HP Route Analytics Management Software

Software Version: 9.00

User's Guide

Document Release Date: March 2010
Software Release Date: March 2010

invent

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

Trademark Notices

Linux is a U.S. Registered trademark of Linus Torvalds.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Unix® is a registered trademark of The Open Group.

Documentation Updates

This guide's title page contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

http://h20230.www2.hp.com/selfsolve/manuals

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

You can visit the HP Software Support web site at:

http://www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software Support Online provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the HP Software Support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

Contents

1 Introduction

This chapter introduces the Route Analytics Management Software route analytics tools.

Chapter contents:

About Route Analytics Management Software

The Route Analytics Management Software (RAMS) is an IP Route Analytics tool that listens to routing protocols and builds a real-time routing topology map. The map enables you to visualize and understand the dynamic operation of your network. RAMS Traffic also collects and aggregates traffic data, enabling you to view traffic flows on top of the routing topology.

RAMS offers the following powerful contributions to network planning and analysis:

- **Unified, real-time routing topology view**. View complex topologies hierarchically or by protocol, autonomous system (AS), Interior Gateway Protocol (IGP) area, or Border Gateway Protocol (BGP)/Multiprotocol Label Switching (MPLS) virtual private network (VPN). The History Navigator window lets you play back a history of your routing topology changes.

- **Monitoring and alerts**. Monitor vital service parameters such as network churn and prefix flaps, watch for changes in specific end-to-end service paths and prefixes, and look for degrading redundancy. You can also raise alerts on all watched parameters to head off costly outages.
- **Interactive analysis**. Perform before and after comparisons and detailed event analysis using a comprehensive routing base and complete event history to rapidly establish the cause of the problem.
- **Planning support**. Display network activity patterns to help optimize performance and minimize unnecessary transit fees or bandwidth costs. You can simulate a link failure or change link metric costs, to see how your routing topology responds to specific failures or upgrades. You can also import and export these simulated changes to manage multiple routing scenarios using external editors.
- **Reports**. View trends and identify emerging issues before they become problems. You can generate graphical user interface (GUI)-based reports for any recorded time period to obtain key information about network health.
- **IPv4 and IPv6**. View and report on IPv6 prefixes, if the topology supports IPv6. IPv6 is supported for BGP and IS-IS.

How Route Analytics Management Software Appliances Operate

Route Analytics Management Software appliances physically connect to the network directly to one of the routers on the network or through a switch or hub. The appliances then establish communication with several routers in the network through the routing protocol over this single physical connection. It is only necessary for the appliance to listen to link-state routing protocols such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS) in one location, because each router knows of all adjacencies in the network. Link-state routers send periodic update messages that communicate network information to each other, and to the appliance.

OSPFv3 adds the following additional information:

- **Routers**—The appliance gathers information on the R and V6 bit from Options field. For further information, see RFC 5340, section 2.7 (Packet Format Changes).
- **IPv6 prefixes**—The appliance collects the LA and NU bits from the Prefix Options, as described in RFC 5340, section A.4.1.1.

Unlike links between OSPF and IS-IS routers, BGP peerings may not follow physical paths. BGP routers and their peerings are discovered indirectly by receiving routes with a next hop attribute that contains the address of a BGP router. Beyond the physical connection between a BGP router and a peer, the existence of a BGP peering is inferred if it is advertising prefixes.

When you first connect the appliance to the network, it usually acquires the topology in a matter of minutes; however, the process can take up to one hour for an Enhanced Interior Gateway Routing Protocol (EIGRP) network.

The appliance then maintains a real-time topological view of the entire network. You can view and manage the network from your desktop computer through the graphical user interface.

Key Route Analytics Management Software Components

RAMS includes the following components:

HP RAMS Route Recorder — An appliance that records routing data and stores it in a real-time database. The recorder can concurrently monitor most major routing protocols (OSPF, IS-IS, BGP, and EIGRP) across multiple domains and ASs from a single appliance.

HP RAMS Flow Collector — An appliance that collects traffic flow information exported from the routers and NetFlow recorders, and stores this information in a database. (RAMS Traffic only)

The Flow Collector is supported only on appliance models with two disk volumes . See *HP Route Analytics Management Software Appliance Setup Guide* for more information.

HP RAMS Flow Analyzer — An appliance that correlates traffic and routing data and then uses the combined data to produce reports. (RAMS Traffic only)

HP RAMS Modeling Engine — An appliance that creates a synthesized view of data collected across the network. The Modeling Engine presents this data in a graphical user interface accessible from your desktop, providing a single, cohesive view of network activity.

The size and distribution of the network and the number of supported concurrent users determines the needed number and type of appliances. In distributed networks, a single Modeling Engine can support multiple, geographically distributed Route Recorders. In RAMS Traffic deployments, you should install separate appliances for the Modeling Engine, Flow Analyzer, Flow Collectors, and Route Recorders.

With RAMS Traffic, you can monitor and record network events in different parts of the network with multiple Route Recorder units. The distributed Route Recorders collect routing data locally, from the area where they are installed, through generic route encapsulation (GRE) tunnels, or both. A centralized Modeling Engine retrieves the recorded data from each recorder. Users can then monitor network-wide routing from the Modeling Engine. Users can also archive network-wide data from a central location, and obtain reports from every Route Recorder in the configuration when they access the Modeling Engine.

When there are multiple Route Recorders in a distributed RAMS Traffic deployment, you can configure each appliance to record data per protocol or per multiple protocols, per area or per area within a protocol, or in any combination thereof. For a description of recorder configuration, see the "Configuration and Management" chapter in the *HP Route Analytics Management Software Administrator Guide*.

The next figures show how data flows through the network. RAMS is shown in Figure 1 and RAMS Traffic is shown in Figure 2.

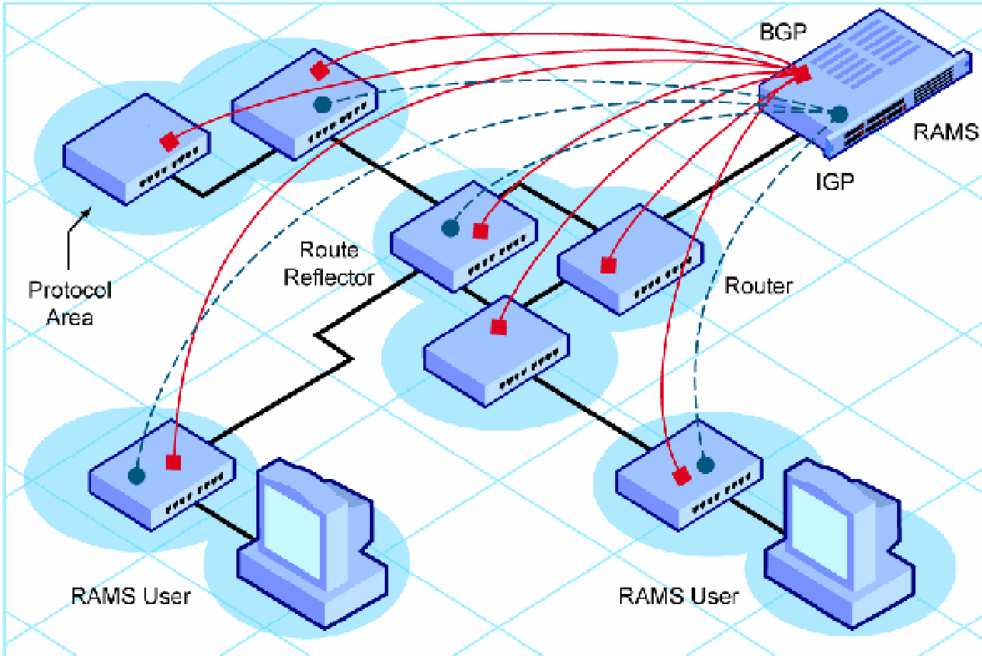

Figure 1 RAMS Data Flow

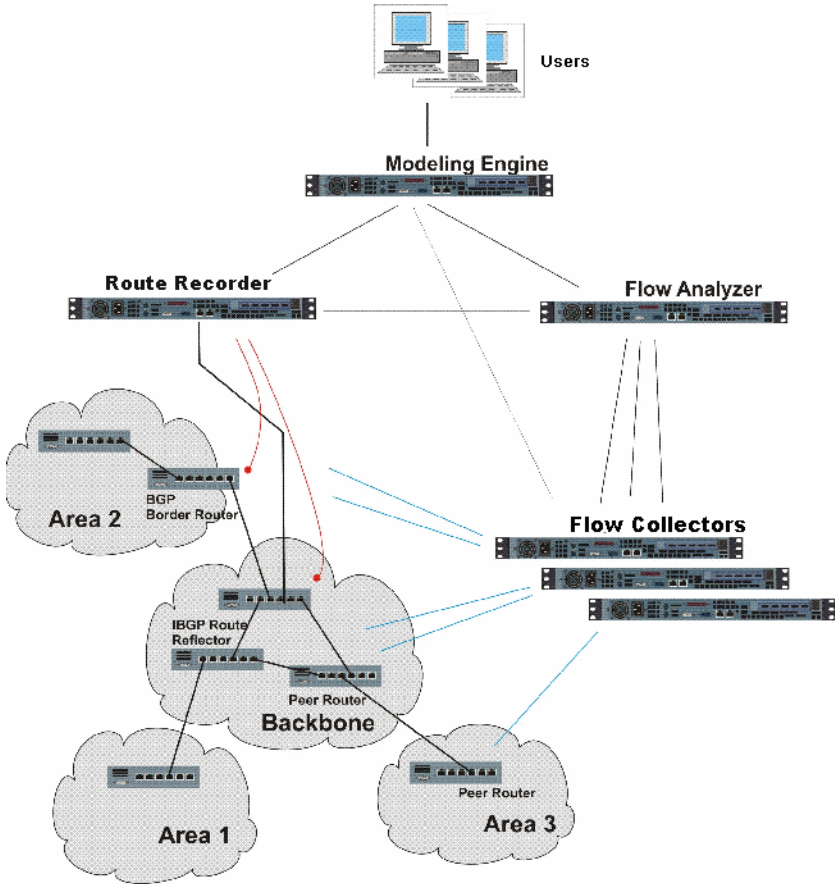

Figure 2 RAMS Traffic Data Flow

In a distributed environment where multiple appliances are installed on the network, you must designate the Modeling Engine as the master appliance during the configuration process (with the Master Capability license key). The Route Recorders in the deployment act as clients. (For RAMS Traffic, the Route Recorder, Flow Collectors, and Flow Analyzer are all client units.) From the master, you view and configure clients for recording. You also manage licenses for the entire configuration from the master.

In RAMS Traffic, Flow Collectors are located near the router where they collect traffic flow data. The recorders aggregate this data before providing it to the Flow Analyzer. The Flow Collector receives routing data from the Route Recorder and traffic data from the NetFlow exporters to aggregate this data.

The Flow Analyzer receives data from one or multiple Flow Collectors, and combines the data to create a network-wide report. The Modeling Engine queries the routing and traffic databases of each appliance to create a synthesized view of both route and flow across the network, then updates the topology map with this data whenever the routing topology changes, thereby providing an accurate, real-time view of how the network is directing traffic.

Accessing and Using Route Analytics Management Software

You can connect to the appliance using either of the following methods:

- A web browser for accessing the Administration web pages. Use the web interface to perform tasks such as database management, report creation, software updates, and recorder configuration.
- A Virtual Network Computing (VNC) or X Window System client for displaying the Route Analytics Management Software application. Network engineers and operators use the VNC or X client interface to view the routing topology map and analyze network activity.

Both types of viewers accommodate remote access, so you can view and manage one or more units from any desktop computer connected to the network, providing that it has a web browser and a VNC or X Window System client installed. See Chapter 2, "Viewers" for instructions on setting up client access.

Web Interface

After you log into the appliance through a web browser, the Home page opens. The top of Home page contains the following navigational links, which provide access to each area of the web interface:

- **Administration** — Connects you to the Administration pages, where you perform administrative tasks such as user management and software updates.

- **Recorder Configuration** — Connects you to the Recorder Configuration page. In a distributed system, you configure recorders and analyzers on the Recorder Configuration page of the master unit. On client units, the Recorder Configuration page is view-only and shows just that client's branch of the configuration tree.
- **Reports Portal** — Connects you to the reports pages of the recorder, where you can run reports detailing recorder activity for IGP and BGP protocols. In a deployment with multiple Route Recorders, you can use the Reports Portal of the centralized Modeling Engine to obtain network-wide reports from a single location.
- **Support** — Connects you to a page providing links to documentation in PDF format, as well as links to the Self Service Support site and software downloads.

You will also find a link to **Logout** of the web interface. To log back in, you must re-enter your user name and password.

Client Application Interface

After you launch the application in a VNC or X Window System viewer, you open a routing topology map, which is a real-time graphical representation of the network. There are three display modes for viewing and manipulating the topology map:

- **Monitoring mode** — In this mode, the topology is currently being recorded and updates to the routing database are shown on the topology map as they occur.
- **Planning mode** — In this mode, planning features are enabled for the topology map.
- **Analysis mode**— In this mode, only previously recorded information in the routing database is shown on the topology map.

See Opening a Routing Topology Map on page 42 for instructions on opening the maps. Monitoring mode is available only with databases that are configured for recording data.

In Planning mode and Analysis modes, you can focus on snapshots of network activity that is meaningful to your network planning and analysis. For example, you can view network data for the last hour, the entire month, or create a customized time range reflecting the state of the network from 11 a.m. to 2 p.m.

Topologies normally open in Monitoring mode, with the following exception: In RAMS Traffic, if you are opening up a topology including a traffic database, the topology automatically opens in Analysis mode with the selected time set to the latest available traffic data. Due to the inherent delay of NetFlow sampling, aggregation and buffering, traffic data is typically delayed by 20 minutes from real time. If you are only interested in routing data, you can open the topology in Monitoring mode by deselecting the traffic databases in the Open Topology dialog box.

To change modes, click the mode icon in the lower left corner of the window and select the desired mode.

In addition to the main topology map window, the following tools are available:

- **History Navigator** — Allows you to replay and analyze historical data. This tool is useful in investigating the cause of past events and helps network engineers plan for better performance in the future.
- **Planning Reports** — Allows you to view a table listing all edits made to the topology map in Planning mode. This tool also provides an analysis of how the edits theoretically affect network traffic.
- **Capacity Reports** — Allows you to view an estimate of future traffic demands based on past data. This allows you to plan network expansion to meet future demands. (RAMS Traffic only)
- **Traffic Reports** — Allows you to view reports based on traffic collected by Flow Collectors, then correlated and analyzed by the Flow Analyzer. (RAMS Traffic only)
- **Path Reports** — Allows you to generate reports to analyze network connectivity and optimize routing performance.
- **IGP and BGP Reports** — Allows you to view IGP and BGP protocol routing data collected from the Route Recorders. In a distributed deployment with multiple Route Recorders, connect to the centralized Modeling Engine to view consolidated reports containing IGP and BGP protocol data collected across the network.

Before proceeding with this document, you should make sure that the RAMS appliance is installed and networked as described in the *RAMS Appliance Setup Guide*.

2 Viewers

This chapter describes how to download and install viewers to access the Route Analytics Management Software client application.

Chapter contents:

Understanding Viewer Options

The X Window System or AT&T Lab Virtual VNC viewer is required to run the client application.

- The X Window System allows you to run an application on a remote computer located anywhere with access to the Internet, and display the application windows on your local computer. Accessing the system using the X Window System works best with high-speed, low-delay Internet connections.
- VNC makes allows you to display a desktop remotely over the Internet using a wide variety of operating systems. Accessing the system through VNC may give better performance than the X Window System over Internet connections with high delay and/or low bandwidth, such as Digital Subscriber Line (DSL) and dial-up connections.

Installing an X Window Server

To use the X Window System, your computer must run an X server to receive and display the output of the remote application. For privacy and security, you must use secure shell (SSH) to connect the session.

Using X Window System Software for MS Windows

The following X Window System products with SSH are available for Microsoft Windows:

- Xming
- Xwin-32
- Xmanager

The third-party X software included with Route Analytics Management Software comes with a 30-day evaluation license. To continue to use the software after this initial 30-day period, you must purchase a license from StarNet Communications Corporation.

The X-Win 32 evaluation from StarNet supports anti-aliased fonts, which allow the BGP root cause analysis and RIB visualizations to display correctly.

To download and install Xming for Windows, perform the following steps:

1. Open a web browser and go to **http://www.straightrunning.com/XmingNotes**
2. Choose **Xming** under Public Domain Releases.
3. Follow the installation steps, selecting **Full Installation**. For efficient operation, choose the options shown in Figure 3.

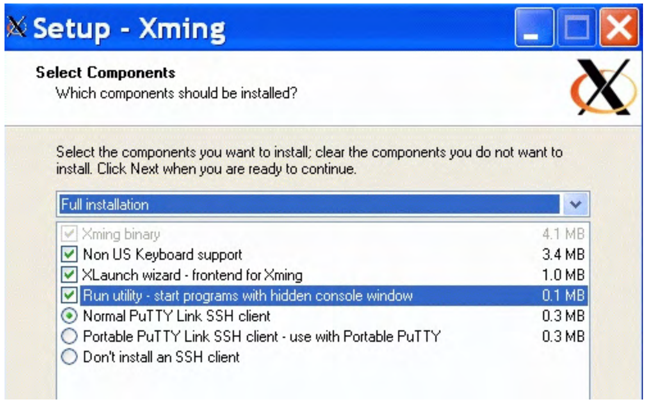

Figure 3 Xming Options

4 Include Normal PuTTY Link SSH client in the installation along with the other default selections

5 When installation is complete, run **XLaunch**.

For display settings, you can choose multiple windows, one window, full screen, or one window without title bar. The multiple window option best fits for most purposes. You can have multiple views of each subset of topology if available under a hierarchical model.

6 Choose a display option and click **Next**.

7 Choose **Start a program** and click **Next**.

8 Choose **Using PuTTY** in the Run Remote area. Enter the IP address in the Connect to computer field, the user name (admin or op) in the Login as user field, and the password (admin or op) in the Password field. Click **Next**.

9 You do not need to enter the parameter settings. Click **Next**.

10 To save the configuration for fast subsequent access, click **Save Configuratio**n. Choose a name for the configuration and click **Save** to save to your desktop or elsewhere.

11 Click **Finish**. You will be connected to the appliance.

To download and install X-Win32 for Windows, perform the following steps:

1 Open a web browser and log into the appliance.

2 Click **Support** on the top navigation bar to open the Support page.

3 Click **Link to StarNet Communications Corp. for X-Win32 Evaluation**.

StarNet's Download X-Win32 Evaluation page opens.

4 Fill out the form shown on-screen, and click **Send Email**.

After sending the form, in return you will receive a 30-day license key and download instructions for X-Win32.

To download and install Xmanager for Windows, perform the following steps:

1 Open a web browser and log into the appliance.

2 Click **Support** on the top navigation bar to open the Support page.

3 Click **Xmanager 30-Day Evaluation** to download an evaluation copy of X-Win32.

4 Select **Save this file to disk,** and then click **OK** to save the X-Win32 executable file to the specified local directory.

5 Open the downloaded .exe file. The Welcome screen appears.

6 Follow the on-screen instructions to install X-Win32.

To start X-Win32, perform the following steps:

1 Double-click the X-Win32 icon on your desktop.

2 Open the window from the X-Win32 folder.

3 Enter the connection details in the window that appears:

- Name: Enter a name for the session.
- Host: Enter either the hostname or IP address of the appliance.
- Protocol: Select SSH. To ensure privacy and security, only SSH connections are accepted.
- User name and Password: Enter your appliance user name and password.

The first time the SSH connection is initiated, you may see a security warning. Click **Yes** to save the host key and continue.

4 Click **Save** to save the connection details.

5 Click **Shortcut** to create a shortcut on the Windows desktop for easy repeat access.

6 Click **Run** to start an X session and open the application.

If you would like to view a demo of the Xmanager setup, select the **Xmanager Setup** (ShockWave Demo) link and follow the screens.

Using X Window Server for UNIX Platforms

The X Window System is included with Linux and Solaris platforms.

SSH is required to run the system through the X Window System.

To run the X Window System on Red Hat Linux, perform the following steps:

1 Start a graphical user interface such as XDE or Gnome on your desktop.

2 From the shell in a terminal window, open an SSH connection to the appliance. Enter the following command:

```
ssh X userid@unit
```

For example:

```
ssh X op@10.0.0.24
```

3 Enter your appliance user password when prompted.

The application opens on the desktop.

To run the X Window System on Solaris, perform the following steps:

1 Start a graphical user interface such as OpenWindows or Common Desktop Environment (CDE) on your desktop.

2 Open an SSH connection to the appliance in a terminal window (CDE) or shelltool (OpenWindows). Enter the following command at the shell prompt:

```
ssh -X userid@unit
```

For example:

```
ssh -X op@10.0.0.24
```

3 Enter your appliance user password when prompted.

The application opens on the desktop.

The X Window System may work on other platforms, but other platforms are not tested or fully supported. For more information on these platforms, go to http://www.packetdesign.com.

Installing VNC Viewer

To use VNC, you must install a VNC viewer (client) on your computer. The VNC viewer connects to the VNC server running on the appliance. Before starting the VNC viewer on the desktop, configure and start the VNC server as described in the Administration chapter of *HP Route Analytics Management Software Administrator Guide*

Downloading VNC

The following versions of the VNC viewer are available on the Support page:

- Windows 9x, NT, 2000, XP
- Linux (x86)
- Macintosh (OS9)
- Macintosh (OS X)
- Solaris (sparc)

Downloading and Installing VNC on Windows

To download and install NVC, perform the following steps:

1 Open a web browser and log into the appliance.

2 Click **Support** on the top navigation bar to open the Support page.

3 On the Support page, click the link for the appropriate version of the VNC viewer.

4 Select **Save this file to disk**, and then click **OK**.

This saves the VNC viewer to a local directory.

5 The downloaded VNC file is compressed. Before installing it, decompress it with an application such as WinZip.

6 Run the VNC viewer.exe file to install VNC.

To start VNC, perform the following steps:

1 Double-click the VNC icon to open the Connection Details dialog box.

2 If this is a first-time installation, click **Options**, adjust options as needed, and then click **OK**:

- Choose **Tight Encoding** to improve performance.
- Choose **Full-screen mode** to eliminate scroll bars on the VNC viewer and window frame. This prevents the taskbar and minimized icons on the desktop from being scrolled off-screen.

> When the VNC display is in full-screen mode, the Windows taskbar will not be visible. Press **Ctrl-Esc Esc** to make the Windows taskbar visible, then right-click the VNC icon to see the menu.

3 Enter the appliance IP Address or hostname in the VNC Server text box followed by **:1**.

Example: **192.168.1.5:1**

4 Click **OK** to start the VNC viewer.

> If a “Failed to connect to server” warning appears, the VNC server is not running or the system is in single operator mode and another operator is already accessing it. Contact the appliance administrator to resolve the problem.

5 Enter the VNC authentication password. For instructions on configuring the password, see the “Administration” chapter in the *HP Route Analytics Management Software Administrator Guide*).

6 Click **OK** to start the VNC viewer.

7 To save the optional settings for VNC, right-click the VNC icon in the Windows taskbar and select **Save connection info as**.

Downloading and Installing VNC on Linux

To download VNC, perform the following steps:

1 Click the link for the appropriate version of the VNC viewer.

The File Download window opens.

2 Select **Save to disk**, choose the location for the file, and then click **OK**.

To decompress VNC, perform the following steps:

1 Open the console and log in as root.

2 Change to the directory where the TightVNC rpm is saved.

3 Enter the following command:

rpm -U vnc-3.3.3r2+tight1.2.4-1.i386.rpm

When the installation completes, the shell prompt reappears.

4 To verify the installation, enter the following command:

rpm -qa | grep vnc

To start VNC, perform the following steps:

1 Enter the following command at the command line, where a.b.c.d is the appliance IP address or hostname:

vncviewer a.b.c.d:1

or

vncviewer -fullscreen a.b.c.d:1

The warning “vncviewer: ConnectToTcpAddr: connect: Connection refused: appears if you omit “:1” at the end of the IP address or if the VNC server is not running. In the latter case, contact the administrator to start the VNC Server from the Administration page.

2 At the password prompt, enter the VNC authentication password and click **OK**. For instructions on configuring the password, see the “Administration” chapter in the *HP Route Analytics Management Software Administrator Guide*).

This starts the VNC viewer.

> If the system is in Single Operator mode and another operator is already accessing it, the following message appears on the console: “vncviewer: VNC server closed connection.” Contact the appliance administrator to resolve the problem. If shared access to the VNC desktop is appropriate, ask the appliance administrator to change the setting to Multiple Operators and restart the VNC serer.

3 To exit the VNC viewer in full-screen mode, use the F8 key to bring up the menu and choose **Quit viewer**.

Installing SVG Plug-In

Adobe offers a free Scalable Vector Graphic (SVG) plug-in that you can download from the following URL: **http://www.adobe.com/svg/viewer/install/main.html**).

The Adobe plug-in is compatible with various browsers on Linux, Mac OS X and Microsoft Windows platforms.

Select **Install SVG Plug-In** and follow the steps provided onscreen to install the plug-in.

View System Information

Selecting this link displays the currently configured settings for your appliance.

➤ To view System Information, you must log in as the administrator.

Opening the Client Application

Use the procedures in this section to open the client application in MS Windows using X Window server software or VNC server. For instructions on starting the application in Linux or Solaris, see Using X Window Server for UNIX Platforms on page 33.

X Windows

Follow these steps only if you have already downloaded, installed, and initially run one of the X Windows options. For instructions on setting up those options, see Using X Window System Software for MS Windows on page 30.

To open the client application using Xming, perform the following steps:

1. Click the shortcut for your saved XLaunch configuration.
2. Enter your password, as prompt, and click **Finish**.
3. The client application opens.

To open the client application using Xwin-32, perform the following steps:

1. Click the Xwin-32 shortcut that you saved.
2. Enter your password, as prompt, and click **Finish**.
3. The client application opens.

To open the client application using X Manager, perform the following steps:

1. Click the X Manager shortcut that you saved.

2 Enter your password, as prompt, and click **Finish**.

3 The client application opens.

VNC

VNC viewer behavior depends upon on the VNC display that is specified.

For VNC display 1, the application is started when the VNC server is started on the web page. When the first connection is made to the VNC server, the application is already running. If the connection is ended, the session persists. When a new connection is made, the application will be in the state in which it was left at the end of the first session.

If the application is closed using the Quit menu command or by closing the main window, then the next time VNC is opened the desktop appears without the main window. In this case, perform the following steps:

1 Left click in the background of the VNC desktop or click **Start** from the taskbar at the bottom of the VNC desktop.

2 Click the product name. The main window opens.

When you start the VNC server, the application is automatically started on the VNC desktop and opens as soon as the VNC viewer connects to the server.

When the VNC viewer makes a connection to one of the VNC displays 2-10, a new instance of the VNC server is started and the application is automatically started on the VNC desktop. If the application is closed using the Quit menu command or by closing the main window, then the session ends. If a new session is connected, a new instance of the application is started.

The size of the VNC window depends upon how the individual VNC display (1-10) is configured. See the “Administration” chapter in the *HP Route Analytics Management Software Administrator Guide*.

3 The Routing Topology Map

This chapter describes how to use the routing topology map to monitor your network.

Chapter contents:

Working with Routing Topology Maps

The Routing Topology Map window provides an overall view of the currently-running network, including any tactical changes made during outage repairs that might not be reflected in network design documents.

