select one or more sites. If you select both, MATE uses the object type from which you select the context menu. It does not use both simultaneously.

Step 2 Right-click on one of the objects. For nodes, select the Filter to->Shortest Paths-><selection> context menu. For sites, select Filter to intersite Shortest Paths-><selection>. The selection is any one of the following.

- Between selected nodes/sites (applicable only for multiple selections)
- With selected nodes/sites as source
- With selected nodes/sites as destination
- Through all selected nodes/sites
- Through any selected nodes/sites

Step 3 The Shortest Paths table is populated according to your selections: shortest IGP, shortest TE, or shortest latency. To view a path in the plot, select it from this table. The Type column identifies the path type.

Figure 8-1 Shortest IGP, Latency, and TE Paths

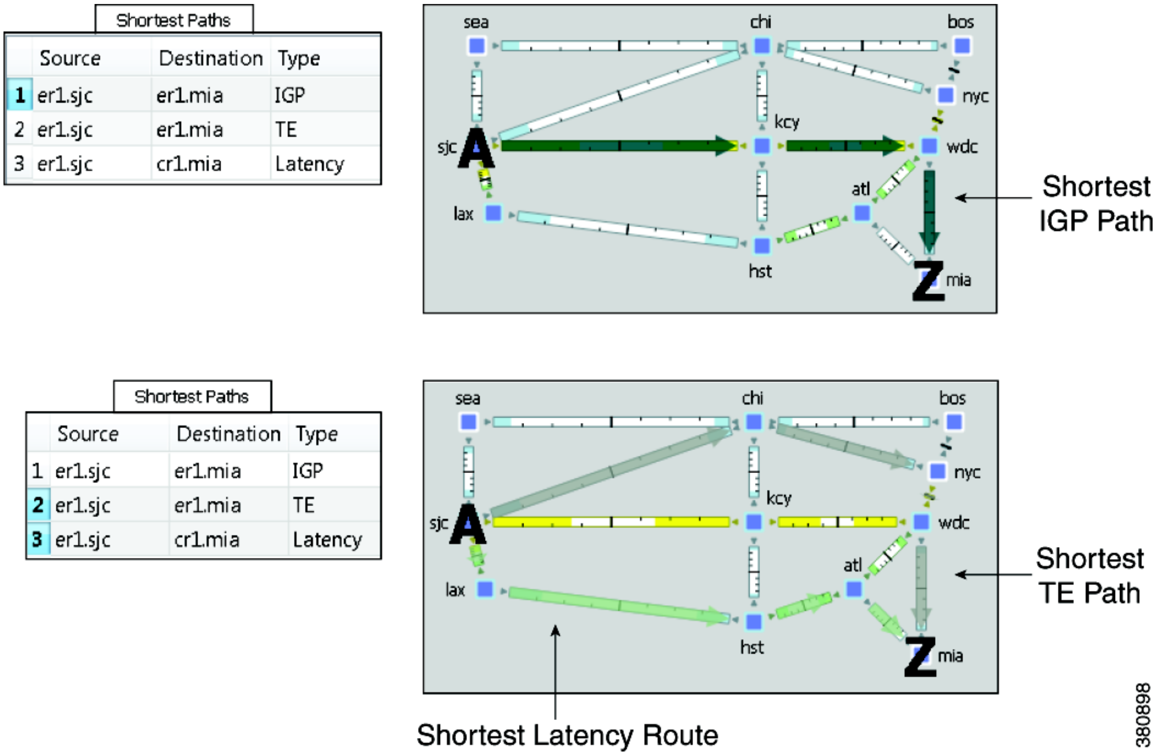

Shortest Paths

	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 IGP Path

Shortest TE Path

Shortest Latency Route

Plot Shortest Path

From the Shortest Paths table, right-click a path and select Plot Shortest Path from the context menu. From here, you can see each node in the shortest path, as well as the traffic utilization on interfaces between each node in the path. Clicking on the Select Shortest Paths button highlights the path in the network plot.

Figure 8-2 Shortest Path Plot

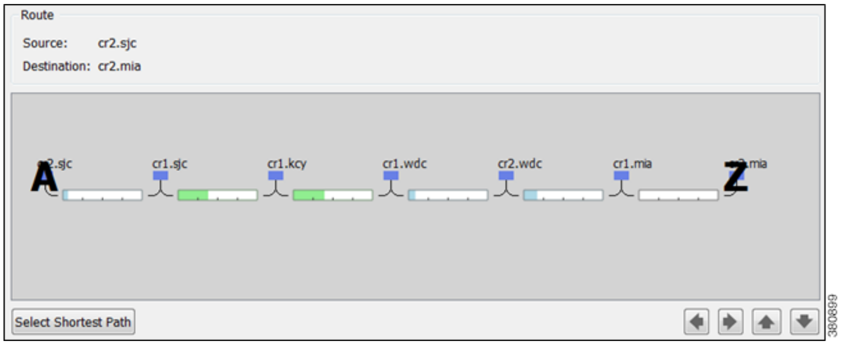

Route

Source: cr2.sjc

Destination: cr2.mia

Multi-topology Routing

MATE Design supports simulation of multi-topology routing. Interfaces can be assigned to one or more IGPs. Demands and LSPs can be assigned to a specific IGP, and will only route through interfaces belonging to that IGP. This multi-topology simulation uses these rules.

- When an interface is common to more than one topology, all IGP properties of the interface (including the metric) must be the same in both topologies.
- The IGPs must all be defined to be of the same type, for example OSPF or IS-IS. This is not a restriction in practice unless either the OSPF topology or the IS-IS topology uses multiple areas.
- Demands and LSPs defined with a specific topology can only route over circuits belonging to that topology. A circuit belongs to a topology if either one of its interfaces does.
- Demands and LSPs with no defined topology have no restrictions on routing. They are routed using the default topology, to which all interfaces belong.

Edit Topologies

To create, rename, or delete topologies, select the Edit->Topologies menu; the Edit Topologies dialog box appears (Figure 8-3).

To...	Do This...	Enter...	Click...
Create a new topology	Click New	A topology name	OK, and then Done
Rename an existing topology	Select the topology, and then click Rename	A different topology name	OK, and then Done
Remove a topology	Click Delete		Done

Figure 8-3 Edit Topologies

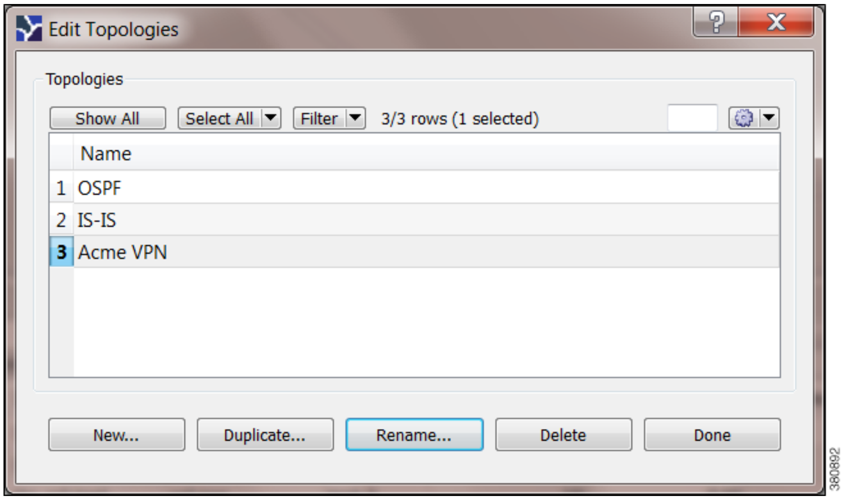

3/3 rows (1 selected)

	Name
1	OSPF
2	IS-IS
3	Acme VPN

Add Topologies to Objects

Follow these steps to associate interfaces, demands, or LSPs with topologies. Once you have associated the topologies, you can view them by showing the Topology column.

Step 1 Select one or more interfaces, demands, or LSPs from their respective tables.

Step 2 Right-click one of the selected objects, and select Properties from the context menu.

Step 3 Select the desired topology or topologies, as follows.

Interfaces	Demands and LSPs
1. Select the Advanced tab. **2.** In the Topologies section, click the Edit button. **a.** For each topology that you want on this interface, click the F (false) to change it to T (true). **b.** Click OK. **3.** Click OK.	**1.** Select the desired topology from the Topology drop-down list. You can associate one topology to each demand or LSP. **2.** Click OK.

By selecting from a similar Topology drop-down list, you can also add a topology to LSP meshes when creating them.

IGP Metric Initializer

The IGP Metric Initializer is a quick way of setting IGP metrics based on circuit latency values. MATE Design calculates the metric by multiplying each circuit's Delay Sim value by the multiplier you specify. You also set a minimum value for the resulting IGP metric. MATE Design then uses the higher of the two values.

Example: If the circuit Delay Sim value is 15, the multiplier value is 2, and the minimum value is 1, the resulting IGP metric is 30 (Delay Sim * Multiplier). However, if the minimum value were 50, the resulting IGP metric would be 50.

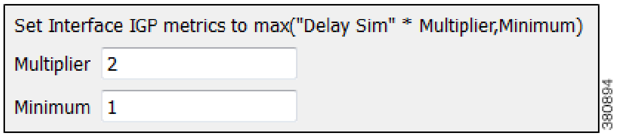

Set Interface IGP metrics to max("Delay Sim" * Multiplier,Minimum)

Multiplier 2

Minimum 1

Since MATE Design updates all metrics on all selected interfaces, you can modify interfaces individually or in groups. For example, you might select all core routers in the Interface table and then simultaneously calculate their IGP metrics.

Step 1 Make your selections.

- Select one or more interfaces.
- Select one or more circuits. This updates IGP metrics for both associated interfaces.
- If you make no selections, IGP metrics are calculated for all interfaces.

Step 2 Either select the Initializers->IGP Metric menu, or right click an interface or circuit, and select Initialize IGP Metric from the context menu.

Step 3 Enter a multiplier and minimum value.

Step 4 Click OK.

CHAPTER 9

MPLS Simulation

This chapter describes how MATE 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.

Terms

- RSVP LSPs— LSPs that are established through RSVP-TE. These are commonly known as *MPLS TE tunnels*. WAE Collector discovers RSVP LSPs, and you can also create them using the MATE GUI. They are identified with a Type property of RSVP.
- FRR LSP—Fast Reroute LSPs, which provide node or link protection to RSVP LSPs in the event of failures. WAE Collector discovers FRR LSPs, and you can also create them using the MATE GUI. They are identified with the Metric Type property of FRR.
- SR LSPs—Segment routing LSPs, which do not use RSVP for routing. You can create SR LSPs using the MATE GUI. They are identified with a Type property of SR.
- LDP LSPs—LSPs that are established through LDP. WAE Collector discovers traffic from LDP LSPs, but not the LSPs themselves. MATE does not distinguish them from RSVP LSPs.
- PCEP LSPs—LSPs that are discovered through Open Daylight PCEP and added to the set of LSPs discovered through Collector. They are identified with a PCEP property of T (true).
- LSP paths—Primary and secondary routes established by the LSP, and are often called *MPLS TE tunnel paths*. They have a numerical Path Option property that identifies their priority.

Creating and Visualizing LSPs

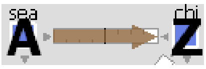

To simulate MPLS LSPs in the network, you must first set up the LSPs to be routed. In the MATE GUI, enter LSPs into the plan file in the following ways.

- 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 on one or more LSPs and select the option from the context menu.

Figure 9-1 LSP Plot

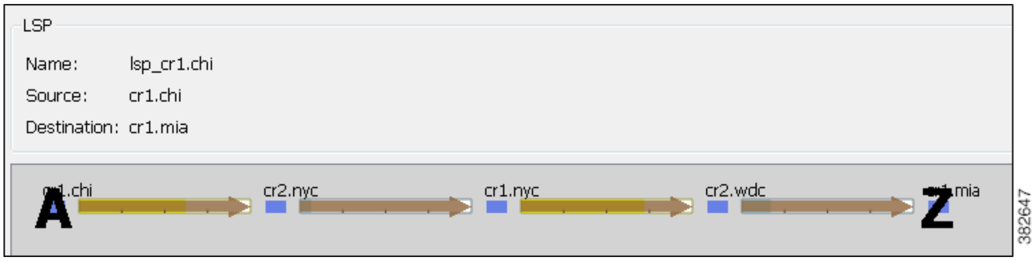
LSP
Name: lsp_cr1.chi
Source: cr1.chi
Destination: cr1.mia

Routing RSVP LSPs

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

Dynamic LSP Routing and CSPF

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

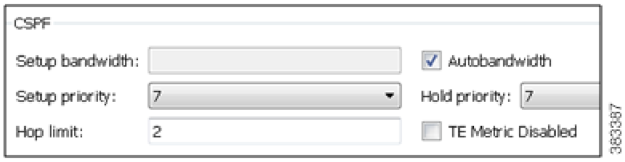
CSPF
Setup bandwidth: ☑ Autobandwidth
Setup priority: 7 Hold priority: 7
Hop limit: 2 □ TE Metric Disabled

- 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.
- 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.

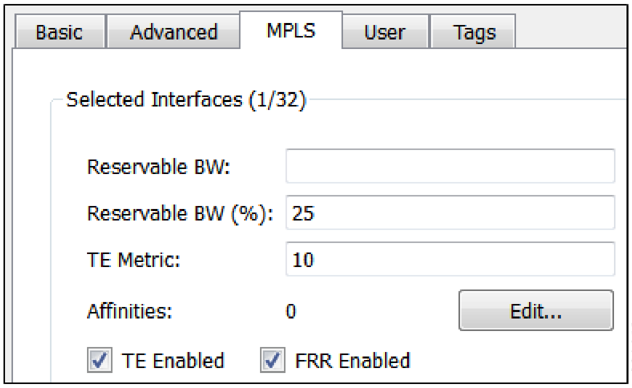

Selected Interfaces (1/32)

Reservable BW:

Reservable BW (%): 25

TE Metric: 10

Affinities: 0 Edit...

☑ TE Enabled ☑ FRR Enabled

381119

- 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.

    Capacity Sim * (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 traff 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 on 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

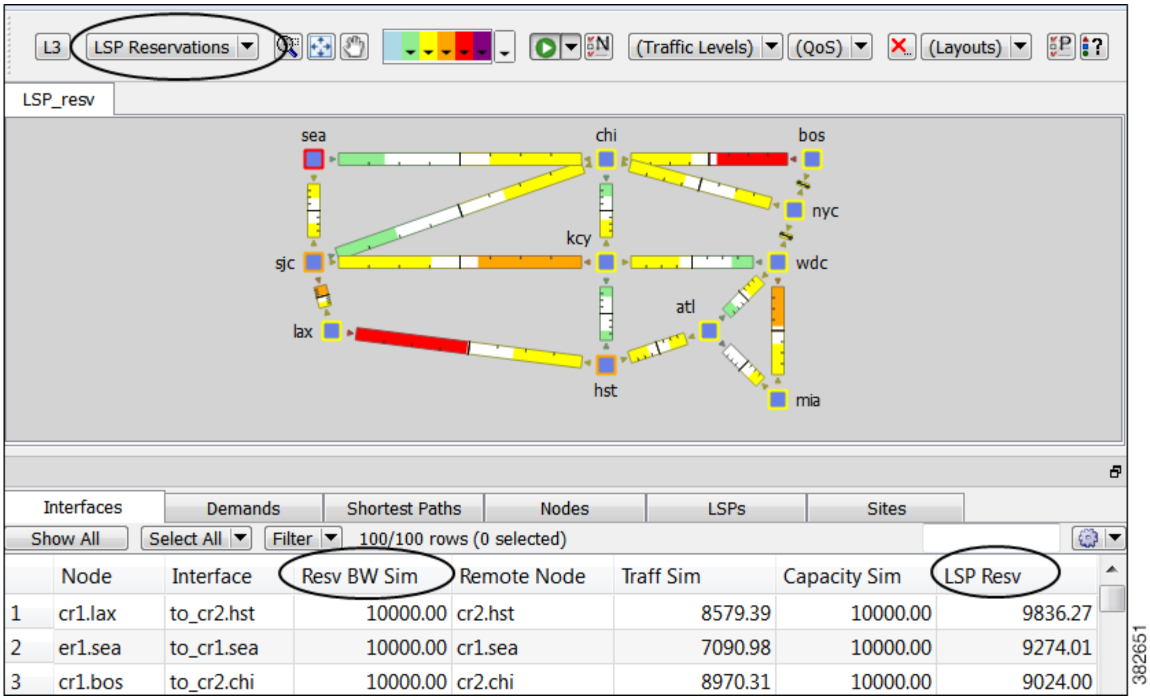

100/100 rows (0 selected)

	Node	Interface	Resv BW Sim	Remote Node	Traff Sim	Capacity Sim	LSP Resv
1	cr1.lax	to_cr2.hst	10000.00	cr2.hst	8579.39	10000.00	9836.27
2	er1.sea	to_cr1.sea	10000.00	cr1.sea	7090.98	10000.00	9274.01
3	cr1.bos	to_cr2.chi	10000.00	cr2.chi	8970.31	10000.00	9024.00

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.

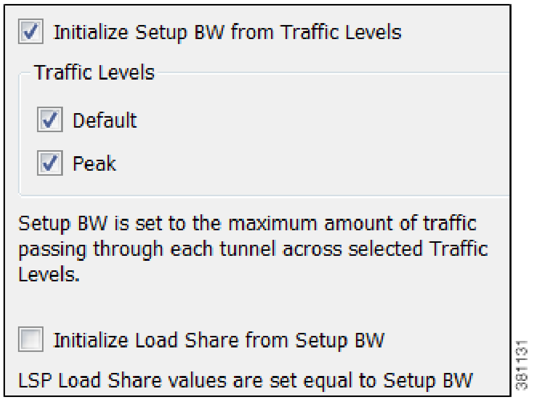

☑ Initialize Setup BW from Traffic Levels

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 ore more LSPs from the LSPs table.

Step 2 Either select the Initializers->LSP Setup BW menu, or right-click on 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.

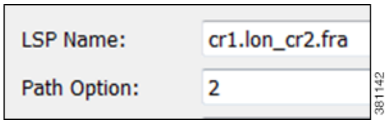

LSP Name: cr1.lon_cr2.fra

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.

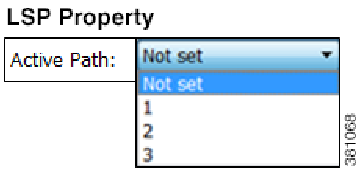

Active Path: Not set

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.

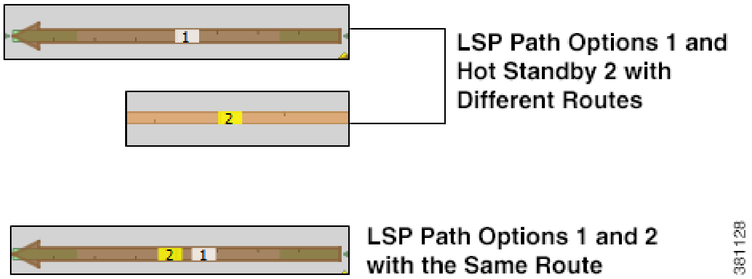
1
2
LSP Path Options 1 and Hot Standby 2 with Different Routes
2 1
LSP Path Options 1 and 2 with the Same Route
381128

Note This 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

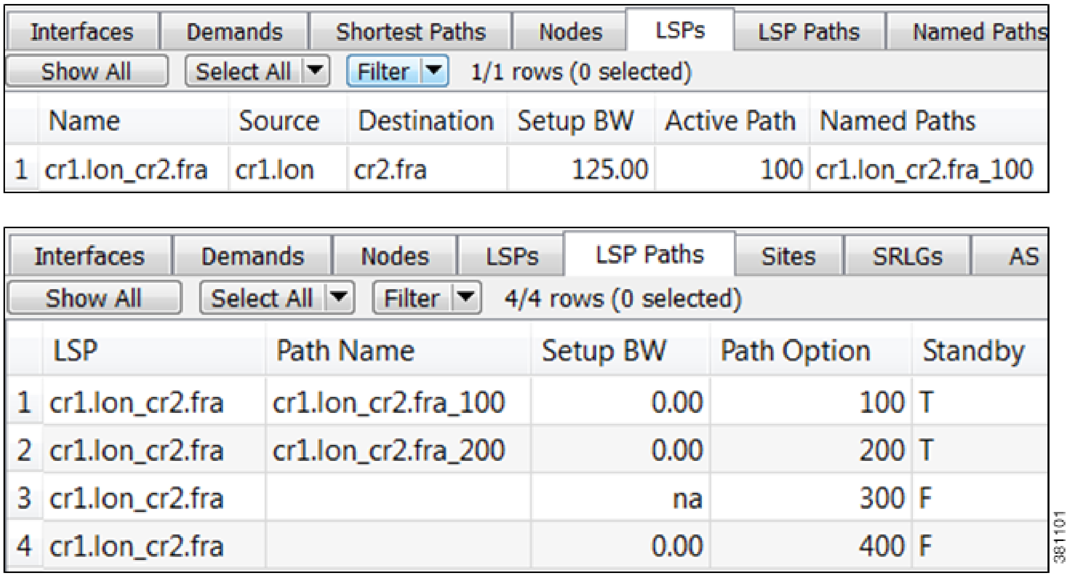

	Name	Source	Destination	Setup BW	Active Path	Named Paths
1	cr1.lon_cr2.fra	cr1.lon	cr2.fra	125.00	100	cr1.lon_cr2.fra_100

	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), MATE 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.

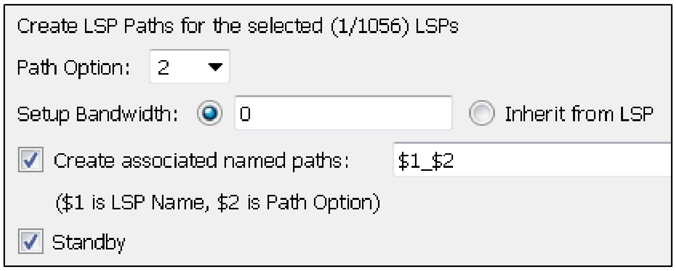

Create LSP Paths for the selected (1/1056) LSPs

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 on 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 though 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.

Strict

Loose

Exclude

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 MATE 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

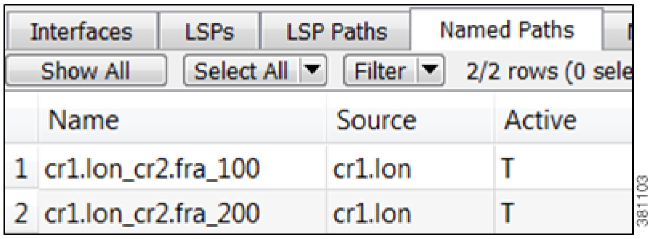

	Name	Source	Active
1	cr1.lon_cr2.fra_100	cr1.lon	T
2	cr1.lon_cr2.fra_200	cr1.lon	T

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

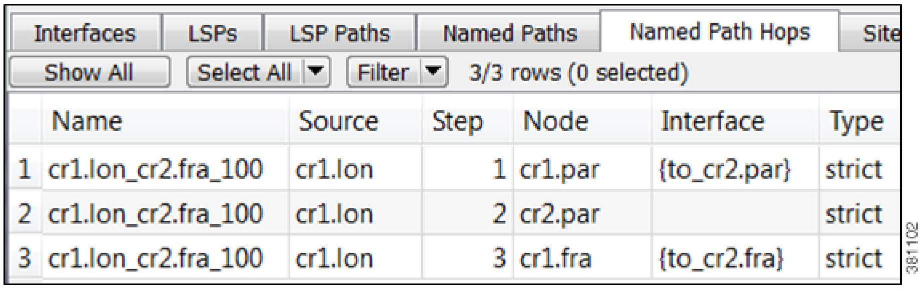

	Name	Source	Step	Node	Interface	Type
1	cr1.lon_cr2.fra_100	cr1.lon	1	cr1.par	{to_cr2.par}	strict
2	cr1.lon_cr2.fra_100	cr1.lon	2	cr2.par		strict
3	cr1.lon_cr2.fra_100	cr1.lon	3	cr1.fra	{to_cr2.fra}	strict

Create Named Paths and Their Hops

Named paths can be created in the MATE GUI in several ways.

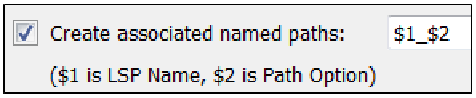
☑ Create associated named paths: $1_$2
($1 is LSP Name, $2 is Path Option)

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

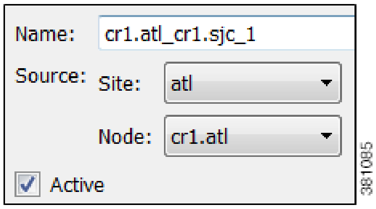
Name: cr1.atl_cr1.sjc_1
Source: Site: atl
Node: cr1.atl
☑ Active

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.

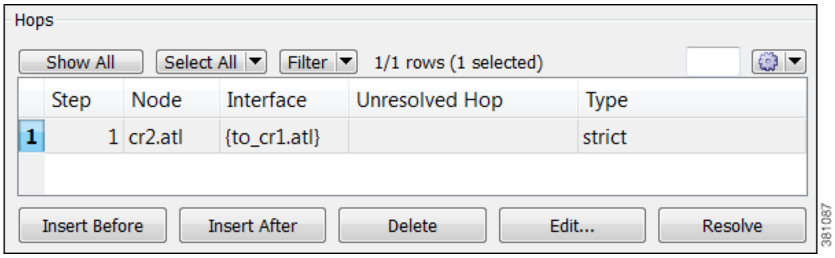

1/1 rows (1 selected)

	Step	Node	Interface	Unresolved Hop	Type
1	1	cr2.atl	{to_cr1.atl}		strict

Step 3 Hop options

- To add a new hop, select an existing hop (if applicable) from the Hops list, and then click on 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.

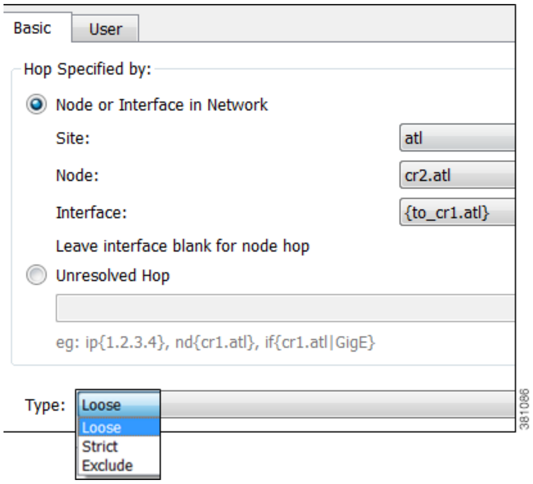

● Node or Interface in Network

Site: atl

Node: cr2.atl

Interface: {to_cr1.atl}

Leave interface blank for node hop

○ Unresolved Hop

eg: ip{1.2.3.4}, nd{cr1.atl}, if{cr1.atl|GigE}

Type: Loose

Step 5 Click OK.

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***

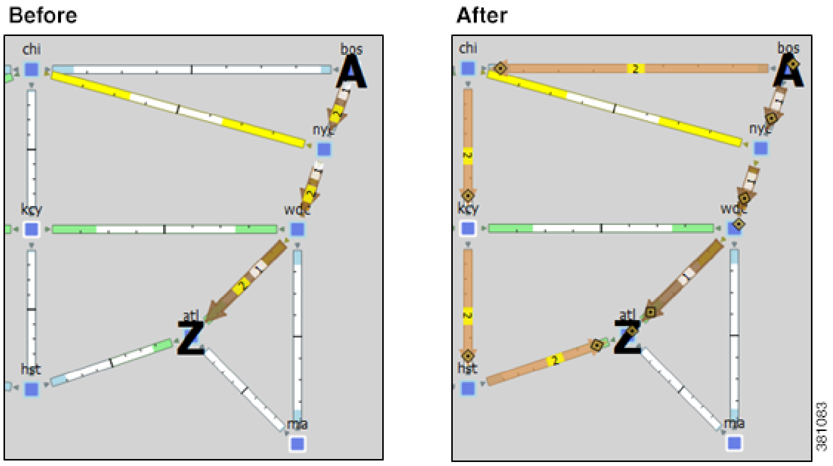

MATE Design creates named paths and named path hops for the LSP paths. After running the initializer, you can verify the named paths and their hops in the Named Paths and Named Path Hops table, respectively.

- 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.

Pre-requisites

- 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.

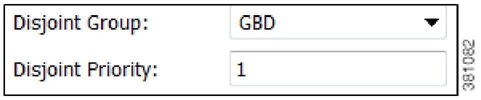
Disjoint Group: GBD
Disjoint Priority: 1

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 on 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.

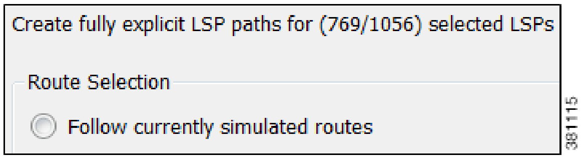
Create fully explicit LSP paths for (769/1056) selected LSPs

Route Selection

○ Follow currently simulated routes

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

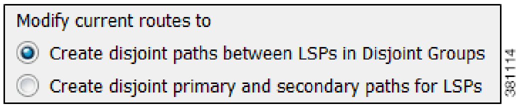
Modify current routes to

● Create disjoint paths between LSPs in Disjoint Groups

○ Create disjoint primary and secondary paths for 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, MATE 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.

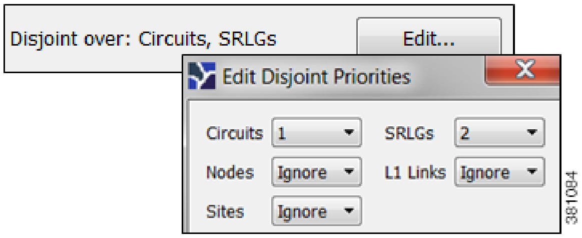
Disjoint over: Circuits, SRLGs

Edit Disjoint Priorities

Circuits: 1
SRLGs: 2
Nodes: Ignore
L1 Links: Ignore
Sites: Ignore

Step 4 Click OK.

Actual Paths

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

Example: MATE 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 on 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.

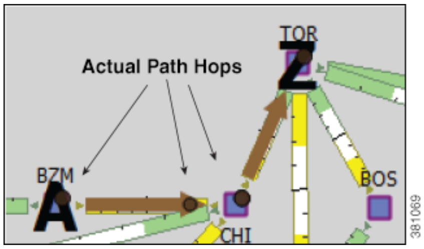

Column	Description
Actual	The simulation can only use actual paths if they are resolved. If the network discovery was incomplete, this might not be possible. • T = true; the simulation converted the actual path hops into a complete path from source to destination of the LSP • F = false; the simulation did not convert the actual path hops into a complete path • na = not applicable; no actual path was available.
Route	• 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

MATE 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.

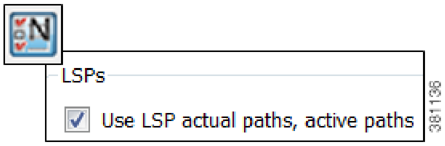
LSPs

☑ Use LSP actual paths, active 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.

MATE 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, like Gold, Silver, or Bronze.

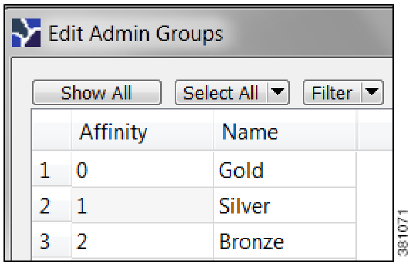

	Affinity	Name
1	0	Gold
2	1	Silver
3	2	Bronze

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.

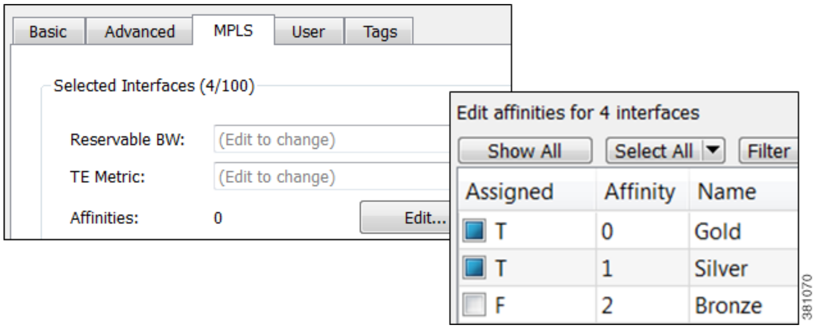

Selected Interfaces (4/100)

Reservable BW: (Edit to change)

TE Metric: (Edit to change)

Affinities: 0

Assigned	Affinity	Name
☑ T	0	Gold
☑ T	1	Silver
☐ F	2	Bronze

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 on the selections or select the Edit->Properties menu; the Properties dialog box appears.

Step 3 Select the MPLS tab.

a. Click the Affinities Edit button; the Edit Affinities dialog box is displayed.

b. 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.

c. 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.)

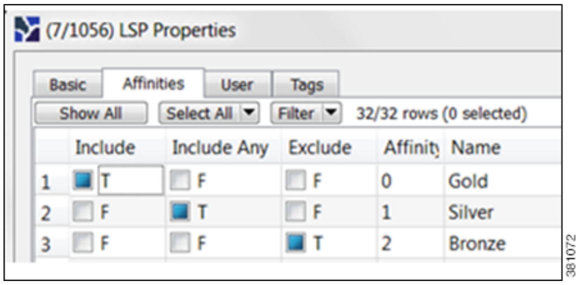

	Include	Include Any	Exclude	Affinity	Name
1	☑ T	☐ F	☐ F	0	Gold
2	☐ F	☑ T	☐ F	1	Silver
3	☐ F	☐ F	☑ T	2	Bronze

Step 1 From either the LSPs table, select one or more LSPs to which you are associating the affinities.

Step 2 Either double-click on 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 MATE, 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 ABRs[1]. (See 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.

[1] 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.

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 9-8 shows that by deselecting this network option, the inter-area LSP that does not have hops set at ABRs can dynamically route.

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

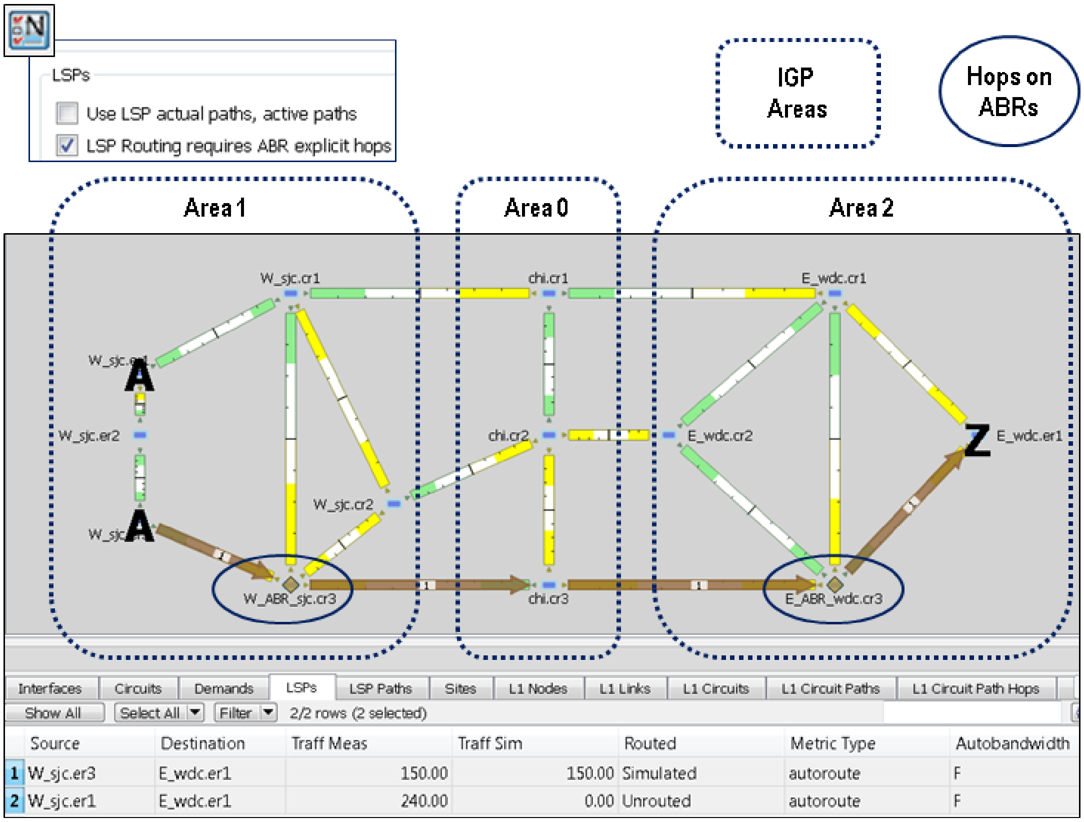

LSPs

☐ Use LSP actual paths, active paths

☑ LSP Routing requires ABR explicit hops

2/2 rows (2 selected)

	Source	Destination	Traff Meas	Traff Sim	Routed	Metric Type	Autobandwidth
1	W_sjc.er3	E_wdc.er1	150.00	150.00	Simulated	autoroute	F
2	W_sjc.er1	E_wdc.er1	240.00	0.00	Unrouted	autoroute	F

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

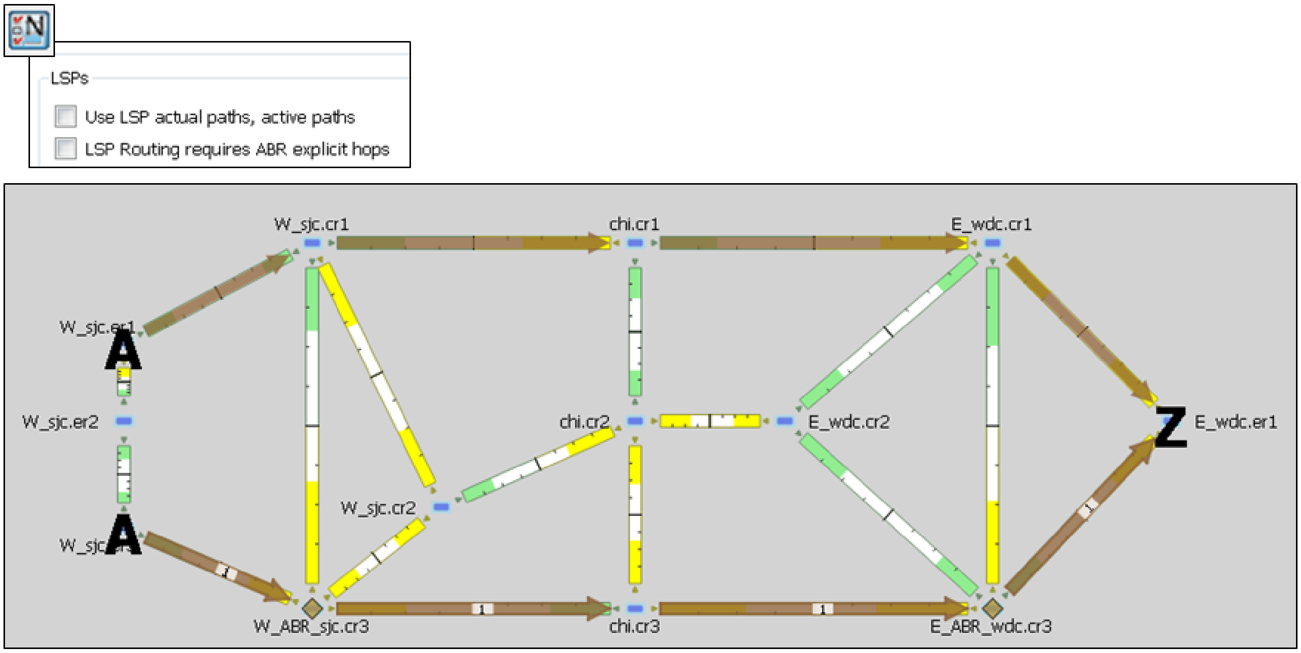

LSPs

□ Use LSP actual paths, active paths

□ LSP Routing requires ABR explicit hops

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

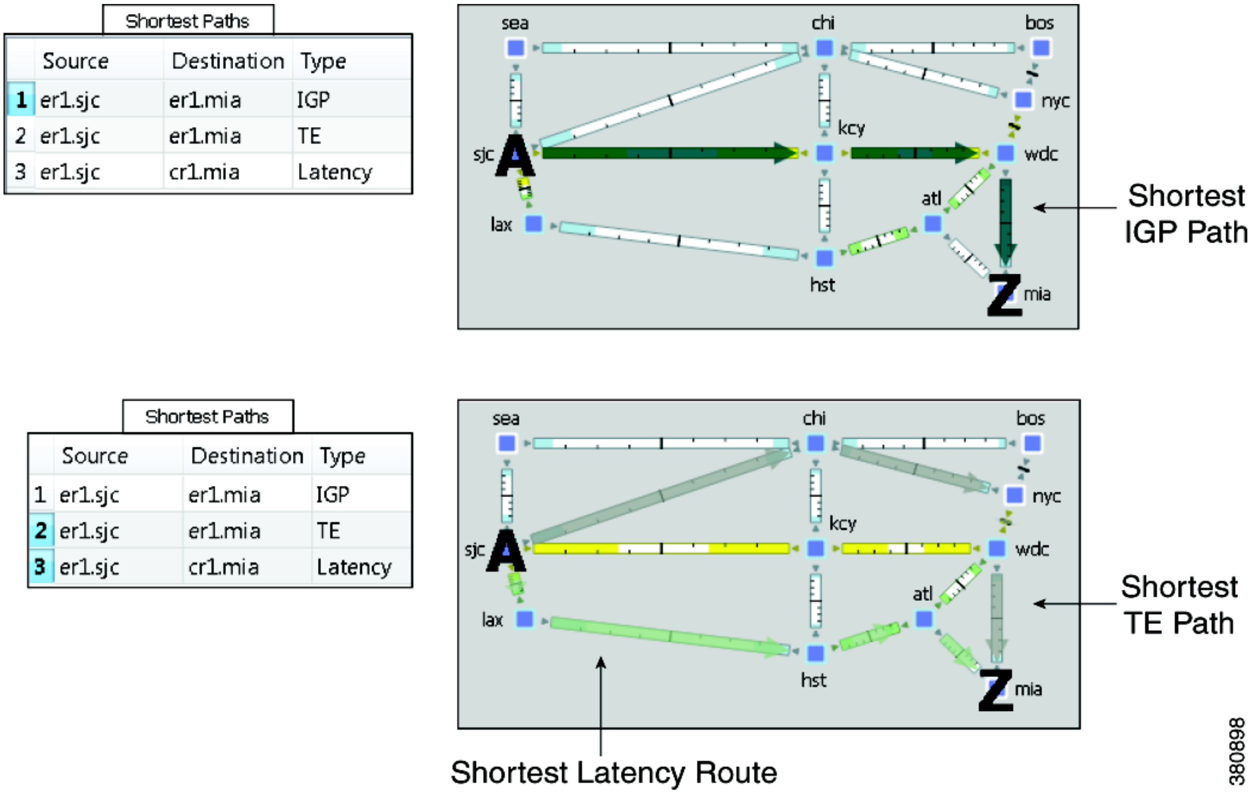

Shortest Paths

	Source	Destination	Type
1	er1.sjc	er1.mia	IGP
2	er1.sjc	er1.mia	TE
3	er1.sjc	cr1.mia	Latency

Shortest IGP Path

Shortest Paths

	Source	Destination	Type
1	er1.sjc	er1.mia	IGP
2	er1.sjc	er1.mia	TE
3	er1.sjc	cr1.mia	Latency

Shortest TE Path

Shortest Latency Route

Path Latency Calculations

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

- By default, MATE 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 MATE 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 MATE 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.

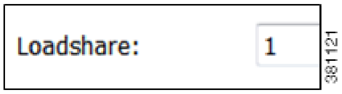
Loadshare: 1

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.

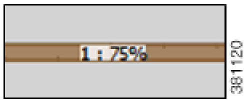

Note that this percentage may or may not be the same as the Loadshare value in the LSPs table. This is because the Loadshare value is relative to the other Loadshare values for parallel LSPs. For this reason, if manually editing the Loadshare property, it is recommended that you consider the entire set 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

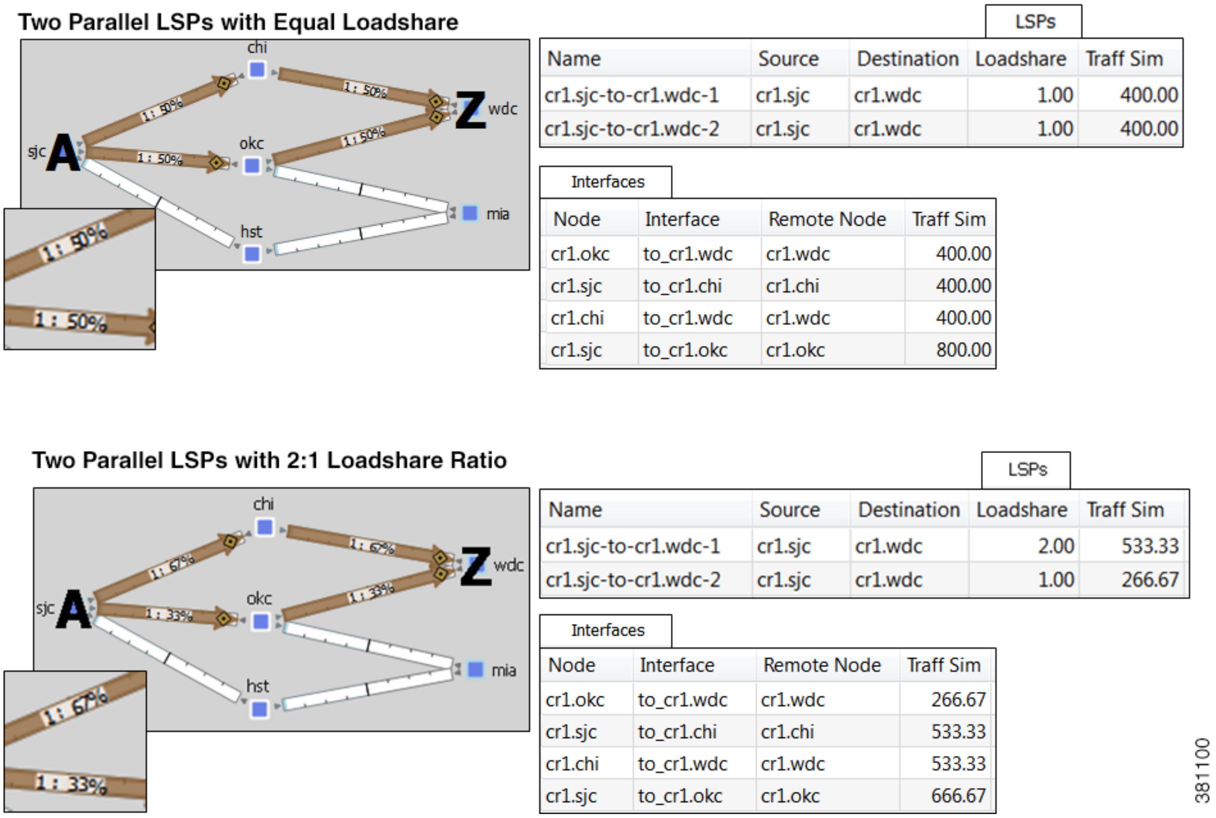

Two Parallel LSPs with Equal Loadshare

LSPs

Name	Source	Destination	Loadshare	Traff Sim
cr1.sjc-to-cr1.wdc-1	cr1.sjc	cr1.wdc	1.00	400.00
cr1.sjc-to-cr1.wdc-2	cr1.sjc	cr1.wdc	1.00	400.00

Interfaces

Node	Interface	Remote Node	Traff Sim
cr1.okc	to_cr1.wdc	cr1.wdc	400.00
cr1.sjc	to_cr1.chi	cr1.chi	400.00
cr1.chi	to_cr1.wdc	cr1.wdc	400.00
cr1.sjc	to_cr1.okc	cr1.okc	800.00

Two Parallel LSPs with 2:1 Loadshare Ratio

LSPs

Name	Source	Destination	Loadshare	Traff Sim
cr1.sjc-to-cr1.wdc-1	cr1.sjc	cr1.wdc	2.00	533.33
cr1.sjc-to-cr1.wdc-2	cr1.sjc	cr1.wdc	1.00	266.67

Interfaces

Node	Interface	Remote Node	Traff Sim
cr1.okc	to_cr1.wdc	cr1.wdc	266.67
cr1.sjc	to_cr1.chi	cr1.chi	533.33
cr1.chi	to_cr1.wdc	cr1.wdc	533.33
cr1.sjc	to_cr1.okc	cr1.okc	666.67

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

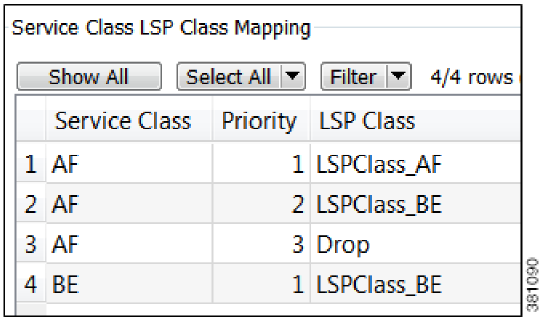

Service Class LSP Class Mapping

	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 foregoes 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.

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

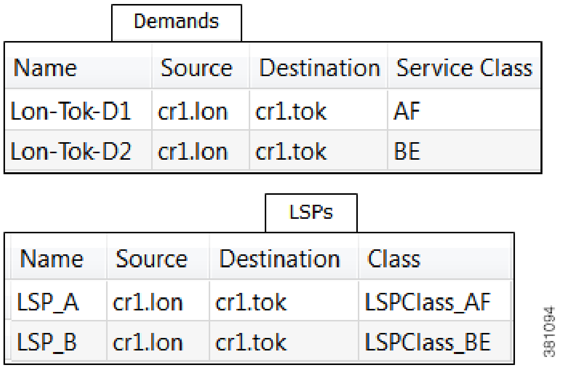

Demands

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

- 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

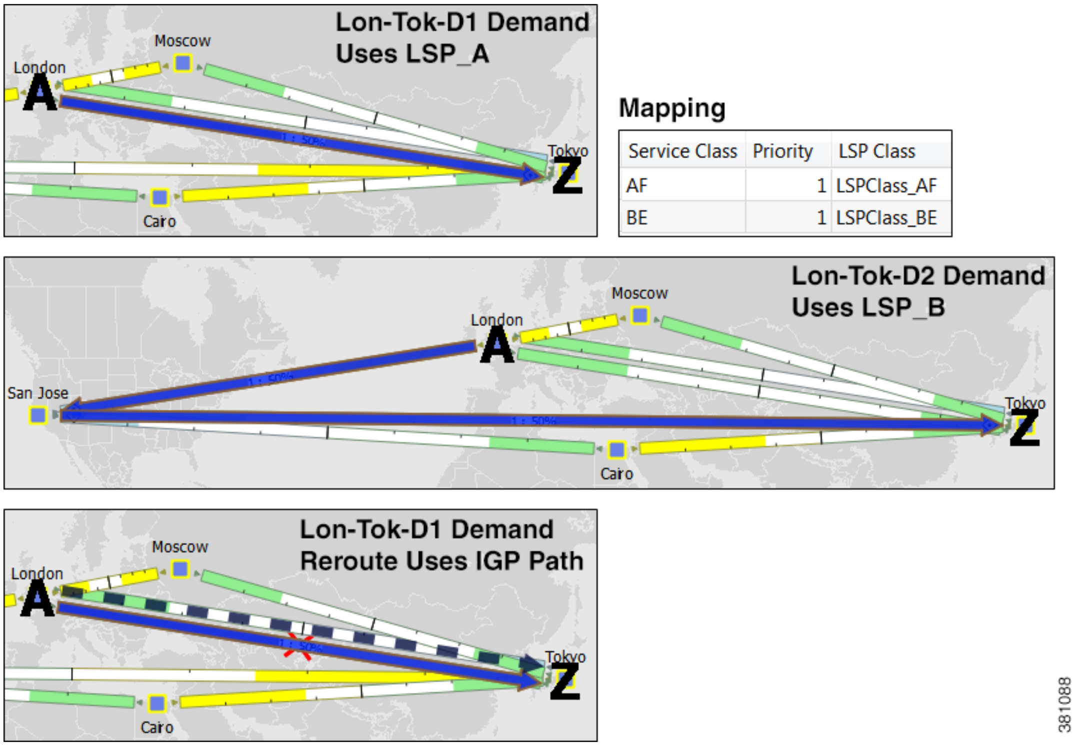

Mapping

Service Class	Priority	LSP Class
AF	1	LSPClass_AF
BE	1	LSPClass_BE

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

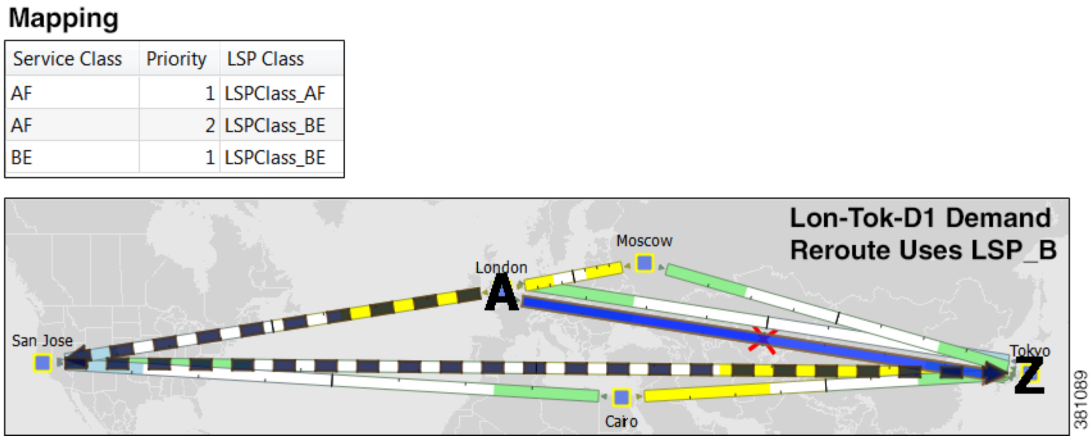

Mapping

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

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

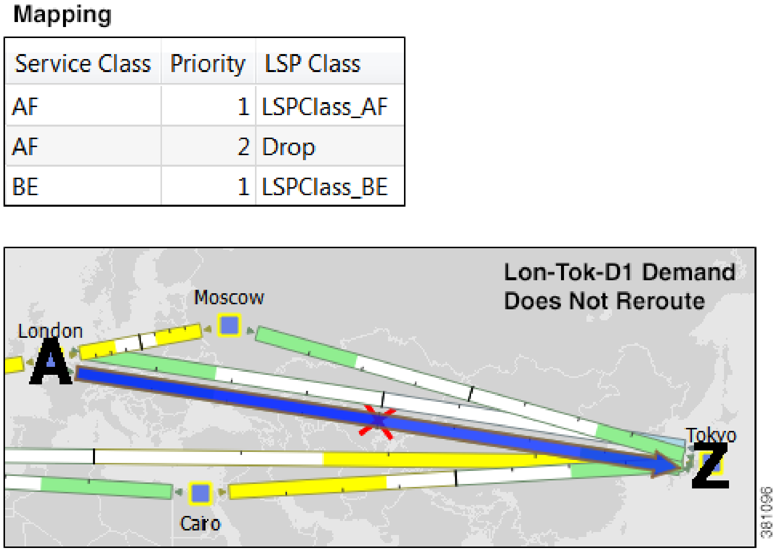

Mapping

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 on 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

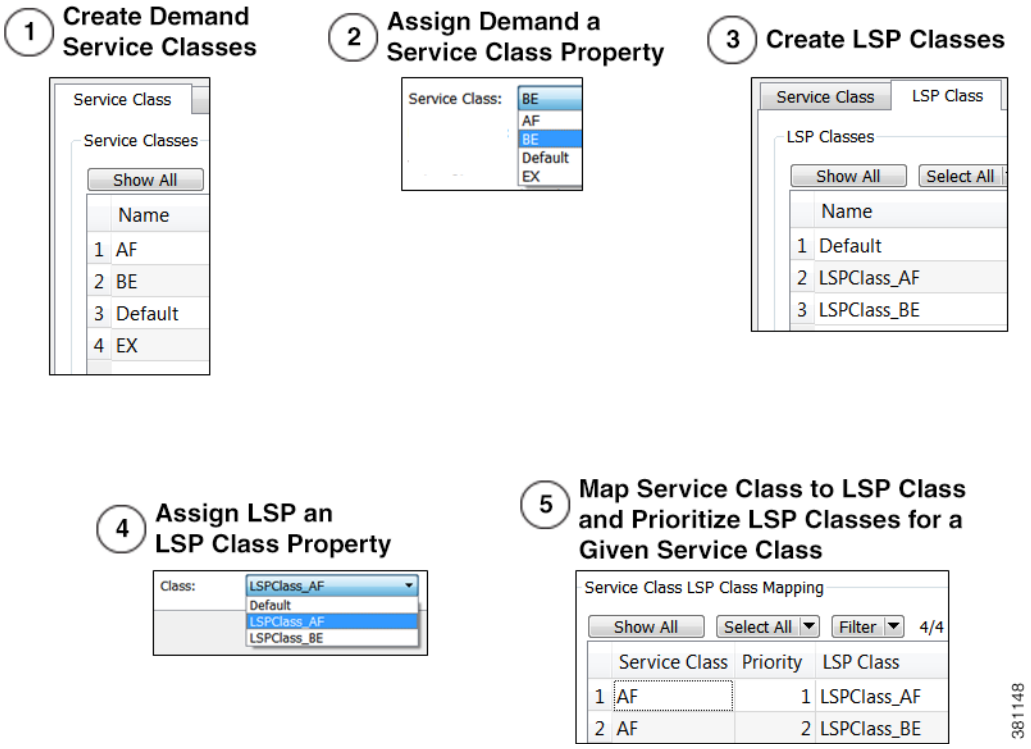

	Service Class	Priority	LSP Class
1	AF	1	LSPClass_AF
2	AF	2	LSPClass_BE

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.

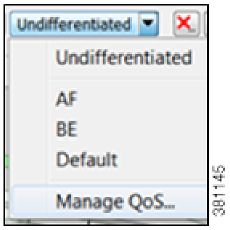

Step 2 Select the LSP Class tab.

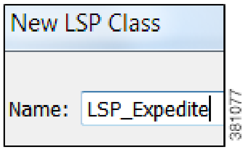

Name: LSP_Expedite

Step 3 In the LSP Class 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.

Assign LSP Classes to LSPs

Step 1 Select one or more LSPs.

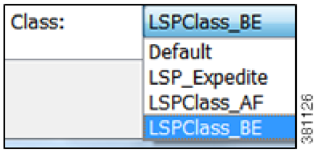
Class: LSPClass_BE

Step 2 Right-click on 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.

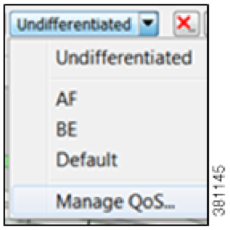

Step 2 Select the LSP Class tab.

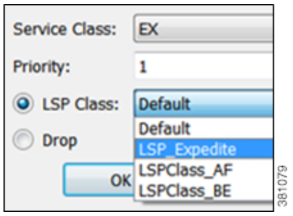
Service Class: EX
Priority: 1
● LSP Class: Default
○ Drop

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

MATE 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 MATE 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.

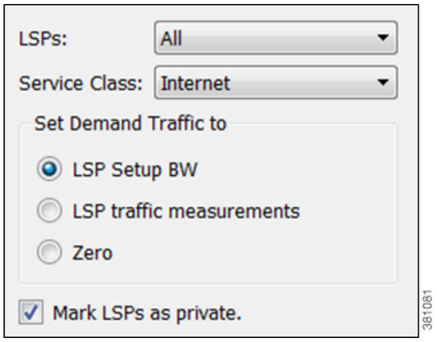

LSPs: All
Service Class: Internet
Set Demand Traffic to
● LSP Setup BW
○ LSP traffic measurements
○ Zero
☑ Mark LSPs as private.

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.

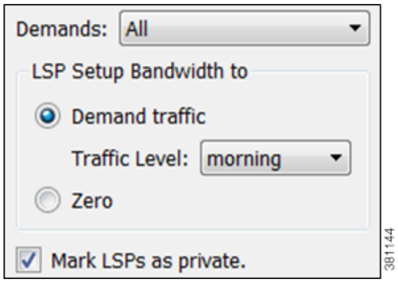

Demands: All

LSP Setup Bandwidth to

● Demand traffic

Traffic Level: morning

○ Zero

☑ Mark LSPs as private.

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 on 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 on 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.

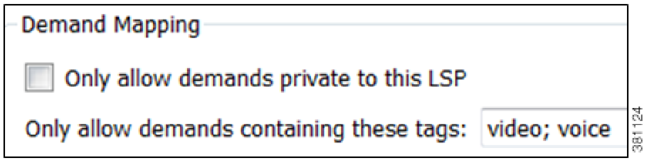

Demand Mapping

□ Only allow demands private to this LSP

Only allow demands containing these tags: video; voice

Simulation Modes

MATE 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, MATE 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

MATE Design simulates a Fast Reroute convergence mode using FRR LSPs. If a failure occurs in the route of a protected RSVP LSP, a pre-signaled 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 MATE 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.

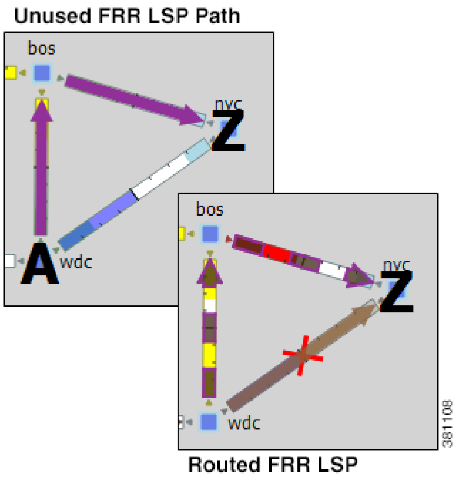

FRR LSP Routing

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.

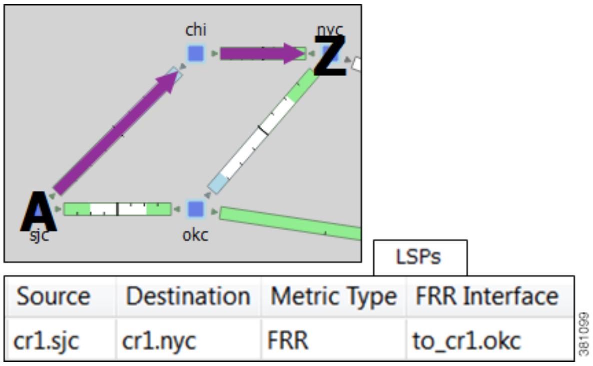

Source	Destination	Metric Type	FRR Interface
cr1.sjc	cr1.nyc	FRR	to_cr1.okc

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

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

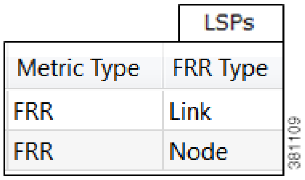

Metric Type	FRR Type
FRR	Link
FRR	Node

- Link-protection LSPs—Protect an LSP from a circuit failure by routing an FRR LSP around the failed circuit. The source node of the FRR LSP is the local node of the failed circuit, and if configured correctly, the destination of the FRR LSP is the remote node of the failed circuit. Figure 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 9-17 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

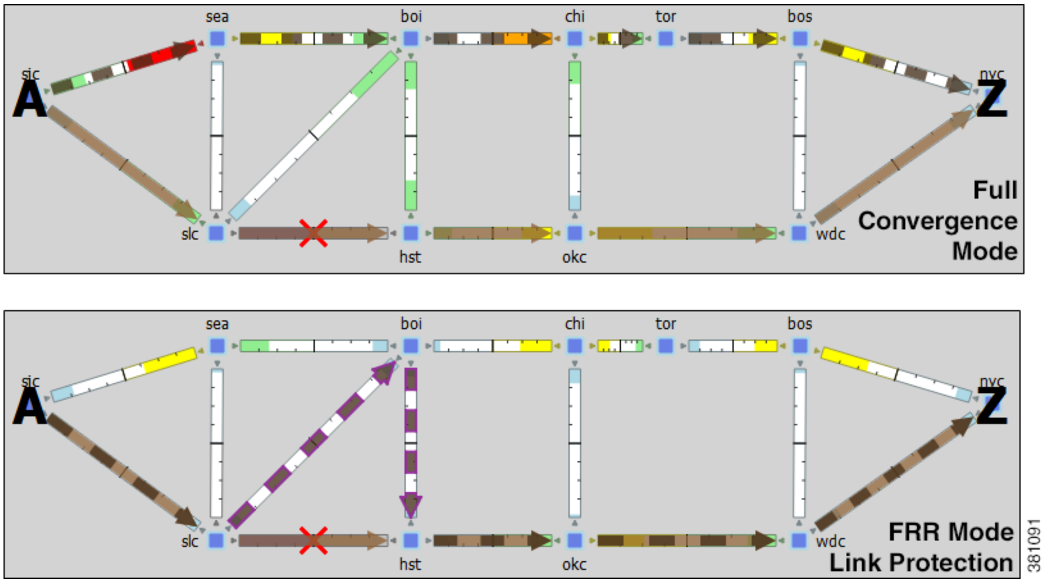

Figure 9-17 Example Full Convergence vs. Fast Reroute with Node Protection

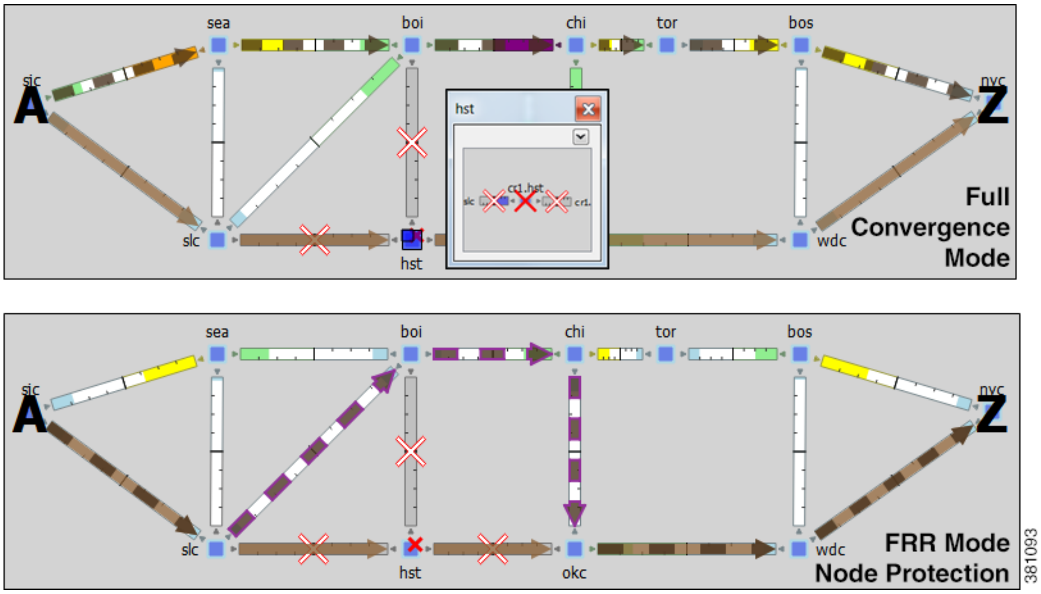

SRLG Protection

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 MATE, 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

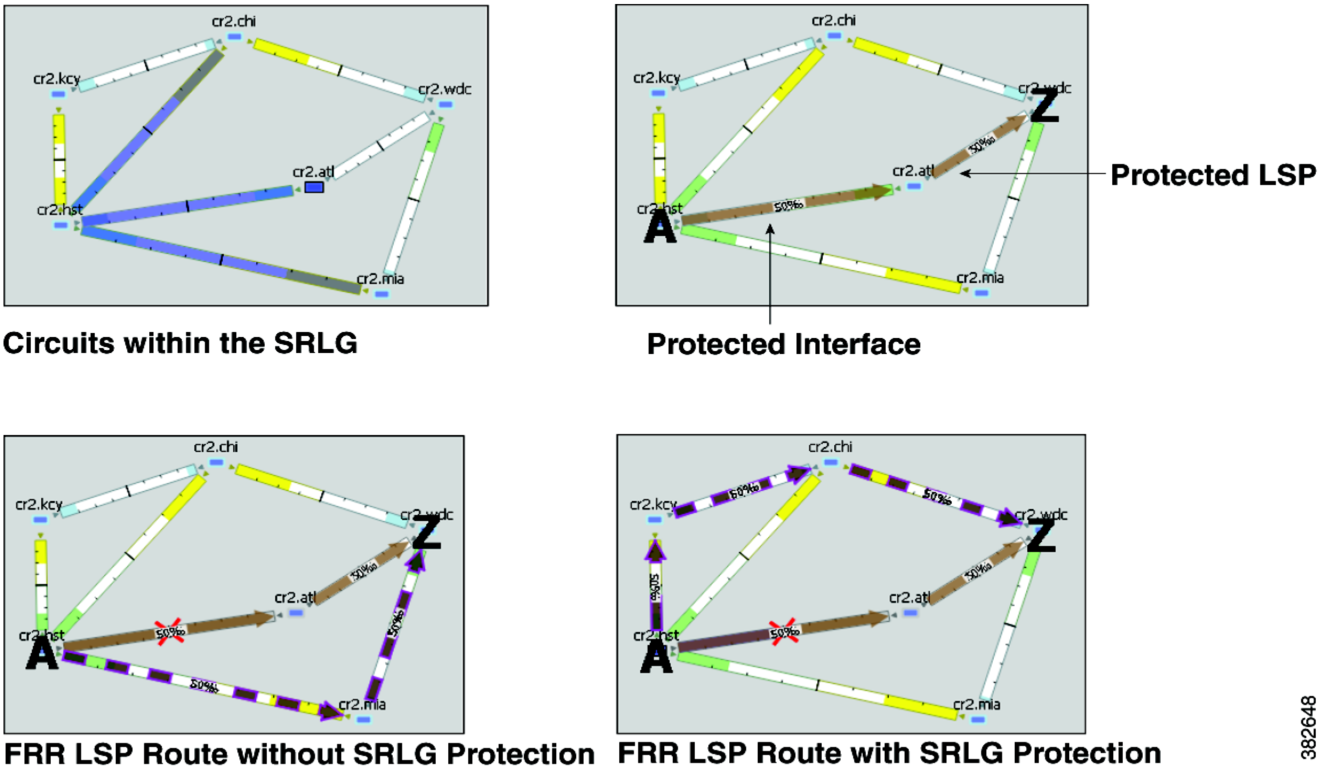

Circuits within the SRLG

Protected Interface

FRR LSP Route without SRLG Protection

FRR LSP Route 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, MATE 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, MATE 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

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

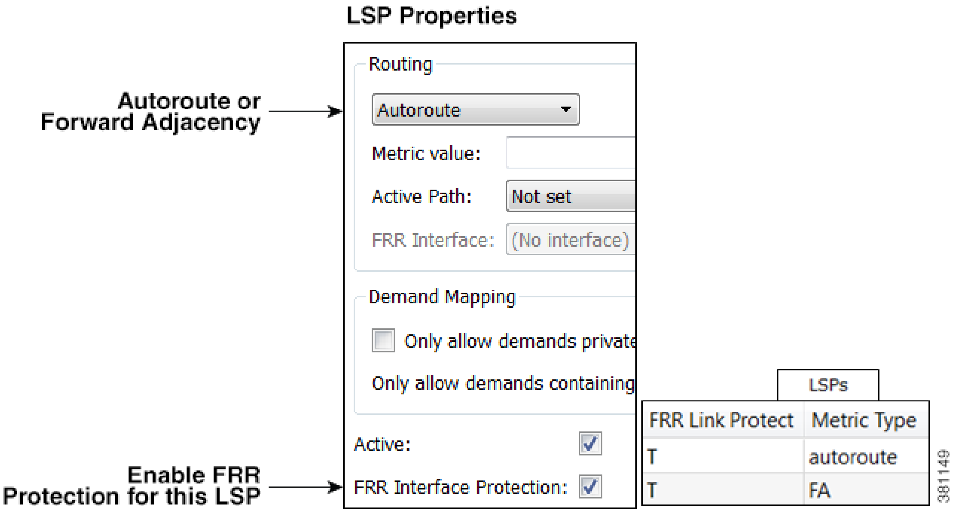

LSP Properties

Routing

Autoroute

Metric value:

Active Path: Not set

FRR Interface: (No interface)

Demand Mapping

☐ Only allow demands private

Only allow demands containing

Active: ☑

FRR Interface Protection: ☑

LSPs

FRR Link Protect	Metric Type
T	autoroute
T	FA

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.