Each stored database has a routing topology that you open to view router status and links, spot outages, identify routing failures, and uncover potential configuration errors that may result in service outages following maintenance activities. The routing topology map lets you view the routing events that led to failure and perform forensic analysis. Its accurate, vendor-independent view of the routing network can help you identify implementation or interoperability issues that are not easily isolated using other tools.

Chapter 2, "Viewers" explains how to start the application in the X Window System or in a VNC viewer. The next step is to select and open a routing topology.

Opening a Routing Topology Map

When a topology initially loads, the appliance uses a randomizing process to place the nodes on the routing topology map. The placement is not geographical. After you save the layout the first time, the database loads more quickly when it is reopened.

IPV6 menu items are displayed only if an IPv6 topology is loaded.

To open a routing topology map, perform the following steps:

1 Open the client application, as described in Chapter 2, "Viewers"

2 In the main window, select **Topology > Open Topology** to display the topology list.

Database names shown in green are configured for recording data, and database names shown in black are inactive.

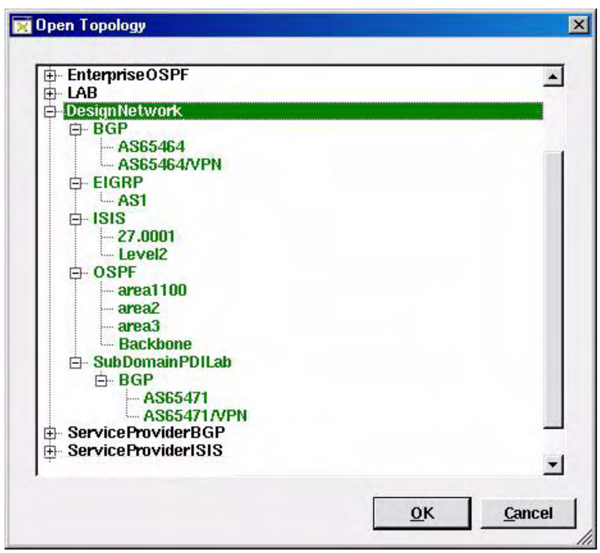

Figure 4 Opening a Topology

3 Select databases from the list. To identify a range, press **Shift** as you click names. To add or remove individual items, press **Ctrl** while you click the name. Selecting a group name selects all items within the group.

In RAMS Traffic, if you plan to perform BGP-specific analysis operations, such as the Root Cause Analysis, deselect any traffic databases in the Open Topology list. Traffic data is not relevant for BGP-specific analysis, and loading traffic information from the database can slow analysis.

4 Click **OK**.

After the database loads, the topology map opens in the main window. The database size determines the amount of time it takes to render the topology map.

In RAMS Traffic, any topology that contains traffic data automatically opens in Analysis mode. This is due to inherent delays in NetFlow aggregation and buffering. If you are interested only in routing data, open the topology in Monitoring mode by deselecting any topologies that contain traffic in the Open Topology dialog box.

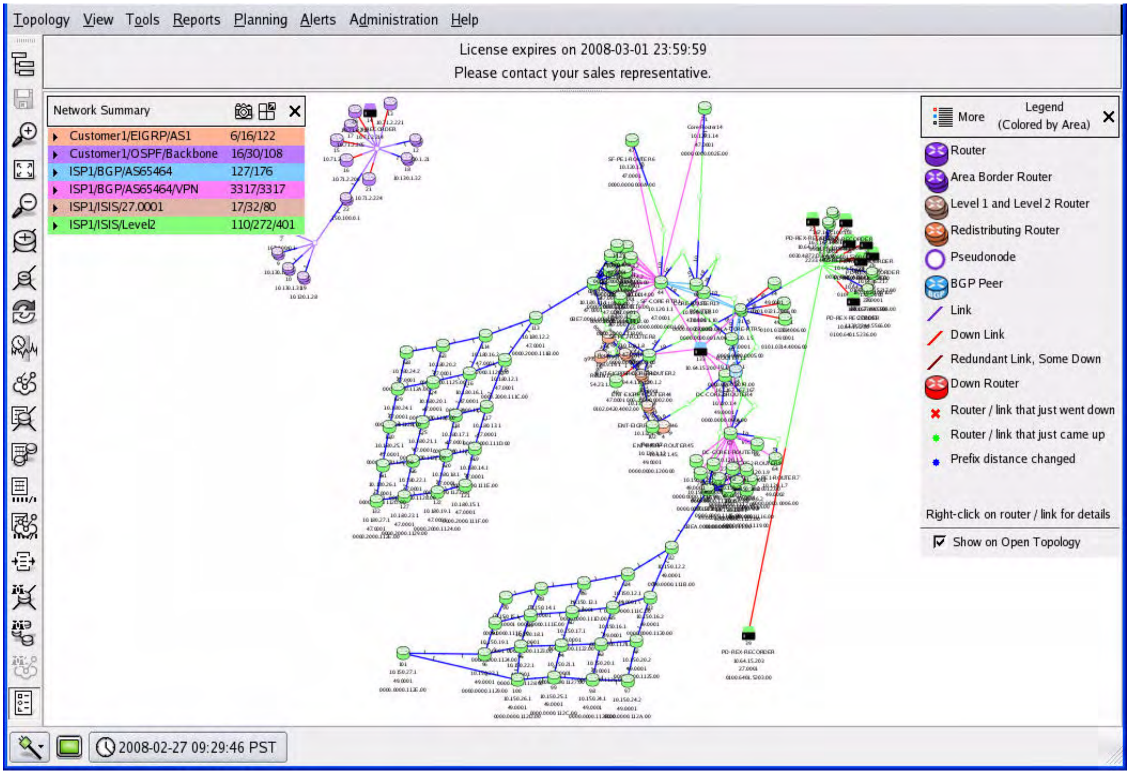

Figure 5 Routing Topology Map Main Window

Symbols and Colors

The routing topology map consists of node symbols connected by links. The symbol used for each node depends on its function. See Legend Panel on page 46 for the default shape and color associations. The shapes and colors are customizable.

The nodes and links in each routing area of the network are shown in different colors. You can change or turn off element coloring on the map by selecting **View > Color Modes** and then choosing one of the following items:

- Color by Area
- Color by Traffic Bitrate (RAMS Traffic only)

- Color by Traffic Utilization (RAMS Traffic only)
- Uncolor
- Color by Metric

When coloring is turned off (uncolor), nodes are black and links are grey. For additional information, see Colors on page 75.

Links and Peerings

IGP links and BGP peerings are shown as colored lines connecting nodes. Each link on the topology map is divided in half to represent the two directions of communication between a pair of nodes. The half adjacent to a node represents the outbound direction from that node. Each half can function separately up or down.

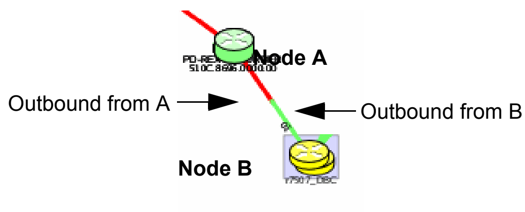

Figure 6 Link Colors

If an IGP has two or more adjacencies between a pair of nodes, which is usually due to the presence of multiple physical links, then the link is displayed on the map as two parallel lines. If some of the adjacencies are down and some are up, then one of the lines will be red and the other will not. If all of the adjacencies are down, then both lines will be red.

By contrast with IGP, peerings between BGP protocol routers and their clients are implied and may not follow physical paths. The system discovers and monitors BGP connections indirectly by receiving routes from BGP protocol routers for which the next hop attribute contains the address of another BGP router.

Because BGP peerings can be too dense on the map to be useful, the default setting is for the peerings not to be displayed. See Auto-Hide on page 79.

Legend Panel

The legend is displayed in the upper-right corner of the routing topology map (Figure 7). You can move the legend by placing the cursor in the top area and holding it down while you move. Selecting a symbol in the legend highlights the symbols on the topology map.

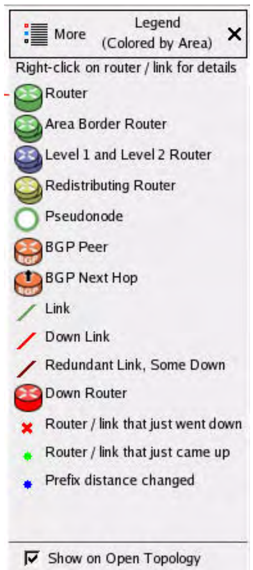

Figure 7 Topology Map Legend

The content in the legend is dynamic. You can change the colors as described in Colors on page 75. Use the options shown in Table 1 to control display of the legend area.

Table 1 Legend Appearance Options

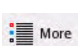	**More/Less** button	Toggles between less and more detailed symbol descriptions.

Table 1 Legend Appearance Options

	**Close** button	Closes the Legend. The next time you open a topology, the legend reappears unless you deselect the **Show on Open Topology** checkbox. You can open the legend panel at any time from the Help menu or by clicking the **Legend** button, as described in Main Window Toolbar on page 52.
Show on Open Topology checkbox		Determines whether the legend is displayed when you open a topology. The checkbox is enabled by default.

Network Summary Panel

The Network Summary panel (Figure 8) allows you to view current network counts and analyze the effects of network changes. By default, the panel is displayed when you open a topology map.

- If you do not want to display the panel when you open a topology map, choose **Administration > Options > Miscellaneous**, deselect the **Show Network Summary on Open Topology** check box, and click **Close**.
- If you close the Network Summary panel, reopen it by choosing **View > Show Network Summary**.

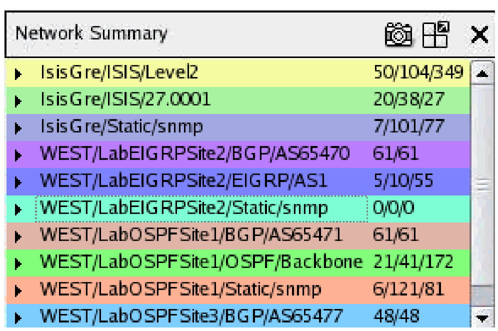

Figure 8 Network Summary Panel

Expand the detail information for a specific protocol by clicking on the protocol, as shown in Figure 9.

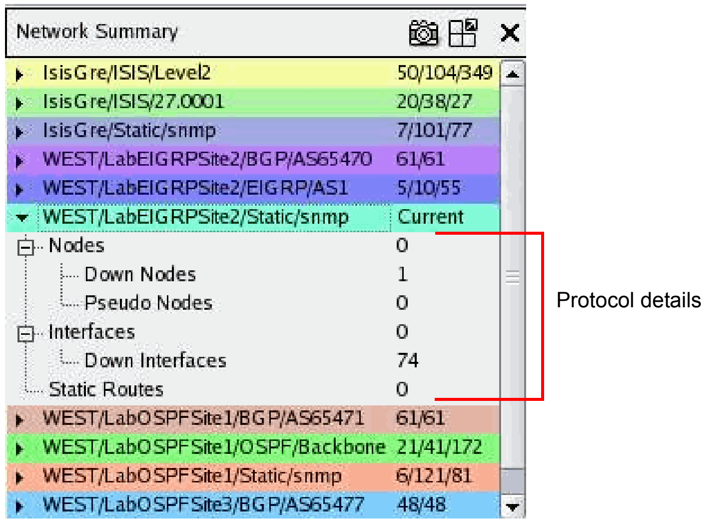

Figure 9 Network Summary Panel with Expanded Row

Table 2 lists the information in the Network Summary panel. The information depends on protocols.

Table 2 Network Summary Panel Information

Item	Description	Protocols
Total nodes	Total number of network nodes	IGP, EIGRP, IS-IS, OSPF
Down nodes	Number of nodes that are not operational	IGP, EIGRP, IS-IS, OSPF
Isolated nodes	Number of nodes with no links to other nodes	IGP, EIGRP, IS-IS, OSPF
Total links	Total number of network links	IGP, EIGRP, IS-IS, OSPF
Down links	Number of links that are not operational	IGP, EIGRP, IS-IS, OSPF
IPv4 Prefixes IPv6 Prefixes	Total number of network prefixes	IGP, EIGRP, IS-IS, OSPF (IS-IS and BGP only for IPv6)

Table 2 Network Summary Panel Information

Item	Description	Protocols
Down IPv4 Prefixes Down IPv6 Prefixes	Number of prefixes that are not operational	IGP, EIGRP, IS-IS, OSPF (IS-IS and BGP only for IPv6)
Unique IPv4 Prefixes	Number of distinct IPv4 prefixes	IGP, EIGRP, IS-IS, OSPF (IS-IS and BGP only for IPv6)
Down Unique IPv4 Prefixes	Number of distinct IPv4 prefixes that are not operational	IGP, EIGRP, IS-IS, OSPF (IS-IS and BGP only for IPv6)
Active prefixes	Number of prefixes that have at least one operational route	BGP, VPN
Active routes	Number of routes that are operational	BGP, VPN
Baseline up routes	Number of operational routes in the baseline	BGP, VPN
Baseline down routes	Number of non-operational routes in the baseline	BGP, VPN
Next hops	Number of unique BGP next hops for all routes	BGP, VPN
MP Next hops	Number of unique BGP multi-protocol (MP) next hops for all routes	BGP, VPN
Neighbor ASs	Number of unique neighbor ASs	BGP, VPN
Pseudo nodes	Number of nodes that are pseudonodes	BGP, IS-IS, OSPF
Overloaded	Number of nodes indicated as overloaded	BGP, IS-IS

Table 2 Network Summary Panel Information

Item	Description	Protocols
IPv6 Capable Nodes	Number of IPv6 capable nodes in the topology	IS-IS
Total OSI prefixes	Number of nodes with Open Systems Interconnection (OSI) prefixes	BGP, IS-IS
Down OSI prefixes	Number of OSI prefixes that are not operational	BGP, IS-IS

Table 3 shows the buttons that are displayed at the top of the Network Summary panel.

Table 3 Network Summary Buttons

	**Snapshot Statistics**	Saves the current Network Summary counts and displays the statistics in a new column.
	**Tear Off**	Releases the Network Summary panel from the main window so that you can move it around your desktop.
	**Put Back**	Reattaches the Network Summary panel to its original location in the main window.

Main Window Status Bar

The status bar at the bottom of the main topology map window indicates the current mode, recorder status, and the date and time that you are viewing on the map. When applicable, a status message is shown on the right of the date and time.

To change the date and time, click the field. **OK** and **Cancel** buttons are displayed. Make changes and Click **OK**.

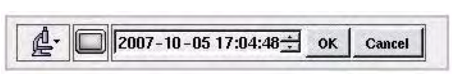

Figure 10 Main Window Status Bar

In Analysis mode, you can click the date/time area to display the playback controls. Choose a date and time and a step size in sections for the playback, and click **OK**. Then click the right-facing arrow to begin the playback.

Figure 11 Main Window Status Bar

In RAMS Traffic, the time in the status bar is based on the time of the traffic data. For current data, there is typically a delay, the length of which depends on network complexity.

Change modes by clicking the icon for the current mode and selecting the desired mode. Table 4 shows the icons.

Table 4 Mode Icons

	**Monitoring mode**	The topology is currently being recorded and updates to the routing database are shown on the topology map as they occur.
	**Analysis mode**	Only previously recorded information in the routing database is shown on the topology map.
	**Planning mode**	Planning features are enabled for the topology map. This mode activates another set of buttons in the main window.

An indicator to the right of the mode button  provides status based on color (Table 5). Click the status indicator to obtain additional information on the recorder and peers.

Table 5 Color Status

Green	The system is running and recording to the database, and adjacencies between the system and peer routers in all areas are up.

Table 5 Color Status (cont'd)

Blue	At least one of the following applies: •Data is being recorded, but EIGRP topology exploration is in progress, so changes in the topology will not be shown until completion. The time on the recorder is not synchronized with the time in the appliance. •Collector operations are in process. •The database has recorded events that are in the future relative to the current time of day for the GUI.
Yellow	Data is being recorded, but adjacencies to peer routers in some areas are down.
Red	No data is being recorded because the recorder is not running.
Gray	This is a historical database to which no new data is being recorded.
Purple	There is a database replication-related error.

Main Window Toolbar

The toolbar on the left side of the main routing topology map window contains buttons that control the map display and provide shortcuts to menu options. You can move the toolbar to the right, top or bottom of the window by dragging the dimpled strip at the top of the toolbar. Table 6 shows the buttons on the toolbar.

Table 6 Main Window Toolbar Buttons

	**Show / Hide Topology Hierarchy**	Insert a pane that shows a tree view of the routing databases on the map. See "Understanding the Topology Hierarchy" on page 3-20.
	**Save Layout**	Save the current map layout (disabled if no changes). If this is the first time you have saved the layout, the system prompts you to name the layout and specify it as your default layout. You can view saved layouts from the Reports Portal in the web application. See the *HP Route Analytics Management Software Administrator Guide*.

Table 6 Main Window Toolbar Buttons (cont'd)

	Zoom In / Zoom Out	Increase or decrease the size of images and text shown in the topology map. When zooming, you can use the scroll wheel on your mouse to scroll up and down within the zoomed topology map.
	Reset View to 1:1	Restore the topology map to the default viewing scale.
	Node Size Up / Node Size Down	Increase or decrease node size and labels.
	Relayout	Generate a new randomized layout of the nodes.
	History Navigator	Open the History Navigator window for the current topology. See Chapter 4, "The History Navigator"
	Routing Reports	Open the VPN Routing Reports window. See Viewing VPN Routing Reports on page 364.
	List Routers	Open the List of All Routers report to generate a list of all routers in the current topology map. See Router List on page 108.
	List Links	Open the List of All Links report to generate a list of all links in the current topology map. See Links List on page 109.
	List IPv4 Prefixes	Open the List of All Prefixes report to generate a list of all IPv4 prefixes in the current topology map. See Prefix List on page 112.
	List IPv6 Prefixes	Open the List of All Prefixes report to generate a list of all IPv6 prefixes in the current topology map. See Prefix List on page 112.
	List VPN Prefixes	Opens the List VPN Prefixes report. See Finding a VPN Route By Prefix on page 137.
	RIB Browser	Open the Routing Information Base (RIB) browser. See RIB Browser on page 180.
	Find Router	Open the Find Router or LAN node search to search for routers by IP address (IPv4 or IPv6) or router name. See Router List on page 108.

Table 6 Main Window Toolbar Buttons (cont'd)

	List/Find Paths	Open the List or Find Paths window to highlight a path between a router and an Internet prefix or domain name. See Finding a Route By Prefix on page 136.
	List/Find VPN Paths	Open a window to find VPN paths in the topology. See Finding a VPN Route By Prefix on page 137.
	Show Legend	Open or close the legend. See Legend Panel on page 46.

Duplicate functionality is available with some buttons and their associated menu items. The descriptions in this chapter uses the button options, when available.

When you zoom in on the topology map, an overview area opens in the lower right corner of the window to show the zoom area within the context of the full map (Figure 12). If you notice that interaction with the map becomes slower when this area is open, click the arrow in the lower right corner to hide the area. Click the modified icon to reopen the area.

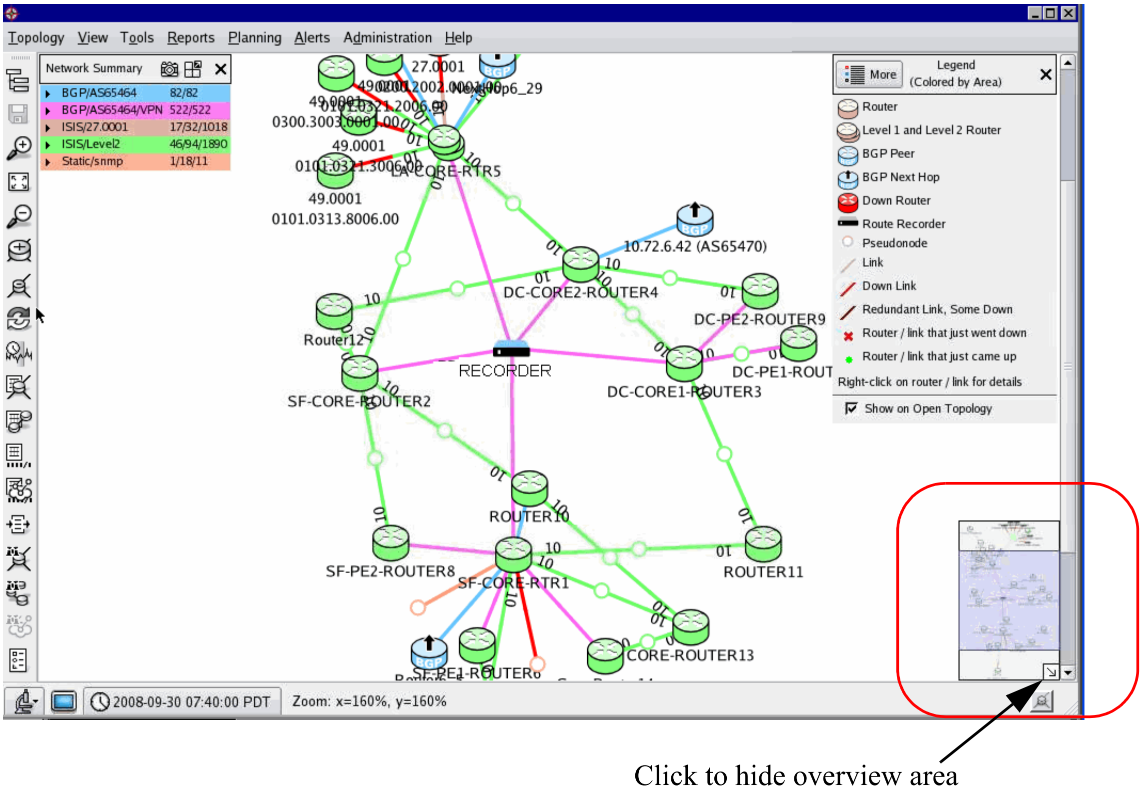

Figure 12 Routing Topology Map Overview Area Shown and Hidden

In Planning mode, a second toolbar is displayed on the right side of the routing topology map window. See The Planning Toolbar on page 252 for information on the toolbar icons.

Keyboard Shortcuts

Table 7 lists keyboard shortcuts for frequently-used menu selections.

Table 7 Keyboard Shortcuts

Shortcut	**Description**
Ctrl-F	Find router
Ctrl-O	Open topology
Ctrl-Q	Quit
Ctrl-S	Save layout
Ctrl-W	Close topology

Main Window Menus

Use the main menu bar to open routing topology maps and monitor and analyze routing data. This section describes the options available for each menu.

Some menu items are for specific protocol families and are shown only if your appliance is licensed for that protocol.

Topology

Table 8 describes the items on the Topology menu.

Table 8 Topology Menu Items

Item	Description
Open Topology	Open a routing topology database to display the topology map.
Close Topology	Close a routing topology database.
Analysis Mode Planning Mode Monitoring Mode	List the available modes, except the mode that you are currently in. **Note:** When you open a static-only database, Planning mode is not available.
Go To Time	View the topology map as it appeared at a particular time in history. (Analysis mode only)
Go To Latest Time	Activate the OK and Cancel buttons for the time field at the bottom of the window. This allows you to change the current time.
Load Layout	Load a saved layout to reposition the nodes. If needed, you can save the same layout multiple times under different names.
Reload Layout	Restore the node positions according to the previously loaded layout (disabled if no layout is loaded or no nodes have been moved).
Relayout	Create a new randomized placement of the nodes on the topology map. This can help you find a preferred orientation to name and save.
Delete Layout	Delete a named layout.
Save Layout	Save the current map layout (disabled if no changes). If this is the first time you save the layout, the system prompts you to name the layout and allows you to set it as the default.

Table 8 Topology Menu Items (cont'd)

Item	Description
Save Layout As	Open the Save Layout dialog box to save the current layout under a new name. You can also set the new layout as the default.
Select Layout Background	Open the Select Layout Background dialog box. Imported image files are shown in a list of available files.
Quit	Close the client application. The system application continues to run.

View

Table 9 describes the items on the View menu.

Table 9 View Menu Items

Item	Description
Show/Hide Topology Hierarchy	Show or hide a tree view of the routing databases on the left side of the routing topology map window. See Understanding the Topology Hierarchy on page 69.
Show/Hide Network Summary	Show or hide current counts of network elements. See Network Summary Panel on page 47.
Show/Hide Toolbar	Show or hide the toolbar on the left edge of the main window.
Color Modes	Modify color characteristics on the routing topology map. The color modes specifying bitrate and utilization are disabled in Monitoring mode. See Symbols and Colors on page 44.
Node Labels	Choose node label options. See Node/Link Labels on page 74.
Link Labels	Choose link label options. See Node/Link Labels on page 74.
Hide Nodes	Hide the nodes that match a specified pattern from view on the routing topology map (see Hiding Nodes on page 87).
Hide Leaf Nodes	Remove nodes that are on the edges of the map and have only a single link. Each time you select **Hide Leaf Nodes**, additional nodes are removed. See Hiding Nodes on page 87 for additional information.

Table 9 View Menu Items (cont'd)

Item	Description
Show / Hide Failed Nodes	Show or hide failed nodes.
Show / Hide Unconnected Nodes	Show or hide unconnected nodes.
Unhide All Nodes	Show any hidden (trimmed) nodes.
Highlight / Unhighlight All Paths	Turn path highlights on or off.
Zoom In	Increase image and text size.
Rest View to 1:1	Restore the topology map to the default viewing scale.
Zoom Out	Reduce image and text size.
Node Size Up	Increase node and label size. Node placement is saved when you save a layout.
Node Size Down	Decrease node and label size. Node placement is saved when you save a layout.
Node Size Reset	Reset nodes and accompanying text to the default size. This function does not affect zoom level.
Group Edit Mode	Select to enable or disable this mode. Group Edit Mode is also enabled automatically when you create a group on the routing topology map. For more information, see Creating Groups on the Routing Topology Map on page 115.

Tools

Table 10 describes the items on the Tools menu.

Table 10 Tools Menu Items

Item	Description
Find Router	Search for routers by IP address or router name. See Router List on page 108.
List/Find Paths	Highlight a path between a router and an Internet prefix or domain name. See Finding a Route By Prefix on page 136.
List/Find VPN Paths	Find VPN paths in the topology. See Finding a VPN Route By Prefix on page 137.
Highlight by Exit Router	Find the set of exit routers toward a specified prefix, Network Service Access Point (NSAP) Address, IP address or domain name, and color code all routers to indicate which exit router each will use. The border routers that act as exit routers to the specified destination flash between the coded color and yellow. See Viewing Exit Routers on page 90. To highlight exit routers, the listed network must include a route to the desired destination.
Prefix Diagnostics	Configure and view reports to help troubleshoot prefix issues. See Prefix Diagnostics Reports on page 98.
List Routers	Generate a list of all routers in the current routing topology map. See Router List on page 108.
List Links	Generate a list of all links in the current routing topology map. See Links List on page 109.
List Interfaces	Generate a list of interfaces discovered by static protocol. See Interfaces List on page 111.
List Prefixes List IPv4 Prefixes List IPv6 Prefixes	Generate a list of all IPv4 or IPv6 prefixes in the current topology map. See Prefix List on page 112. The type of prefix is not specified if the topology contains only IPv4 prefixes.

Table 10 Tools Menu Items (cont'd)

Item	Description
List VPN Prefixes	Generate a list of all of prefixes that are part of a VPN. See VPN Prefixes List on page 114.
List Flows	Highlight a specified flow. (RAMS Traffic only)
Topology Diagnostics	Show diagnostic information (for EIGRP topologies only). See Diagnosing EIGRP Topology Errors on page 141.

Reports

Table 11 describes the items on the Reports menu.

Most report tables support column sorting. You can resort data by clicking any column heading in the report. Click again to change the sort order (descending/ ascending).

Table 11 Reports Menu Items

Item	Description
Routing Reports	Open the VPN Routing Reports window. See "Viewing VPN Routing Reports" on page 7-8.
Traffic Reports	Access traffic reports and import data. See Chapter 9, "Traffic Flows and Reports," for more information. (RAMS Traffic only)
Path Reports	View computation paths between pairs of routers and generate reports for analysis. See Chapter 10, "Path and Routing Analysis Reports," for more information.
Routing Analysis Reports	Open the Routing Analysis Reports window (Analysis mode only). See Using Routing Analysis Reports on page 471.
History Navigator	Open the History Navigator window to analyze past routing data. See Chapter 4, "The History Navigator," for more information.
Route Cause Analysis	Open the Root Cause Analysis window. See Root Cause Analysis on page 167.

Table 11 Reports Menu Items (cont'd)

Item	Description
RIB Visualization	Open the RIB Visualization window. See RIB Visualization on page 175.
RIB Browser	Open the RIB browser to view link and route information. See RIB Browser on page 180.
Flow Record Browser	Display aggregated flow information. See Flow Record Browser on page 199. (RAMS Traffic only)
Events Monitor	Open the events list. See Understanding the Events List on page 200.
Syslog Messages	Open the list of syslog messages.
Correlate Time Series	Display a graph of external time series data in correlation with the routing history. See Correlating Time Series Data on page 219.

Planning

Table 12 describes the items on the Planning menu. For additional information on Planning menu items, see Chapter 10, "Path and Routing Analysis Reports"

Table 12 Planning Menu Items

Item	Description
Add Router	Place a new node on the topology map.
Add Peering	Create a peering relationship between two nodes on the topology map.
Add IPv4 Prefix	Apply an IPv4 prefix to a router on the topology map. For BGP routers, you can add prefixes manually or by selecting a filtering method for the prefix.
Add IPv6 Prefix	Apply an IPv6 prefix to a router on the topology map. For BGP routers, you can add prefixes manually or by selecting a filtering method for the prefix.

Table 12 Planning Menu Items (cont'd)

Item	Description
Add VRF	Add Virtual Routing and Forwarding (VRF), which allows multiple instances of a routing table to co-exist within the same router at the same time.
Add Traffic Flow	Create on the routing topology map. (RAMS Traffic only)
Add VPN Customer	Add a full mesh, or hub and spoke VPN customer to the network, including VRF routes. and traffic flows. (RAMS Traffic only)
Edit BGP Prefix	Change or remove IPv4 prefix attributes on one or more nodes.
Edit BGP IPv6 Prefix	Change or remove IPv6 prefix attributes on one or more nodes.
Edit Traffic Flows	Make changes to IPv4 flows or VPN flows. (RAMS Traffic only)
Edit Router	Change the overload bit for an IS-IS node.
Edit VRF	Modify the Virtual Routing and Forwarding router tables.
Down Router	Change the state of a node from up to down, simulating what would happen if the selected router should fail.
Down Peering	Change the state of a peer relationship from up to down, simulating what would happen if the selected peering fails. This action brings down all the peerings in the table or only selected relationships.
Down IPv4 Prefix	Change the state of one or more IPv4 prefixes from up to down on a router, simulating what would happen if the selected prefixes fail.
Down IPv6 Prefix	Change the state of one or more IPv6 prefixes from up to down on a router, simulating what would happen if the selected prefixes fail.
Down VRF	Change the state of one or more customer sites from up to down on a router, simulating what would happen if the selected custom site fails.
Analyze Edits	Updates all traffic and routing edits simultaneously.
Reports	Open the Planning Reports window. See Working with Planning Reports on page 285.

Table 12 Planning Menu Items (cont'd)

Item	Description
Undo All Edits	Cancels edits without opening the Planning reports or Show Edits table. See Show Edits on page 317.
Capacity Planning	Open the Capacity Planning window. (Analysis mode only) See Working with Capacity Planning Tools on page 304. (RAMS Traffic only)
Show Edits	Display recent edits to the topology map. See Show Edits on page 317 (RAMS only).

Alerts

Table 13 describes the items on the Alerts menu. For additional information on Alerts menu items, see Chapter 12, "Alerts"

Table 13 Alerts Menu Items

Item	Description
View Alert Status	View the current state of all defined alerts.
Configure Alerts	Set up alerts based on alarm and inform criteria.
Dispatch Specifications	Determine the notification method to distribute information about network and traffic status.
Suppression Specifications	Determine the minimum level of activity that is considered typical or recurring and does not require notification.

Administration

Table 14 describes the items on the Administration menu.

Table 14 Administration Menu Items

Item	Description
Assign Names	
Routers	Customize router names and map them to the routers IP address. See Assigning Router Names on page 104.
IPv4 Prefixes	Assign names to prefixes to more easily identify them in the reachability reports. See Assigning IPv4 or IPv6 Prefix Names on page 107.
IPv6 Prefixes	Assign names to prefixes to more easily identify them in the reachability reports. See Assigning IPv4 or IPv6 Prefix Names on page 107.
ASs	Customize AS names. See You can highlight invisible links in yellow on the topology map by clicking Highlight. on page 148.
Groups	
Routers	Mange router groups. See Creating Groups Using the Menu on page 122.
Links	Manage link groups. Link groups are used to generate a watchlist on an alert; for example, for adjacency state change. See Creating Groups Using the Menu on page 122”.
IPv4 Prefixes	Manage IPv4 prefix groups. See Creating Groups Using the Menu on page 122.
IPv6 Prefixes	Manage IPv6 prefix groups. See Creating Groups Using the Menu on page 122.
Paths	Manage source-destination paths and generate an alert if there is a change. See Creating Groups Using the Menu on page 122.
VPN Customers	Manage VPN customer groups. See Creating Groups Using the Menu on page 122.

Table 14 Administration Menu Items (cont'd)

Item	Description
VPN Route Targets	Manage VPN route target (RT) groups. See Creating Groups Using the Menu on page 122.
VPN Prefixes	Manage VPN prefix groups. See Creating Groups Using the Menu on page 122.
VPN	
Router Location Names	View and add router location names for XML customer reports. See Generating VPN Customer Traffic Reports on page 361.
VPN Customers	Import and configure customer configuration mappings. See Creating Customer and RT Associations on page 357.
Traffic	
Set Interface Capacities	Set the interface capacities (data rate) that the system uses to compute utilization percentages. See Setting Interface Capacities on page 402.
Traffic Change Reports	Configure settings for traffic reports. See Top Changes Report on page 407.
MPLS WAN	
Unmatched Serial Interfaces	Add or modify PE addresses for unmatched MPLS WAN interfaces. See Configuring Unmatched Interfaces on page 489.
Reachability Ranges	Configure ranges for the MPLS WAN reachability reports. See Modifying Reachability Ranges on page 491.
VPN Connections	Group customer edge (CE) and provider edge (PE) connections into VPNs for the MPLS WAN feature. See Setting Up the VPN Connection Configuration on page 491.
Expected Prefixes	Modify the list of expected prefixes for the MPLS WAN feature. See Identifying Expected Prefixes on page 494.
Static VPN Connections	Set up static VPN connections for the MPLS WAN feature. See Setting Up the VPN Connection Configuration on page 491.

Table 14 Administration Menu Items (cont'd)

Item	Description
Other Menu Options	
Combine Routers	Manually correct problems that occur with consolidation of routers on the routing topology map. See Combining Routers on page 90.
Saved Filters	View and modify custom filters. See Creating Custom Filters on page 101.
Assign BGP ASs to Routers	Assign routers manually to a BGP AS, usually for a BGP confederation. See Assign and Verify BGP AS Assignments to Routers on page 151.
Options	Set preferences. See "Applying Configuration Options" on page 71.

Help

Table 15 describes the items on the Help menu.

Table 15 Help Menu Items

Item	Description
User's Guide	Display the PDF version of the User's Guide in a new window.
Developer's Guide	Display the PDF version of the Developer's Guide in a new window.
Administrator's Guide	Display the PDF version of the Administrator's Guide in a new window.
Show / Hide Legend	Open or close the legend that defines each symbol on the routing topology map.
About ...	Display software version information.

Topology Map Layouts and Background Images

You can customize the appearance of the routing topology map in the following ways:

- Save a topology map under multiple names with different layouts. Changes, such as resizing or hiding/unhiding nodes in a layout, are preserved when you save the layout.
- Customize a layout by changing node placement, topology symbols, and colors. The changes you make are specific to your login name and are view-only for other users. If you load and modify a layout created by another user, your changes are saved under your login name and do not change the other user's layout.
- Apply background images to topology maps. For example, you can use a map that shows the geographic location of network routers and then arrange nodes based on physical or logical groupings, such as per building or per lab. You can import background images in BMP, JPG, PNG, SVG, or XPM format using the Layout Backgrounds window. See the "Configuration and Management" chapter in the *HP Route Analytics Management Software Administrator Guide*. Image files are stored in a database and are accessible from any device on the network.

To apply a background image to the routing topology map, perform the following steps:

1 Choose **Topology > Select Layout Background**.

A list of available image files is displayed (Figure 13).

Figure 13 Available Background Images

2 Choose an image and click **OK**.

The image is now included as the background for the routing topology map.

To remove a background image from the topology map, perform the following steps:

1 Select **Topology > Select Layout Background**.

2 Choose **<blank>** and click **OK**.

Understanding the Topology Hierarchy

The topology hierarchy shows the routing databases in a hierarchical tree view and allows you to work with a subset of the topologies in the map. This is useful for large networks that contain numerous IGP areas and BGP ASs.

Next to each branch is a status light. The color of each light indicates the state of the individual recorder. You can hover the cursor over each light for a status message related to the recorder state.

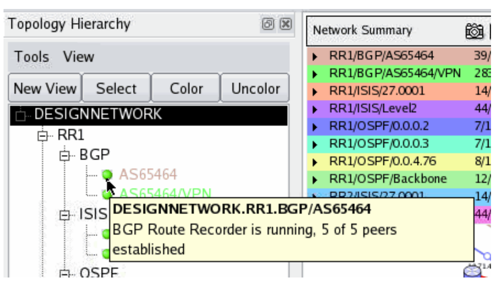

Figure 14 Status Message

The Tools and View menus within the topology hierarchy pane contain a subset of the items on the similarly named menus on the main window, but they operate only on the areas selected in the tree structure. For example, use the topology hierarchy Tools menu to list the routers in just one area.

The system also lets you view each individual area or AS in a separate window. Click **New View** to open a new topology map window that contains only the selected areas.

To display or hide the topology hierarchy, perform the following steps:

1 Choose **View > Show Topology Hierarchy**.

The hierarchy is displayed on the left side of the routing topology map.

2 Click the docking icon to dock or undock the hierarchy window.

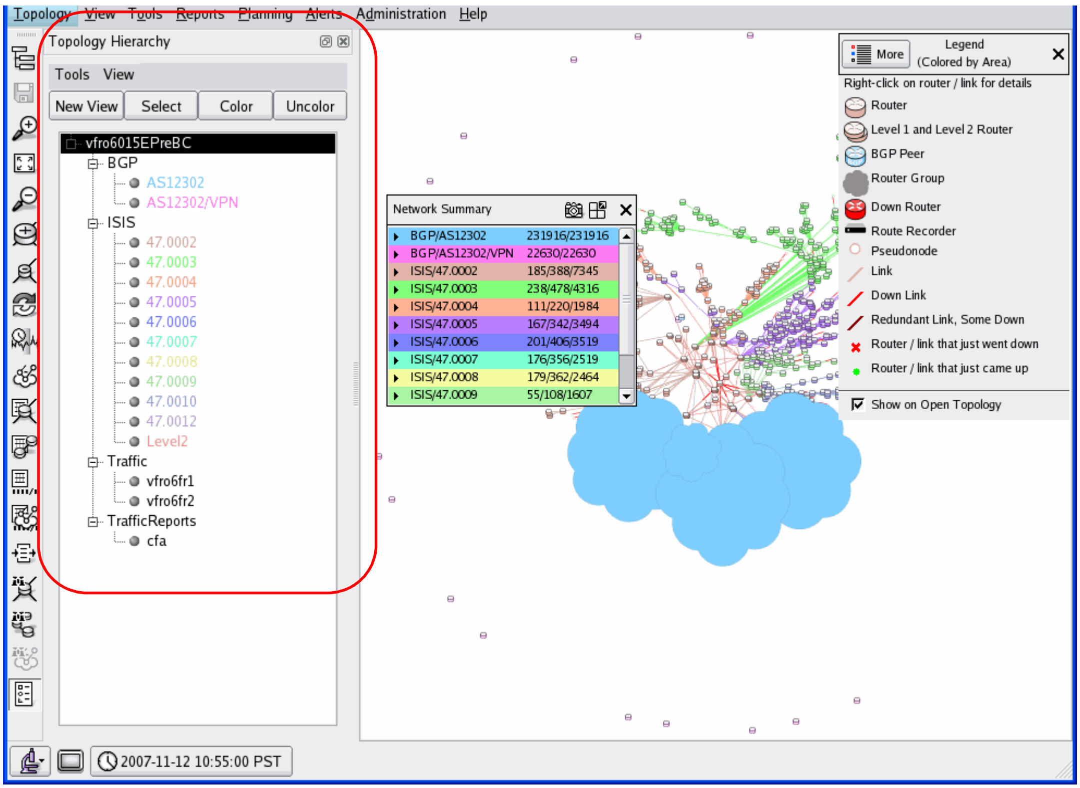

Figure 15 Topology Hierarchy

To view an individual area or AS, perform the following steps:

1 Click **Show Topology Hierarchy**.

2 Choose the desired areas in the tree structure.

3 Click **Select** in the topology hierarchy, and then click **Zoom In** or **Zoom Out** to expand the selected area of the topology map to fill the window. You can also click **New View** in the topology hierarchy pane to open a new window containing just the selected area.

Viewing Network Anomalies

You can use the routing topology map to see possible points of failure, failed links and routers, and other anomalies. For example, when a link goes down, that link turns red.

To identify network anomalies at a glance, perform the following steps:

1 Look for a link that is shown in red, meaning that the link is down. When a link is represented by two parallel lines where one is red and the other is not, it means that the link represents multiple parallel physical links, only some of which are down.

2 Look for parts of the network that route through a single router or LAN.

Applying Configuration Options

Configuration options are available to control the look and feel and behavior of the routing topology map.

To apply configuration options, perform the following steps:

1 Choose **Administration > Options** to open the configuration options window.

Configuration option categories

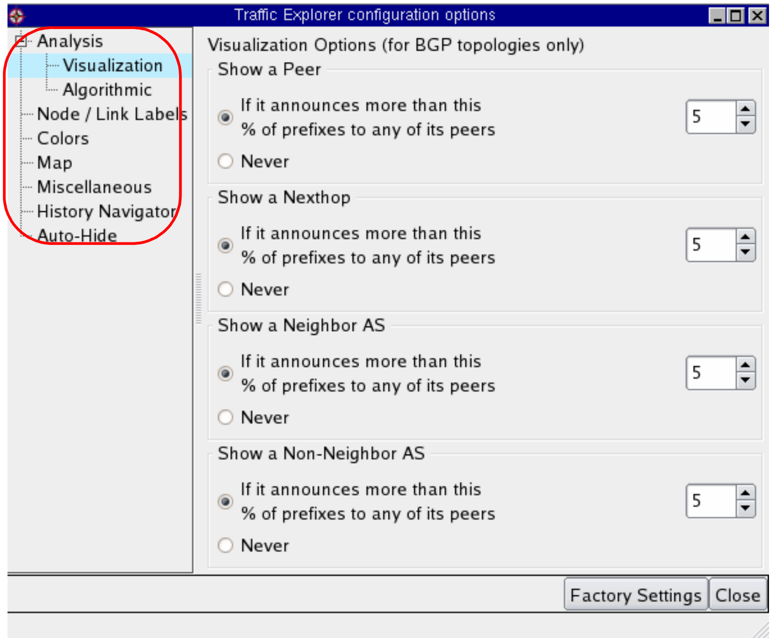

Figure 16 Configuration Options Window

2 Choose any of the categories from the left menu and make changes as described in the following sections:

- Analysis—see page 73
- Node/Link Labels—see page 74
- Colors—see page 75
- Map—see page 76
- History Navigator—see page 78
- Auto-Hide—see page 79

3 To save changes, click **Close**. To restore all options to their default values, click **Factory Settings**.

Analysis

The visualization and algorithmic analysis options control thresholds and level of detail displayed on the map.

Visualization Options

You can customize the level of detail that appears in visualizations and animations of the BGP RIB, as described in RIB Visualization on page 175. Visualization analysis controls whether a network entity appears in a RIB visualization window or a root cause analysis animation window. For each option, you can choose to always include it, include it if it announces more than a specified percentage of prefixes to any of its peers, or never include it. The Always option is disabled if choosing it could create a visualization that is too big or crowded to read. The following options are supported:

- **Show a Peer—**Include a peer if it announces 5 percent (default) or more of the total number of prefixes.
- **Show a Nexthop—**Include a nexthop if it announces 5 percent (default) or more of the total number of prefixes.
- **Show a Neighbor AS—**Includes the neighbor AS if it announces 5 percent (default) or more of the total number of prefixes.
- **Show a Non-Neighbor AS—**Includes the non-neighbor AS if it announces 5 percent (default) or more of the total number of prefixes.

Algorithmic Analysis Options

You can set the thresholds that are used in root cause analysis. In each case, a higher value decreases the level of detail and a lower value increases it. See Root Cause Analysis on page 167.

The following algorithmic analysis control thresholds are used in root cause analysis:

- Show if more than this percentage of prefixes on an edge is flapping. Default is 10 percent.
- Show if more than this percentage of total prefixes shifted from an edge. Default is 1 percent.

Node/Link Labels

Node labels determine the node details that are displayed on the routing topology map. The options depend on the protocol family. Each option is displayed only if it applies to currently licensed protocols. For example, router names and system IDs do not appear for an OSPF network.

Link labels determine information that is displayed for links on the routing topology map. When link labels are shown, a label is added to each half of the link, one above and one below.

Changes that you make to node labels in the Administration Options window take effect as the default settings the next time you load the topology. To change node labels immediately, choose **View > Node Label** or **View > Link Labels**, and select the type of label to apply.

You can choose from the following node and link label options.

Node labels:

- **Automatic**—Automatically chooses a label for the node. For example, if no router name is available for the node, the system uses the Domain Name Service (DNS) Name node label. In automatic mode, node labels are prioritized by router names, DNS names, IP address, and system IDs. Automatic is enabled by default; selecting any of suboptions disables automatic mode.
- **DNS Names**—Includes the router DNS name. If no DNS name resolution is performed, selecting this option initiates DNS name resolution for all routers. In a large network, this can take time.
- **IPv4 Addresses—**Includes the router IPv4 address.
- **IPv6 Addresses—**Includes the router IPv6 address.
- **Router Names**—Includes the name of the router obtained from the protocol (if available).
- **System IDs (IS-IS only)**—Includes IS-IS System IDs when IS-IS is present on the routing topology map. If you choose **Show IS-IS NSAP Addresses** from the Miscellaneous options (as described in Map on page 76), the label NSAP Address is used.

- **Area IDs**—Includes IS-IS Area IDs when IS-IS is present on the topology map.
- **Label Routers Only**—Does not label the pseudonodes that represent LANs.
- **ID Numbers**—Includes an internal ID number for the router that may be useful as a shorthand reference.

Link labels:

- **Automatic**—Selects the Utilizations mode when the opened topology includes traffic, otherwise selects the Metrics mode.
- **Metrics**—Shows the IGP metric for each link, provided that the link represents only a single protocol adjacency.
- **Utilizations**—When the opened topology includes traffic, this option shows the traffic utilization level in Analysis and Planning modes but nothing in Monitoring mode. When the opened topology does not include traffic, this option is not shown in the menu and no link labels are shown on the map.
- **No Labels**—No link labels are shown.

Colors

The following color options are available for the routing topology map:

- **Color Allocation Order**—Changes the default colors for routers and links when you click and drag color samples in the chart. For example, to make orange the first color for an element on the topology map, click the orange color sample (by default, in position 5) and drag it to position 1 on the chart. The color formerly in position 1 replaces the color in position 5. Click **Set** to save changes or **Default** to return to the default settings.
- **Color links with metrics greater than**—Assigns colors to links that have metric values greater than the specified value. Use this option to find links with very high metrics.
- **Traffic Coloring Options**—Allows you to enter values to determine how bitrate and utilization are shown in color. (RAMS Traffic only)

- **Default Color Mode**—Allows you to manipulate the colors that are used in Analysis, Planning, and Monitoring modes. The following options appear in a drop-down list for each mode:

  > Monitoring mode can only use Color by Area and No Color.

  — **Color by Area**—The nodes and links of each area appear in distinct colors, which are determined by the Color Allocation Order chart, when a topology is loaded.

  — **Color by Traffic Bitrate**—Traffic flows are colored according to bitrate when a topology is loaded. Colors are determined by the Traffic Coloring Options chart. For example, flows with a high bitrate appear in red. (RAMS Traffic only)

  — **Color by Traffic Utilization**—Traffic flows are colored according to utilization levels when a topology is loaded. Colors are determined by the Traffic Coloring Options chart. For example, flows with high utilization appear in red. (RAMS Traffic only)

  — **No Color**—All areas and traffic flows appear with black nodes and gray links when a topology is loaded.

Map

You can set the following miscellaneous preferences, which take effect when you close the Options window:

- **Show link metrics on the map**—Displays metrics for the links on the map.
- **Show BGP Peerings**—If selected, shows all of the BGP peerings. You must reload the topology to view the changes.
- **Show BGP Next Hops**—If selected, shows all of the BGP next hops. You must reload the topology to view the changes.
- **Show Static Next Hops**—If selected, shows all of the static next hops. You must reload the topology to view the changes.

Because BGP peerings can be too dense on the map to be useful, the default setting is for the peerings not to be displayed. Changes to Show BGP Peerings and Show BGP Next Hops require a topology reload to display the changes.

If the system has a valid MPLS WAN license, then the Show BGP Next Hops and Show Static Next Hops fields, which are required to see the PE routers and PE/CE paths on the routing topology map, are automatically enabled and not displayed in this window. For more on the MPLS WAN feature, see MPLS WAN on page 479.

- **Hide links in group on the map—**Hides the links for the selected group.
- **Hide name suffix**—Determines how DNS names appear on the routing topology map. When DNS names of nodes appear, this suffix (if present) is trimmed from the names to reduce crowding of the layout. You can supply multiple DNS suffixes that are separated by a space, a comma, or a comma followed by a space. The default string is *mycompany.com*.
- **Path Highlight: ECMP degree**—Limits the number of Equal Cost Multi-Path (ECMP) routes to compute and appear when you highlight a path between two routers or use the **List/Find Path** option from the Tools menu. The default is 4096; set this value to 1 to disable ECMP routes.

Miscellaneous

The following miscellaneous preferences take effect when you close the Options window:

- **Font size**—Determines the font size for map labels.
- **Default layout on Open Topology**—Specifies the name of the saved layout that is used when loading a topology. This field is empty by default. You can enter a name or automatically set one by saving a topology layout choosing **Save Layout** and enabling the **Use as default layout** checkbox.
- **Show legend on Open Topology**—Automatically shows the legend panel when you load a topology. See Legend Panel on page 46 for information about the legend panel.

- **Show network summary on Open Topology**—Automatically shows the Network Summary when you load a topology. See "Network Summary Panel" on page 47 about this feature.
- **Show ISIS NSAP Addresses**—Appears only when the loaded topology contains IS-IS protocol routers. When this option is enabled, NSAP addresses appear on the topology map rather than the system ID for the IS-IS router. See "Node/Link Labels" on page 74.
- **Event panel default time interval**—Sets the time period (default 600 seconds) for routing events using the **Events** button on a Node Info panel or Link Info panel when there is no other time-range-based window opened. For example, if this value is set to 300 seconds, the list includes events that occurred in the past 300 seconds.
- **Max number of bars in bar chart view—**Sets the maximum number of bars that is displayed in bar chart views.
- **Display protocol packet events**—Display events indicating receipt of protocol packets including BGP updates, Link State Packets (LSPs), or Link State Advertisements (LSAs).

History Navigator

You can set the following default options for the History Navigator window.

- Choose the graphs that appear in the History Navigator window (only the Events graph is enabled by default).
  - **Events**—Show the number of routing protocol changes that occurred in the network between recorded snapshots. Examples include a neighbor adjacency going down or a new prefix being announced.
  - **Routers**—Show the number of physical entities in the network. For OSPF, this includes AS Border Routers in other areas that are visible within the viewed area.
  - **Links**—Show the sum of router-to-router links plus the number of router-to-prefix links. Does not apply to BGP protocols.
  - **IPv4 Prefixes/IPv6 Prefixes**—Show the cumulative number of IPv4 or IPv6 prefixes available in the network. Applies only to IS-IS and BGP.
  - **Routes**—Show the number of routes advertised in the network. Does not apply to IGP protocols.

- **ES Neighbors**—Show the number of ES neighbors.
- **Prefix Neighbors**—Show the number of prefix neighbors.
- **BGP Updates**—Show the number of announced and withdrawn packets found on the network in the preceding 10 minutes. Announced packets are represented by blue lines and re-announced packets are represented in dark yellow lines.
- **ISIS Activity**—Show LSP activity for IS-IS domains. New activity appears in blue, refreshed activity appears in dark yellow. (IS-IS only)
- **OSPF Activity**—Show LSA activity for OSPF domains. New activity appears in blue, refreshed activity appears in dark yellow.
- **EIGRP Activity**—Show updated and inferred activity for EIGRP domains. New activity appears in blue, refreshed activity appears in dark yellow.
- **Interfaces**—Show router interfaces discovered by static protocol.
- **Traffic—**Include traffic in the graph. (RAMS Traffic only)

- **Value of playback step in seconds**—Set the default value for a single step forward or backward in time during History Navigator playback. Default is 600 seconds.
- **Max number of data points in event graph**—Limit the event graph to the specified number of data points. Default is 25,000 data points.
- **Monitoring mode Update interval**—Sets the number of seconds between updates. Default is 10 seconds.

Auto-Hide

The system cannot determine whether a node or link that goes down has temporarily failed or is decommissioned. Use the following auto-hide options to determine when nodes and links that are down should be hidden or removed.

- **Seconds to auto-hide detached nodes—**Determines the number of seconds after which detached nodes are hidden. A node can become detached from the rest of the network when all of its links are auto-hidden (for example, if the node has failed temporarily or is decommissioned). The default value is **-1** (disabled).

- **Seconds to auto-hide pseudonodes with one attachment**—Determines the number of seconds after which pseudonodes are hidden. This condition may be caused by a problem in the router implementation of IS-IS. When the pseudonode for a network changes, the designated router of the old pseudonode does not flush its attachment to the old pseudonode.

  When setting this value, consider how many seconds should pass before auto-hiding the pseudonode when the number of attached links is reduced to one. The default value is **-1** (disabled).

- **Seconds to auto-hide failed links**—Determines the number of seconds after which a link that fails or is permanently decommissioned from service is removed from the map. These links are normally shown in red. The default value is 43,200 seconds (12 hours). To disable this option, change the value to **-1**. The link appears again if you use the History Navigator window to view a time when the link was up.