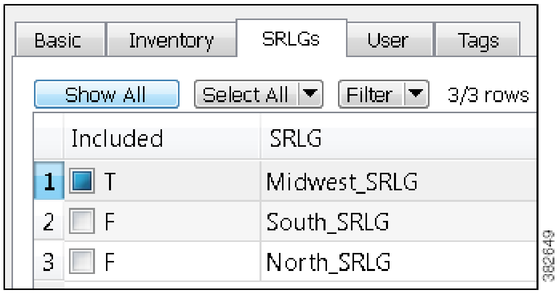

	Included	SRLG
1	■ T	Midwest_SRLG
2	□ F	South_SRLG
3	□ F	North_SRLG

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.

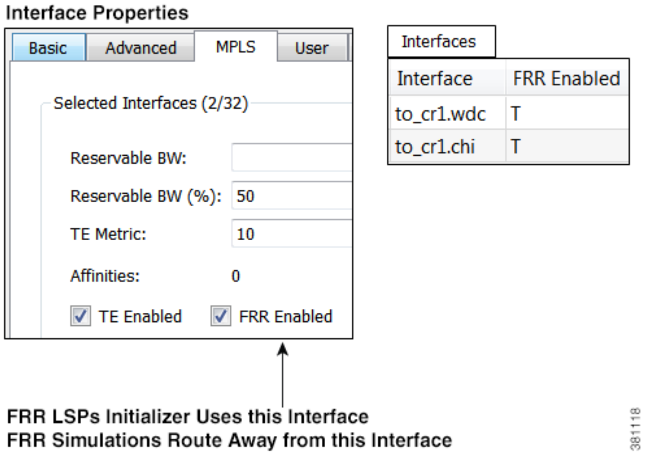

Selected Interfaces (2/32)

Reservable BW:

Reservable BW (%): 50

TE Metric: 10

Affinities: 0

☑ TE Enabled ☑ FRR Enabled

Interface	FRR Enabled
to_cr1.wdc	T
to_cr1.chi	T

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—MATE 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 9-20).
- Node protection—MATE 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 9-21).

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

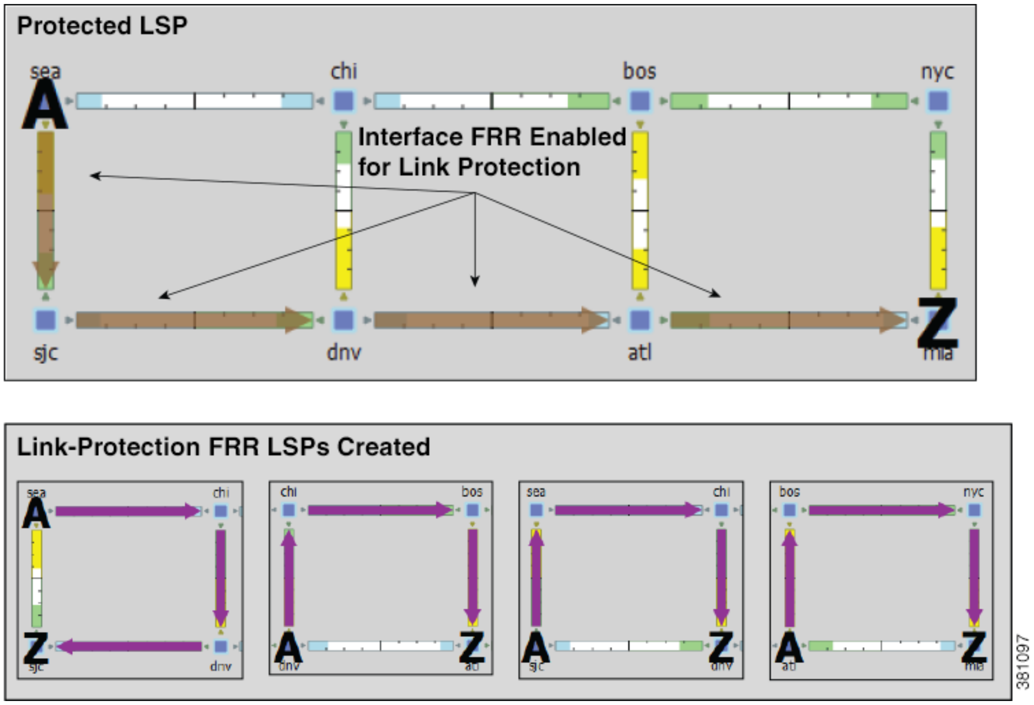

Figure 9-21 Example Node-Protection FRR LSPs Automatically Created

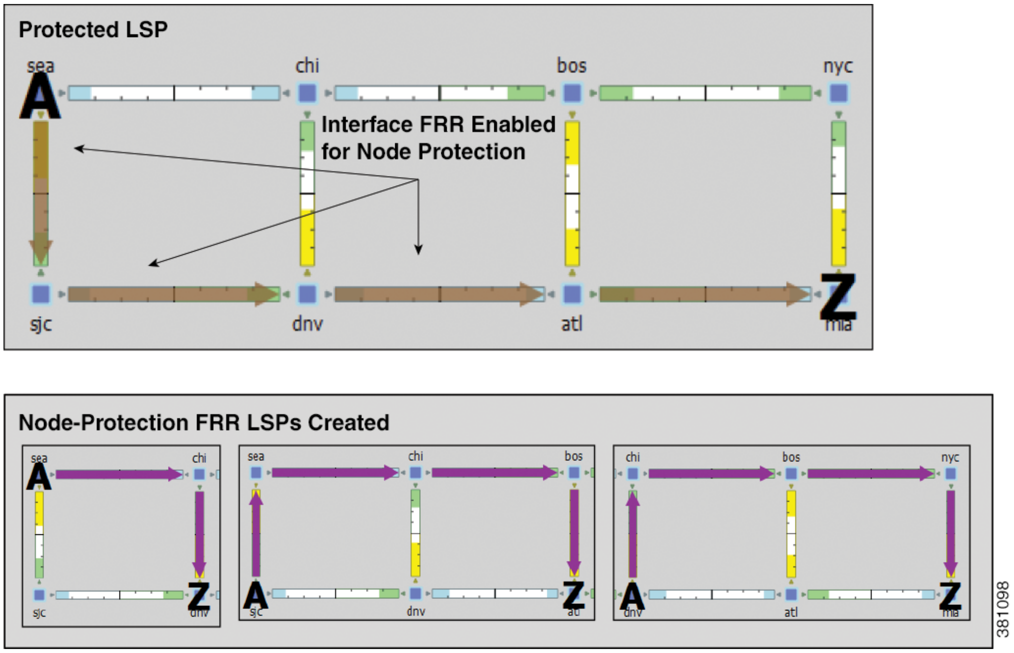

Run FRR LSPs Initializer

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

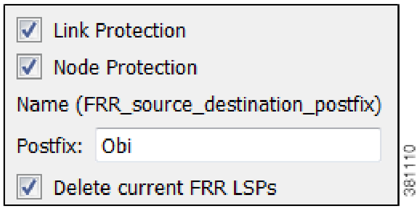

☑ Link Protection

☑ Node Protection

Name (FRR_source_destination_postfix)

Postfix: Obi

☑ Delete current FRR LSPs

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 on the LSP and select Properties.

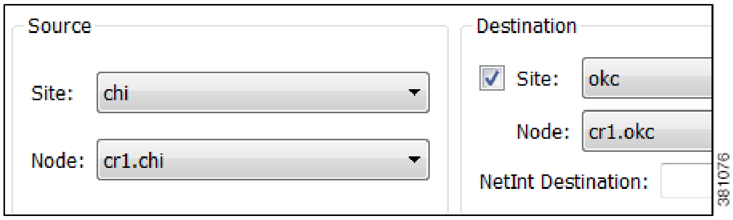

Source

Site: chi

Node: cr1.chi

Destination

☑ Site: okc

Node: cr1.okc

NetInt Destination:

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.

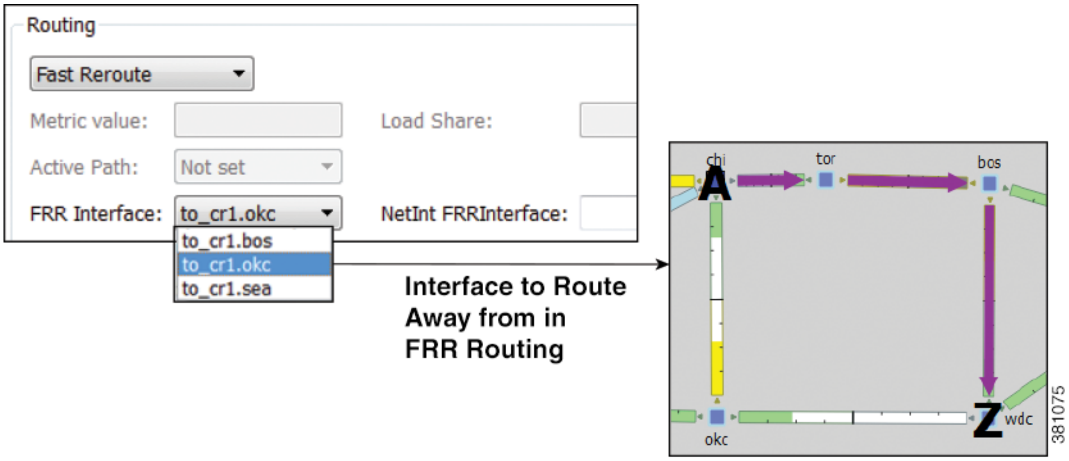

Routing

Fast Reroute

Metric value:

Load Share:

Active Path: Not set

FRR Interface: to_cr1.okc

NetInt FRRInterface:

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.

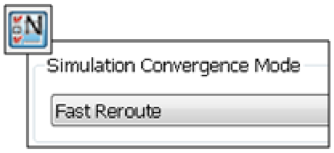

Simulation Convergence Mode

Fast Reroute

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

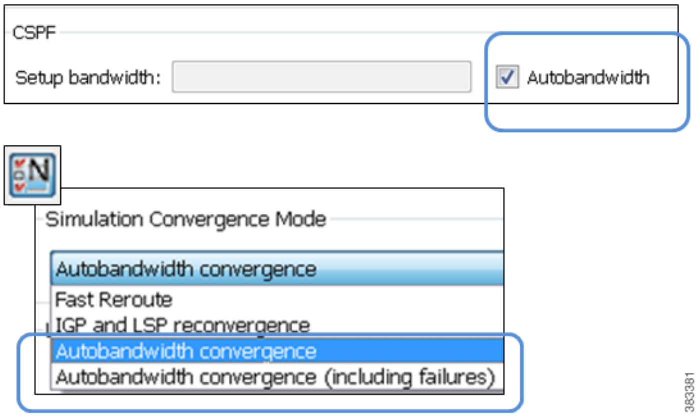

CSPF

Setup bandwidth: ☑ Autobandwidth

Simulation Convergence Mode

Autobandwidth convergence

- Fast Reroute
- IGP and LSP reconvergence
- Autobandwidth convergence
- Autobandwidth convergence (including failures)

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.

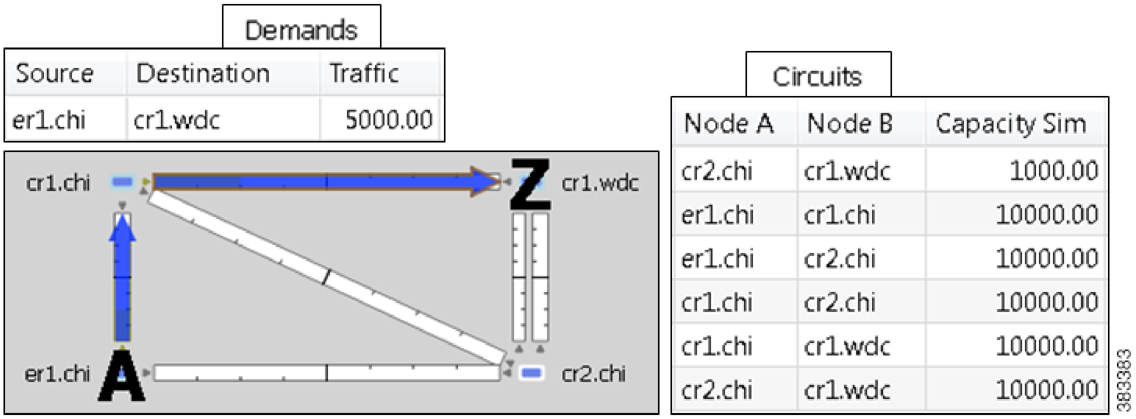

Demands

Source	Destination	Traffic
er1.chi	cr1.wdc	5000.00

Circuits

Node A	Node B	Capacity Sim
cr2.chi	cr1.wdc	1000.00
er1.chi	cr1.chi	10000.00
er1.chi	cr2.chi	10000.00
cr1.chi	cr2.chi	10000.00
cr1.chi	cr1.wdc	10000.00
cr2.chi	cr1.wdc	10000.00

- 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 9-23 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 9-24 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

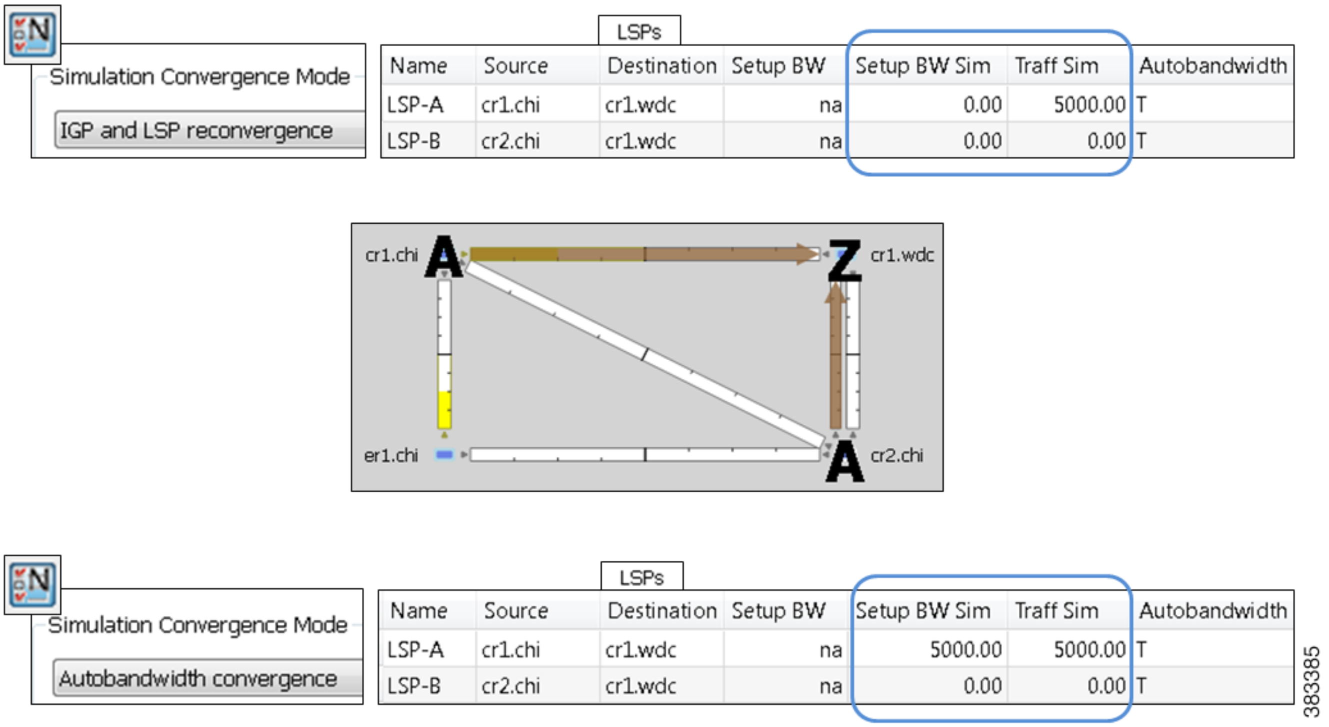

Simulation Convergence Mode: IGP and LSP reconvergence

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

Figure 9-23 Example Autobandwidth Convergence with Failures

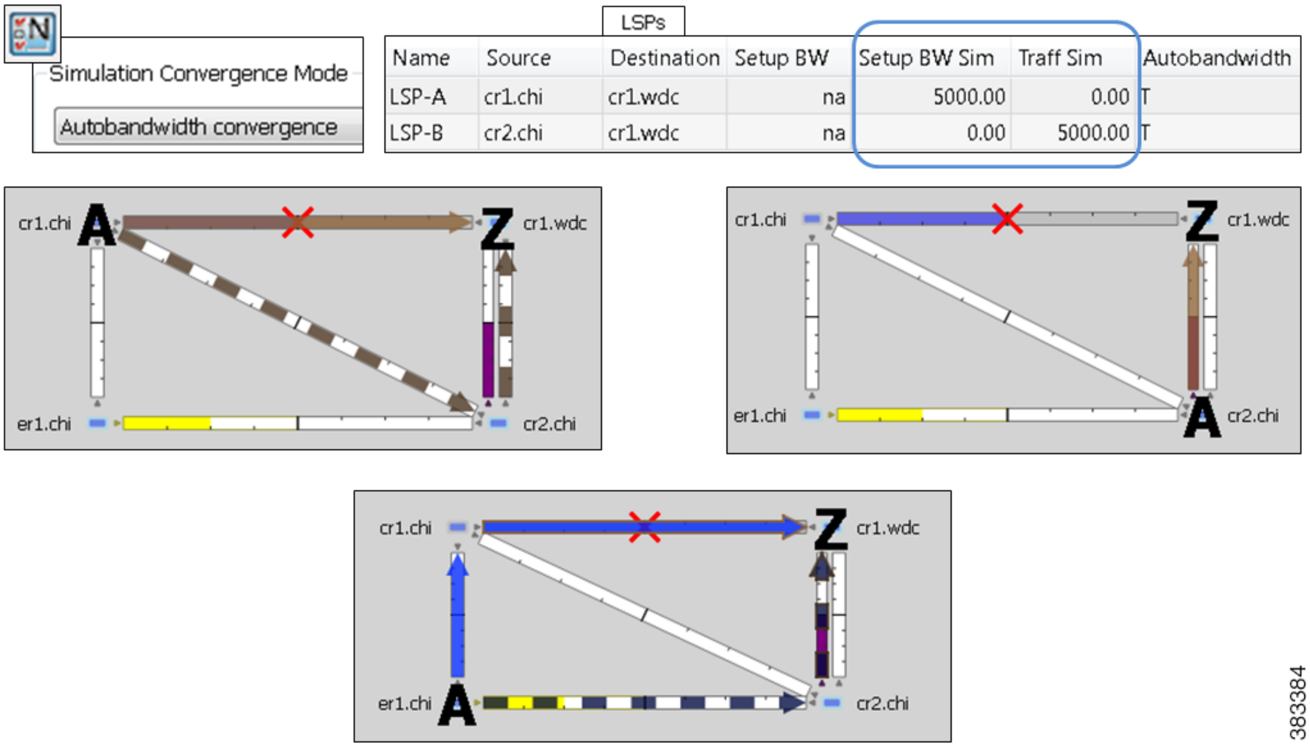

Simulation Convergence Mode

Autobandwidth convergence

LSPs

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

Figure 9-24 Example Autobandwidth Convergence Including Failures

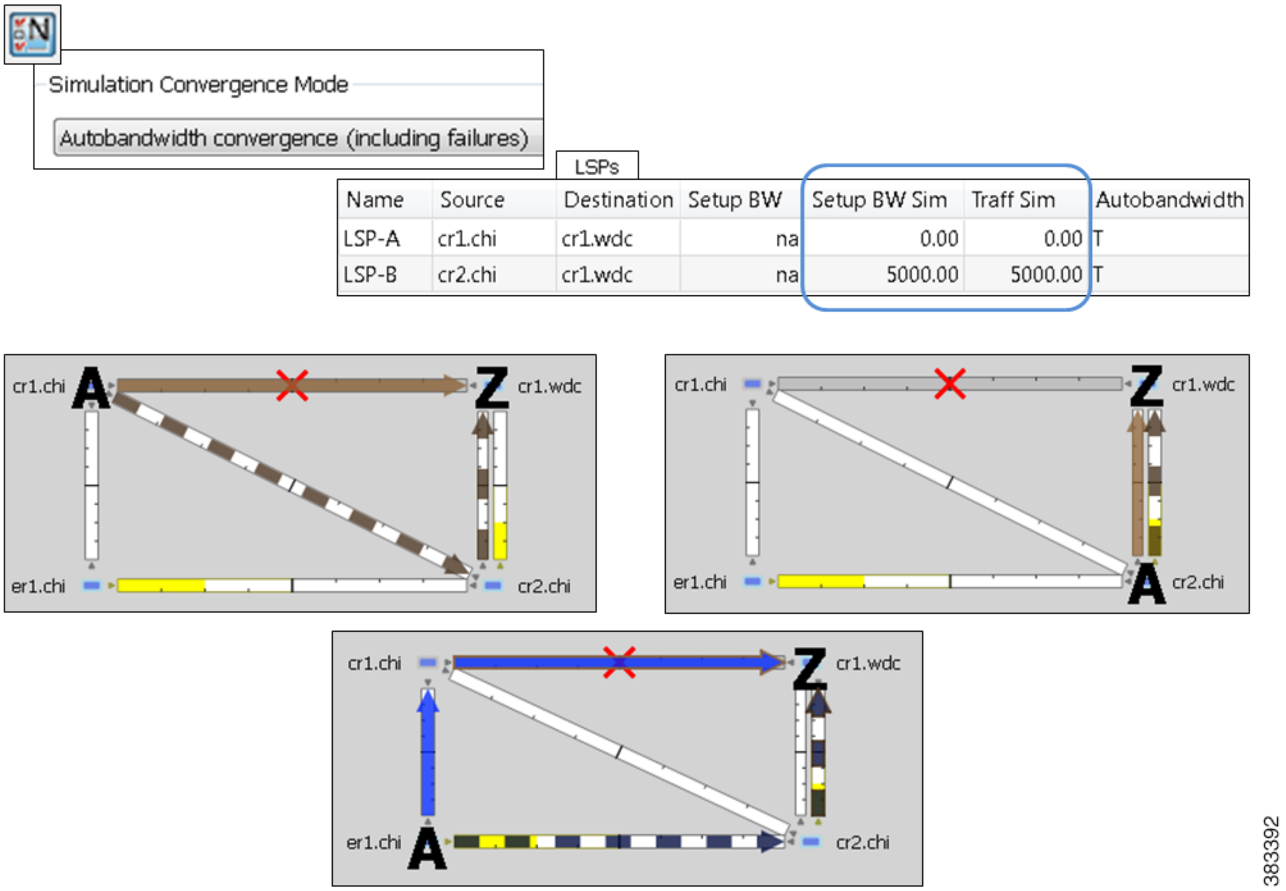

Simulation Convergence Mode

Autobandwidth convergence (including failures)

LSPs

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

Advanced RSVP LSP Simulations

Ignore Reservable Bandwidths for Capacity Planning

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

When you enable this routing mode, MATE 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 *MATE 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 9-25 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

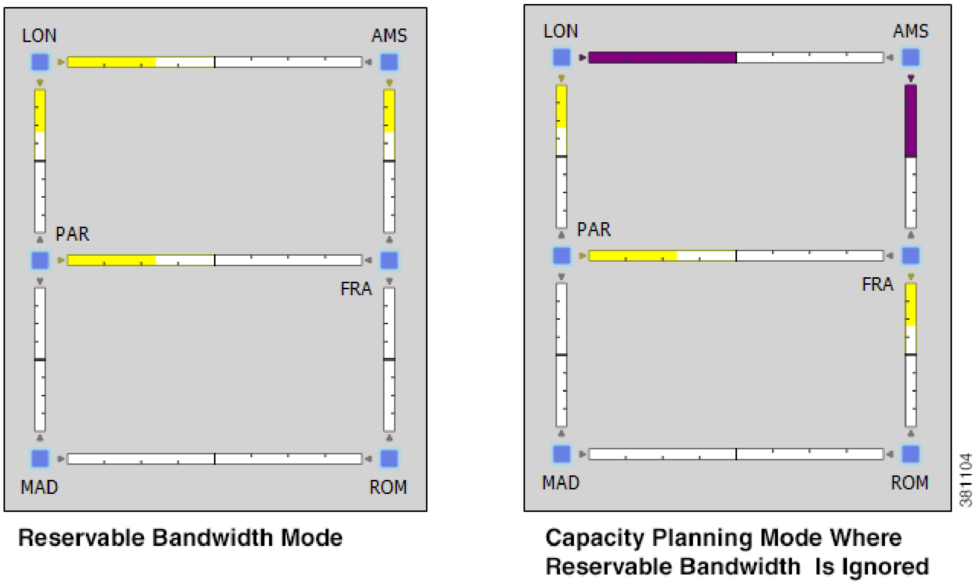
LON
AMS
PAR
FRA
MAD
ROM
LON
AMS
PAR
FRA
MAD
ROM
381104
Reservable Bandwidth Mode
Capacity Planning Mode Where Reservable Bandwidth Is Ignored

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

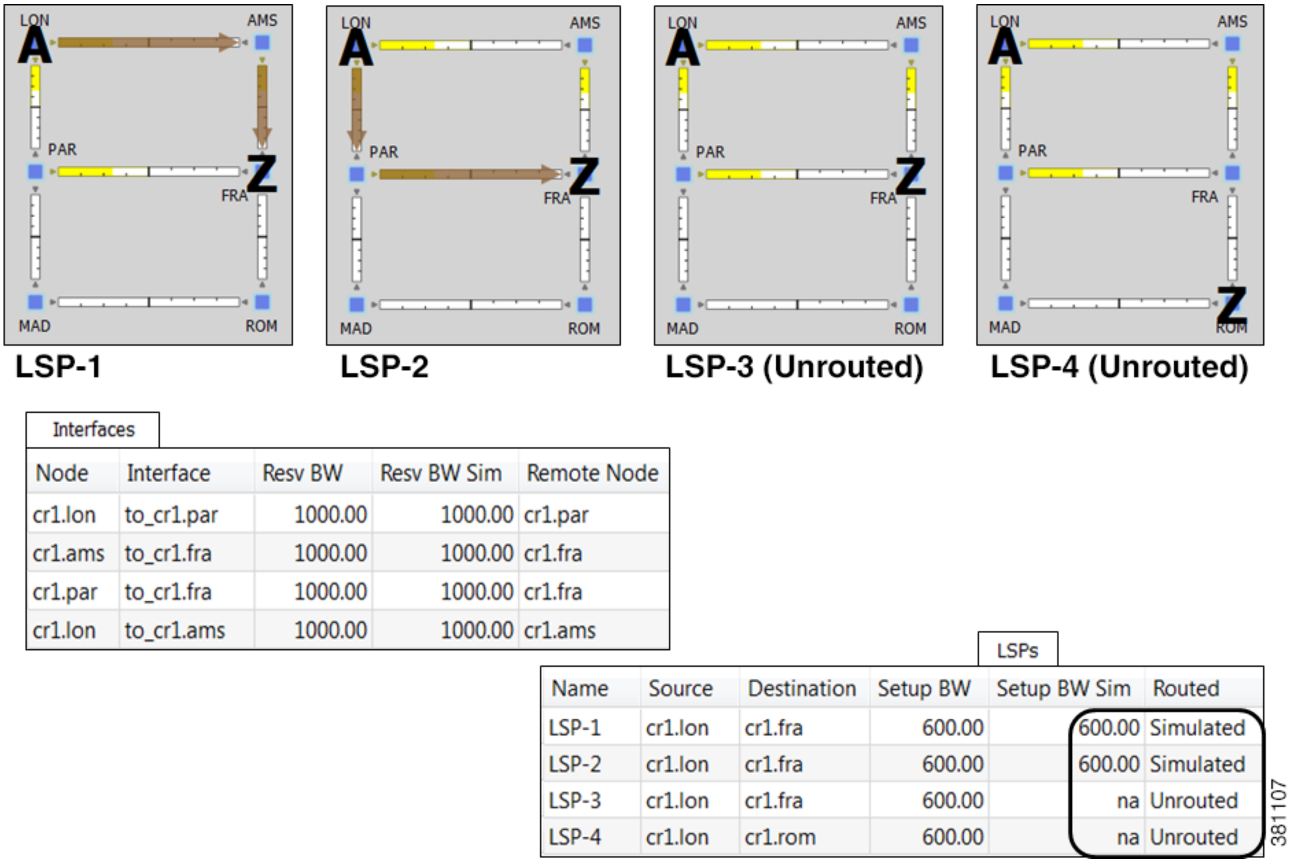

LSP-1 LSP-2 LSP-3 (Unrouted) LSP-4 (Unrouted)

Interfaces

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

LSPs

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

In Figure 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

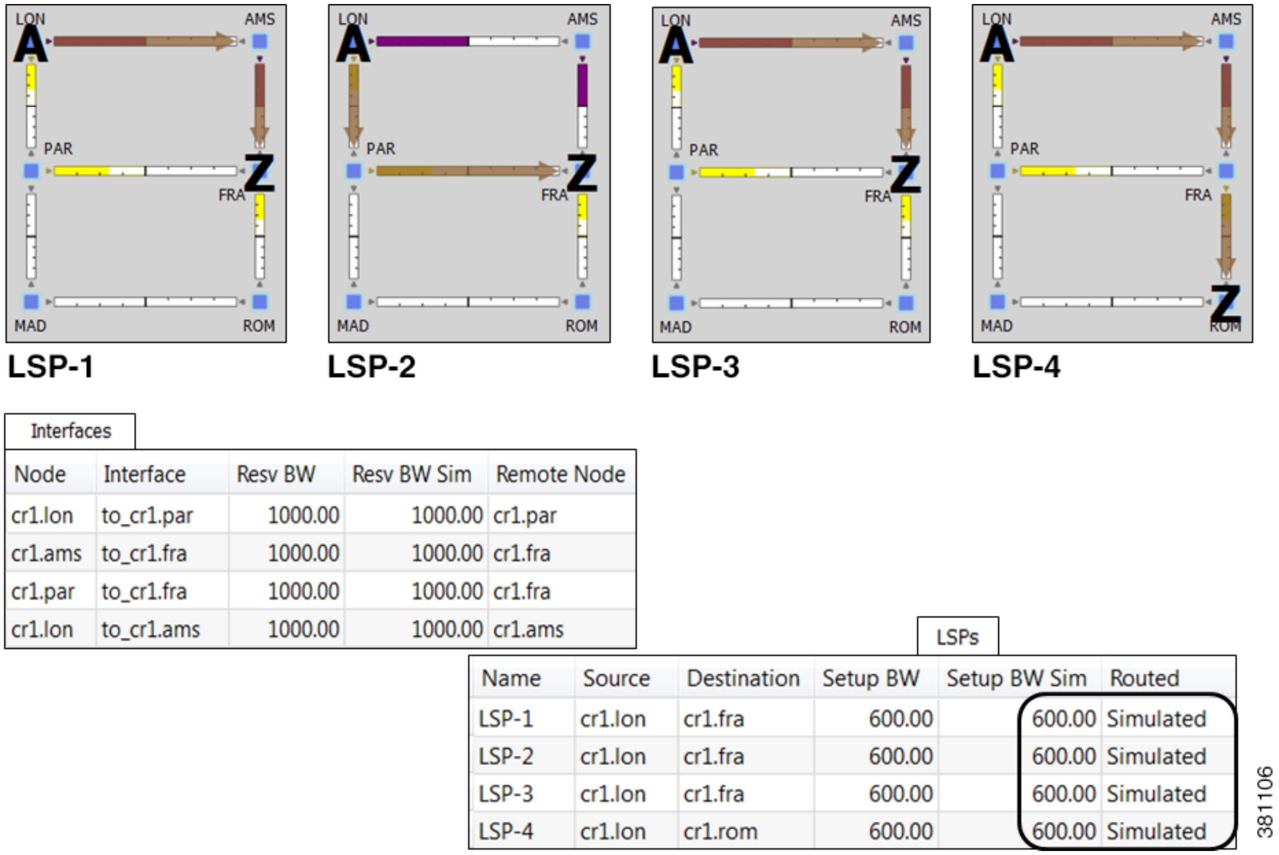

Interfaces

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

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.

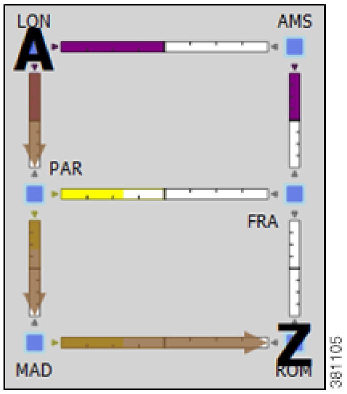

LSP-4

Enable Capacity Planning Mode for LSPs

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

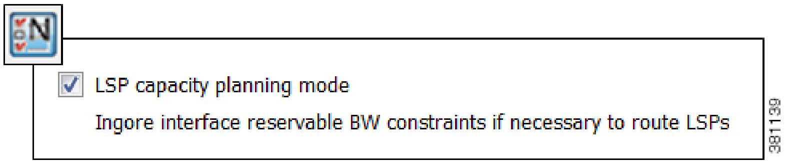

☑ LSP capacity planning mode

Ingore interface reservable BW constraints if necessary to route LSPs

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 MATE 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 9-28 P2MP LSP with chi as the Source for Three Destinations

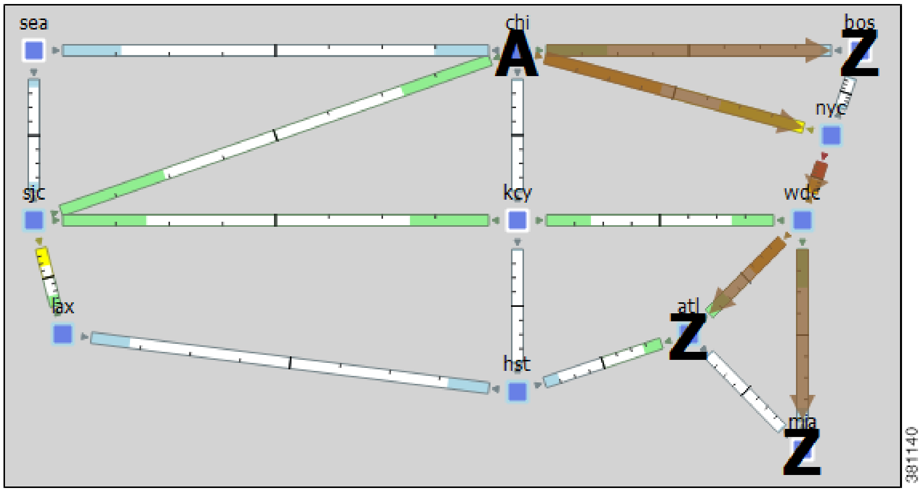

P2MP LSP Bandwidth

All sub-LSPs within a given P2MP LSP share bandwidth and therefore, common circuits between these sub-LSPs have the same bandwidth, rather than an aggregate. Common practice is to set all sub-LSPs within a P2MP LSP to the same bandwidth. If you do not, the bandwidth for the shared circuit will be the highest bandwidth set.

To see the reserved LSP bandwidth, show the LSP Resv column in the Interfaces table. You can also view the associated traffic by selecting LSP Reservations in the Network Plot pull-down menu located in the Visualization toolbar. For more information, see the View LSP Reservations section.

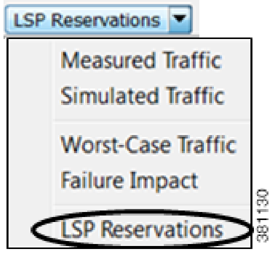

Example: If A-C LSP and A-D LSP each have a bandwidth 400 Mbps, the A-B circuit has an aggregate bandwidth of 800 Mbps (Figure 9-29). However, if A-C and A-D are sub-LSPs within a P2MP LSP, the A-B circuit has a bandwidth of 400 Mbps (Figure 9-30).

Figure 9-29 Bandwidth Behavior without a P2MP LSP

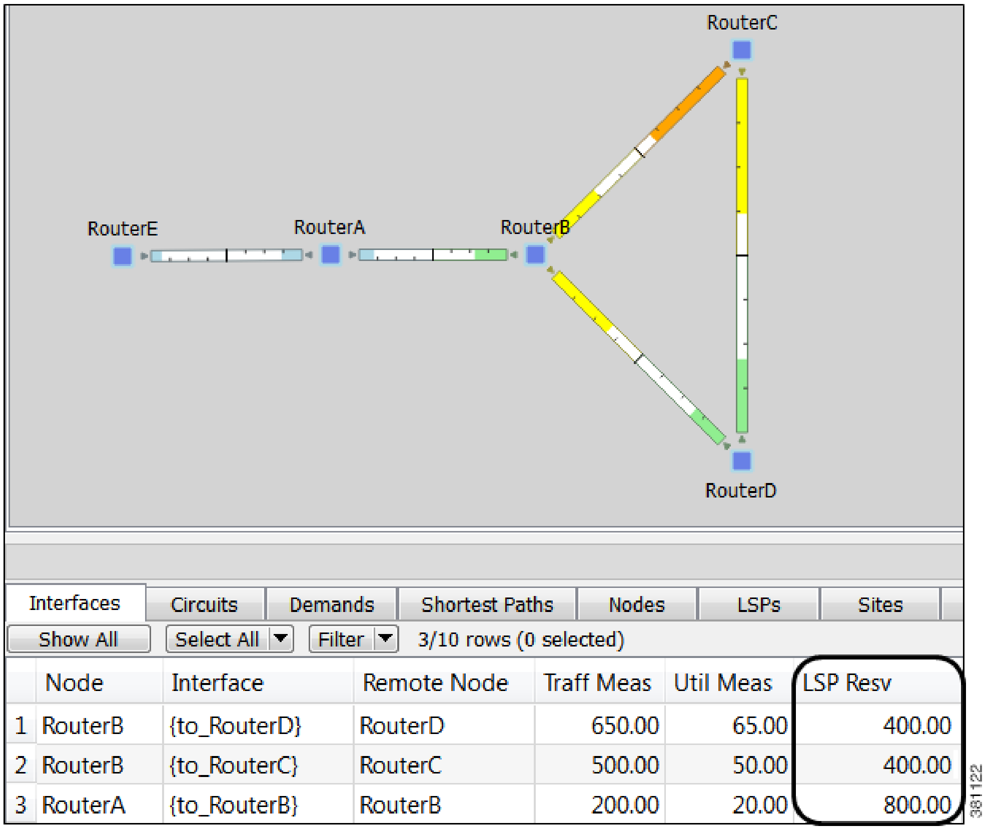

	Node	Interface	Remote Node	Traff Meas	Util Meas	LSP Resv
1	RouterB	{to_RouterD}	RouterD	650.00	65.00	400.00
2	RouterB	{to_RouterC}	RouterC	500.00	50.00	400.00
3	RouterA	{to_RouterB}	RouterB	200.00	20.00	800.00

Figure 9-30 Bandwidth Behavior with a P2MP LSP

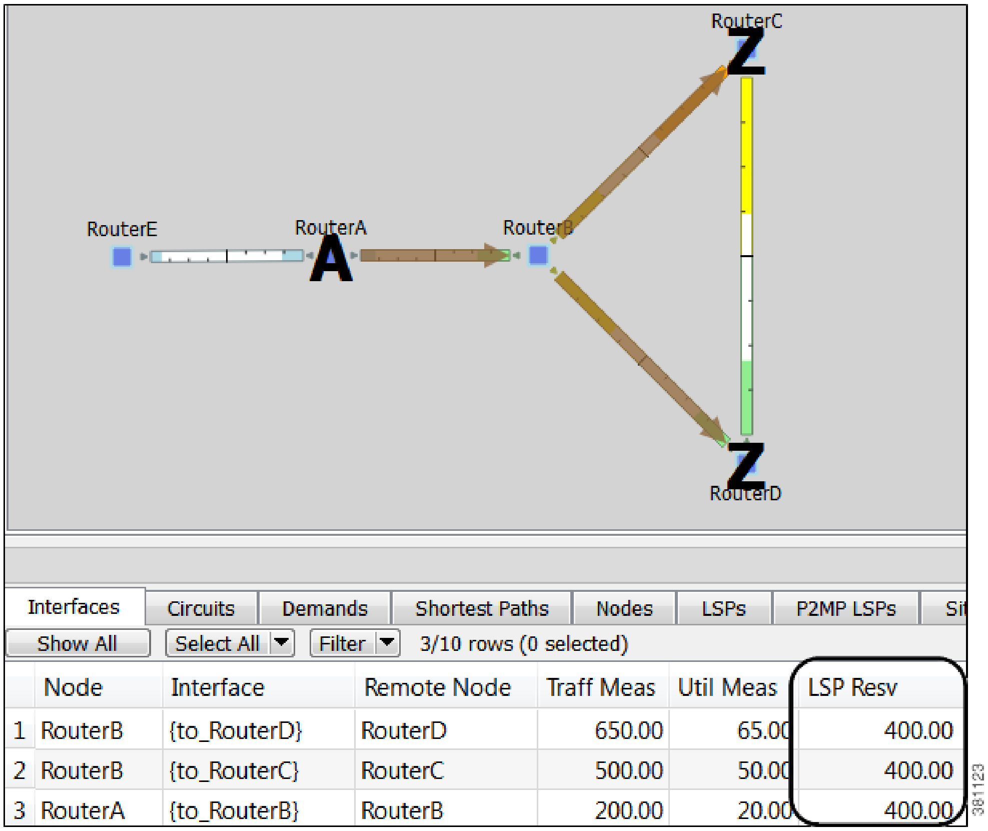

	Node	Interface	Remote Node	Traff Meas	Util Meas	LSP Resv
1	RouterB	{to_RouterD}	RouterD	650.00	65.00	400.00
2	RouterB	{to_RouterC}	RouterC	500.00	50.00	400.00
3	RouterA	{to_RouterB}	RouterB	200.00	20.00	400.00

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 PSMP 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

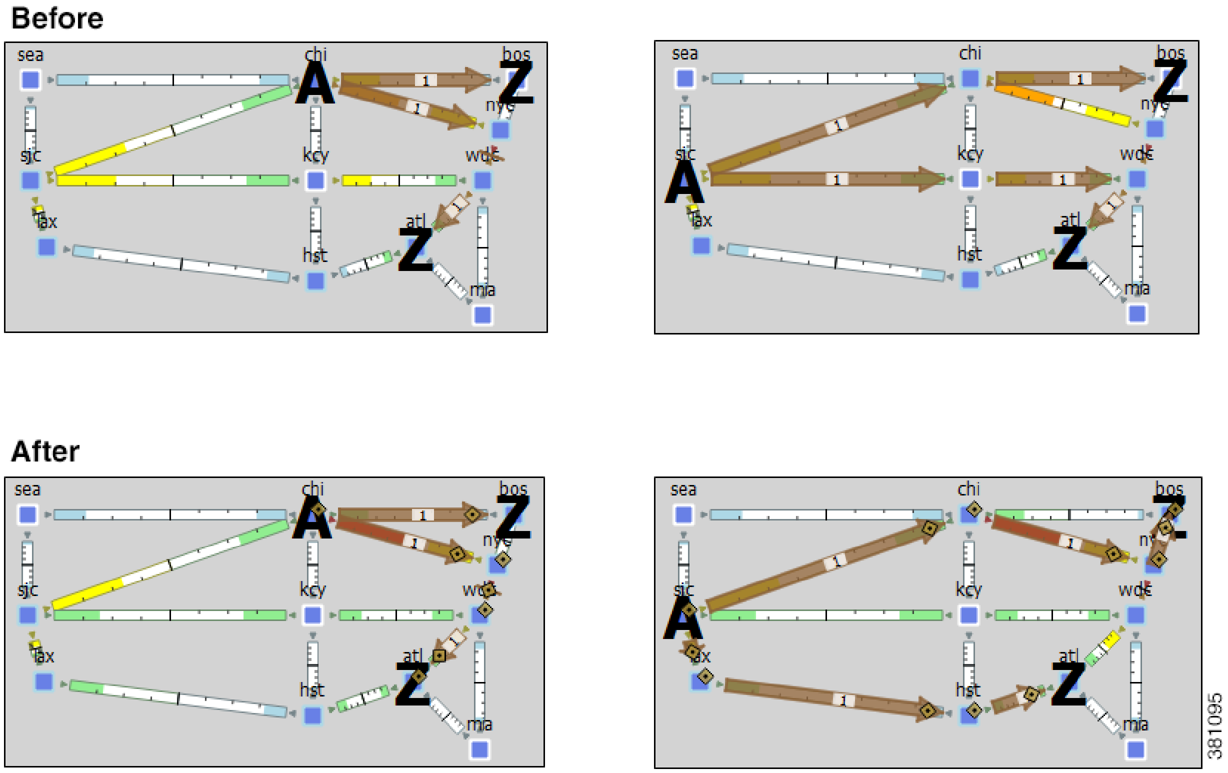

MATE Design creates named paths and named path hops for the sub-LSP paths. After running the initializer, you can verify the named paths and their hops in the Named Paths and Named Path Hops table, respectively.

- 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>.

Pre-requisites

- The plan file must already contain the LSP paths. See the Create LSP Paths section.

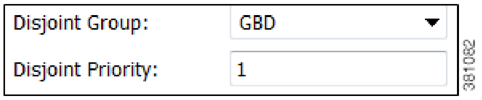
Disjoint Group: GBD
Disjoint Priority: 1

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

Step 1 From the P2MP LSPs table, 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 he named paths.

Step 2 Select the Initializers->P2MP LSP Explicit Paths menu, or right-click on 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.

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

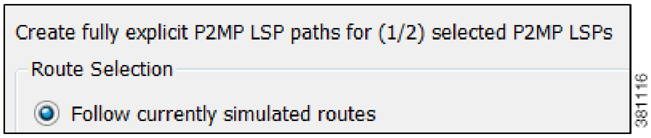

Create fully explicit P2MP LSP paths for (1/2) selected P2MP LSPs

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.

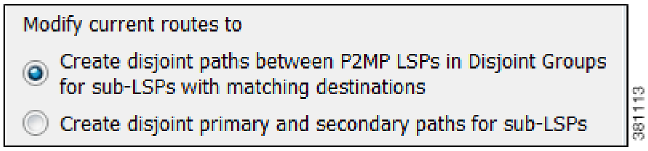

Modify current routes to

● 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, MATE 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.

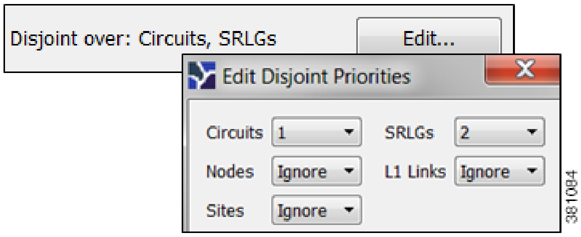

Disjoint over: Circuits, SRLGs

Edit Disjoint Priorities

Circuits 1 SRLGs 2

Nodes Ignore L1 Links Ignore

Sites Ignore

Step 4 Click OK.

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.

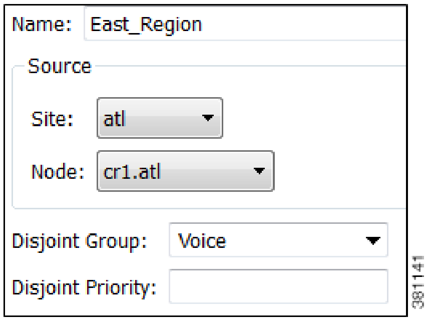

Name: East_Region

Source

Site: atl

Node: cr1.atl

Disjoint Group: Voice

Disjoint Priority:

- **a.** 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.
- **b.** Enter a unique P2MP LSP name. If you do not enter a name, MATE Design creates one using the format p2mp_lsp_<source_node>.
- **c.** Select both the source site and node. This selection is the root for all sub-LSPs contained within this P2MP LSP.
- **d.** Optional: Enter or select a disjoint group and disjoint priority for MATE 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.

- **a.** Double-click on the sub-LSPs from the LSP table. The LSP Properties dialog box appears.
- **b.** 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 on 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.
- 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.

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.

MATE 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, MATE creates and routes RSVP LSPs. However, you can create Segment Routing (SR) LSPs 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 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 (MATE 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 and/or interface segments.

- 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.
- 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 MATE, 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 9-33 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.

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 on an SR LSP and select Filter to Segment Lists. You can then right-click on the segment list and select Filter to Segment List Hops, or double-click the segment list.

Figure 9-32 Example Node Segment

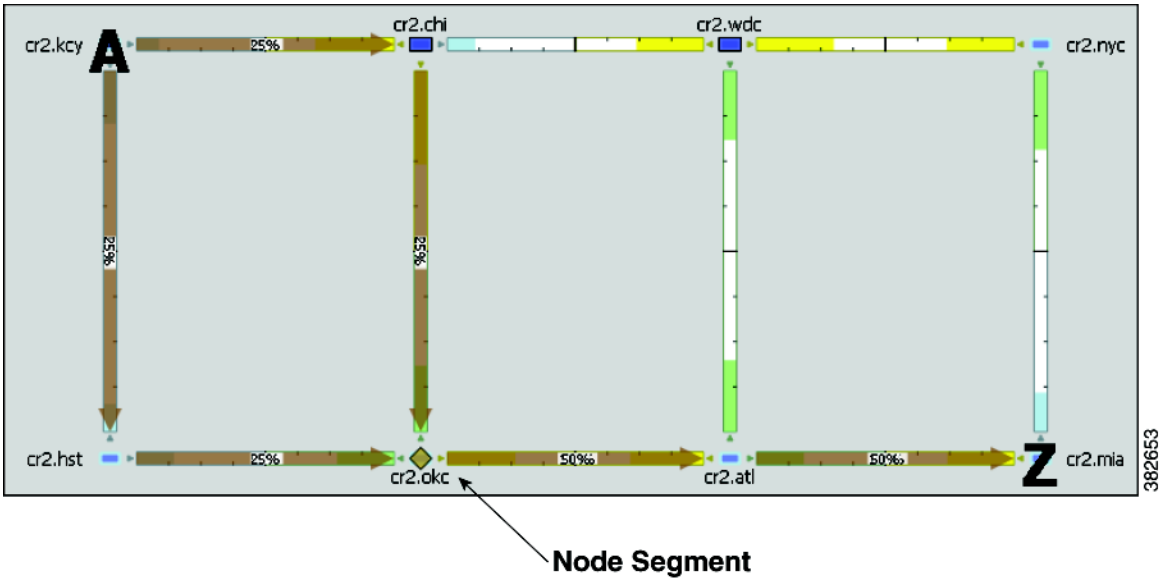

Figure 9-33 Example Interface Segment

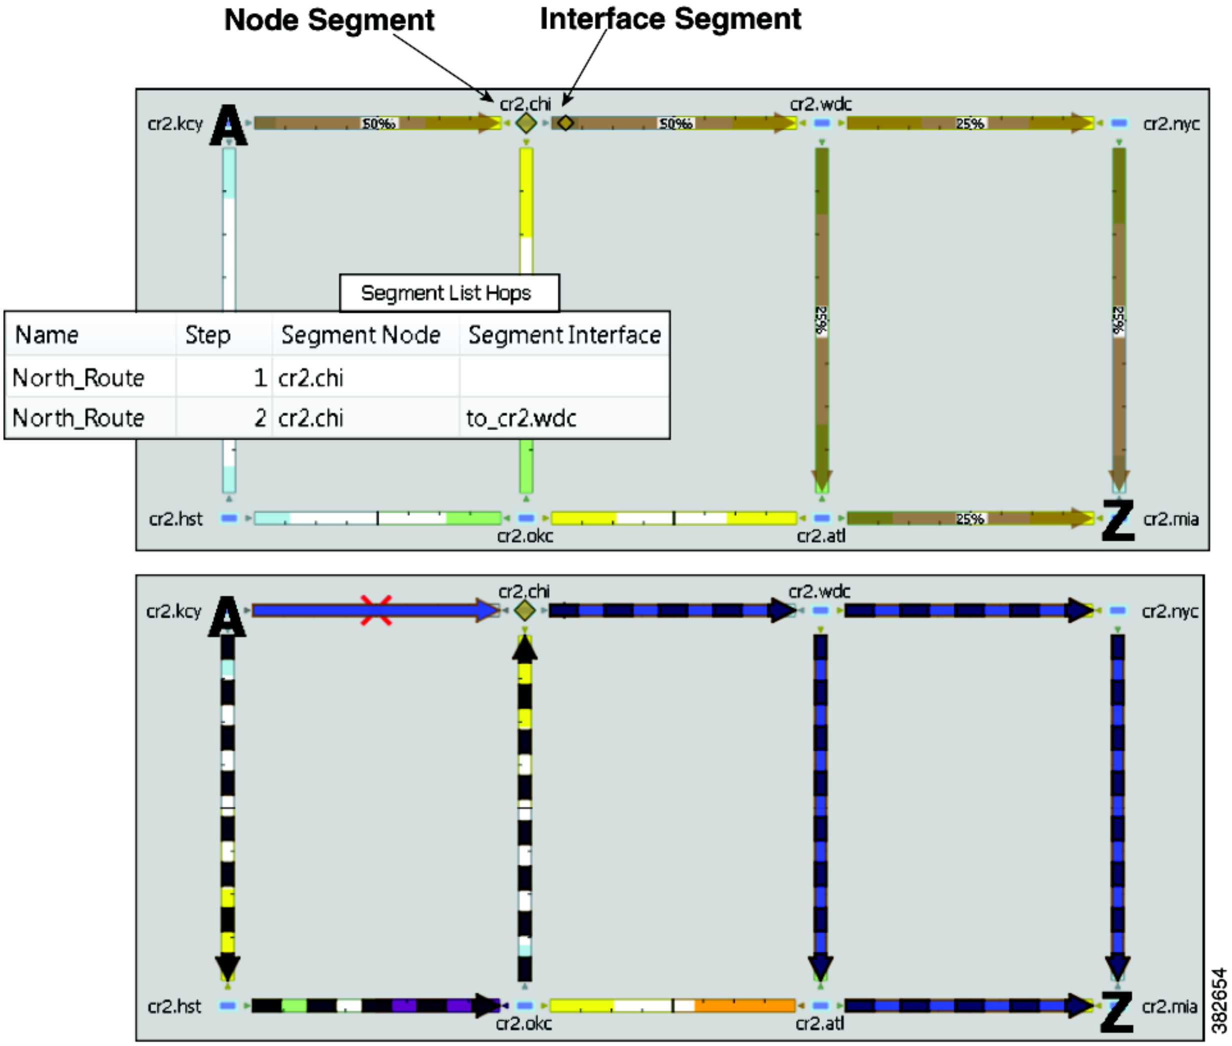

Segment List Hops

Name	Step	Segment Node	Segment Interface
North_Route	1	cr2.chi	
North_Route	2	cr2.chi	to_cr2.wdc

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 9-34 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 9-34 Example SR LSP Paths

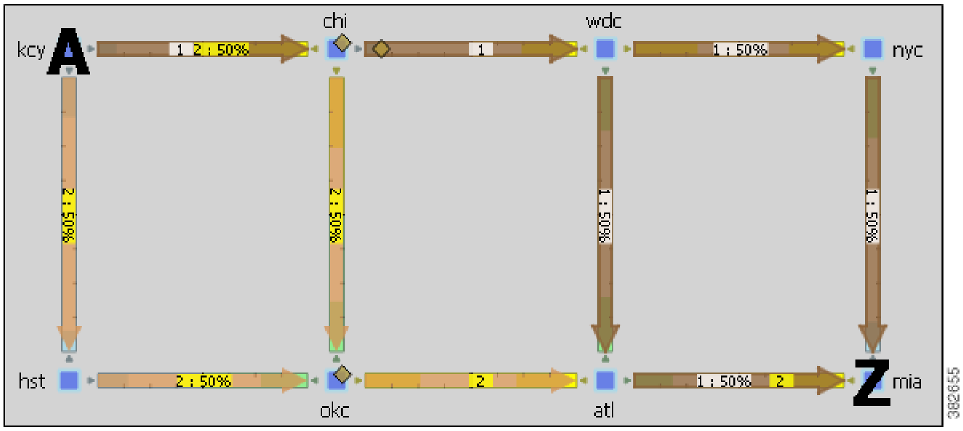

SR LSP Routing

Assuming the common LSP mechanisms, such as autoroute and forwarding adjacencies have determined that a demand will enter an LSP, and assuming it is 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.

- 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 IGP path to that node.

  If it is using an interface segment, the demand takes the shortest path the source node of that interface, and then routes to that interface's destination node.

- 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 on 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.

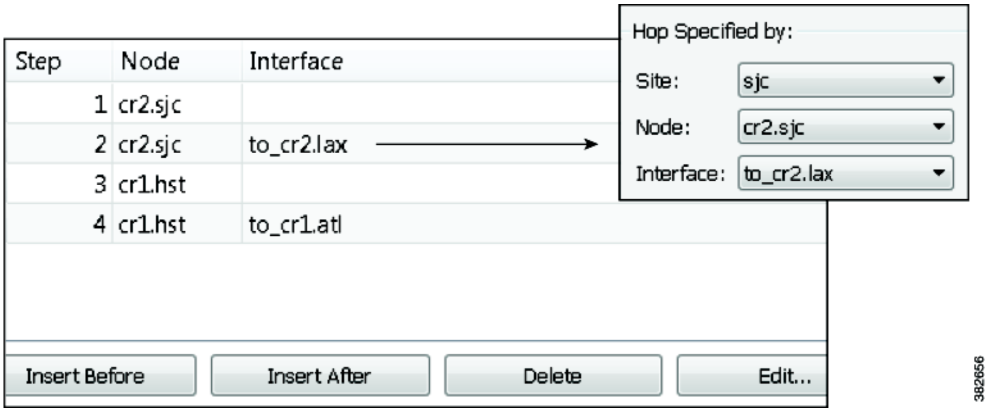

Step	Node	Interface
1	cr2.sjc	
2	cr2.sjc	to_cr2.lax
3	cr1.hst	
4	cr1.hst	to_cr1.atl

Hop Specified by:

Site: sjc

Node: cr2.sjc

Interface: to_cr2.lax

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, and interface. For node segments, do not select an interface.

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

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

Network Options for MPLS Simulations

MATE 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

For More Information ...	See ....
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

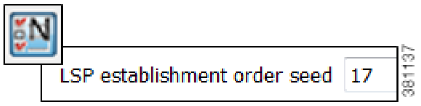
LSP establishment order seed 17

MATE 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. MATE 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, MATE 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

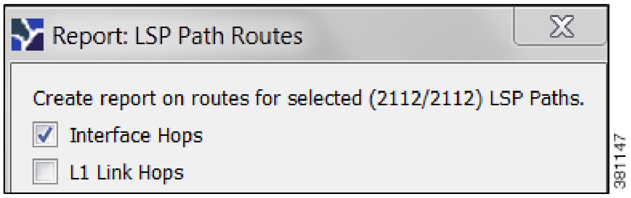
Report: LSP Path Routes

Create report on routes for selected (2112/2112) LSP Paths.

☑ Interface Hops

□ L1 Link Hops

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 on 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, MATE 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, MATE 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 on the selection in the table and select Simulation Diagnostics.

CHAPTER 10

Quality of Service Simulation

Quality of Service (QoS) is a means of ensuring high-quality performance for critical applications. The concept is that since requirements of some users and services are more critical than others, some traffic requires preferential treatment.

Using MATE Design QoS features, you can ensure these mission-critical service levels are met without reactively expanding or over-provisioning the network. These QoS features are available for undifferentiated traffic, for service classes alone, for interface queues alone, and for service classes that are mapped to queues.

Undifferentiated traffic—Aggregate traffic on an interface.

Service class—A user-defined classification of traffic that is not discovered by Collector. Examples could include voice, video, and data. Service classes apply to the entire network unless you map them to specific queues.

Queue—In live networks, traffic waits in conceptual lines (queues) and then is forwarded over an interface on a per-queue basis according to QoS parameters. Similarly in MATE, each queue has a set of user-defined QoS parameters (interface queue properties) that specify how these queues are prioritized and what percentage of traffic they carry. An interface may contain zero or more queues that are discoverable by Collector. You can also manually create and configure them. The traffic per queue is also discovered.

QoS Parameters

In MATE, QoS requirements are defined by policies and interface queue properties (Figure 10-1).

Policy—Maximum percentage of traffic capacity that can be utilized by either a service class or by undifferentiated traffic. There are two policies: one for normal operation and one for worst-case scenarios. Policies set on service classes do not affect QoS requirements of any other service class. Nor would this parameter have any effect on live network behavior.

In addition to setting policies for the network, you can further refine the policy to a group of interfaces (called *interface policy group*). For instance, you might need to model behavior for service classes offered only on interfaces of a specific capacity, or you might need to observe service class traffic on only one area of the network.

Interface queue properties—Configured parameters that would affect routing behavior in a live network. In MATE Design, the interface queue properties are priority, weight, and police limit.

- The **priority** identifies the precedence of the queue. For example, traffic in a priority 1 queue is routed before traffic in a priority 2 queue. Queues with the same priority evenly share the capacity based on weighted-round robin (WRR) calculations. You can change this behavior using the weight and police limit parameters. There are an unlimited number of priorities, though most networks only use no more than three. By default, queues do not have priorities.
- The **weight** is the percentage of preference given to queues of an equal priority level, which enables the network to fairly distribute the load amongst available resources. For example, if 10 Gbps were passing through a 10GbE interface on two priority 1 queues, by default 5 Gbps would pass through each queue. However, if you set the weight of one queue to 75% and the other to 25%, the distribution would be 7 Gbps and 2.5 Gbps, respectively. By default, all queues have a weight of 100%.
- The **police limit** is the maximum percentage of available capacity permitted through a queue of a given priority level, thereby preventing traffic from higher priority queues from starving lower priority queues. For instance, if the interface is a 20GbE and a priority 1 queue has a police limit of 40%, then only 8 Gbps of interface traffic can go through this queue. By default, all queues have a police limit of 100%. To see examples of this *starvation*, refer to the examples in the Policies and QoS Bound Calculations section where you can see that lower priority queues received zero traffic due to priority settings.