  Because links between BGP routers and their clients are inferred (as described in Links and Peerings on page 45), the system cannot conclusively determine if a peering has failed or is simply inactive. If the system does not receive routing information about a peering within a certain amount of time, the peering is marked **inactive** (rather than **down**) and is not removed from the routing topology map.

- **Seconds to auto-hide failed nodes**—Determines the number of seconds after which a node that fails or is permanently decommissioned from service is removed from the map. The default value is 43,200 seconds (12 hours). To disable this option, change the value to **-1**.
- **Seconds to auto-hide failed links to pseudonodes**—Determines the number of seconds after which a failed link with one end that is a pseudonode is removed from the map. The default value is 600.
- **Seconds to auto-hide failed pseudonodes**—Determines the number of seconds after which a pseudonode that fails is removed from the map. The default value is 600.

Working with Router Information and Layout

This section describes options for viewing and arranging routers (nodes) on the routing topology map:

- Viewing Node and Link Details on page 81
- Hiding Nodes on page 87
- Finding a Router on page 89
- Viewing Exit Routers on page 90
- Working with Tables on page 93
- Creating Custom Filters on page 101
- Assigning Router Names on page 104
- Viewing Current Network Inventory on page 107

To access these options, open the routing topology map, as described in Opening a Routing Topology Map on page 42.

Viewing Node and Link Details

In addition to the IP address and name labels, the system stores details about the nodes and links on the topology map. To access the extra detail about a particular node or link, right-click the object to open an information panel.

The title bar of the information panel shows the name of the node or link and the buttons shown in Table 16 .

Table 16 Information Title Bar Buttons

	**Tack**	Keeps the informational panel open. If you do not click the **Tack** icon, the current information panel closes when you click anywhere outside the panel.
	**Close**	Closes the detail panel.

Node Information Panel

Right-click a node in the routing topology map to access the node information panel (Figure 17). If the node is a pseudonode, the corresponding Designated Router (DR) is highlighted on the topology map. The information that is presented depends up on the node and associated protocol. For example, Figure 17 shows information for IS-IS on the left and OSPFv3 on the right.

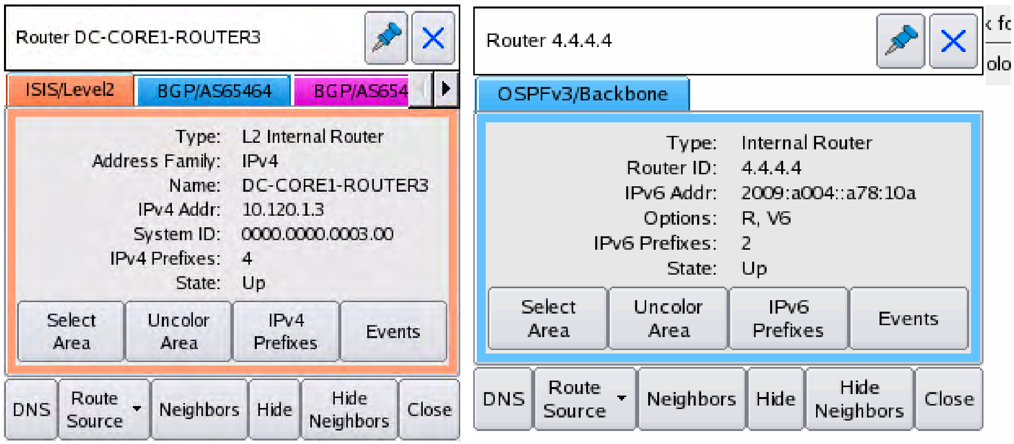

Figure 17 Node Information Panel

Each tab shows details for an instance of the node in a particular topology area (for example, OSPF area or IS-IS level). The color of the tab is the same as the color of the nodes and links in that area on the map, and the label identifies the protocol and area identifier. If the Collector is configured to include the node, an additional tab is presented to display some of that information.

The details available for a router depend on the protocol, but typically include the type of the node and one or more identifiers. In addition, the number of IPv4 and IPv6 prefixes the router announces and the up or down state of the router are shown. The tab label text is red if the router state is down. In the case of IS-IS nodes, the up state (overloaded) appears if the node is overloaded. An overloaded router remains active, but transit use is restricted.

A row of buttons on each tab provides access to functions specific to the particular instance (or routing process) on the node. The list of buttons depends on the current mode. The following buttons appear in all modes:

- **DNS**—Resolves the address of this router into a DNS name. If the resolution succeeds, the DNS name for the router appears on an additional line in the node information panel. This button is not present for pseudonodes. If an IPV6 address is available, then DNS tries to resolve using that address. If an IPv6 address is not available, IPv4 is used.
- **Route Source**—Sets this node as the starting point for path highlighting. After a node is set as the route source, the text on this button changes to Route Destination when the node information panel appears for another node. The path between the two nodes highlights if you click the button. For an example, see Highlighting the IP Route Between Two Points in the Network on page 134. This button is not present for pseudonodes.
- **Neighbors**—Provides a list of neighboring routers and the interface addresses and metrics of the links connecting them.

> For IS-IS routers that do not enable Traffic Engineering (TE) extensions, the interface addresses is not known. If a single /30 or /31 prefix is in common between the adjacent routers, the prefix appears in place of the source and destination interface addresses.

- **Hide**—Hides the selected node. To view the node, choose **View > Unhide All Nodes** or click **Unhide All Nodes** in the lower-right corner of the Routing Topology Map window. Click the button again to hide the nodes.
- **Hide Neighbors**—Hides the neighbors of the selected node.
- **Close**—Closes the information panel. The panel also closes if you click on the map.

In Monitoring and Analysis modes, the following buttons are displayed:

- **Select Area**—Selects all of the nodes within the routing area that contains this node and draws a bounding box around those nodes on the topology map.
- **Uncolor/Color Area**—Colors or removes colors from the area of the topology map that contains this node.
- **IPv4 Prefixes/IPv6 Prefixes**—Shows a list of the IPv4 or IPv6 prefixes announced by this router.
- **OSI Prefixes**—Shows a list of the OSI prefixes announced by this router. This option appears only when OSI IS-IS is detected.

- **Events**—Shows a list originated by this router or for which this router is the neighbor. It also shows information regarding a neighbor announcing this router, including the parameters of the connection. If a time range is selected in the History Navigator window, the same time range is used for this list; otherwise, events occurring in the last 10 minutes are listed.

In Planning mode, the following additional buttons are displayed:

- **Down**—Brings down the node for planning purposes.
- **Set Overloaded**—For IS-IS nodes only. The node is up, but not able to send data through the network. The node can receive traffic but routes are not permitted to go through it.

In RAMS Traffic, if there is traffic information for the node, a Traffic tab opens in addition to the protocol tabs, as shown in Figure 18.

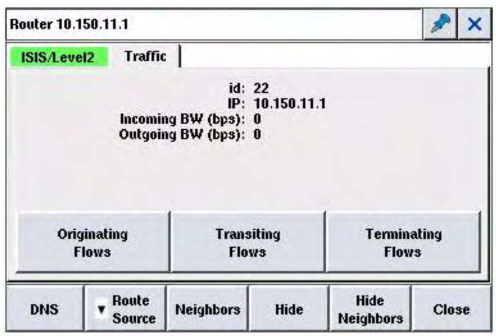

Figure 18 Node Traffic Tab

The body of the tab shows the IP address of the router, and information about ingress and egress flows. For example, Egress Flows: 12 indicates that the router is exporting 12 traffic flows. Egress bandwidth (BW) in bps: 1.99 indicates that these 12 flows have a total bitrate of 1.99 K.

Click **Ingress Flows** and **Egress Flows** to show a table that lists detailed information for each flow moving into or out of the router.

Link Information Panel

Right-click a link in the routing topology map to open a link information panel similar to the example shown in Figure 19.

Figure 19 Link Information Panel

Each link has two halves that represent the two directions of communication. The router interface corresponding to the half that was clicked appears on the left side of the link information panel. The exception is links to pseudonodes. For pseudonodes, the interface for the real router appears on the left and the pseudonode on the right. The direction is indicated by the node names in the title bar of the panel.

The link information panel can have one or more tabs. The first tab indicates the instance of the link in a particular topology area (for example, EIGRP area). If there is more than one tab, it indicates that the routers at the two ends of the link participate in more than one routing protocol (or multiple instances of the same protocol).

You can select a top-level tab to show the link details corresponding to the protocol instance of that tab. The color of each tab is the same as the color of the nodes and links in that area on the map. The tab label tells the protocol and area identifier.

An inner level of tabs is included if the link on the map represents multiple, parallel, physical links between the two routers. The inner tab labels indicate the interface of the router on the left side of the arrow in the title bar (for example, 10.72.4.42/24). The body of the tab provides details about the interfaces on the routers at each end of the link, including the following information:

- **Interface**—Address of the interface for the routers at either end of the link.

> For IS-IS routers that do not enable TE extensions, the interface addresses will not be known. If there is a single /30 or /31 prefix in common between the adjacent routers, the prefix will appear in place of the source and destination interface addresses.

- **Metric**—Metric value for each interface. The link metric value helps determine the best path to take through the network, and is based on bandwidth and delay, among other factors. For EIGRP protocol links, the metric is shown as two components that are calculated from inverse bandwidth and delay.
- **State**—Indication of whether a link is up, down, or inactive.The state Inactive applies to BGP peerings only. Unlike the physical or virtual links that connect link-state routers (OSPF and IS-IS), peerings between BGP protocol routers and their clients are inferred. If the system fails to receive information about BGP peering within a certain time frame, the link is marked Inactive, as the system cannot determine conclusively if the peering is down.
- **BW** (for EIGRP)—Bandwidth of the data traveling across the link.
- **Delay** (for EIGRP)—Time required to move packets from the source to destination.

  A row of buttons on the tab provides access to functions specific to the particular instance
  (or routing process) on the node, including the following:

- **Uncolor/Color AS**—Color or remove color from the AS of the routing topology map that contains this link.
- **Select AS**—Select all of the ASs within the routing area that contains this link. A bounding box is drawn around the ASs.
- **Events**—List routing events related to adjacency changes on this link. If a time range is selected in the History Navigator window, the same time range is used for this list; otherwise, events occurring in the last 10 minutes are listed.
- **Close**—Close the information panel. The panel also closes if you click on the map.

In RAMS Traffic, if there is traffic information for the link, a Traffic tab appears next to the protocol tabs as shown in Figure 20.

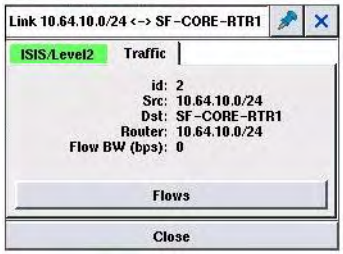

Figure 20 Link Traffic Tab

The Traffic tab provides the IP address of the router (Rtr), the number of flows on the link (Flows), and the total rate of the flows (Flow BW) in kilobits per second (kbps).

You can click Flows to show the Details by Flow table.

Hiding Nodes

You can instruct the system to show a particular class of routers on the topology map based on the naming conventions established when the routers were named. If the core router names have a common text string in their name, for example the letters core-gw, you can use the Hide Nodes option to show only these routers.

You can also hide nodes through leaf trimming. With this option, all routers on the edge of the network with a single link to the network are hidden from view. This operation can be repeated multiple times. Edge nodes with only one link are removed each time you select **Hide Leaf Nodes**.

To hide the listed nodes, perform the following steps:

1 Select **View > Hide Nodes** to open the hide Nodes window, as shown in Figure 21.

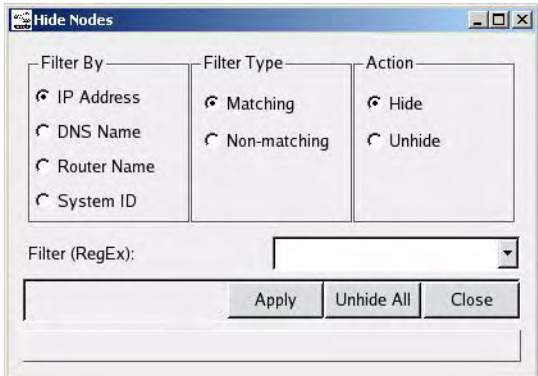

Figure 21 Hide Nodes

2 Choose the desired filtering option.

3 Click **Matching** to select routers that match the criteria of the filter or **Non-matching** to select routers that do not match the criteria of the filter.

4 Click **Hide** in the Action section to remove nodes or click **Unhide** to restore them.

5 In the **Filter (RegEx)** field, enter a regular expression to select the class of routers to be matched. For example, `-core-gw$` matches routers whose names end with -core-gw. The string `^10\.251\.` matches addresses that begin with 10.251. Matching is case-sensitive.

> The syntax of extended regular expressions is explained in Regular Expressions on page 226. The syntax is not the same as shell or file manager pattern patching, so a pattern like *-core-gw is not correct.

6 Click **Apply** to save your changes.

7 Click **Close**.

To hide leaf nodes, perform the following steps:

1 Open the routing topology map.

2 Select **View > Hide Leaf Nodes** to hide the nodes on the edge of the network. Figure 22shows the result of the Hide Leaf Nodes operation performed twice in succession. Hidden nodes are saved with the topology layout.

Figure 22 Hiding Leaf Nodes - Before and After

To restore hidden nodes, click **Unhide All Nodes**.

This operation does not hide all edge routers since edge routers often have multiple redundant connections to the network core.

Finding a Router

You can locate a router on the map according to name or address. For instructions on scanning the router list, see Router List on page 108.

To find a router by name or IP address, perform the following steps:

1 Select **Tools > Find Router**.

2 In the Search For field, enter the IP address, name or system ID of a router, or the prefix of a LAN pseudonode (Figure 23). If the name entered is the initial portion of multiple router names, then all those routers will be matched.

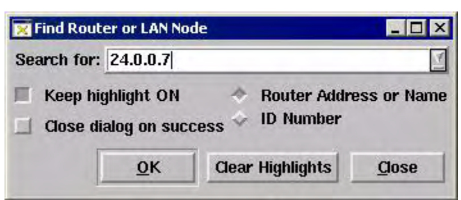

Figure 23 Find Router

3 Click **OK**. The router flashes yellow on the routing topology map.

4 To highlight multiple routers at the same time, select **Keep highlight ON** and deselect **Close dialog on success**, and then repeat the previous steps.

Viewing Exit Routers

In a large network with multiple exit routers to the Internet, it is often useful to see which exits are being taken from various points in the network. This information can help you balance flows to achieve optimum performance or minimize monthly cost in transit fees and bandwidth costs.

The system can calculate the IGP path from each router to its nearest exit router if there is a default route or if external routes are redistributed from BGP. Each router is then color-coded by exit router and exit routers are highlighted in flashing yellow. Alternatively, to highlight the path from a single router to its exit router, see Finding a Route By Prefix on page 136.

To view the exit routers from all routers in the network, perform the following steps:

1 Choose **Tools > Highlight by Exit Router**.

   The Highlight By Exit Router search opens.

2 Enter the desired Internet prefix, NSAP address, or domain name in the Prefix DNS Name field. You can enter an IPv4 or IPv6 prefix.

3 Click **OK**.

Combining Routers

The appliance supports the ability to manually correct problems with consolidation of router nodes on the map in the event that the automatic procedures fail. The Combine Routers feature allows combining of two map nodes into one or separating one node into two. The following examples illustrate the two cases:

- In a network where BGP and OSPF are recorded, the appliance uses heuristics to determine if a particular BGP peer is the same router as one of the routers that is learned in the OSPF topology. If the heuristics fail, then separate router nodes are shown on the map for the BGP and OSPF instances of the router. The Combine Routers feature permits the manual merging of the two nodes into one so that routing functions correctly.

- In a network resulting from the merger of two companies, there could be two different routers in two separate IS-IS areas that have been assigned the same System ID. The heuristic combines those two routers into one node, which is not correct. The Combine Routers feature can be used to separate the two routers.

In most networks, the heuristics produce the correct consolidation of all the protocol instances of each router based on the various addresses and prefixes associated with the protocol instances, so you do not need to use the Combine Routers feature. In addition, if the Collector recorder has been configured to obtain router information with SNMP or CLI and if the routers can be successfully accessed by these methods, then the appliance can automatically detect and correct consolidation errors by identifying each router according to its interface addresses and other details. For information on configuring the Collector, see the "Configuration" chapter in the *HP Route Analytics Management Software Administrator Guide*.

In the Combine Routers window (Figure 24),the parent rows in the table represent the multi-protocol nodes on the map and the child rows of the table represent the individual protocol router instances that are combined into the node. You can join some or all of the protocol instances in one node with those of another node, or you can split some protocol router instances out of one node to make a new one if the existing consolidation is incorrect.

To combine or separate routers manually:

1. Choose **Administration > Combine Routers**

   The system presents a warning that changing router combinations can affect core functionality such as path finding, traffic volume calculations, and displaying nodes on the routing topology map.

2. Click **OK** to open the Combine Routers table. The parent rows in the table represent nodes on the map and the child rows of the table represent the individual protocol router instances.

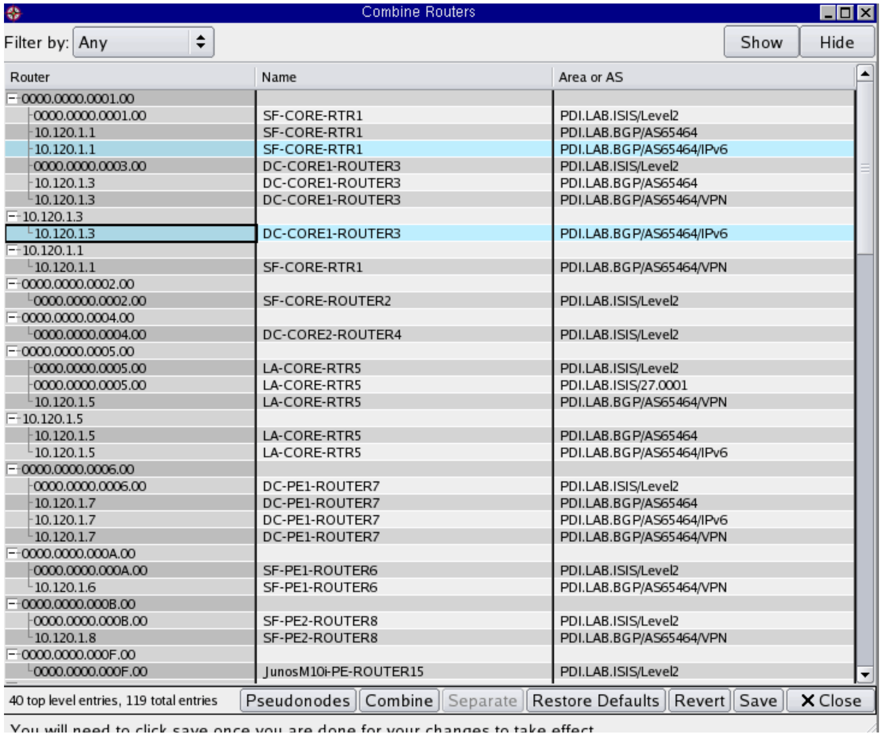

Figure 24 Combine Routers Table

3 Perform any of the following tasks.[1] The indicated buttons are activated when appropriate selections are made from the table. Multiple rows can be selected in a range with Shift-click or individually with Ctrl-click.

- Use the Filter by drop-down list to show only specified entries in the table. See Creating Custom Filters on page 101.
- To combine different protocol instances into a single parent router, select the individual protocol instances under the different parent routers, and click **Combine**.
- To separate previously-combined protocol instances into separate routers, select the protocol instances to be removed from a particular parent router, and click **Separate**.

1. The Pseudonodes button is not currently used.

- To return to the default organization of the table, click **Restore Defaults**.
- To return to the last saved configuration, click **Revert**.
- To save the changes, click **Save**.
- To close the window, click **Close**.

Working with Tables

Information in the client application is often presented in table format. Many of the tables support right-click column options, as shown in Figure 17.

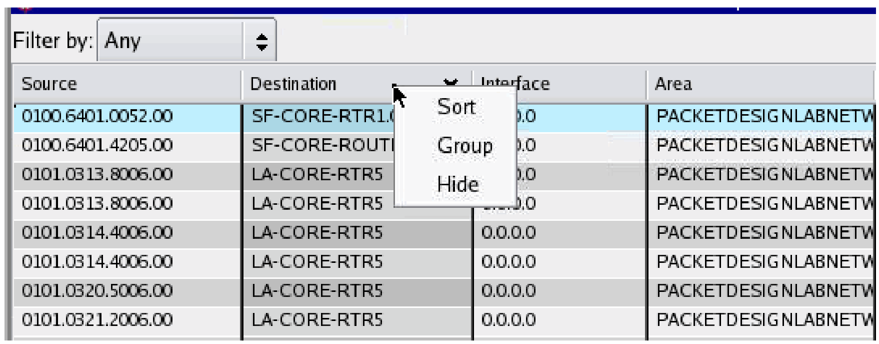

Figure 25 Right-Click Column Options

The following options may be supported, depending upon the table:

- Sort—Select to sort the column in ascending order and select again to sort in descending order.
- Group—Combine all entries with the same name in a tree view. Click - to hide the sub-entries and **+** to show the sub-entries (Figure 17). Note that only the sub-entries have data in the other table columns.

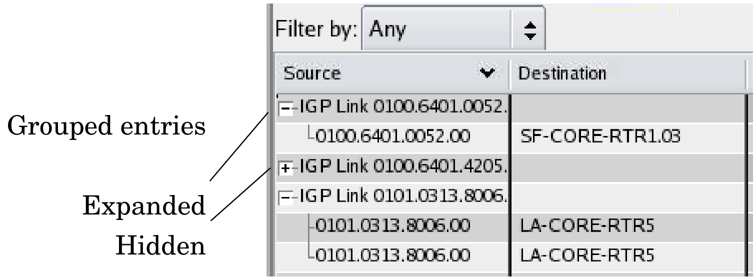

Figure 26 Grouping Column Entries

- Ungroup—Restore the ungrouped view.
- Collapse All—Show only the top level groups.
- Expand All—Show all of the grouped entries.
- Hide—Conceal the selected column.
- Show—Display the previously hidden column. (Choose **Show** and select the column name.)

Setting Up Prefix Diagnostics

The Prefix Diagnostics window allows you to generate reports to help troubleshoot issues with prefixes. While much of the information in the Prefix Diagnostics reports is available in other existing reports, Prefix Diagnostics provide a centralized view of the prefix information.

Prefix diagnostics work only on routing topologies, not traffic topologies. You must be in Analysis mode to use this feature. You can specify IPv6 prefixes if IPv6 is licensed and the recorder topology includes IPv6 prefixes.

To configure and view prefix diagnostics, perform the following steps:

1 In Analysis mode, choose **Tools > Prefix Diagnostics**.

The Configure Prefix Diagnostics Report window opens.

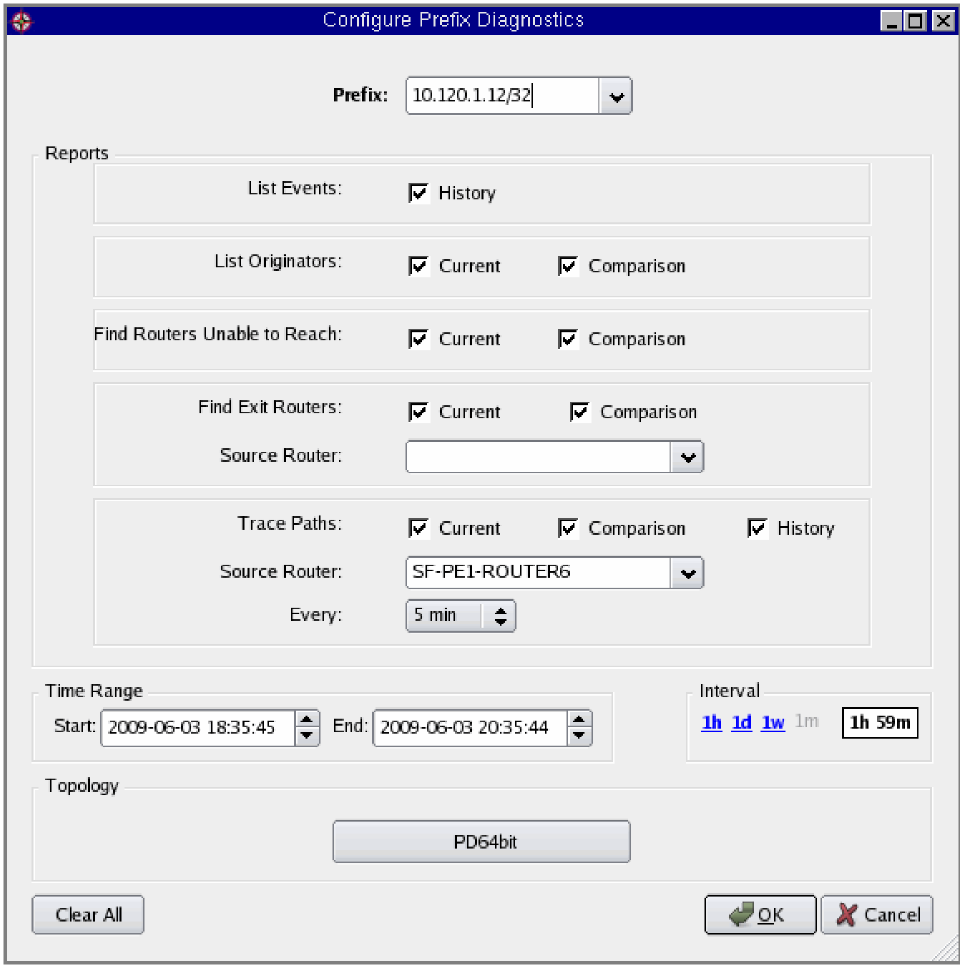

Figure 27 Prefix Diagnostics Configuration

2 Enter or select the prefix of interest from the drop-down list. You can specify IPv6 prefixes if IPv6 is licensed and the recorder topology includes IPv6 prefixes.

3 To restrict the report topology to a subset of the current topology, click the button with the topology name above the bottom row of buttons, and make a selection in the window that opens. Restricting the topology allows you to focus on specific areas of interest.

> Topology restrictions automatically apply to all settings in the report.

4 Select check boxes to specify the reports and associated options.

Reports:

- List Events—List of events that happened for this prefix.
- List Originators—List of the routers that originate this prefix.
- Find Routers Unable to Reach—List of the routers that cannot reach the specified prefix.
- Find Exit Routers (include source router)—List of the last router or routers in the topology that are in the path (from the source, if specified) to the selected prefix. If a source is not specified, then all exit routers are listed.
- Trace Paths (include source router)—List of the paths from the specified source to the selected prefix. Includes a step size for the history report.

Report options (availability depends on the particular report):

- Current—Reports are generated for the specified end time.
- Comparison—Reports present information for the start and end times. In these reports, the Before column refers to the starting time and the After column refers to the ending time.
- History—Reports present information for the full time period between the starting and ending time.

> Click **Clear All** if you want to clear all of the entries on the page.

5 Specify times in the Start and End fields. The scale units, 1h (one hour), 1d (one day), 1w (one week), and 1m (one month) adjust the start time such that the interval between the start and end times is the selected period (such as one hour or one day).

6 Click **OK** to accept the settings and open the report window.

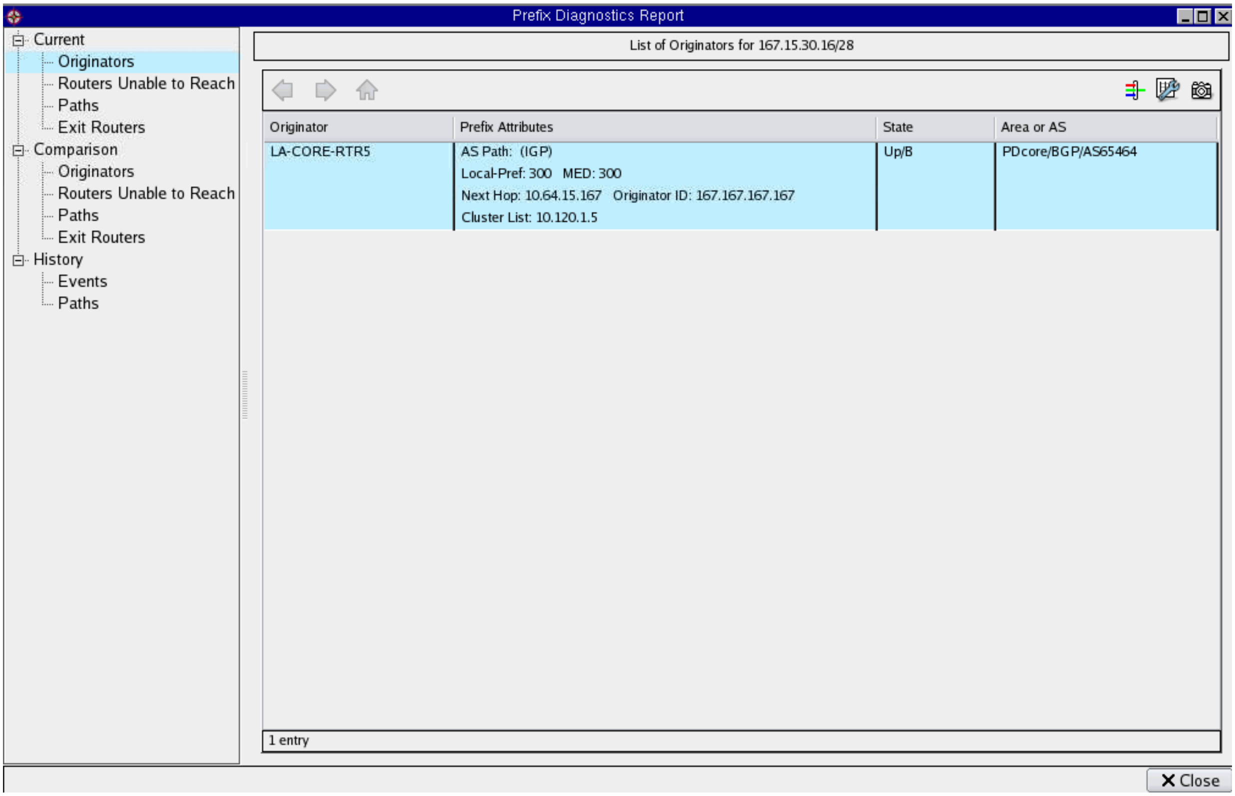

Figure 28 Prefix Diagnostics Report

7 The side menu lists all of the possible reports, however, only the reports that were selected in the configuration window are active; the others are grayed out.

Prefix Diagnostics Reports Buttons

The following buttons are available in Prefix Diagnostics reports, depending on the specific report selection.

Table 17 Prefix Diagnostics Report Buttons

	**Drill-down**	If available, this button allows to see finer detail within a set of data.
	**Go back one drill-down**	During a drill-down, goes back one drill-down level.
	**Go forward one drill-down**	During a drill-down, goes forward one-drill down level.
	**Go back to top; undo all drill-downs**	Goes to the highest point in the drill-down hierarchy, and "unrolls" the drill-down view from the window.
	**Configure**	Opens the configuration window. See Setting Up Prefix Diagnostics on page 94.
	**Advanced Filter**	Allows you to define advanced filters. See Advanced Filtering on page 397 and Using Filters on page 221 for more information.
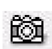	**Snapshot**	Opens a new window containing a snapshot of the current window.

Prefix Diagnostics Reports

The following reports are available in the Prefix Diagnostics Report window:

- Originators Report on page 99
- Routers Unable to Reach Report on page 99
- Paths Report on page 100
- Exit Routers on page 101
- Events Report on page 101

For Comparison reports, the Before column refers to the Start time and the After column refers to the End time. History reports cover the full time period between the starting and ending time.

Originators Report

The Originators report is available with the Current and Comparison options, and includes the following columns:

- Originator—Name or IP address of the originating router.
- Prefix Attributes—Attributes of the prefix as announced by the router. For comparison reports, the Before and After values refer to the start and end times.
- State—State of the prefix. For comparison reports, the Before and After values refer to the start and end times.
- Area or AS—Area of the router.

Routers Unable to Reach Report

The Routers Unable to Reach report is available with the Current and Comparison options, and includes the following columns:

- Router—Name or IP address of the router that cannot reach the prefix.
- Unable to Reach—Comparison report only. Indication of whether the router was able to reach the specified prefix at the start (Before column) and end (After column) times.
- Last Hop Router—Router from which the router could not continue to the specified prefix. For comparison reports, the Before and After values refer to the start and end times.
- Last Hop Area—Area in which the last hop router is located. If the router has multiple areas, the following criterion is applied to determine the area. IGP areas (EIGRP AS) is preferred over BGP and if there are multiple IGP areas, then backbone (OSFP) or Level 2 (IS-IS) is preferred over other areas. For comparison reports, the Before and After values refer to the start and end times.

Paths Report

This report extends the information that is included in the Find or List Paths report (see Finding a Route By Prefix on page 136).

For Comparison reports, the Before information is grouped in the first set of rows, and the After values are grouped in the second set of rows. To quickly find the Before and After rows, right click in the Path column header and choose **Collapse All**. You can then click the **+** sign for either link to expand those entries.

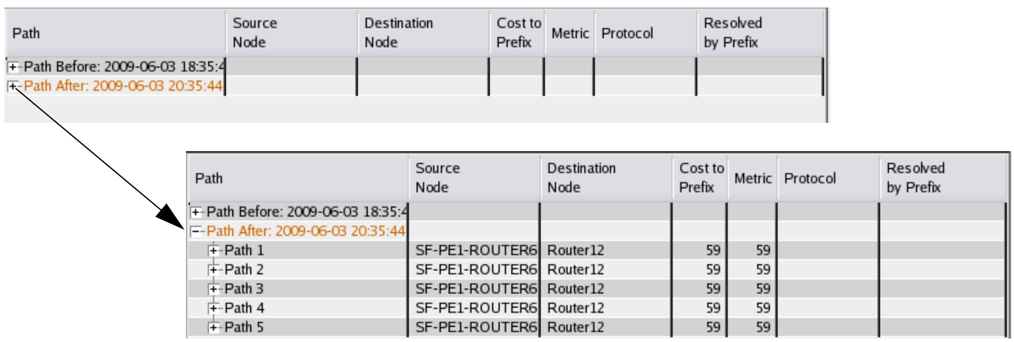

Path	Source Node	Destination Node	Cost to Prefix	Metric	Protocol	Resolved by Prefix
+ Path Before: 2009-06-03 18:35:4						
+ Path After: 2009-06-03 20:35:44						

Path	Source Node	Destination Node	Cost to Prefix	Metric	Protocol	Resolved by Prefix
+ Path Before: 2009-06-03 18:35:4						
− Path After: 2009-06-03 20:35:44						
+ Path 1	SF-PE1-ROUTER6	Router12	59	59		
+ Path 2	SF-PE1-ROUTER6	Router12	59	59		
+ Path 3	SF-PE1-ROUTER6	Router12	59	59		
+ Path 4	SF-PE1-ROUTER6	Router12	59	59		
+ Path 5	SF-PE1-ROUTER6	Router12	59	59		

Figure 29 Expanding Before and After Entries for the Paths Prefix Diagnostics Report

For the History path report, the number of paths may be very large. To reduce the length of the report while preserving the important information, each path is shown the first time it is identified, but subsequently, the path is shown only if there has been a change.

The Paths report is available with the Current, Comparison, and History options, and includes the following columns:

- Path—Name of the originating router.
- Source Node—Name or IP address of the source router.
- Destination Node—Name or IP address of the destination router.
- Cost to Prefix—Router's calculated cost to reach the destination prefix.
- Metric—Link metric.
- Protocol—Protocol over the path.
- Resolved by Prefix—Prefix by which each next hop was resolved.

Exit Routers

The Exit Routers report is available with the Current and Comparison options, and includes the following columns: If the source router is selected, only that router is included as a source.

- Router—Name or IP address of the source router.
- Exit Router—Last hop router in the selected topology.
- Hops—Number of hops from originating to exit router. If an ECMP path exists, then the number of hops is the maximum number of hops among the possible ECMP paths.
- Cost to Prefix—Router's calculated cost to reach the destination prefix.

Events Report

The Events report contains historical information for the interval between the start and end times and includes only the observations that correspond to events involving the specified prefix.

The contents and treatment of time ranges are the same as for the Event Details report in the History Navigator. See Event Details on page 204.

The report includes the following columns:

- Time—Date and time of the event.
- Router—IP address of the originating router.
- Operation—Type of event.
- Neighbor/Prefix—Neighboring router for neighbor operations or the prefix for prefix operations.
- Attributes—Affected attributes of the router or prefix.
- Area or AS—Area or AS in which the router is located.

Creating Custom Filters

The Custom Filter Repository allows you to create and save custom filters that you can apply to other filter windows.

In other windows that you use filters, the custom filters you create appear in the Filter By Expression field.

For additional information about using filters, see Using Filters on page 221 in Chapter 4, “The History Navigator”

To access the custom filter repository, perform the following steps:

1 Choose **Administration > Saved Filters** to open the Custom Filter Repository window. (Figure 30).

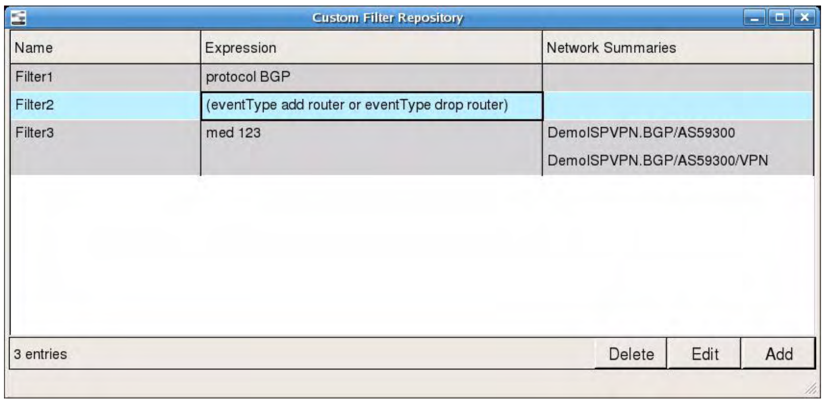

Figure 30 Custom Filter Repository

The following columns are displayed:

- **Name**—Custom filter name
- **Expression**—Filter expression associated with the custom filter name.
- **Network Summaries**—BGP and VPN network summary in which the filter is used. You can filter the active routes in a BGP or VPN topology and view the results in its Network Summary.

Only BGP or VPN-specific filters are intended for showing results in their respective network summaries. Using other filters (for example, IGP filters) will result in empty counts in the Network Summary column.

To add a custom filter, perform the following steps:

1 Choose **Administration > Saved Filters**.

2 Click **Add** from the Custom Filter Repository.

3 In the Name field, enter the name of the filter (Figure 31).

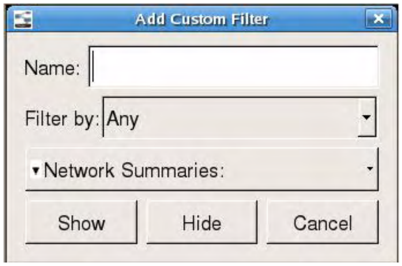

Figure 31 Add Custom Filter

4 Choose the type of filter from the drop-down list and then select the desired option from the options that are displayed when you choose the filter type.

5 Select the network summaries that you want to include.

6 Click **Show** to create a custom filter that accepts all items meeting the filter conditions, or click **Hide** to create a custom filter that rejects all items meeting the filter conditions. Clicking **Show** or **Hide** also saves the filter. Click **Cancel** to stop the operation and close the window.

To edit a custom filter, perform the following steps:

1 Choose **Administration > Saved Filters**.

2 Choose the filter you want to edit, and click **Edit**.

3 Modify the filter expression in the Filter field.

You can edit only the filter expression, not the filter name.

4 Click **Show** to have the filter accept all items meeting the filter conditions, or click **Hide** to have the filter reject all items meeting the filter conditions. Clicking **Show** or **Hide** also saves the filter. Click **Cancel** to stop the operation and close the window.

To delete a custom filter, perform the following steps:

1 Select **Administration > Saved Filters**.

2 Select the filter you want to delete and click **Delete**.

3 Click **Yes** to confirm.

Assigning Router Names

Use the Router Name feature to control how the routers in your system are identified. By default, the IP Address of the router is used to identify the router.

To assign router names, perform the following steps:

1 Choose **Administration > Assign Names > Routers** to open the Router Names window (Figure 32).

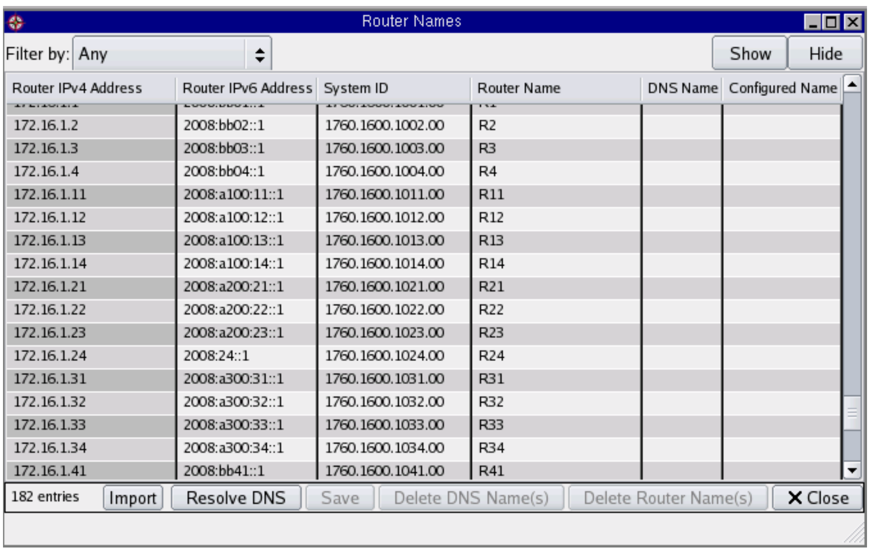

Figure 32 Router Names Window

The table contains the following columns:

- **IP Address—**IP Address for a particular router.
- **System ID—**System ID received from the router (empty if there is no system ID).

The System ID column will appear only if IS-IS is detected.

- **Router Name—**Router name derived from the routing protocol.

- **DNS Name—**Name resolved with the DNS server using the router's IP address.
- **Configured Name—**User-defined name that identifies the router.

Router names are prioritized in the following order:

1. **User-configured name** (highest priority)
2. **Protocol-delivered router name**
3. **DNS name**

Use the following Filter by options to determine which routers are shown or hidden:

- **Any—**Show all rows.
- **Router IP Address—**Enter or choose the IP address of the router. If you select the router's IP address and click **Show**, the system lists only the row with the IP address that you entered. If you click **Hide**, the system shows all the router IP addresses except the one that you entered.
- **Router Names—**Enter or choose the name of the routers. Choose any of the following options:
  - **Substring—**Filters the routers with the given string as a substring for either Router Name, DNS Name, or Configured Name.
  - **Exact Match—**Filters the routers with the given string as an exact match either of its Router Name, DNS Name, or Configured Name.
  - **Begins With—**Filters the routers with the beginning string used for router name, DNS name, or user-specified name.

The following buttons are included in the Router Names window:

- **Show—**Show all the router entries to view.
- **Hide—**Hide router entries.
- **Save—**Save information. This button is active only when you have modified information on the screen.
- **Import—**Edit multiple router names**.**
- **Close—**Exit the Router Names window.

To change a router name, perform the following steps:

1 Choose **Administration > Assign Names > Routers** to open the Router Names window.

2 Select the router.

3 Enter the new name in the Configured Name column.

4 Click **Save**.

Changing Multiple Router Names

To change the names of multiple routers, perform the following steps:

1 Choose **Administration > Assign Names > Routers** to open the Router Names window.

2 Click **Import**.

3 Enter router names in the format shown in Figure 33.

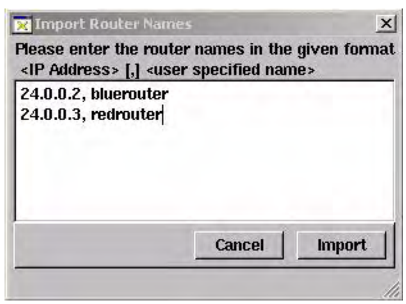

Figure 33 Import Router Names Dialog B

4 Click **Import**.

If you attempt to import a router name that is not known or in an incorrect format, the error message "Discarded Invalid Import Entries" will appear at the bottom of the topology window.

Assigning IPv4 or IPv6 Prefix Names

Use the Prefix Name feature to identify prefixes by name in MPLS VAN WAN reachability reports.

To assign prefix names, perform the following steps:

1 Choose **Administration > Assign Names > IPv4 Prefix Names** or **Administration > Assign Names > IPv4 Prefix Names** to open the Prefix Names window (Figure 3-24).

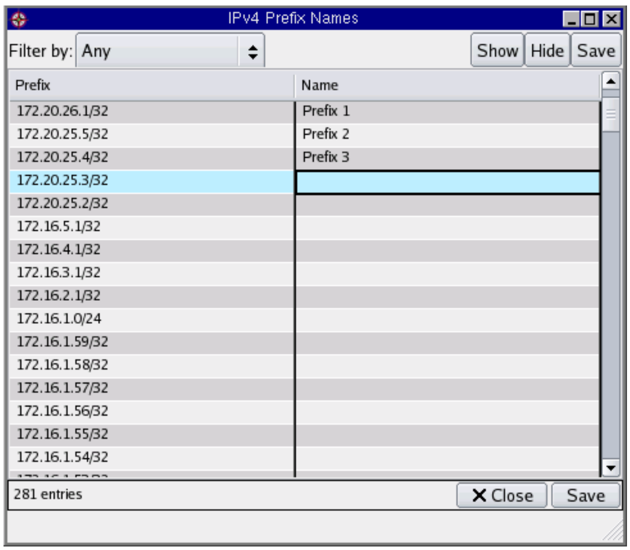

Figure 3-1 Prefix Names Window (IPv4 Shown)

2 Double-click an entry in the Name column and enter the name.

3 Click **Save**.

Viewing Current Network Inventory

The application can show the current list of routers, links, and prefixes in each IGP area and AS of a network. You can sort these lists by prefix, AS, attributes, status, etc., and trace any entry in any back to the associated router in the topology map with a single click.

Router List

Open the List of All Routers window to perform the following operations:

- Verify if a particular router is currently up and running.
- Verify that a router appears in the correct IGP area or AS.
- Verify that a router is configured as expected, including that the correct IP address is associated with it and that it has the correct name (for IS-IS or EIGRP).
- List the hardware type and software version of each router (EIGRP only).
- View the total number of routers currently in the network.
- Verify the health of the network visually without sifting through hundreds of syslog entries.

To find a router using the List of All Routers window, perform the following steps:

1 Choose **Tools > List Routers** to open the List of All Routers window.

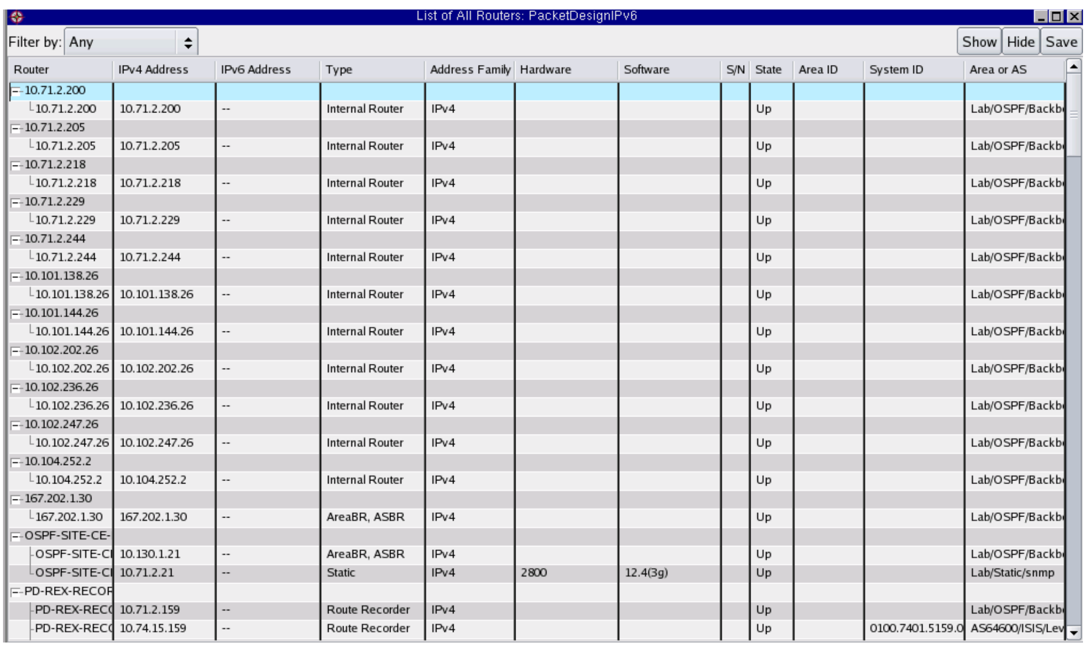

List of All Routers: PacketDesignIPv6

Filter by: Any

Show Hide Save

Router	IPv4 Address	IPv6 Address	Type	Address Family	Hardware	Software	S/N	State	Area ID	System ID	Area or AS
10.71.2.200											
10.71.2.200	10.71.2.200	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.71.2.205											
10.71.2.205	10.71.2.205	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.71.2.218											
10.71.2.218	10.71.2.218	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.71.2.229											
10.71.2.229	10.71.2.229	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.71.2.244											
10.71.2.244	10.71.2.244	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.101.138.26											
10.101.138.26	10.101.138.26	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.101.144.26											
10.101.144.26	10.101.144.26	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.102.202.26											
10.102.202.26	10.102.202.26	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.102.236.26											
10.102.236.26	10.102.236.26	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.102.247.26											
10.102.247.26	10.102.247.26	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
10.104.252.2											
10.104.252.2	10.104.252.2	--	Internal Router	IPv4				Up			Lab/OSPF/Backb
167.202.1.30											
167.202.1.30	167.202.1.30	--	AreaBR, ASBR	IPv4				Up			Lab/OSPF/Backb
OSPF-SITE-CE-											
OSPF-SITE-CI	10.130.1.21	--	AreaBR, ASBR	IPv4				Up			Lab/OSPF/Backb
OSPF-SITE-CI	10.71.2.21	--	Static	IPv4	2800	12.4(3g)		Up			Lab/Static/snmp
PD-REX-RECOF											
PD-REX-RECO	10.71.2.159	--	Route Recorder	IPv4				Up			Lab/OSPF/Backb
PD-REX-RECO	10.74.15.159	--	Route Recorder	IPv4				Up		0100.7401.5159.0	AS64600/ISIS/Lev

Figure 34 List of All Routers

2 Use the Filter by drop-down list at the top of the window to filter the list as needed (see Using Filters on page 221). You can filter or search by IPv4 or IPv6 address.

3 To identify a router on the map, click anywhere in the row corresponding to the router. The row is highlighted in the List of All Routers window and the router flashes yellow on the routing topology map.

Links List

Open the List of All Links window to display the number and current state (up or down) of all routing adjacencies in the network, along with their link metrics and the router interface addresses.

For IS-IS routers that do not enable TE extensions, the interface addresses is not known. If there is a single /30 or /31 prefix in common between the adjacent routers, the prefix appears in place of the source and destination interface addresses.

To view the list of all links, perform the following steps:

1 Choose **Tools > List Links** to open List of All Links window (Figure 35).

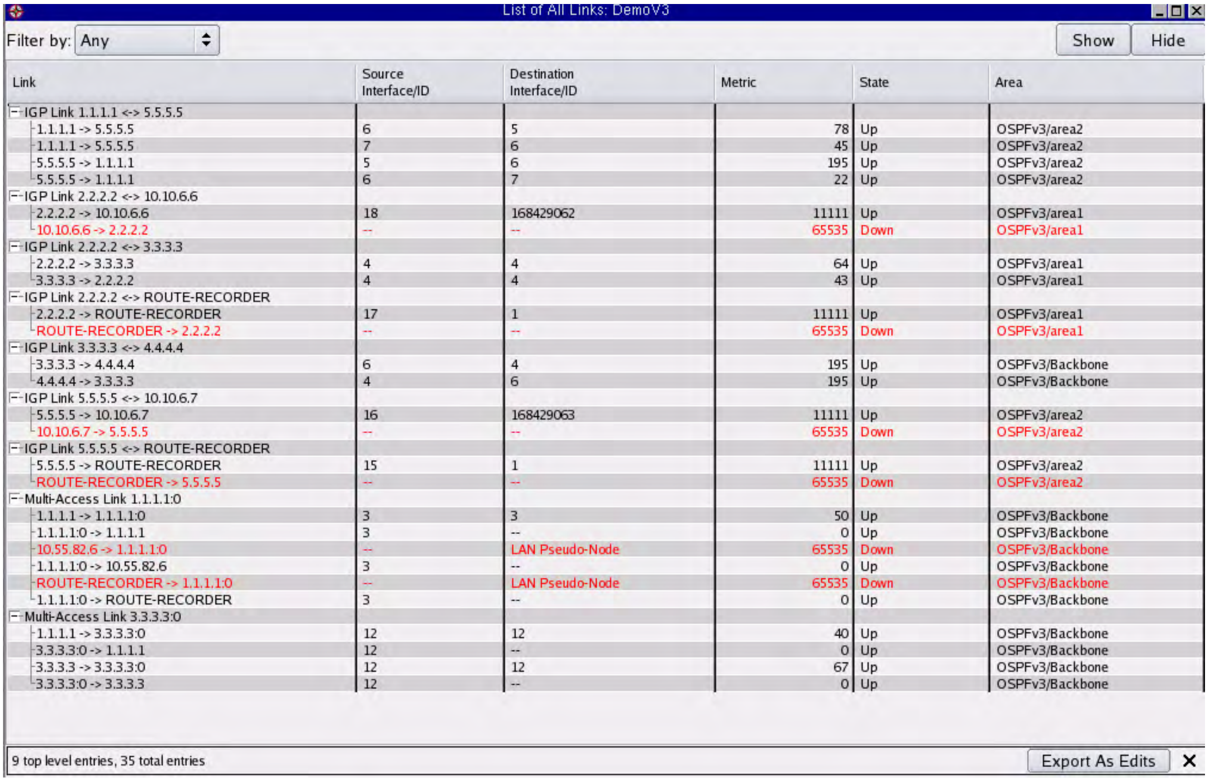

List of All Links: DemoV3

Filter by: Any — Show — Hide

Link	Source Interface/ID	Destination Interface/ID	Metric	State	Area
IGP Link 1.1.1.1 <-> 5.5.5.5					
1.1.1.1 -> 5.5.5.5	6	5	78	Up	OSPFv3/area2
1.1.1.1 -> 5.5.5.5	7	6	45	Up	OSPFv3/area2
5.5.5.5 -> 1.1.1.1	5	6	195	Up	OSPFv3/area2
5.5.5.5 -> 1.1.1.1	6	7	22	Up	OSPFv3/area2
IGP Link 2.2.2.2 <-> 10.10.6.6					
2.2.2.2 -> 10.10.6.6	18	168429062	11111	Up	OSPFv3/area1
10.10.6.6 -> 2.2.2.2	--	--	65535	Down	OSPFv3/area1
IGP Link 2.2.2.2 <-> 3.3.3.3					
2.2.2.2 -> 3.3.3.3	4	4	64	Up	OSPFv3/area1
3.3.3.3 -> 2.2.2.2	4	4	43	Up	OSPFv3/area1
IGP Link 2.2.2.2 <-> ROUTE-RECORDER					
2.2.2.2 -> ROUTE-RECORDER	17	1	11111	Up	OSPFv3/area1
ROUTE-RECORDER -> 2.2.2.2	--	--	65535	Down	OSPFv3/area1
IGP Link 3.3.3.3 <-> 4.4.4.4					
3.3.3.3 -> 4.4.4.4	6	4	195	Up	OSPFv3/Backbone
4.4.4.4 -> 3.3.3.3	4	6	195	Up	OSPFv3/Backbone
IGP Link 5.5.5.5 <-> 10.10.6.7					
5.5.5.5 -> 10.10.6.7	16	168429063	11111	Up	OSPFv3/area2
10.10.6.7 -> 5.5.5.5	--	--	65535	Down	OSPFv3/area2
IGP Link 5.5.5.5 <-> ROUTE-RECORDER					
5.5.5.5 -> ROUTE-RECORDER	15	1	11111	Up	OSPFv3/area2
ROUTE-RECORDER -> 5.5.5.5	--	--	65535	Down	OSPFv3/area2
Multi-Access Link 1.1.1.1:0					
1.1.1.1 -> 1.1.1.1:0	3	3	50	Up	OSPFv3/Backbone
1.1.1.1:0 -> 1.1.1.1	3	--	0	Up	OSPFv3/Backbone
10.55.82.6 -> 1.1.1.1:0	--	LAN Pseudo-Node	65535	Down	OSPFv3/Backbone
1.1.1.1:0 -> 10.55.82.6	3	--	0	Up	OSPFv3/Backbone
ROUTE-RECORDER -> 1.1.1.1:0	--	LAN Pseudo-Node	65535	Down	OSPFv3/Backbone
1.1.1.1:0 -> ROUTE-RECORDER	3	--	0	Up	OSPFv3/Backbone
Multi-Access Link 3.3.3.3:0					
1.1.1.1 -> 3.3.3.3:0	12	12	40	Up	OSPFv3/Backbone
3.3.3.3:0 -> 1.1.1.1	12	--	0	Up	OSPFv3/Backbone
3.3.3.3 -> 3.3.3.3:0	12	12	67	Up	OSPFv3/Backbone
3.3.3.3:0 -> 3.3.3.3	12	--	0	Up	OSPFv3/Backbone

9 top level entries, 35 total entries — Export As Edits

Figure 35 List of All Links Window

2 Use the Filter by drop-down list at the top of the window to filter the links appearing in the window (see Using Filters on page 221).