For More Information ...	See ...
Create service classes	Create Service Classes
Create and editing policies	Create and Edit Service Class Policies
Create interface policy groups	Create or Edit Policy Groups for Interfaces
Create Interface queues	Create Queues
Assign queues to interfaces	Assign Queues to Interfaces
Edit interface queue properties	Edit Interface Queues Properties (QoS Requirements)
Map service classes to queues	Map Service Classes to Queues
Import QoS model into the plan fie	*MATE Integration and Development Guide*

Figure 10-1 Policies and Interface Queue Parameters

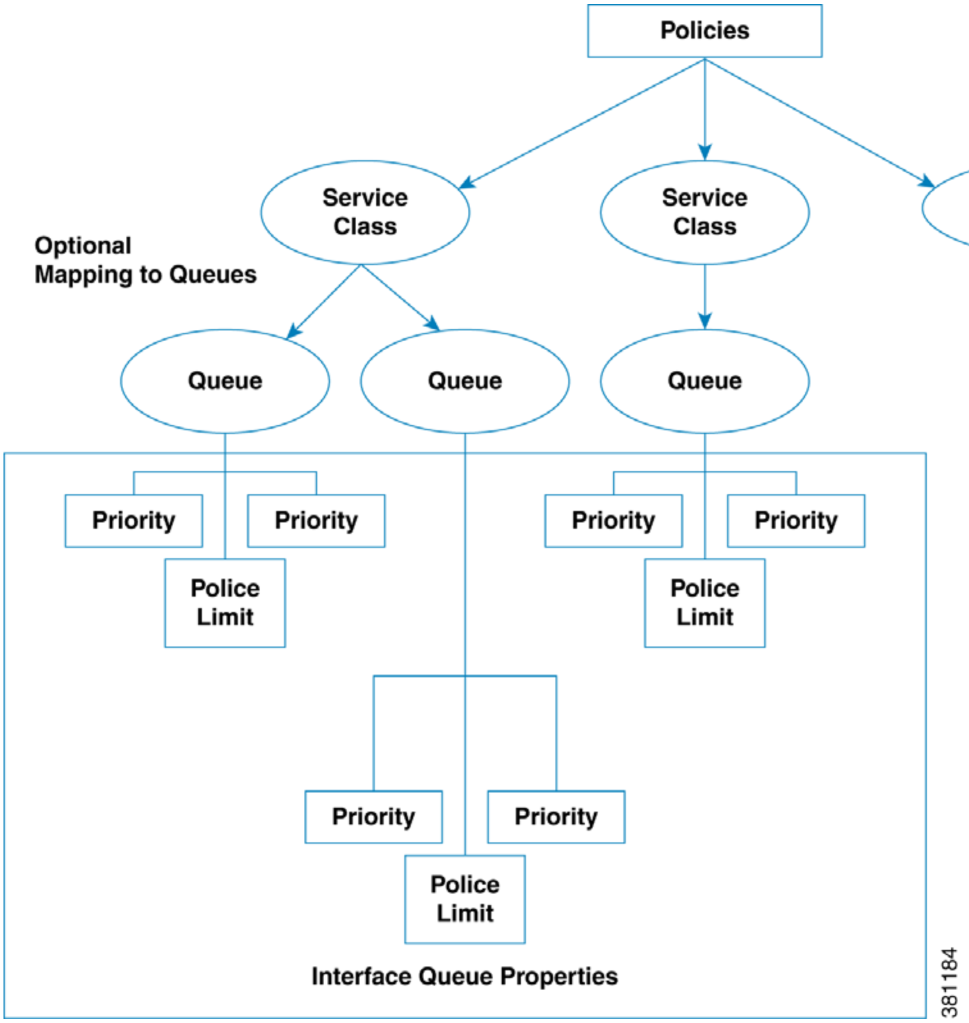

QoS Bound and QoS Violations

MATE Design uses the concepts of *QoS bound* and *QoS violation* as a way of identifying whether QoS parameters are being met or surpassed, thus better enabling you to plan for service requirements across the network. Policies and queue properties determine the QoS bound calculation. In turn, this calculation determines whether there is a violation.

QoS Bound—Maximum interface capacity available without violating these QoS requirements. A separate QoS bound is calculated for both policy and interface queue properties.

QoS Bound For ...	Calculation Based On ...
Interface Queues	Combination of interface queue properties, or in live networks, it is the capacity percentage that is discovered.
Service class	Policy

QoS Bound For ...	Calculation Based On ...
Service class mapped to queues	The lower of these two calculations is used. • Policy for service class • Queue properties for queues
Undifferentiated traffic	Policy

In the plot, the QoS bound on an interface is a combination of the color and white (Figure 10-2). The interface capacity that is not available because it exceeds the QoS bound is in gray. Columns that convey the QoS bound information are QoS Bound Meas, QoS Bound Meas (%), QoS Bound Sim, and QoS Bound Sim (%).

Figure 10-2 QoS Bound and Available Interface Capacity

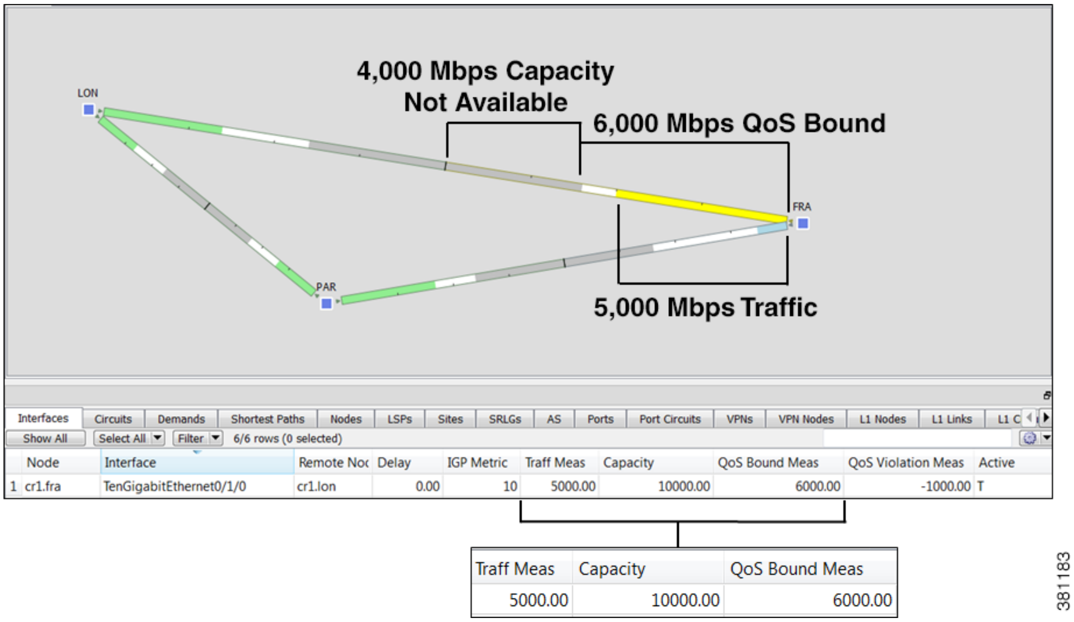

6/6 rows (0 selected)

	Node	Interface	Remote Noc	Delay	IGP Metric	Traff Meas	Capacity	QoS Bound Meas	QoS Violation Meas	Active
1	cr1.fra	TenGigabitEthernet0/1/0	cr1.lon	0.00	10	5000.00	10000.00	6000.00	-1000.00	T

Traff Meas	Capacity	QoS Bound Meas
5000.00	10000.00	6000.00

QoS bound calculations are a set of decisions being made to determine how to raise traffic on the queues until that traffic cannot be raised any further. This capacity, or the reason the traffic cannot be raised further, is defined both by the QoS parameters and the amount of traffic. For example, when traffic arrives at Queue X, MATE Design fixes the traffic on all other queues and then determines how it can raise the traffic on Queue X until some other traffic blocks it.

For those queues that do not reach full capacity, unused queue capacity is made available for other queues.

QoS Violation—Total traffic minus the capacity permitted for the queue (QoS bound). A violation occurs if the maximum QoS capacity allotted through policies and interface queue properties is exceeded. If the number appearing in the QoS Violation column is positive, then the allotted capacity has been surpassed. If the number is negative, the allotted capacity has not been reached. In the plot, traffic exceeding the QoS bound appears in red and white stripes on the interface in violation (Figure 10-3).

Figure 10-3 Example QoS Violation

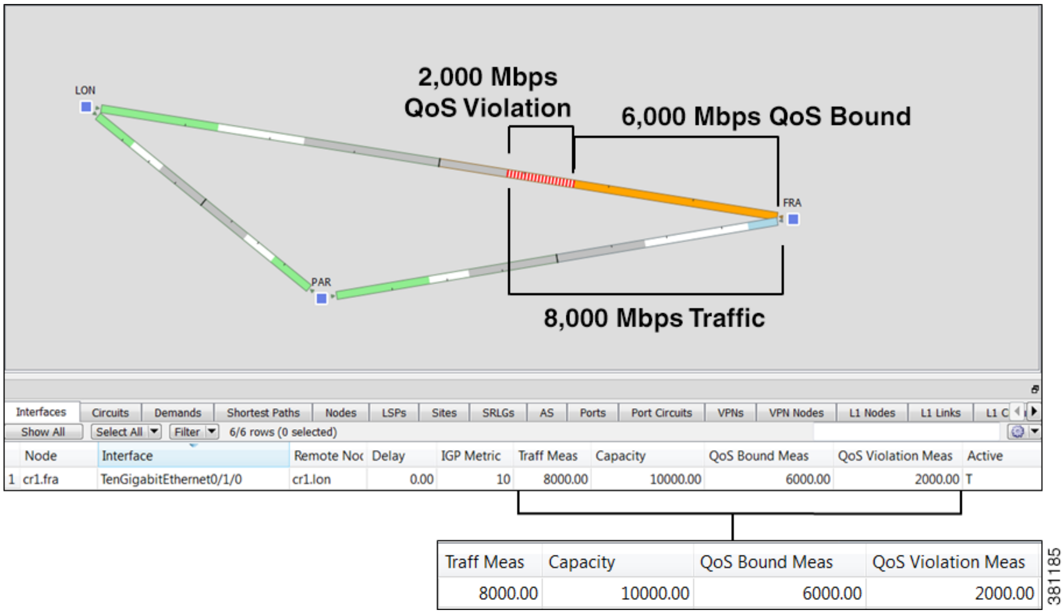

	Node	Interface	Remote Noc	Delay	IGP Metric	Traff Meas	Capacity	QoS Bound Meas	QoS Violation Meas	Active
1	cr1.fra	TenGigabitEthernet0/1/0	cr1.lon	0.00	10	8000.00	10000.00	6000.00	2000.00	T

Traff Meas	Capacity	QoS Bound Meas	QoS Violation Meas
8000.00	10000.00	6000.00	2000.00

Policies and QoS Bound Calculations

If no other QoS parameters are set via the interface queue properties of priority, weight, and police limit, the QoS bound is equivalent to the policy set.

Table 10-1 Example Policies and QoS Bound Calculations

Example Configuration	QoS Bound	QoS Violation (Positive # = Violation)
Interface capacity = 10,000 Mbps Undifferentiated traffic = 5,000 Mbps Normal operation policy = 60%	6,000 Mbps (60%)	-1,000 Mbps (-10%) Since this number is negative, there is no capacity violation (Figure 10-2).
Interface capacity = 10,000 Mbps Undifferentiated traffic = 8,000 Mbps Normal operation policy = 60%	6,000 Mbps (60%)	2,000 Mbps (20%) Since this number is positive, there is a capacity violation (Figure 10-3).

Interface capacity = 10,000 Mbps
Voice traffic = 6,000 Mbps
Video traffic = 2,000 Mbps
Voice normal operation policy = 90%
Video normal operation policy = 60%
	Voice = 9,000 Mbps (90%)
	Video = 6,000 Mbps (60%)
		Voice = -3,000 Mbps (30%)
		Video = -4000 Mbps (40%)

Interface Queue Properties and QoS Bound Calculations

MATE Design simultaneously calculates QoS bound for each queue in the interface. In so doing, MATE Design uses the interface queue parameters (priority, weight, and police limit) and the traffic measured or simulated for all queues in the interface. Priority is always considered first. If there are queues of equal priority, then weight is applied next.

- Queues with priority 1 share all available interface capacity. Their weight and police limits further refine how much each priority 1 queue can use (their QoS bound). Each priority 1 queue can borrow available capacity from other priority 1 queues up to the limit of their QoS bound.
- The available capacity for priority 2 queues is the total interface capacity less all capacity consumed by priority 1 queues. The process then begins again for all priority 2 queues. Their weight and police limits determine their QoS bound, and priority 2 queues can borrow capacity from each other up to the limits set by the QoS bound.
- This process continues for each successive priority level. Traffic that is outside any QoS bound is dropped to the lowest priority of all traffic on the interface.

For discovered networks with measured traffic, if no MATE Design QoS parameters are set, the QoS bound is based on whatever capacity percentages the live network has for each queue.

Priority

Provided policies are not set that further affect the QoS bound, a queue's QoS bound is calculated as follows.

- Priority 1 QoS bound = 100% of the interface capacity
- Priority 2 QoS bound = Total interface capacity - amount of traffic consumed by priority 1 queues
- Priority 3 QoS bound = Total interface capacity - amount of traffic consumed by (priority 1 + priority 2 queues)
- QoS bound for each succeeding priority follows this same pattern where the traffic consumed by all higher priority queues is subtracted from the total interface capacity.

Table 10-2 Examples of Priority QoS Bound Calculations

Example Configuration	QoS Bound	QoS Violation (Positive # = violation)	QoS Bound Calculations

Interface capacity = 20,000 Mbps
EF traffic = 6,000 Mbps; priority = 1
BE traffic = 3,000 Mbps; no priority set
	EF = 20,000 Mbps
	BE = 14,000 Mbps
		EF = -14,000 Mbps
		BE = -11,000 Mbps
			EF = Total interface capacity because it is the only priority 1 queue
			BE = 20,000 (interface capacity) - 6,000 (consumed by higher priority queues)

Interface capacity = 10,000 Mbps
EF Traffic = 6,000 Mbps; priority = 1
BE Traffic = 5,000 Mbps; priority = 2
	EF = 10,000 Mbps
	BE = 4,000 Mbps
		EF = -4,000 Mbps
		BE = 1,000 Mbps
			EF = Total interface capacity because it is the only priority 1 queue
			BE = 10,000 (interface capacity) - 6,000 (consumed by higher priority queues)

Weight

The weight identifies the forwarding precedence for queues of equal priority. If weights for queues of the same priority do not add up to 100%, weights are converted proportionally so they do add up to 100%.

Table 10-3 Examples of Weight QoS Bound Calculations

Example Configuration	QoS Bound	QoS Violation (Positive # = Violation)	QoS Bound Calculations

Interface capacity = 10,000 Mbps
AF1 traffic = 3,000 Mbps; priority = 1; weight = 100%
AF2 traffic = 6,000 Mbps; priority = 1; weight = 100%
	AF1 = 5,000 Mbps
	AF2 = 7,000 Mbps
		AF1 = -2,000 Mbps
		AF2 = -1,000 Mbps
			AF1 = Half of capacity for priority 1 queues since both queues have equal weights
			AF2 = 5,000 (half of capacity) + 2,000 (unused AF1 capacity)

Interface capacity = 10,000 Mbps
AF1 = 5,000 Mbps; priority = 1; weight = 60%
AF2 traffic = 6,000 Mbps; priority = 1; weight = 40%
	AF1 = 6,000 Mbps
	AF2 = 5,000 Mbps
		AF1 = -1,000 Mbps
		AF2 = 1,000 Mbps
			AF1 = 60% of capacity for all priority 1 queues
			AF2= 10,000 (interface capacity) - 5,000 (consumed by AF1 queue)

Police Limits

Priority 1 queues have 100% of the interface traffic, and thus starve out the remaining queues. To prevent this queue starvation, use police limits to configure how much of the maximum percentage should be available for a given priority level.

Table 10-4 Examples of Police Limit QoS Bound Calculations

Example Configuration	QoS Bound	QoS Violation (Positive # = Violation)	QoS Bound Calculations

Interface capacity = 10,000 Mbps
EF traffic = 1,000 Mbps; priority = 1; police limit = 50%
BE traffic = 2,000 Mbps; priority = 2
	EF = 5,000 Mbps
	BE = 9,000 Mbps
		EF = -4,000 Mbps
		BE = -7,000 Mbps
			EF = 50% of total interface capacity
			BE = 10,000 (interface capacity) - 1,000 (capacity consumed by EF)

Example Configuration	QoS Bound	QoS Violation (Positive # = Violation)	QoS Bound Calculations

Interface capacity = 10,000 Mbps
EF traffic = 1,000 Mbps; priority = 1; police limit = 5%
BE traffic = 2,000 Mbps; priority = 2
	F = 500 Mbps
	BE = 9,500 Mbps
		EF = 500 Mbps
		BE = -7500 Mbps
			EF = 5% of total interface capacity
			BE = 10,000 (interface capacity) - 500 (capacity consumed by EF)

Interface capacity = 10,000 Mbps
EF = 3,000 Mbps; priority = 1; police limit = 20%
AF1 traffic = 4,000 Mbps; priority = 2; police limit = 75%
AF2 traffic = 2,500 Mbps; priority = 2; police limit 25%
	EF = 2,000 Mbps
	AF1 = 6,000 Mbps
	AF2 = 4,000 Mbps
		EF = 1,000 Mbps
		AF1 = -2,000 Mbps
		AF2 = -1,500 Mbps
			EF = 20% of total interface capacity
			AF1 = 75% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF])
			AF2 = 10,000 (interface capacity) - 2,000 (capacity consumed by EF) - 4,000 (capacity consumed by AF1)

Interface QoS Bound Calculations Using Multiple QoS Parameters

MATE Design calculates a QoS bound for interface queues based on all three parameters if they are all configured: priority, weight, and police limits.

Table 10-5 Examples of Interface QoS Bound Calculations Using Multiple QoS Parameters

Example Configuration	QoS Bound	QoS Violation (Positive #= Violation)	QoS Bound Calculation

Interface capacity = 10,000 Mbps
EF = 3,000 Mbps; priority = 1; police limit = 20%
AF1 traffic = 4,000 Mbps; priority = 2; weight = 75%
AF2 traffic = 2,500 Mbps; priority = 2; weight = 25%
	EF = 2,000 Mbps
	AF1 = 6,000 Mbps
	AF2 = 4,000 Mbps
		EF = 1,000 Mbps
		AF1 = -2,000 Mbps
		AF2 = -1,500 Mbps
			EF = 20% of total interface capacity
			AF1 = Maximum of these two values.
			• 75% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF])
			• 8,000 (available capacity) - 2,500 (AF2 traffic)
			AF2 = Maximum of these two values.
			• 25% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF])
			• 8,000 (available capacity) - 4,000 (AF1 traffic)

Service Class QoS Bound Calculations Using Multiple QoS Parameters

If service classes have policies and they are mapped to queues, MATE Design calculates a QoS bound for both. MATE Design then uses the lowest value of the two so as to enforce restrictions in the strictest possible manner.

Example:

Interface capacity = 10,000 Mbps

QoS bound for service class = 50%, or 5,000 Mbps based on policy

QoS bound for EF queue = 7,500 Mbps based on combined parameters of priority, weight, and police limit

The QoS bound for this service class is 5,000 Mbps because the policy QoS bound calculation is lower.

View Queue and Service Class Information

Table 10-6 View Queue and Service Class Information

To View	Show or Select
Queue information	Show the Interface Queue table. Select the queue from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns display traffic data specific to the queue type selected.
Per-queue traffic in the Interfaces table	Select the queue from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns display data specific to the queue type selected.
Per-service-class traffic	Select the service class from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns in the Interfaces table display data specific to the service class selected.
Service class demands	Show the Service Class column in the Demands table.

View QoS Bounds and QoS Violations

As Figure 10-2 and Figure 10-3 demonstrate, the QoS bounds and QoS violations appear in the plot view. Table 10-7 lists the available column options to display numeric values of the QoS bound calculations. For information on QoS values as they relate to VPNs, see the VPN Simulation chapter.

Table 10-7 View QoS Bounds and QoS Violations

To View	Show this Column in Interfaces, Circuits, or Interface Queues Table
Measured Traffic	
Maximum capacity before a QoS bound is violated under normal operations	QoS Bound Meas
QoS bound as a percentage of total interface capacity	QoS Bound Meas (%)
QoS violations under normal operations; if the number is positive, there is a violation	QoS Violation Meas
QoS violation as a percent of the total interface capacity	QoS Violation Meas (%)
Simulated Traffic	
Maximum capacity before a QoS bound is violated under normal operations	QoS Bound Sim
QoS bound as a percentage of total interface capacity	QoS Bound Sim (%)

To View	Show this Column in Interfaces, Circuits, or Interface Queues Table
QoS violations under normal operations; if the number is positive, there is a violation	QoS Violation Sim
QoS violation as a percent of the total interface capacity	QoS Violation Sim (%)
Worst-Case Traffic	
Maximum capacity before a QoS bound is violated under worst-case operations	WC QoS Bound
WC QoS bound as a percentage of total interface capacity	WC QoS Bound (%)
QoS violations under worst-case operations; if the number is positive, there is a violation	WC QoS Violation
WC QoS violation as a percent of the total interface capacity	WC QoS Violation (%)
Service class causing the worst-case utilization	WC Service Class

Create Service 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 toolbar.

Step 2 Click New. The New Service Class dialog box that appears contains queues if they have already been discovered or if they have been manually created. Otherwise, the dialog box is empty. For instructions on how to create queues, see Create Queues.

a. Enter a unique name.

b. Optional: If queues exist and if you want to map this new service class to one or more queues, select them from the list. If queues do not exist, but you want them, you must manually create the queues and then return to this dialog box to select them. See Map Service Classes to Queues.

c. Click OK.

Step 3 Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the new service classes are not saved.

Figure 10-4 Create Service Classes

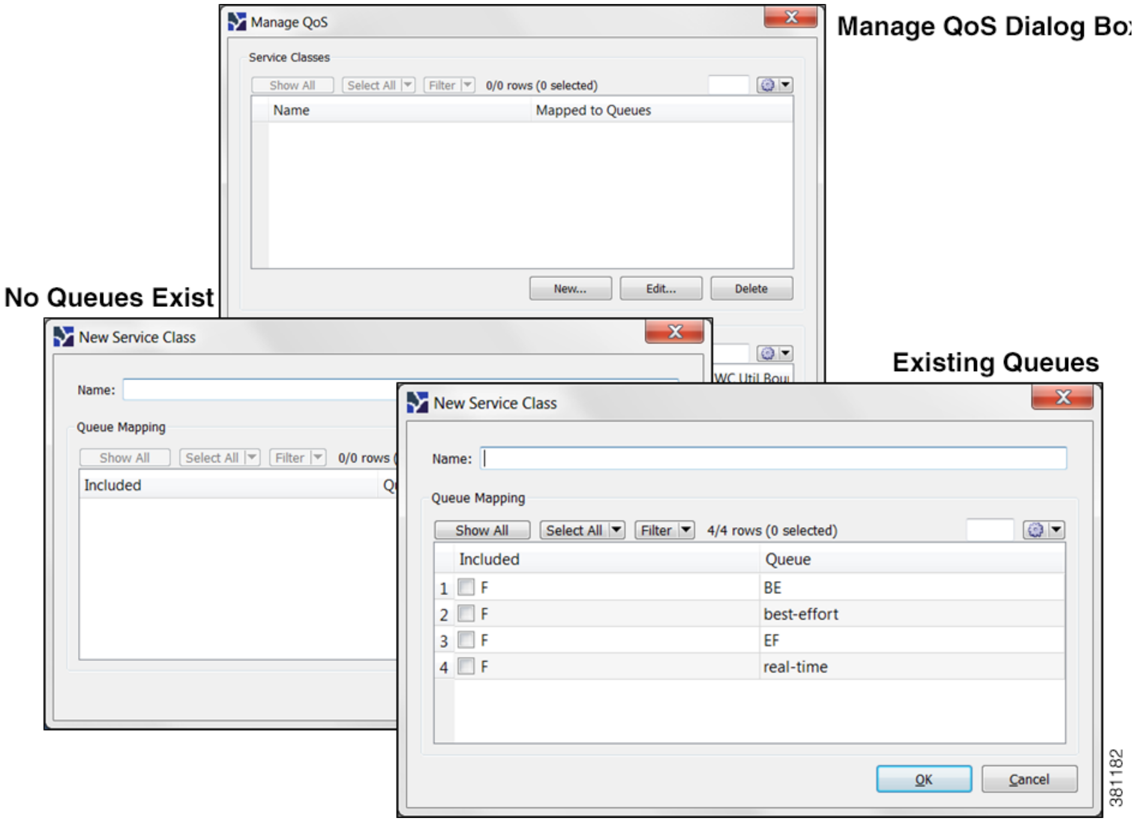

Manage QoS Dialog Bo

No Queues Exist

Existing Queues

	Included	Queue
1	□ F	BE
2	□ F	best-effort
3	□ F	EF
4	□ F	real-time

Create Queues

MATE Design discovers queues. However, you can manually add them. Once discovered or created, queues appear in the Interface Queues table.

Step 1 By default, queues are assigned to all interfaces. If you want this new queue to be assigned to specific interfaces, you must first select one or more interfaces.

Step 2 Open a New Interface Queues Properties dialog box in one of two ways.

- Select the Insert->Interface Queues menu.
- Right click in the plot area then select New->Interface Queues.

Step 3 Enter the queue name.

Step 4 Optional: Enter the queue properties of priority, weight, and police limit. For information on how these queue properties behave, click here: Interface Queue Properties and QoS Bound Calculations.

Step 5 Click OK. The new queue appears as an option in the QoS drop-down menu in the toolbar, as well as in the Manage QoS dialog box.

Step 6 Optional: Map a service class to the queue. For further instructions, click here: Map Service Classes to Queues.

Another way to create queues is to edit the name of an existing one. For instructions, see Assign Queues to Interfaces.

Assign Queues to Interfaces

The easiest way to assign queues to interfaces is to select the interface before creating the queue (see Create Queues). However, you can follow these steps to re-assign the queue to different interfaces.

Step 1 From the Interface Queues table, select one or more queues.

Step 2 Double click; a Properties dialog box appears.

Step 3 From the Node list, select the node associated with the interface.

Step 4 From the Interface list, select the interface to which you are assigning the queue.

Step 5 Optional: You can change the queue name or add a new one.

Step 6 Optional: Specify QoS requirements in the Priority, Weight, and Police Limit fields. By default, queues do not have a priority and both their weight and police limits are 100%. For information on how these properties behave, see Interface Queue Properties and QoS Bound Calculations.

Step 7 Click OK.

Map Service Classes to Queues

To map service classes to queues, those queues must first exist either because they were discovered by Collector or because you manually added them. See Assign Queues to Interfaces.

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 toolbar.

Step 2 From the Service Classes list, select one service class.

Step 3 Click Edit. A Service Class dialog box appears.

a. Select one or more queues.

b. Click OK.

Step 4 Repeat #2 and #3 for each service class to which you are mapping queues.

Step 5 Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the above changes are not saved.

Create or Edit Policy Groups for Interfaces

Creating a policy group for interfaces enables you to set policies for the group in the Manage QoS dialog box.

Step 1 Select one or more interfaces.

Step 2 Double-click to open the Properties dialog box.

Step 3 Select the Advanced tab.

Step 4 In the QoS Policy Group field, you have three options

- To add this interface to an existing policy group, select it from the drop-down menu.
- To change the name of an existing policy, select it and then type over it.
- To add a new policy group, select the empty row in the drop-down menu and then type the name of the new group.

Step 5 Click OK.

Step 6 To assign a service class to this policy group, use the Manage QoS option in the QoS drop-down in the toolbar. For further instructions, click here: Create and Edit Service Class Policies.

Create and Edit Service Class Policies

You can configure policies for undifferentiated traffic and for service classes. For information on how policies behave, see Policies and QoS Bound Calculations.

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 toolbar.

Step 2 In the Service Class Policies section, click New, or select a service class row and click Edit. A Service Class Policy dialog box appears.

a. If creating a policy for undifferentiated traffic, select that option. If creating a policy for an existing service class, select it from the Service Class drop-down menu.

b. Optional: To apply this service class mapping to a group of interfaces, select from or enter the name in the Interface Policy Group drop-down menu. You can enter a name that does not exist and then create the policy group for a set of interfaces. For further instructions, click here: Create or Edit Policy Groups for Interfaces.

c. In the Normal Operation field, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service class under normal conditions.

d. In the Worst-Case field, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service under worst-case operating conditions.

e. Click OK. These values now appear in the Manage QoS dialog box where you can edit them as needed.

Step 3 Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the above changes are not saved.

Figure 10-5 Manage QoS dialog box

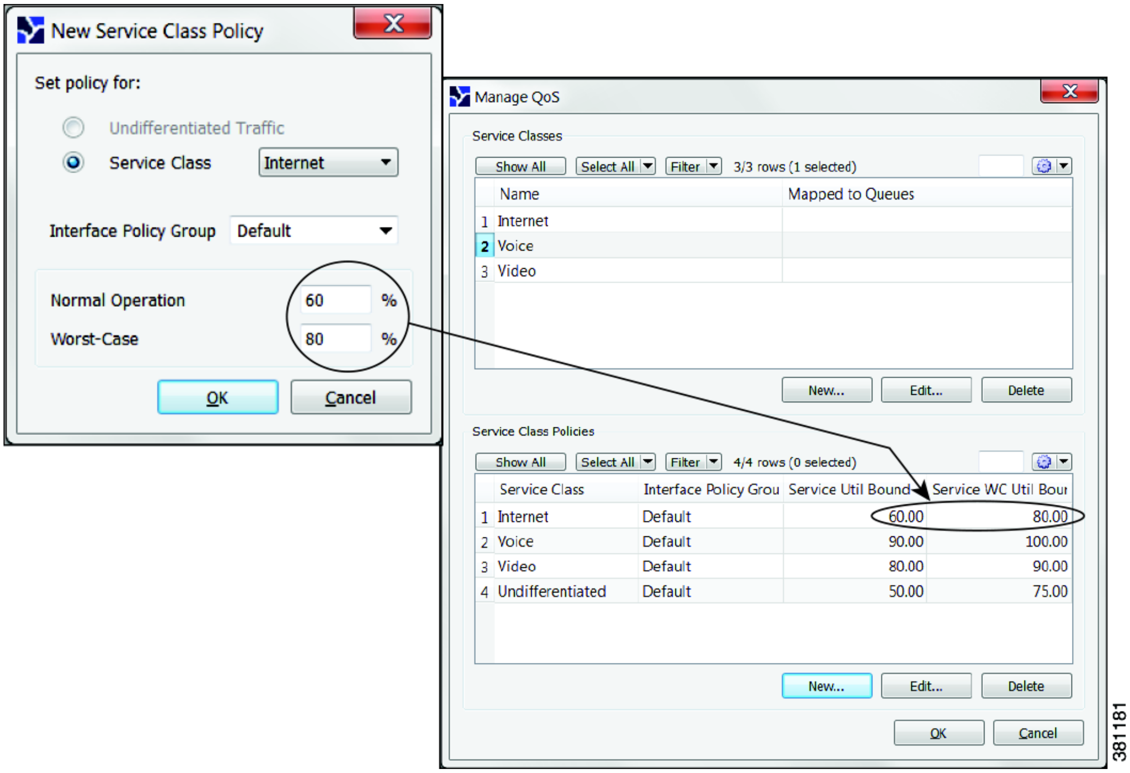

New Service Class Policy

Set policy for:

○ Undifferentiated Traffic
● Service Class: Internet

Interface Policy Group: Default

Normal Operation: 60 %
Worst-Case: 80 %

Manage QoS

Service Classes — 3/3 rows (1 selected)

	Name	Mapped to Queues
1	Internet	
2	Voice	
3	Video	

Service Class Policies — 4/4 rows (0 selected)

	Service Class	Interface Policy Grou	Service Util Bound	Service WC Util Bour
1	Internet	Default	60.00	80.00
2	Voice	Default	90.00	100.00
3	Video	Default	80.00	90.00
4	Undifferentiated	Default	50.00	75.00

Edit Interface Queues Properties (QoS Requirements)

The three parameters are priority, weight, and police limit. For information on how these queue properties behave, click here: Interface Queue Properties and QoS Bound Calculations.

Step 1 In the Interface Queues table, select one or more queues.

Step 2 Right click. The Interface Queues Properties dialog box appears.

Step 3 Change one or more fields to create the desired QoS requirement.

Step 4 Click OK.

CHAPTER 11

BGP Simulation

This chapter describes how MATE Design models multi-AS networks and simulates basic BGP routing. MATE Design does not directly emulate BGP routing configurations, such as local prefs and MEDs. Rather, it provides a high-level modeling of typical peering policies, such as standard customer, transit, and settlement-free arrangements for service providers. This model enables you to quickly and easily evaluate the effects of peering locations and basic policy variations.

Additionally, you can extend these high-level models to significantly more complex policy-based routing situations using *external endpoints* as demand sources and destinations.

For More Information...	See...
Demands	Traffic Demand Modeling chapter
External endpoints	Advanced Routing with External Endpoints chapter

Internal and External AS Types

To model a multi-AS network, each node is assigned an AS, and each AS is defined as either *internal* or *external*. A typical multi-AS model in MATE Design consists of the following.

- A single internal AS representing the full topology of your network.
- Individual peering nodes of neighboring external AS's.
- Peering circuits connecting the internal AS to the nodes in the external AS's.

Thus, generally, there are a large number of external AS's in the plan file, but usually only one or a small number of internal AS's. All nodes in an external AS are typically placed in the same site, although you can place them in any site. Figure 11-1 shows an example model with two internal AS's and four external AS's.

AS names and their types are defined in the AS Properties dialog box and listed in the AS table. Nodes are assigned to AS's in the node Properties dialog box.

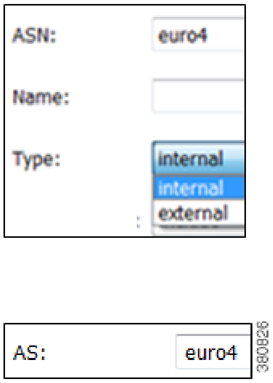

ASN: euro4

Name:

Type: internal

AS: euro4

Figure 11-1 Example Multi-AS Model

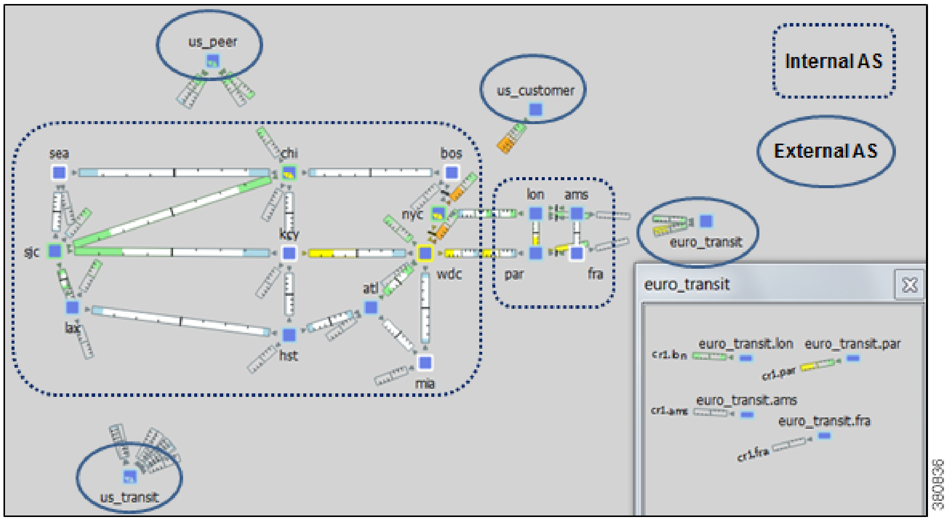

For More Information...	See...
Trim AS's from plan, but keep AS traffic	Trim AS's
Creating AS's	Create AS's
Assigning names to AS's	Assign Names to ASNs

Inter-AS Routing

Route Selection between Internal AS's

Demands routed within a single AS have a specified source node and destination node where traffic originates and terminates. Demands routed between two connecting internal AS's are specified in the same manner: with a source node in the first AS and a destination node in the second.

MATE Design routes within an AS, to and from the border exit point, are determined by the IGP protocols. The selection of border exit point is modeled by the Routing Policy, which is set to either Shortest Exit or Respect MEDs. This property is set in the Edit AS Relationships dialog box, which is accessed through the AS Properties dialog box.

- Shortest exit—The border exit node is selected, which is closest to the source node, within the IGP of the source AS. If there is a tie, the exit node with the lowest BGP ID is used.
- Respect MEDs—The border exit node is selected, which is closest to the destination node, within the IGP of the destination AS. If there is a tie, the exit node with the lowest BGP ID is chosen.

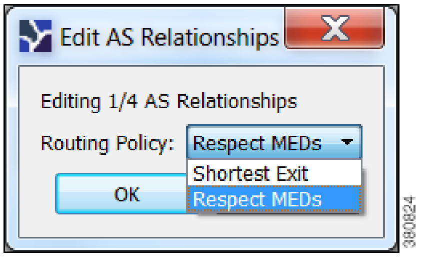

Edit AS Relationships

Editing 1/4 AS Relationships

Routing Policy: Respect MEDs

Figure 11-2 shows these modeled as two internal AS's: AS A and AS B. The same demand (A1 to B4) is routed using each of the two different routing policies.

Figure 11-2 Route Selection Based on Routing Policy

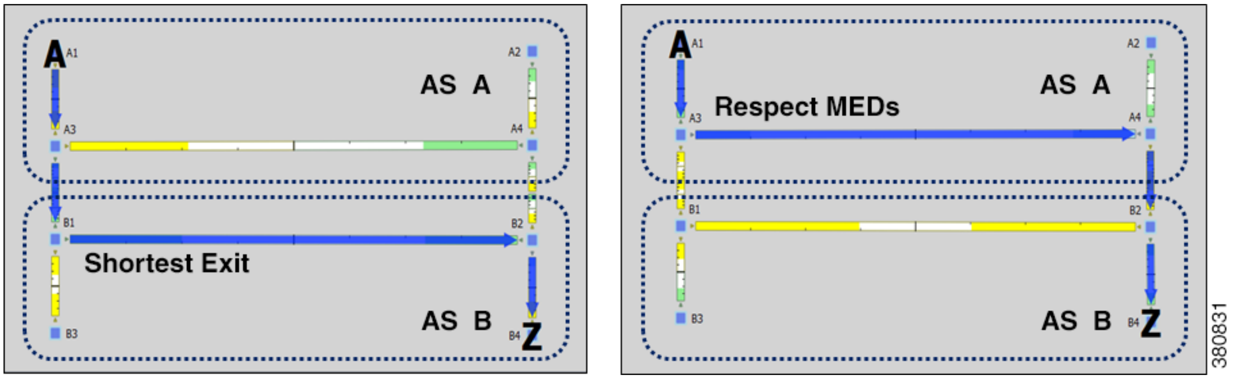

For More Information...	See...
Configure AS routing policies	Edit AS Routing Policies

Routes Between External and Internal AS's

Table 11-1 lists typical routing configurations that can be constructed by applying different combinations of Routing Policies for traffic in both directions between two AS's.

- In a peer relationship, routing in both directions is Shortest Exit, which means each controls its own border exit points.
- For a customer relationship, the customer determines the border exit points for traffic in both directions.
- For a transit relationship, the transit AS provides paid transit to the internal AS, so the internal AS determines all border exit points.

Table 11-1 Typical AS Routing Configurations

Type	Policy To	Policy From
Peer	Shortest Exit	Shortest Exit
Customer	Respect MEDs	Shortest Exit
Transit	Shortest Exit	Respect MEDs

Like traffic routed within an AS, traffic routed between AS's is represented by demands. However, for demands from and/or to external AS's, the external AS is defined as the source or destination of the demand. Optionally, the specific node in the external AS from which the traffic enters or exits the internal AS is also specified.

Failover between nodes in the external AS can be modeled. For example, if the traffic is sourced from an external AS and if the peering circuit from which traffic is entering the internal AS under normal operation fails, the traffic can enter the internal AS from a different interface or peering node in the same external AS. In the Demands table, the sources and destinations are represented as follows.

AS{<ASN>}:if{node_name|interface_name}

Example: AS{33287}:if{cr01.newyork.ny|POS3/7/0/0}

For more detailed information on demand sources and destinations, see the Traffic Demand Modeling chapter.

The AS that controls the routing chooses which peering node to use. If the internal AS controls the routing, then because the topology of the internal AS is known, you can simulate the routing to the peering node. However, since MATE Design has limited knowledge of the external AS topology, if the external AS controls the routing, you cannot predict how traffic will be distributed among the exit points.

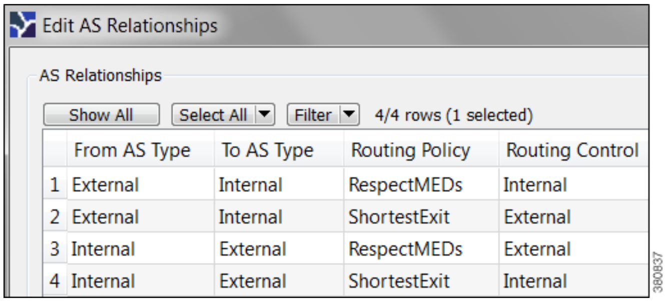

	From AS Type	To AS Type	Routing Policy	Routing Control
1	External	Internal	RespectMEDs	Internal
2	External	Internal	ShortestExit	External
3	Internal	External	RespectMEDs	External
4	Internal	External	ShortestExit	Internal

The AS that controls the routing is determined by the AS type, direction of the demand, and the Routing Policy property as described in Table 11-2.

Table 11-2 Determining the AS that Controls the Routing

Direction	Routing Policy	AS with Routing Control
External AS to Internal AS (Ingress)	Respect MEDs	Internal
	Shortest Exit	External
Internal AS to External AS (Egress)	Respect MEDs	External
	Shortest Exit	Internal

Two AS's can be in one of four different routing relationships to one another, depending on which of the two routing policies is chosen in each direction (Table 11-3).

- If traffic is routed to an external AS when it has control and there is no knowledge of its topology, a set of demands is created from the source in the internal AS (or from another external AS), each with a destination set to one of the border nodes in the external AS. This way, any division of traffic between the exit points can be modeled.
- If traffic is routed to an external AS when an internal AS has control, a single demand is created from the source to the AS itself. MATE simulations determine the correct exit point for this single demand based on the source.
- If traffic is to be routed from an external AS when it has control, then a demand is created from each node in the external AS to each node in the internal AS.
- If traffic is to be routed from an external AS when an internal AS has control, then a demand is created to each node in the internal AS using the external AS as the source. The demand originates from one or multiple nodes in the external AS, depending on the topology and the metric cost to reach the destination node. For example, a single demand from an external AS to a specific node could be sourced from two different nodes in the external AS, each carrying 50% of the demand traffic.

Table 11-3 Effects of Routing Policy and Routing Control

Direction	Routing Policy	AS with Routing Control	Demand Source or Destination Endpoint in Remote AS	Number of Demands
External AS to Internal AS (Ingress)	Respect MEDs	Internal	Entire external AS	One only
	Shortest Exit	External	Border nodes	One for each node
Internal AS to External AS (Egress)	Respect MEDs	External	Border nodes	One for each node
	Shortest Exit	Internal	Entire external AS	One only

External Meshes

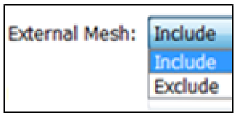
External Mesh: Include
Include
Exclude

An external mesh consists of two or more external AS's with a Type property of External. An internal AS typically restricts advertisement of BGP routes for some external AS's to other external AS's. For example, destinations reachable through the transit network would not be advertised to a peer, or vice versa. In MATE Design, these restrictions are represented by the absence of demands between the two external AS's.

Each AS has a property called External Mesh, which MATE Design uses when inserting demand meshes into a plan. Demands are created for external AS's only if one or both AS's have External Mesh set to Include. If both AS's are set to Exclude, no demands are created for the external AS. For example, in Figure 11-3 the peer and transit AS's are both set to Exclude, so no demands are created between those AS's. All other external AS demands are included in the demand mesh. Table 11-4 shows the External Mesh settings for common AS relationships.

The External Mesh property is set in the AS Properties dialog box.

Table 11-4 External Mesh Settings for Common AS Relationships

Relationship	External Mesh Setting	Result
Peer	Exclude	Demands permitted to/from customers only
Customer	Include	Demands permitted to/from all external AS's
Transit	Exclude	Demands permitted to/from customers only

For internal AS's, the External Mesh property is ignored. More complex route advertisement policies cannot be represented by these simple External Mesh settings. In this case, demand mesh creation must be performed in several steps, possibly using a script and the `dmd_mesh_creator` CLI tool.

Figure 11-3 External Mesh Control

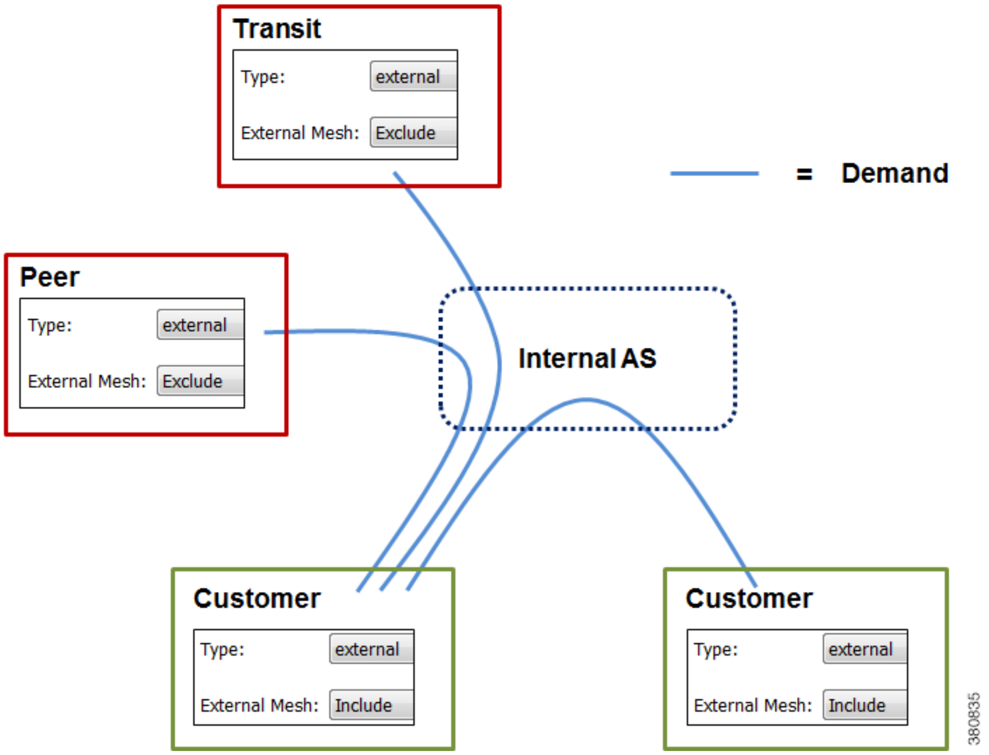

BGP Simulations

As with all MATE Design simulations, AS routing uses demands. An IP simulation for a particular failure scenario and traffic level performs these steps.

Step 1 Demands are routed using the established LSPs (if applicable) and using the specified BGP protocols given the specified failure scenarios.

Step 2 Interface utilizations are calculated from the demand traffic using the specified traffic level.

MATE Design allows routes to be calculated between selected nodes even if no demands are present. In this case, only the first step applies.

BGP demands do not failover between external AS's. That is, all traffic to or from an external AS behaves the same under peering failures to an external AS. You can change this default behavior using external endpoints to simulate specific external AS nodes where traffic goes in and out of the network, as well as set priorities so that if one traffic source or destination goes down, the traffic can still be sourced from or delivered to another external AS node.

For More Information...	See...
Demands	Traffic Demand Modeling chapter
External endpoints	Advanced Routing with External Endpoints chapter

BGP Routing Details

BGP Multihop

Collector automatically constructs BGP pseudonodes where necessary when BGP multihops are detected.

MATE Design models the nodes in external AS's that are directly connected, for example through eBGP, to nodes in internal AS's. One exception is that you can model BGP multihops by setting the node Type property to psn (pseudonode), such as might occur at a peering exchange. This pseudonode can represent the switch that connects a number of external AS nodes to the same internal AS node. In this instance, multiple external AS nodes are connected by circuits to a BGP psn node, and this node is connected to a node in the internal AS.

Note In all cases, eBGP multipaths across parallel border circuits is assumed.

BGP Load Balancing

BGP load balancing to an external AS uses eBGP multipaths or eBGP multihops. MATE Design models these two eBGP load balancing designs in the same manner, though in the GUI they are identified only as multipaths. BGP multipath options are disabled by default.

To set BGP multipath options globally, follow these steps.

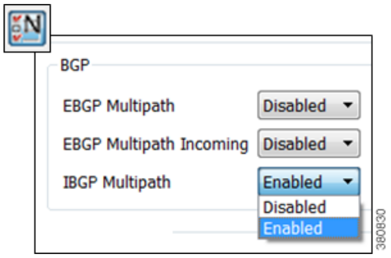

BGP

EBGP Multipath: Disabled

EBGP Multipath Incoming: Disabled

IBGP Multipath: Enabled

Step 1 Select the Edit->Network Options menu or the Network Options icon from the Visualization tool bar.

Step 2 Select the Protocols tab.

Step 3 For each BGP multipath option that you want enabled, select Enabled, and then click OK.

- EBGP Multipath—Turns on eBGP multipath within the internal AS's. Demand routings through the internal AS to an external AS are divided among external routes with equal-cost BGP exit routes.

- EBGP Multipath Incoming—Turns on eBGP multipath in all external AS's. Demand routings from external AS's into an internal AS are divided among external routes with equal-cost BGP exit routes.
- IBGP Multipath—Turns on iBGP multipath within the internal AS's. Demand routes through an internal AS to an external AS are divided among internal paths to equal-cost BGP exit routes.

BGP Next Hop

In networks, there are two common configurations for the BGP next-hop IGP metric used in the path selection. One is to set the next-hop self on the iBGP peers (next-hop self = on). The other is to configure IGP metrics on eBGP interfaces, and to inject the interface prefix into the IGP database by setting the interface to be a passive IGP interface (next-hop self = off).

MATE does not have an explicit next-hop self setting so it simulates paths as if next-hop self is off. That is, the IGP metric of the egress peering interface is included in the IGP distance to the peering router and is used in the iBGP path selection. However, next-hop self to an external AS can effectively be simulated by setting the metrics on all egress interfaces to that external AS to 0. (You can set the IGP metric in either the Interfaces Properties or Circuit Properties dialog box accessed through the context menu.)

Trim AS's

A plan file created through Collector BGP discovery typically has more external AS's in the plot than you are interested in viewing or simulating. Some AS's might have very low traffic levels and others might have only one circuit connecting into the internal AS. You can trim these unwanted AS's to simplify the network. Trimming removes the AS from the plan file, but retains the traffic in and out of the AS for simulation. That is, its interface traffic is still included in the total traffic of connected nodes.

Typically, plan files are trimmed before demands are created. However, you can optionally trim demands in a plan file. If any demand is sourced from a trimmed external AS, the demand source is moved to the first hop on the demand path in the remaining network. Likewise, if a demand has a destination that is a trimmed external AS, the demand destination is moved to the last hop in the remaining network. If two or more demands in the same service class are trimmed so that their resulting source and destination nodes are equal, these demands are aggregated into one demand with their traffic summed. Demands split by ECMPs are converted into multiple demands, each with traffic divided proportionally to the split ECMP.

Note that multicast demands are removed, and cannot be trimmed.

Step 1 Select one or more AS's to trim from the AS table.

For example, sort the AS's by traffic level and then select all AS's that are below the level of interest. Alternatively, filter the table for traffic using the less than operator (<) and the traffic level, and then select the filtered AS's.

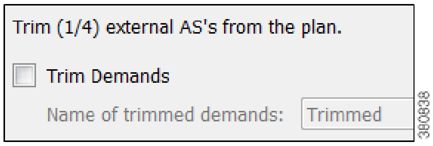

Trim (1/4) external AS's from the plan.

□ Trim Demands

Name of trimmed demands: Trimmed

Step 2 Right-click on one of the selected AS's and select Trim from the context menu.

Step 3 Optional: Select the option to trim demands.

By default, the moved demands are named "Trimmed," though you can enter a different name.

Step 4 Click Yes to confirm the action.

Create AS's

Follow these steps to create an empty AS. After creating the AS, you still need to associate nodes with it and create the relationship between this AS and others. See Associate Nodes with an AS and Edit AS Routing Policies, respectively.

Step 1 Either select the Insert->AS menu, or right-click in the empty plot and select New->AS. The AS Properties dialog box appears.

Step 2 Configure the AS properties.

- ASN—AS number, which is a text string that can be a number or name.
- Name—AS name.
- Type—Internal AS's have a full topology, and external AS's have a collapsed topology with just border nodes and a virtual node.
- External Mesh—When creating a demand mesh, this option tells MATE whether to create external meshes. When one or both AS's are set to Include, MATE creates a mesh between the external AS's (default). If both are set to Exclude, no demands are created.
- Description—A text description of the AS.

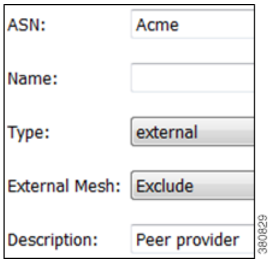
ASN: Acme

Name:

Type: external

External Mesh: Exclude

Description: Peer provider

Step 3 The change the routing policy, click Edit AS Relationships. (The default is Respect MEDs.) For information on how these options affect MATE simulation, see the Route Selection between Internal AS's section.

Step 4 Click OK. This creates an empty AS.

Associate Nodes with an AS

Step 1 Select one ore more nodes, and double-click on any on of them to open a node Properties dialog box.

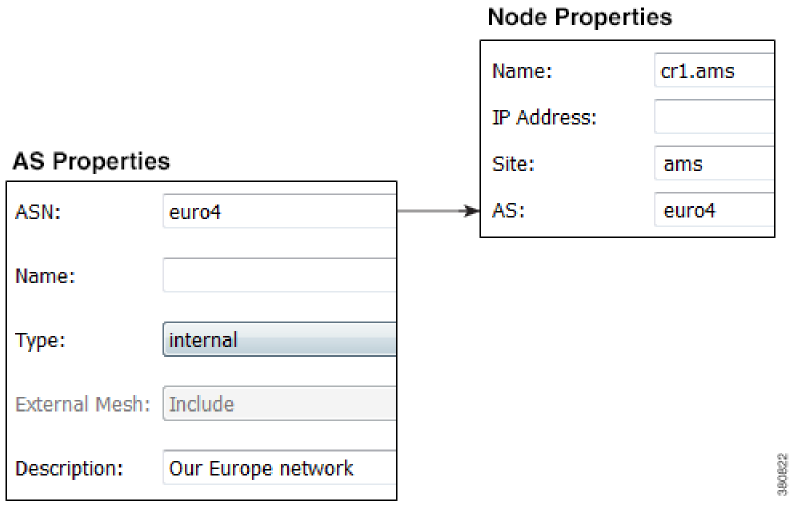

AS Properties

ASN:	euro4
Name:	
Type:	internal
External Mesh:	Include
Description:	Our Europe network

Node Properties

Name:	cr1.ams
IP Address:	
Site:	ams
AS:	euro4

Step 2 In the drop-down AS list, select the AS to which you want to assign the nodes.

Step 3 Click OK.

Edit AS Routing Policies

The way to create AS relationships is to set the Routing Policy. For a description of AS relationships, see the Inter-AS Routing section.

Step 1 Select one or more AS's from the AS table.

Step 2 Follow one of these methods to access the Edit AS Relationships dialog box.

- Select the Edit->AS Relationships menu.
- Right-click on one of the selected AS's, and select Edit AS Relationships from the context menu.
- Right-click on one of the selected AS's, and select Properties from the context menu. Click on Edit AS Relationships.

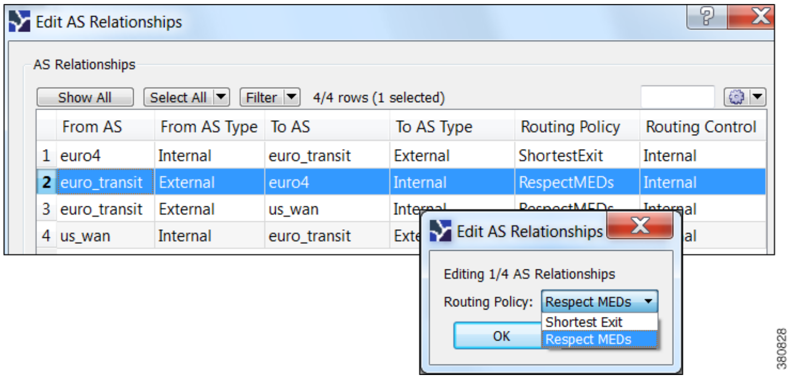

4/4 rows (1 selected)

	From AS	From AS Type	To AS	To AS Type	Routing Policy	Routing Control
1	euro4	Internal	euro_transit	External	ShortestExit	Internal
2	euro_transit	External	euro4	Internal	RespectMEDs	Internal
3	euro_transit	External	us_wan	Internal	RespectMEDs	Internal
4	us_wan	Internal	euro_transit	Exte		al

Editing 1/4 AS Relationships

Routing Policy: Respect MEDs

Step 3 Select the AS pair that you want to configure. There is a separate line for each direction in the relationship so you can configure them independently.

Step 4 Click Edit.

a. Set the Routing Policy to Respected MEDs (default) or Shortest Exit.

b. Click OK.

Step 5 Click OK again.

Assign Names to ASNs

MATE Design can assign names to ASNs by looking them up in the `$CARIDEN_HOME/etc/ASNs.txt` file. This file contains a list of all assigned ASNs with their names and descriptions.

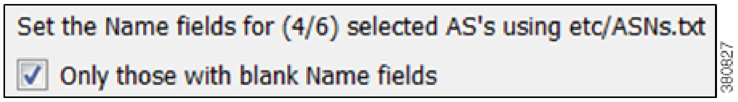

Set the Name fields for (4/6) selected AS's using etc/ASNs.txt

☑ Only those with blank Name fields

Step 1 Select one or more AS’s from the AS table. If no AS’s are selected, names are assigned to all AS’s in the plan.

Step 2 Either select the Initializers->Assign Names to ASNs menu, or right-click on a selected AS and select Assign Names to ASNs from the context menu.

Step 3 To assign names only if the ASN does not have one, select the “Only those with blank Name fields” option.

Step 4 Click OK.

CHAPTER 12

Advanced Routing with External Endpoints

To model basic IGP routing, demands are sourced or destined for nodes within the topology. To model basic inter-AS routing, the sources and destinations are neighboring external AS's, or a combination of the external AS and the peering node in that AS. However, more complex routing situations require the use of *external endpoints* as the source or destination. External endpoints can contain multiple member nodes and AS's, and you can specify when traffic enters or exits from each of them individually. This allows you to simulate routing within and between AS's where multiple traffic entry and exit points are used simultaneously. You can also prioritize where the traffic fails over to other nodes and AS's.

There are numerous use cases both for IGP and inter-AS routing.

- Simulate content caching failovers for in-network source of demands that are backed up by another in-network source. If connectivity is lost to the first, the traffic is then sourced from the second.
- Simulate edge routing with a single entry point into the network edge and a specific failover point. Alternatively, you could model multiple entry points, depending on which is closest to the destination.
- Simulate complex BGP routing policies from a transit provider. For example, a transit entry location and failover location can be specified per destination.
- Simulate failover between peering AS's, for example from one single-connection transit provider to another.

For More Information...	See...
Demands	Traffic Demand Modeling chapter
AS's and routing policies	BGP Simulation chapter

Routing with External Endpoints

An external endpoint is a MATE Design object that identifies specific entry (source) or exit (destination) points for demands. These are identified in the External Endpoints table by a name and an optional Tag property.

Each external endpoint consists of one or more members that are defined as nodes, external AS's, or a combination of an external AS and external node. As such, by setting a demand's source or destination to an external endpoint, you can simulate traffic going from multiple sources to a single destination, from a single source to multiple destinations, or multiple sources going to multiple destinations. Because of this flexibility, they are useful for specifying particular secondary entry and exit points in the event of failures.

For More Information...	See...
Create external endpoints and their members	Create External Endpoints and Their Members section

External Endpoint Members

Each member is assigned properties that prioritize traffic entering and exiting that member, in what order to failover to another member, and how to distribute traffic for members of equal priority. These properties are set when you create the member (in its Properties dialog box), and all members are listed in the External Endpoint Members table.

- Member Endpoint—Defines whether the member is a node, AS, or external node via an AS.
- Priority—The sequential order in which external endpoint members are used in the simulations should failures occur.
- Equal priority routing—If members have the same priority, this property identifies how the traffic is distributed.

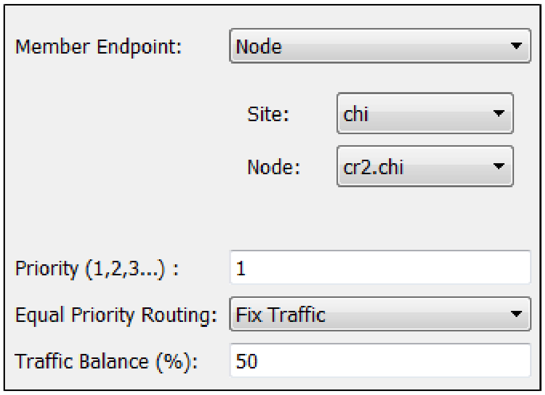

Member Endpoint: Node

Site: chi

Node: cr2.chi

Priority (1,2,3...) : 1

Equal Priority Routing: Fix Traffic

Traffic Balance (%): 50

Note Members within an external endpoint that have the same priority must either all be Shortest Path or none of them be Shortest Path.

- Shortest Path—Use the member that results in the use of shortest path between source and destination.

  If a member is using a shortest path, an “na” shows in the Traffic Balance (%) column in the External Endpoint Members table.
- Fix Traffic—Set the traffic across members of equal priority as defined in the Traffic Balance (%) field.
- Deduce Traffic—Behaves the same as Fix Traffic in that it sets traffic across members of equal priority as defined in the Traffic Balance (%) field. However, upon running Demand Deduction, the Traffic Balance (%) field is updated based on the measured traffic in the network. Note that Demand Deduction only estimates the traffic balances for external endpoints with a priority that is in use in the current no-failure simulation.Thus, Deduce Traffic is usually set to Priority 1.

Routing Simulations

Note Although this section describes demands as being sourced from an external endpoint, the same methodology holds true if a demand’s destination is an external endpoint.

If a demand’s source is defined as an external endpoint, then the following selection of external endpoint members ensues.

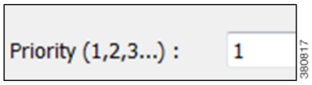

Priority (1,2,3...) : 1

Step 1 Members with the highest priority (lowest number) are used as the demand’s source. For example, if the external endpoint has two members with a priority of 1, the demand is sourced from both members provided they are available.

If one or more of the members are not available, the traffic from the unavailable members is evenly redistributed to the other top priority members.

Step 2 If none of the top priority members are available to source the traffic but there are next-priority members available, step 1 is repeated for the next priority external endpoint members. Only if all members with the same priority fail does the traffic get routed according to the next priority in the sequence.

Note that if a failure occurs that does not affect the external endpoint member’s ability to send or receive traffic, then traffic is rerouted as usual without a need to use the additional members.

If an external endpoint member is an external AS, with or without a node specified, then the routing from or to that member is determined by the BGP routing policy determined by the AS relationships. The distribution of traffic between external endpoints with the same priority is the same as for node members. For more information about AS routing, see the BGP Simulation chapter.

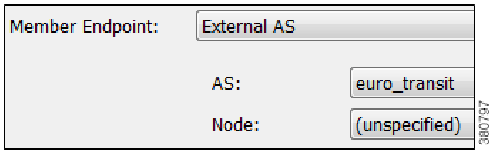

Member Endpoint: External AS

AS: euro_transit

Node: (unspecified)

Traffic Distribution

The traffic distribution through these demands is based on the Equal Priority Routing property, and if applicable, the Traffic Balance (%) property used to define the external endpoint members.

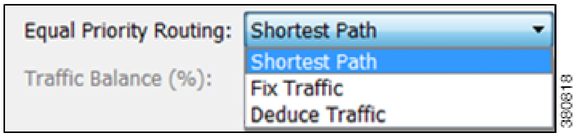

Equal Priority Routing: Shortest Path

Traffic Balance (%):

- If there is only one member and it is defined as Shortest Path, the demand takes the shortest path as defined by the IGP metrics.

  Of the routable demands, if the Traffic Balance (%) values are all empty, the traffic is routed and equally load balanced across the demands with the shortest IGP paths. Note in the case of multiple internal AS's, the shortest IGP route is the shortest route in the first AS the demand enters.
- If multiple members of the same priority are set to Shortest Path, the demand takes the path with the shortest IGP path. If all interfaces between the source members and the destination have the same shortest IGP paths, then the traffic is load balanced equally across them.
- If one or more members of the same priority have their Equal Priority Routing property set to Fix Traffic or Deduce Traffic, the demand traffic is split according to each member’s Traffic Balance (%) value.

  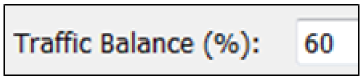

  Traffic Balance (%): 60

  - If the Traffic Balance percentages across sources with the same priority sum to less than 100%, the overall demand traffic is decreased to that percentage.
  - If an external endpoint member of the same priority fails, the traffic on the remaining members increases in proportion, so that the same amount of traffic is still routed.

Example: Node A is failed. Nodes B, C, and D each have a priority of 2 and are each a Fix Traffic type. Their traffic balances are 20%, 20%, and 40%, respectively. The demand has 1,000 Mbps of traffic.

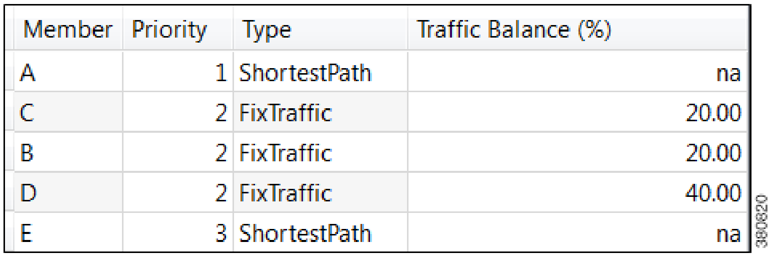

Member	Priority	Type	Traffic Balance (%)
A	1	ShortestPath	na
C	2	FixTraffic	20.00
B	2	FixTraffic	20.00
D	2	FixTraffic	40.00
E	3	ShortestPath	na

380820

- Since Node A failed, the demand routes 200 Mbps traffic through node B, 200 Mbps through C, and 400 Mbps through D, totaling 800 Mbps.
- If Node D fails, then the demand routes 400 Mbps traffic through Node B and 400 through C. If Node B fails too, the entire 800 Mbps is routed through C.
- If all three priority 2 members fail, then 1000 Mbps is routed through node E, which is the priority 3 member.

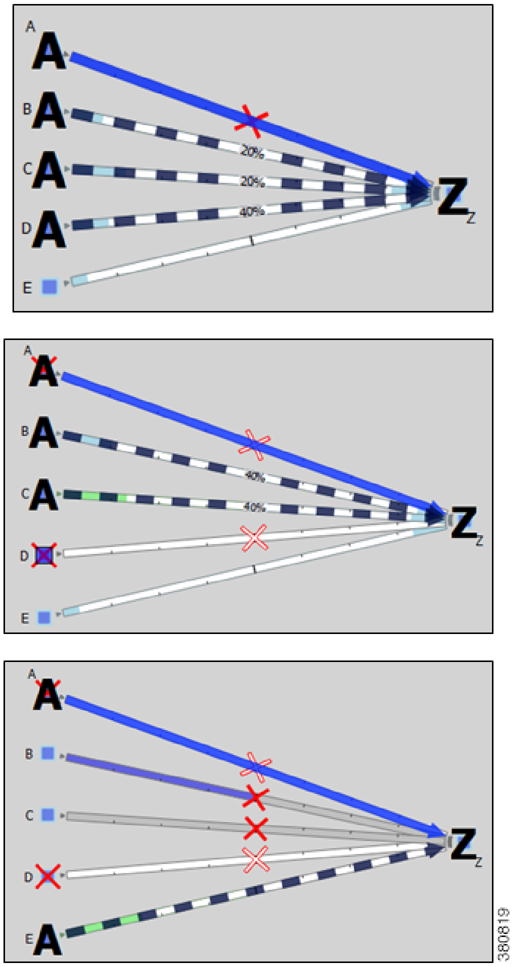

States

External endpoint members can be failed, set to an inactive state, and included in Simulation Analysis to calculate a worst-case failure analysis. Additionally, they can be included in SRLGs, which in turn can also be failed, set to inactive, and included in Simulation Analysis. Table 12-1 shows the icons used for these states.