3 Perform any of the following operations:

- Select a link in the list. This causes the link to flash yellow on the map.
- Copy a single row of the table by pressing **Ctrl+C**, or copy the entire table by pressing **Ctrl+A**. The data is copied to the clipboard, from which you can paste it into a text file. This operation captures all of the data from one or more rows.

- Export the table by clicking **Export**. This operation copies a subset of the data in each row in the format that is required for import into the Router/Link Edits window.
- If the routing topology has changed since the report was opened, refresh the contents by clicking **Refresh**.

Interfaces List

Open the Interfaces List to display the router interfaces discovered by the Collector (see the "Configuring the Route Recorder for the Collector" section in the *HP Route Analytics Management Software Administrator Guide*.

To view the interfaces list, perform the following steps:

1 Choose **Tools > List Interfaces** to open List of Interfaces window (Figure 35).

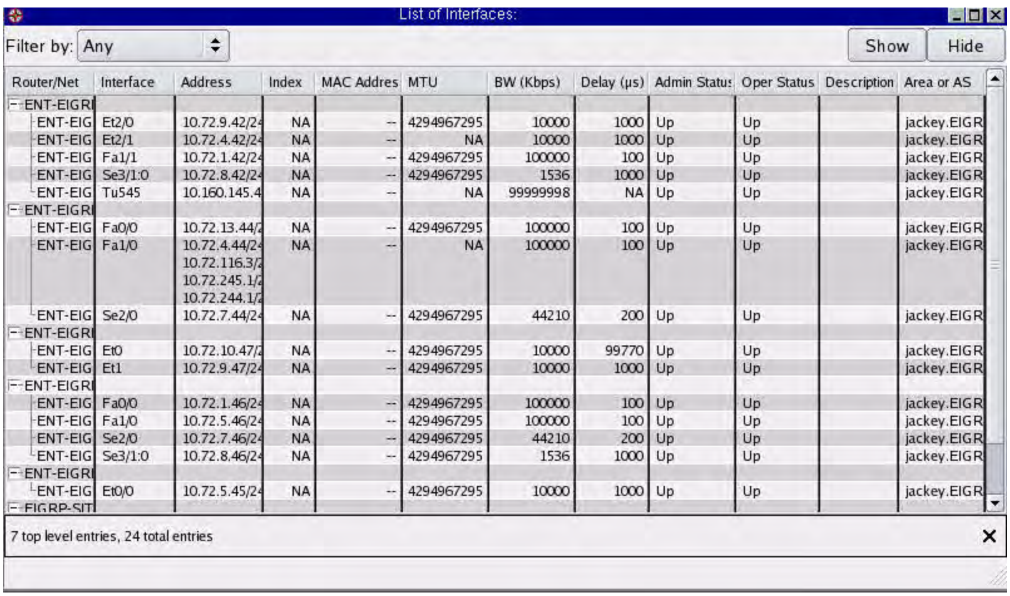

Figure 36 List of Interfaces Window

2 Use the Filter by drop-down list at the top of the window to filter the links appearing in the window (see Using Filters on page 221).

Prefix List

Open the List of IPv4 or IPv6 Prefixes window to obtain the following information:

- Determine if a prefix is currently advertised or not, and by which routers.
- See what metric is advertised with each prefix.
- Verify that area border routers (ABRs) are properly advertising prefixes.
- View the routers in a network that are advertising default routes.
- List the advertised BGP prefixes and the list of attributes associated with each BGP prefix.

To view the list of prefixes, perform the following steps:

1 Choose one of the Prefixes items from the Tools menu to open the List of All Links window (Figure 37).

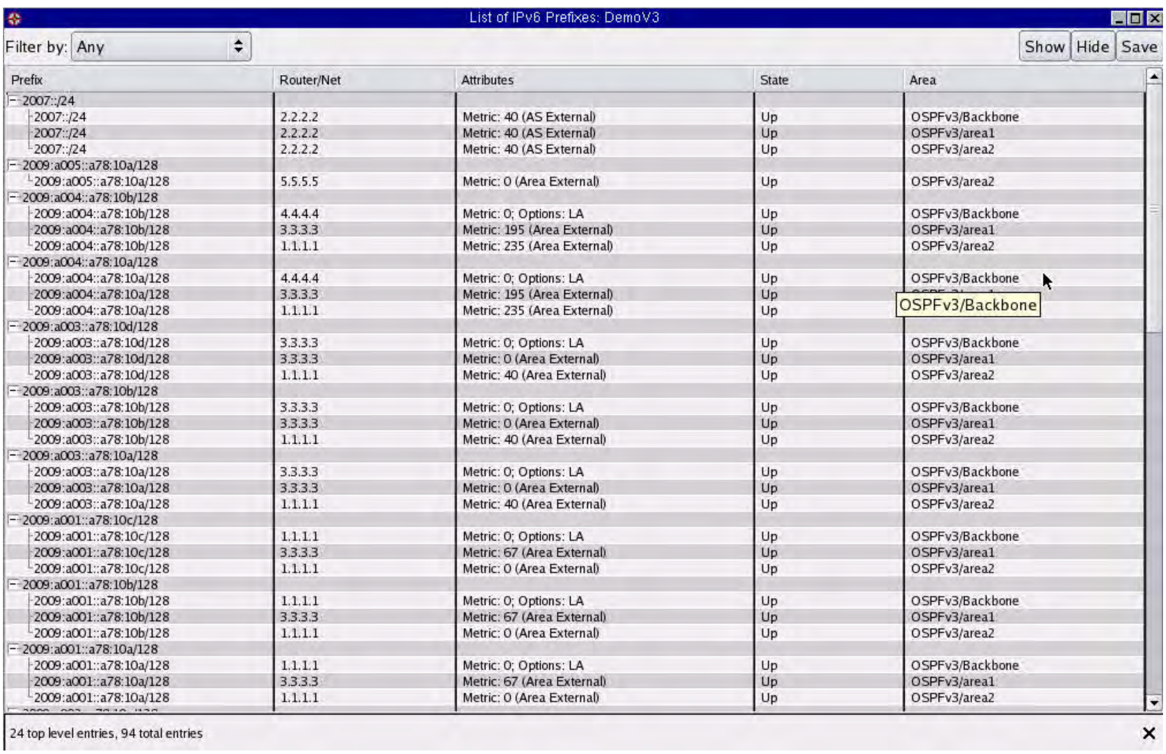

Prefix	Router/Net	Attributes	State	Area
2007::/24				
2007::/24	2.2.2.2	Metric: 40 (AS External)	Up	OSPFv3/Backbone
2007::/24	2.2.2.2	Metric: 40 (AS External)	Up	OSPFv3/area1
2007::/24	2.2.2.2	Metric: 40 (AS External)	Up	OSPFv3/area2
2009:a005::a78:10a/128				
2009:a005::a78:10a/128	5.5.5.5	Metric: 0 (Area External)	Up	OSPFv3/area2
2009:a004::a78:10b/128				
2009:a004::a78:10b/128	4.4.4.4	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a004::a78:10b/128	3.3.3.3	Metric: 195 (Area External)	Up	OSPFv3/area1
2009:a004::a78:10b/128	1.1.1.1	Metric: 235 (Area External)	Up	OSPFv3/area2
2009:a004::a78:10a/128				
2009:a004::a78:10a/128	4.4.4.4	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a004::a78:10a/128	3.3.3.3	Metric: 195 (Area External)	Up	[illegible]
2009:a004::a78:10a/128	1.1.1.1	Metric: 235 (Area External)	Up	
2009:a003::a78:10d/128				
2009:a003::a78:10d/128	3.3.3.3	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a003::a78:10d/128	3.3.3.3	Metric: 0 (Area External)	Up	OSPFv3/area1
2009:a003::a78:10d/128	1.1.1.1	Metric: 40 (Area External)	Up	OSPFv3/area2
2009:a003::a78:10b/128				
2009:a003::a78:10b/128	3.3.3.3	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a003::a78:10b/128	3.3.3.3	Metric: 0 (Area External)	Up	OSPFv3/area1
2009:a003::a78:10b/128	1.1.1.1	Metric: 40 (Area External)	Up	OSPFv3/area2
2009:a003::a78:10a/128				
2009:a003::a78:10a/128	3.3.3.3	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a003::a78:10a/128	3.3.3.3	Metric: 0 (Area External)	Up	OSPFv3/area1
2009:a003::a78:10a/128	1.1.1.1	Metric: 40 (Area External)	Up	OSPFv3/area2
2009:a001::a78:10c/128				
2009:a001::a78:10c/128	1.1.1.1	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a001::a78:10c/128	3.3.3.3	Metric: 67 (Area External)	Up	OSPFv3/area1
2009:a001::a78:10c/128	1.1.1.1	Metric: 0 (Area External)	Up	OSPFv3/area2
2009:a001::a78:10b/128				
2009:a001::a78:10b/128	1.1.1.1	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a001::a78:10b/128	3.3.3.3	Metric: 67 (Area External)	Up	OSPFv3/area1
2009:a001::a78:10b/128	1.1.1.1	Metric: 0 (Area External)	Up	OSPFv3/area2
2009:a001::a78:10a/128				
2009:a001::a78:10a/128	1.1.1.1	Metric: 0; Options: LA	Up	OSPFv3/Backbone
2009:a001::a78:10a/128	3.3.3.3	Metric: 67 (Area External)	Up	OSPFv3/area1
2009:a001::a78:10a/128	1.1.1.1	Metric: 0 (Area External)	Up	OSPFv3/area2

Figure 37 List of Prefixes Window for IPv6

2 Use the Filter by drop-down list at the top of the report to filter the prefixes appearing in the window (see Using Filters on page 221).

3 Perform any of the following operations:

- Select a router in the list, and the node will flash yellow on the map.
- If the routing topology has changed since the report was opened, refresh the contents by clicking **Reload**.

To list prefixes for a node, perform the following steps:

1 Locate the node on the routing topology map.

2 Right-click the node to open the node information panel.

3 Choose one of the Prefixes options.

The List of Prefixes window opens to show all prefixes that are advertised by the node you selected. For each prefix, all nodes that advertise the prefix are listed.

The names and addresses in the Router/Net column are the routers that are advertising the prefix. The way that the routers are listed depends on which protocol is in use. In OSPF, the pseudonode advertises the prefix of a LAN. In IS-IS, the designated router advertises the prefix of a LAN. For a point-to-point link there is no pseudonode, so both routers advertise the prefix. An EIGRP network does not have pseudonodes, so all prefixes are advertised by routers.

OSI Prefixes List

Open the List of OSI Prefixes window to obtain the following information:

- Determine if a prefix is currently advertised or not, and by which routers.
- Determine the metric that is advertised with each prefix.

To view the List of OSI Prefixes/ES Neighbors, perform the following steps:

1 Choose **Tools > List of OSI Prefixes** to open the List of OSI Prefixes window.

2 Use the Filter by drop-down list at the top of the report to filter the OSI Prefixes and ES Neighbors.

VPN Prefixes List

Open the List of VPN Prefixes window to obtain the following information:

- Determine if a prefix is currently advertised or not, and by which routers.
- Determine the metric that is advertised with each prefix.
- List the VPN routers in a network. (To list the routers, sort on the Router/Net column, right-click the column header and choose **Group**, and then right-click and choose **Collapse all**).

To view the list of VPN prefixes, perform the following steps:

1 Choose **Tools > List VPN Prefixes** to open the List of VPN Prefixes window.

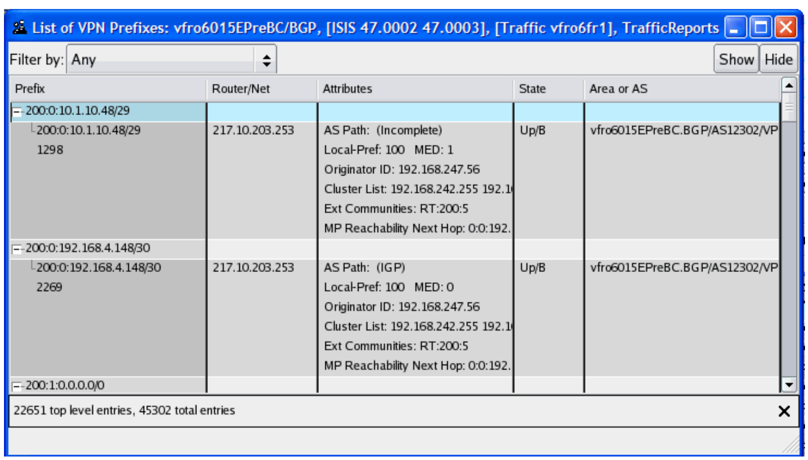

Figure 38 List of Prefixes Window

2 Use the Filter by drop-down list at the top of the report to filter the list.

Understanding Topology Groups

Topology groups are collections of network elements that are treated as a single unit in alert watchlists. A watchlist is a set of routers associated with a specific alert (see Creating New Alerts on page 508).

In addition to their role in alert watchlists, router groups allow you to simplify the topology map display by showing the routers in a group as a cloud, which you can operate on as a single entity or expand to display the individual routers.

Child groups are groups that are nested within another group.

See Administration on page 65 for a list of the group types that you can create from the Administration menu.

Example: The router group NewYork is created with routers A, B, and C, and the router group California is created with routers D, E, and F. With these groups in place, you can define alert watchlists that focus on each of these areas. If you want to create another watchlist that includes both of these groups, you can define the group UnitedStates and add NewYork and California as child groups.

Creating Groups on the Routing Topology Map

You can create groups directly on the routing topology map.

Be sure to save the routing layout before making any changes to the layout.

To create a router group directly on the topology map, perform the following steps:

1 Click, hold, and drag the cursor diagonally from an open area on the routing topology map that encompasses the routers you want to group. Release the mouse button.

   A bounding box appears around the routers and the selected routers change color, as shown in the following figure.

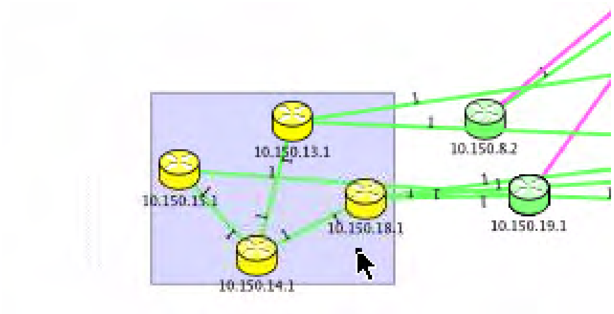

Figure 39 Router Group Bounding Box on the Topology Map

Only the routers completely within the bounding box are included in the counts shown in the Selection menu. Each direction of a link is counted separately.

2 With the cursor placed within the bounding box, right-click to open the Selection panel (Figure 40).

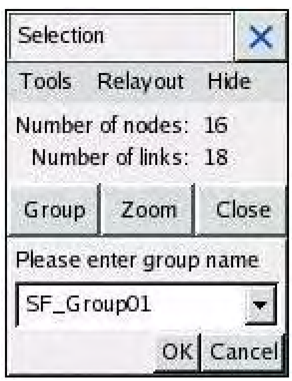

Figure 40 Selection Panel

3 Click **Group**.

The Selection panel expands to display a field for creating a group name.

4 Enter a name and click **OK**.

You can resize a group cloud that is open. Double-click the cloud, if necessary, to show the routers in the group. by (Figure 42). Click once to display a blue bounding box with a small rectangle in the lower right corner. Drag the rectangle to expand or shrink the cloud. Dragging horizontally or vertically changes the aspect ratio, while dragging diagonally preserves the aspect ratio.

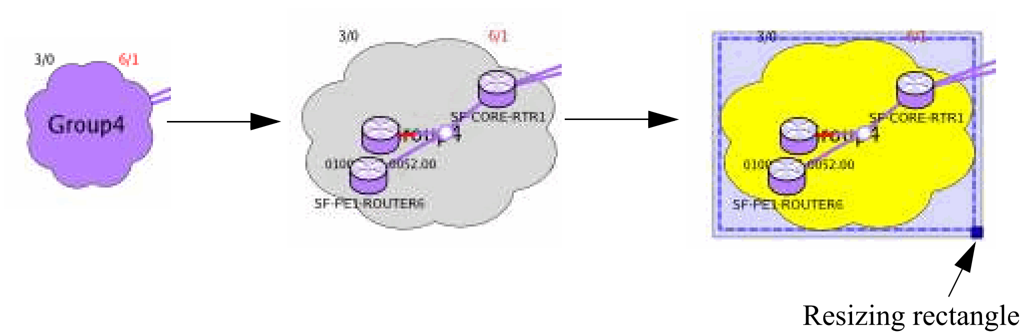

Figure 41 Enlarging a Group Cloud

Group Edit mode controls whether you can move nodes into and out of a group on the routing topology map. Group Edit mode is enabled automatically when you create a group on the map, and an edit icon appears in the lower right corner of the routing topology map window. Click the icon to leave Group Edit mode, and click the icon again if you want to reenter Group Edit mode.

You can also enable or disable the mode explicitly by choosing **View > Group Edit Mode** to toggle the check mark on or off.

When group edit mode is enabled, you can move nodes in and out of open groups and use the Ungroup and Destroy buttons in the Group information panel. When group edit mode is disabled, you cannot move nodes in and out of open groups or use the Ungroup and Destroy buttons in the Group information panel. If you try to do move nodes into or out of a group, the nodes are bounced back to their original positions.

In addition to creating a single group on the topology map, you can:

- Create multiple router groups.
- Nest a group a within another group.
- Include a router in multiple groups.

- View the routers within a group.

You can also create a group that contains nodes that are not adjacent to each other.

To select nodes for a group individually, perform the following steps:

1 Click any node on the routing topology map.

A bounding box is created around the node.

2 Press and hold the Ctrl key and then click other nodes. The bounding box expands to contain the additional nodes.

3 With the elements selected, right-click and choose **Group**.

4 Enter a name and click **OK**.

The selected area changes to a cloud. Links and element locations will change on the topology map to accommodate the disparate nodes.

To view the routers within a group on the topology map, perform the following steps:

1 Double-click a grouping cloud.

The cloud expands to show its contents, as shown in the following figure.

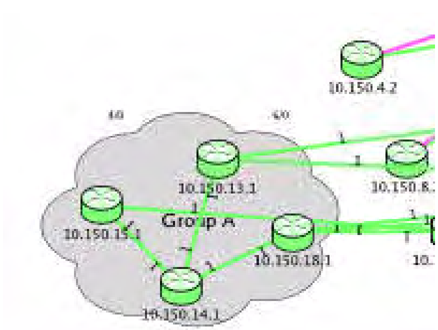

Figure 42 Viewing Elements Within a Cloud

2 Right-click the cloud to open the cloud panel, as shown in (Figure 43):

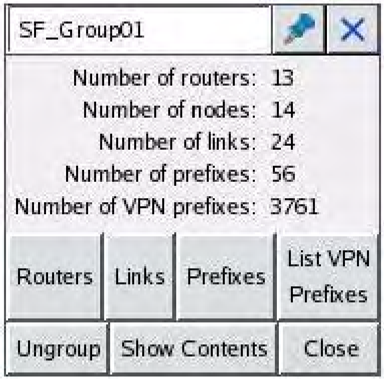

Figure 43 Cloud Panel

With the cloud expanded, you can:

- Move elements to different positions within the cloud by clicking and dragging elements within the cloud.
- Drag elements out of the cloud to remove them from the group or drag routers in to extend the group.
- Drag elements from one cloud to another to change the group membership or visibility. If an element is in multiple router groups, you can drag it from one group to another. When you drag an element from one group to another, the system prompts you to choose whether to remove the element from the first group (answer **Yes** to the prompt) or to leave the element in both groups but change the visibility to the second group (answer **No** to the prompt).

The general status of network elements within a cloud is indicated by the numbers that appear above it. By gliding the cursor over a cloud you can also obtain pop-up statistics that include the group name and a breakdown of the included elements with the up or down status (Figure 44).

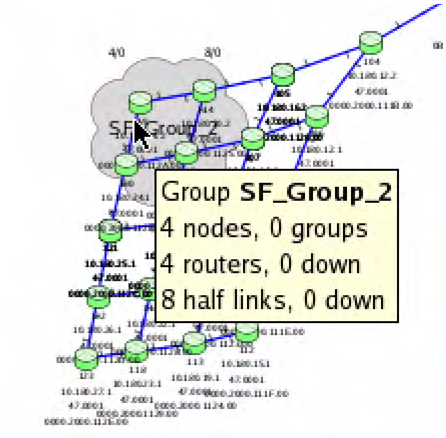

Figure 44 Status in Cloud

3 Perform any of the following operations on the cloud:

- **Routers**—Click the **Routers** button to open the List of Routers in Group report (Figure 45).

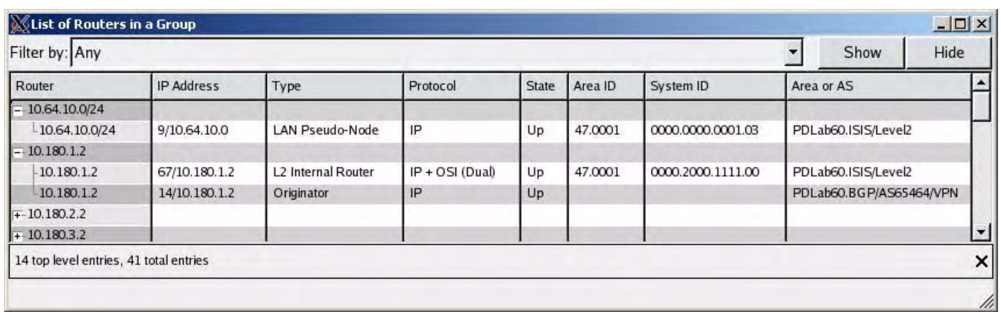

List of Routers in a Group

Filter by: Any — Show — Hide

Router	IP Address	Type	Protocol	State	Area ID	System ID	Area or AS
− 10.64.10.0/24							
10.64.10.0/24	9/10.64.10.0	LAN Pseudo-Node	IP	Up	47.0001	0000.0000.0001.03	PDLab60.ISIS/Level2
− 10.180.1.2							
10.180.1.2	67/10.180.1.2	L2 Internal Router	IP + OSI (Dual)	Up	47.0001	0000.2000.1111.00	PDLab60.ISIS/Level2
10.180.1.2	14/10.180.1.2	Originator	IP	Up			PDLab60.BGP/AS65464/VPN
+ 10.180.2.2							
+ 10.180.3.2							

14 top level entries, 41 total entries

Figure 45 List of Routers in a Group

- **Link** —Click the **Links** button to open the List of Links in Router Group report (Figure 46).

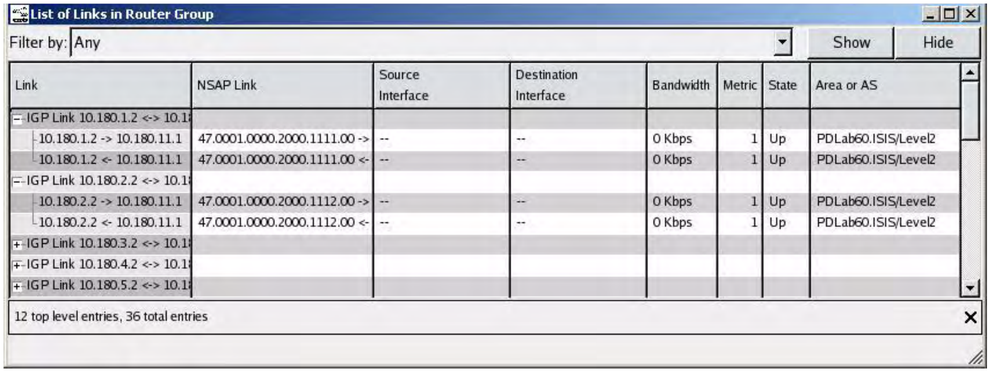

Figure 46 List of Links in a Group

- **Prefixes**—Click the **Prefixes** button to open the List of Prefixes in Router Groups report (Figure 47).