If the external endpoint member that you fail or set to inactive is an AS, then entry through any of its nodes is not possible. The same is true if an SRLG is failed or set to inactive and if it contains an external endpoint member that is an AS.

Table 12-1 L1 Link Failure and Inactive States

External Endpoint Member	Description
If the external endpoint member is a node, this icon appears under that node in the plot. If the member is an external AS or an external AS via a node, this icon appears under all nodes in the AS.	
✖➔	Failed external endpoint member.
✕➔	External endpoint member is not operational because it is contained in a failed SRLG.
✖➔	External endpoint member is not active.
✕➔	External endpoint member is not operational because it is contained in an SRLG that is not active.

In-Network Routing Examples

Content from Multiple Data Centers

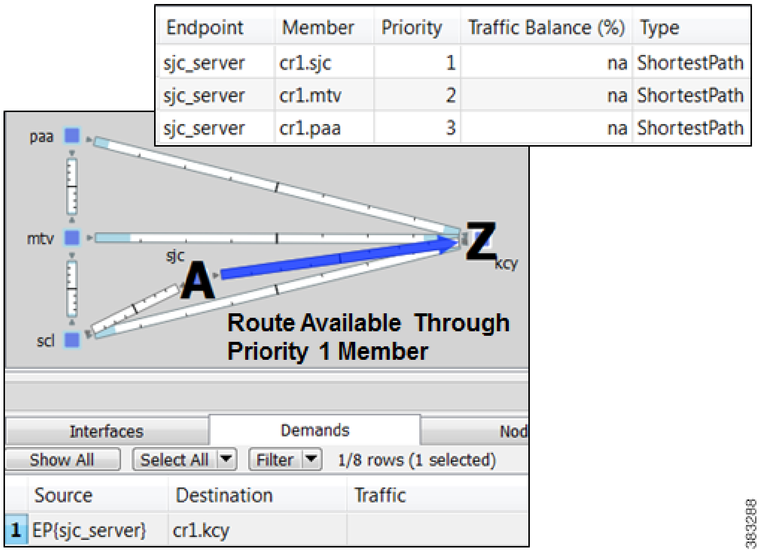

Endpoint	Member	Priority	Traffic Balance (%)	Type
sjc_server	cr1.sjc	1	na	ShortestPath
sjc_server	cr1.mtv	2	na	ShortestPath
sjc_server	cr1.paa	3	na	ShortestPath

Route Available Through Priority 1 Member

	Source	Destination	Traffic
1	EP{sjc_server}	cr1.kcy	

This example demonstrates servers that cache content for customers in Kansas City (cr1.kcy). The primary server is located in San Jose (cr1.sjc), with backup servers in Mountain View (cr1.mtv) and Palo Alto (cr1.paa). Each of the servers is a member of an external endpoint (sjc_server) that is used as the source of the demand to cr1.kcy.

- Member cr1.sjc has a priority of 1, and it takes the shortest path.
- Member cr1.mtv has a priority of 2, and it takes the shortest path.
- Member cr1.paa has a priority of 3, and it takes the shortest path.

Figure 12-1 shows how the failure scenario works.

- When the priority 1 server node fails (cr1.sjc), content continues to be delivered using cr1.mtv as the source.
- If cr1.mtv were also to fail, then cr1.paa would source the traffic to cr1.kcy.

Figure 12-1 Example Failover for Multiple Data Centers

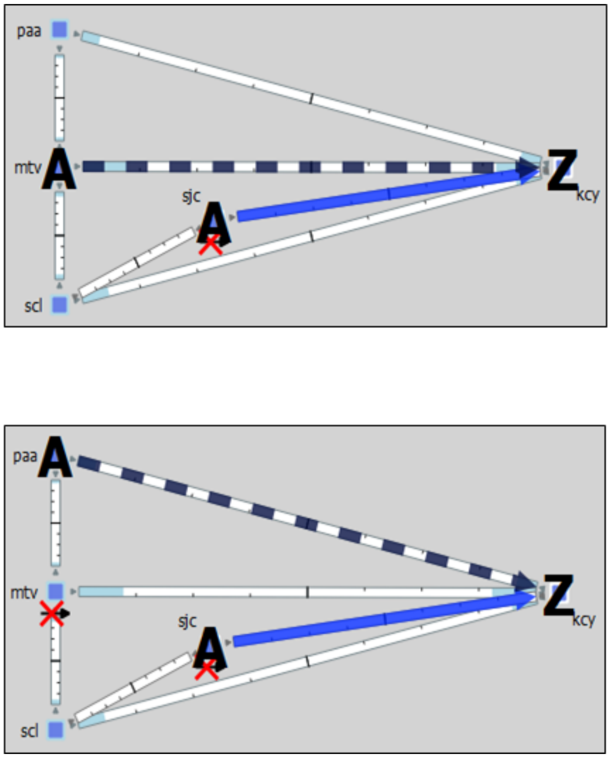

If priority 1 member fails, traffic is sourced from priority 2

If both priority 1 and priority 2 members fail, traffic is routed on the priority 3 member

External Source

One method of modeling an external demand source with multiple entry points is to construct an edge node, connecting it to these entry points, and giving the interfaces equal metrics. Then source the demand from this edge node.

Alternatively, you can construct these multiple entry points using a single external endpoint as the demand source, thus reducing the complexity of the topology.

In this example, a demand has an external source with two entry points that come into the network at Chicago. Figure 12-2 demonstrates two ways to model this source. One uses additional topology (edge node) to source the demand. The other uses a single external endpoint that has two members (cr1.chi and cr2.chi), both set to priority 1 with shortest path traffic.

Figure 12-2 Example External Demand Source

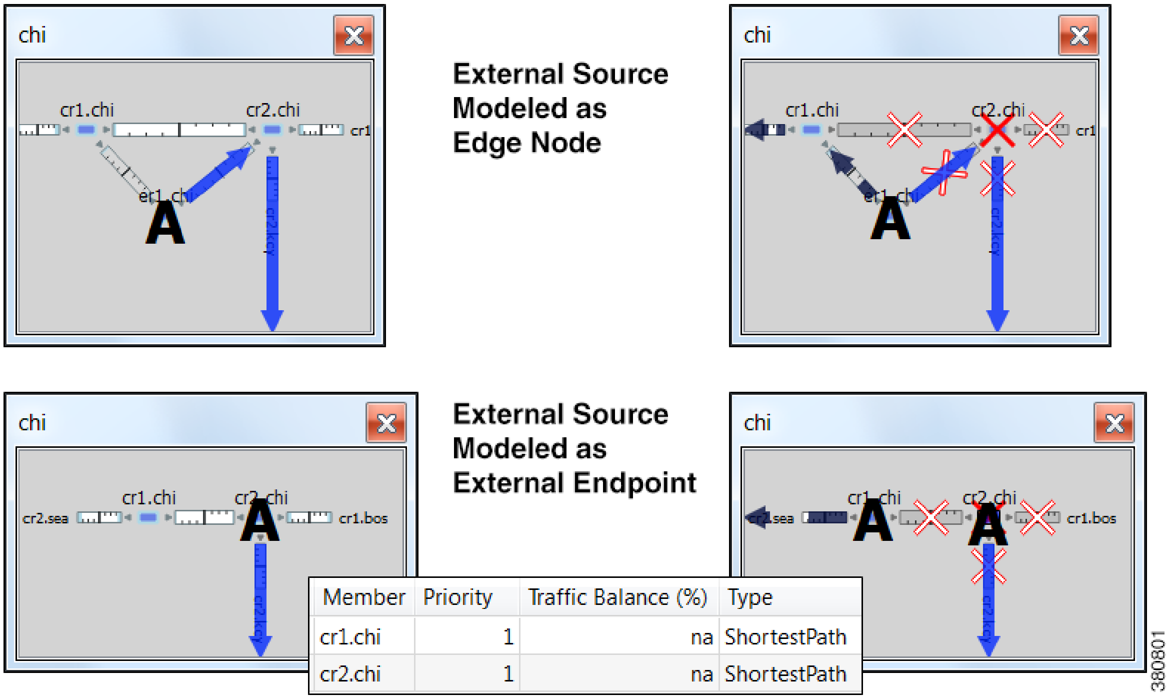

Member	Priority	Traffic Balance (%)	Type
cr1.chi	1	na	ShortestPath
cr2.chi	1	na	ShortestPath

Intra-AS Routing Examples

Failover to Nodes in External AS's

This example demonstrates a content provider (Acme) who has most of its incoming Internet traffic sourced through a transit provider, and this traffic is then distributed to Acme's customers. Acme has identified specific policies for each market using communities, and these policies determine, per location, which transit provider peering points are used for primary routing and failover. The network is set up as follows (Figure 12-3).

- The San Diego, Boise, Cheyenne, and Seattle nodes each belong to the same transit provider, AS 2.
- The remaining nodes belong to Acme's internal AS (AS 1), and these are Acme's customers.
- The San Diego node (pr1.sdg) in AS 2 is sending traffic to three key customers located in Phoenix (cr1.pho), Salt Lake City (cr1.slc), and Denver (cr1.dnv).

Figure 12-3 Example Transit Peer AS 2 San Diego Sends Traffic to Customers in an Internal AS

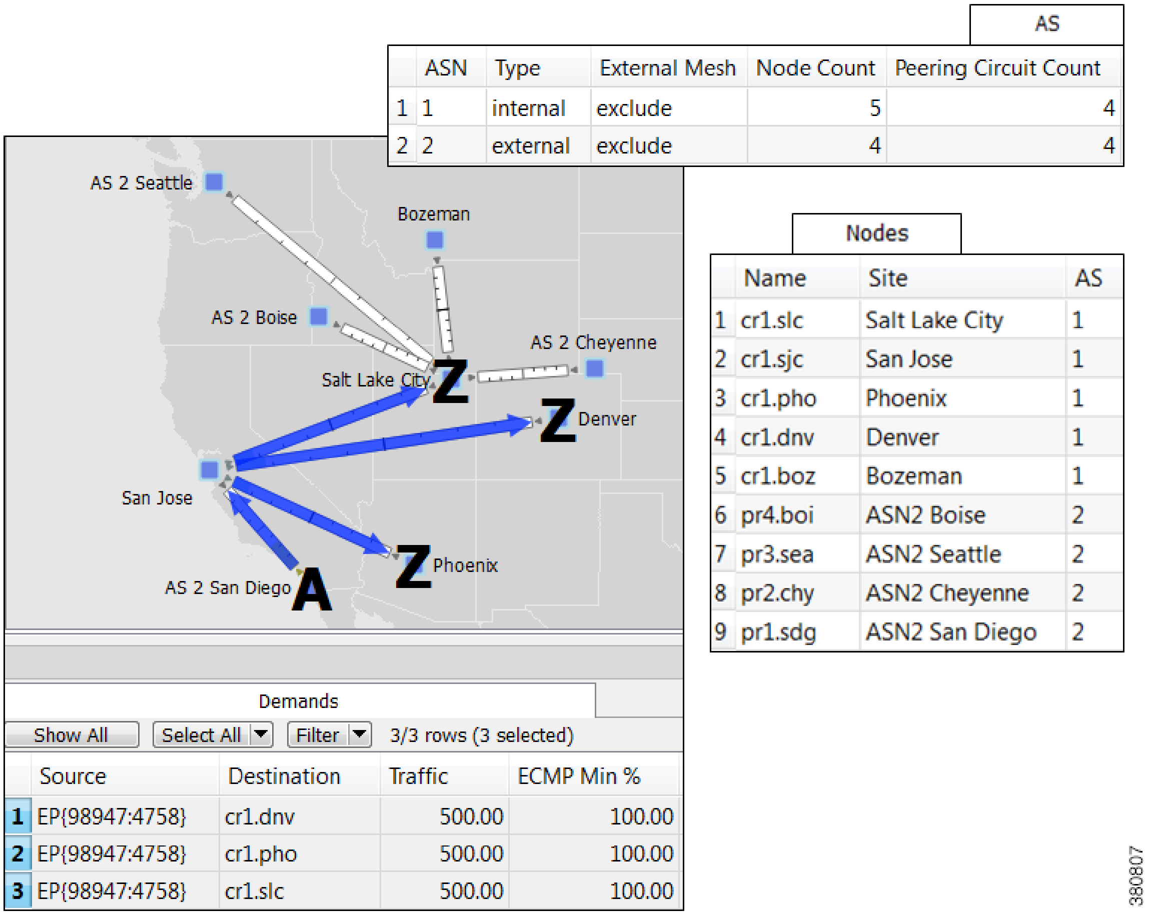

AS

	ASN	Type	External Mesh	Node Count	Peering Circuit Count
1	1	internal	exclude	5	4
2	2	external	exclude	4	4

Nodes

	Name	Site	AS
1	cr1.slc	Salt Lake City	1
2	cr1.sjc	San Jose	1
3	cr1.pho	Phoenix	1
4	cr1.dnv	Denver	1
5	cr1.boz	Bozeman	1
6	pr4.boi	ASN2 Boise	2
7	pr3.sea	ASN2 Seattle	2
8	pr2.chy	ASN2 Cheyenne	2
9	pr1.sdg	ASN2 San Diego	2

Demands

3/3 rows (3 selected)

	Source	Destination	Traffic	ECMP Min %
1	EP{98947:4758}	cr1.dnv	500.00	100.00
2	EP{98947:4758}	cr1.pho	500.00	100.00
3	EP{98947:4758}	cr1.slc	500.00	100.00

In MATE Design, Acme uses an external endpoint named after community 98947:4758. It contains the following members, all of which are in the same AS 2 transit provider.

- pr1.sdg set to priority 1 with shortest path traffic
- pr2.chy set to priority 2 with shortest path traffic
- AS 2, which means the remaining nodes in the AS, set to priority 3 with shortest path traffic

As Figure 12-3 shows, multiple demands are created using this external endpoint (98947:4758) as the source, and each sending traffic to a different customer destination. Figure 12-4 focuses on the traffic destined for Denver (cr1.dnv), showing that the failover scenario behaves as follows.

- If a failure prevents the peer pr1.sdg node from sending traffic to cr1.pho, a failover occurs to pr2.chy within that same peer so that there is no service disruption.
- If failures occur such that neither pr1.sdc, nor pr2.chy can send traffic to cr1.dnv, the remaining nodes in the AS source the traffic. Although this is set to use the shortest path, since there are multiple paths, the traffic is distributed evenly.

Note, the same result could have been achieved by setting the priority 1 to an interface going from San Jose to pr1.sdg and by setting the priority 2 to an interface going from Salt Lake City to pr2.chy.

Figure 12-4 Example Failover Within the Same External AS

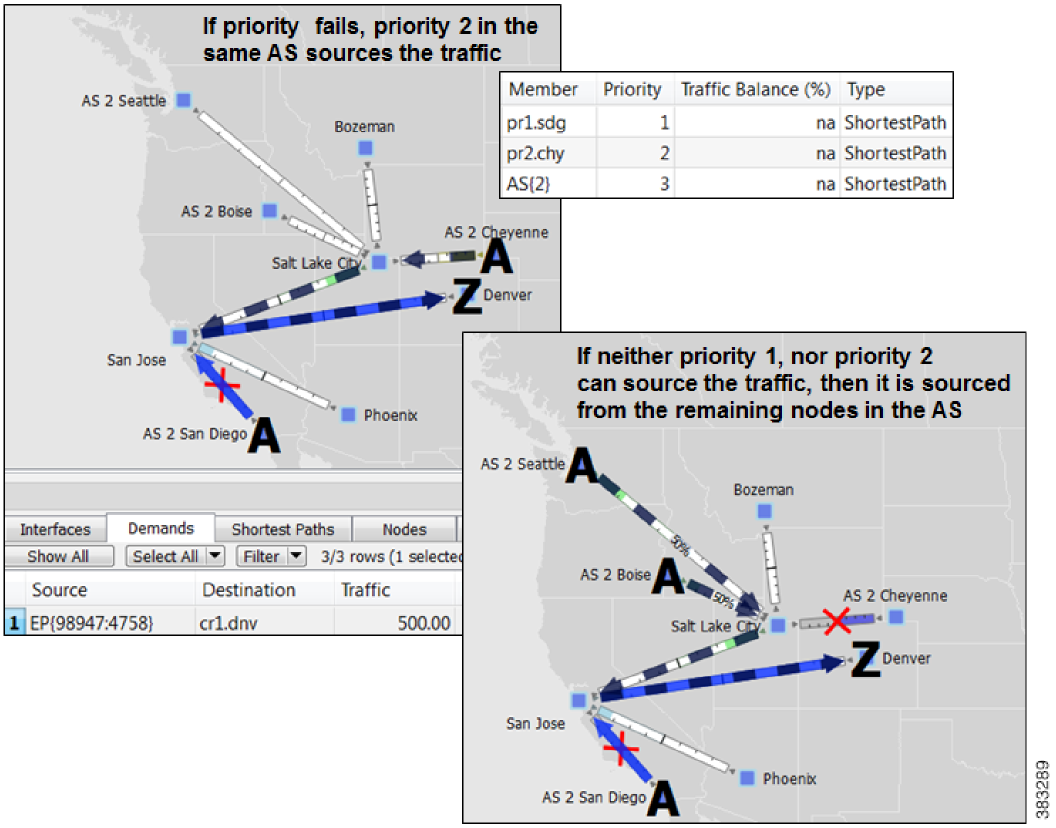

If priority fails, priority 2 in the same AS sources the traffic

Member	Priority	Traffic Balance (%)	Type
pr1.sdg	1	na	ShortestPath
pr2.chy	2	na	ShortestPath
AS{2}	3	na	ShortestPath

	Source	Destination	Traffic
1	EP{98947:4758}	cr1.dnv	500.00

If neither priority 1, nor priority 2 can source the traffic, then it is sourced from the remaining nodes in the AS

Failover from an External AS to Another

In this example, a small ISP (AcmeNet) buys transit from two other service providers. AcmeNet peers with AS 2 in Phoenix and AS 3 in Denver. AcmeNet has a customer base in Phoenix (cr1.pho). The customer has a preferred access to the Internet through the San Francisco (pr1.sfo) in AS 2 (Figure 12-5), and the secondary access is through Denver (pr2.dnv) in AS 3 (Figure 12-6). In MATE Design, AcmeNet simulates this network by creating an external endpoint (Internet) containing two different external AS's as members.

- AS 2 is priority 1 with shortest path
- AS 3 is priority 2 with shortest path

The demand is sourced from this external endpoint named Internet.

Figure 12-5 Example Transit Peer ASN2 San Francisco Sends Traffic to Customers in an Internal AS

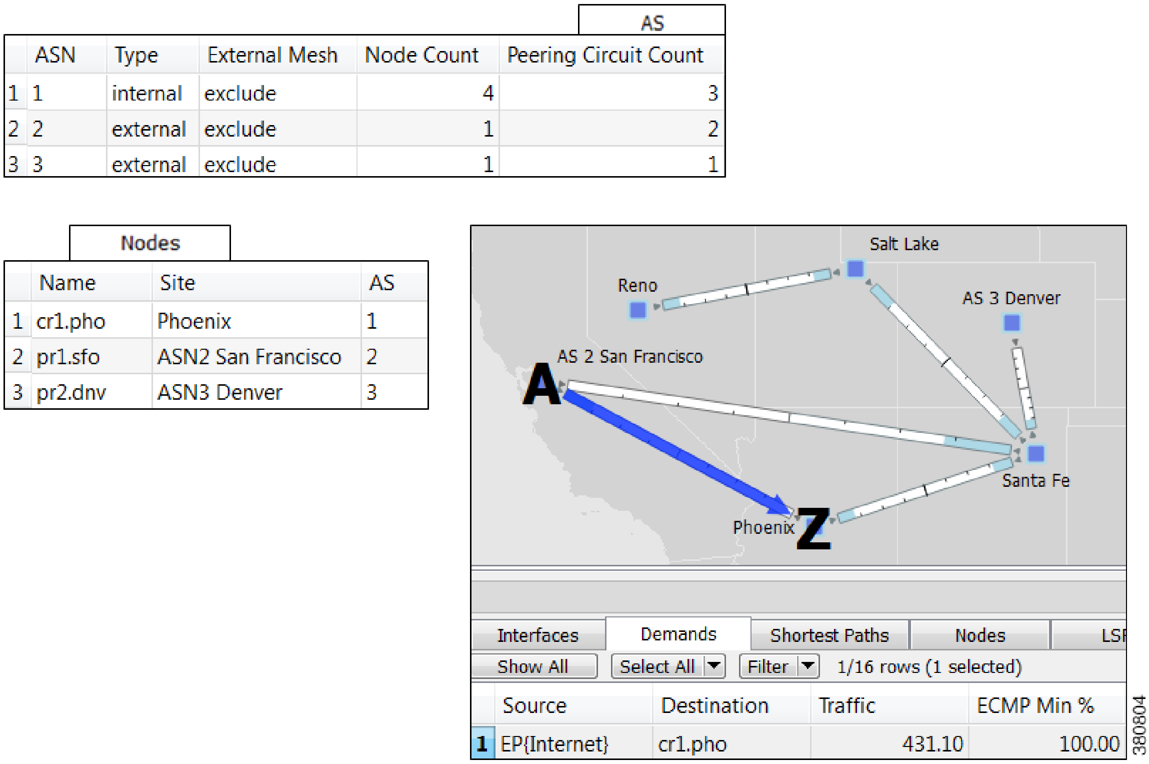

AS

	ASN	Type	External Mesh	Node Count	Peering Circuit Count
1	1	internal	exclude	4	3
2	2	external	exclude	1	2
3	3	external	exclude	1	1

Nodes

	Name	Site	AS
1	cr1.pho	Phoenix	1
2	pr1.sfo	ASN2 San Francisco	2
3	pr2.dnv	ASN3 Denver	3

	Source	Destination	Traffic	ECMP Min %
1	EP{Internet}	cr1.pho	431.10	100.00

Figure 12-6 Example Failover to a Different External AS

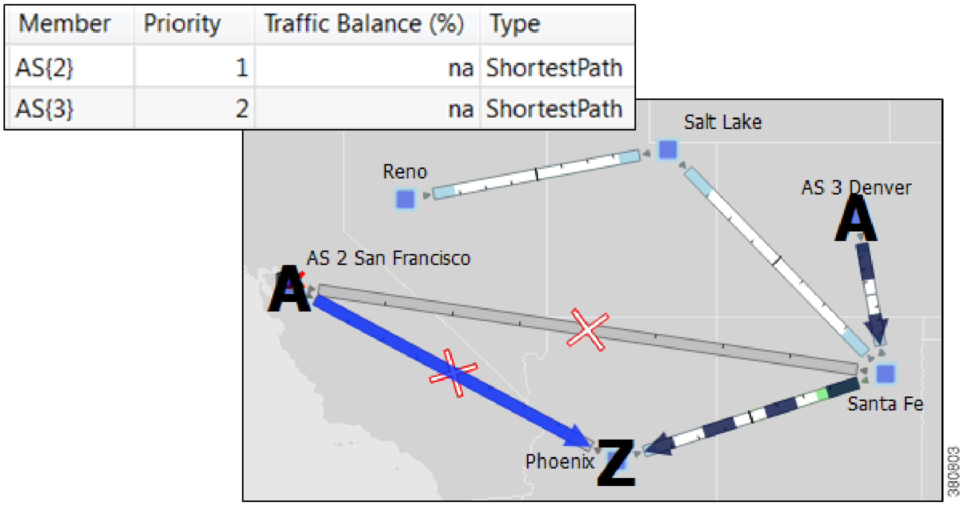

Member	Priority	Traffic Balance (%)	Type
AS{2}	1	na	ShortestPath
AS{3}	2	na	ShortestPath

Create External Endpoints and Their Members

The easiest and recommended method of creating external endpoint members is to do so while creating its associated external endpoint, as follows.

Step 1 Either right-click in an empty area and select New->Demands->External Endpoints, or select the Insert->Demands->External Endpoints menu. An external endpoint Properties dialog box opens.

To edit existing external endpoints, double-click on it from the External Endpoints table.

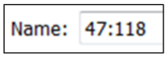
Name: 47:118

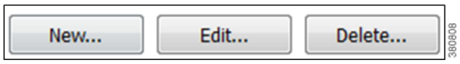

Step 2 Enter a unique external endpoint name in the Name field.

Step 3 To add a new member, click the New button.

To edit an existing member, select it from the Members table and click Edit.

An external endpoint member Properties dialog box appears.

a. Select the member type as node or external AS.

- For nodes, select its site and node name.
- For interfaces, select its site, node, and the interface. Using an interface enables you to specify the exact interface on which the demand traffic is going into or out of a node.
- For AS's, select its AS name, and select either the "unspecified" option or node name within the AS. The "unspecified" option routes traffic evenly throughout the entire AS.

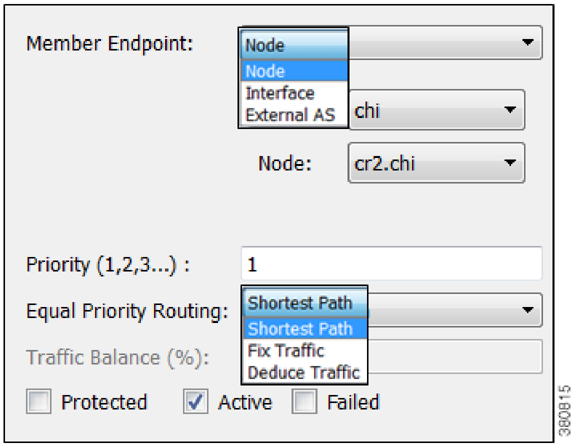
Member Endpoint: Node
chi
Node: cr2.chi
Priority (1,2,3...) : 1
Equal Priority Routing: Shortest Path
Traffic Balance (%):
□ Protected ☑ Active □ Failed

Note For information on how the following external endpoint member fields affect routing, see the Routing with External Endpoints section.

b. Enter a priority number to determine the sequence in which the member is used with other members in this external endpoint. Multiple members can have the same priority.

c. Select whether to use a shortest path route, or to direct a specified percentage of traffic across members that have the same priority (Fix Traffic or Deduce Traffic).

d. If balancing traffic, enter the percentage of traffic this member should route.

If you set Deduce Traffic, this value is updated by Demand Deduction provided there are no members with a higher priority that are in use under normal (non-failure) operation.

e. Click OK in the external endpoint member Properties dialog box.

Step 4 Click OK in the external endpoint Properties dialog box.

CHAPTER 13

VPN Simulation

The MATE Design virtual private network (VPN) model is a representation of a virtual subnetwork within the plan file. Viewing and simulating VPN within MATE Design facilitates numerous network tasks and can answer many questions. The following are a few.

- Which VPNs are on my network? Where and how are they configured?
- Which VPNs are using congested interfaces?
- Which VPNs will experience congestion under any of a given list of failure scenarios?
- Which failures scenarios cause the worst-case congestion or latency for a VPN?

There are many varieties of VPNs. For example, there are Layer 2 (L2) VPNs and Layer 3 (L3) VPNs, each having different categories within it, and there are vendor-specific VPN implementations as well. Each VPN type has its own specific configuration and terminology. The MATE Design VPN model supports a number of these VPN types based on either route-target or full-mesh connectivity.

VPN Model

VPN Objects

Object	Description	Examples
VPNs	A set of VPN nodes that can exchange data with each other.	• Layer 2 VPN: The VPN represents an individual VPLS containing VSIs. • Layer 3 VPN: The VPN represents sets of VRFs associated with a set of VPN nodes that forward traffic between themselves. Often, this set of VRFs signify a single customer or service.
VPN nodes	Connection points in a VPN. They exist on standard nodes, and each node can contain multiple VPN nodes. A VPN node can be in only one VPN.	• Layer 2 VPN: The VPN node represents the VSIs (virtual switch instances) configured on each router. • Layer 3 VPN: The VPN node represents the VRF instances configured on each router.

VPN Topology and Connectivity

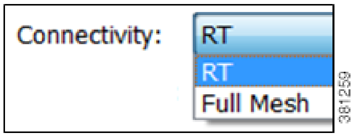

MATE Design VPN topologies route connections are established through route targets (RTs) or through a full mesh of VPN nodes. This Connectivity property is set in the VPN Properties dialog box.

Knowing a VPN's topology and connectivity enables MATE Design to calculate which demands between VPN nodes carry traffic for a particular VPN, and thus which interfaces carry traffic for that VPN. In turn, MATE Design can calculate the vulnerability of a VPN to certain failure and congestion scenarios.

A demand is associated with a VPN, meaning it carries traffic for that VPN, if each of the following points are true.

- The two VPN nodes are in the same VPN.
- The demand is in the same service class as the VPN.
- Only for VPNs with RT connectivity, the RT Export property of one VPN node must match the RT Import property of another VPN node.

Once demands are associated with the VPN, this configuration simulates the associated access circuits exchanging traffic as if they were on the same LAN.

Note that a demand associated with a VPN can additionally contain other traffic that is for that VPN.

Connectivity	Description
Full mesh	Full-mesh connectivity is a complete mesh of connections between VPN nodes in a VPN such that they can all communicate with one each other. This connectivity is typical in a VPLS, where all VSIs identify one another based on a common AGI.
Route targets	Route targets model the more complex connectivity used in Layer 3 VPNs, such as hub-and-spoke networks. Here, the VRFs exchange data with one another based on the matching of RT Export and RT Import properties set for each VPN node. Having an import/export pair does not create bidirectional communication. Rather, traffic flows in the opposite direction of the routed advertisements. For example, if node A's RT Import matches node B's RT Export, then traffic can flow from node A to B. For traffic to flow from node B back to node A, node B must have an RT Import that matches an RT Export of node A. It is this combination of matching imported and exported RTs that define which VPN nodes can exchange data, and it is the VPN name that identifies the VPN itself. 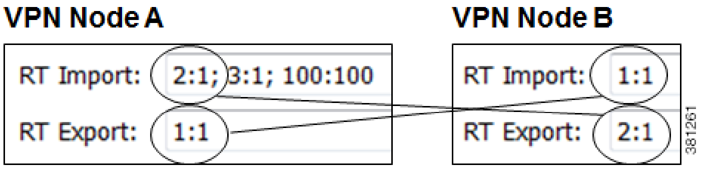 VPN Node A RT Import: 2:1; 3:1; 100:100 RT Export: 1:1 VPN Node B RT Import: 1:1 RT Export: 2:1 381261

VPNs

Each VPN consists of a set of VPN nodes that can exchange data within it. VPNs have three key properties that uniquely identify them and define how the traffic within them is routed.

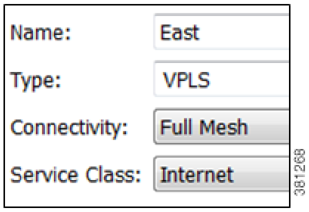

Name: East
Type: VPLS
Connectivity: Full Mesh
Service Class: Internet

- Name—Unique name of the VPN.
- Type—The type of VPN. You can select from the defaults, which are VPWS, VPLS, or L3VPN, or you can enter a string value to create a new one. Once entered, the new VPN type appears in the drop-down list and is available for other VPNs and VPN nodes.
- Connectivity—Determines how MATE Design calculates connectivity and associated demands for VPNs.
  - Full Mesh—Connectivity is between all nodes in the VPN. MATE Design ignores the RT Import and RT Export properties of the VPN nodes.
  - RT—Connectivity is based on the RT Import and RT Export properties of its VPN nodes.

VPNs Table

The VPNs table lists the VPN properties described above, its associated service class, traffic, and the number of VPN nodes within that VPN (Table 13-1). For information on QoS measurements, see the Quality of Service Simulation chapter. For information on the Worst-Case columns not listed here, see Table 13-3.

Note Since the traffic and QoS calculations are based on all interfaces within the VPN for the service class specified for that VPN, the plot view might differ from the table. For example, the plot view could show Internet traffic while a VPN carrying voice traffic is selected.

Table 13-1 VPNs Table Columns for Normal Operation

Columns	Description
	Note: All traffic and QoS violations are based on traffic carried on all interfaces used by the VPN for the service class defined for that VPN.
Service Class	Service class associated with this VPN. All values within the table are associated with this service class.
Num Nodes	Number of VPN nodes in this VPN.
Util Meas	The maximum measured utilization of all interfaces used by this VPN.
Util Sim	The maximum simulated utilization of all interfaces used by this VPN.
Total Src Traff Meas	Total amount of measured source traffic on this VPN.
Total Dest Traff Meas	Total amount of measured destination traffic on this VPN.
QoS Violation Sim	Maximum QoS violation under normal operations for all simulated traffic for all interfaces used by this VPN. If the number is positive, there is a violation.
QoS Violation Sim (%)	QoS violation as a percent of the total simulated interface capacity.

Columns	Description
QoS Violation Meas	Maximum QoS violation under normal operations for all measured traffic for all interfaces used by this VPN. If the number is positive, there is a violation.
QoS Violation Meas (%)	QoS violation as a percent of the total measured interface capacity.
Latency	Maximum latency of all demands used by this VPN.
Tags	User-defined identifiers that makes it easy to group VPNs.

VPNs are not selectable from the network plot; you can only select and filter to VPNs through tables. When selected, all VPN nodes within the VPN are highlighted in the plot (Figure 13-1).

For More Information ...	See ...
Create VPNs based on existing VPN nodes	Create VPNs from Existing VPN Nodes
Create empty VPNs to which you can later add VPN nodes	Create VPNs
Add VPN nodes to VPNs	Add VPN Nodes to VPNs

Identify Interfaces Used by VPNs

To view which interfaces are associated with a VPN, right-click on a VPN in the VPNs table and select Filter to Interfaces from the context menu. his is useful for viewing the VPN topology in the network plot. If you then select all of these filtered interfaces, you can see the VPN outlined in the network plot.

To view which VPNs are associated with an interface, right-click on an interface in the Interfaces table and select Filter to VPNs from the context menu. This is useful for determining which VPNs would be affected should a circuit fail or go down for maintenance.

Note Utilization measurements might be different between the tables since the VPN table calculates measurements only for the service class associated with that VPN.

VPN Nodes

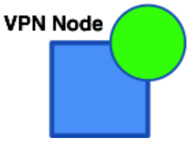

VPN nodes are defined by several properties that together, determine which VPNs the nodes belong to and how the demands are routed. The following are required properties.

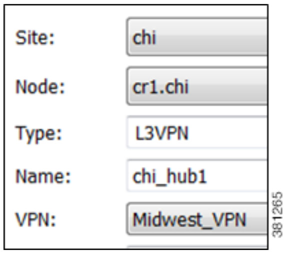

Site: chi
Node: cr1.chi
Type: L3VPN
Name: chi_hub1
VPN: Midwest_VPN

- Node—Name of the node on which the VPN node resides. This node name corresponds with one in the Nodes table.
- Type—The type of VPN. You can select from the defaults, which are VPWS, VPLS, or L3VPN, or you can enter a string value to create a new one. Once entered, the new VPN type appears in the drop-down list and is available for other VPN nodes and VPNs.
- Name—Name of the VPN node.
- VPN—Name of the VPN in which this VPN node resides. The drop-down lists shows existing VPNs of the same type set in the Type field. You can create a VPN node without setting its VPN, but without this selected, the VPN node is not included in simulations as a member of any VPN.

To simulate RT connectivity, you must set the VPN Connectivity property to RT and then set the RT Import and RT Export properties on the individual VPN nodes within it.

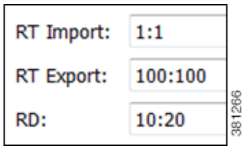

RT Import: 1:1
RT Export: 100:100
RD: 10:20

- RT Import and RT Export—The pairing of RT values identify which VPN nodes connect with each other. For more information, see the VPN Topology and Connectivity section.
- Optional: RD—Route distinguisher (RD) uniquely identifies routes within a VRF as belonging to one VPN or another, thus enabling duplicate routes to be unique within a global routing table.

VPN Nodes Table

The VPN Nodes table lists the VPN node properties described above, as well as columns that identify the VPN nodes' relationship within the VPN and its traffic. Table 13-2 describes the columns.

Table 13-2 VPN Nodes Table

Columns	Description
Total Connect	Number of VPN nodes that are connected to this VPN node as defined by the RT Import and RT Export pairings. These may or may not be in the same VPN.
VPN Connect	Number of VPN nodes that are connected to this VPN node and are in the same VPN as defined by the VPN column.
Num VPN Nodes	Number of nodes in the VPN that this VPN node belongs to as defined by the VPN column. This value is "na" if the VPN node does not belong to a VPN.
Src Traff Meas	Total amount of measured traffic entering the VPN at this node (source traffic).
Dest Traff Meas	Total amount of measured traffic leaving the VPN at this node (destination traffic).
Tags	User-defined identifier that makes it easy to group VPN nodes into a single VPN. If you give a VPN node a tag, then when you create a VPN later, you can identify its VPN nodes using tags.

VPN nodes are not selectable from the network plot; you can only select and filter to them through tables.

Once selected from the VPN Nodes or VPNs table, the associated site and the nodes within that site appear with a green circle on it (Figure 13-1).

Figure 13-1 VPN Nodes within a VPN

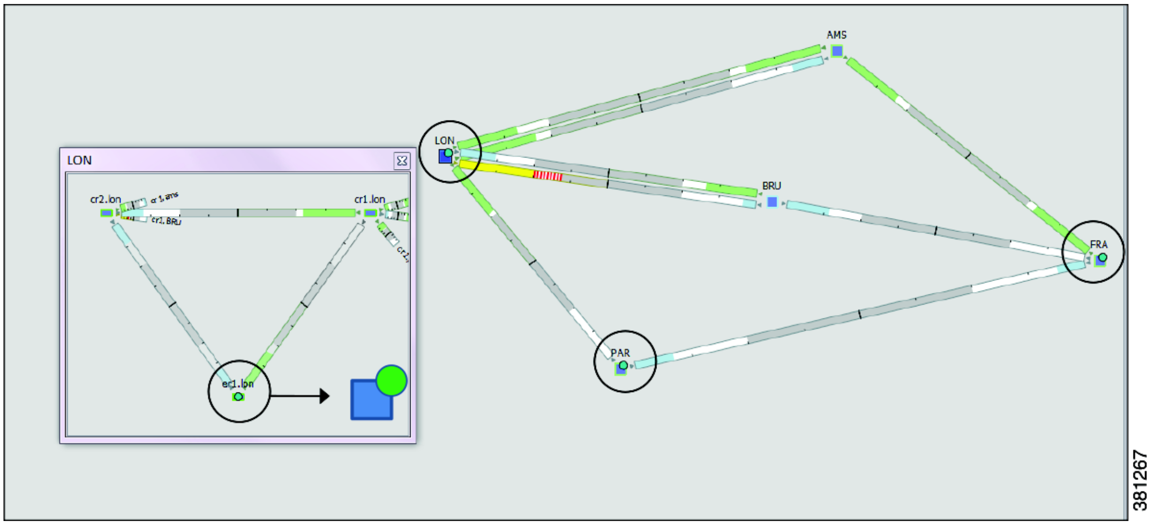

For More Information ...	See ...
Create VPN nodes	Create VPN Nodes
Add VPN nodes to VPNs	Add VPN Nodes to VPNs

Layer 3 VPN Example

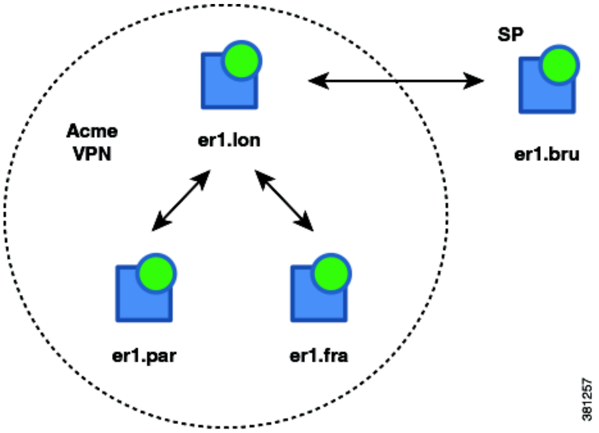

This example illustrates a scenario where the Acme manufacturing company has three offices, but permits the two branch (er1.par and er1.fra) offices to exchange data only with headquarters (er1.lon).

Additionally, headquarters communicates with an SP VPN node (er1.bru) that is not in the Acme VPN. Figure 13-2 shows the footprint of the Acme VPN and the RTs set for all VPN nodes in this example.

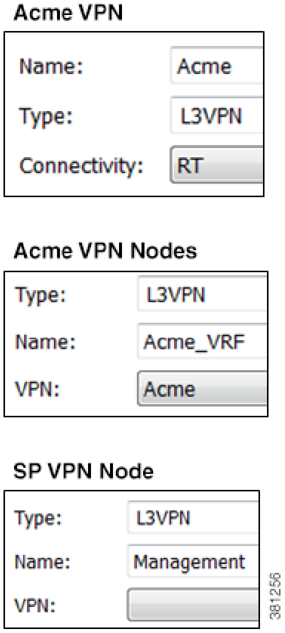

Acme VPN

Name: Acme

Type: L3VPN

Connectivity: RT

Acme VPN Nodes

Type: L3VPN

Name: Acme_VRF

VPN: Acme

SP VPN Node

Type: L3VPN

Name: Management

VPN:

- The VPN is named Acme, and it is set to a Connectivity of RT and a Type of L3VPN.

  In turn, each branch office is set to the Acme VPN, with a Type of L3VPN.
- To exchange data with two other VPN nodes in the Acme VPN, headquarters (er1.lon) imports the offices' exported route targets of 2:1 (er1.par) and 3:1 (er1.fra).
- In turn, headquarters (er1.lon) exports a route target of 1:1.

  All three of these other VPN nodes import it (both offices and the SP VPN node).

  Since the SP VPN node (er1.bru) is not in the Acme VPN, its communication with er1.lon is not within the context of that VPN.

The VPN footprint in Figure 13-2 shows that if the circuit between er1.fra and er1.bru became congested or failed, the VPN would be impacted. A failure of the circuit between the two branch offices however, would not impact it. This failure is illustrated in Figure 13-3, which shows that none of the demands associated with the VPN are rerouted.

Figure 13-2 Example RT Connectivity and Acme VPN Footprint

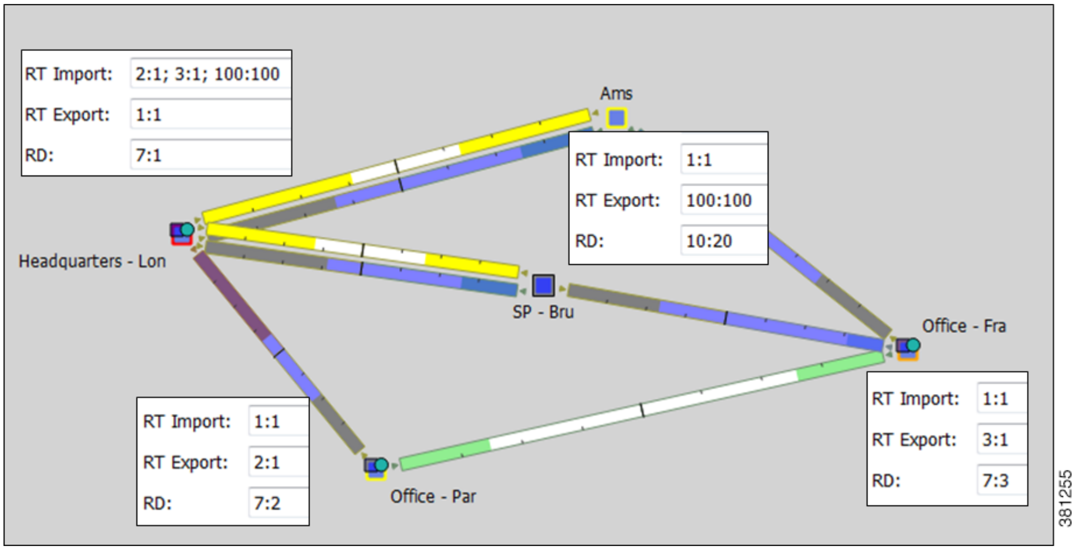

RT Import: 2:1; 3:1; 100:100

RT Export: 1:1

RD: 7:1

RT Import: 1:1

RT Export: 100:100

RD: 10:20

RT Import: 1:1

RT Export: 2:1

RD: 7:2

RT Import: 1:1

RT Export: 3:1

RD: 7:3

Figure 13-3 Example Failure Between Branch Offices in the Acme VPN

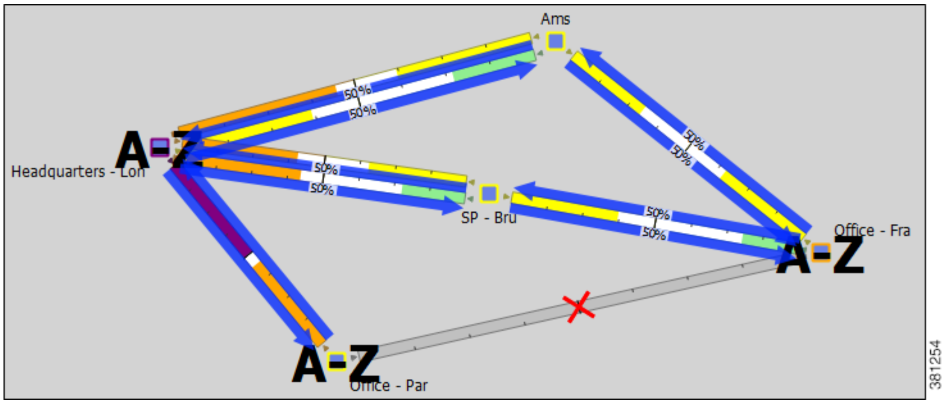

For this example, Figure 13-4 illustrates the filtering of VPN nodes to its associated Acme VPN and the filtering of the Acme VPN to its associated demand traffic. It also shows the calculations of the Total Connect and VPN Connect columns in the VPN Nodes table.

- The Total Connect for the VPN node residing on er1.lon, headquarters is highest because it exchanges data with three other VPN nodes.

  Each of the offices and the service provider VPN node have only 1 in the Total Connect column because they each exchange data only with (have RT pairings with) headquarters.

- The VPN Connect for the VPN node residing on er1.lon, headquarters is highest because it exchanges data with and is in the same VPN as the two offices; all three VPN nodes share the same VPN name.

  Each of the offices have 1 in the VPN Connect column because they are communicating with only one VPN node in the same VPN.

  The service provider VPN node (er1.bru) has 0 VPN Connects because it does not reside in a defined VPN.

Figure 13-4 VPN Nodes Filtered to Acme VPN, and Acme VPN Filtered to Demands

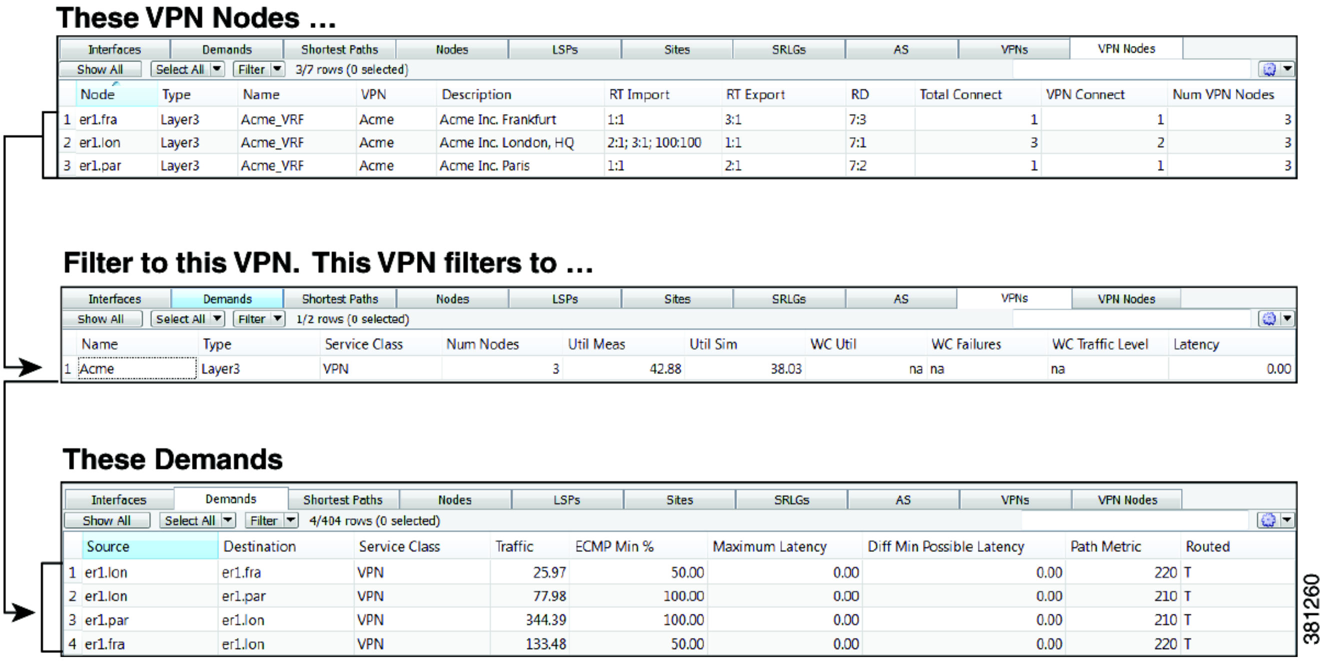

	Node	Type	Name	VPN	Description	RT Import	RT Export	RD	Total Connect	VPN Connect	Num VPN Nodes
1	er1.fra	Layer3	Acme_VRF	Acme	Acme Inc. Frankfurt	1:1	3:1	7:3	1	1	3
2	er1.lon	Layer3	Acme_VRF	Acme	Acme Inc. London, HQ	2:1; 3:1; 100:100	1:1	7:1	3	2	3
3	er1.par	Layer3	Acme_VRF	Acme	Acme Inc. Paris	1:1	2:1	7:2	1	1	3

	Name	Type	Service Class	Num Nodes	Util Meas	Util Sim	WC Util	WC Failures	WC Traffic Level	Latency
1	Acme	Layer3	VPN	3	42.88	38.03	na	na	na	0.00

	Source	Destination	Service Class	Traffic	ECMP Min %	Maximum Latency	Diff Min Possible Latency	Path Metric	Routed
1	er1.lon	er1.fra	VPN	25.97	50.00	0.00	0.00	220	T
2	er1.lon	er1.par	VPN	77.98	100.00	0.00	0.00	210	T
3	er1.par	er1.lon	VPN	344.39	100.00	0.00	0.00	210	T
4	er1.fra	er1.lon	VPN	133.48	50.00	0.00	0.00	220	T

VPN Simulation Analysis

When you run a Simulation Analysis, you have the option to record worst-case utilization and latency for VPNs in the VPNs table (Table 13-3). You can then right-click on a VPN to fail it to its worst-case utilization or to fail it to its worst-case latency.

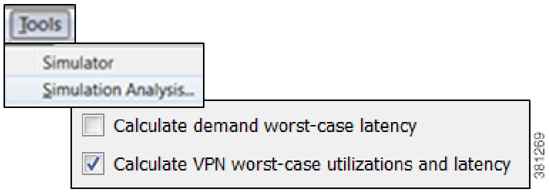

□ Calculate demand worst-case latency
☑ Calculate VPN worst-case utilizations and latency

For information on running Simulation Analysis and failing VPNs to their worst-case utilization, see the Simulation Analysis chapter.

Table 13-3 Simulation Analysis Columns in the VPNs Table

Columns	Description
	Note: All calculations are based on traffic carried on all interfaces used by the VPN for the service class defined for that VPN.
WC Util	Worst-case VPN utilization over all failure scenarios.
WC Failures	Failures causing the worst-case utilization of the VPN.
WC Traffic Level	Traffic level causing the utilization of the interface identified in the WC Util column.
WC QoS Violation	Highest worst-case QoS violation for all interfaces used by this VPN. A QoS violation is equal to the worst-case traffic minus the worst-case capacity permitted (worst-case QoS bound).

Columns	Description
WC QoS Violation (%)	Highest worst-case QoS violation for all interfaces in this VPN expressed as a percentage of total capacity.
WC Latency	Maximum VPN latency over failure scenarios considered.
WC Latency Failures	Failures causing the worst-case VPN latency.

Create VPN Nodes

Step 1 Either right-click in an empty plot area and select New->VPN->VPN Node from the context menu, or select the Insert->VPN->VPN Node. A properties dialog box appears.

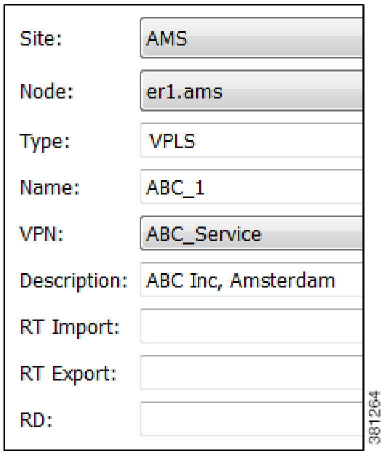

Site: AMS
Node: er1.ams
Type: VPLS
Name: ABC_1
VPN: ABC_Service
Description: ABC Inc, Amsterdam
RT Import:
RT Export:
RD:

Step 2 In the Site and Name fields respectively, select the site in which the VPN node will exist, and select the node on which the VPN node is being configured.

Step 3 Either select a VPN type or enter a string value for a new one. The defaults are VPWS, VPLS, and L3VPN.

Step 4 In the Name field, enter the name of the VPN node. This does not have to be unique.

Step 5 In the VPN drop-down list, select the VPN to which you are adding this VPN node. If you do not see the VPN that you expect to see, verify you correctly selected the type.

Step 6 Optional: Enter a description that further identifies the VPN node. A customer name, for example, might be helpful.

Step 7 If the Connectivity for the VPN is RT, enter the applicable route targets in the RT Import and RT Export fields. All VPN nodes with the same import RT as another VPN node’s export RT can receive traffic from that VPN node. Those VPN nodes with the same export RT as another VPN node’s import RT can send traffic to that VPN node. For more information, see the VPN Topology and Connectivity section.

Step 8 Optional: Enter a route distinguisher in the RD field.

Step 9 Click. OK.

Create VPNs

You can create VPNs from existing VPN nodes or you can create VPNs and then later add VPN nodes with them.

Create VPNs from Existing VPN Nodes

When you create VPNs from existing VPN nodes, all VPN nodes are assigned to these newly created VPNs and the existing VPNs become empty. This is because VPN nodes can belong to only one VPN at a time.

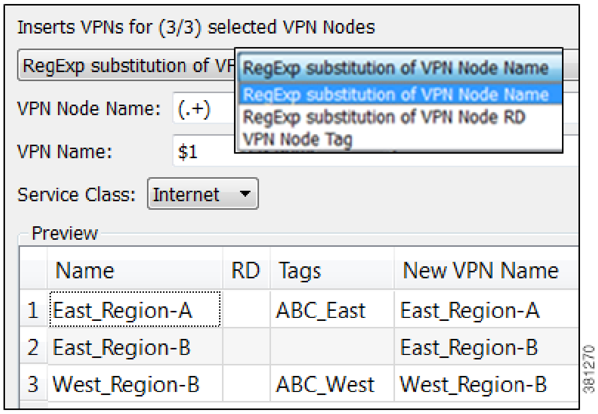
RegExp substitution of VPN Node Name
RegExp substitution of VPN Node Name
RegExp substitution of VPN Node RD
VPN Node Tag
381270

Inserts VPNs for (3/3) selected VPN Nodes

VPN Node Name: (.+)

VPN Name: $1

Service Class: Internet

Preview

	Name	RD	Tags	New VPN Name
1	East_Region-A		ABC_East	East_Region-A
2	East_Region-B			East_Region-B
3	West_Region-B		ABC_West	West_Region-B

Step 1 If you are creating a VPN for specific nodes, select VPN nodes from the VPN Nodes table.

Step 2 Either right-click in an empty plot area and select New->VPN->VPNs from VPN Nodes from the context menu, or select the Insert->VPN->VPNs from VPN Nodes. A properties dialog box appears.

Step 3 In the drop-down list, select the method for creating the VPN: VPN node name, RD, or VPN node tag.

Step 4 If applicable, enter the VPN node name or VPN node RD, and enter the VPN name. These two fields work together to create and name the VPN. Both fields use regular expressions. The $ in the VPN Name field identifies which parenthetical expression in the VPN Node Name or VPN Node RD field to use. For instance, $2 means use the second set of parenthesis from which to create the VPN name.

- The default is a regular expression that matches the entire VPN node name and to create one VPN for each unique VPN node name. That is, the default in VPN Node Name is (.+) and the default VPN Name is $1, which creates a VPN with a name that is identical to each VPN node (or all VPN nodes if none are selected).

  If your convention is to use the same VRF name or the same service ID for every VPN node, this default works well. If, however, the VPN name is encoded in the VRF name or service ID, use a regular expression to isolate the part of the VPN node name that is to be used.

  Example: By adding characters before or after the parenthesis, you can create a set of VPNs that are similar to VPN node names.

  Selected VPN node names: AG-VPN-AMS and AG-VPN-FRA

  VPN Node Name: AG-(.+)

VPN Name: $1

Results in two VPNs: VPN-AMS and VPN-FRA

Example:

Selected VPN node names: vrf_AKD_V001_Amsterdam, vrf_AKD_V001_Paris, and vrf_AKD_V001_Frankfurt

VPN Node Name: (vrf)_(.+)_(V[0-9]+)_(.+)

VPN: $2

Results in one VPN: AKD

- If you selected to create a VPN from VPN node RDs (step #3), MATE Design uses regular expressions for both the VPN Node RD and VPN Name fields.

  Example: Create a VPN named "7" from three existing VPN nodes with RDs of 7:1, 7:2, and 7:3.

  VPN Node RD: (.+):(.+)

  VPN:$1

- If you selected to create a VPN from VPN node tags, MATE Design uses a tag to create the new VPN. If a VPN node has more than one tag, only the first tag listed is used. (To create VPN node tags or to change the order of their appearance, use the VPN Node Properties dialog box. Access it by double-clicking on one or more VPN nodes.)

Step 5 To see a list of VPN nodes that will be included in the VPN and the VPN names being created, click the Update Preview button.

Step 6 Select the service class for this VPN.

Step 7 Click OK.

Create VPNs

Step 1 Either right-click in an empty plot area and select New->VPN->VPN from the context menu, or select the Insert->VPN->VPN. A New VPN dialog box appears.

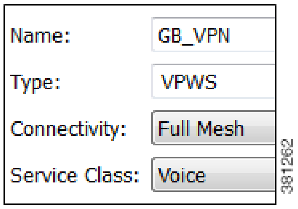

Name: GB_VPN

Type: VPWS

Connectivity: Full Mesh

Service Class: Voice

Step 2 Enter a unique name for the VPN.

Step 3 Select the VPN type: VPWS, VPLS, or L3VPN.

Step 4 Select the service class for this VPN.

Step 5 Click OK.

Step 6 If desired, add VPN nodes to the newly created VPN. See the Add VPN Nodes to VPNs section.

Add VPN Nodes to VPNs

Step 1 Select one or more VPN nodes from the VPN Nodes table.

Step 2 Right-click on any of the selected VPN nodes, and select Properties from the context menu.

Step 3 In the drop-down list of the Properties dialog box, select the VPN to which you are adding the selected VPN nodes.

Step 4 Click OK.

Multicast Simulation

MATE Design supports source-specific multicast (SSM), which is a method of delivering multicast packets to receivers only from a specified source address that is requested by the receiver. By limiting the source, SSM reduces resource requirements and improves security.

SSM Parameters

MATE Design supports SSM multicast, which is specified by an (S,G) parameter per multicast flow. The (S,G) pair and is labeled using a dot-decimal notation similar to IP addresses (for instance, 1.1.1.1, 2.2.2.2). Each (S,G) pair is the name of a multicast flow, and each flow is listed in the Multicast Flows table.

Discovered Versus Simulated Multicast Flows

How you work with multicast flows within MATE Design differs depending on whether the multicast flow was discovered by Collector or whether you are simulating it (Table 14-1).

Table 14-1 Discovered Versus Simulated Multicast Flows

	Discovered Multicast Flow	Simulated Multicast Flow
Creation	MATE discovers multicast flows ((S,G) pairs), and multicast traffic is derived from the Multicast Flow Traffic table (accessed from the Multicast Flows tab after selecting the Edit->Traffic menu).	First, manually create the multicast flow, defining both the source (S) and the destinations (G). Next, create a demand that links the source to these multicast destinations, allowing the routing to be simulated through the network.
Hops	Each discovered multicast flow includes multicast flow hops, which are node-interface combinations through which the multicast path flows. The Multicast Demand initializer uses multicast flow hops to determine multicast sources and destinations.	When you create a demand, it determines the path to take.
External Hops	MATE discovers multicast flow hops on interfaces that are external to the plan file. These are interfaces from a plan node to an external node, or from an external node to a plan node. The Multicast Demand initializer uses multicast flow external hops to determine multicast sources and destinations.	When you create a demand, it determines the path to take, but multicast external flow hops are not identified.

	Discovered Multicast Flow	Simulated Multicast Flow
Destinations	MATE does not identify a list of multicast destinations for each flow. The Multicast Demand initializer generates destinations from the Multicast Flow Hops and Multicast Flow External Hops tables.	Specify destinations (nodes, interfaces, external AS's, or external endpoints) when you create multicast flows. When creating multicast demands, specify these as multicast destinations.
Applicable Tables	• Multicast Flows • Multicast Flow Hops • Multicast Flow External Hops • Multicast Flow Traffic (accessed from the Multicast Flows tab after selecting the Edit->Traffic menu)	• Multicast Flows • Multicast Flow Destinations • Demands
Plot Displays	• Selecting a multicast flow from the Multicast Flow table highlights a pink path on the plot (Figure 14-1). • Selecting a multicast flow hop from the Multicast Flow Hops table highlights a pink path on the plot (Figure 14-2).	• Selecting multicast demands from the Demands table highlights the multicast path in blue on the plot and labels the source as A and the destinations as Z. • Selecting multicast flow destinations from the Multicast Flow Destinations table, labels the final destinations with a Z in the plot.

Figure 14-1 Multicast Flows

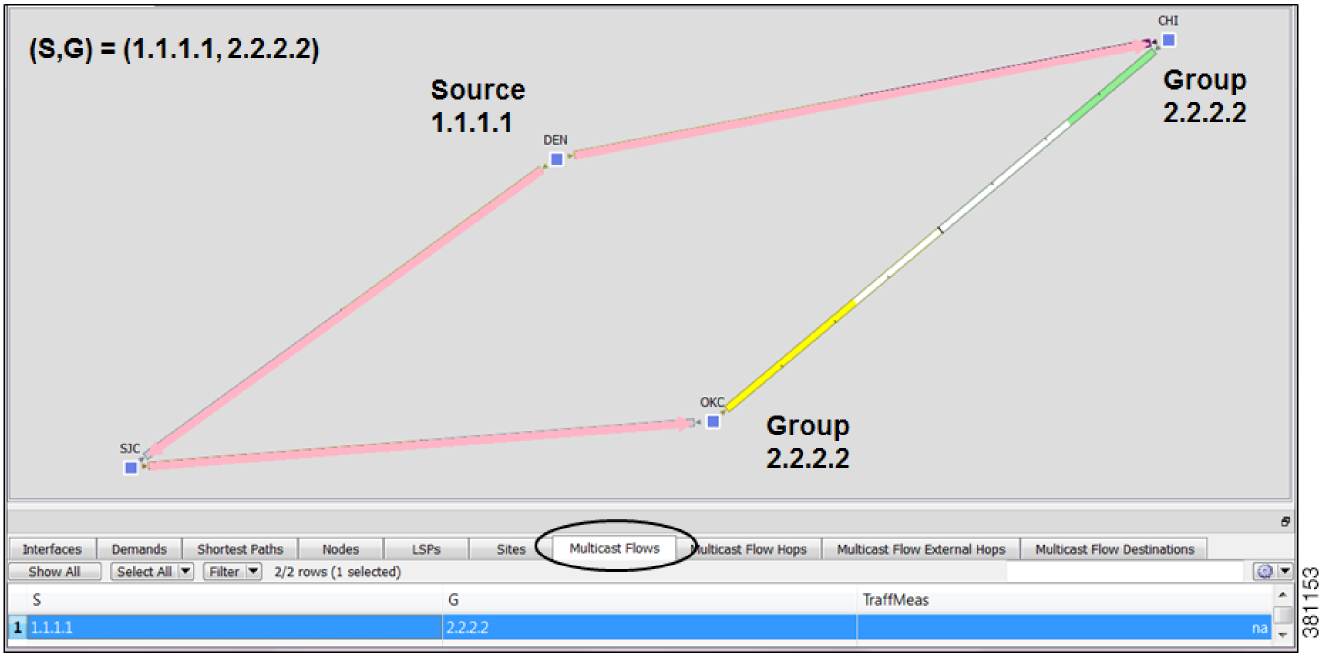

	S	G	TraffMeas
1	1.1.1.1	2.2.2.2	na

Figure 14-2 Multicast Flow Hops

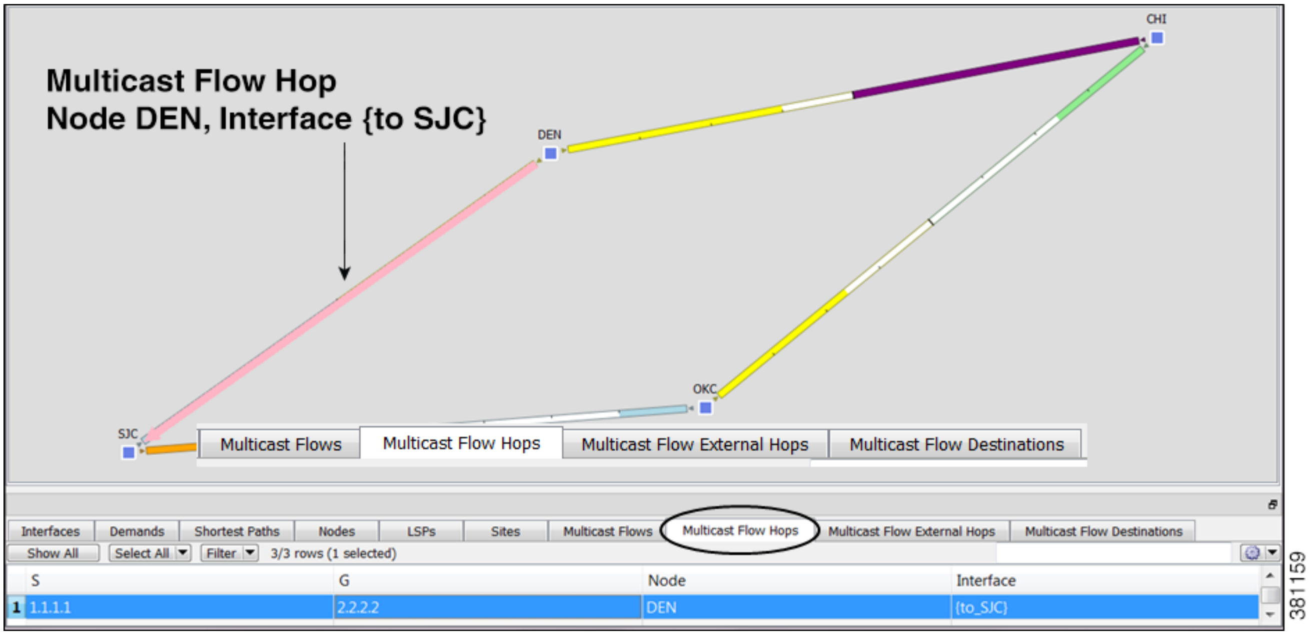

3/3 rows (1 selected)

	S	G	Node	Interface
1	1.1.1.1	2.2.2.2	DEN	{to_SJC}

Global Multicast Simulations

Flow Hops

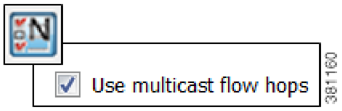

☑ Use multicast flow hops

If a plan file contains multicast information discovered by Collector, then it includes the current hops taken by multicast flows. You can specify that MATE Design multicast simulations should follow these flow hops if possible. This is useful, for example, when calculating incremental routing changes on the current network state, such as those caused by a failure. For planning purposes, when current state is not relevant, you can change this behavior to disregard the multicast flow hops.

MATE Design uses network state in its multicast simulation. You can set simulations to take into account multicast flow hops, and

Step 1 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 multicast flow hops in simulation, select or deselect the “Use multicast flow hops” option accordingly.

Step 4 Click OK.

Cisco Next Hop

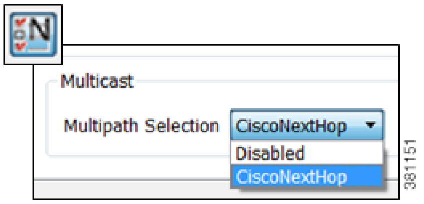

Multicast

Multipath Selection: CiscoNextHop

For manually inserted demands, all SSM demand traffic for the (S,G) pair (multicast flow) goes through any interface that is traversed by that demand. For more information, see the Simulated Multicast Demands section.

However, you can set MATE Design to use Cisco next hops, which are calculated using a hash on S and G for a multicast flow (S,G). The hash calculation is different in IOS and IOS XR. The default behavior is that of IOS. The IOS XR hash is used on all nodes whose OS field starts with IOS XR.

Step 1 Select the Edit->Network Options menu or click the Network Options icon in the Visualization toolbar.

Step 2 Select the Protocols tab.

Step 3 Select CiscoNextHop from the Multicast, Multipath Selection list.

Step 4 Click OK.

Multicast Demands

For discovered multicast flows, you can insert demands manually or use the initializer.

For simulated multicast flows, create demands manually. For information on demands in general, see the Traffic Demand Modeling chapter.

To find demands associated with a multicast flow, right-click on either a multicast flow or multicast flow destination, and select Filter to Demands. Conversely, to find multicast flows associated with demands, right-click on a demand and select Filter to -> Multicast Flows.

Initialized Multicast Demands

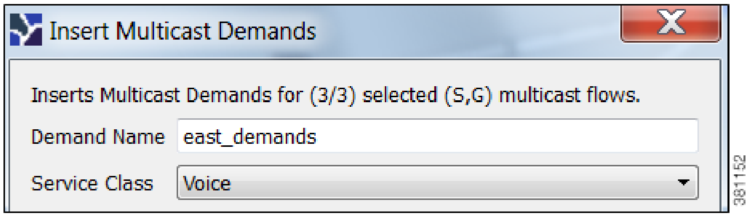

Insert Multicast Demands

Inserts Multicast Demands for (3/3) selected (S,G) multicast flows.

Demand Name: east_demands

Service Class: Voice

The initializer creates simulated multicast flows corresponding to the discovered multicast flows. The demand traffic for all traffic levels existing is set to the traffic measured in the multicast flow. Following is how MATE Design determines the source and destination.

- **Source**—MATE Design first looks for an external endpoint with the same source (S) name as the source of the discovered multicast flow. If this source does not exist, MATE Design infers it from the multicast flow hops that have traffic going out of them, but none coming in. The inferred source can be a node or external endpoint.

- **Destination**—MATE Design first looks for multicast flow destinations with the same (S,G) name as those in the discovered multicast flow. If these multicast node destinations do not exist, MATE Design infers them from the multicast flow hops that have traffic going into them, but none coming out. MATE Design also infers multicast external flow hops if traffic is leaving a node in a direction that is outside the plan. These destinations are populated in the Multicast Flow Destinations table.

To run this initializer, select the Initializers->Multicast Demands menu.

Simulated Multicast Demands

For manually inserted demands, all SSM demand traffic for the (S,G) pair (multicast flow) goes through any interface that is traversed by that demand. SSM demands are routed to take the shortest path from the destination to the source, rather than from the source to the destination, as unicast demands do. By default, multicast multipath is disabled. If two paths of equal cost exit from a node on the route back to the source, the path is chosen based on these parameters.

- The remote interface with the highest IP address is used.
- If IP addresses are not available, the router name with the lowest lexicographical name is used.

However, you can set MATE Design to use the Cisco next-hop multicast, multipath selection method. For information, see the Cisco Next Hop section.

The sources of these multicast demands can be nodes, interfaces, external AS's, or external endpoints. Using interfaces enables you to specify the exact interface on which demand traffic is entering. Using external endpoints enables you to model situations where the multicast source is outside of the nodes in the plan, and there is more than one possible entry point of the traffic flow into and through the interfaces in the plan.

In this example, the selected demand has a node source of DEN, and a multicast destination containing two nodes: CHI and OKC. The source node (S) is marked by A, and the destinations in the receiver group (G) are marked by Z.

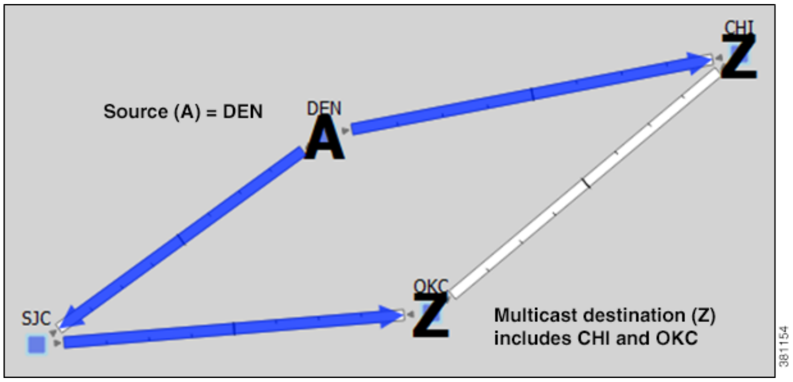

You can model a source that is outside the plan using external endpoints. In this example, the demand's source is an external endpoint that contains SJC and DEN as members. The multicast destination includes both CHI and OKC.

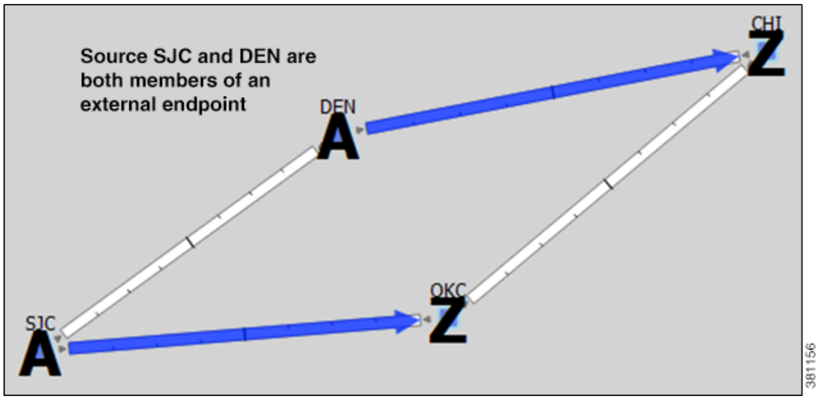

For instructions on how to create multicast demands, see Create Demands for Multicast Flows.

View Multicast Flows

To highlight discovered multicast flows and multicast flow hops in the plot, select them from their respective tables. To view sources and destinations of multicast demands, select the demand from its table. The source is identified in the plot by A and the destinations are identified by Z.

View	Show this Table
Multicast flows, including both source and receiver names	Multicast Flows
Discovered multicast flow hops, including (S,G) name, hop node, and hop interface	Multicast Flow Hops
Discovered multicast flow external hops (destinations that are inferred as external to the plan), including (S,G) name, direction of the outbound interface	Multicast Flow External Hops
User-created multicast flow destinations, including (S,G) name and destination node	Multicast Flow Destinations
Multicast demands	Demands

Create Multicast Flows

The easiest and recommended method of creating multicast flow destinations is to do so while creating its associated multicast flows, as follows.

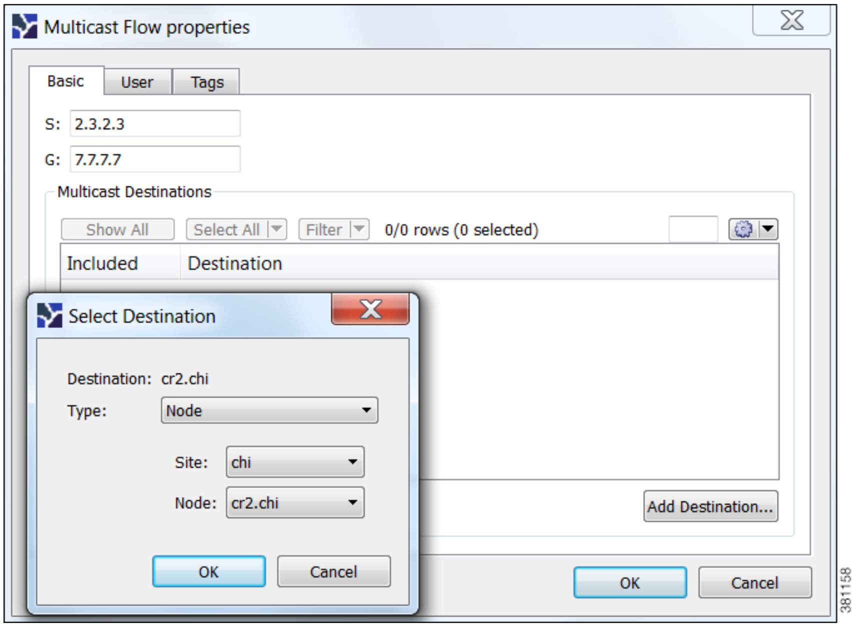

Multicast Flow properties

S: 2.3.2.3

G: 7.7.7.7

Multicast Destinations

0/0 rows (0 selected)

Included	Destination

Select Destination

Destination: cr2.chi

Type: Node

Site: chi

Node: cr2.chi

Step 1 Either select the Insert->Multicast Flow menu, or right-click on the plot and select New->Multicast Flow from the context menu. A Multicast Flow dialog box appears.

Step 2 In the S field, identify the source node by entering its dot-decimal name. For example, 2.3.2.3.

Step 3 In the G field, identify the collective name for all receiver nodes (destinations) by entering its dot-decimal name. For example, 7.7.7.7.

Step 4 Add destinations to the receiver group.

a. Click on Add Destination. A Select End Point dialog box appears.

b. In the Type list, select either a node, interface, external AS, or external endpoint.

c. If you are adding a node to the receiver group, select both the site and node name of the destination point in the multicast flow.

If you are adding an interface, select the site, the node, and the interface.

If you are adding an external AS, select the external AS name and node for the destination point.

If you are selecting an external endpoint, select its name.

d. Click OK. These values now appear in the Multicast Flow dialog box where you can further disable or enable them from participating in the multicast flow.

Step 5 Repeat steps 4A-4D for each destination in the multicast flow.

Step 6 Click OK in the Multicast Flow dialog box. Note if you close without clicking OK again, the above changes are not saved.

After creating the multicast flow, to make it meaningful you need to create demands for the multicast flows. See Create Demands for Multicast Flows.

Create Demands for Multicast Flows

Step 1 Either select the Insert->Demand menu, or right-click on the plot and select New->Demand. A New Demand dialog box appears.

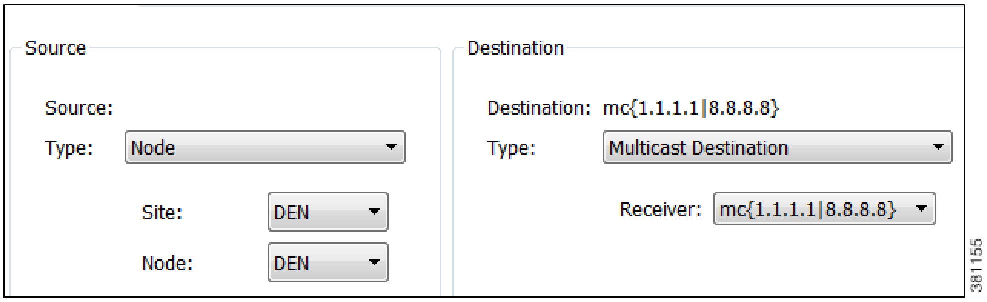
Source

Source:

Type: Node

Site: DEN

Node: DEN

Destination

Destination: mc{1.1.1.1|8.8.8.8}

Type: Multicast Destination

Receiver: mc{1.1.1.1|8.8.8.8}

Step 2 In the Name field, enter the name of the demand.

Step 3 In the Source area, define the source (S) of the multicast flow.

a. From the Type list, select the source as a node, interface, external AS, or external endpoint.

b. For node sources, select both the site and the node.

For interface sources, select the site, node, and the interface.

For external AS's, select both the external AS and the node within the plan file through which this external traffic flows.

For external endpoints, select its name.

Step 4 In the Destination area, define the receiver group (G) of the multicast flow.

a. Select Multicast Flow Destination from the Type list.

b. Select the receiver that identifies the simulated multicast flow (S,G) from the Receiver list.

Step 5 Optional: Complete all other fields as needed.

Step 6 Click OK.

Toggle Multicast Flow Destinations to On and Off

Step 1 In the Multicast Flows table, double-click on the multicast flow. A Multicast Flow Properties dialog box appears with a list of existing destinations. Each destination appears with a T next to it if it is currently part of the receiver group or an F it exists, but is not part of the receiver group.

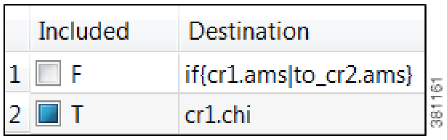

	Included	Destination
1	□ F	if{cr1.ams|to_cr2.ams}
2	☑ T	cr1.chi

Step 2 To enable or disable the multicast flow destination, click on its associated T/F value.

Step 3 Click OK in the Multicast Flow dialog box. Note if you close without clicking OK again, the change is not saved.

CHAPTER 15

Layer 1 Simulation

Although MATE Design focuses on IP/MPLS Layer 3 (L3) network views and traffic, you can also create and view Layer 1 (L1) topologies, thus giving a more complete view of the network. Using the L3 and L1 topologies together, you can specify how L3 circuits route over L1 links. These Layer 3-to-Layer 1 associations enable you to simulate individual L1 link and L1 node failures, or to include L1 links and L1 node failures in a Simulation Analysis to determine the effects on L3 traffic. Additionally, you can include L1 links and L1 nodes in SRLGs. With this interaction of layers, you can better design the network to meet performance requirements.

Note Aside from this chapter, throughout this guide the term *circuit* means a circuit with L3 properties. In this chapter, for purposes of clarity, *L3 circuit* and *L1 circuit* refer to objects defined in the Circuits and L1 Circuits tables, respectively, whereas *circuit* refers to the actual network element.

L1 and L3 Views

Figure 15-1 shows an example L1 view in comparison to an L3 view. Notice the different uses of circuits. In the L3 view, L3 circuits are relevant to IP/MPLS routing and show traffic. The L1 view shows the physical topology with no traffic; if selected, it also shows the L1 circuit paths taken across the L1 links. L3 is the default view unless the plan file does not contain L3 nodes.

To alternate between these two views, use one of these methods.

- Click the Layer View button in the top, left of the Visualization toolbar. This button is a toggle between Layer 3 (L3) and Layer 1 (L1), and tells you which layer you are viewing. L1
- Select the View->Layer 1 View or View->Layer 3 View menus.
- Use keyboard shortcuts Ctrl-1 for L1 and Ctrl-3 for L3. (On a Mac, use the Alt-1 and Alt-3.)

Note that to more easily find relationships between L3 and L1 objects, including interfaces, L3 circuits, L3 nodes, ports, port circuits, L1 nodes, L1 circuits, and L1 links, right-click on the object and select the appropriate filter from the context menu.

Figure 15-1 Comparison of L1 and L3 Views

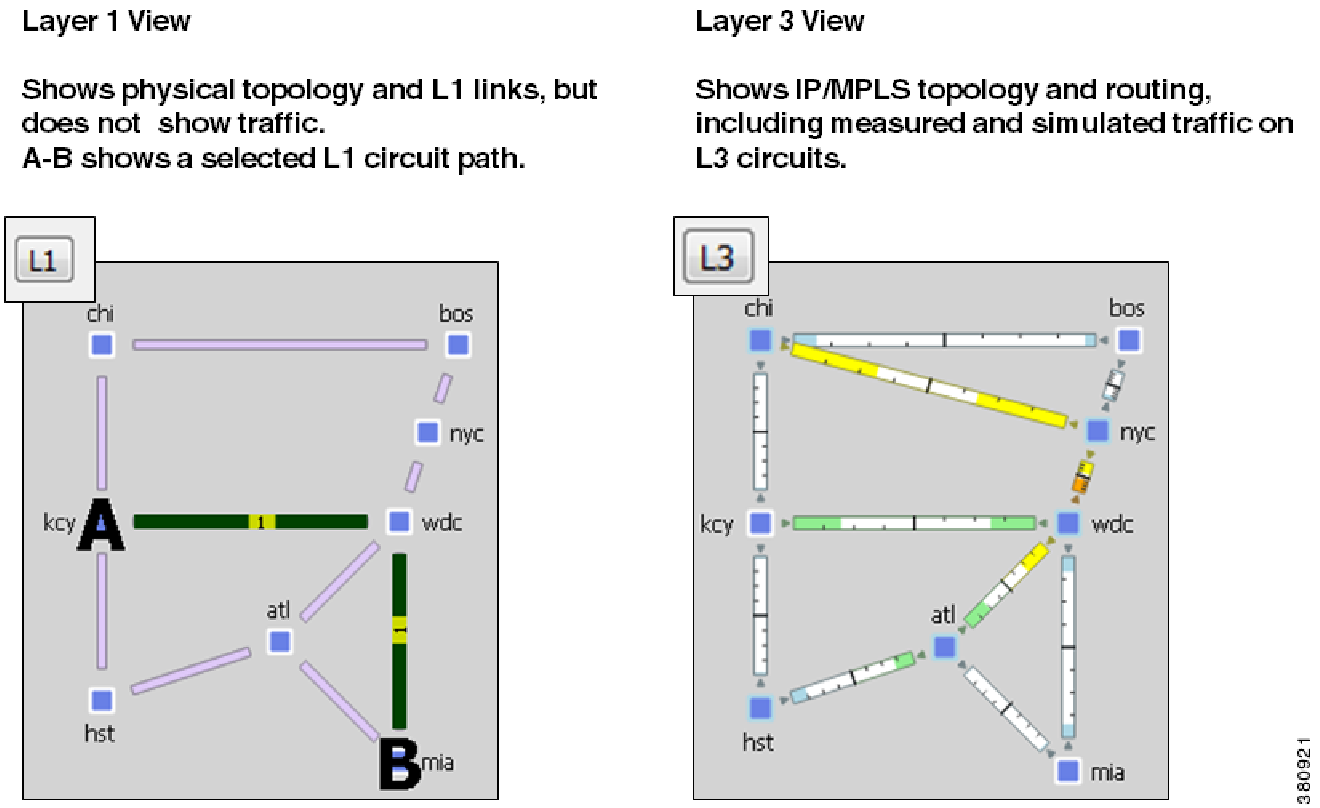

Layer 1 View

Shows physical topology and L1 links, but does not show traffic.
A-B shows a selected L1 circuit path.

Layer 3 View

Shows IP/MPLS topology and routing, including measured and simulated traffic on L3 circuits.

L1 Topology

An L1 topology consists of a unique set of objects (Table 15-1) that are distinct from IP/MPLS L3 topology and objects. However, two of the objects, L1 circuits and L3 circuits, represent the same circuit.

- An L3 circuit contains L3 properties of the circuit.
- An L1 circuit contains L1 properties, specifically L1 routing, of the circuit.

Collector does not discover L1 objects, but you can easily create the objects and generate an L1 topology in MATE Design using these methods.

- Import L1 topology from a another plan file. This method creates the same L1 topology and the same L1 link, L1 circuit, and L3 mappings as in the source file from which you are copying. Note that you can also copy L1 topology from a template plan file.
- Construct an L1 topology based on all or part of an existing L3 topology in the same plan file. This method creates an L1 topology that mirrors the L3 one, and automatically generates the mappings between L1 links, L1 circuits, and L3 circuits based on the selected sites or L3 nodes. It also automatically generates one standby L1 circuit path per each L1 circuit and sets its Path Option to 1. L1 link metrics are all set to 1.

- Individually creating the L1 nodes, L1 links, L1 circuits, L1 circuit paths, and L1 circuit path hops. You can add both L1 nodes and L1 links to SRLGs for simulation purposes. This is typically done to represent ducts containing multiple fiber segments, where a break in the duct could cause multiple L1 link failures, and thus potentially multiple L3 circuit failures for each of these L1 links.
- Create the L1 topology directly in the tables in the plan file, such as by editing the L1 tables using the Plan Tables as Database tool under the Edit menu or by creating the table in the plan file using a text editor.

For More Information ...	**See ...**
Import L1 model (objects and tags)	*MATE Integration and Development Guide*
Create L1 topology based on L3 topology using the Insert tool	Create L1 Topology from L3 Topology section
Copy L1 topology from a template	*MATE GUI Visualization Guide*
Modify how the L1 topology appears in the plot	*MATE GUI Visualization Guide*

L1 Objects

Table 15-1 L1 Objects

L1 Object	**Description**
L1 node	Models L1 devices and is listed in the L1 Nodes table.
L1 link	Models the fiber segments connecting the L1 nodes. These are listed in the L1 Links table and are highlighted in blue upon selecting them. L1 circuits are routed over L1 links, which is how these L1 links can be mapped to the IP/MPLS topology in the L3 view. For more information, see L1 Simulations and Relationships Between L1 and L3.
L1 circuit	Models the L1 properties of a circuit, such as which L1 nodes are connected to it and the physical path and distance. These are listed in the L1 Circuits table and are highlighted in green upon selecting them. An L1 circuit can be associated with either an L3 circuit or L3 port circuit. For more information, see Relationships Between L1 and L3.
L3-L1 link	Models the physical relationship between an L3 node and an L1 node.
L1 circuit path	Models the paths available for the L1 circuits. These are listed in the L1 Circuit Paths table and are highlighted in green upon selecting them. The number on the highlighted path indicates the Path Option number.
L1 circuit path hop	Models the L1 node or L1 links in the L1 circuit path. These are listed in the L1 Circuit Path Hops table.
Actual L1 circuit path hop	Discovered L1 circuit path hops. Although Collector does not discover these, if you have a discovered set of L1 circuit path hops, you can manually add them using the Actual tab in the L1 circuit path Properties dialog box. You can also edit the .txt plan file to add an <ActualL1CircuitPathHops> table.

For More Information ...	See ...
Create or edit L1 nodes	Create L1 Nodes section
Create or edit L1 links	Create L1 Links section
Create or edit L1 circuits	Create L1 Circuits section Create L1 and L3 Mappings section
Create or edit L3-L1 links	Create L3-L1 Links between Nodes section
Create L1 circuits that are automatically mapped to L3 circuits	Create L1 and L3 Mappings section
Associate port circuits or ports with L1 circuits or L1 nodes, respectively	Create L1 and L3 Mappings section
Create L1 circuit paths	Create L1 Circuit Paths
Create explicit L1 circuit path hops	Individually: Create Individual L1 Circuit Path Hops section Automatically: Explicit L1 Circuit Path Initializer section
Set L1 link metrics using an initializer tool	Layer 1 Link Metric Initializer section
Initialize the distance and delay of L1 links	Simulation chapter
Consider L1 node and L1 links when running the Metric Optimization tool	Metric Optimization chapter
Consider L1 node and L1 links when creating secondary paths using the Explicit LSP Optimization tool	LSP Optimization chapter

Sites

All sites are listed in the Sites table whether they contain L1 nodes, L3 nodes, or both. To determine how many L1 nodes, if any, are in a site, show the L1 Node Count column. By default, sites containing only L1 nodes appear only in the L1 view, and sites containing only L3 nodes appear only in the L3 view. You can change this default in the Plot Options dialog box. Empty sites appear in both views.

For More Information ...	See ...
Show or hide sites containing only L1 nodes or only L3 nodes (for the entire L1 or L3 view)	Edit L1 Plot Options
Hide individual nodes or sites	*MATE GUI Visualization Guide*

L1 Simulations

When reading the L1 tables, the derived Established and Routed properties indicate simulation capabilities.

- Established—Identifies whether an L1 circuit or L1 circuit path is routed when there are no failures. Note that while an L1 circuit or L1 circuit path has Established equals F (false), it might be routable under failures (Figure 15-15).

- Routed—Identifies whether and how an L1 circuit or L1 circuit path is routed.
  - Unrouted
  - Actual (if being routed through actual L1 circuit path hops)
  - Simulated (if being routed and not using actual L1 circuit path hops)

Active L1 Path Selection

Since multiple L1 circuit paths can be routed for a single L1 circuit, MATE Design uses the Active Path Sim property to determine which L1 circuit path to use to set the derived properties for each L1 circuit, such as Distance Sim and Delay Sim. The Active Path Sim property is copied from the user-specified Active Path property if it has been set and if it is routable. Figure 15-5 shows an example of the Active Path Sim changing due to a change in an L1 circuit Feasibility Limit property.

Standby Paths and Path Options

The user-specified L1 circuit path's Standby property sets the path to standby (T) or non-standby (F). The user-configurable Path Option property is an integer that identifies the order in which the routing of L1 circuit paths is explored. The lower the value, the higher the priority to attempt to route the L1 circuit path.

If other L1 circuit paths must be explored, the standby L1 circuit path with the lowest Path Option value is used provided it is routable. If it is not routable, attempts are made for each Path Option with the next lowest value (next highest priority). If none of the standby L1 circuit paths are routable, MATE Design then follows the same procedure with non-standby L1 circuit paths, starting with the one with the highest priority Path Option.

L1 Circuit Path Routing Order

- MATE Design always attempts to route L1 circuit paths that have their Standby property set to T (true), regardless of whether they are set as the active path. If they are routable, they are routed in the order of this priority.
  - Actual L1 circuit path hops have first priority. Note that while these are not discovered by Collector, if they are discovered through another means, you can manually add them using the Actual tab of the L1 circuit paths Properties dialog box or by adding an <ActualL1CircuitPathHops> table to the .txt plan file.
  - L1 circuits paths with explicit Lambdas or Lambdas defined through L1 circuit path hops are then routed. For information on lambdas, see the Lambda Assignments section.
  - Thereafter, L1 circuit paths that do not have Lambda properties set explicitly or through L1 circuit path hops are routed.
- After attempting to route standby L1 circuit paths, MATE Design determines whether and how to route non-standby L1 circuit paths (Standby property set to F (false)).
  - MATE Design determines if the non-standby L1 circuit path should route by checking the following L1 circuit properties in this order.
    - If the L1 circuit Active Path property is set to this L1 circuit path, MATE Design attempts to route the non-standby L1 circuit path.

- If the number of routed L1 circuit paths is less than the associated L1 circuit's Min Routed Paths property, then MATE Design attempts to route the non-standby L1 circuit paths.

– Non-standby L1 circuit paths are routed based on their Path Option setting (lower numbers are routed first). For each of these paths, the route order is based on the same criteria as described above for standby L1 circuit paths (actual L1 circuit path hops, then Lambda values, and then L1 circuit path hops that do not have Lambda values).

Within each of the above routing instances, MATE Design routes the L1 circuit path based on the following criteria regardless of whether the L1 circuit path is standby or non-standby.

- Routes are based on the lowest Setup Priority value and then deterministically. For example, amongst L1 circuit paths with actual L1 circuit path hops, those with higher Setup Priority are routed first. See the L1 Circuit Path Setup Priority section.
- L1 circuit paths that do not require automatic wavelength assignments to route have higher priority than those that do. For information on the Auto Lambda property, see the Lambda Assignments section.

View L1 Circuit Paths

To view the actual paths or simulated paths of L1 circuits, select the View->L1 Circuits menu. Note that given the priority of routing, if the actual L1 circuit path hops path is valid, then the actual and simulated paths are the same. You can also choose to view the shortest latency path from this menu.

Rerouted simulated L1 circuit paths show automatically when they are rerouted due to failures, regardless of whether you are showing simulated paths.

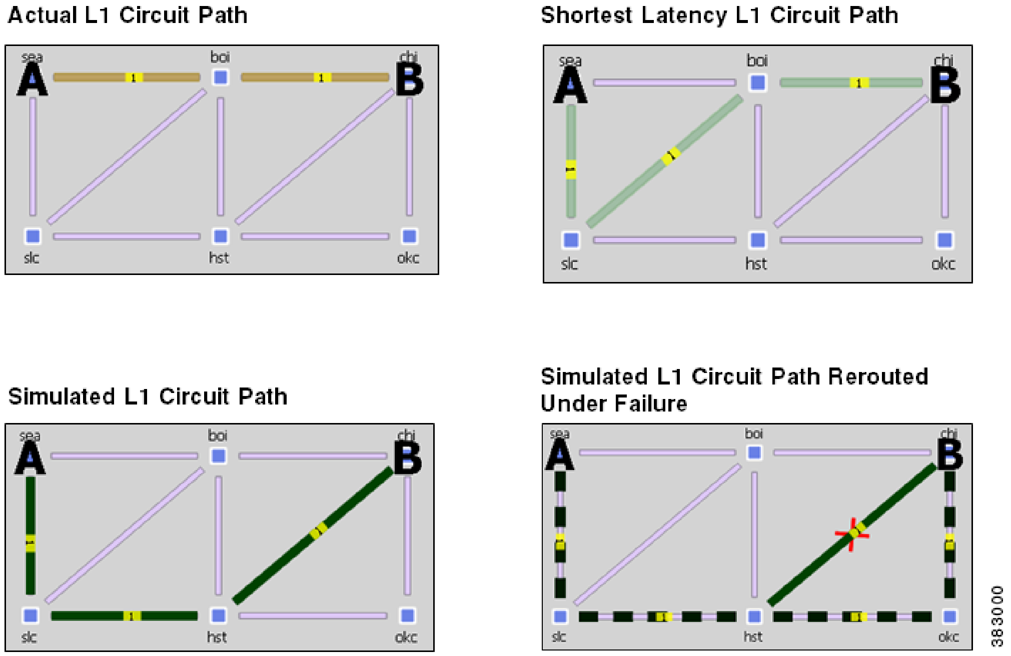
Actual L1 Circuit Path
Shortest Latency L1 Circuit Path
Simulated L1 Circuit Path
Simulated L1 Circuit Path Rerouted Under Failure
sea
boi
chi
A
B
slc
hst
okc
1
383000

L1 Circuit Path Hops

You can force an L1 circuit path to choose a different route by setting one or more node hops or link hops.

Node Hops	Link Hops
• Loose—The path must pass through the node, but the node does not have to be connected directly to the previous hop. • Strict—The node must be directly connected to the previous node. • Exclude—This node cannot be in the path.	• Strict—The L1 link must be in the path in the specified order. • Exclude—This L1 link cannot be in the path.

Generally, node hops are used for simplicity, but link hops are useful for when you need to distinguish between two parallel L1 links. To do this, set a loose hop on the L1 node preceding the link hop, and then set a strict hop on the parallel L1 link over which you want to route this L1 circuit.

You can set L1 circuit path hops in an L1 circuit path Properties dialog box. Since MATE Design always uses the shortest path, remember to set as many hops as needed to achieve the desired path, and put them in the order of the desired path. The L1 Circuits Path Hops table identifies the path and the steps.

Note Since L1 circuits paths are bidirectional, to simulate L1 circuit path hops MATE Design automatically adds the initial L1 node in the path as a strict hop.

Example: Figure 15-2 shows an example of a default shortest path being used, as well as an alternate circuit path created using a loose nyc node hop. It also demonstrates that you could achieve the same path by setting only an exclude link hop on wdc-bos.

Figure 15-2 Shortest L1 Circuit Paths Versus Paths Determined by L1 Circuit Path Hops

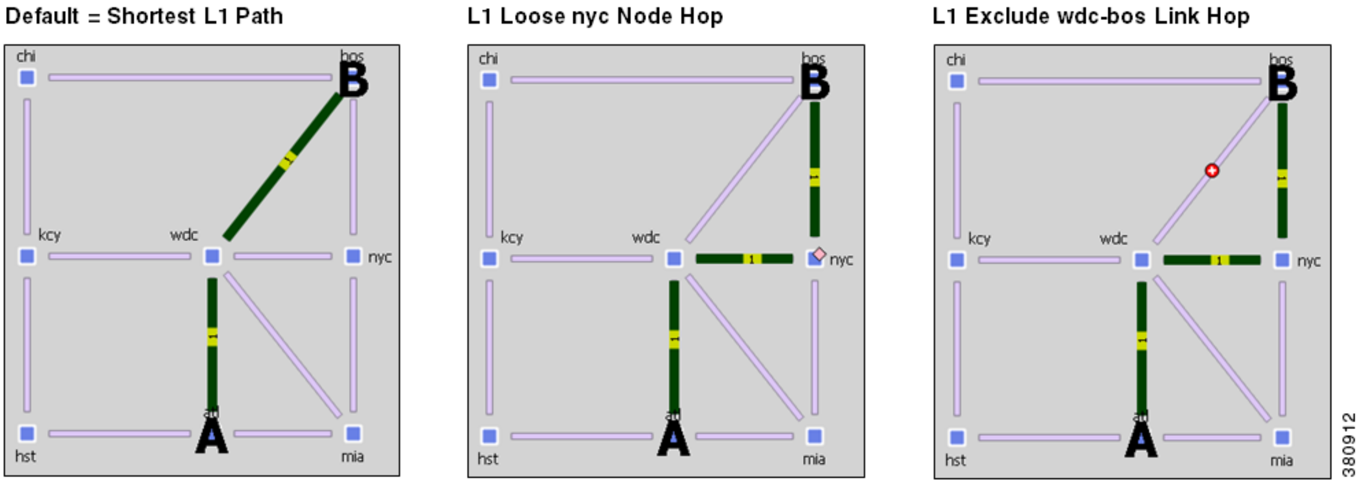

For More Information ...	See ...
Automate the process of setting up L1 circuit hops	Explicit L1 Circuit Path Initializer
Set L1 circuit hops per L1 node or L1 link	Create Individual L1 Circuit Path Hops

Explicit L1 Circuit Path Initializer

The Explicit L1 Circuit Paths Initializer automates the process of creating explicitly routed L1 circuit paths. It can optionally use disjoint groups or create disjoint paths in the process. Two L1 circuit paths are *disjoint* if they do not route over common L1 objects. This is useful, for instance, when planning L1 circuit paths, particularly when multiple L1 circuit paths connect the same pair of L1 nodes. Using disjointness on these L1 circuit paths can help improve availability.

If any L1 circuit path is unrouted, no L1 circuit path hops are created for it. If the simulated Lambda of a L1 circuit path is different than the user-specified Lambda on the path, then the initializer sets the Lambda value of the L1 circuit path to the simulated Lambda. To retain existing L1 circuit path hops, select the “Respect existing L1 circuit path hops” option.

The initializer creates L1 circuit paths based on either current routes or on modified routes.

- Follow current simulated routes—Use only the existing simulated L1 circuit path routes and do not create disjoint groups
- Create disjoint paths between L1 Circuits in Disjoint Groups—Create disjoint L1 circuit paths for L1 circuits that are in the same disjoint group. For example, if two L1 circuits are in the same disjoint group and each has two path options, all four L1 circuit paths will become disjoint from each other.

  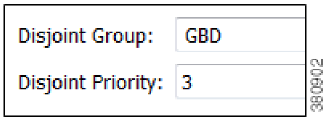
  Disjoint Group: GBD
  Disjoint Priority: 3

  The pre-requisite for this option is that the L1 circuits must first be added to the disjoint groups in the L1 circuits’ Properties dialog boxes. Here, you can also assign priorities to L1 circuits within these groups. Higher priority L1 circuits are assigned shorter path routes based on L1 link metrics. The higher the number, the lower the priority. For information on entering and selecting fields in this Properties dialog box, see the Create L1 Circuits section. The L1 Circuit table lists these properties, if applicable.
- Create disjoint paths within L1 Circuits—Create disjoint paths within L1 circuits. For instance, if two L1 circuits each have two path options, each set of L1 circuit paths will be disjoint from each other, but they may or may not be disjoint from the L1 circuit paths belonging to the other L1 circuit.

  These objects are configurable and can be L1 links, L1 nodes, sites, and SRLGs. These objects have a disjointness priority specified, which is used if complete disjointness is not possible.

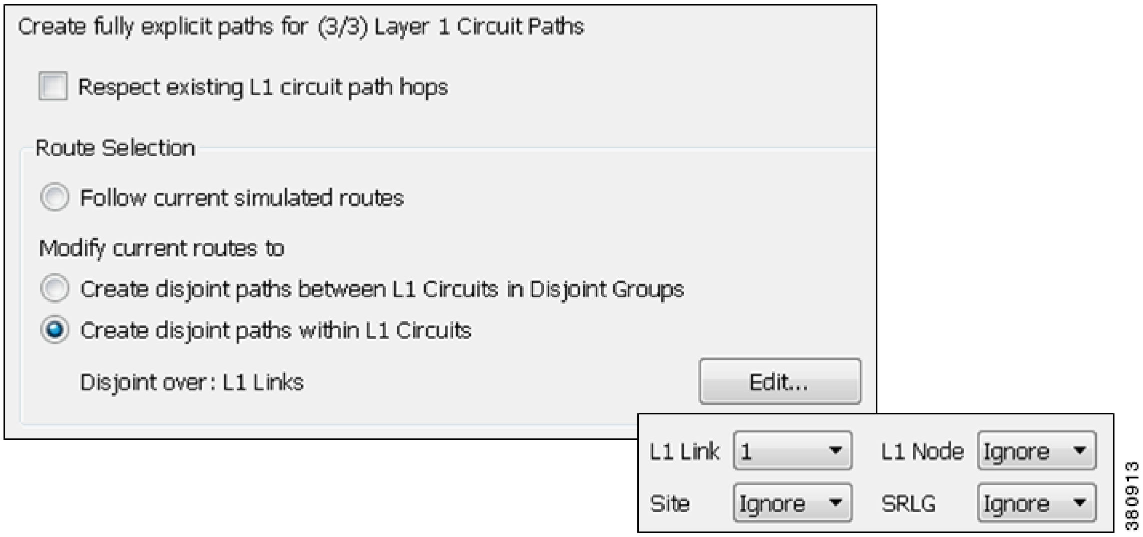

Create fully explicit paths for (3/3) Layer 1 Circuit Paths

☐ Respect existing L1 circuit path hops

Route Selection

○ Follow current simulated routes

Modify current routes to

○ Create disjoint paths between L1 Circuits in Disjoint Groups

● Create disjoint paths within L1 Circuits

Disjoint over: L1 Links

L1 Link: 1
L1 Node: Ignore
Site: Ignore
SRLG: Ignore

The L1 Circuit Path Hops table lists the newly created the L1 circuit path hops. See the L1 Circuit Path Hops section for a description of these fields.

Step 1 Select one or more L1 circuit paths. If you select none, the default is to create L1 circuit path hops for all L1 circuit paths.

Step 2 Either select the Initializers->Explicit Layer 1 Circuit Paths menu, or right-click on an L1 circuit path and select Explicit Layer 1 Circuit Paths. The Layer 1 Circuit Paths Initializer appears.

Step 3 Select whether to use the existing L1 circuit path hops.

Step 4 Select how to create the explicitly routed L1 circuit path hops.

- Follow current simulated routes
- Modify existing routes to use disjoint L1 circuit paths, select whether to create the disjoint paths for each L1 circuit or to create disjoint paths between path options within the same L1 circuit.

Step 5 To modify the disjoint priorities that determine which L1 objects are considered in the paths, click Edit. Then select the priorities from the drop-down L1 object menus, and click OK. For example, if you want L1 links and SRLGs to be equally considered, but not nodes or sites, you could set both L1 links and SRLGs to 1, and set site and L1 nodes to Ignore.

Step 6 Click OK.

L1 Circuit Route Reports

You can create reports on L1 circuits based on the L1 circuit path identified as the Active Path Sim for each L1 circuit. Additionally, the report contains all other columns in an L1 Circuits table, though they are initially hidden.

Step 1 Select the Tools->Reports->Layer 1 Circuit Routes menu.

Step 2 Click OK.

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

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

L1 Link Metrics

When L1 circuit paths are routed over L1 links, by default they are established along the shortest path between the L1 nodes. This shortest path is defined as the path with the lowest metric, and the metric of an L1 circuit path is defined as the sum of the L1 link metrics in the path. To move L1 circuit paths away from an L1 link, raise the metric on that L1 link.

Since you can simultaneously edit multiple L1 link Metric properties, this method is useful for altering numerous L1 circuit paths at a time. For exact path specification per L1 circuit path, however, you need to set explicit L1 circuit path hops.

L1 Link
Metric: 7.3

If two shortest L1 circuit paths have the same metric, MATE Design chooses one arbitrarily provided no other properties are impacting the route.

Note that in contrast to an interface IGP metric, this L1 metric does not represent an actual configurable quantity in the network.

Example: Figure 15-3 shows an example L1 circuit from sjc to kcy. When the sjc-slc L1 link has a metric of 2.7, the shortest L1 circuit path is sjc-slc-kcy. By raising the sjc-slc metric, the shortest L1 circuit path becomes sjc-sea-slc-kcy.

Figure 15-3 Example Effect of L1 Link Metric on Routing L1 Circuits

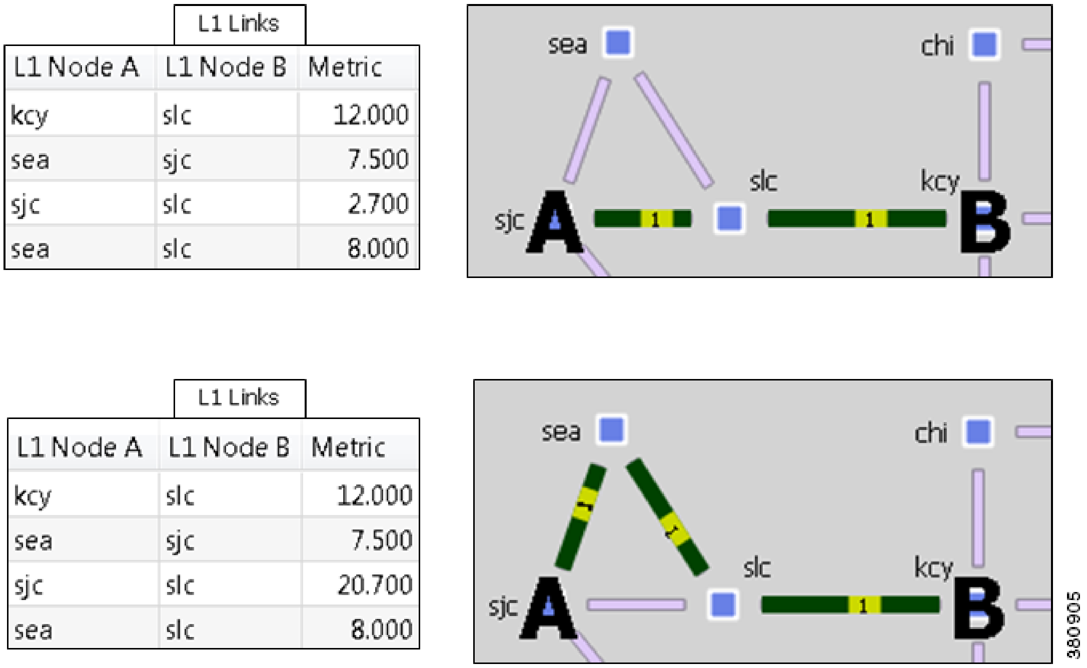

L1 Links

L1 Node A	L1 Node B	Metric
kcy	slc	12.000
sea	sjc	7.500
sjc	slc	2.700
sea	slc	8.000

L1 Links

L1 Node A	L1 Node B	Metric
kcy	slc	12.000
sea	sjc	7.500
sjc	slc	20.700
sea	slc	8.000

Layer 1 Link Metric Initializer

The L1 Link Metric Initializer sets L1 metrics based on L1 link distances. MATE Design calculates the metric by multiplying each L1 link's Distance value by the multiplier you specify. You also set a minimum value for the resulting metric. MATE Design then uses the higher of the two values.

Example: If the L1 Distance value is 7, the multiplier value is 3, and the minimum value is 1, the resulting L1 metric is 21 (Distance * Multiplier). However, if the minimum value were 25, the resulting L1 metric would be 25.

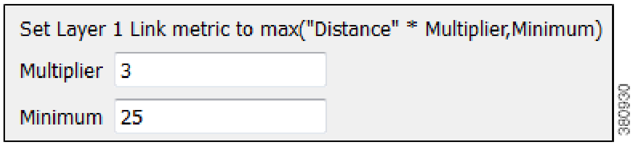
Set Layer 1 Link metric to max("Distance" * Multiplier,Minimum)

Multiplier 3

Minimum 25

Step 1 Make your selections.

- Select one or more L1 links.
- If you make no selections, metrics are calculated for all L1 links.

Step 2 Either select the Initializers->Layer 1 Metric menu, or right click an L1 link, and select Initialize L1 Metric from the context menu.

Step 3 Enter a multiplier and minimum value.

Step 4 Click OK.

L1 Circuit Path Setup Priority

The Setup Priority property identifies the order in which L1 circuits paths are routed when they share the same L1 link. The lower the number, the higher the priority. Figure 15-6 and Figure 15-8 show examples of rerouting L1 circuit paths based on Setup Priority.

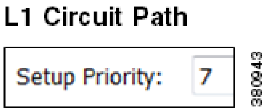
L1 Circuit Path

Setup Priority: 7

To set the Setup Priority, right-click on one or more L1 circuit paths, and select Properties from the context menu. This value appears in the L1 Circuit Paths table in the Setup Priority column.

L1 Circuit Wavelengths

The L1 circuit lambda (wavelength) properties include those that identify transmission feasibility and lambda assignments. Using these properties, you can answer such questions as the following.

- Can I add an L3 circuit that uses a specific L1 circuit?
- How will my network choose a new L1 circuit path under L1 failure?
- Can I modify an existing L1 circuit for an L3 circuit?
- How many available wavelengths are there in my L1 topology?
- What happens to the network when there are L1 failures? Does optical restoration select a new L1 circuit path, and if so, which path?

Note The terms *lambda* and *wavelength* are interchangeable.

Transmission Feasibility

In a real network, the quality of an L1 circuit deteriorates as it passes through L1 links. Using feasibility properties, MATE Design enables you to simulate this weakening of L1 circuits to determine if they have degraded to the point of being unroutable.

Note Although the feasibility properties are predominantly used for tracking the quality of an L1 circuit's wavelength as it passes through L1 links, it is an abstract property. You could optionally use it for other purposes, such as for representing latency restrictions. For example, for an L1 circuit, the stricter the requirements on latency, the lower the Feasibility Limit you would set. In this example the Feasibility Metric of L1 links would be set to their respective Delay values.

L1 Object	Feasibility Property	Description
L1 link	Feasibility Metric	Identifies how much an L1 circuit is degraded when it passes through that L1 link.
L1 circuit	Feasibility Limit	Identifies the acceptable quality of the L1 circuit.
	Feasibility Metric Sim	This is copied from the Feasibility Metric Sim of the active L1 circuit path.
	Feasible Transmission	Identifies whether the L1 circuit can be transmitted. This is copied from the Feasible Transmission of the active L1 circuit path. If all L1 circuit paths have a higher Feasibility Metric Sim than the L1 circuit's Feasibility Limit, this value is F (false), and the L1 circuit cannot route. This value is set to "na" if there is no Feasibility Limit.
L1 circuit path	Feasibility Metric Sim	The sum of the Feasibility Metric of the L1 links over which the L1 circuit path is routed.
	Feasible Transmission	Identifies whether the L1 circuit path can be transmitted. If the L1 circuit's Feasibility Limit is more than the sum of the Feasibility Metrics for the L1 links that the L1 circuit path traverses, this value is T (true) and the L1 circuit path can route. This value is set to "na" if there is no L1 circuit Feasibility Limit. **Note** This property is not applicable if an L1 circuit is using actual L1 circuit path hops.

Routing Based on Feasibility Properties

The Feasibility properties affect routing as follows.

- If there is no Feasibility Limit, the shortest path is based on the L1 link Metric property, and Feasible Transmission is set to "na."
- If the L1 circuit has a Feasibility Limit, then the shortest L1 circuit path is based on the L1 links that result in a L1 circuit path's Feasibility Metric Sim value that is lower than the L1 circuit's Feasibility Limit. The Feasible Transmission of the L1 circuit path is set to T (true).
- If the L1 circuit path Feasibility Metric Sim is higher than the L1 circuit's Feasibility Limit, then it is not routed, and the Feasible Transmission is set to F (false) except when actual L1 circuit path hops exist.

Example: Figure 15-4 shows a series of changing L1 link Feasibility Metric values and the effects on the L1 circuit paths when the L1 circuit has a Feasibility Limit of 20.

- Row 1 shows that all Feasibility Metrics are 10, so the L1 circuit path routes directly from sea to boi.
- Row 2 shows that the L1 link between sea and boi has increased its Feasibility Metric to 23, thus exceeding the L1 circuit's Feasibility Limit. The L1 circuit path reroutes across two links whose total Feasibility Metric total is 20 as shown in the Feasibility Metric Sim column.
- Row 3 shows the Feasibility Metric on the slc-boi L1 link was raised to 15, thus making it impossible for the sea-boi L1 circuit path to route because the sum of the L1 link Feasibility Metrics is now higher than the L1 circuit's Feasibility Limitation. The L1 circuit path's Feasibility Metric Sim changes to "na" and its Feasible Transmission changes to F (false).

Figure 15-4 Example Effects of L1 Link Feasibility Metric on Routing L1 Circuit Paths

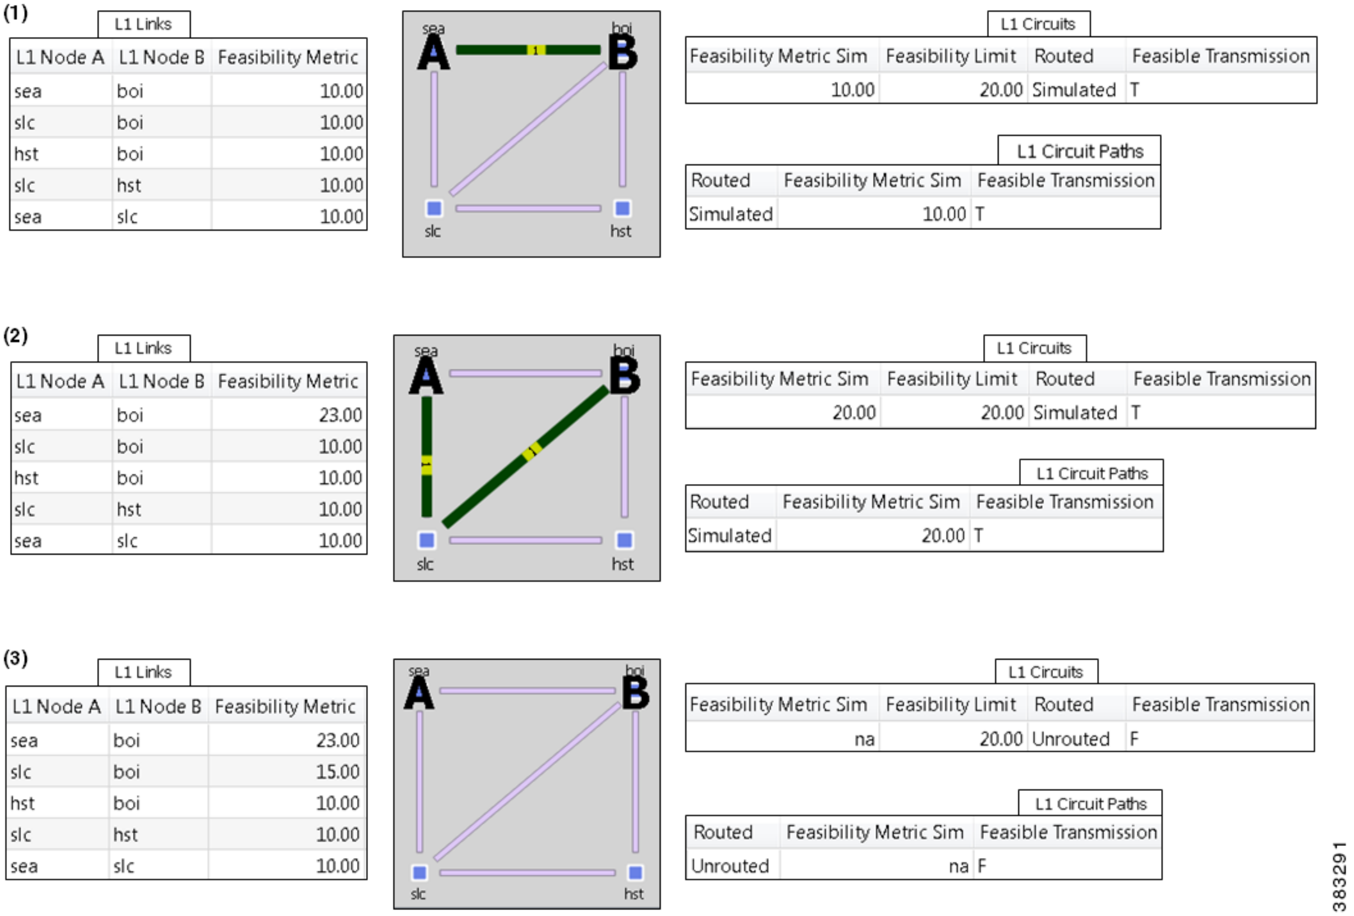

(1)

L1 Links

L1 Node A	L1 Node B	Feasibility Metric
sea	boi	10.00
slc	boi	10.00
hst	boi	10.00
slc	hst	10.00
sea	slc	10.00

L1 Circuits

Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
10.00	20.00	Simulated	T

L1 Circuit Paths

Routed	Feasibility Metric Sim	Feasible Transmission
Simulated	10.00	T

(2)

L1 Links

L1 Node A	L1 Node B	Feasibility Metric
sea	boi	23.00
slc	boi	10.00
hst	boi	10.00
slc	hst	10.00
sea	slc	10.00

L1 Circuits

Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
20.00	20.00	Simulated	T

L1 Circuit Paths

Routed	Feasibility Metric Sim	Feasible Transmission
Simulated	20.00	T

(3)

L1 Links

L1 Node A	L1 Node B	Feasibility Metric
sea	boi	23.00
slc	boi	15.00
hst	boi	10.00
slc	hst	10.00
sea	slc	10.00

L1 Circuits

Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
na	20.00	Unrouted	F

L1 Circuit Paths

Routed	Feasibility Metric Sim	Feasible Transmission
Unrouted	na	F

Example: Figure 15-5 demonstrates the effects of changing the L1 circuit's Feasibility Limit. All L1 links have a Feasibility Metric value of 10. The L1 circuit has two L1 circuit paths: Path Option 1 uses a route of sea-slc-hst-boi, and Path Option 2 uses sea-boi.

- Row 1 shows an L1 circuit Feasibility Limit of 35. The Feasibility Metric Sim of Path Option 1 is 30 since it routes over three L1 links, each having a Feasibility Metric of 10. Path Option 2 has a Feasibility Metric Sim of 10. Since the Feasibility Metric Sim value of each L1 circuit path is under the L1 circuit's Feasibility Limit, each one can route. The Active Path Sim for the L1 circuit is based on the Path Option with the highest priority, which is Path Option 1. All Feasible Transmission values are T (true).

- Row 2 shows that the L1 circuit's Feasibility Limit was lowered to 20. Now the Feasibility Metric Sim of Path Option 1 is greater than the Feasibility Limit, and therefore cannot route. The L1 circuit now uses Path Option 2 as its Active Path Sim.
- Row 3 shows that the L1 circuit has lowered its Feasibility Limit to 5. Now the Feasibility Metric Sim of Path Option 2 is greater than the Feasibility Limit, and therefore cannot route. All Feasible Transmission values become F (false), and all Feasibility Metric Sim values become "na," as does the L1 circuit's Active Path Sim value.

Figure 15-5 Example Effects of L1 Circuit Feasibility Limit on Routing L1 Circuit Paths

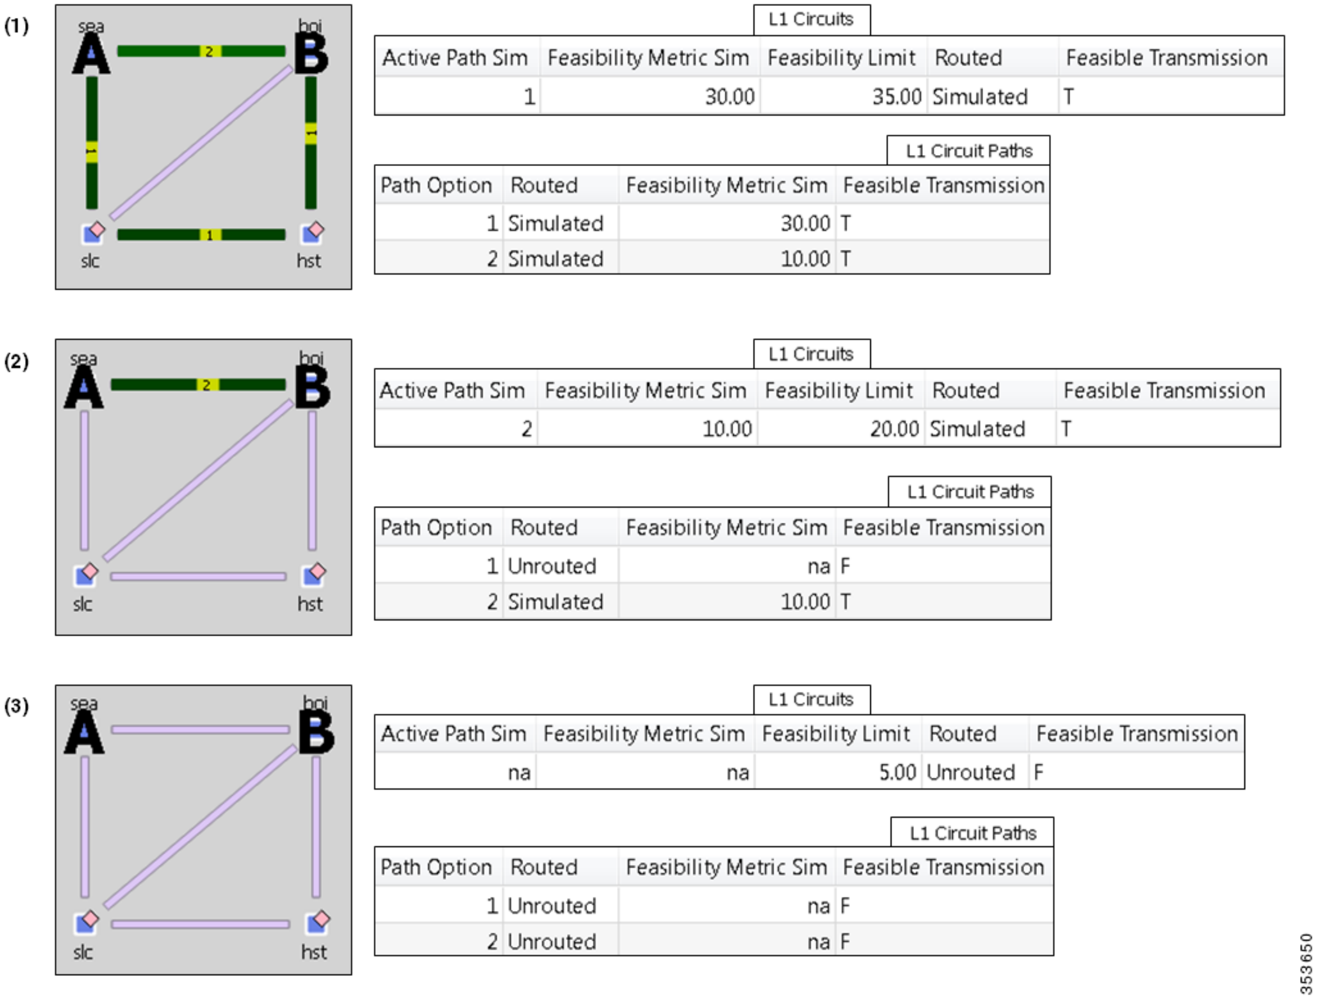

(1)

L1 Circuits

Active Path Sim	Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
1	30.00	35.00	Simulated	T

L1 Circuit Paths

Path Option	Routed	Feasibility Metric Sim	Feasible Transmission
1	Simulated	30.00	T
2	Simulated	10.00	T

(2)

L1 Circuits

Active Path Sim	Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
2	10.00	20.00	Simulated	T

L1 Circuit Paths

Path Option	Routed	Feasibility Metric Sim	Feasible Transmission
1	Unrouted	na	F
2	Simulated	10.00	T

(3)

L1 Circuits

Active Path Sim	Feasibility Metric Sim	Feasibility Limit	Routed	Feasible Transmission
na	na	5.00	Unrouted	F

L1 Circuit Paths

Path Option	Routed	Feasibility Metric Sim	Feasible Transmission
1	Unrouted	na	F
2	Unrouted	na	F

Lambda Assignments

Since multiple L1 circuit paths are transmitted in the same fiber, it is required that they each occupy a different lambda as they cross an L1 link . The lambda corresponds to a specific frequency range (wavelength) in the optical fiber.

In MATE Design, the Lambda property identifies the L1 circuit path wavelength. To simulate changes in the Lambda values (regenerators) along an L1 circuit path, you can manually set Lambda values for each L1 circuit path hop.

- Each L1 circuit path traversing an L1 link must have a unique Lambda while on that L1 link. If more than one L1 circuit path has the same Lambda as they cross the L1 link, lambda blocking occurs. That is, the L1 circuit path with the highest Setup Priority property (lowest number) is routed, and the others are blocked.
- Each L1 circuit path must use the same Lambda value on all L1 links that it traverses unless Lambda values are set on the L1 circuit path hops. This requirement is called the *lambda continuity constraint*.

L1 Object	Feasibility Property	Description
L1 circuit	Lambda Sim	Simulated wavelength number, which is copied from the Lambda Sim of the active L1 circuit path. This value is set to “na” if no L1 circuit path can be routed.
	Lambda Blocking	True (T) means the L1 circuit cannot be routed because there is lambda blocking on all of its L1 circuit paths.
L1 circuit path	Lambda	The Lambda value of the L1 circuit path. • If there is no Lambda value entered, then the Lambda value of the most recently used L1 circuit path hop Lambda value is copied here. • If there are no L1 circuit path hops or if there are no Lambdas assigned to L1 circuit path hops, then this value is set to “na.” The maximum Lambda value (integer) allowed for an L1 circuit path is limited by the lowest Resv Num L1 Circuit Paths Sim value of all L1 links in its path. For more information, see the Reservable L1 Circuits Paths on L1 Links section.
	Lambda Sim	Simulated wavelength number, which is copied from one of these values in this order of priority. • Lambda Sim value of the first actual L1 circuit path hop • Lambda Sim value of the first L1 circuit path hop • Lambda value of this L1 circuit If Auto Lambda is invoked, then MATE Design changes the Lambda Sim in order to route the L1 circuit path. This value is set to “na” if the L1 circuit path cannot be routed.
	Lambda Blocking	True (T) means this L1 circuit path cannot be routed either because it has the same Lambda value as another L1 circuit path going over an L1 link, or the Lambda Sim exceeds the Resv Num L1 Circuits Paths Sim on any L1 link in its path.
	Auto Lambda	If lambda blocking occurs and if Auto Lambda is T (true), the L1 circuit path attempts to use another lambda if it is available. This property is useful for simulating optical restoration when there are L1 failures.

L1 Object	Feasibility Property	Description
L1 circuit path hop	Lambda	The user-specified wavelength of the L1 circuit hop. The Lambda is used for all of the subsequent hops if no other change is specified. This value is set to "na" if it is not explicitly set.
	Lambda Sim	Simulated wavelength number, which is copied from one of these values in this order of priority. • The hop's Lambda value. • The Lambda of the first preceding L1 circuit path hop that has a Lambda value. • The Lambda property of its L1 circuit path. When you manually create L1 circuit path hops, MATE Design automatically creates the first L1 circuit path hop to be the first L1 node in the path. It also copies the L1 circuit path's Lambda Sim value to this hop. If Auto Lambda is used, then this value is set to the Lambda Sim value of the L1 circuit path. This value is set to "na" if no Lambda values have been specified for this L1 circuit path hop or if the L1 circuit path cannot be routed.

Routing Based on Lambdas

Lambda properties affect L1 circuit path routing as follows.

- If the Lambda is not set and cannot be derived, or if there is no lambda blocking, the L1 circuit path is routed.
- If there is lambda blocking, MATE Design checks the Auto Lambda property for each L1 circuit path.
  - If F (false), the L1 circuit path is not routed.
  - If T (true), MATE Design attempts to route the L1 circuit path on the first Lambda that is continuously available on all L1 links over which it is routed. If no such Lambda is found, then the L1 circuit path is not routed.

Note that the Export Lambda Utilization tool enables you to create and export a file that identifies all of the L1 circuit Lambda Sim values being used on each L1 link. This tool is accessed through the File->Export->Lambda Utilization menu. For information, see the *MATE Integration and Development Guide*.

Example: Figure 15-6 shows two L1 circuit paths attempting to route over a common L1 link (chi-nyc).

- Path Option 1 uses sea-boi-chi-nyc.
- Path Option 2 uses sea-slc-hst-okc-chi-nyc.

Both have a Lambda value of 7. Both have Auto Lambda set to F (false). Path Option 1 can route due to its lower Setup Priority value. Path Option 2 cannot route due to Lambda Blocking being T (true). Notice that the Established column is F (false) and the Routed column is Unrouted.

Example: Figure 15-7 shows that Path Option 2 can route after its Auto Lambda property is set to T (true). As a result, its Lambda Sim automatically changes to an unused Lambda value so that there is no lambda blocking. The Lambda Blocking column changes to F (false), the Established to T (true), and the Routed to Simulated. The reason Path Option 1 can still route, even though its Auto Lambda is still set to F, is because it has a lower Setup Priority value. Thus, MATE Design uses Lambda 7 for this path.