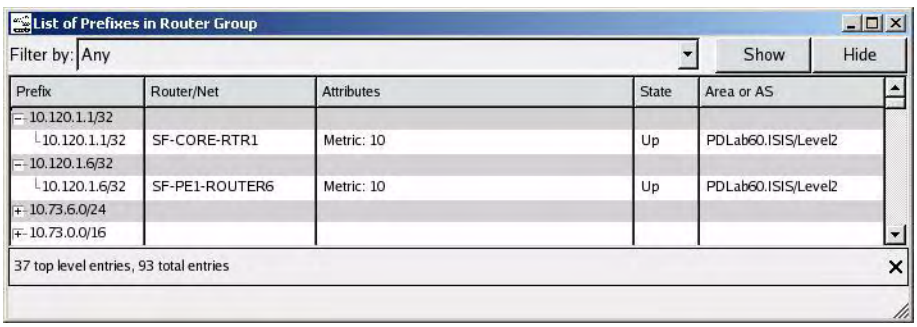

Figure 47 List of Prefixes in Router Groups

- **List VPN Prefixes**—Click the **VPN Prefixes** button to open the List of VPN Prefixes for Router Group report (Figure 48).

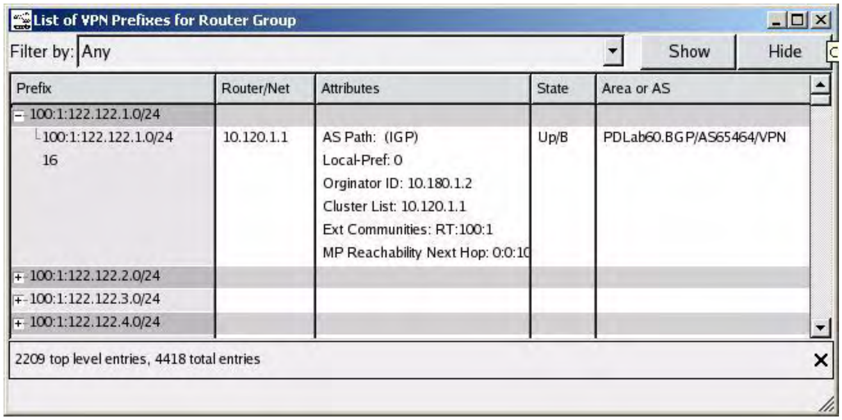

Figure 48 List of VPN Prefixes for Router Group

- **Ungroup**—Remove the grouping on the map or delete (destroy) the group. When you ungroup the cloud on the map, the elements retain a group identity in the Router Groups report but the cloud does not appear on the topology map. When you delete the group, it is removed from the Router Groups report and the topology map.
- **Show Contents**—View the network elements within the cloud. This is the same as double-clicking the cloud.
- **Close**—Close the cloud panel. Double-click the expanded cloud to close it from the network element view.
- **Hide Contents**—Close the information panel for the opened group.

Creating Groups Using the Menu

You can create router groups, link groups, path, or prefix groups from the Administration menu. Groups are useful in configuring alerts, applying filters, and focusing in on specific areas within the topology map.

The following group types are accessible from the Administration menu in the client application:

- Router Groups
- Link Groups

- Prefix Groups/IPv4 Prefix Groups/IPv6 Prefix Groups
- Path Groups
- VPN Customer Groups
- VPN RT Groups
- VPN Prefix Groups
- VPN Site Groups

A router can belong to multiple groups, both of which are on the routing topology map. However, only one of the groups can be visible at a time. You can specify the group that will display the individual router.

The following options are supported for most of the group types (see Working with Groups on page 125):

- **New Group**—Create a new group.
- **Edit Group**—Modify an existing group.

  You cannot modify groups that the system has automatically created, including area groups and MPLS WAN sites. on MPLS WAN sites, see MPLS WAN on page 479.

- **Copy**—Make a copy of the group or group element to use in another group.
- **Move**—Move the group or group element to another group.
- **Delete**—Remove a group or group element.
- **Show area groups (router groups only)**—Create a router group with a structure that matches the topology hierarchy. See Understanding the Topology Hierarchy on page 69 for information on the topology hierarchy.
- **Show site groups (router groups only)**—Display the site groups on the map. See MPLS WAN on page 479.

- **Tooltips**—Move the cursor over a group name to view details about the group (Figure 49). Tooltips are available in the group window and in the dialog box that opens when you move a group or group element.

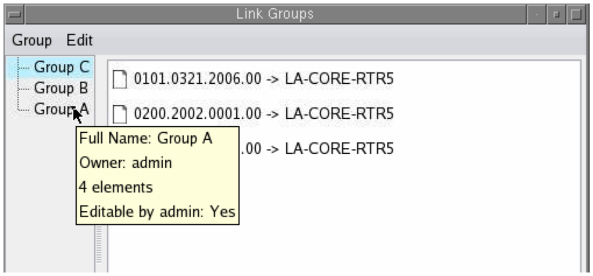

Figure 49 Group Tooltip

- **Router group checkbox** (router groups only, see Figure 50)—Select the check box for a router group to display that group on the map.

Click the docking icon to dock or undock the Router Groups window in the routing topology maps window.

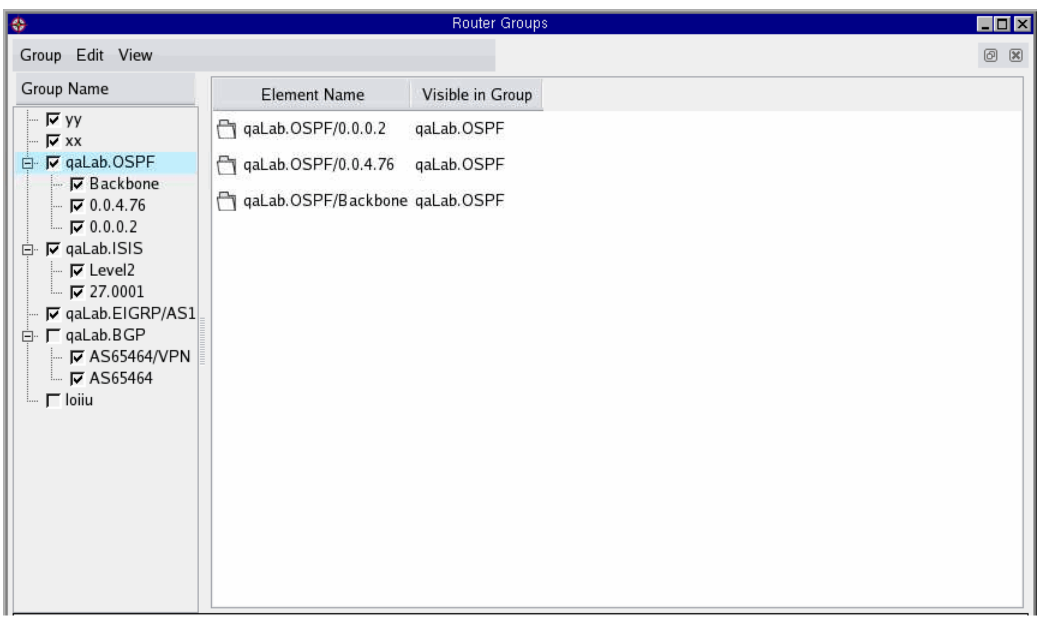

Figure 50 Router Groups

Working with Groups

You can create, edit, copy, move, or delete groups, and show all of the area or site groups.

Changes that you make to groups are reflected immediately on the routing topology map. The system prevents you from making changes if you do not have permission to do so.

To create a new group, perform the following steps:

1 Choose **Administration** and select one of the Groups menu items. If a group does not exist, a pop-up window indicates how to create a new group. Click **OK** to close the pop-up.

2 If a group of this type does not already exist, a pop-up window indicates how to create a new group. Click **OK** to close the pop-up.

3 Choose **Group > New Group**.

4 Enter a group name, as shown in Figure 51.

5 If the IPv6 is supported, choose the type of prefix (IPv4, IPv6, or NSAP).

6 For most of the group types, the first tab that opens shows a list of IP addresses or names on the left. Highlight the items that you want to add to the group, and click **->** to move them to the selection area on the right (Figure 51).

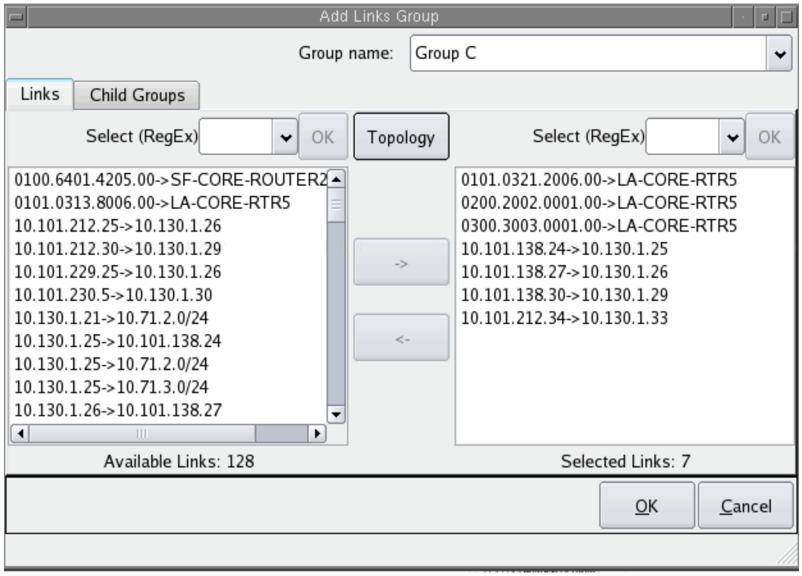

Figure 51 Creating a New Group

For path groups, enter the starting and terminating IP addresses for each path, and click **Save** to add the path to the group (Figure 52).

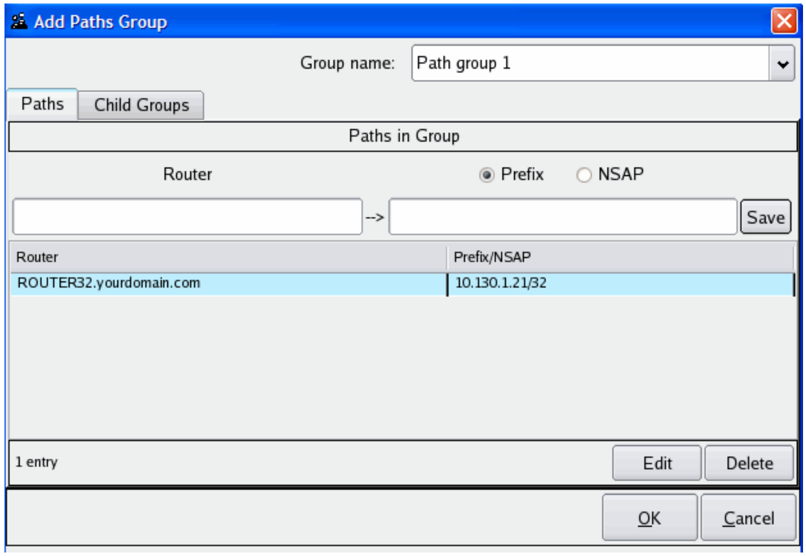

Figure 52 Creating a New Path Group

7 If you want to select only specific members of the list, enter the common characters in the Select (RegEx) field and click **OK** to the right of the field. For example, if you want only the IP addresses that start with 192, enter 192 in the Select (RegEx) field, click **OK** to the right of the field, and the click the arrow to move the selected IP addresses from the Available Items field to the Selected Items field.

➤ The syntax of extended regular expressions is explained in Regular Expressions on page 226. The syntax is not the same as shell or file manager pattern patching, so a pattern like *-core-gw is not correct.

8 Click the **Child Groups** tab.

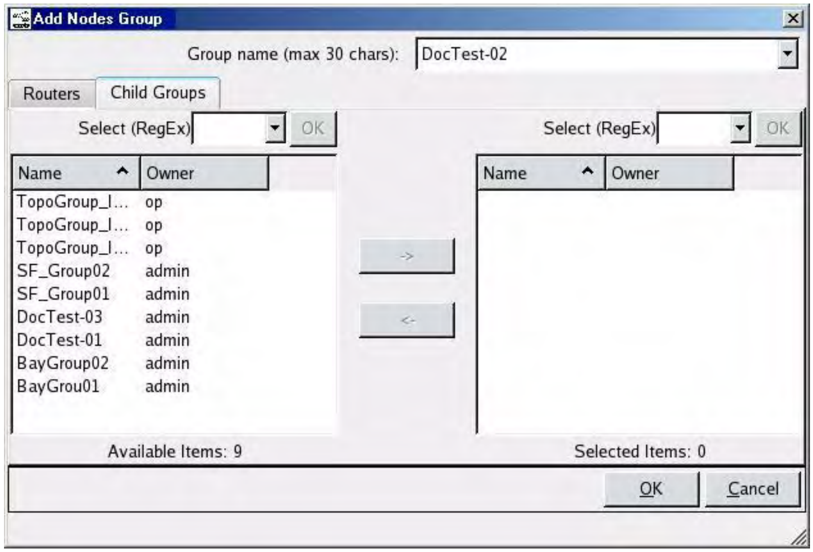

Figure 53 Child Groups

9 The Child Groups tab lists items associated with the group type. If you want to select only specific members of the list, enter the common characters in the Select (RegEx) field and click **OK** by the field.

For example, if you only want IP addresses starting with 192, enter **192** in the Select (RegEx) field, click **OK** by the field, and the click the arrow to move the selected IP addresses from the Available Items field to the Selected Items field.

10 After adding the appropriate child groups, click **OK** to save your settings.

To edit a group, perform the following steps:

You cannot modify groups that the system has automatically created, including area groups and groups based on MPLS WAN sites. The editing menu options are disabled if you do not have permission to edit a group.

1 Choose **Administration** and select one of the Groups menu items to open the group window.

2 Right-click the group name and choose **Edit Group**.

3 Set up the group elements as you would when adding a new group. See Creating Groups Using the Menu on page 122.

4 Click **OK**.

To make an element visible in a specified router group, perform the following steps:

1 Choose **Administration > Groups > Routers**.

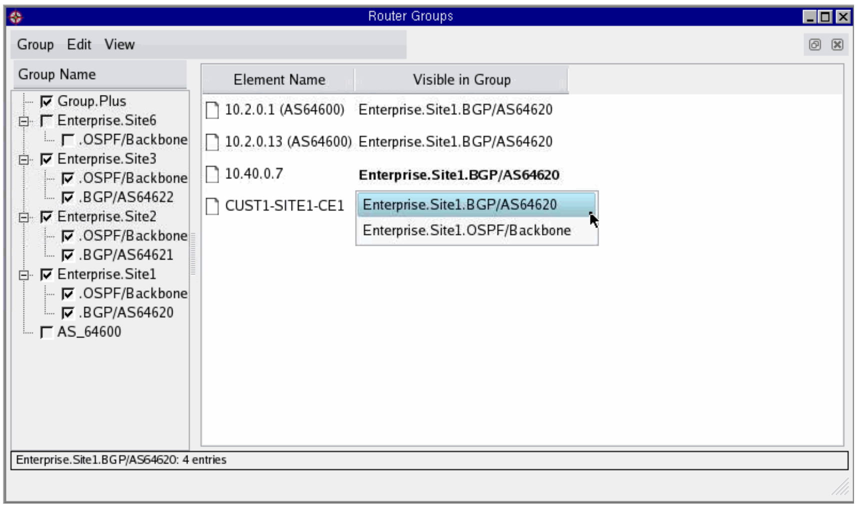

Figure 54 Changing Router Group Visibility

2 Select the group of interest. You can expand or collapse the hierarchy view on the left site of the window by choosing **View > Expand All** or **View > Collapse All**.

The group elements are displayed on the right. The Visible in Group column displays the group in which the element currently appears on the routing topology map. If the entry is in **bold** text, the element is also a member of one or more additional groups that is visible on the map.

3 To change the group in which a member is visible, double-click the bold entry to display selection arrows. Click again and select the desired group.

The element is displayed in the new group.

To move or copy a group within the group hierarchy, perform the following steps:

1 Choose **Administration** and select one of the Groups menu items to open the group window.

2 Use one of these options to move the group:

- Right-click the group you want to move and choose **Move** or **Copy**. In the dialog box that opens, choose the group that will include the selected group as a subgroup or choose **Top Level**, and click **OK**. The group is moved or copied to the new location (Figure 55).

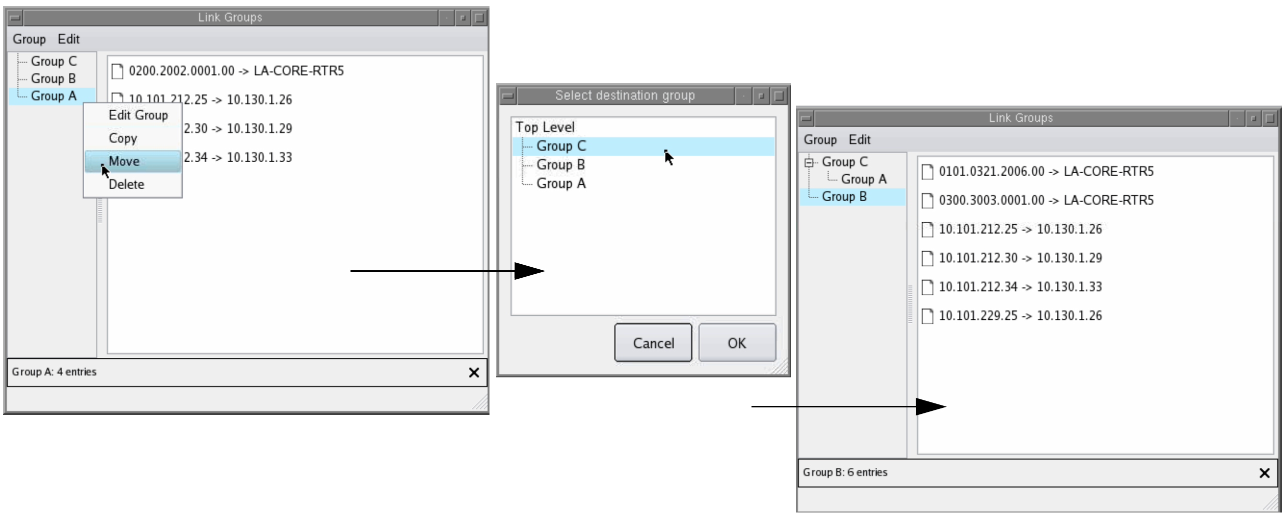

Figure 55 Moving Groups

- Drag and drop the group to a new location in the hierarchy.

The system prevents you from dragging and dropping an item if you do not have permission to do so.

To copy or move a group element within the group hierarchy, perform the following steps:

1 Choose **Administration** and select one of the Groups menu items to open the group window.

2 Choose one of these options to copy or move an individual group member (such as an IP address) to another location:

- Right-click the item you want to move and choose **Move** or **Copy**. In the dialog box that opens, choose the group that will include the selected item, and click **OK**. The item is copied or moved to the new group.
- Drag and drop the item to a different location in the hierarchy.

To delete a group or group element, perform the following steps:

Only the group owner or an administrator can delete a group.

1 Choose **Administration** and select one of the Groups menu items to open the group window.

2 Select the element or group, right-click, and choose **Delete**.

3 Click **OK** to confirm.

Removing a router from a group does not delete the router from your topology. You can add a router

To create a single router group that matches the structure of the topology hierarchy, perform the following steps:

1 Choose **Administration > Groups > Router**.

2 Choose **Group > Group by Area**.

3 Enter a name and click **Apply**.

A new router group is created with a structure that matches the topology hierarchy.

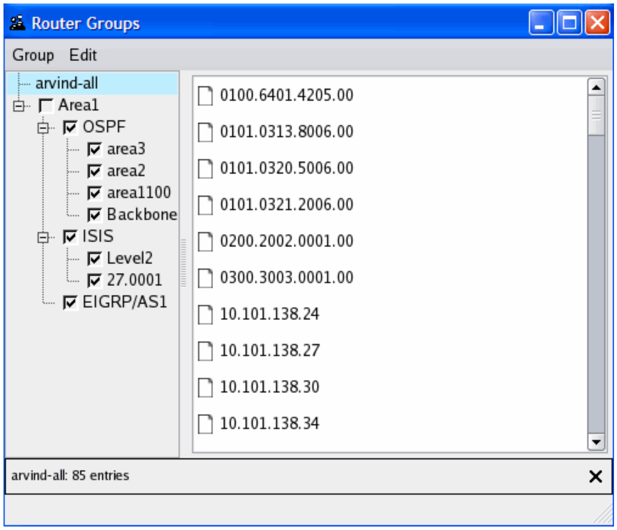

Figure 56 Grouping by Area

Hiding Forwarding Adjacencies in Link Groups

Some networks contain routers that include forwarding adjacency (FA) configuration. FAs allow administrators to specify LSP tunnels as links in the IGP network. You can create a forwarding adjacency between routers regardless of their location in the network, and routers with forwarding adjacency can be multiple hops apart.

FAs may crowd the routing topology map, making it more difficult to focus on other areas of interest. To simplify the routing topology display, you can add all the FAs to a single link group and then hide the group when you want to focus on other areas of interest.

To create a link group containing all the FAs in the network, perform the following steps:

1 Choose **Administration > Groups > Links**.

2 Choose **Group > New Group**.

3 Enter a name for the group.

4 Click **Select FA t**o highlight all of the FAs in the Links list.

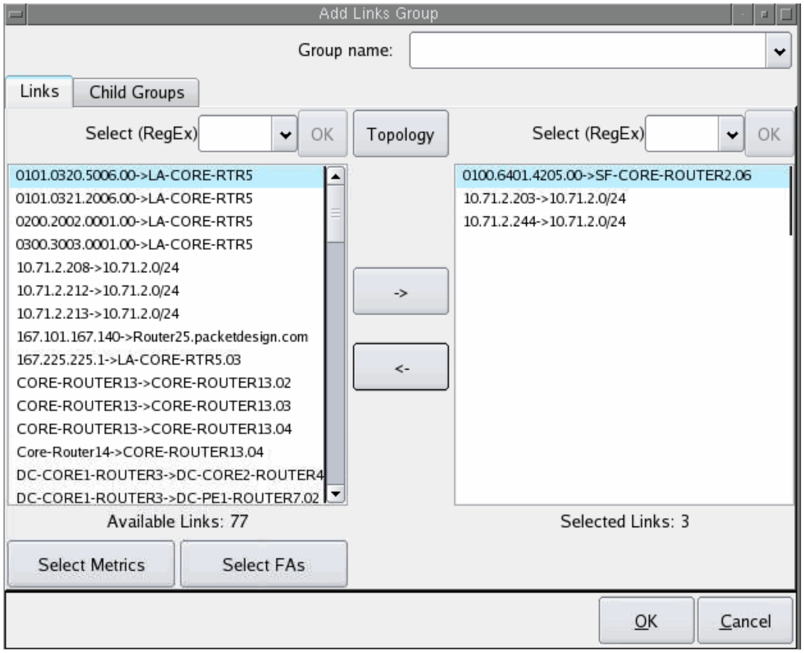

Figure 57 Link Group - Select FA Option

5 Click the right facing arrow to move the highlighted link to the selection area.

6 Click **OK**.

7 A new group is now created containing all the FAs. To hide the group on the routing topology map, right-click the group cloud and select **Hide**.

Understanding Network Routes

This section describes client application support for viewing and managing routes through the network:

Highlighting the IP Route Between Two Points in the Network

Viewing the IP paths taken by traffic from a source router to a destination router in a multidomain network is useful for network planning. You can identify paths for which it is critical to avoid delays caused by rerouting due to router failures, such as a voice over IP (VoIP) service path between an IP private branch exchange (PBX) and a Public Switched Telephone Network (PSTN) media gateway.

The system can quickly show the path that is resolved between two nodes in a network at the current time or at any point in recorded topology history.

Paths are shown in the following colors:

- Forward path: Green
- Reverse path: Orange

Nodes are shown in the following colors:

- Unidirectional: source: Green, destination: Yellow
- Bidirectional: source: Green, destination: Orange

You can use Find or List Paths window to list each segment of the path along with the link metric and the prefix by which each next hop was resolved. See Finding a Route By Prefix on page 136.

Select the path using the information panels for the source and destination nodes on the routing topology map. Figure 58 shows the node information panels for the source and destination of a route.

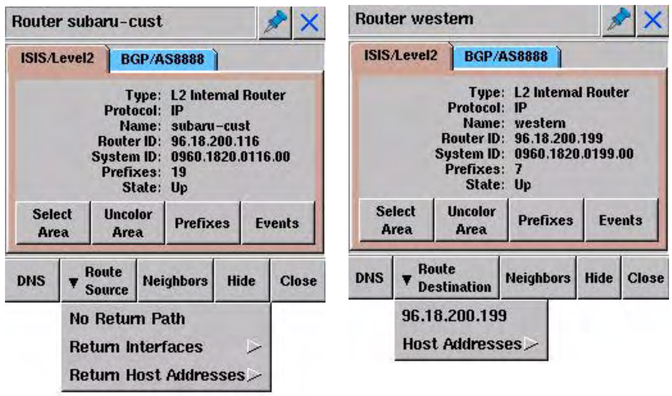

Figure 58 Source and Destination Node Information Panels

To view the path between two routers, perform the following steps:

1 Right-click the source node on the routing topology map to open the node information panel.

2 Click **Route Source**, and select one of the following from the drop-down list:

- **No Return Path—**Provides route source only.
- **Return Interfaces—**Specifies return interface for route source.
- **Return Host Addresses—**Provides the host address.

The route from one router toward another router is highlighted in yellow on the routing topology map. The return route is highlighted in orange. IPv6 addresses are shown, if available.

3 Right-click the destination node on the routing topology map.

4 Click **Route Destination** in the node information panel that opens, and then select an interface from the drop-down list, or select **Host Addresses**.

The route from one router toward another router is highlighted in yellow on the routing topology map. The return path is highlighted in orange. IPv6 addresses are shown, if available.

5 To see the path details, click **List/Find Paths**.

The route may not be complete between the two points if the destination address falls within a prefix that is not routable in the topology known to the system, or if the address resolves to a summary prefix such as the default route. Consequently, the route from point B to point A might be incomplete even if the route from point A to point B is complete.

Finding a Route By Prefix

In addition to finding paths between pairs of nodes on the routing topology map, the system can also find the route from a router to any prefix internal or external to the network for which a route exists.

To find a route using a prefix, perform the following steps:

1 Choose **Tools > List/Find Paths** to open the The Find or List Paths window.

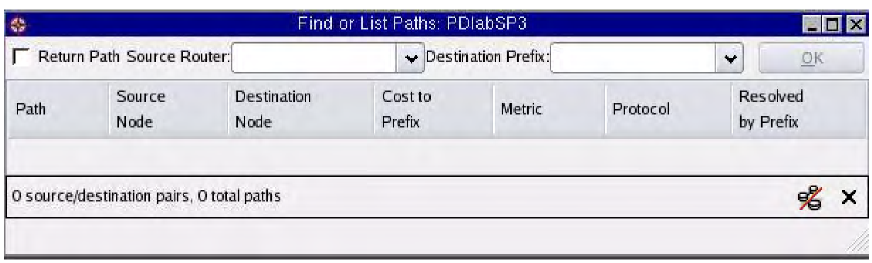

Figure 59 Find or List Paths

2 Enter the IPv4 or IPv6 address or System ID of the source router or name in the Source Router field.

3 Enter the destination IPv4 or IPv6 address, Internet prefix, or domain name in the Destination Prefix field.

If IPv6 is supported and a domain name is entered for the destination, the DNS lookup requests an IPv6 address if the source is specified as an IPv6 address or if the source router identified by name does not support IPv4. Otherwise, the DNS lookup requests an IPv4 address.