Figure 15-6 Example Lambda Blocking

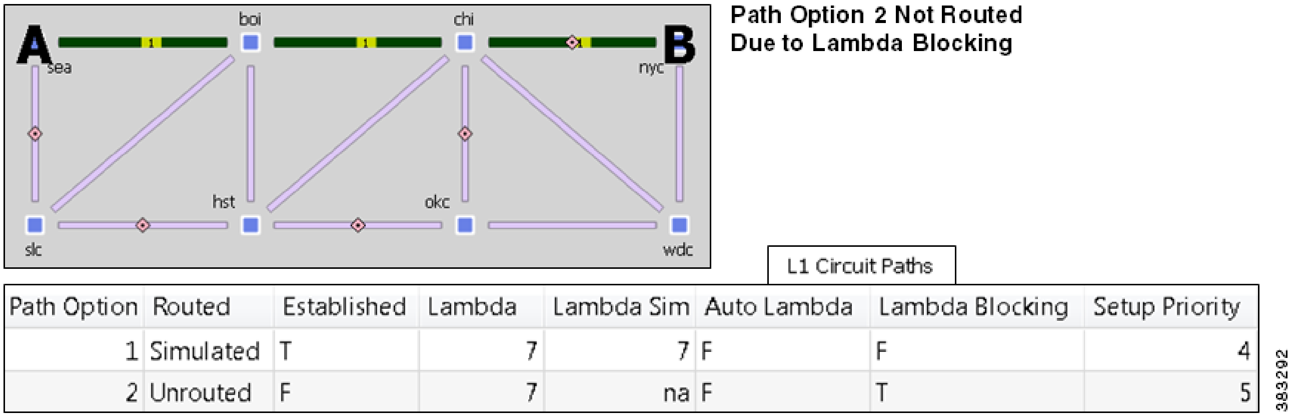

Path Option 2 Not Routed
Due to Lambda Blocking

Path Option	Routed	Established	Lambda	Lambda Sim	Auto Lambda	Lambda Blocking	Setup Priority
1	Simulated	T	7	7	F	F	4
2	Unrouted	F	7	na	F	T	5

Figure 15-7 Example Auto Lambda

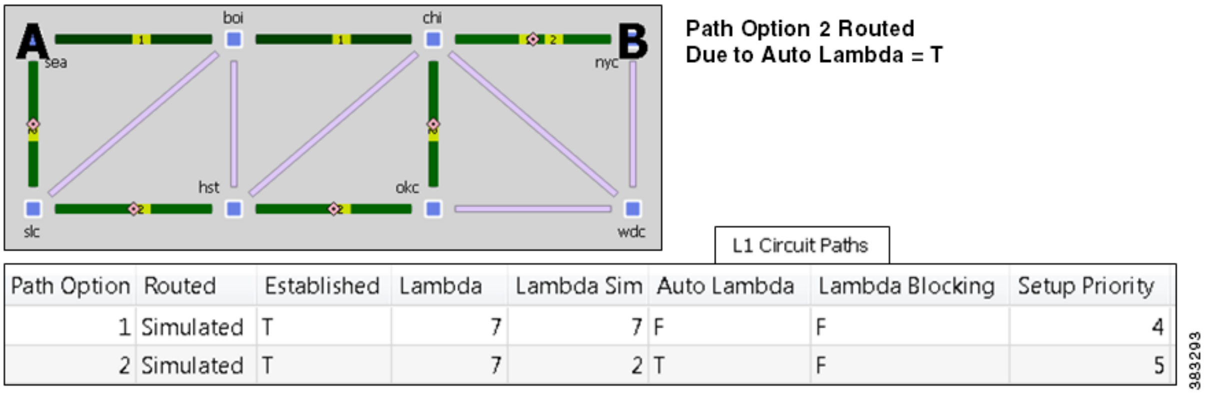

Path Option 2 Routed
Due to Auto Lambda = T

Path Option	Routed	Established	Lambda	Lambda Sim	Auto Lambda	Lambda Blocking	Setup Priority
1	Simulated	T	7	7	F	F	4
2	Simulated	T	7	2	T	F	5

Reservable L1 Circuits Paths on L1 Links

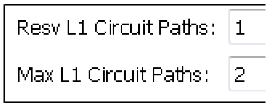

Resv L1 Circuit Paths: 1
Max L1 Circuit Paths: 2

Reserving L1 circuit paths provides another mechanism for influencing the routing of L1 circuit paths. The properties used to reserve L1 circuit paths are the L1 link's Max L1 Circuit Paths and Resv L1 Circuit Paths. To set these properties, right-click on one or more L1 links, and select Properties from the context menu.

L1 Link Properties	Description
Max L1 Circuit Paths	Total number of L1 circuit paths that the L1 link can carry. If not specified, the L1 link has no L1 circuit path restrictions.
Resv L1 Circuit Paths	Number of L1 circuit paths that can be reserved to be routed over the L1 link.

L1 Link Properties	Description
Resv L1 Circuit Paths Sim	(Derived) If a Resv L1 Circuit Paths value is entered, this is copied to the Resv L1 Circuit Paths Sim column. If the Resv L1 Circuit Paths is empty, the Max L1 Circuit Paths value is copied to the Resv L1 Circuit Paths Sim column. The maximum Lambda value (integer) allowed for an L1 circuit path is limited by the lowest Resv Num L1 Circuit Paths Sim value of all L1 links in its path. For more information on lambdas and the potential for the Resv L1 Circuit Path blocking the route of an L1 circuit, see the Lambda Assignments section.
Avail L1 Circuit Paths	(Derived) Number of additional L1 circuit paths that are available to route through the L1 link. This does not include the paths that are currently being routed over the L1 link. The calculation is Resv L1 Circuit Paths Sim minus Num L1 Circuit Paths.
Util Sim	Num L1 Circuit Paths as a percentage of Max L1 Circuit Paths.

- The number of L1 circuit paths using an L1 link cannot exceed the L1 link Resv L1 Circuit Paths value. If there is no Resv L1 Circuit Paths value, then the number of L1 circuit paths using an L1 link cannot exceed the Max L1 Circuit Paths value.
- If multiple L1 circuit paths are attempting to use the same L1 link, then the L1 circuit path with the lowest Setup Priority value is used.

Example: Figure 15-8 shows an slc-hst L1 link with no Max L1 Circuit Paths property. One L1 circuit from slc to wdc filters to three L1 circuit paths that use this slc-hst L1 link.

- By changing the L1 link’s Max L1 Circuit Paths property to 2, only the two L1 circuit paths with the lowest Setup Priority value can route over the L1 link. These are Path Options 1 and 2.
- By adding an L1 link’s Resv L1 Circuit Paths value of 1, only Path Option 1, which has the lowest Setup Priority, can route over this slc-hst L1 link.

Figure 15-8 Example Maximum L1 Circuit Paths and Resv L1 Circuit Paths

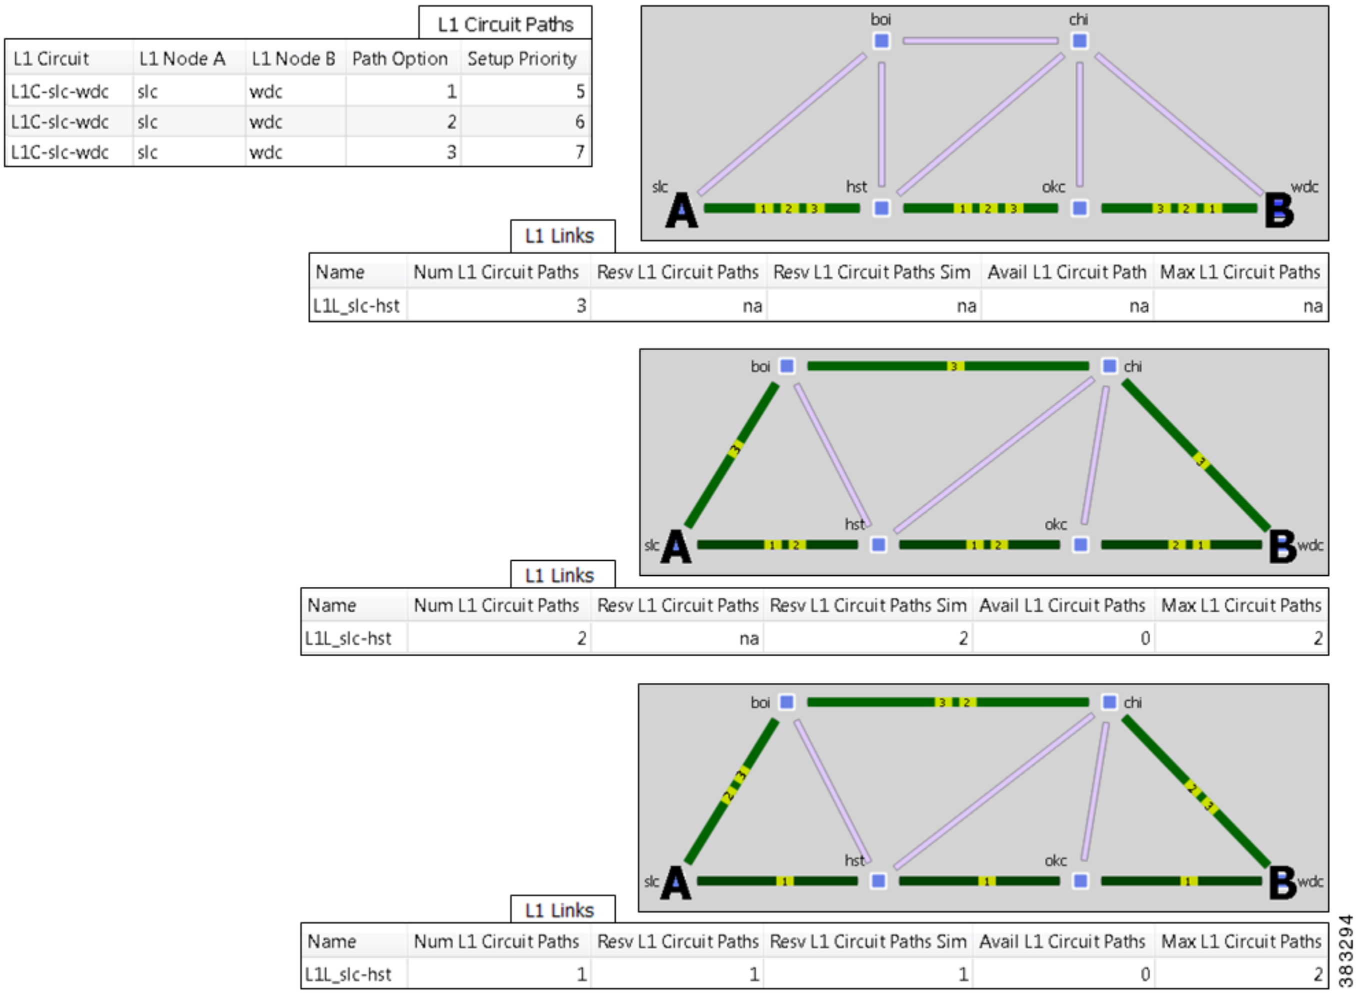

L1 Circuit Paths

L1 Circuit	L1 Node A	L1 Node B	Path Option	Setup Priority
L1C-slc-wdc	slc	wdc	1	5
L1C-slc-wdc	slc	wdc	2	6
L1C-slc-wdc	slc	wdc	3	7

L1 Links

Name	Num L1 Circuit Paths	Resv L1 Circuit Paths	Resv L1 Circuit Paths Sim	Avail L1 Circuit Path	Max L1 Circuit Paths
L1L_slc-hst	3	na	na	na	na

L1 Links

Name	Num L1 Circuit Paths	Resv L1 Circuit Paths	Resv L1 Circuit Paths Sim	Avail L1 Circuit Paths	Max L1 Circuit Paths
L1L_slc-hst	2	na	2	0	2

L1 Links

Name	Num L1 Circuit Paths	Resv L1 Circuit Paths	Resv L1 Circuit Paths Sim	Avail L1 Circuit Paths	Max L1 Circuit Paths
L1L_slc-hst	1	1	1	0	2

Relationships Between L1 and L3

Although L1 circuits and L3 circuits represent the same circuit, they are distinct objects in different tables and views. However, you can map the L1 circuits to the L3 circuits so as to see the effect of failed and inactive L1 links on the L3 network. By mapping L1 to L3 circuits, you are effectively adding L1 properties to a circuit whose L3 properties are defined in the Circuits table. Note that you need not map all L1 circuits to L3 circuits and vice versa.

Similarly, you can map L1 circuits to circuits within a LAG (port channel). Again, this has the effect of adding L1 properties to a port circuit whose L3 properties are defined in the Port Circuits table.

Selecting an L3 circuit or port circuit that has an L1 circuit mapped to it highlights the associated L1 path in green in the L1 view.

Example: This example demonstrates the dynamics between an L1 link and L3 circuits. Figure 15-9 shows an L1 link named kcy-slc that is filtered to its associated L1 circuits. Each L1 circuit path has one L1 circuit path (labeled “1” on the path) where the end L1 nodes are labeled A and B. Figure 15-10 shows the association of these L1 circuits with L3 circuits.

This L1 Link ...	Is Used By These L1 Circuits ...	Which Are Mapped to These L3 Circuits
kcy-slc	chi-sea	cr1.chi-cr2.sea
	chi-sjc	cr2.chi-cr1.sjc
	kcy-sjc	cr1.sjc-cr1.kcy

For More Information ...	See ...
Map L3 circuits to L1 circuits	Map an L1 Circuit to an L3 Circuit
Map L3 circuits to L1 port circuits	Create L3-L1 Links between Nodes

Figure 15-9 Relationship Between L1 Links and L1 Circuits

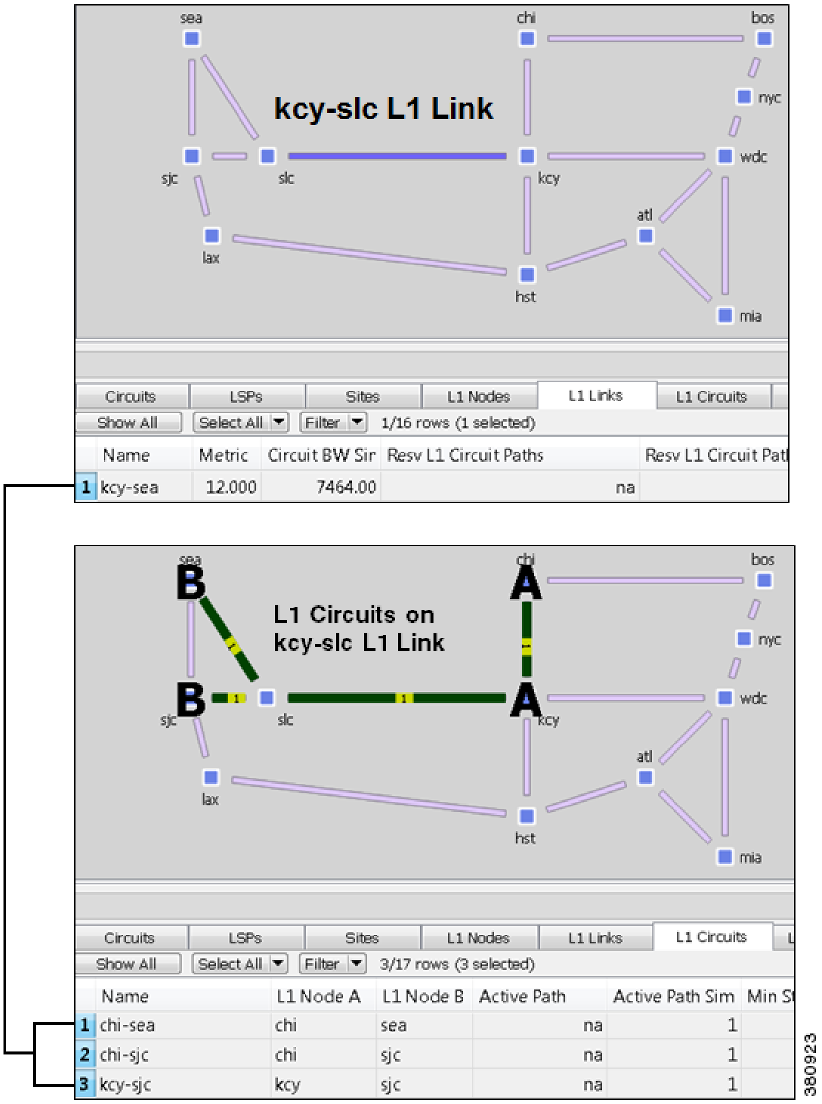

1/16 rows (1 selected)

	Name	Metric	Circuit BW Sir	Resv L1 Circuit Paths	Resv L1 Circuit Pat
1	kcy-sea	12.000	7464.00	na	

3/17 rows (3 selected)

	Name	L1 Node A	L1 Node B	Active Path	Active Path Sim	Min S
1	chi-sea	chi	sea	na	1	
2	chi-sjc	chi	sjc	na	1	
3	kcy-sjc	kcy	sjc	na	1	

Figure 15-10 Relationship Between L1 Circuits and L3 Circuits

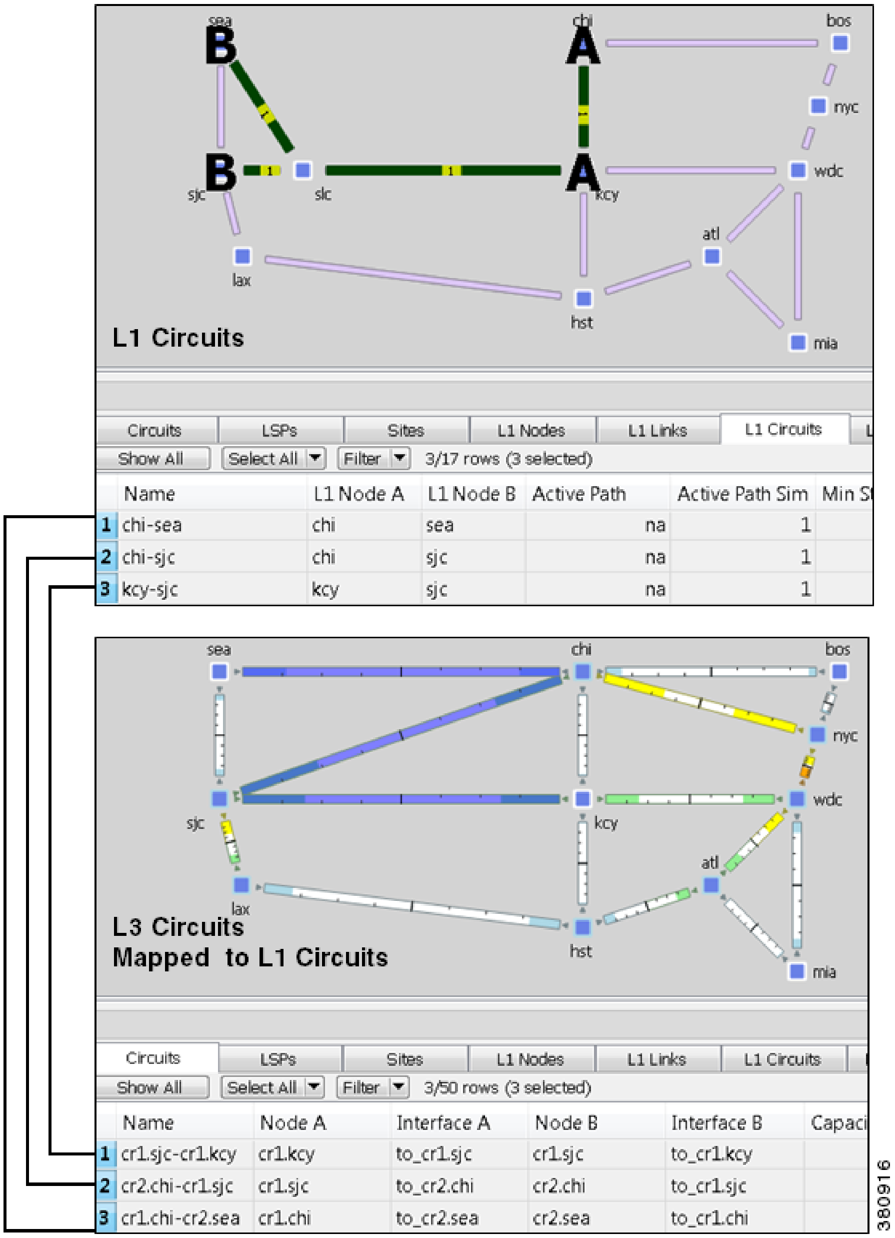

	Name	L1 Node A	L1 Node B	Active Path	Active Path Sim	Min S
1	chi-sea	chi	sea	na	1	
2	chi-sjc	chi	sjc	na	1	
3	kcy-sjc	kcy	sjc	na	1	

	Name	Node A	Interface A	Node B	Interface B	Capaci
1	cr1.sjc-cr1.kcy	cr1.kcy	to_cr1.sjc	cr1.sjc	to_cr1.kcy	
2	cr2.chi-cr1.sjc	cr1.sjc	to_cr2.chi	cr2.chi	to_cr1.sjc	
3	cr1.chi-cr2.sea	cr1.chi	to_cr2.sea	cr2.sea	to_cr1.chi	

Delay and Distance

Both Delay and Distance are properties that can be set in the L1 link and L1 circuit path Properties dialog boxes. To automate the process of setting these, use the Latency and Distance Initializer. For information, see the Simulation chapter.

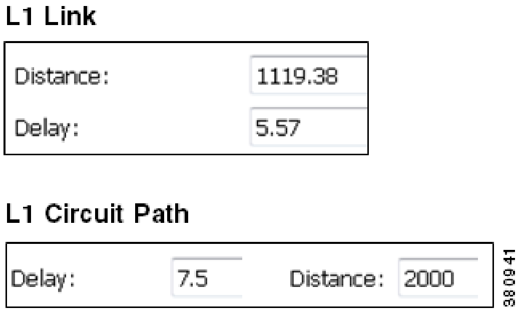

L1 Link

Distance: 1119.38

Delay: 5.57

L1 Circuit Path

Delay: 7.5 Distance: 2000

Table	Column	Description
L1 Links	Delay	One-way transmission latency over the L1 link in milliseconds (ms).
	Distance	Distance of the L1 link in kilometers (km).
L1 Circuits	Delay Sim	• (Derived) This valued is copied from the L1 circuit path Delay Sim column of the active path. • If no L3 Delay is explicitly set for the L3 circuit, the Delay Sim column in the (L3) Circuits table is copied from the Delay Sim column of the mapped L1 circuit, which is based on the L1 circuit path being used as the active path.
	Distance Sim	• (Derived) This valued is copied from the L1 circuit path Distance Sim column of the active path. • If no L3 Distance is explicitly set for the L3 circuit, the Distance Sim column in the (L3) Circuits table is copied from the Distance Sim column of the mapped L1 circuit, which is based on the L1 circuit path being used as the active path.
L1 Circuit Paths	Delay	One-way transmission latency over the L1 circuit path in milliseconds (ms).
	Delay Sim	(Derived) If the L1 circuit path Delay value is entered, it is copied to the Delay Sim column. Otherwise, this is set to the sum of the delays of the L1 links over which this L1 circuit path is routed.
	Distance	Distance of the L1 circuit path in kilometers (km).
	Distance Sim	(Derived) If the L1 circuit path Distance value is entered, it is copied to the Distance Sim column. Otherwise, this is set to the sum of the distances of the L1 links over which this L1 circuit path is routed.

Example: Figure 15-11 shows how L1 link Delay and Distance properties affect L1 circuit path properties (and thus, L1 circuit properties), and how they can potentially affect L3 circuit properties.

- One L1 circuit, which is mapped to an L3 circuit, has two L1 circuit paths.
- Since the active L1 circuit path (Path Option 2) has neither a Delay, nor a Distance set, its Delay Sim and Distance Sim columns are derived by adding together the Distance and Delay values of the L1 links over which it traverses. The combined Delay of these three L1 links is 6.880 ms, and the combined Distance is 900 km. The L1 circuit inherits these in its Distance Sim and Delay Sim columns.

- Since the L3 circuit mapped to the L1 circuit does not have a Delay or Distance property, its Delay Sim and Distance Sim values are copied from the L1 circuit Delay Sim and Distance Sim columns.
- Note that if the L1 circuit's active path were to change to 1 (Path Option 1), the L1 circuit and its mapped L3 circuit would have a Distance Sim of 1400 km and a Delay Sim of 10.880.

Figure 15-11 Example Effect of L1 Link Delay and Distance

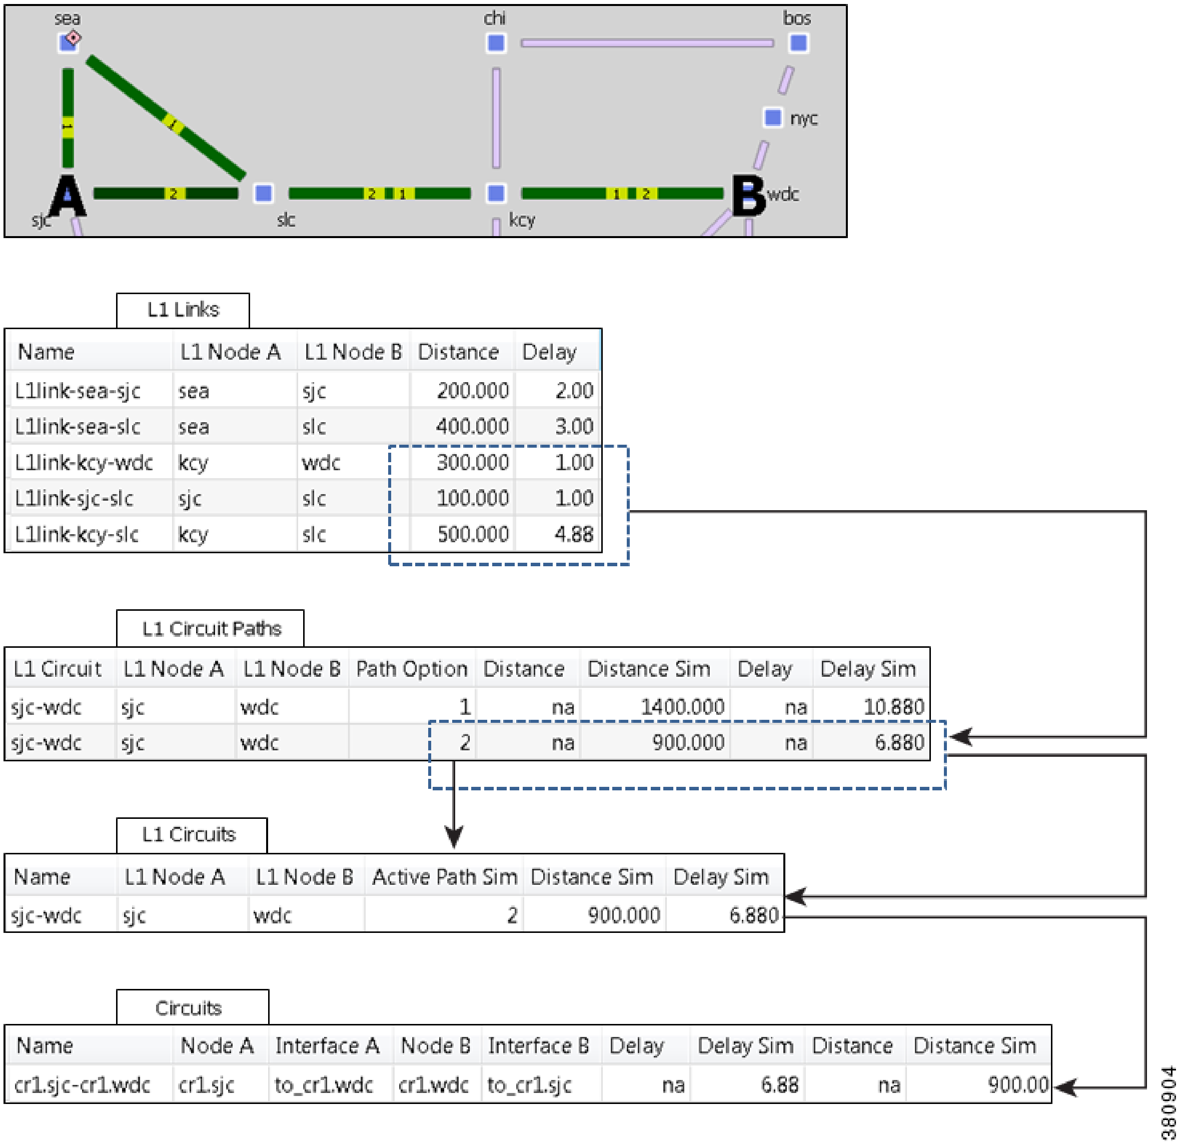

L1 Links

Name	L1 Node A	L1 Node B	Distance	Delay
L1link-sea-sjc	sea	sjc	200.000	2.00
L1link-sea-slc	sea	slc	400.000	3.00
L1link-kcy-wdc	kcy	wdc	300.000	1.00
L1link-sjc-slc	sjc	slc	100.000	1.00
L1link-kcy-slc	kcy	slc	500.000	4.88

L1 Circuit Paths

L1 Circuit	L1 Node A	L1 Node B	Path Option	Distance	Distance Sim	Delay	Delay Sim
sjc-wdc	sjc	wdc	1	na	1400.000	na	10.880
sjc-wdc	sjc	wdc	2	na	900.000	na	6.880

L1 Circuits

Name	L1 Node A	L1 Node B	Active Path Sim	Distance Sim	Delay Sim
sjc-wdc	sjc	wdc	2	900.000	6.880

Circuits

Name	Node A	Interface A	Node B	Interface B	Delay	Delay Sim	Distance	Distance Sim
cr1.sjc-cr1.wdc	cr1.sjc	to_cr1.wdc	cr1.wdc	to_cr1.sjc	na	6.88	na	900.00

L1 Failures and Simulation Analysis

Both L1 links and L1 nodes can be failed, set to an inactive state, and included in Simulation Analysis to calculate a worst-case failure analysis. Additionally, they can both be included in SRLGs, which in turn can also be failed and included in Simulation Analysis. These capabilities enable you to assess the effect of L1 outages on the IP/MPLS network, thus improving your ability to design around and plan for potential service disruptions.

Failure and Inactive States

Just as L1 links and L1 nodes can be set to failed states, they can also be set to inactive states, similarly to L3 circuits and nodes in the L3 topology. Table 15-2 and Table 15-3 show the icons used in the L1 view for these states. Note that sites containing failed or inactive L1 links and L1 nodes indicate as such.

Table 15-2 L1 Link Failure and Inactive States

L1 Link	Description
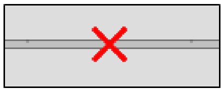	L1 link failed.
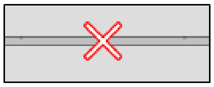	L1 link is not operational because of one of the following reasons. - L1 node failed that is connected to the L1 link. - A site failed that contains the L1 link. - An SRLG failed that contains the L1 link. - An SRLG failed that contains an L1 node connected to the L1 link.
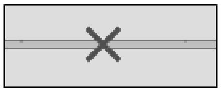	L1 link is inactive.
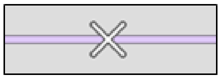	L1 link is inactive for one of the following reasons. - An L1 node that is connected to the L1 link is inactive. - An inactive SRLG contains the L1 link.

Table 15-3 L1 Node Failure and Inactive States

L1 Node	Description
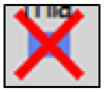	L1 node failed.
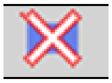	Node is not operational because it is contained in a failed site or SRLG.

L1 Node	Description
☒	L1 node is inactive.
☒	L1 node is not operational because an SRLG containing it is inactive.

For More Information ...	See ...
Fail objects and recover them from failure	Simulation chapter
Set objects to active or inactive	
Run Simulation Analysis	Simulation Analysis chapter
Fail objects to worst-case utilization	
Protect objects nodes from inclusion in Simulation Analysis	

L1 Failure Effects on L3

L3 circuits are rerouted due to L1 failures if the L3 circuit is mapped to an L1 circuit that would traverse a failed L1 node or L1 link under normal operations.

Example: Drawing upon the example L1 link, L1 circuits, and L3 circuits relationships shown in Figure 15-9 and Figure 15-10, Figure 15-12 shows the effects of failing the kcy-slc L1 link. It also shows the effect of an slc L1 node failure. If an L1 node fails, the connected L1 links are non-operational. In this example, both the L1 link and L1 node failure have the effect of making the same L3 circuits non-operational.

Example: Drawing upon the example in Figure 15-12, Figure 15-13 demonstrates what would happen if the kcy-sjc L1 circuit path were rerouted. The L1 circuit path takes a southern route, and the L3 circuit continues to route traffic. The L3 view continues to show active routing, and by selecting the mapped L3 circuit, you can switch to the L1 view to see both the original (solid green) and rerouted, rerouted L1 path (dashed lined). Note that the L3 circuit Delay Sim value (in the Circuits table) might change due to the rerouting of the L1 circuit path.

Figure 15-12 Example Effect of Failing an L1 Link or L1 Node on L3 Circuits

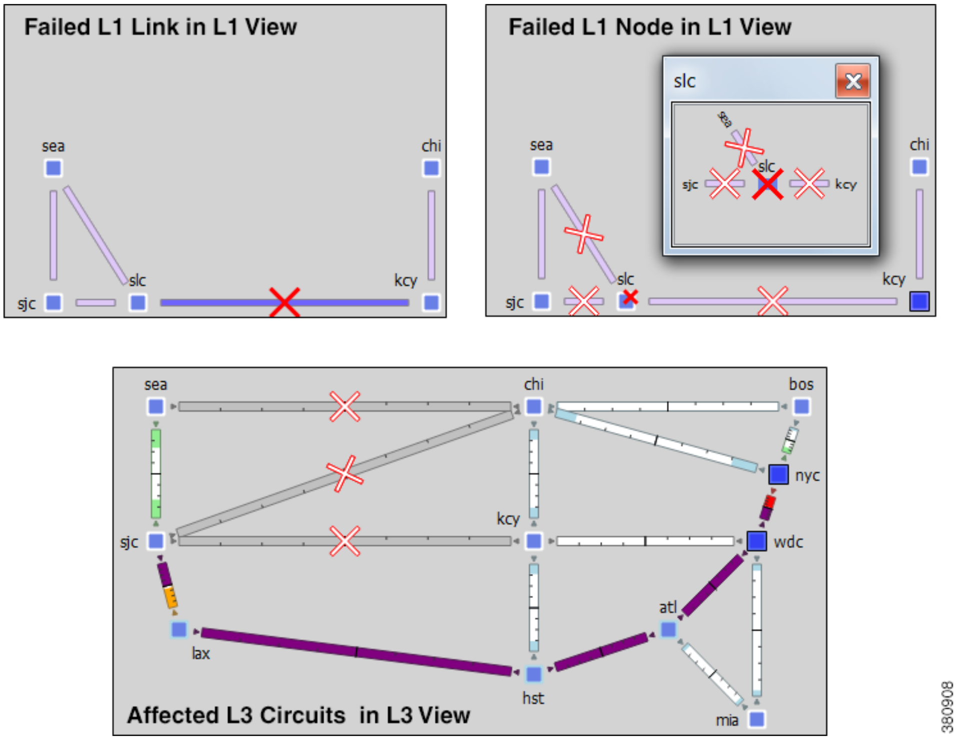

Figure 15-13 Example Effect of Rerouting L1 Circuit Paths

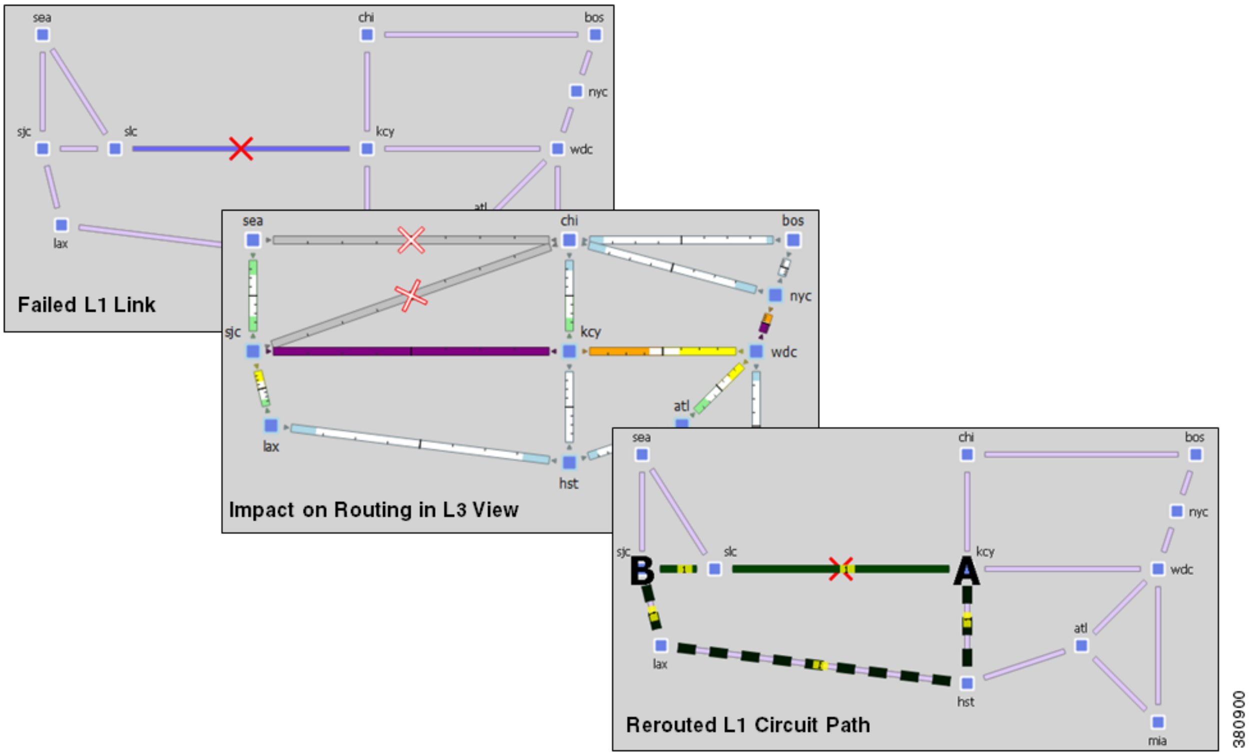

Resiliency under L1 Failures

L1 circuit paths can be rerouted around failures using either standby (static) or non-standby (dynamic) L1 circuit paths. If no restrictions have been specified with respect to path hops, this reroute takes the best path available based on L1 link metrics.

- Standby paths are always established (available) to route. If there are multiple standby L1 circuit paths, then MATE Design attempts to route each available standby L1 circuit path based on their Path Option values (the lower the value, the higher the priority). In this way you can simulate optical circuit protection, such as might be configured for 1+1 protection (one primary and one backup path).
- Non-standby paths are dynamically available. That is, they are available only if MATE Design has attempted to route all standby paths and if the number of routed standby paths is less than the L1 circuit's Min Routed Paths property. Thereafter, MATE Design attempts to route each non-standby L1 circuit path in the order of its Path Option value. This behavior models optical circuit restoration.

For details on how L1 circuit paths are routed, see the L1 Simulations section.

Example: The L1C_chi-atl L1 circuit has three L1 circuit paths (Figure 15-14).

- Path Option 1 is standby, and has no hops.
- Path Option 2 is standby, and has a strict kcy hop, and an exclude hst hop.
- Path Option 3 is non-standby and has no hops.

The non-standby Path Option 3 is not routed until both standby paths are unroutable due to failure. Notice the Routed column changes for all three paths, indicating which paths are routable or simulated. Also notice that the Established column remains consistent. This is because whether or not an L1 circuit path is available for routing under normal operation has no bearing on whether it is routable under failure.

Figure 15-14 Example of Dynamic Non-Standby Path Routing Under Failure

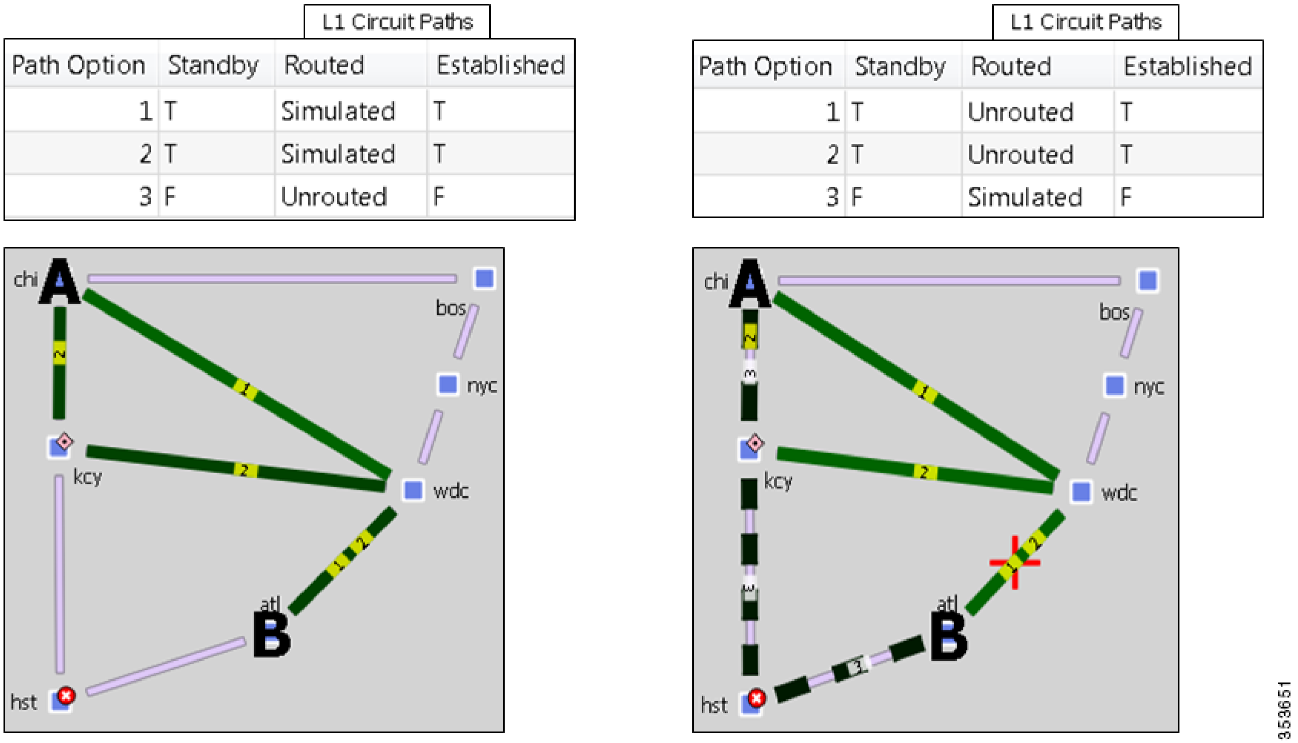

L1 Circuit Paths

Path Option	Standby	Routed	Established
1	T	Simulated	T
2	T	Simulated	T
3	F	Unrouted	F

L1 Circuit Paths

Path Option	Standby	Routed	Established
1	T	Unrouted	T
2	T	Unrouted	T
3	F	Simulated	F

Minimum Routed Paths

The number of L1 circuit paths that MATE Design attempts to route is defined by the L1 circuit's Min Routed Paths property. If the number of routed standby paths is less than the Min Routed Paths property, then MATE Design attempts to route non-standby paths. This can be useful when modeling 1+1+R resilience schemes.

Example: Figure 15-15 shows an L1 circuit with Min Routed Paths set to 1. It has three L1 circuit paths; two are standby and one is non-standby. When one standby L1 circuit path cannot be rerouted under failure, there is no attempt to route the non-standby path. By changing the Min Routed Path to 2, the non-standby path is routed.

Figure 15-15 Example Resilience Scheme Using Minimum Routed Paths Property

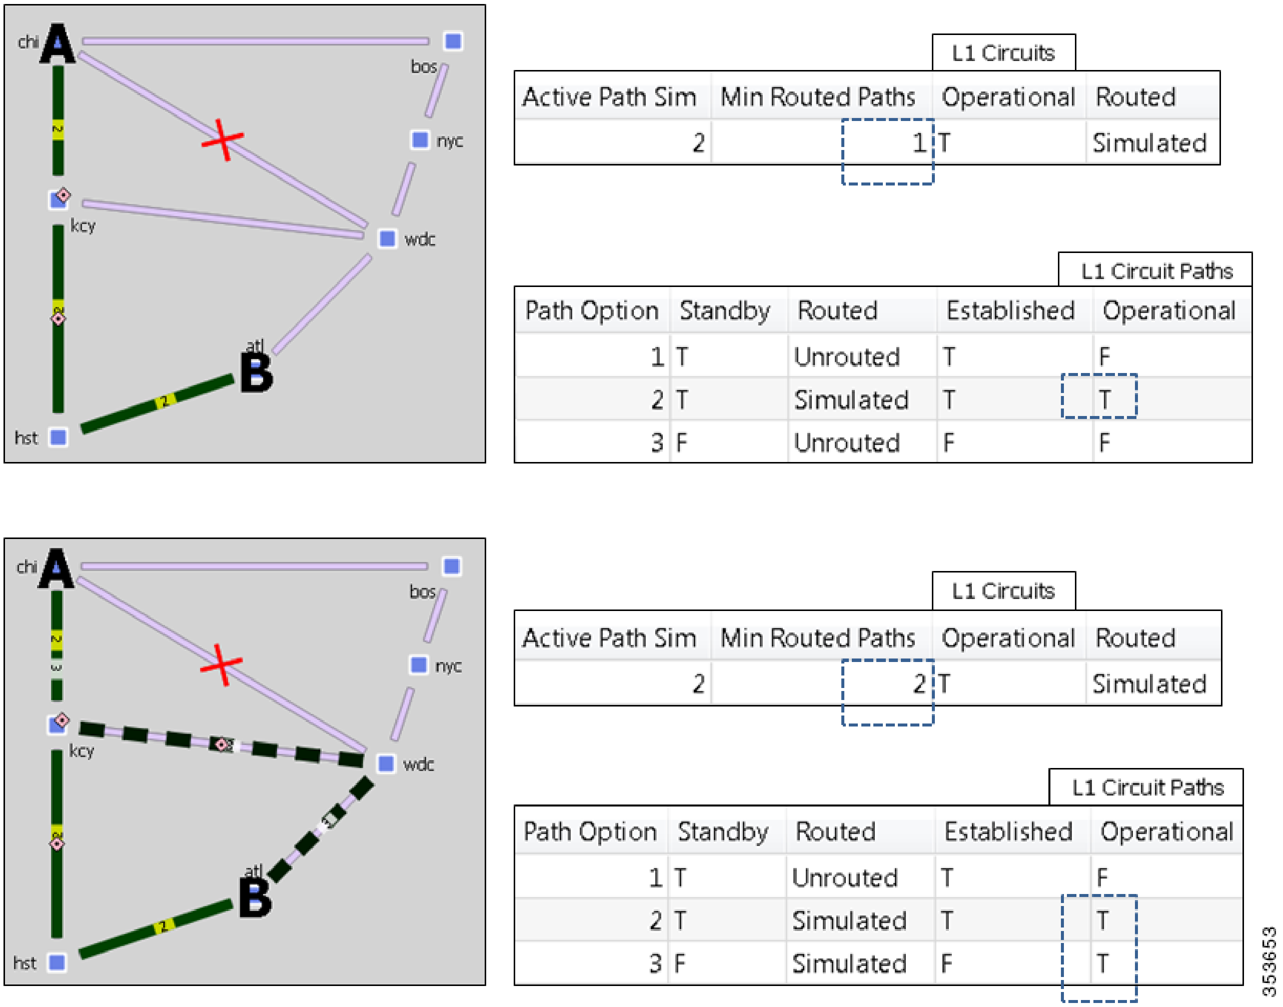

L1 Circuits

Active Path Sim	Min Routed Paths	Operational	Routed
2	1	T	Simulated

L1 Circuit Paths

Path Option	Standby	Routed	Established	Operational
1	T	Unrouted	T	F
2	T	Simulated	T	T
3	F	Unrouted	F	F

L1 Circuits

Active Path Sim	Min Routed Paths	Operational	Routed
2	2	T	Simulated

L1 Circuit Paths

Path Option	Standby	Routed	Established	Operational
1	T	Unrouted	T	F
2	T	Simulated	T	T
3	F	Simulated	F	T

Inactive State

Setting an L1 link or L1 node to inactive reroutes L1 circuit paths using that L1 link or L1 node. It has no effect on the operational state of L3 circuits unless the L1 circuit path cannot be rerouted, in which case the associated L3 circuit becomes operationally down. Note that the L3 circuit Delay Sim value (in the Circuits table) might change due to the rerouting of an L1 circuit path.

Example: Figure 15-16 shows a rerouted hst-wdc L1 circuit path when an L1 link that it traverses is set to inactive. The same L1 circuit path reroute occurs as a result of the atl L1 node being set to inactive.

Figure 15-16 Example Reroute of L1 Circuit when L1 Link or L1 Node is Inactive

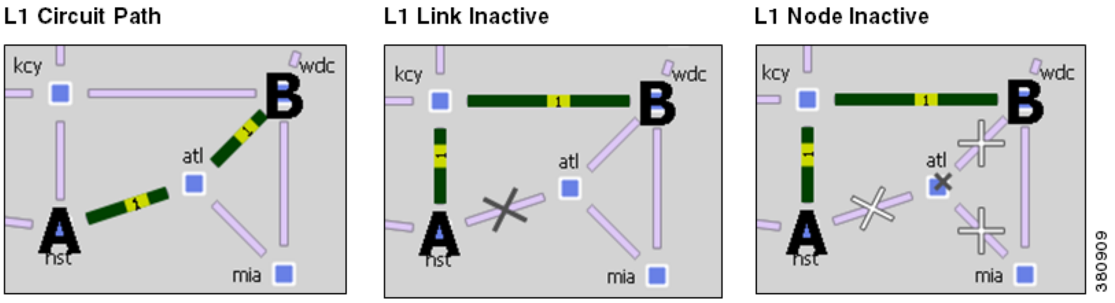

L1 Objects in Simulation Analysis

Like L3 objects, L1 links and L1 nodes can be included in Simulation Analysis either as single objects or as part of SRLGs. If an L1 link or L1 node is the cause of the worst-case utilization, then the L3 view reflects that utilization in the Worst-Case Traffic view.

After running a Simulation Analysis, you can view the L1 failure scenarios that caused the worse-case utilization for an interface. Right-click on the interface or circuit, and select Fail to WC. A dialog box appears from which you can select the SRLG, L1 link, or L1 node failure scenario you want to observe. Upon clicking OK, the network plot changes to show this particular failure scenario. If you are in the L3 view, switch views to see the failure in the L1 view.

For More Information ...	See ...
Run a Simulation Analysis on L1 links and L1 nodes	Simulation Analysis chapter
Fail L1 links and L1 nodes to worst-case utilization	
Protect L1 links and L1 nodes from inclusion in Simulation Analysis	

Create L1 Topology

Create L1 Topology from L3 Topology

MATE Design provides a tool that constructs a complete L1 topology using an existing L3 topology as a model. You have an option to create the topology based on sites or L3 nodes. MATE Design then shows the same site in the L1 view and constructs the following L1 topology.

- Creates L1 nodes.
  - Create L1 nodes from sites—MATE Design creates one L1 node within each selected site regardless of the number of L3 nodes within that site. The L1 nodes are assigned the same name as their sites. No L1 nodes are created for L3 nodes that are not residing in sites.
  - Create L1 nodes from L3 nodes—MATE Design creates one L1 node for each L3 node selected. The L1 nodes are assigned the same name as the L3 nodes. If the L3 node resides within a site, the L1 node is placed there as well.

- Creates L1 links between all L1 nodes whose corresponding L3 nodes have connectivity between them.
- Creates an L1 circuit and L1 circuit path between each pair of L1 nodes connected by L1 links.
- Creates the mapping between L1 circuits and L3 circuits. For more information on these mappings, see Relationships Between L1 and L3.

Follow these steps to create an L1 topology from an L3 topology. You can be in either an L1 or L3 view while performing these steps.

Step 1 If you want to create an L1 topology based on only specific L3 objects, select those sites or L3 nodes from the L3 view. If you do not select sites or nodes, the topology is based on all sites or all L3 nodes, depending on your selection in step #3.

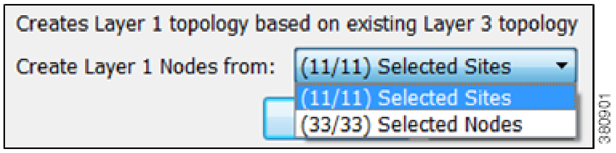

Creates Layer 1 topology based on existing Layer 3 topology

Create Layer 1 Nodes from: (11/11) Selected Sites

Step 2 Either select the Insert->Layer 1->L1 Topology from L3 Topology menu, or right-click on the plot and select New->Layer 1->L1 Topology from L3 Topology from the context menu. A dialog box appears asking you to use selected sites or nodes.

Step 3 Select either sites or nodes, and click OK.

Create L1 Nodes

Step 1 Either select the Insert->Layer 1->L1 Node menu, or right-click on the plot and select New->Layer 1->L1 Node from the context menu. An L1 Node dialog box appears.

Step 2 Enter a unique L1 node name.

Step 3 Either select the site to which you want this L1 node to belong, or select None if you do not want the l1 node to reside in a site.

Step 4 Optional: Enter a description.

Step 5 Optional: Select the Identification tab to enter more descriptive information.

Step 6 Click OK.

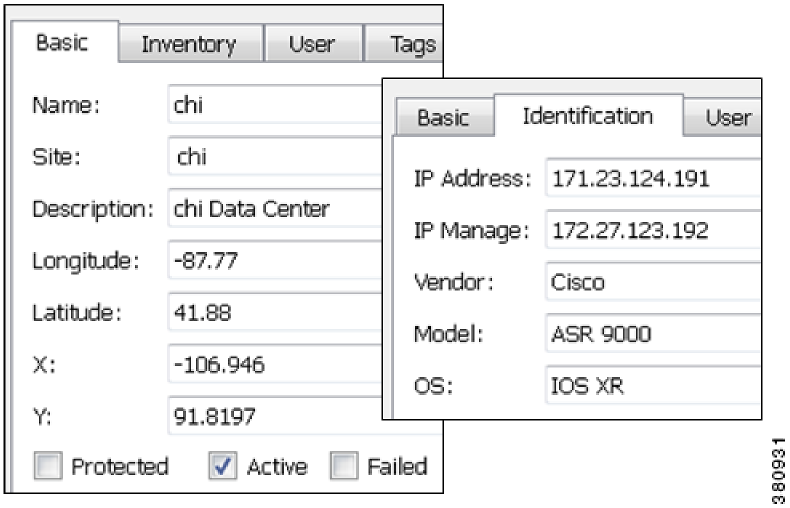

Name: chi
Site: chi
Description: chi Data Center
Longitude: -87.77
Latitude: 41.88
X: -106.946
Y: 91.8197
□ Protected ☑ Active □ Failed

IP Address: 171.23.124.191
IP Manage: 172.27.123.192
Vendor: Cisco
Model: ASR 9000
OS: IOS XR

Create L1 Links

Each L1 link name and L1 node combination must be unique.

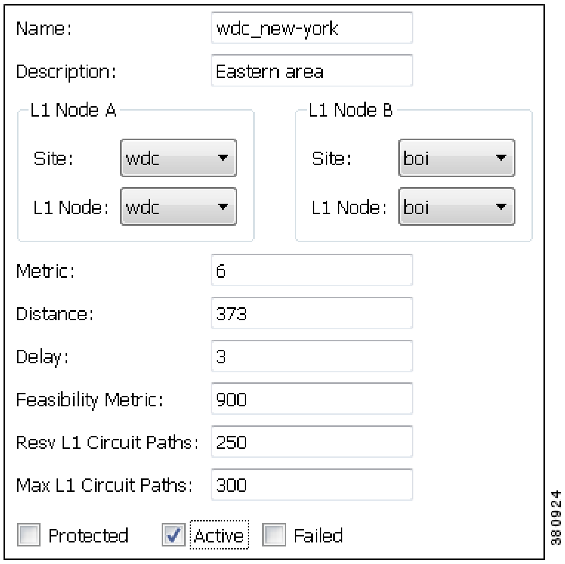
Name: wdc_new-york
Description: Eastern area

L1 Node A
Site: wdc
L1 Node: wdc

L1 Node B
Site: boi
L1 Node: boi

Metric: 6
Distance: 373
Delay: 3
Feasibility Metric: 900
Resv L1 Circuit Paths: 250
Max L1 Circuit Paths: 300
□ Protected ☑ Active □ Failed

Step 1 Either select the Insert->Layer 1->L1 Link menu, or right-click on the plot and select New->Layer 1->L1 Link from the context menu. An L1 Link dialog box appears.

Step 2 Enter an L1 link name and description. If the L1 link has the same source and destination as another L1 link, then it must have a unique name. Otherwise, it is not required.

Step 3 Required: Select the two nodes you want to connect with this L1 link by first selecting the site and then selecting the node.

Step 4 Optional selections are as follows.

Field	For Information ...
Metric	L1 Link Metrics section
Distance, Delay	Delay and Distance section
Feasibility Metric	L1 Circuit Wavelengths section
Resv L1 Circuits, Max L1 Circuits	Reservable L1 Circuits Paths on L1 Links section
Protected, Active, Failed	Simulation and Simulation Analysis chapters

Step 5 Click OK.

Create L1 Circuits

Each L1 circuit name and L1 node combination must be unique.

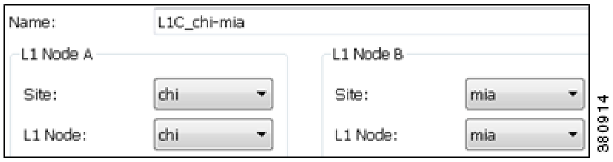

Name: L1C_chi-mia

L1 Node A
Site: chi
L1 Node: chi

L1 Node B
Site: mia
L1 Node: mia

Step 1 Either select the Insert->Layer 1->L1 Circuit menu, or right-click on the plot and select New->Layer 1->L1 Circuit from the context menu. An L1 Circuit dialog box appears.

Step 2 Enter an L1 circuit name. If the L1 circuit has the same source and destination as another L1 circuit, then it must have a unique name. Otherwise, it is not required.

Step 3 Required: Select the two L1 nodes you want to connect with this L1 circuit by first selecting the site (if applicable) and then selecting the L1 node.

Step 4 Optional selections are as follows.

Field	For Information ...
Disjoint Group, Disjoint Priority	Explicit L1 Circuit Path Initializer
Bandwidth	See below
Feasibility Limit	L1 Circuit Wavelengths
Active Path	Active L1 Path Selection
Min Routed Paths	Minimum Routed Paths

Bandwidth—Enter a value for the bandwidth required by the L1 circuit. This value is not used for routing. It is copied to the BW Sim column in the L1 Circuits table. If no value is entered, the BW Sim column is set to the simulated capacity of the L3 circuit that is mapped to this L1 circuit. The total BW Sim for all L1 circuits routing over an L1 Link is listed in the Circuit BW Sim column of the L1 Links table.

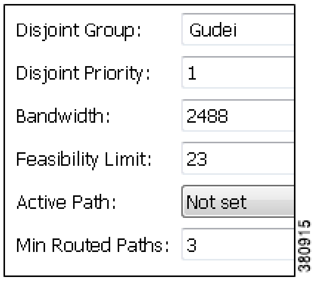

Disjoint Group: Gudei
Disjoint Priority: 1
Bandwidth: 2488
Feasibility Limit: 23
Active Path: Not set
Min Routed Paths: 3

Step 5 Click OK.

Create L1 Circuit Paths

Step 1 Select one or more L1 circuits. If you do not select any, L1 circuit paths are created for all of the L1 circuits.

Step 2 Access an L1 circuit path in one of these ways.

- Select the Insert->Layer 1->L1 Circuit Paths menu.
- Right-click on an L1 circuit in the L1 Circuits table and select Insert L1 Circuit paths from the context menu.
- Right-click on the plot and select New->Layer 1->L1 Circuit Paths from the context menu.

Step 3 Select or enter a path option.

Step 4 Select whether the L1 circuit path is a Standby path, and click OK.

Edit L1 Circuit Paths

If editing L1 circuit paths, you can change the L1 circuit to which the L1 circuit path is mapped, including the sites and L1 nodes. Other options include the following.

Field	For Information ...
Path Options	Standby Paths and Path Options
Delay, Distance	Delay and Distance
Setup Priority	L1 Circuit Path Setup Priority
Auto Lambda, Lambda	L1 Circuit Wavelengths
Standby	Standby Paths and Path Options and Resiliency under L1 Failures
Hops	Explicit L1 Circuit Path Initializer and Create Individual L1 Circuit Path Hops

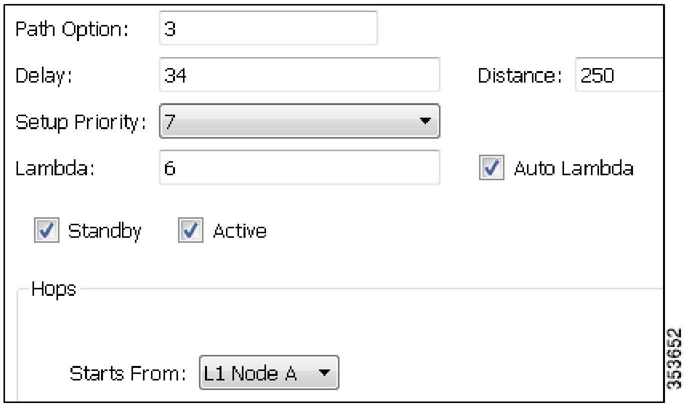
Path Option: 3
Delay: 34
Distance: 250
Setup Priority: 7
Lambda: 6
☑ Auto Lambda
☑ Standby ☑ Active
Hops
Starts From: L1 Node A

Create Individual L1 Circuit Path Hops

To manually create L1 circuit path hops, the L1 nodes and/or L1 links must already exist. Alternatively, you can use the Explicit L1 Circuit Path Initializer to set explicit L1 circuit paths.

Step 1 Either right-click on an L1 circuit path in the L1 Circuits Paths table and select the Properties menu, or double-click on an L1 circuit path. An L1 Circuit Paths dialog box appears.

Step 2 In the Hops area Starts From drop-down list, select the first L1 node hop in the path as either A or B.

Step 3 If you are editing an existing hop, double-click on it from the Hops list.

If you are creating a hop on this path for the first time, click Insert After in the Hops area.

If you are creating a hop on this path and other hops already exist, select an existing hop and then click either Insert Before or Insert After, depending on the required order of the hops. The L1 Circuit Hops dialog box appears.

Step 4 Create or edit the node or link hop as follows.

L1 Circuit Node Hops	L1 Circuit Link Hops
1. In the Node Hop area, select the site and its L1 node that you want to set as the hop. 2. Select the type as loose, strict, or exclude. 3. Click OK.	1. Click the Link Hop option. 2. Select both endpoints (site and node) of the L1 link hop you are setting. 3. Select the L1 link name you are creating. Only those L1 links with the endpoints you selected are listed. 4. Select the type as strict or exclude. 5. Click OK.
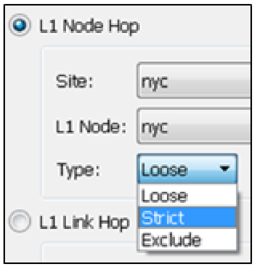 ● L1 Node Hop Site: nyc L1 Node: nyc Type: Loose 380934	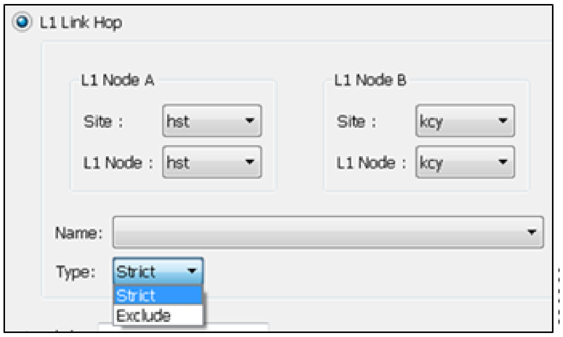 ● L1 Link Hop L1 Node A Site : hst L1 Node : hst L1 Node B Site : kcy L1 Node : kcy Name: Type: Strict

Step 5 Optional: Add a lambda value for this L1 circuit path hop.

Step 6 Repeat steps #3 and #5 for each hop in the path. Since MATE Design always chooses the shortest path, you must enter enough hops to create an alternate, longer path.

Step 7 Click OK.

Step 8 In the L1 circuit path's dialog box, verify that the list of finished hops is in the order of the desired path, and then click OK.

Duplicate L1 Objects

To duplicate an L1 link, L1 circuit, or L1 circuit path right-click on one, and select Duplicate from the context menu. Duplicating an L1 circuit or L1 circuit path also duplicates its L1 circuit path hops.

When duplicating an L3 circuit, you have the option to duplicate its associated L1 circuit if one exists and create an automatic mapping to it. If you select no, the newly duplicated L3 circuit maps to the existing L1 circuit.

Create L1 and L3 Mappings

Create L3-L1 Links between Nodes

You can model the physical relationship between an L3 node and an L1 node to create an L3-L1 link, which is an object listed in the L3-L1 Links table. Creating these relationships can greatly expedite the process of creating L1 circuits from L3 circuits. For information on how these are used in that process, see the Map Multiple L1 Circuits to L3 Circuits section.

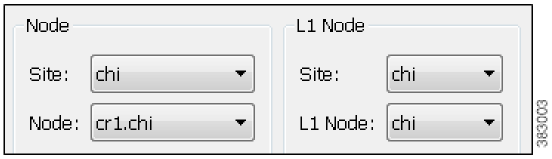

Node
Site: chi
Node: cr1.chi

L1 Node
Site: chi
L1 Node: chi

Note For L3-L1 links to be valid for creating L1 circuits, there must a pair of them. That is, there must be one entry for each L3 end node of the L3 circuit. The L1 circuit that is created uses the L1 nodes to which those L3 nodes are mapped.

Example: This L3-L1 Links table shows an example of three L3-L1 links, as follows.

- L3 node cr1.hst is connected to (linked to) the L1 node hst.
- L3 node cr1.okc is connected to the L1 node okc.
- L3 node cr2.wdc is connected to the L1 node wdc.

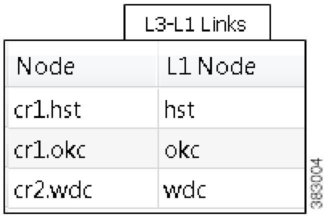

L3-L1 Links

Node	L1 Node
cr1.hst	hst
cr1.okc	okc
cr2.wdc	wdc

Step 1 Either select the Insert->Layer 1->L3-L1 Link menu, or right-click in an empty area and select New->Layer 1->L3-L1 Link from the context menu.

Note that pre-selecting an L3 node or L1 node has no effect on the selections opened in the dialog box.

Step 2 Select the L3 node and the site in which it is contained, if applicable.

Step 3 Select the L1 node and the site in which it is contained, if applicable.

Step 4 Click OK.

Map Multiple L1 Circuits to L3 Circuits

Note For L3-L1 links to be valid for creating L1 circuits, there must a pair of them. That is, there must be one entry for each L3 end node of the L3 circuit. The L1 circuit that is created uses the L1 nodes to which those L3 nodes are mapped.

The "L1 Circuits for L3 Circuits" tool creates L1 circuits based on L3 circuits and associates them with the L3 circuits using the following sequence of priorities. For each L1 circuit created, an associated L1 circuit path is also created. Nested objects are taken into account.

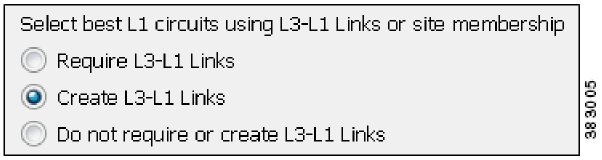
Select best L1 circuits using L3-L1 Links or site membership

○ Require L3-L1 Links

● Create L3-L1 Links

○ Do not require or create L3-L1 Links

- If the L3 circuit already has an associated L1 circuit, one is not created.
- MATE Design checks the L3-L1 Links table for valid entries.
  - If multiple pairs of L3-L1 links exist that could create the L1 circuit, one is randomly used.
  - If no valid L3-L1 link entry exists and if "Require L3-L1 Links" is selected, no L1 circuit is created.
  - If no valid L3-L1 link entry exists and if "Do not require or create L3-L1 Links" is selected, for each L3 circuit MATE Design checks for sites that contain both the L3 end node and at least one L1 node.
    - If multiple L1 nodes are co-located within the same sites as the L3 nodes, then one is randomly selected.
    - If this co-location is not found, MATE Design checks for L1 nodes within the parent sites of the sites containing the L3 end nodes. If none are found, an L1 circuit is not created.
- If "Create L3-L1 Links" is selected, then the L3-L1 links are created that map the respective L3 and L1 nodes.

If needed, you can later make changes to the L1 circuit's endpoints and the L3-L1 links that are created by editing their respective Properties dialog boxes.

Examples

Figure 15-17 shows a plan file containing six L3 circuits and four L3-L1 links. All sites except sea contain one L3 node and one L1 node. The sea site contains two L3 nodes and one L1 node. Refer to the diagram for details of the objects' relationships and to see the objects in the L3-L1 Links table. There are no L1 circuits.

Figure 15-17 Topology and Existing L3-L1 Links

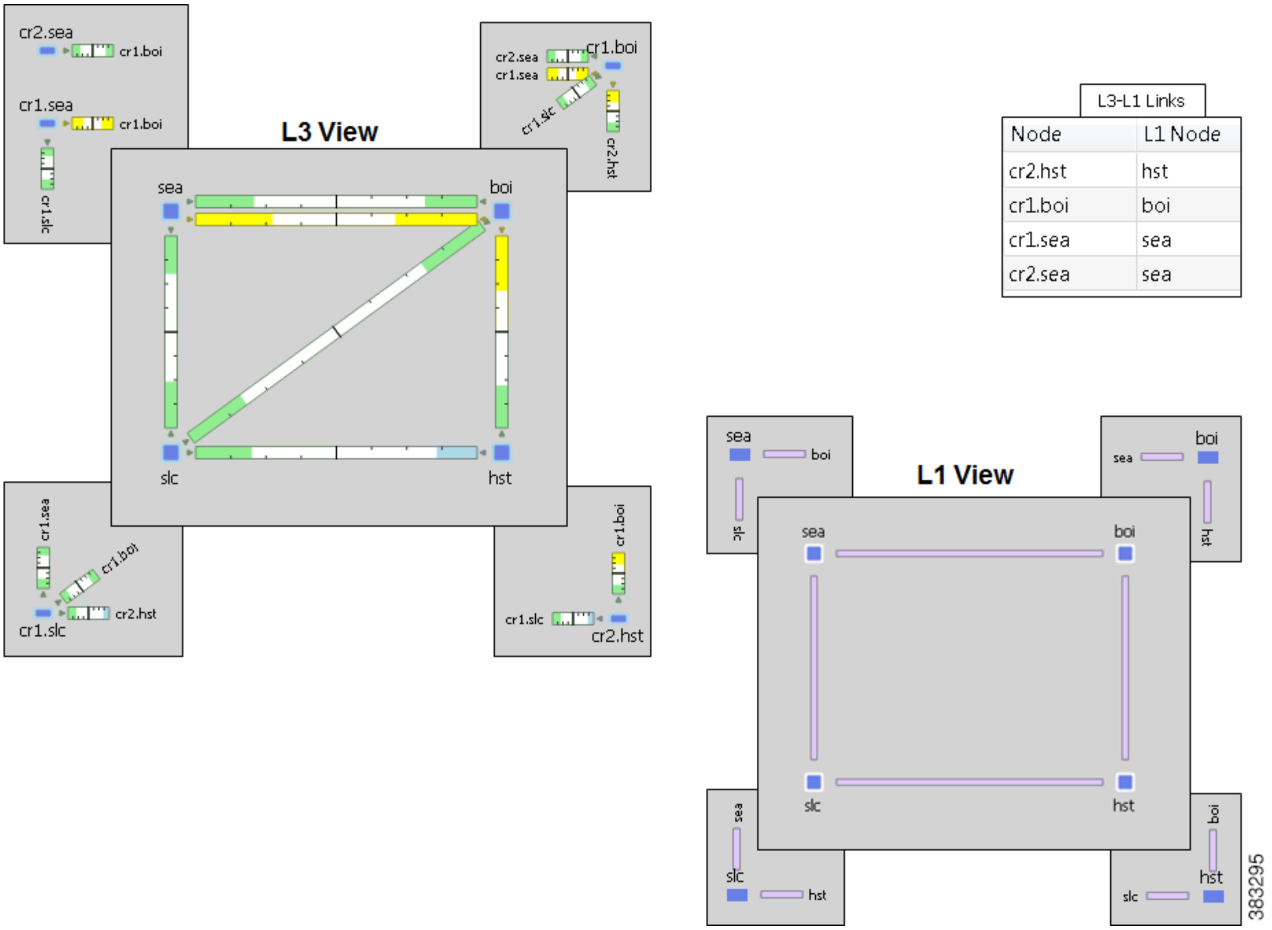

L3-L1 Links

Node	L1 Node
cr2.hst	hst
cr1.boi	boi
cr1.sea	sea
cr2.sea	sea

If you select the Require L3-L1 Links option, MATE Design creates three L1 circuits: one between hst and boi, and two between from sea and boi.

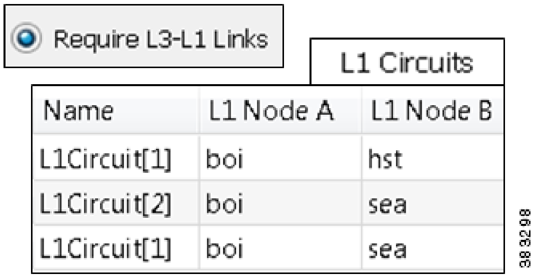

● Require L3-L1 Links

L1 Circuits

Name	L1 Node A	L1 Node B
L1Circuit[1]	boi	hst
L1Circuit[2]	boi	sea
L1Circuit[1]	boi	sea

If you select not to require or create L3-L1 Links, MATE Design creates six L1 circuits: three using the L3-L1 Links table and the remaining three by determining that the source and destination sites both contain L3 and L1 nodes.

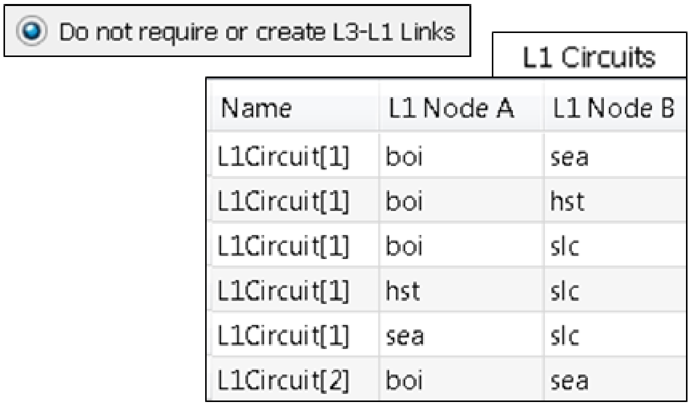

● Do not require or create L3-L1 Links

L1 Circuits

Name	L1 Node A	L1 Node B
L1Circuit[1]	boi	sea
L1Circuit[1]	boi	hst
L1Circuit[1]	boi	slc
L1Circuit[1]	hst	slc
L1Circuit[1]	sea	slc
L1Circuit[2]	boi	sea

Using the same topology as in Figure 15-17 but with no pre-existing L3-L1 Links, if you select the Create L3-L1 Link option, MATE Design creates the same six L1 circuits as if you had selected no options. However, it also creates five L3-L1 Links: one for each L3 node having an L3 circuit for which an L1 circuit was created.

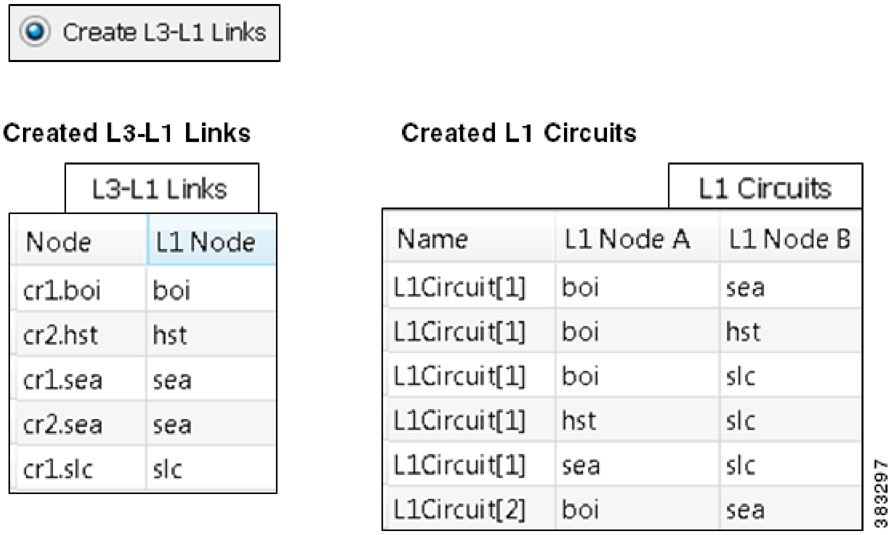

● Create L3-L1 Links

Created L3-L1 Links

L3-L1 Links

Node	L1 Node
cr1.boi	boi
cr2.hst	hst
cr1.sea	sea
cr2.sea	sea
cr1.slc	slc

Created L1 Circuits

L1 Circuits

Name	L1 Node A	L1 Node B
L1Circuit[1]	boi	sea
L1Circuit[1]	boi	hst
L1Circuit[1]	boi	slc
L1Circuit[1]	hst	slc
L1Circuit[1]	sea	slc
L1Circuit[2]	boi	sea

Process for Creating Multiple L1 Circuits from L3 Circuits

Step 1 In this process, you can optionally use L3-L1 links. If you are using them, verify you have the set of L3-L1 links you need, and if not, create them. For information, see the Create L3-L1 Links between Nodes section.

Step 2 If you want to create L1 circuits for only selected L3 circuits, first select those L3 circuits. Otherwise, L1 circuits are created for as many L3 circuits as possible.

Step 3 Either right-click in an empty area, and select Layer 1 -> L1 Circuits for L3 Circuits from the context menu, or select the Insert-> Layer 1->L1 Circuits for L3 Circuits menu.

Step 4 Select either to require the use of existing L3-L1 Links, to create them, or neither.

Step 5 Click OK.

Map an L1 Circuit to an L3 Circuit

Step 1 Either right-click on an L3 circuit in the Circuits table and select Properties, or double-click on an L3 circuit.

Step 2 Select the Advanced tab.

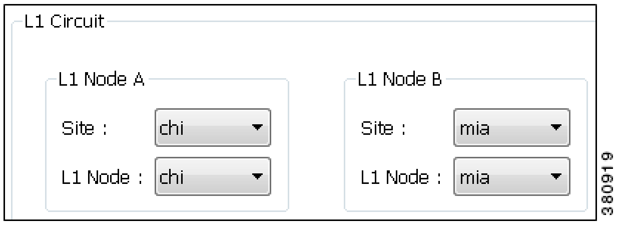

L1 Circuit

L1 Node A
Site : chi
L1 Node : chi

L1 Node B
Site : mia
L1 Node : mia

Step 3 In the L1 Circuit area, select both L1 Node A and L1 Node B site and nodes.

Step 4 Either create and map a new L1 circuit to the L3 circuit, or map an existing one to the L3 circuit.

- Map existing L1 circuit—Choose the Select option, and select the L1 circuit from the drop-down list.
- Create L1 circuit—Choose the Create option, and enter a new name for the L1 circuit.

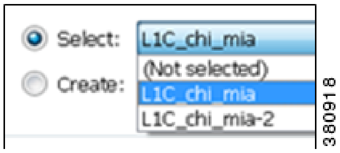

● Select: L1C_chi_mia
○ Create:

(Not selected)
L1C_chi_mia
L1C_chi_mia-2

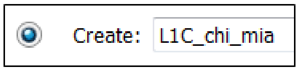

● Create: L1C_chi_mia

Step 5 Click OK.

Map an L1 Node to an L3 Port

Step 1 Either right-click on port in the Ports table and select Properties, or double-click on a port.

Step 2 Select the Layer 1 tab.

Step 3 If applicable, select the L1 node's site.

Step 4 Select the L1 Node A, and click OK.

Map an L1 Circuit to an L3 Port Circuit

Step 1 Either right-click on port circuit in the Port Circuits table and select Properties, or double-click on a port circuit.

Step 2 Select the Layer 1 tab.

Step 3 In the L1 Circuit area, select both L1 Node A and L1 Node B site and nodes.

Step 4 Either create and map a new L1 circuit to the L3 circuit, or map an existing one to the L3 circuit.

- Map existing L1 circuit—Choose the Select option, and select the L1 circuit from the drop-down list.
- Create L1 circuit—Choose the Create option, and enter a new name for the L1 circuit.

Step 5 Click OK.

Visualize L1 Network

Many of the visualization tools available for the L3 view are also available for L1, as follows. These tools are described in the *MATE GUI Visualization Guide*.

- Assign Sites to Nodes—Automatically create sites based on L1 node names and a set of parameters that define which part of the node name to use.
- Arrange Sites and Nodes—Arrange sites and nodes using tools that offer a variety of alignment options or arrange sites to specific geographic locations. You can also show and hide sites, as well as change where their site name is displayed in the plot.
- Set Plot Backgrounds—Set the background for the network plot, including schematic view, outline of a geographic background, and detailed geographic background.
- Create or Edit Layouts—Create multiple Design layouts for a single plan file, each containing its own set of visual properties.

Edit L1 Plot Options

The Plot Options dialog box has a L1 tab that enables you to change the link appearance, as well as change which sites appear in the L1 view. Since these are plot options, they apply to the current layout and are saved on a per-layout basis when you save the plan file. You can also save them as user defaults. For information on user and system defaults, see the *MATE GUI Visualization Guide*.

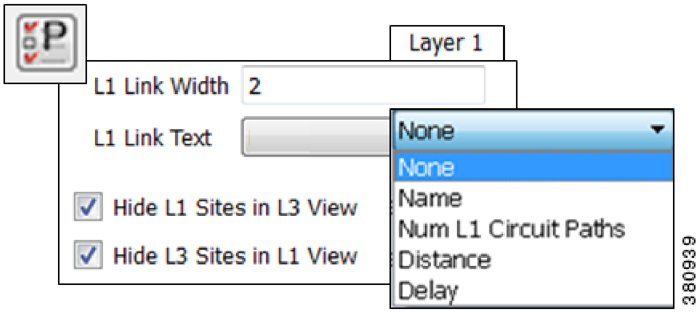

L1 Link Width 2

L1 Link Text

☑ Hide L1 Sites in L3 View

☑ Hide L3 Sites in L1 View

Step 1 Click on the Plot Options icon in the top, right of the Visualization toolbar, or select the Edit->Plot->Plot Options->Edit Options menu.

Step 2 Select the Layer 1 tab.

Step 3 Edit Layer 1 fields as follows, and then click OK.

Change L1 Link Appearance	Show or Hide Sites
• L1 Link Width—The L1 link width is based on its capacity (the higher the capacity, the wider the circuit). Use this option to set the minimum width and the capacity it represents. • L1 Link Text—By default, there is no text on L1 links. Use this option to display the link name, distance, delay, or the number of L1 circuit paths using the link.	By default, sites containing only L1 nodes are hidden in the L3 view, and sites containing only L3 nodes are hidden in the L1 view. Use these selections to toggle whether to hide sites containing only L1 or L3 nodes.

CHAPTER 16

Metric Optimization

The Metric Optimization tool optimizes metrics in the network core. It replaces the IGP metrics of core interfaces in a network plan with new metrics chosen to maximize the throughput achievable by the resultant routes, under a specified set of failure scenarios.

The Tactical Metric Optimization tool enables you to quickly fix local congestion by doing as few modifications as possible on selected interfaces to push utilizations under a given level.

Note The core of the network must be connected, and all nodes must be contained in one AS. That is, there must be a path between any two core nodes that passes only through other core nodes in the same AS, and not through edge nodes. If this is not the case, Metric Optimization signals an error

Core Versus Edge

MATE plan files allow nodes to be defined as either core nodes or edge nodes. Note that these definitions are not related to any explicit router configurations. A similar classification of circuits is determined implicitly by the nodes they connect. If a circuit is attached to one or more edge nodes, it is defined to be an edge circuit. Otherwise, it is a core circuit. Edge metrics and core metrics mean metrics of edge or core circuits, respectively.

This core and edge distinction helps define certain policies followed by MATE Design tools. In Metric Optimization, these are as follows.

- Only core metrics are changed.
- The objective is to minimize utilization in the core. Utilization in the edge is not taken into account.
- The routes selected by the metrics are restricted by the tool so that any route between a source and a destination will enter the core at most once, and leave the core at most once. That is, *edge leakage* of routes is prevented.

A consequence of these policies is that larger edge metrics in the plan file to be optimized might result in better optimized core routes. With larger given edge metrics, the Metric Optimization tool has more flexibility in how it can set the core metrics without leakage of routes into the edge.

Metric Optimization

There are many trade-offs in the optimization performed by Metric Optimization. The main trade-off is between performance under normal operation versus performance under failure. Metrics can be selected to do well under the former while ignoring the latter, or to attempt to balance performance between the two.

Metric Optimization optimizes undifferentiated traffic, but does not optimize traffic per service class or per interface queue. If there are policies set on the undifferentiated traffic, then MATE Design attempts to minimize both the interface Utilization and the policy violations.

To open the tool, select the Tools->Metric Optimization menu.

For More Information...	See...
IGP simulations	IGP Simulation chapter
Demands and demand latency bounds	Traffic Demand Modeling chapter
Simulation Analysis and protecting objects from being included in the analysis	Simulation Analysis chapter
Layer 1 objects and simulation	Layer 1 Simulation chapter

Normal (Non-Resilient) Options

The optimization objectives for normal operation are as follows, in order of priority.

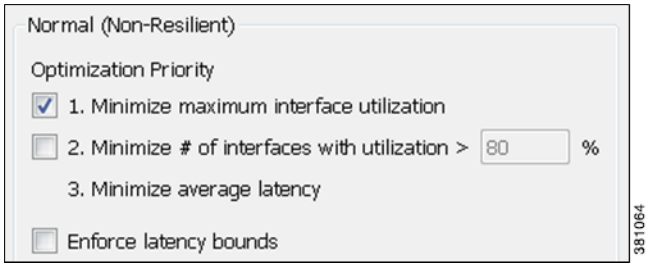

Normal (Non-Resilient)

Optimization Priority

☑ 1. Minimize maximum interface utilization

□ 2. Minimize # of interfaces with utilization > 80 %

3. Minimize average latency

□ Enforce latency bounds

- Minimize maximum interface utilization—This default is usually the primary objective, and concentrates on reducing the utilization of the interface with the highest utilization.
- Minimize # of interfaces with utilization > *x* % —Often a small number of interfaces are bottlenecks in the network, and their utilizations cannot be reduced through rerouting. However, it is still desirable to reduce the utilization of other, under-utilized interfaces. This option attempts to reduce all interface utilizations to below a specified utilization level, or, if this is not possible, to minimize the number of interfaces over this level.
- Minimize average latency—If after the optimizations flexibility is available in the choice of metrics, then MATE Design uses this flexibility to minimize the average latency of the routes across the network. This is not user-configurable and is always on.
- Enforce latency bounds—If selected, the choice of paths for each demand is restricted, if possible, to those demands with latency below the latency bound. For information on setting latency bounds, see the Traffic Demand Modeling chapter.

Failure (Resilient) Options

For planning purposes, the most benefit from Metric Optimization is that it is able to optimize routing over a large number of failure scenarios. This capability controls that optimization, as well as controls the trade-off between normal and failure optimizations.

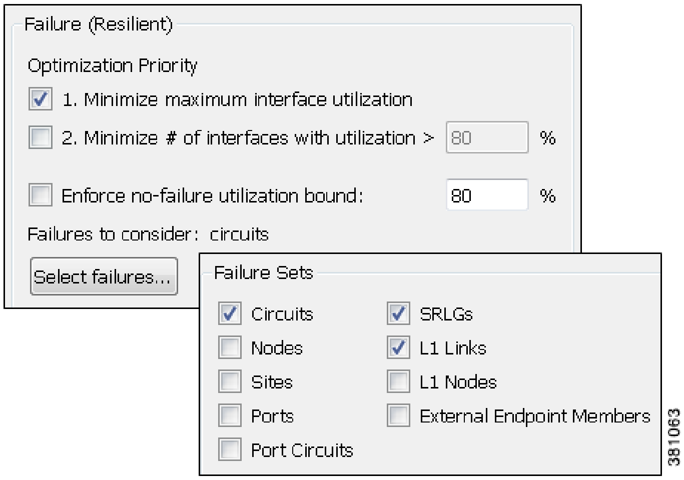

Failure (Resilient)

Optimization Priority

☑ 1. Minimize maximum interface utilization

□ 2. Minimize # of interfaces with utilization > 80 %

□ Enforce no-failure utilization bound: 80 %

Failures to consider: circuits

Select failures...

Failure Sets

☑ Circuits ☑ SRLGs

□ Nodes ☑ L1 Links

□ Sites □ L1 Nodes

□ Ports □ External Endpoint Members

□ Port Circuits

- Minimize maximum interface utilization—If selected, MATE Design selects metrics to minimize the maximum interface utilization over all interfaces and all failure scenarios. If not selected, only normal operation is optimized, and the remainder of the failure section is ignored.
- Minimize # of interfaces with utilization > *x* % —The number of interfaces with utilization (under any failure scenario) above a specified percentage is minimized.
- Enforce no-failure utilization bound—If not selected, maximum failure utilization is minimized with no consideration for utilization under normal operation. It is often convenient to select this option and place a bound on the utilization under normal operation, which Metric Optimization then satisfies, if possible.
- Failures to consider—This option determines the failure scenarios over which optimization is performed. You can select any combination of failures.
  - All unprotected circuit, node, Layer 1 (L1) node, and L1 link failures
  - All SRLG failures

  This option does not consider port or port circuit failures. For information on protecting objects from being included in this analysis, see the Simulation Analysis chapter.

Other Options

You can use these options to fine-tune the optimization.

- Optimization Type—MATE Design can perform two different types of metric optimizations: global or incremental.

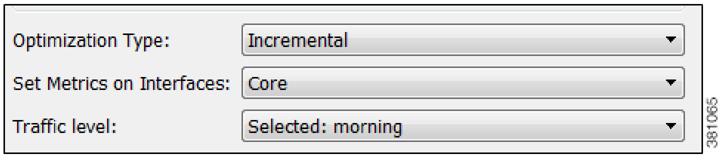

Optimization Type: Incremental

Set Metrics on Interfaces: Core

Traffic level: Selected: morning

  - Incremental optimization uses the current routes as a basis from which improvements can be made. The target routes are not changed as radically as when using global optimizations, and therefore the number of metric changes from the metric targets is typically much smaller using the incremental option.

  - Global optimization creates optimal routes from scratch, and targets the old metrics afterward. Because MATE Design can choose routes globally, this option tends to perform better under some circumstances, particularly when the target routes are poorly chosen. However, global optimization typically takes longer than incremental optimization.

- Set Metrics on Interfaces—By default, MATE Design sets or changes metrics on any of the core interfaces in the network to achieve its optimization goals. Alternatively, as provided by this option, Metric Optimization restricts metric changes to the following interfaces: all, core, selected in table, or tag name.

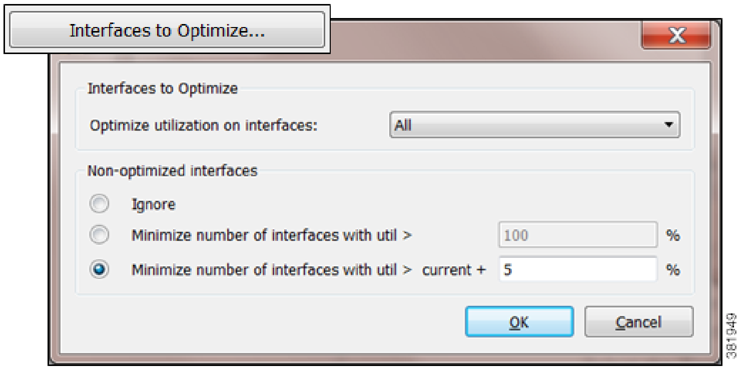
Interfaces to Optimize...

Interfaces to Optimize
Optimize utilization on interfaces: All

Non-optimized interfaces
○ Ignore
○ Minimize number of interfaces with util > 100 %
● Minimize number of interfaces with util > current + 5 %

OK Cancel

- Interfaces to Optimize—By default, MATE Design optimizes all interfaces. You can restrict the optimized interfaces using this option and selecting the desired value in the Optimize utilization on interfaces drop-down list: all, core, selected in table, or tag name.

  For those interfaces that are not optimized, you can set the following.

  - Ignore
  - Minimize number of interfaces with util > ___ %
  - Minimize number of interfaces with util > current + ___ %

- Traffic Level—The traffic level over which the optimization is performed.

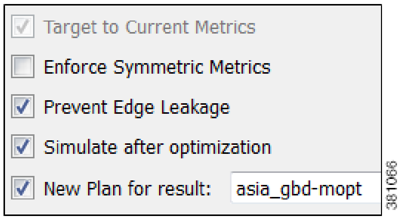
☑ Target to Current Metrics
□ Enforce Symmetric Metrics
☑ Prevent Edge Leakage
☑ Simulate after optimization
☑ New Plan for result: asia_gbd-mopt

- Target to Current Metrics—There are many sets of metrics that will result in equivalent sets of routes. (Doubling all metrics, for example, does not change routes.) MATE Design selects between these sets of metrics by metric targeting. Metrics can be selected to match as closely as possible the current metrics in the network. This simplifies the changeover from one set of metrics to the next.

  Targeting optimized metrics to the current metrics in the plan is performed by default. Various targeting refinements are available as CLI options to specify preferred upper and lower bounds on metric selections.

- Enforce Symmetric Metrics—This option sets the two metrics on each interface of each circuit equal to one another. This constraint can reduce optimization performance. However, it ensures that demands from A (source) to Z (destination) and from Z to A (in the core) take the same path, and so the chances of a session between A and Z being interrupted by a failure is minimized.

- Prevent Edge Leakage—This default option restricts how MATE Design can change demand routes. If a demand starts by routing only through the core, or from the edge through the core to the edge, this option prevents the routing from being changed so that it routes through the edge in the middle of its path. This is usually a desired routing policy restriction.

- Simulate after optimization—This default option performs a simulation analysis after optimizing metrics. The simulation analysis allows a detailed view, for example, of the utilizations under failure resulting from the new metric settings.

- New Plan for Result—This default option creates a new plan as a result of running Metric Optimization. Unless a name is specified, MATE Design attaches the suffix *-mopt* to the plan file name. If not selected, MATE Design changes the current plan file with the updated information.

Tactical Metric Optimization

The Tactical Metric Optimization tool is a reduced version of the Metric Optimization tool that minimizes the number of changes and runs faster. The options for this tool are a subset of those for the Metric Optimization tools.

Select the Tools->Tactical Metric Optimization menu.

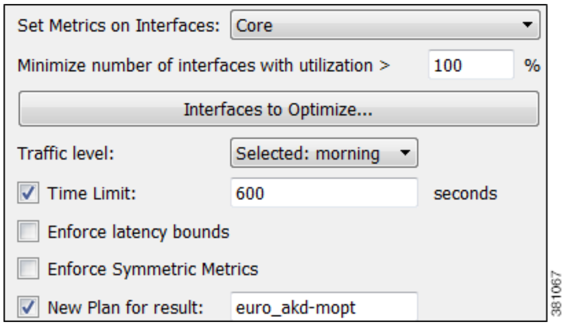

Set Metrics on Interfaces: Core

Minimize number of interfaces with utilization > 100 %

Traffic level: Selected: morning

☑ Time Limit: 600 seconds

☐ Enforce latency bounds

☐ Enforce Symmetric Metrics

☑ New Plan for result: euro_akd-mopt

Report

Metric Optimization writes a Metric-Opt report to the Design History section of the Report window. To access this information later, select the Window->Reports menu.

Metric Optimization Report

- Options—Summarizes the options used to call Metric Optimization. Some of these options (for example, the Output file and Report file) are only relevant when using the CLI tool.
- Number of demands—Total number of demands in the plan.
- Objectives—Summarizes the optimization objectives in order of priority.

Warnings

Warnings detected by Metric Optimization might lead to misleading or undesirable results.

- No Improvement in Optimization

  ```
  No routing improvement found: metrics unchanged
  ```

  Metric Optimization was unable to improve on the metrics in the network. If the incremental option was specified and latency bounds were enforced, this could be because the target metrics violated latency bounds in ways that incremental optimization could not fix.

- Optimization Exit Diagnostics

  ```
  Optimization performance may be limited by low edge metrics
  ```

  During the optimization, some desirable routes could not be selected because the given edge metrics would have caused edge leakage.

- Optimization Constraints Violated

- `<n> \{intra. inter\site core metrics exceed \crlf \{intra, inter\site-metric \{upper, lower\-bound`

  One or more intrasite or intersite core metric bounds have been violated.

- Maximum normal utilization exceeds specified bound

  The original plan had normal utilization above the specified bound, and Metric Optimization was unable to reduce the normal utilization below this bound. If this occurs, Metric Optimization tries to set normal utilization as low as possible. No worst-case optimization is performed because reducing utilization under the normal scenario is the highest priority.

- <n> demands with non-zero bandwidth exceed latency bounds

  Since enforcing latency bounds is the highest optimization priority, this situation should not happen unless a latency bound is less than the shortest possible latency for a demand.

- <n> demands with zero bandwidth exceed latency bounds

  Metric Optimization does not always enforce latency bounds for zero bandwidth demands. These warnings can occasionally occur even if it is possible to find lower latency paths for the signaled demands.

- Unrouted Demands
  - <n> unroutable demands under normal operation

    There are <n> demands for which no route exists between source and destination, even under the No Failure scenario.

  - <n> unroutable demands under <m> circuit failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> circuit failure scenarios. (Different demands might be unroutable under different circuit failures.)

  - <n> unroutable demands under <m> SRLG failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> SRLG failure scenarios. (Different demands might be unroutable under different SRLG failures.)

  - <n> unroutable demands under <m> source/dest node failure scenarios

    There are <n> demands whose source and/or destination nodes fail under <m> failure scenarios. These demands clearly cannot be routed in these circumstances. (Different demands might be unroutable under different node failures.)

  - <n> unroutable demands under <m> non-source/dest node failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> node failure scenarios, where the nodes that fail are intermediate nodes in the path of the demands, and are not source/destination nodes of the demands. (Different demands might be unrouted under different node failures.)

Routing Summary

Summary statistics are displayed of the routes in network, both before (in parentheses) and after the optimization.

- Core/Edge Max Utilization—The maximum percentage utilization over any interface in the core or edge network. The normal utilization is under the normal (no-failure) scenario, and the worst-case utilization is maximized over all failure scenarios in the failure set.

- Latency—Median and average latencies over all routed demands, both in milliseconds and as a percentage of the latency of the smallest latency route possible.

  The latency of each demand's routing is calculated as a percentage of the latency of the shortest possible route, by latency, for that demand. Under the percentage column, the median and average of these percentages are displayed.
- Num of routes away from shortest path—The number of routes (out of all the demands routed) that do not follow the shortest path, by latency. This statistic is an indicator of how much the routes have been affected by utilization minimization as the primary criterion for optimization.

Metrics

The metric targeting choice is shown, together with the number of metrics different from the target metrics. Lists of interfaces for which metrics differ from the target metrics are displayed, together with both the target metric and the optimized metric. The metrics are separated into those on intrasite and intersite interfaces.

LSP Optimization

This chapter describes three LSP optimization tools.

- The Explicit LSP Optimization tool optimizes the placement of primary and secondary paths for selected LSPs.
- The Tactical Explicit LSP Optimization tool is a subset of the functions in the Explicit LSP Optimization tool for quicker mitigation of local congestion.
- The LSP Loadshare Optimization tool optimizes LSP load balancing percentages between parallel LSPs.

Explicit LSP Optimization

The Explicit LSP Optimization tool minimizes congestion by optimizing the placement of primary and secondary paths for selected LSPs. By default, MATE Design minimizes the utilization across primary paths under normal operation and creates disjoint secondary paths so that a single failure cannot disrupt both paths simultaneously.

Select the Tools->Explicit LSP Optimization menu.

- Objectives for optimizing explicit LSP paths are specified in the Primary Paths, Secondary Paths, and Tertiary Paths sections of the dialog box.
- Selections and Results, such as which interfaces and LSPs to operate on, whether to tag the new LSP routes, and whether to open plan files are specified in the lower, right section of the dialog box.
- Further Optimization Details, such as how to handle non-optimized interfaces and whether to maintain or discard LSP properties, are set by clicking the Details button.

The default is to optimize utilization on all interfaces using all LSPs on those interfaces. Another default is to remove CSPF constraints, such as affinities and hop limits. For example, MATE Design sets the setup bandwidth to 0, thus providing the greatest flexibility when setting these explicit paths. Note that you can change these and other default behaviors.

Upon completion, MATE Design tags the LSPs with *expopt* and generates a new plan file named with an *-eopt* suffix. This plan file opens, showing the LSPs table that is filtered to these rerouted (and newly tagged) LSPs. Saving this plan file then simplifies the process of identifying which LSPs to reconfigure in the network.

MATE Design also writes a report containing the results of the optimization. To access this information later, select the Window->Reports menu.

Primary Paths

This section of the dialog box enables you to specify objectives for primary path selections.

- Primary Paths—Define whether to reroute primary LSP paths.
  - Optimized—Create optimized explicit primary paths based on the objectives selected.
  - Keep—Route LSPs along the existing primary paths.

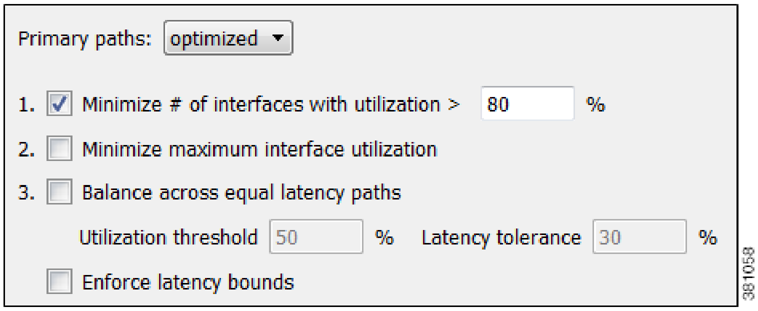

Primary paths: optimized ▾

1. ☑ Minimize # of interfaces with utilization > 80 %
2. □ Minimize maximum interface utilization
3. □ Balance across equal latency paths

Utilization threshold 50 % Latency tolerance 30 %

□ Enforce latency bounds

Primary paths are optimized using the following three objectives in order of priority. The first two objectives (1 and 2) move LSPs away from their shortest latency path in an attempt to reduce the utilizations of the most highly utilized interfaces in the network.

- 1. Minimize # of interfaces with utilization > _[80]_ x% —Specify a percentage and minimize the number of selected interfaces with utilizations over that percentage under normal operation (default).
- 2. Minimize maximum interface utilization—Route primary paths so that the maximum interface utilization over all selected interfaces is minimized under normal operation.
- 3. Balance across equal latency paths—Balance utilizations over lower utilized interfaces. For example, use this option to balance utilizations on parallel interfaces between two nodes if the two interfaces have the same latency.
  - Utilization threshold—Keep the number of interfaces with utilization greater than this value as low as possible without significantly increasing the latencies of the primary paths.
  - Latency tolerance—Permit this percentage of additional latency.
- Enforce latency bounds—Enforce latency bounds that can be specified for some or all demands in the plan file. If selected, this option takes precedence over all of the above objectives. For information on setting latency bounds, see the Traffic Demand Modeling chapter.

Secondary Paths

This section enables you to specify objectives for secondary path selections.

- Secondary Paths—Define whether and how to route secondary paths.
  - Optimized—Create optimized explicit secondary paths based on the objectives selected.
  - Dynamic—Route secondary paths dynamically. No explicit hops will be created for the path.
  - None—No secondary paths are created; existing paths are removed.
- Hot Standby—Set the secondary path to be a *hot* standby, which means it is established at the same time as the primary path, rather than after the primary path fails.

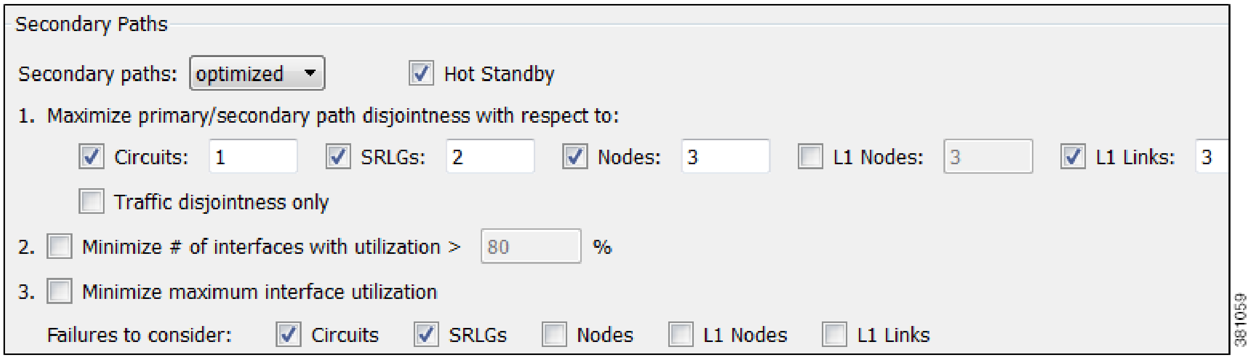

Secondary Paths

Secondary paths: optimized ▾ ☑ Hot Standby

1. Maximize primary/secondary path disjointness with respect to:

☑ Circuits: 1 ☑ SRLGs: 2 ☑ Nodes: 3 □ L1 Nodes: 3 ☑ L1 Links: 3

□ Traffic disjointness only

2. □ Minimize # of interfaces with utilization > 80 %

3. □ Minimize maximum interface utilization

Failures to consider: ☑ Circuits ☑ SRLGs □ Nodes □ L1 Nodes □ L1 Links

For optimized secondary paths, the following three objectives are used in order of priority listed.

- 1. Maximize primary/secondary path disjointness with respect to—Define primary and secondary paths for each LSP that are disjoint with respect to circuit, SRLG, node, and Layer 1 (L1) link failures, depending on what is selected.
  - Circuits—No circuit is used by both the primary and secondary paths (default = 1).
  - SRLGs—No SRLG is used by both the primary and secondary paths (default = 2).
  - Nodes—No node is used by both the primary and secondary paths (default = 3).
  - L1 Nodes—No L1 node is used by both the primary and secondary paths (default = 3).
  - L1 Links—No L1 link is used by both the primary and secondary paths (default = 3).
  - Traffic disjointness only—A path is acceptable even if it uses similar circuits, SRLGs, or nodes as other paths provided there is no traffic routed over the LSP when failures occur.

  You can specify the degree of disjointness of primary and secondary paths. The lower the number, the higher the disjointness priority. For example, if it is important that the paths are node disjoint and the SRLG disjointness are less important, you could change the setting to circuits 1, nodes 2, and SRLGs 3.

  Note that the network topology sometimes makes it impossible to fulfill all the selected disjointness requirements. In this case, paths are selected that are maximally disjoint. That is, they are disjoint for as many circuits, SRLGs, nodes, L1 nodes, and L1 links as possible.
- Options 2 and 3 both operate on the selected failure scenarios (circuits, SRLGs, nodes, L1 nodes, and L1 links) listed in the “Failures to consider” options listed under 3. These choices are the failure scenarios over which the simulation is performed at the end of the optimization. Note that this selection of failure scenarios is distinct from the failure scenarios selected for the disjointness objective 1.
  - 2. Minimize # of interfaces with utilization > ___ % —Minimize the number of interfaces with utilizations over the specified percentage across all selected failure scenarios.
  - 3. Minimize maximum interface utilization—Minimize the maximum interface utilization over all interfaces and over all selected failure scenarios.

Tertiary Paths

This section specifies the tertiary path configuration.

- Tertiary Paths—Define whether to create tertiary paths.
  - Dynamic—Create dynamic tertiary paths.
  - None—No tertiary paths are created.

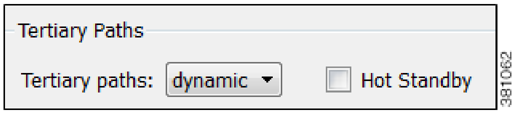
Tertiary Paths

Tertiary paths: dynamic ▾ □ Hot Standby

- Hot Standby—Set the tertiary path to be a hot standby path, which means it is brought up with the primary path, rather than after the primary path fails.

Selections and Results

These options identify which interfaces are used, whether LSPs are tagged, and whether to create new plan files containing the results.

The default is to optimize LSP paths on all selected interfaces, and by default, all interfaces are considered selected. You can, however, optimize a limited number of interfaces or specify a set of LSPs by pre-selecting them or by using tags.

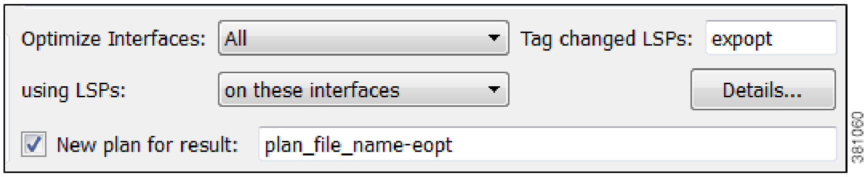
Optimize Interfaces: All ▾ Tag changed LSPs: expopt

using LSPs: on these interfaces ▾

☑ New plan for result: plan_file_name-eopt

- Optimized Interfaces—Identify which interfaces to optimize: all interfaces, those selected prior to opening the dialog box, or those with specific tags.
- Using LSPs—Specify which LSPs are candidates for rerouting. The default is all the LSPs currently routed on the selected interfaces.
- Tag changed LSPs—Create and add tags to all LSPs rerouted during optimization. Using tags makes it easier to find the LSPs that need to be reconfigured in the actual network according to the optimization results. By default, newly optimized LSPs are tagged as *expopt*.
- New plan for result—This default option creates a new plan with the results of the optimization. Unless a name is specified, MATE Design attaches an *-eopt* suffix to the current plan file name. If not selected, MATE Design changes the current plan file with the updated information.

Further Optimization Details

The Details button opens a dialog box wherein you can identify less frequently used parameters, such as traffic levels and how to treat non-optimized interfaces.

- Non-Optimized Interfaces—You can specify whether to ignore non-optimized interfaces or set the acceptable level of utilization for them.

  If setting an acceptable utilization level and if both options are selected, MATE Design uses the higher of the two. These settings are calculated on a per-interface basis.

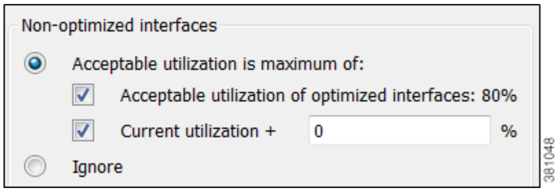
Non-optimized interfaces

● Acceptable utilization is maximum of:

☑ Acceptable utilization of optimized interfaces: 80%

☑ Current utilization + 0 %

○ Ignore

  - Acceptable utilization of optimized interfaces: ___% —This value is the same as the utilization threshold set in the

Primary Path section, option 1 (Minimize # of interfaces with utilization > ___%, where the default is 80). Therefore to change this value, you must change it in the Primary Path section.

If using the Tactical Explicit LSP Optimization tool, this field is equivalent to, and thus only changeable in, the Acceptable Utilization ___% field.

- Current utilization + ___% —Current utilization of non-optimized interfaces plus the added percentage.

Example: There are two non-optimized interfaces: cr1.chi_cr1.mia has a utilization of 60% and cr2.sjc_cr2.okc has a utilization of 78%.

The acceptable utilization settings for non-optimized interfaces are as follows.

- The primary path utilization threshold is 80%.
- The current utilization has 5% added to it.

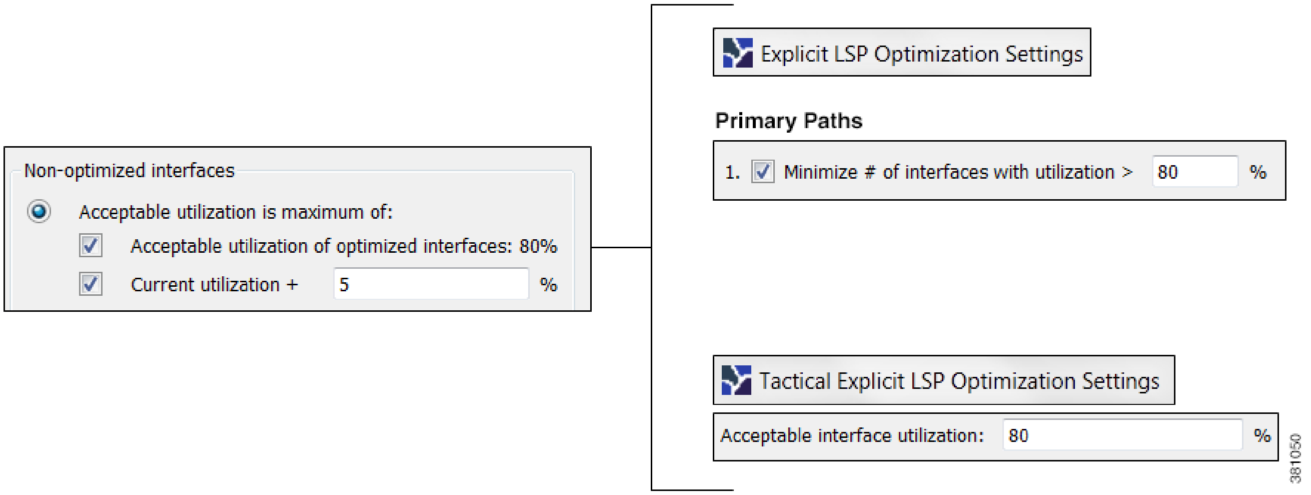

Non-optimized interfaces

● Acceptable utilization is maximum of:

☑ Acceptable utilization of optimized interfaces: 80%

☑ Current utilization + 5 %

Explicit LSP Optimization Settings

Primary Paths

1. ☑ Minimize # of interfaces with utilization > 80 %

Tactical Explicit LSP Optimization Settings

Acceptable interface utilization: 80 %

Result: The maximum utilization for each interface is individually calculated. The acceptable level of utilization for cr1.chi_cr1.mia is 80%, and the acceptable level of utilization for cr2.sjc_cr2.okc is 83% (78 + 5).

- LSP Configuration, Create Fully Explicit Named Paths For—By default, MATE Design creates named paths across the LSPs that were selected if a reroute is required to achieve the optimization objectives. For instance, if a selected dynamically routed LSP has an acceptable route, MATE Design does not create a named path for it.

  To change this default, select "all selected LSPs." MATE Design then creates and routes fully explicit named paths for all selected LSPs.

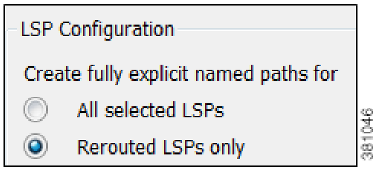

LSP Configuration

Create fully explicit named paths for

○ All selected LSPs

● Rerouted LSPs only

- LSP Configuration, In Addition—By default, MATE Design sets the Setup BW to zero, providing the greatest flexibility when creating explicit routes. As well, all affinities and hop limits are removed, and the setup and hold priorities are set to 7. These changes apply only to the LSPs with newly created or changed explicit named paths.

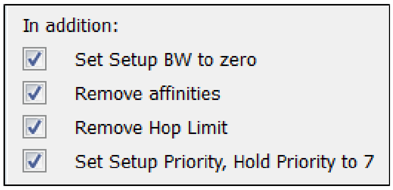

In addition:
☑ Set Setup BW to zero
☑ Remove affinities
☑ Remove Hop Limit
☑ Set Setup Priority, Hold Priority to 7

  You can turn off these defaults, and you can also set them individually after the optimization is performed using the LSP Paths Properties dialog box. (In the LSP Paths table, right-click on an LSP path and select Properties.) If these defaults are turned off, the original parameters are preserved.

- Traffic Level—Select the traffic level you want used in the utilization calculations and optimizations. For information on traffic levels, see the Traffic Demand Modeling chapter.

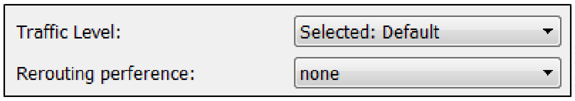

Traffic Level: Selected: Default
Rerouting perference: none

- Rerouting Preference—By default, preferences on which LSPs to reroute are not based on traffic volume. You can select high-traffic or low-traffic options to sequentially give priority to LSPs with higher and lower traffic, respectively.

Tactical Explicit LSP Optimization

The Tactical Explicit LSP Optimization tool is a reduced version of the Explicit LSP Optimization tool. This tool optimizes only primary paths using the minimum number of path changes required to bring utilizations below an acceptable level. It is useful when you need to reduce congestion in a specific area of the network with a limited number of LSP reconfigurations.

As described in the Selections and Results section, if you need to target specific interfaces or LSPs to confine the optimization to problem areas., then select those interfaces or LSPs first. Then select the Tools->Tactical Explicit LSP Optimization menu.

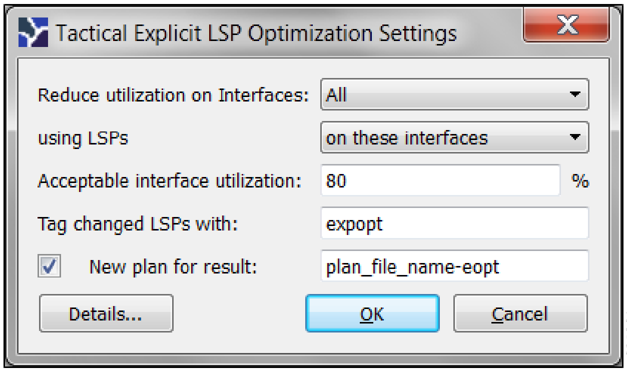

Tactical Explicit LSP Optimization Settings

Reduce utilization on Interfaces: All
using LSPs: on these interfaces
Acceptable interface utilization: 80 %
Tag changed LSPs with: expopt
☑ New plan for result: plan_file_name-eopt

- Acceptable interface utilization___% — Maximum allowable traffic utilization on interfaces used in the optimization.
- Details button—See the Further Optimization Details section.
- All other options—See the Selections and Results section.

LSP Loadshare Optimization

The LSP Loadshare Optimization tool automates the process of finding and setting the most favorable loadshare ratios across parallel LSPs to balance traffic and avoid congestion. This optimization is based on which of the following options are selected. Only interfaces using parallel LSPs are included in the optimization.

- Minimize max interface util—Minimize the maximum interface utilization over all interfaces on the LSP routes. MATE Design attempts to minimize the number of loadshare parameter changes required to do so.

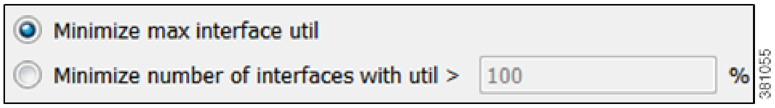

● Minimize max interface util
○ Minimize number of interfaces with util > 100 %

- Minimize number of interfaces with util > ___ %—Minimize the number of interfaces that have a utilization greater than the specified value. This is a looser constraint than minimizing the maximum interface utilization. Therefore, MATE Design has the opportunity to modify as few Loadshare values as possible, thus reducing the amount of reconfiguration required.

- Number of flow bins—Routers typically cannot divide flows arbitrarily between parallel LSPs, but instead allocate them to a fixed number of "bins" of approximately equal size. The bins are then divided between the parallel LSPs. This option enables you to specify the total number of bins, which in turns places a constraint on the traffic division.

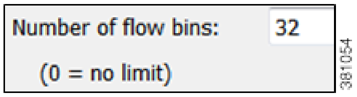

Number of flow bins: 32
(0 = no limit)

Upon completion, MATE Design tags the LSPs with LSPLoadshare and generates a new plan file named with an -Lshare suffix. This plan file opens, showing the LSPs table that is filtered to these rerouted (and newly tagged) LSPs. Saving this plan file then simplifies the process of identifying which LSPs to reconfigure in the network. Both the tags and the new plan file names are editable.

Also upon completion a report is generated, identifying how many LSPs and interfaces were affected, the number of bins used, the resulting maximum interface utilization, and if applicable, the number of interfaces with utilization over the specified threshold.

For More Information...	See....
Comprehensive description of how loadshare simulations work, and how to set the Loadshare property manually	MPLS Simulation chapter

Optimized Loadshare Examples

Minimization Example: The base plan for this example, Acme_Network, contains two sets of parallel LSPs, each with a Loadshare value of 1 and each using strict hops on named paths to reach its destination. All of these LSPs have a Traff Sim of 3,000 Mbps. The interfaces to which these LSPs filter all have a Traff Sim value of 3,000 Mbps, except for sjc-to-okc, which has 6,000 Mbps of simulated traffic (Figure 17-1).

- Using the Acme_Network plan file (Figure 17-1), if you minimize the maximum interface utilization across all LSPs, all four Loadshare parameters change, and the maximum interface utilization is 40% (Figure 17-2). The default plan file name and default LSP tags are used, thus naming the new plan Acme_Network-Lshare and the tags as LSPLoadshare.
- Using the Acme_Network plan file (Figure 17-1), if you minimize the number of interfaces with utilization greater than 50%, only two of the Loadshare values change (Figure 17-3). Notice in the figure how Loadshare values are relative to the other Loadshare values for parallel LSPs, and not to all LSPs in the table. Also, only the affected LSPs are tagged.

Figure 17-1 Example Acme Network Before LSP Loadshare Optimization

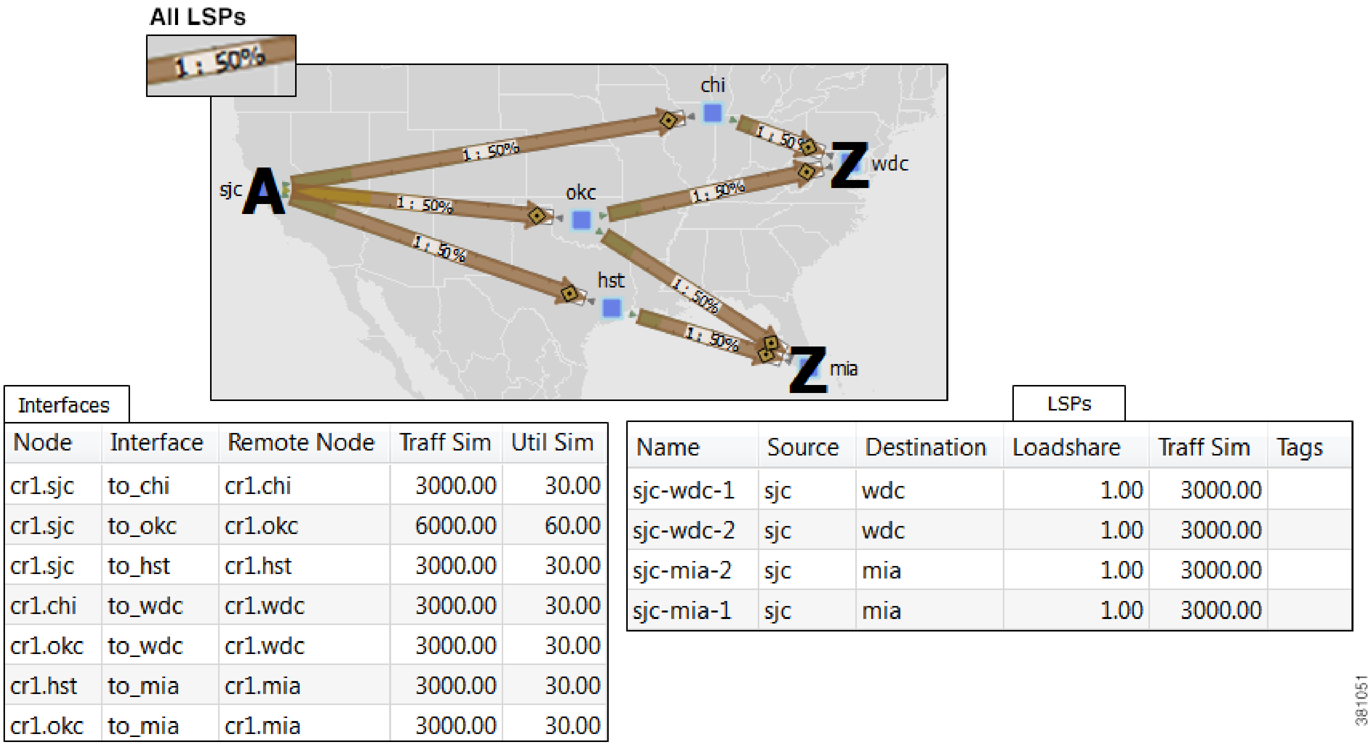

Node	Interface	Remote Node	Traff Sim	Util Sim
cr1.sjc	to_chi	cr1.chi	3000.00	30.00
cr1.sjc	to_okc	cr1.okc	6000.00	60.00
cr1.sjc	to_hst	cr1.hst	3000.00	30.00
cr1.chi	to_wdc	cr1.wdc	3000.00	30.00
cr1.okc	to_wdc	cr1.wdc	3000.00	30.00
cr1.hst	to_mia	cr1.mia	3000.00	30.00
cr1.okc	to_mia	cr1.mia	3000.00	30.00

Name	Source	Destination	Loadshare	Traff Sim	Tags
sjc-wdc-1	sjc	wdc	1.00	3000.00	
sjc-wdc-2	sjc	wdc	1.00	3000.00	
sjc-mia-2	sjc	mia	1.00	3000.00	
sjc-mia-1	sjc	mia	1.00	3000.00	

Figure 17-2 Example Acme Network After Minimizing Maximum Interface Utilization

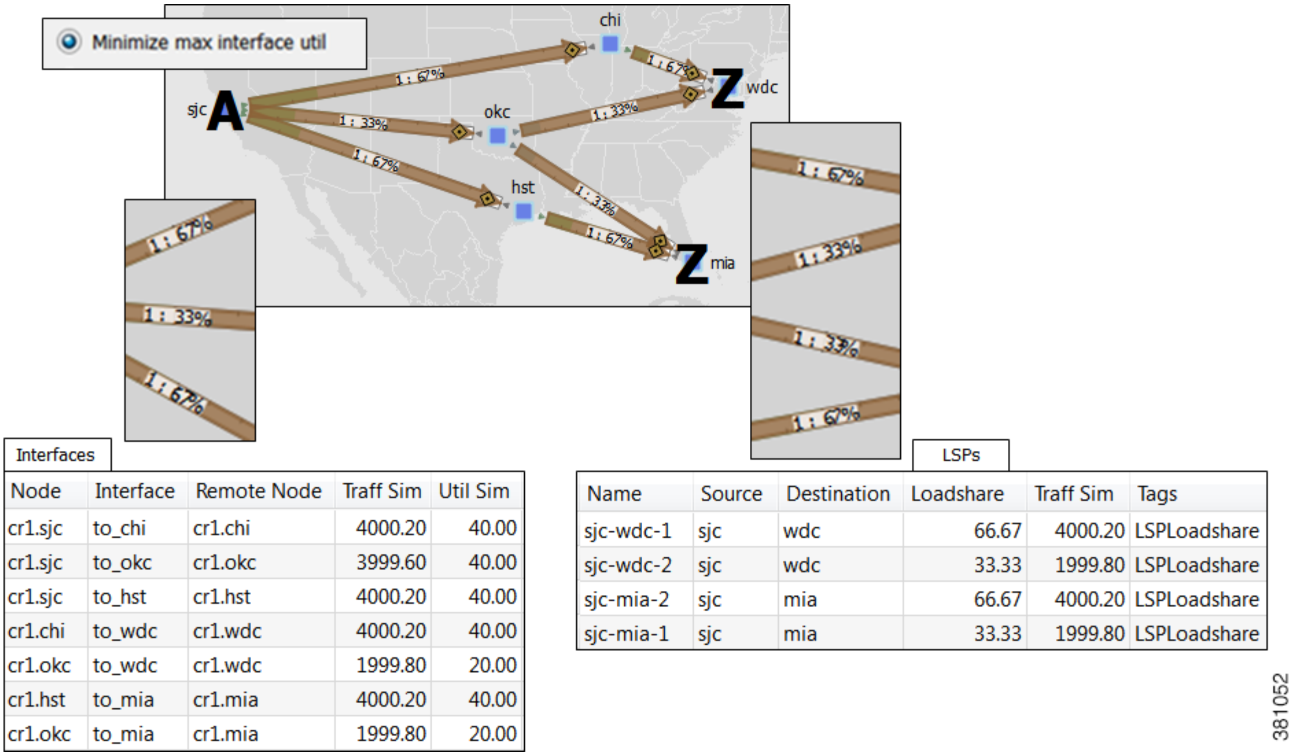

Node	Interface	Remote Node	Traff Sim	Util Sim
cr1.sjc	to_chi	cr1.chi	4000.20	40.00
cr1.sjc	to_okc	cr1.okc	3999.60	40.00
cr1.sjc	to_hst	cr1.hst	4000.20	40.00
cr1.chi	to_wdc	cr1.wdc	4000.20	40.00
cr1.okc	to_wdc	cr1.wdc	1999.80	20.00
cr1.hst	to_mia	cr1.mia	4000.20	40.00
cr1.okc	to_mia	cr1.mia	1999.80	20.00

Name	Source	Destination	Loadshare	Traff Sim	Tags
sjc-wdc-1	sjc	wdc	66.67	4000.20	LSPLoadshare
sjc-wdc-2	sjc	wdc	33.33	1999.80	LSPLoadshare
sjc-mia-2	sjc	mia	66.67	4000.20	LSPLoadshare
sjc-mia-1	sjc	mia	33.33	1999.80	LSPLoadshare

Figure 17-3 Example Acme Network After Minimizing Number of Interfaces with Utilization > 50%

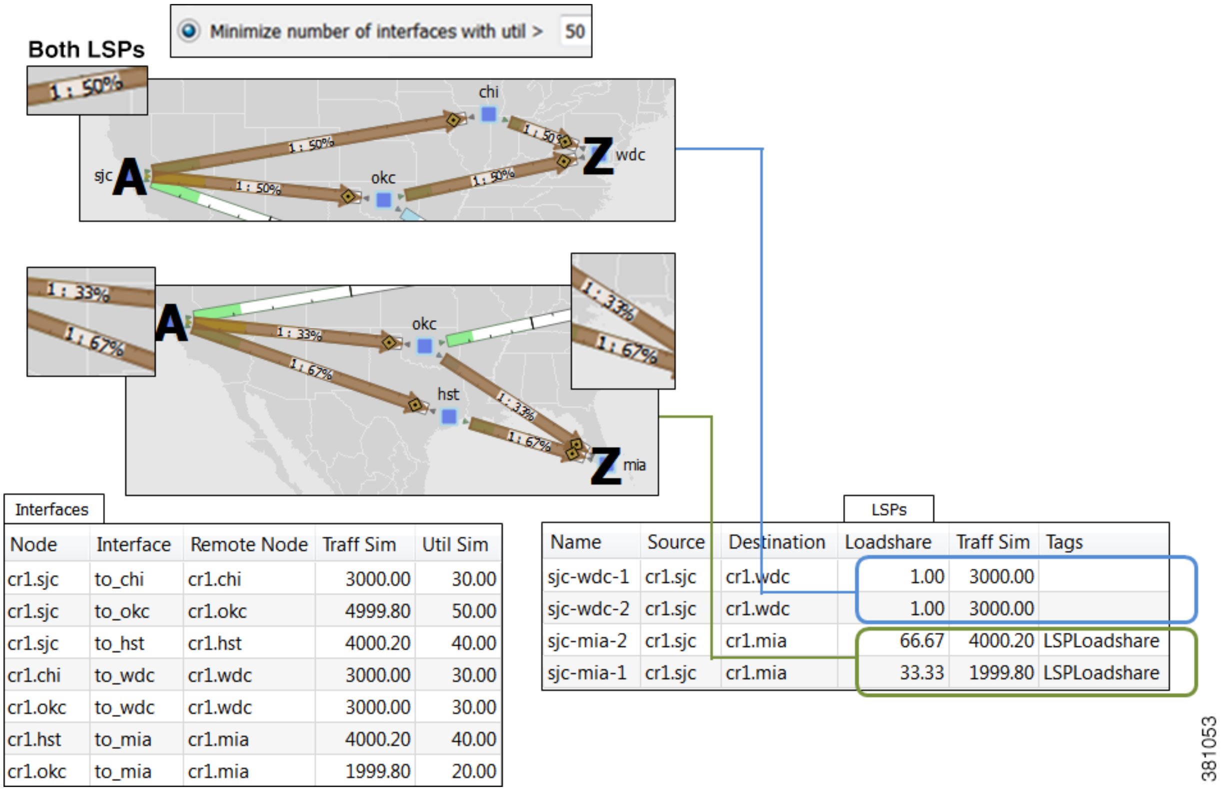

● Minimize number of interfaces with util > 50

Interfaces

Node	Interface	Remote Node	Traff Sim	Util Sim
cr1.sjc	to_chi	cr1.chi	3000.00	30.00
cr1.sjc	to_okc	cr1.okc	4999.80	50.00
cr1.sjc	to_hst	cr1.hst	4000.20	40.00
cr1.chi	to_wdc	cr1.wdc	3000.00	30.00
cr1.okc	to_wdc	cr1.wdc	3000.00	30.00
cr1.hst	to_mia	cr1.mia	4000.20	40.00
cr1.okc	to_mia	cr1.mia	1999.80	20.00

LSPs

Name	Source	Destination	Loadshare	Traff Sim	Tags
sjc-wdc-1	cr1.sjc	cr1.wdc	1.00	3000.00	
sjc-wdc-2	cr1.sjc	cr1.wdc	1.00	3000.00	
sjc-mia-2	cr1.sjc	cr1.mia	66.67	4000.20	LSPLoadshare
sjc-mia-1	cr1.sjc	cr1.mia	33.33	1999.80	LSPLoadshare

Bin Example: In this example, a node uses a maximum of 32 bins, and the optimal traffic allocation is 45% of the traffic through LSP A and 55% of the traffic through LSP B. The node that sources these two LSPs, however, cannot do this exact split. The optimization divides the traffic into 32 bins, each with the same amount of traffic in them. Thus, each bin has 3.125% of the traffic (100% of the traffic divided by 32). The optimization also determines how the node should split the traffic. In this case, the split is to give 43.75% (which is 14 bins of 3.125% each) to LSP A and 56.25% (which is 18 bins of 3.125% each) to LSP B. Thus, it optimizes and distributes 32 (14+18) bins of traffic. This is as close as possible to the optimal 45%/55% split using 32 bins.

Run LSP Loadshare Optimization Tool

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

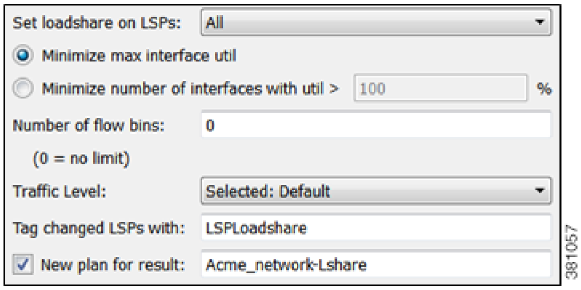

Set loadshare on LSPs: All

● Minimize max interface util

○ Minimize number of interfaces with util > 100 %

Number of flow bins: 0

(0 = no limit)

Traffic Level: Selected: Default

Tag changed LSPs with: LSPLoadshare

☑ New plan for result: Acme_network-Lshare

Step 2 Select the Tools->LSP Loadshare Optimization menu.

Step 3 Select which LSPs on which you want to optimize the Loadshare value.

Step 4 Select which method of optimization to use.

- Minimize max interface util
- Minimize number of interfaces with util > ___ %

Step 5 If distributing the traffic to bins, enter the total number of flow bins to use.

Step 6 Optional: Select which traffic level to optimize.

Step 7 Optional: Change the LSP tag name.

Step 8 Optional: Select whether to create a new plan file based on the results, and change how to name it.

CHAPTER 18

Changeover

MATE Design enables you to model the topology, routings, and utilizations of a currently operational network. It also allows exploration of modifications to that network. For example, interface metrics can be modified or explicit LSP routings changed to obtain a new routing configuration.

The Changeover tool provides a step-by-step sequence of routing configuration changes to move a network safely from an initial configuration to a pre-specified final configuration. MATE Design selects the order of these changes so that there is as little as possible congestion in the network during intermediate configurations, and so that this congestion lasts for as few intermediate steps as possible.

Only certain configuration changes are allowed between the initial and final plan.

- Changes to interface metrics
- Changes to LSP configurations
- Changes of circuits and node from the active to the inactive state, and vice-versa

An individual step in the changeover sequence consists of one of the following.

- A single metric change on a specific interface
- LSP configuration changes on all LSPs sourced from a specific node
- The activation or inactivation of a specific circuit or node

Changeover Options

To run Changeover, both the initial and final network plans should be open in MATE Design, and the desired plan should be selected to view.

Select the Tools->Changeover to open a dialog box with the following options.

- Initial Plan—Name of the initial plan, selected from open plan files.

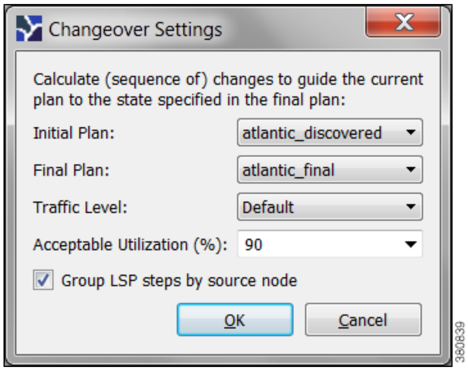

Changeover Settings

Calculate (sequence of) changes to guide the current plan to the state specified in the final plan:

Initial Plan:	atlantic_discovered
Final Plan:	atlantic_final
Traffic Level:	Default
Acceptable Utilization (%):	90

☑ Group LSP steps by source node

- Final Plan—Name of the final plan, selected from open plan files.
- Traffic Level—Changeover monitors maximum interface utilization levels for all steps in the sequence. The utilizations are calculated using this traffic level.
- Acceptable Utilization—The percentage of maximum acceptable utilization of any interface at any step during the changeover sequence. Changeover attempts to keep utilizations below this level, although it is not always possible. Often the utilization spikes upward during the last few steps of the sequence, if, for example, new ECMP paths must be put in place. MATE Design attempts to keep the number of high-utilization steps to a minimum.
- Select whether to group the LSP steps by source node.

Report

Changeover creates a Changeover report in the Reports window, containing Summary, Steps, and Utilization sections.

- To access this information later, select the Window->Reports menu.
- To export the steps taken to change over to a new plan file, right-click on Steps in the Reports window, and select Export Tables.
- To export the utilization values to a table, right-click on Utilizations in the Reports window, and select Export Tables.

Summary

The Summary section contains a list of the options used to run Changeover. It also contains a summary of the differences between the initial and final plans, the number of steps taken, and the number of steps resulting in intermediate configurations with utilizations in excess of the acceptable utilization.

Steps

The Steps section contains the following columns, some of which might need to show since they are hidden.

- Step—Steps of the Changeover sequence.
- Action—Action taken at the step.
  - MetricChange—An interface metric is changed.
  - CapacityChange—An interface capacity is changed.
  - TunnelsFromRouter—LSPs from a specified router are to be reconfigured.
  - ActivateNode—A node (router) is activated.
  - DeactivateNode—A node (router) is de-activated.

  - ActivateInterface—An interface is activated.
  - DeactivateInterface—An interface is de-activated.
- Node
  - If Action is MetricChange, the node containing the interface of the metric to change.
  - If TunnelsFromRouter, the node at which the LSPs are to be reconfigured.
  - If ActivateNode or DeactivateNode, the relevant node.
  - If ActivateInterface or DeactivateInterface, the node containing the relevant interface.
- RemoteNode—If Action is MetricChange, the remote node containing the interface of the metric to change.
- Circuit—If Action is MetricChange, the circuit for the metric change.
- Interface
  - If Action is MetricChange, the interface name.
  - If TunnelsFromRouter, ActivateNode, or DeactivateNode, this is na (not applicable).
  - If ActivateInterface or DeactivateInterface, the relevant interface name.
- MaxUtil—The maximum interface utilization at this step.
- UtilLocation—The interface on which the maximum utilization occurs.
- ChangeTo—If Action is MetricChange, the original metric. Otherwise, this is set to “na.”
- ChangeFrom—If Action is MetricChange, the metric to be configured at this step. Otherwise, this is set to “na.”

Utilizations

The Utilizations section contains all utilizations for all interfaces in the network, for each step in the Changeover sequence.

CHAPTER 19

Reports

MATE Design enables you to generate reports that compare IP/MPLS topology and configuration information, demand routings, LSP routings, or traffic between two plan files. You can also generate a complete plan comparison. Following are a few examples of how these reports could be used.

- Plan for upgrades—Compare plan files before and after topology and configuration changes are made to a current network. For example, you could compare the original plan file to a proposed new plan in which circuits or nodes were added or upgraded.
- Mitigate congestion—Compare plans before and after making configuration changes to mitigate congestion due to a failure or planned maintenance, such as metric changes or LSP additions and reroutings.
- Validate simulations—Compare the simulated traffic under failure in one plan to the measured traffic after a failure has occurred in another plan to validate the accuracy of the simulation.
- Plan audit—Review any changes between two plan files.
- Identify demand and LSP reroutes—Compare plans to determine which demands or LSPs experienced routing changes, for example as the result of an interface metric change or capacity upgrades.

Upon comparing plan files, a report of the results automatically opens. You can access this report later by selecting the Window->Reports menu.

For More Information ...	**See ...**
Reports on LSP path routes	MPLS Simulation chapter
Reports on L1 circuit routes	Layer 1 Simulation chapter
Reports on demand groupings and demand grouping costs	Forecasting Traffic chapter Cost Modeling chapter
Reports on infrastructure costs	Cost Modeling chapter

Plan Comparison Reports

Comparison reports enable you to compare objects between two plan files to determine

- Which objects are present in one plan, but not in the other, and
- Which objects are present in both plans and which, if any, have different properties.

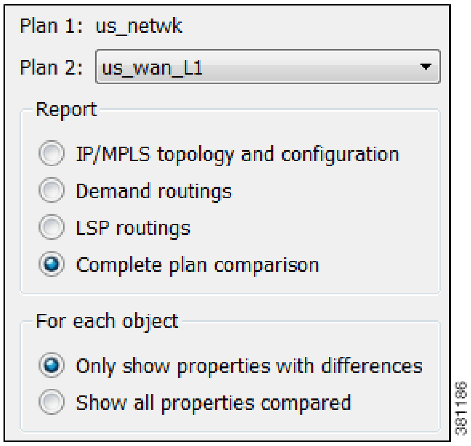

Plan 1: us_netwk

Plan 2: us_wan_L1

Report

○ IP/MPLS topology and configuration

○ Demand routings

○ LSP routings

● Complete plan comparison

For each object

● Only show properties with differences

○ Show all properties compared

You can run four types of plan comparison reports, as follows.

- IP/MPLS topology and configuration—-Objects describing the topology (such as nodes, circuits, and interfaces), configured objects (such as LSPs and LSP paths), and related properties (such as IGP metrics and capacity). For a complete list of compared objects and properties, see Table 19-1.
- Demand routings—Demand paths and the properties that indicate routing changes, such as Path Length and Max Latency. For a complete list of compared properties, see Table 19-2.
- LSP routings—Properties of LSPs that indicate routing changes, such as TE Path Metric and Active Path Sim. For a complete list of compared properties, see Table 19-3.
- Complete plan comparison—All tables in the table schema are compared, except for internal simulation caching tables. For a complete list of these tables, refer to the `$CARIDEN_HOME/docs/table_schema` file.

Report Columns

Each MATE table has key columns that uniquely identify objects. For example, the key columns in the Interfaces table are the Node and Interface columns, which represent the name of the node containing the interface and the name of the interface itself. In a Plan Comparison report, an object in one plan is identified with an object in another plan only if the key columns of the two objects match. That is, key columns determine if an object is present in both plans or in only one of the plans.

The Plan Comparison report displays three other types of columns. Table 19-1, Table 19-2, and Table 19-3 list the properties reported upon and their associated column type.

- Information only (Info)—There are no comparisons made. Information is reported for the plan file from which you are running the report (Plan 1).
- Differences (Diff)—Specific properties are compared for each object that has matching key columns across the two plan files. Objects are then identified as belonging in either Plan 1 only (file from which you are running the report), in Plan 2 only (plan file to which you are comparing), or in both Plan 1 and Plan 2. If the object exists in both plan files, but has different properties, a Diff column shows T (true) if there are any differences or F (false) if there are none.

  Differences are based on the current state of the open plan files. For instance, if a circuit is currently failed, the demand routings change to route around the failure.
- Summary Differences (Summ Diff)—These are not columns within MATE tables. They are T/F (true or false) values, depending on some differences in the table objects that are not represented in the table columns. For example, if a common LSP path is configured differently between plans, this

difference would be represented in a summary difference column in the LSP section. If demand routings are different, this would be represented in a summary difference column in the Demands section.

Note For easy reference, the following tables list columns in the order of type first: Key, Info, Diff, and Summ Diff. Properties within each column type are alphabetical.

Table 19-1 Plans Comparison Report: IP/MPLS Topology and Configuration

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
Interfaces	Name	Key	
	Node	Key	
	Remote Node	Info	
	Affinities	Diff	
	Area	Diff	
	Capacity	Diff	
	Circuit	Diff	
	Description	Diff	
	FRR Protect	Diff	
	IGP Metric	Diff	
	IP Address	Diff	
	PC Min BW	Diff	
	PC Min Links	Diff	
	Resv BW	Diff	
	TE Metric	Diff	
Circuits	InterfaceA	Key	
	InterfaceB	Key	
	NodeA	Key	
	NodeB	Key	
	Name	Info	
	Active	Diff	
	Capacity	Diff	
	Delay	Diff	
Nodes	Name	Key	
	Active	Diff	
	AS	Diff	
	BGP ID	Diff	
	ECMP Max	Diff	
	IP Address	Diff	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	IP Manage	Diff	
	Model	Diff	
	OS	Diff	
	Vendor	Diff	
LSPs	Node	Key	
	Source	Key	
	Active	Diff	
	Destination	Diff	
	Exclude	Diff	
	FRR Link Protect	Diff	
	Hold Pri	Diff	
	Hop Limit	Diff	
	Include	Diff	
	Include Any	Diff	
	Load Share	Diff	
	Metric	Diff	
	Metric Type	Diff	
	Setup BW	Diff	
	Setup Pri	Diff	
	Unresolved Destination	Diff	
	LSPs Path Diff	Summ Diff	T (true) if there are any differences in the LSP Paths table for this LSP. If two LSPs have differences in their named path hops or LSP paths, these differences are propagated to this column.
LSP Paths	Node	Key	
	Path Option	Key	
	Source	Key	
	Exclude	Diff	
	Hold Pri	Diff	
	Hop Limit	Diff	
	Include	Diff	
	Include Any	Diff	
	Path Name	Diff	
	Setup BW	Diff	
	Setup Pri	Diff	
	Standby	Diff	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	Named Path Diff	Summ Diff	T (true) if there are any differences in the Named Paths table for this LSP path.
Named Paths	Active	Key	
	Name	Key	
	Source	Key	
	Named Path Hops Diff	Summ Diff	T (true) if there are any differences in the Named Path Hops table for this named LSP path.
Named Path Hops	Name	Key	
	Source	Key	
	Step	Key	
	Interface	Diff	
	IP Address	Diff	
	Node	Diff	
	Type	Diff	
	Unresolved Hop	Diff	

Table 19-2 Plans Comparison Report: Demand Routings

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
Demands	Destination	Key	
	Name	Key	
	Service Class	Key	
	Source	Key	
	Destination Site	Info	
	Source Site	Info	
	Active	Diff	
	ECMP Min %	Diff	
	Max Latency	Diff	
	Path Metric	Diff	
	Path Diff	Summ Diff	T (true) if there are any differences in routing of the demand.

Table 19-3 Plans Comparison Report: LSP Routings

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
LSPs	Name	Key	
	Source	Key	
	Active Path Sim	Diff	
	Destination	Diff	
	Routed	Diff	
	Shortest TE Path	Diff	
	TE Path Metric	Diff	
	Actual Path Diff	Summ Diff	T (true) if there are any differences in routing of the actual LSP path.
	Simulated Path Diff	Summ Diff	T (true) if there are any differences in routing of the LSP.

Report Sections

The opening section, Summary, gives a numeric summary of the number of objects in each plan file, as well as the number of objects in both plan files and of these in both, the number of objects with property differences (Figure 19-1).

Example: Nodes A, B, and C are in Plan 1, and B, C, and D are in Plan 2. Nodes B and C have matching key columns. The B nodes have identical properties, but the C nodes have different IP addresses. Therefore, in the Summary section, the In Both Plans column shows 2, and the Different Properties column shows 1.

An individual section is generated for each object that appears in only one plan file and if there are differences between the properties, a section is generated showing these differences. Thus, there is the potential of generating one to three sections per compared table: one for objects appearing only in Plan 1, one for objects appearing only in Plan 2, and one listing objects existing in both plan files, along with their differences noted with a T (true) or F (false).

Figure 19-1 Example Plan Comparison Report Summary

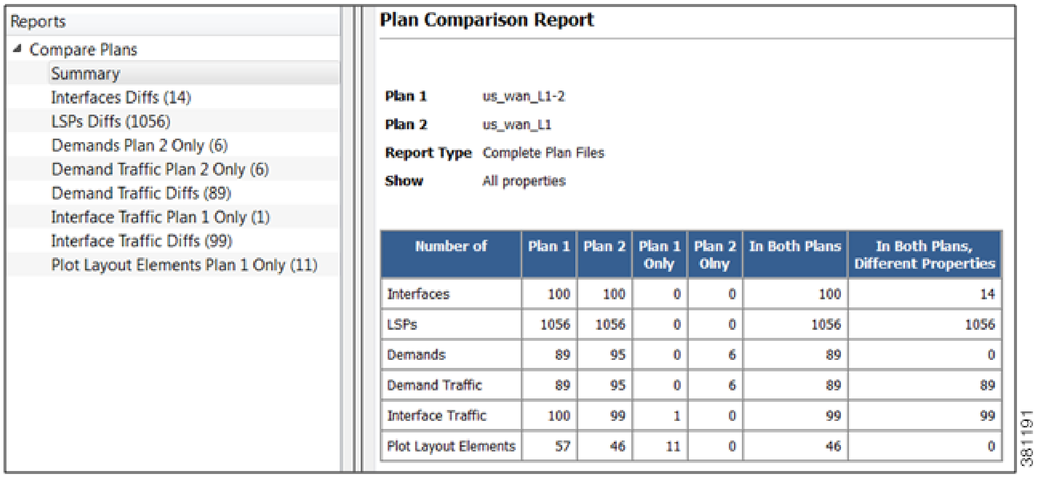
Reports
Compare Plans
Summary
Interfaces Diffs (14)
LSPs Diffs (1056)
Demands Plan 2 Only (6)
Demand Traffic Plan 2 Only (6)
Demand Traffic Diffs (89)
Interface Traffic Plan 1 Only (1)
Interface Traffic Diffs (99)
Plot Layout Elements Plan 1 Only (11)
381191

Plan Comparison Report

Plan 1	us_wan_L1-2
Plan 2	us_wan_L1
Report Type	Complete Plan Files
Show	All properties

Number of	Plan 1	Plan 2	Plan 1 Only	Plan 2 Olny	In Both Plans	In Both Plans, Different Properties
Interfaces	100	100	0	0	100	14
LSPs	1056	1056	0	0	1056	1056
Demands	89	95	0	6	89	0
Demand Traffic	89	95	0	6	89	89
Interface Traffic	100	99	1	0	99	99
Plot Layout Elements	57	46	11	0	46	0

Other than Summary, the report sections and columns (properties) that appear depend on the options selected when generating the report.

- If you select to show only properties with differences, only those columns with differences in properties appear in the report. If only one value is different, all objects are listed (not just the one with the different property).
- If you select to show all properties compared, all compared properties appear in the report whether there are differences between the two plan files or not.

Create Plan File Comparison Reports

Step 1 Open the two plan files you want to compare.

Step 2 In either of these two plan files, select the Tools->Reports->Compare Plans menu. A Compare Plans dialog box appears.

Step 3 Select the plan file to which you are comparing the open plan.

Step 4 Select the type of comparisons you want to generate.

- IP/MPLS topology and configuration
- Demand routings
- LSP routings
- Complete plan files

Step 5 Select whether to show only properties with differences or whether to show all properties compared.

Step 6 Click OK.

Traffic Comparison Reports

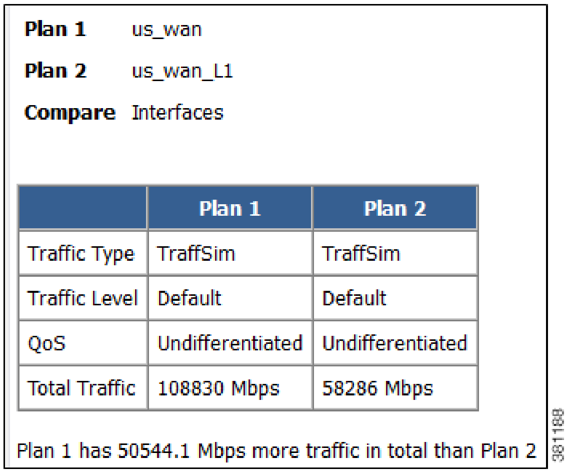

Plan 1 us_wan
Plan 2 us_wan_L1
Compare Interfaces

	Plan 1	Plan 2
Traffic Type	TraffSim	TraffSim
Traffic Level	Default	Default
QoS	Undifferentiated	Undifferentiated
Total Traffic	108830 Mbps	58286 Mbps

Plan 1 has 50544.1 Mbps more traffic in total than Plan 2

Traffic Comparison reports compare traffic values within a single plan or between two plan files. You can compare traffic on one object at a time (nodes, interfaces, circuits, demands, LSPs, multicast flows, and flows). Each object has a different set of comparisons from which to choose. For example, interface traffic and capacity can be compared, and LSP traffic and setup bandwidth can be compared. Another example would be to compare the simulated interface traffic in a plan file before a failure to the measured interface traffic of a plan file in which a failure has occurred. Such a comparison would enable you to determine which interfaces show the greatest differences.

For each plan file compared, the current traffic level and either the selected service class or queue are used. They can be different in each plan.

The opening section, Summary, gives a high-level summary of what is being compared. Additionally, a section of the difference on the selected object and traffic columns is generated and named accordingly.

Create Traffic Comparison Reports

Step 1 If comparing traffic from the open plan file to another, open the other plan file.

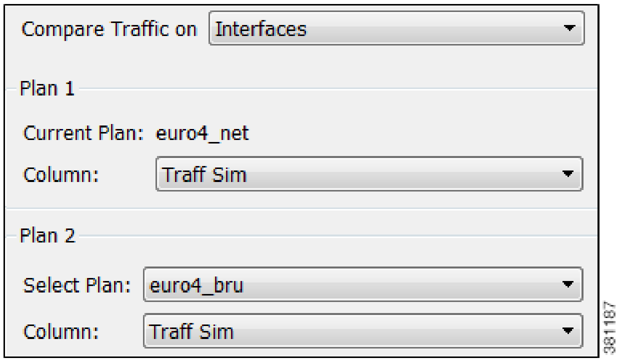

Compare Traffic on: Interfaces

Plan 1

Current Plan: euro4_net

Column: Traff Sim

Plan 2

Select Plan: euro4_bru

Column: Traff Sim

Step 2 In each plan file, select the travel levels and select either the service class or queue traffic you want to compare.

Step 3 Select the Tools->Reports->Compare Traffic menu. This plan file is identified as Plan 1 in the Compare Traffic dialog box that appears.

Step 4 Select the object type on which you want to compare traffic: interfaces, nodes, LSPs, demands, multicast flows, or flows.

Step 5 Select the traffic column in the current plan that you want to compare (Plan 1).

Step 6 The default Plan 2 is the current plan file. To change this, select a different plan file for Plan 2.

Step 7 Select the traffic column being compared for Plan 2.

Step 8 Click OK.

View or Export Reports

Upon generating a report, it automatically opens in a window. From there you can select the various sections of the report from the left navigation pane.

To view a report after having closed it, select the Window->Reports menu. All reports generated on that plan file appear and are selectable from one report window.

If you want to save these reports for later use, you must save the plan file. However, running the report again overwrites the report. To preserve the information, you can export the data by report section. Right-click on any of the sections from the left navigation pane, and select Export Tables.

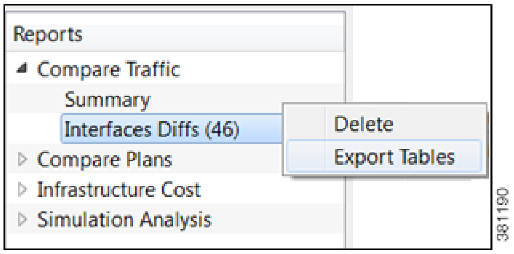

Cost Modeling

MATE Design provides basic cost modeling of network plans. Costs are specified for nodes and circuits, and these can be used to approximate infrastructure costs for the network or parts of the network. This allows for quick assessments of the costs of various upgrade alternatives, as well as for greenfield build-outs.

The specified circuit costs can also be used to estimate demand and demand grouping costs. This is the cost of routing traffic between particular locations or sets of locations in the network, taking into account the amount and cost of the circuit infrastructure used. This can be used to assess the cost of a customer, a particular service, or any other defined traffic source on the network.

Infrastructure Costs

An Infrastructure Cost report can help you approximate the total monetary cost of a network design based on estimated costs assigned to nodes and circuits. These costs are calculated from initial costs and per-period costs, where per-period costs are amortized based on the specified number of periods and interest rates. The report gives an infrastructure cost based on sums of all node and circuit costs, as follows.

- Total initial cost—Sum of all initial costs for each object type.
- Total cost per period—Sum of all per-period costs for each object type before the amortization is applied.
- Total cost—Total initial cost + total cost per period after being amortized to a net present value.
- Total infrastructure cost—Sum of the total node and total circuit costs.

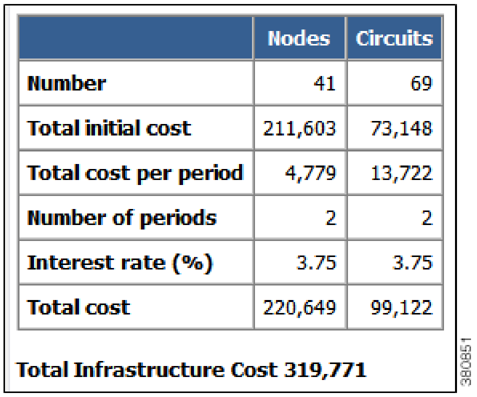

	Nodes	Circuits
Number	41	69
Total initial cost	211,603	73,148
Total cost per period	4,779	13,722
Number of periods	2	2
Interest rate (%)	3.75	3.75
Total cost	220,649	99,122

Total Infrastructure Cost 319,771

Cost Basis

Costs are based on any combination of Cost Initial and Cost Per Period properties that are assigned to nodes and circuits. A typical use case is to set nodes with an initial cost to identify the initial capital investment and to set circuit costs per period to analyze the operational cost of leasing circuits over time.

- Cost Initial—Static, one-time cost to which amortization and interest rates are not applied.
- Cost Per Period—Ongoing costs that are amortized based on the number of periods identified and on the interest rate specified when generating the Infrastructure Cost report.

You can set these properties manually. For circuits, you can also use the Cost Circuit Initializer to set the costs automatically based on circuit distance and capacity.

The total cost for each node and each circuit appears in its respective Cost Initial or Cost Per Period column in the Nodes table and Circuits table, respectively. The sum total of these node and circuit costs appear in the output of the Infrastructure Cost report and are used in its calculations.

An easy way to see their totals, as well as average, minimum, and maximum costs, without running the report is to right-click on the Cost Initial or Cost Per Period column heading.

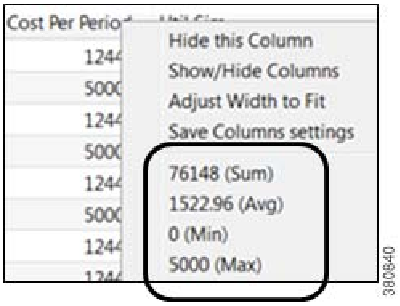

76148 (Sum)
1522.96 (Avg)
0 (Min)
5000 (Max)

Set Cost Properties

You can set these properties manually, as follows.

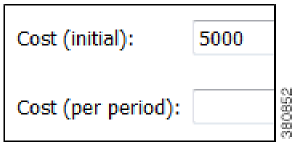

Cost (initial): 5000
Cost (per period):

Step 1 Select one or more nodes or select one or more circuits. Either right-click one of these objects and select Properties, or double-click one of these objects.

Step 2 Select appropriate tab.

- Nodes—Select the Inventory tab.
- Circuits—Select the Advanced tab.

Step 3 Enter values in the Cost (initial) field, Cost (per period) field, or both, and click OK.

Cost Circuit Initializer

The Cost Circuit initializer offers an automated way of setting circuit costs. For example, it enables you to set realistic per-period circuit costs based on distance, capacity, or both.

The total cost for each circuit is the sum of the fixed cost plus the variable cost, where the variable cost is the result of multiplying the per-unit cost times the units selected (distance and/or capacity).

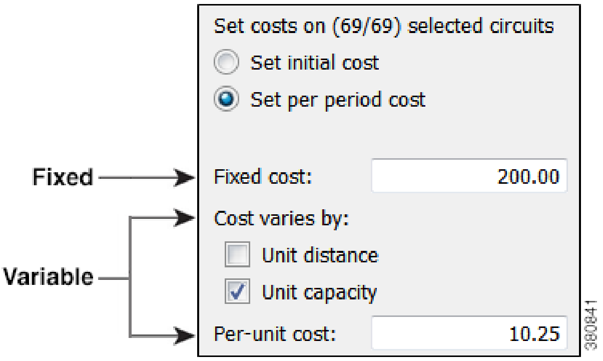

Set costs on (69/69) selected circuits
○ Set initial cost
● Set per period cost
Fixed cost: 200.00
Cost varies by:
□ Unit distance
☑ Unit capacity
Per-unit cost: 10.25

To access this initializer, select the Initializers->Circuit Cost menu. Options are as follows.

- Set cost—Set initial or per-period cost. The options that follow are applied to the choice selected here. Since initial and per-period costs are different values that do not overwrite each other, you can set one and then run the initializer again to set the other.
- Fixed cost—A static cost that is most often used when setting initial costs.
- Variable cost—A cost based on a circuit's distance, capacity, or both. Whichever unit is selected (one or both), the value is multiplied by the "Per-unit cost" value.
  - Per-unit cost—Per-unit monetary cost.
  - Unit Distance—Distance Sim value of the circuit.

    Example: If only "Unit Distance" were selected, if the Distance Sim were 700, and if the "Per-unit cost" were 0.50, the variable value for this circuit would be 700 x 0.50 = 350.00.
  - Unit Capacity—Capacity Sim value of the circuit

    Example: If only "Unit Capacity" were selected, if the Capacity Sim were 400, and if the "Per-unit cost" were 0.25, the variable value for this circuit would be 400 x 0.25 = 100.00.

Example: This example sets a per-period cost for a circuit that has a fixed cost of 250, a per-unit cost of 2, circuit distance of 850, and a capacity of 1,000. Including all of these options in the calculation, the cost of this circuit would be as follows.

Fixed + (Per-unit Cost x Distance Sim x Capacity Sim)

250 + (2 x 850 x 1,000) = 250 + 1,700,250 = 1,700,250

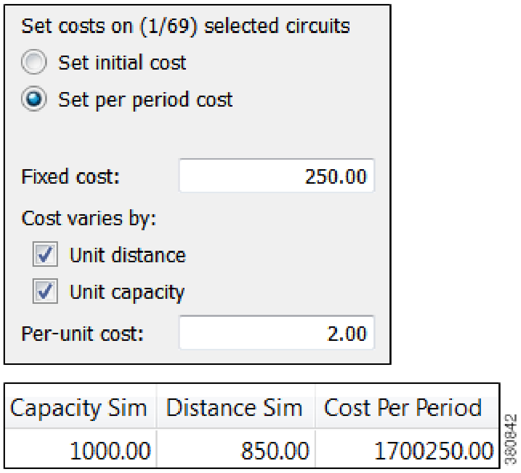

Set costs on (1/69) selected circuits

○ Set initial cost

● Set per period cost

Fixed cost: 250.00

Cost varies by:

☑ Unit distance

☑ Unit capacity

Per-unit cost: 2.00

Capacity Sim	Distance Sim	Cost Per Period
1000.00	850.00	1700250.00

Infrastructure Cost Report

To create an infrastructure cost report, follow these steps.

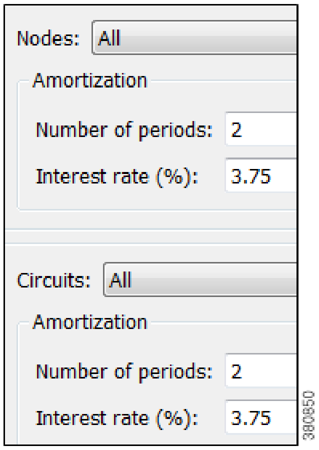

Nodes: All

Amortization

Number of periods: 2

Interest rate (%): 3.75

Circuits: All

Amortization

Number of periods: 2

Interest rate (%): 3.75

Step 1 For any node or circuits you want included in the report, set their initial and/or per-period costs. For more information, see the Set Cost Properties section.

Step 2 If creating a report for selected nodes or circuits, select them before opening the report dialog box.

Step 3 Select the Tools->Reports->Cost Analysis menu.

Step 4 Nodes

- **a.** Select whether to include all nodes, selected nodes, or specifically tagged nodes.
- **b.** If you are calculating per-period costs, set the number of periods and the interest rate.

Note Using tags is a convenient way to determine costs for a specific region of the network or for specific customers.

Step 5 Circuits

- **a.** Select whether to include all circuits, selected circuits, or specifically tagged circuits.
- **b.** If you are calculating per-period costs, set the number of periods and the interest rate.

Step 6 Click OK.

Example Infrastructure Cost Workflow

This example shows the generation of an Infrastructure Cost report based on a specific set of nodes and circuits, both tagged as the East Region. It sets initial costs on the nodes and per-period costs on the circuits based on their distance. The report is for 4 periods, and the interest rate is 5%.

Step 1 Set initial node costs in the node Properties dialog box.

- **a.** Filter to the nodes containing the East Region tag.
- **b.** In this example, there are 20 nodes with this tag. Set the 10 core nodes to an initial cost of 2,000 each and the 10 edge nodes to an initial cost of 1,000 each.

The Nodes table lists these costs in the Cost Initial column. Right-click on this column heading to see the total Cost Initial for all of these nodes. Here, it is 30,000. Since we are only interested in initial node costs, this is the same value that the Infrastructure Cost report will show for total initial node cost.

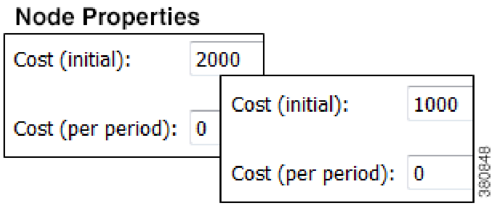

Node Properties

Cost (initial): 2000
Cost (per period): 0

Cost (initial): 1000
Cost (per period): 0

Step 2 Run the Latency and Distance initializer to set Distance Sim values on the circuits (Initializers->Latency and Distance menu). For information, see the Simulation chapter.

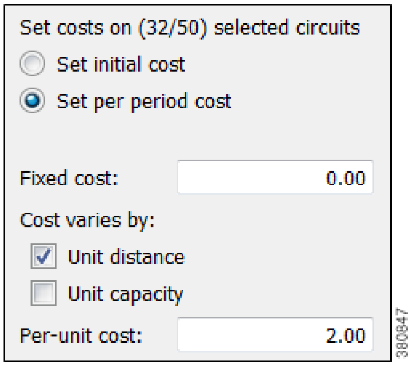

Set costs on (32/50) selected circuits
○ Set initial cost
● Set per period cost

Fixed cost: 0.00

Cost varies by:
☑ Unit distance
□ Unit capacity

Per-unit cost: 2.00

Step 3 Set per-period circuit costs using the Circuit Cost initializer on circuits tagged with East Region.

a. Filter to the circuits containing the East Region tag, and select these circuits.

b. Run the Circuit Cost initializer (Initializers->Circuit Cost menu). For this example, there are no fixed costs, a per-unit cost of 2, and the initializer is only using Distance Sim.

The Circuits table lists these costs in the Cost Per Period column. Right-click on this column heading to see the total Cost Per Period for all of these circuits. Here, it is 35,063.46, and this is the number that is amortized in the Infrastructure Cost report based on 4 periods with a 5% interest rate.

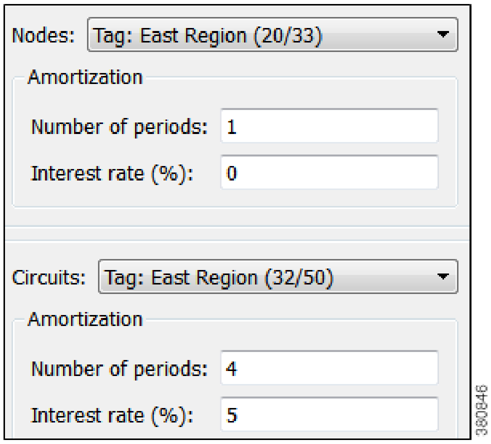

Nodes: Tag: East Region (20/33)

Amortization

Number of periods: 1

Interest rate (%): 0

Circuits: Tag: East Region (32/50)

Amortization

Number of periods: 4

Interest rate (%): 5

Step 4 Now that costs are set up for the region, generate the Infrastructure Cost report.

a. Select the Tools->Reports->Cost Analysis menu.

b. Select the East Region nodes and circuits.

c. Since there are no per-period costs for nodes, no further node options are set.

Set the circuits “Number of periods” to 4, and the “Interest rate (%)” to 5.

The resulting report performs these calculations.

- Sums the total initial costs for all nodes tagged with East Region to determine the total node cost.

- Sums the per-period costs for all circuits tagged with East Region to determine the total per-period circuit cost.
- Amortizes the per-period circuit costs across 4 periods using a 5% interest rate to determine the total circuit cost.
- Adds the total node cost to the total circuit cost to determine the total infrastructure cost.

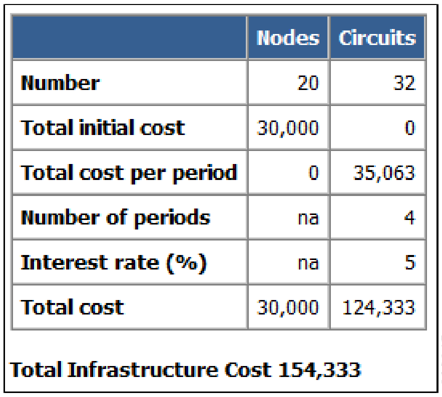

	Nodes	Circuits
Number	20	32
Total initial cost	30,000	0
Total cost per period	0	35,063
Number of periods	na	4
Interest rate (%)	na	5
Total cost	30,000	124,333

Total Infrastructure Cost 154,333

380845

Demand and Demand Grouping Costs

Two estimates of the monetary costs of demands are automatically derived when you set per-period costs for circuits, as follows. These appear as the following column headings in the Demands table.

- Cost by Capacity—Total of the Cost Per Period for the circuits using the demand as routed, scaled by proportion of demand traffic to the total interface capacity.

  It is divided by two to share the circuit cost between the costs of traffic in two directions.

  (Sum Cost Per Period of All Circuits Using the Demand) x (Demand Utilization on the Interface / 2)

  Where, demand utilization on the interface = (Interface Utilization / Capacity)
- Cost by Util—Total of the Cost Per Period for the circuits using the demand as routed, scaled by proportion of demand traffic to the total traffic through the circuit.

  (Sum Cost Per Period of All Circuits Using the Demand) x (Demand Utilization on the Interface) / (Sum of the Utilizations on the Two Interfaces)

  Where, demand utilization on the interface = (Interface Utilization / Capacity)

  If there is only one demand through the circuit, then the full cost of that circuit is allocated to that demand.

Demand grouping costs are calculated when you run the Demand Grouping report (Tools->Reports->Demand Groupings). The Cost by Capacity and Cost by Util are the sum total of the costs of the demands within the demand grouping. For more information on demand groupings, see the Forecasting Traffic chapter.

Example

In this example there are four nodes connected by three circuits.

- Each circuit has 10,000 Mbps capacity.
- Each circuit has a Cost Per Period of 100.
- There are two demands.

  The dnv-mia demand has 2,000 Mbps of traffic.

  The mia-dnv demand has 5,000 Mbps of traffic.
- These two demands are in one demand grouping named Acme_region.

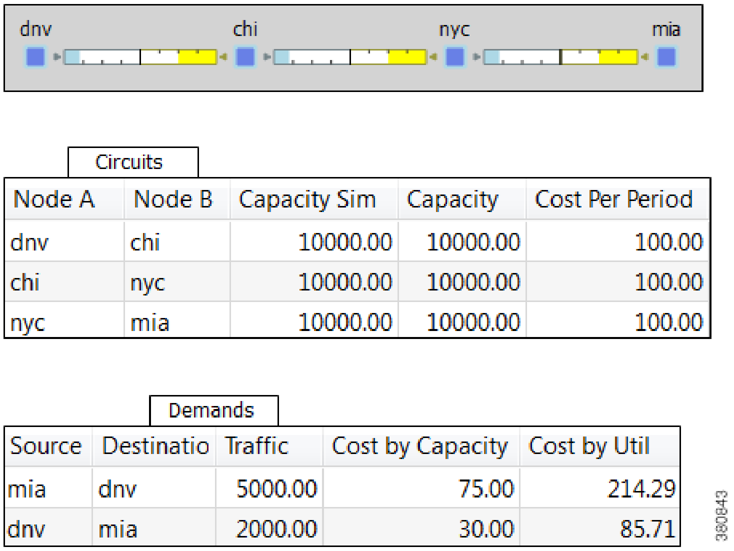

Circuits

Node A	Node B	Capacity Sim	Capacity	Cost Per Period
dnv	chi	10000.00	10000.00	100.00
chi	nyc	10000.00	10000.00	100.00
nyc	mia	10000.00	10000.00	100.00

Demands

Source	Destinatio	Traffic	Cost by Capacity	Cost by Util
mia	dnv	5000.00	75.00	214.29
dnv	mia	2000.00	30.00	85.71

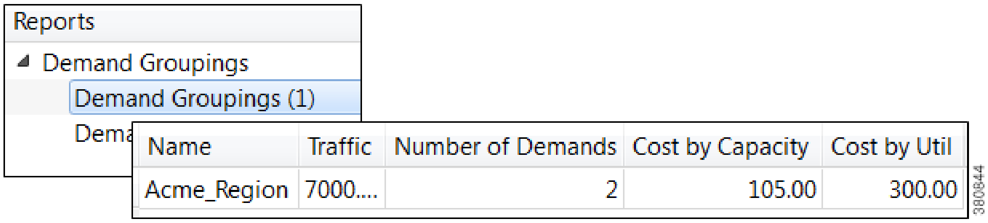

Name	Traffic	Number of Demands	Cost by Capacity	Cost by Util
Acme_Region	7000....	2	105.00	300.00

The Cost by Capacity calculation is as follows.

- dnv-mia demand: (100 + 100 + 100) x ((2,000/10,000)/2) = 30
- mia-dnv demand: (100 + 100 + 100) x ((5,000/10,000)/2) = 75
- Acme_region demand grouping: 30 + 75 = 105

The Cost by Util calculation is as follows.

- dnv-mia demand: (100 + 100 + 100) x 2,000/(2,000 + 5,000) = 85.71
- mia-dnv demand: (100 + 100 + 100) x 5,000/(2,000 + 5,000) = 214.29
- Acme_region demand grouping: 85.71 + 214.29 = 300

CHAPTER 21

MATE GUI and Remote MATE Live

Plan files and templates are stored in archive repositories, which can be installed on a different device than the one on which the MATE GUI is used. From the MATE GUI, you can extract MATE Live .pln plan files and templates from remote archive directories and make whatever modifications are necessary. You can re-insert them as template files into an archive directory used by the MATE Live application.

Note The recommendation is to open the remote plan file, modify it to create the visual template, and save it to the remote archive as a template. When uploading a template to a remote archive, you are overwriting the template that is in the archive.

Prerequisites

- You must have an admin role assigned to you through the web System UI, and you must know that username and password.
- To save templates to a MATE Live archive, that archive must first be configured through the MATE Live UI, Settings->General Settings page.

Step 1 In the MATE GUI, use the File menu to start the process.

- To download a MATE Live plan file or template from a remote archive, select the File->Open from->MATE Live menu.
- To upload a MATE Live template to a remote archive, select the File->Save to->MATE Live menu. You must have an administrative username and password to upload templates.

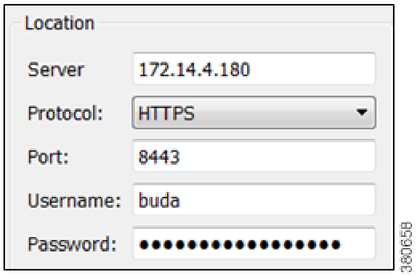

Location

Server: 172.14.4.180

Protocol: HTTPS

Port: 8443

Username: buda

Password: ●●●●●●●●●●●●●●●●●

Step 2 After opening the dialog box, enter the hostname or IP address of the server.

Step 3 Identify how to connect to the server by selecting the appropriate protocol and entering the port number (for example, HTTPS 8443).

Step 4 Enter the username and password that gives you access to the server. Both are case sensitive. If you do not know the password, contact your system administrator. You must have an admin role assigned to you.

Step 5 Optional: Save the data you entered for future use.

- To save all settings except the password, select Save Settings.
- To save the password for future use, select Save Password.

Step 6 If saving a template, click OK and stop here. If opening a plan file or template, continue to step 6.

Step 7 If opening a template, select the Template option.

If opening a plan file, identify which one by selecting the latest or by entering a specific time using a YY/MM/DD HH:MM format.

Click OK.

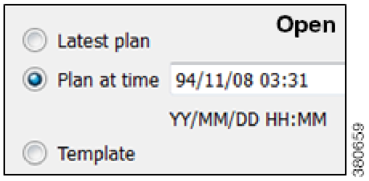
Open

○ Latest plan

● Plan at time 94/11/08 03:31

YY/MM/DD HH:MM

○ Template

Related Topics

- *MATE Integration and Development Guide*
- *Table Schema and CLI Reference*

CHAPTER 22

MATE GUI and Remote MATE Design Archive

Plan files and templates are stored in archive repositories, which can be installed on a different device than the one on which the MATE GUI is used. From the MATE GUI, you can extract .pln plan files from remote archive directories, make whatever modifications are necessary, and re-insert them as .pln files into an archive used by the MATE Design Archive application. You can also open, modify, and then save a remote archive template.

Prerequisite

You must have an admin role assigned to you through the web System UI, and you must know that username and password.

Note The recommendation is to open the remote plan file, modify it to create the visual template, and save it to the remote archive as a template. When uploading a template to a remote archive, you are overwriting the template that is in the archive.

Step 1 In the MATE GUI, use the file menu to start the process.

- To open a MATE Design Archive plan file or template from a remote archive directory, select the File->Open from->Design Archive menu.
- To save a MATE Design Archive plan file or template in a remote archive directory, select the File->Save to->Design Archive menu. You must have an administrative username and password to upload plan files or templates.

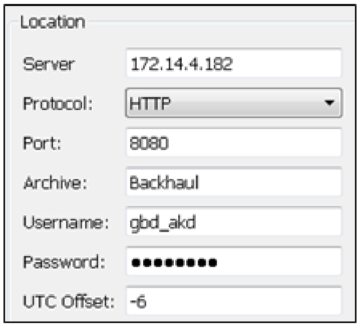

Location

Server: 172.14.4.182

Protocol: HTTP

Port: 8080

Archive: Backhaul

Username: gbd_akd

Password: ●●●●●●●●

UTC Offset: -6

Step 2 After opening the dialog box, enter the hostname or IP address of the server.

Step 3 Identify how to connect to the server by selecting the appropriate protocol and entering the port number (for example, HTTP 8080).

Step 4 Enter the name of the archive you are accessing.

Step 5 Enter the username and password that gives you access to the server. Both are case sensitive. If you do not know the password, contact your web system administrator. You must have an admin role assigned to you.

Step 6 Optional: Save the data you entered for future use.

- To save all settings except the password, select Save Settings.
- To save the password for future use, select Save Password.

Step 7 Enter the UTC offset to apply to the date specified in the "Plan at Time" field. If you do not enter a date, the file is opened or saved according the server's UTC time.

Step 8 If opening a plan file, identify which one by the latest or by a specific time.

If saving a plan file, identify the time as current or as a specific time.

If entering a time, use the YY/MM/DD HH:MM format.

If opening or saving a template, select the Template option.

Then click OK.

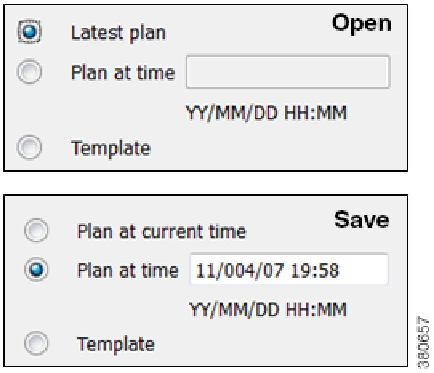

Open
● Latest plan
○ Plan at time
YY/MM/DD HH:MM
○ Template

Save
○ Plan at current time
● Plan at time 11/004/07 19:58
YY/MM/DD HH:MM
○ Template

Upload Templates from MATE Design Archive UI

Step 1 From the MATE Design Archive home page, select the ADMIN link.

Step 2 Select the archive into which you are uploading the template, and then click Edit Configuration.

Step 3 In the Upload Template dialog box, the archive directory name on the server and the current template name appear.

Browse to or enter the name of the local file to be used as the new template, and click Upload.

Related Topics

- *MATE Integration and Development Guide*
- *Table Schema and CLI Reference*

CHAPTER 23

MATE GUI and Remote WAE Core Server

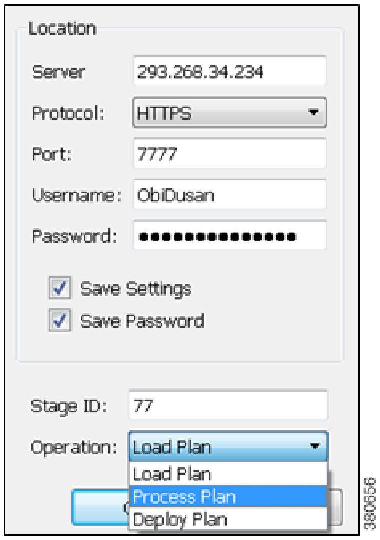

Location

Server: 293.268.34.234

Protocol: HTTPS

Port: 7777

Username: ObiDusan

Password: ●●●●●●●●●●●●●●

☑ Save Settings

☑ Save Password

Stage ID: 77

Operation: Load Plan

From the MATE GUI, you can extract .pln plan files from a remote WAE Core server[1], make whatever modifications are necessary, and re-insert them to the remote server. It is important to note that these plan files are not stored with the plan files generated by the Collector Module.

The Plan Module (on the WAE Core server) contains the *working plan file*, which represents the current state of the network. It can also contain one or more *staging areas*, where each staging area enables you to work on plan files before deploying them.

Saving the plan file to the working area of the Plan Module overwrites the existing plan file. This means that any newly discovered or recently modified plan files are overwritten. For this reason, if you are not deploying the plan file, it is a best practice to access a staging area.

Note It is strongly recommended that you verify a plan file before deploying it to the network.

Prerequisites

- If using the staging area, it must already exist on the Plan Module, and you must know the stage ID.
- The `/opt/cariden/software/wae-core/etc/com.cisco.wano.nsps.nbrs.cfg` file must either have the authenticationEnabled property set to false, or you must know the values entered as the username and password properties.

Step 1 Use the File menu to start the process.

- To open a plan file from the WAE Core server, select the File->Open from->WAE menu.
- To save a plan file to the WAE Core server or to deploy a plan file to the network, select the File->Save to->WAE menu.

Step 2 After opening the dialog box, enter the hostname or IP address of the server.

[1] The Plan Module, Optimization and Prediction Module, and Deployer Module all use the WAE Core server.

Step 3 Identify how to connect to the server by selecting the appropriate protocol and entering the port number (for example, HTTP 7777).

Step 4 If the authentication is enabled, enter the same username and password.

Step 5 Optional: Save the data you entered for future use.

- To save all settings except the password, select Save Settings.
- To save the password for future use, select Save Password.

Step 6 If using a staging area for this access, enter the stage ID number.

Step 7 If saving a plan file, select the operation, as follows.

- Load Plan—Save a plan file to the WAE Core server. No further processing is done to this plan file on the platform.
- Process Plan—Save a plan file to the WAE Core server and simulate failure scenarios such that upon querying a demand, worst-case analysis is automatically performed on interfaces, circuits, and SRLGs.
- Deploy Plan—Compare the plan file with the working plan file that currently resides in the Plan Module, and then deploy the differences (LSPs only) directly to the network. A copy of it is also saved to the working plan file area in the Plan Module.

Step 8 Click OK.

Plot Legend for Design Layouts

This legend is for Design layouts used in the MATE GUI.

Note All traffic is displayed in Mbps. In the network plot, the traffic utilization colors represent outbound traffic.

Legend	**Descriptions**
Interfaces and Circuits	Colors
	Circuit Active, Failed, and Operational States
	Markings and Styles
Sites	Sites
Nodes	Nodes
External Endpoint Members	External Endpoint Members
Demands and Routings	Demands and Routings
LSPs	LSPs
	LSP Paths
	LSP Named Paths
	LSP Actual Paths
	Segment-Routed Paths
Layer 1 View	Layer 1 Links
	Layer 1 Circuit Paths
	Sites in Layer 1 View
	Layer 1 Nodes

Interfaces and Circuits

Colors

Table 24-1 Interface and Circuit Colors

Graphic	Description
The traffic utilization colors are customizable. For information, see the *MATE GUI Visualization Guide*.	
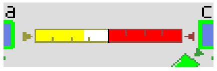	Outbound traffic utilization is depicted by colors. Default colors are light blue, light green, yellow, orange, red, and purple, and left to right, each of these colors represent increased utilization. • In the Measured Traffic and Simulated Traffic views, the color shows the percent of traffic utilized for the available interface capacity. • In the Worst-Case Traffic view, the color represents the worst-case utilization on each interface. This is the highest utilization that a particular interface experiences over all the failure scenarios considered when running a simulation analysis. See the Simulation Analysis chapter in the *MATE Design User Guide*. • In the Failure Impact view, the color represents the failure impact of a circuit. This is defined as the highest utilization experienced by any interface if that circuit were to fail. See the Simulation Analysis chapter in the *MATE Design User Guide*. • In the LSP Reservation view, the color represents the total setup bandwidth that is reserved by LSPs through the interface.
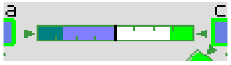	Blue indicates a selected interface or circuit.
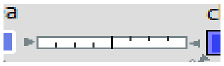	White indicates 0% utilization.
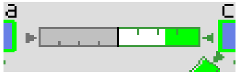	Interfaces are light gray for the following reasons. • The interface has no traffic measurement (Measured Traffic view). • This amount of interface capacity is not available due to a QoS bound (Measured Traffic and Simulated Traffic views). • No simulation analysis was performed (Worst-Case and Failure Impact views). • Circuit is operationally down.
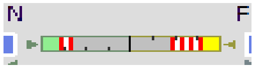	QoS policy violations are indicated by red and white bars. This is displayed when utilization exceeds the QoS bound, which is set through policy or interface queue properties. See the Quality of Service Simulation chapter in the *MATE Design User Guide*.

Graphic	Description
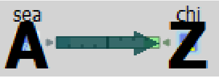	Shortest IGP path between two selected nodes or sites. See the IGP Simulation chapter in the *MATE Design User Guide*.
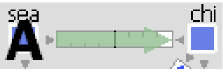	Shortest latency path between two selected nodes or sites. See the IGP Simulation chapter in the *MATE Design User Guide*.
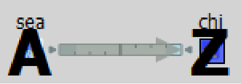	Shortest TE path between two selected nodes or sites. See the IGP Simulation chapter in the *MATE Design User Guide*.

Circuit Active, Failed, and Operational States

Table 24-2 Circuit Active, Failed, and Operational States

Graphic	Description
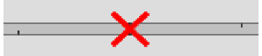	Failed circuit.
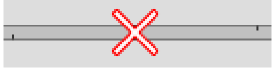	Circuit is not operational for one of the following reasons. • A node connected to the circuit failed. • A site failed that contains a node connected to the circuit. • The circuit is a LAG[a] (port channel) that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to failure of some of its ports or port circuits. • L1 link or L1 node used by the circuit failed. • An SRLG failed that contains the circuit or contains one of the above objects that are required for circuit operation.
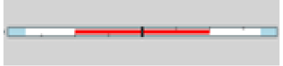	The circuit is a LAG, and one or more ports or port circuits in the LAG failed. The red strikeout line is the percent of unavailable capacity due to the failure.
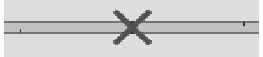	Circuit is inactive.

Graphic	Description
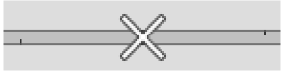	Circuit is not operational for one of the following reasons. • A node connected to the circuit is inactive. • An inactive site contains a node connected to the circuit. • The circuit is a LAG that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to some of its ports or port circuits being inactive. • Mapped L1 circuit cannot be routed. • An SRLG is inactive, and the SRLG contains the circuit or contains one of the above objects that are required for circuit operation.
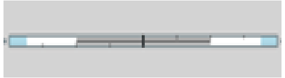	The circuit is a LAG and one or more ports or port circuits in the LAG is inactive. The gray strikeout line is the percent of unavailable capacity due to the inactive ports or port circuits.

[a] A link aggregation group (LAG) is a group of physical ports that are bundled into a single logical interface. It is also known as *bundling* or *trunking*.

Markings and Styles

Table 24-3 Interface and Circuit Markings and Styles

Graphic	Description
Interface styles and colors are customizable. For information, see the *MATE GUI Visualization Guide*.	
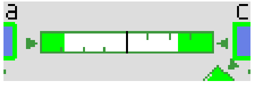	Circuits are rectangles divided in half by a vertical black line. Each half shows an interface. Tick marks show 25%, 50%, 75% utilization levels.
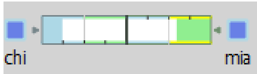	If parallel circuits are grouped for visual display, each side of the circuit shows its constituent interfaces. Here, there are two circuits (two interfaces on each side) grouped by site. If the circuits are between sites, the utilization includes the utilization of nested interfaces, if applicable. • The color fill on each side shows the average utilization of the grouped interfaces, which is weighted against the sum of the grouped capacities. For example, the average utilization for the interfaces from chi to mia is color filled green. • Each interface within a grouped circuit has a thin colored border showing the utilization of that interface. For example, one interface going from mia to chi is bordered by yellow and another by blue, each showing the traffic utilization for the bordered interface.
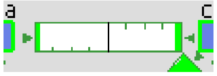	Width of circuit shows capacity. Width increases a fixed amount for every 4x increase in capacity.
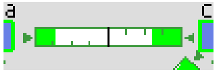	Full-length interface style, which is the default.

Graphic	Description
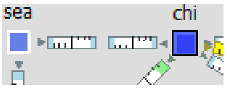	Shortened interface style, showing the same information as a full-length view, just in shorter representations that are repeated at each end of the circuit. In this example, all chi interfaces are shortened at both ends of the circuit.
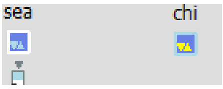	Collapsed interface view, showing the presence of interfaces as up/down arrows inside the site. The color of the arrows show the highest utilization of any of the collapsed interfaces (in this example, blue for sea and yellow for chi). In this example, all chi interfaces are collapsed at both ends of the circuit.

Sites

Table 24-4 Sites

Graphic	Description
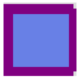	A site is shown as a blue square with a border. The color of the border is the maximum utilization color of all intrasite interfaces within the site. These colors are customizable. For information, see the *MATE GUI Visualization Guide*.
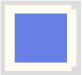	A white border surrounds the site if all intrasite interfaces have zero utilization.
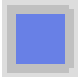	A gray border surrounds the site in the following cases. • There is no measured traffic on intrasite interfaces (Measured Traffic view). • No simulation analysis was performed (Worst-Case and Failure Impact views).
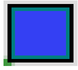	Solid blue indicates the site is selected.
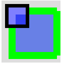	One or more sites, nodes, or circuits within the site are selected.
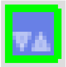	Up/down triangles show that one or more egress interfaces from within the site have been collapsed.
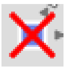	The site failed.
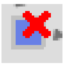	One or more sites, nodes, or circuits within the site failed.
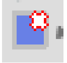	One or more nodes or circuits within the site are not operational because they are contained in a failed SRLG.

Graphic	Description
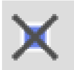	The site is inactive.
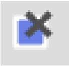	One or more sites, nodes, or circuits within the site are inactive.
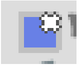	One or more nodes or circuits within the site are not operational because they are contained in an SRLG that is not active.
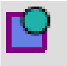	One or more selected VPN nodes exist on nodes within the site.

Nodes

Table 24-5 Nodes

Graphic	Description
	A node is shown as a blue rectangle. The green border surrounds the node if it is the source or destination of simulated or measured traffic, or of LSPs with nonzero bandwidth reservations (Simulated Traffic, Measured Traffic, and LSP Reservation views, respectively).
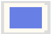	A white border surrounds the node if it is not the source or destination of simulated or measured traffic, or of LSPs with nonzero bandwidth reservations (Simulated Traffic, Measured Traffic, and LSP Reservation views, respectively).
	A node always has a gray border in the Worst-Case Traffic and Failure Impact views.
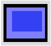	Solid blue indicates the node is selected.
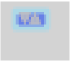	Up/down triangles show that some interfaces from the node have been collapsed.
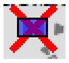	Failed node.
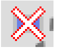	Node is not operational because it is contained in a failed site or SRLG.
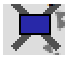	Node is not active.

Graphic	Description
☒	Node is not operational because it is contained in a site or an SRLG that is not active.
■●	One or more selected VPN nodes exist on the node.

External Endpoint Members

Table 24-6 External Endpoints

Graphic	Description
If the external endpoint member is a node, this icon appears under that node in the plot. If the member is an external AS or an external AS via a node, this icon appears under all nodes in the AS. For more information, see the Advanced Routing with External Endpoints chapter in the *MATE Design User Guide*.	
✖➔	Failed external endpoint member.
✖➔	External endpoint member is not operational because it is contained in a failed SRLG.
✖➔	External endpoint member is not active.
✖➔	External endpoint member is not operational because it is contained in an SRLG that is not active.

Demands and Routings

Table 24-7 Demands and Routings

Graphic	Description
For more information, see the Traffic Demand Modeling chapter in the *MATE Design User Guide*.	
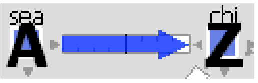	"A" and "Z" mark source and destination site/node of a demand. Blue arrows show demands routed through an interface on a circuit.
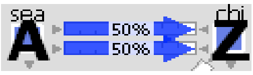	Percentages show the proportion of a split demand (for example from ECMP) through an interface.

Graphic	Description
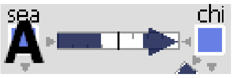	Reroute of a demand under a failure.
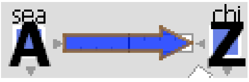	Demand is routed through an LSP.
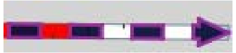	Demand is rerouted through an FRR LSP.
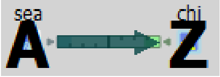	Shortest IGP path of a demand.
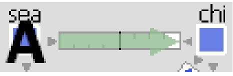	Shortest latency path of a demand.
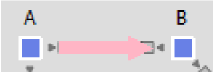	Multicast flow and multicast flow hops. See the Multicast Simulation chapter in the *MATE Design User Guide*.

LSPs

For more information, see the MPLS Simulation chapter in the *MATE Design User Guide*.

Table 24-8 LSPs

Graphic	Description
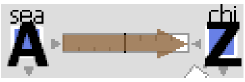	"A" and "Z" mark source and destination site/node of an LSP. Brown arrows show how LSP traffic is routed.
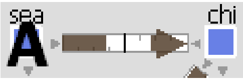	Rerouting of an LSP traffic when a failure occurs in Convergence mode
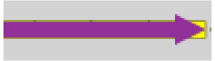	Fast Reroute (FRR) LSP path that would be taken if there were a failure that required its use.
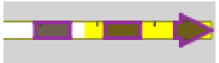	FRR LSP that is currently rerouting LSP traffic around a failure in FRR mode.

Graphic	Description
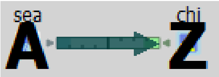	Shortest IGP path of an LSP.
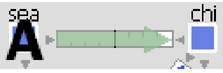	Shortest latency path of an LSP.
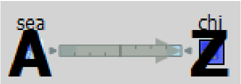	Shortest TE path of an LSP.

LSP Paths

Table 24-9 LSP Paths

Graphic	Description
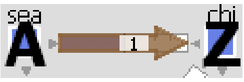	"A" and "Z" mark source and destination site/node of an LSP path. The brown arrow indicates the LSP path is active, that is, carrying traffic. The number in the gray rectangle (here 1) is the path option.
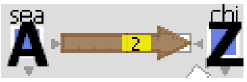	A lighter brown arrow indicates the LSP path is not active and is not carrying traffic. The yellow rectangle indicates a standby path. The number in a yellow rectangle (here 2) is the path option.
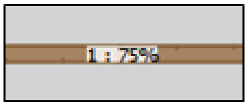	The percent value is the loadshare percentage that this LSP path is carrying for the selected LSP. In this example, the LSP loadshare percentage is 75%.

LSP Named Paths

Table 24-10 LSP Named Paths

Graphic	Description
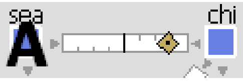	"A" marks the source of a named path. The brown diamond on an interface denotes it as a named path hop. A black dot in the center denotes a strict hop.
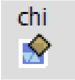	A brown diamond on a site denotes that the site contains a node or interface that is a named path hop. No black dot in the center denotes a loose hop. The site is shown as containing a loose hop unless all hops within the site are set to strict or exclude.
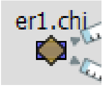	A brown diamond on a node denotes it as a named path hop. No black dot in center denotes a loose hop.

Graphic	Description
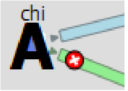	"A" marks the source of a named path. A red circle with a white X on an interface denotes it as a named path hop that is set to *exclude*.
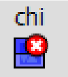	A red circle with a white X on a site denotes that the site contains a named path hop on a node or interface and that it is set to *exclude*.
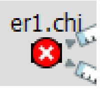	A red circle with a white X on a node denotes it as a named path hop that is set to *exclude*.

LSP Actual Paths

Table 24-11 LSP Actual Paths

Graphic	Description
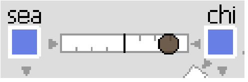	A dark brown circle on an interface denotes an LSP actual path hop on the interface.
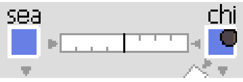	A dark brown circle on site denotes an actual path hop on a node or interface within the site.
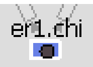	A dark brown circle on a node denotes an actual path hop on the node.

Segment-Routed Paths

Table 24-12 SR LSP Paths

Graphic	Description
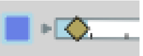	An interface segment in a segment-routed (SR) LSP.
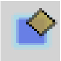	The site contains a node segment or an interface segment.
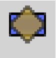	A node segment in an SR LSP.

Layer 1 View

For more information, see the Layer 1 Simulation chapter in the *MATE Design User Guide*.

Layer 1 Links

Table 24-13 Layer 1 Links in L1 View

Graphic	Description
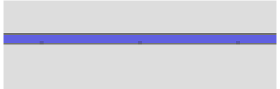	Selected L1 link.
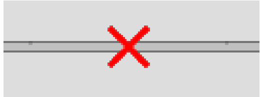	L1 link failed.
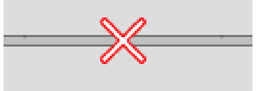	L1 link is not operational because of one of the following reasons. • L1 node failed that is connected to the L1 link. • A site failed that contains the L1 link. • An SRLG failed that contains the L1 link. • An SRLG failed that contains an L1 node connected to the L1 link.
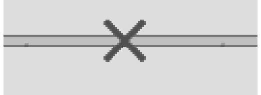	L1 link is inactive.
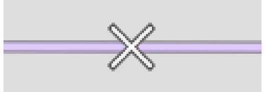	L1 link is inactive for one of the following reasons. • An L1 node that is connected to the L1 link is inactive. • An inactive SRLG contains the L1 link.
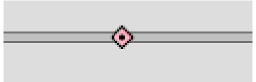	L1 link is a strict L1 circuit path hop.
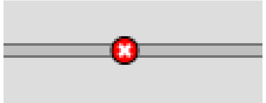	L1 link is an excluded L1 circuit path hop.

Layer 1 Circuit Paths

Table 24-14 Layer 1 Circuit Paths in L1 View

Graphic	Description
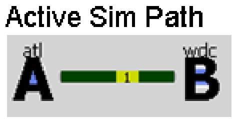 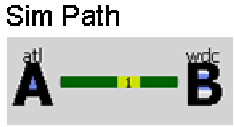	Simulated L1 circuit path from L1 node "A" to L1 node "B" highlights in green when any of the following are selected. • L1 circuit or L1 circuit path • Circuit or associated interface mapped to L1 circuit • Port circuit mapped to L1 circuit The simulated active path appears in a darker green than simulated paths that are not active.
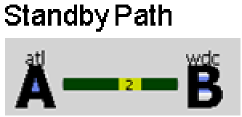 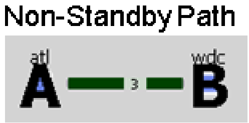	Path options are numerically identified within a rectangular label. Standby paths have a yellow background, and non-standby paths have a gray background. In this example, path option 2 is a standby path, and path option 3 is a non-standby path.
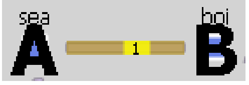	Actual L1 circuit path.
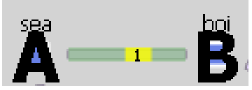	Shortest latency of an L1 circuit path
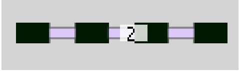	Rerouted L1 circuit path that is in use.

Sites in Layer 1 View

Table 24-15 Sites in Layer 1 View

Graphic	Description
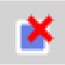	One or more sites, L1 nodes, or links within the site failed.
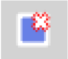	One or more L1 nodes or L1 links within the site are not operational because they are contained in a failed SRLG.
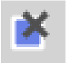	One or more L1 nodes or L1 links within the site are inactive.
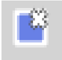	One or more L1 nodes or L1 links within the site are not operational because they are contained in an SRLG that is not active.
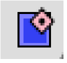	Site contains an L1 node that is a strict L1 circuit path hop.
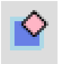	Site contains an L1 node that is a loose L1 circuit path hop.
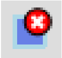	Site contains an L1 node that is excluded as an L1 circuit path hop.

Layer 1 Nodes

Table 24-16 Layer 1 Nodes in L1 View

Graphic	Description
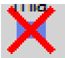	L1 node failed.
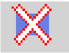	Node is not operational because it is contained in a failed site or SRLG.
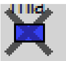	L1 node is inactive.
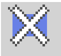	L1 node is not operational because an SRLG containing it is inactive.

Graphic	Description
◈	L1 node is a strict L1 circuit path hop.
◆	L1 node is a loose L1 circuit path hop.
⊗	L1 node is excluded as an L1 circuit path hop.