4 Click **OK**.

The route is calculated to the destination prefix if it is internal or to the nearest exit router if it is external. The segments of the path are listed in the lower section of the Find or List Paths window. The listing includes the link metric, the router's calculated cost to reach the destination prefix, and the prefix by which each next hop was resolved. The forward path is highlighted with green arrows on the routing topology map, and the return path is highlighted with orange arrows.

A path flashes momentarily on the topology map when you select its corresponding entry in the Find or List Paths window.

Multiple paths are shown for equal-cost multipath routes.

If you enter a destination prefix that does not exist in the network, the route might go to the default router and the default router might forward the route to a router outside the topology. In this case, the path may end at a LAN pseudonode.

Finding a VPN Route By Prefix

Open the Find/List VPN Paths window to find VPN paths in the topology.

To find VPN paths, perform the following steps:

1 Choose **Tools > List/Find VPN Paths** to open the Find or List VPN Paths window.

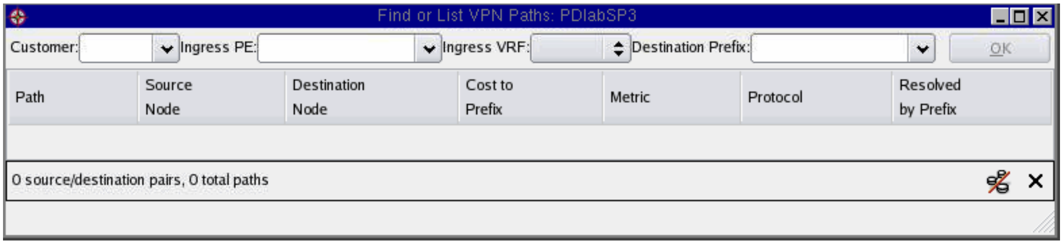

Figure 60 Find or List VPN Paths

2 Enter the customer name in the Customer field. The customer name should be an existing Customer to RT mapping (as defined by **Administration > VPN Customer -RT Mapping Configuration**). The customer name is case-sensitive.

3 Enter the IP address of the ingress PE in the Ingress PE field.

4 Select the desired ingress VRFs.

This step lists all existing VRFs on the selected ingress PE belonging to the specified customer. If the Ingress VRF field remains empty after the Customer and Ingress PE fields are populated, its means that no VRFs on the Ingress PE belong to that customer.

5 Enter the desired prefix of the selected VPN in the Destination Prefix field in IP4 format.

6 Click **OK**.

During route resolution only VPN routes with RTs that match the RT import policy of the Ingress VRF are considered. If a path is found, the segments of the path are listed in the lower section of the Find or List VPN Paths window. The listing includes the link metric, the router's calculated cost to reach the destination prefix, and the prefix by which each next hop was resolved. The paths will flash momentarily on the topology map when its corresponding entry in the Find or List Paths is selected.

Viewing the Highlighted Path Cost for EIGRP

To view the details of a highlighted path, choose **Tools > List/Find Paths**. Figure 61 shows an example of how highlighted paths appear on the topology map, combined with the associated path details, for an EIGRP network. The path source is shown in green, the destination is shown in yellow, and green arrows trace the routes through the network.

The Cost to Prefix column shows the cost that each router along the path calculates to reach the prefix, for applicable protocols, along with one value for the entire path showing what the source router of the path calculates as its cost to reach the prefix.

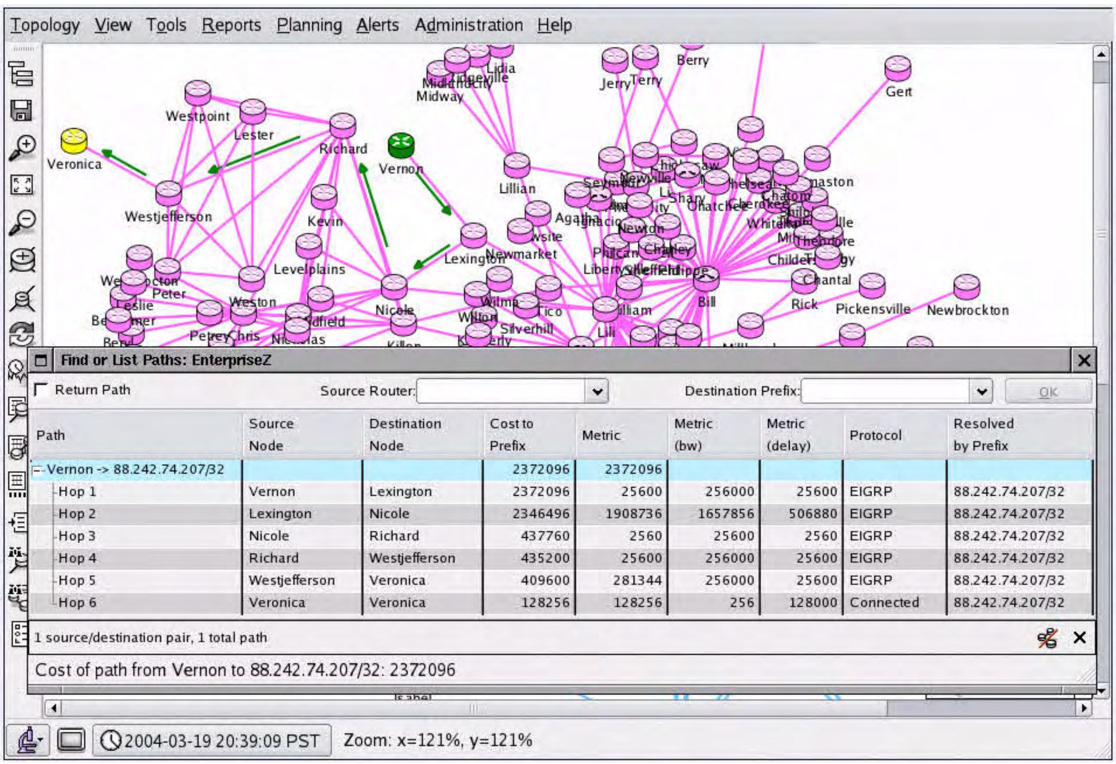

Figure 61 Highlighted Paths

The Metric (bw) and Metric (delay) columns for EIGRP show the cost in EIGRP metric units for the bandwidth and delay values that are configured for the link. The Metric column lists the amount that each link contributes to the overall path distance, or cost. The total path cost for EIGRP is the sum of the delay values for each hop plus the maximum of the bw values.

Because the EIGRP protocol calculates routes from the destination back towards the source, read the Metric column from bottom to top to follow the route from source to destination.

In Figure 61, hop 6 contributes both the bw and delay values to the total cost of 128256. Hop 5 adds the amount by which its bw value is larger than the current maximum of the bw values (256000-256) plus its delay value 25600, for a total of 281344. At hop 4 the bw equals the maximum value, so hop 4 increases the total cost only by its delay value 25600. Hop 3 also adds only its delay of 2560 because its bw value 25600 is smaller than the current maximum. At hop 2, the maximum bw value increases again (1657856-256000) and a delay of 506880 is added. Finally, hop 1 adds just its delay value of 25600 for a total of 2372096.

The Cost to Prefix column shows the distance each hop reported to reach the prefix shown in the Resolved by Prefix column. If the path is resolved using the same prefix at all hops, then the Cost to Prefix value for a hop equals the sum of the Metric column values from that hop to the destination, and the sum of all the values in the Metric column equals the total path cost that is reported in the status bar of the Find or List Paths window.

If there is a change of prefix along the path, for example due to prefix summarization, then the sum of the metrics may be higher or lower than the Cost to Prefix, depending on the distances to the various prefixes as assigned by the routers. It this case, it may not be possible to calculate a valid total cost for the path.

In a multiprotocol network, it is not always possible to calculate the total path cost, because the metrics of different protocols are of different magnitudes. In such cases, the status line indicates that the path cost cannot be calculated.

Diagnosing EIGRP Topology Errors

Topology Diagnostics are available only for EIGRP topologies. Diagnostics allow you to study problems found in the network configuration or in the topology modeling. Choose **Tools > Topology Diagnostics** and select one of the following options:

- **List Topology Errors**—Open a list of messages describing anomalies that were discovered during exploration of the EIGRP topology. Click an entry in the list to highlight the affected routers and links. See List Topology Errors on page 141.
- **List Inaccessible Routers**—Open a list of routers that were not accessible through Telnet during exploration of the EIGRP topology. The list includes a reason, such as authentication failure. Click an entry in the list to highlight the last accessible router on the path toward the inaccessible router. Inaccessible routers are not shown on the topology map. They can cause incorrect routes to appear and reduce your ability of to track changes in the network topology. See List Inaccessible Routers on page 143.
- **List Mismatched Distances**—Open a list of prefixes for which the distance to the prefix reported by a router that peers with the appliance does not match the distance that the system calculates across its model of the topology. This list also shows when the periodic topology explorations start and end, along with summary statistics. See List Mismatched Distances on page 145.
- **Find Invisible Links**—Run a simulation on the topology model to determine if there are any links where a failure will not be immediately detected because the routers that peer with the system will not report an EIGRP distance change. The simulation can take several hours to run on a large network topology. You can cancel it at any time. See Find Invisible Links on page 148.

List Topology Errors

The Route Recorder detects configuration anomalies as it collects information from the routers during its initial exploration of the topology and subsequent periodic re-explorations. These anomalies are stored in the database and shown in the List of Topology Errors report. The report shows only the anomalies that were detected since the start of the last full exploration. These anomalies can indicate router configuration errors that should be corrected.

Figure 62 shows an example of the List of Topology Errors report.

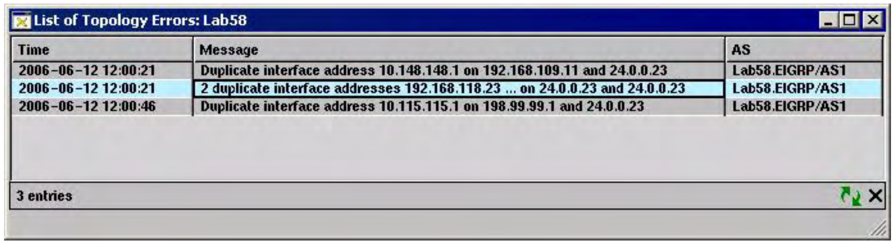

Figure 62 List of Topology Errors Report

The report contains the following columns:

- **Time**—the time when the anomaly was detected.
- **Message**—a description of the problem, which may be any of the following:
  - **Interface mask length mismatch**—Indicates that the address mask length is not the same for the interfaces on the two ends of a link.
  - **Duplicate router ID**—Indicates that two routers are using the same router ID for the EIGRP routing process. The router ID is usually taken from the IP address of a loopback interface or other interface on the router, and should be unique for each router.
  - **Router ID is an interface address on another router**—Indicates that the router ID on one router is the same as an interface address on another router. Because the router ID is normally derived from an interface address on the router and interface addresses are normally unique to one router, this is an anomaly.
  - **Duplicate interface address**—Indicates that one or more interfaces on one router have the same IP addresses as interfaces on another router. The two routers are highlighted.
  - **Potential redistribution error**—Indicates that the prefix is configured for redistribution but the metric is not configured. Applies to the case where an external prefix is advertised by a router but is unreachable from that router.
  - **Variance not supported by the system**—Indicates that the router is configured for equal-cost multi-path routing with a variance value other than one. Only the paths with the lowest metric are included.

— **Router ID unroutable in this AS**—Indicates that the router ID of the indicated router is not an address within any routable prefix in the AS. If the router-to-router path highlighting function is used with this router as the destination, the path will be incomplete.

— **Prefix with delay of 0**—Indicates that the delay component of a prefix metric is zero. This condition can be caused by connected or static routes being redistributed without explicitly specifying the delay. This can result in a routing loop.

— **Routing loop**—Indicates that the Route Recorder can discover a routing loop during topology exploration or while investigating routing changes. Occasionally, a routing loop can persist until manual intervention is taken. The Time column indicates when the loop was detected. Use the cursor in the History Navigator window to show the routing topology at that time, and then click **Events** to look in the All Events list for events related to the prefix that is looping. Also try highlighting the path from one of the peer routers to that prefix. See The History Navigator on page 155.

- **AS**—Includes the AS where the anomaly was detected, for topologies with multiple ASs.

Clicking the table row for an error message highlights the associated routers and links, assuming that those objects are present in the topology at the current time. You can use the History Navigator window to change the current time back to the time of the error message and then use other tables to diagnose the problem. See The History Navigator on page 155.

List Inaccessible Routers

During the EIGRP topology exploration, the system attempts to establish a Telnet/ command line interface (CLI) connection to each router to collect information about neighbors, interface metrics, and external prefix attachments. If the connection to a router fails, the system cannot include that router in the topology, nor learn about other routers connected beyond that router. If a path between two accessible routers passes through an inaccessible router, the system is not able to find that path. It is important to fix router accessibility problems so that the topology is correct.

Figure 63 shows an example of the List of Inaccessible Routers report.

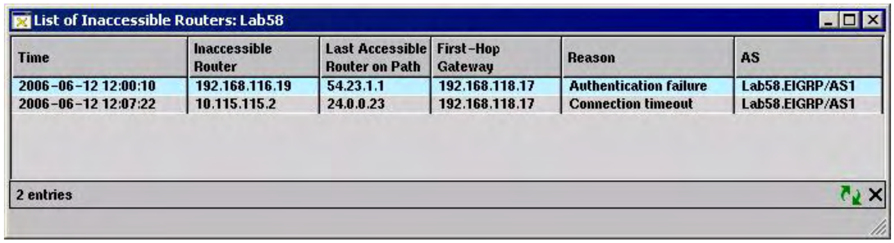

List of Inaccessible Routers: Lab58

Time	Inaccessible Router	Last Accessible Router on Path	First-Hop Gateway	Reason	AS
2006-06-12 12:00:10	192.168.116.19	54.23.1.1	192.168.118.17	Authentication failure	Lab58.EIGRP/AS1
2006-06-12 12:07:22	10.115.115.2	24.0.0.23	192.168.118.17	Connection timeout	Lab58.EIGRP/AS1

2 entries

Figure 63 List of Inaccessible Routers Report

The report contains the following columns:

- **Time**—Indicates the time when the problem was detected.
- **Inaccessible Router**—Identifies the inaccessible router.
- **Last Accessible Router On Path**—Indicates the address of the last router on the path to the inaccessible one. Clicking on the entry in this column highlights that router on the topology map.
- **First-Hop Gateway**—Indicates the gateway (first-hop router) used in attempting the connection if the solution is to specify a different default gateway.
- **Reason**—Indicates a possible reason for failure to access the router:
  - **Authentication failure**—Occurs if the router does not accept the login password or user name/password configured in for the AS.
  - **Invalid input (unauthorized command?)**—Occurs if the Terminal Access Controller Access-Control System (TACACS) account used by the system is not authorized to use one of the commands needed for topology exploration. This error could also occur due to garbled communication.
  - **Connection refused (no VTY?)**—Occurs if too many other Virtual Telnet (VTY) sessions are open at the time the system attempts its connection. No free virtual terminal is available on the router, and the connection is refused. The system attempts several connections with exponentially increasing delay. This error can also occur if the connection is blocked by a firewall or other appliance.
  - **Connection timeout**—Indicates that the router is unreachable and timed out, for example, if the path to the router hits a black hole.

— **Telnet open failed—**Presents additional details, if provided by Telnet.

— **CLI parsing error—**Indicates that the output of the commands issued to the router in a query was not formatted as expected. Report this problem to Technical Support.

— **Problem in recorder—**Indicates that the system was unable to issue the query for some reason. Please report this problem to technical support.

- **AS—**Indicates the AS in which the router resides.

By default, the table is sorted in descending order by time. To sort the table on any other field of information, click the column heading. To change the sort order (descending versus ascending order), click the column heading a second time.

List Mismatched Distances

The List of Mismatched Distances report lists the prefixes whose reported distance (or metric) between the prefix and a peer does not match the calculated distance. When the system calculates metrics across the topology model, those metric values are compared to the metrics reported by appliance peers. If the distance does not match, the prefixes are listed in the List of Mismatched Distances report.

Route Recorder compares these distances at the end of each full topology exploration to provide a measure of the accuracy of the topology model. Ideally, this report should be empty except for the messages telling when the last full topology exploration and subsequent periodic topology explorations began and ended.

Figure 64 shows an example of the List of Mismatched Distances report.

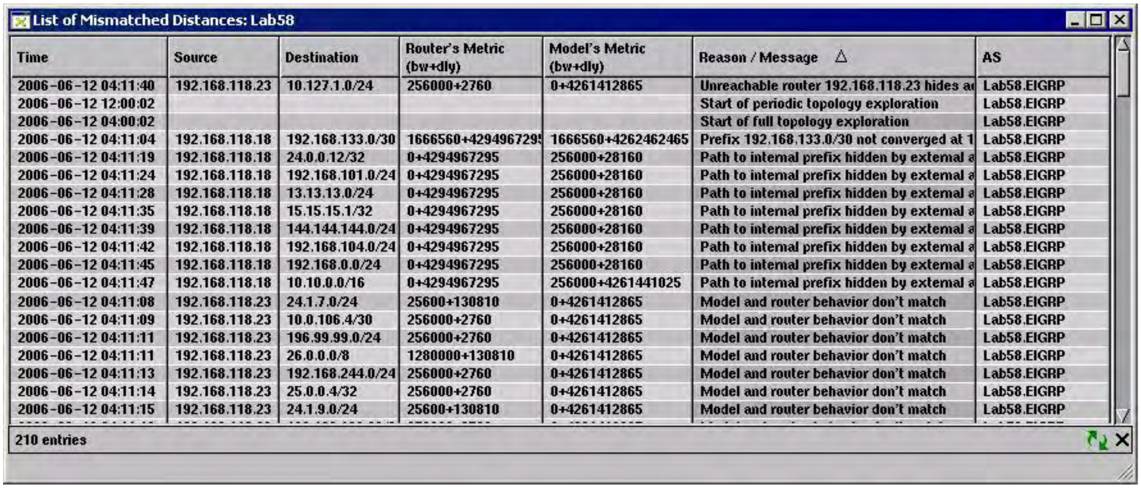

List of Mismatched Distances: Lab58

Time	Source	Destination	Router's Metric (bw+dly)	Model's Metric (bw+dly)	Reason / Message	AS
2006-06-12 04:11:40	192.168.118.23	10.127.1.0/24	256000+2760	0+4261412865	Unreachable router 192.168.118.23 hides a	Lab58.EIGRP
2006-06-12 12:00:02					Start of periodic topology exploration	Lab58.EIGRP
2006-06-12 04:00:02					Start of full topology exploration	Lab58.EIGRP
2006-06-12 04:11:04	192.168.118.18	192.168.133.0/30	1666560+429496729	1666560+4262462465	Prefix 192.168.133.0/30 not converged at 1	Lab58.EIGRP
2006-06-12 04:11:19	192.168.118.18	24.0.0.12/32	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:24	192.168.118.18	192.168.101.0/24	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:28	192.168.118.18	13.13.13.0/24	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:35	192.168.118.18	15.15.15.1/32	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:39	192.168.118.18	144.144.144.0/24	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:42	192.168.118.18	192.168.104.0/24	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:45	192.168.118.18	192.168.0.0/24	0+4294967295	256000+28160	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:47	192.168.118.18	10.10.0.0/16	0+4294967295	256000+4261441025	Path to internal prefix hidden by external a	Lab58.EIGRP
2006-06-12 04:11:08	192.168.118.23	24.1.7.0/24	25600+130810	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:09	192.168.118.23	10.0.106.4/30	256000+2760	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:11	192.168.118.23	196.99.99.0/24	256000+2760	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:11	192.168.118.23	26.0.0.0/8	1280000+130810	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:13	192.168.118.23	192.168.244.0/24	256000+2760	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:14	192.168.118.23	25.0.0.4/32	256000+2760	0+4261412865	Model and router behavior don't match	Lab58.EIGRP
2006-06-12 04:11:15	192.168.118.23	24.1.9.0/24	25600+130810	0+4261412865	Model and router behavior don't match	Lab58.EIGRP

210 entries

Figure 64 List Mismatched Distances Report

The report contains the following columns:

- **Time**—Indicates the time when the anomaly was detected.
- **Source**—Indicates the source IP address involved in the mismatched distance.
- **Destination**—Indicates the source IP address involved in the mismatched distance.
- **Router's Metric**—Indicates the metric for highlighted path cost. See Viewing the Highlighted Path Cost for EIGRP on page 139.
- **Model's Metric**—Indicates the metric for the topology model. See Viewing the Highlighted Path Cost for EIGRP on page 139.
- **Reason/Message**—Indicates the reason that the mismatch may have occurred:
  - **Unreachable router hides actual path**—Indicates that the path goes through a router that the system cannot access. The actual links traversed and their metrics are unknown.

- **Different equal-cost path chosen by model**—Indicates that a different path with equal cost was taken. If there are multiple paths with equal total costs but different bandwidth and delay components of the metric, then the system might choose a different path than the routers actually use. The router algorithm is not always deterministic.
- **Prefix not converged**—Indicates that the system has traced the actual path taken by the routers, and that one of the routers has no route to the prefix in question. This condition usually indicates that EIGRP routing to the prefix has not converged so the distance of the peer router is not valid.
- **Network has routing loop**—Indicates that the system encountered a routing loop when attempting to trace the actual path taken by the routers. This means EIGRP routing to the prefix has not converged and the distance given by the peer router is not valid.
- **Model and router behavior do not match**—Indicates that the system modeling of recorders behavior is not exact. This message may also indicate that a router has become confused and is reporting inconsistent metric information (perhaps due to a bug in the router software). Some router configuration changes, such as changing an access control list (ACL) used in a route filter, do not take effect until the routing process is restarted. The system will see the new value but the router does not use it, causing a distance mismatch. Report this message to Technical Support if it persists across multiple full explorations.
- **Failed to query**—Indicates a possible defect. Report this problem to technical support.

The message inserted at the end of the full exploration describes how many internal and external prefix distances did not match along with a comparison of the total number of distances known from peer routers.

The system periodically re-explores the topology to make sure no changes are missed due to transitions that do not result in an EIGRP update being sent or due to limitations in tracking network dynamics. The period is set as part of recorder configuration and defaults to eight hours. At the end of each periodic topology exploration, a message is added to the List of Mismatched Distances window describing the number of links and prefix attachments that were corrected during the last periodic topology exploration. Ideally, these numbers should be zero.

By default, the table is sorted by time in descending order. To sort the table by any other field, click the column heading. To change the sort order (descending versus ascending order), click the column heading a second time.

Find Invisible Links

The first time this option is selected, the system runs a simulation on its topology model to determine if there are any links where a failure will not be immediately detected because the routers that peer with the system will not report an EIGRP distance change. During the simulation, the system fails each router interface in the topology model one at a time and then checks for a change in the routing distance to any prefix from any of the routers that peer with it. If there is any change, the system will be able to detect a failure of the real interface. If not, the system can only detect the interface failure during the next periodic topology exploration (the default is every eight hours). The most common cause of invisible links is route summarization. Using GRE tunnels to peer with additional routers behind summarization boundaries can increase coverage.

The simulation can take several hours to run on a large network topology, but it can be canceled at any time by clicking **Cancel**. When completed, the results are stored in the database so that they can be viewed again later without waiting for the simulation to run again. If the topology changes or additional peer routers are added, click **Reload** in the Invisible EIGRP Interfaces window to re-run the simulation.

You can highlight invisible links in yellow on the topology map by clicking **Highlight**.

Assigning AS Names

The AS Name feature allows you to assign a name to the AS. This name takes priority over the AS name received from the Whois server. You have the option of keeping the Whois server name as the assigned AS name, or you can enter a new name.

To open the AS Name window, choose **Administration > Assign Names > ASs** (Figure 65).

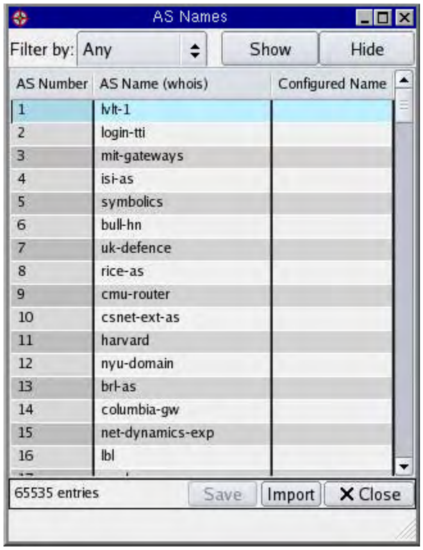

Figure 65 AS Names Window

This window includes the following columns:

- **AS Number—**Provides the number identifying the AS number.
- **AS Name**—Provides the name of the AS derived from the Whois server. (lowest priority)
- **Configured Name**—Enter a unique AS name in this column. (highest priority)

You can use the following filter options and buttons on this page:

- **Filter By—**Use these options show or hide entries in the table:
  - **Any—**Shows all rows.
  - **AS Number—**Shows a drop-down list that allows you to filter by entering the AS Number, while also showing previously used AS Numbers. Simultaneously, an **Options** drop-down list also opens, which allows you to filter with the following choices:

    Greater Than: Provides AS numbers with a greater number than the one entered in the text field.

    Exact Match: Provides only exact matches entered in the text field.

Begins With: Provides matches that begin with numbers entered in the text field.

— **AS Names—**Opens a text field where you can enter AS names. Simultaneously, an **Options** drop-down list also appears, allowing you to filter with the following choices:

Substring: Filters the AS Names with the given string as a substring for its name.

Exact Match: Filters the AS names with the given string as an exact match of its name.

Begins With: Filters the AS names with the beginning string used for the AS name.

- **Show—**Select this to show all AS name entries that you want listed.
- **Hide—**Select this to hide router entries.
- **Save—**Select this to save information you edit. This button is active only when you have edited information on the screen.
- **Import—**Select this when you want to edit several AS names.
- **Close—**Select this to exit the AS Names window.

To change the AS name, perform the following steps, perform the following steps:

1 Choose **Administration > Assign Names > ASs**.

The AS Names window opens.

2 Select the row of the AS name you want to change.

3 Enter the new AS name in the Configured Name field.

4 Click **Save**.

To change the name of multiple AS names, perform the following steps:

1 Choose **Administration > Assign Names > ASs**.

2 Click **Import**.

The Import AS Names dialog box opens, as shown in Figure 66.

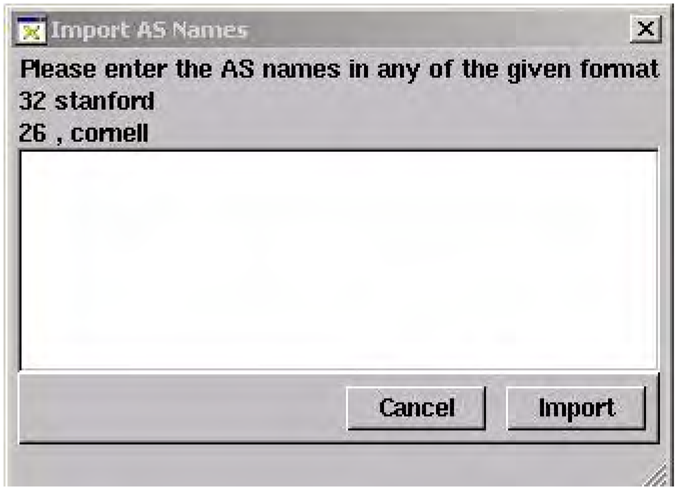

Figure 66 Import AS Names

3 Enter the names of the ASs you want to rename in the format.

4 Click **Import**.

The new AS names now appear in the Configured Name column.

Assign and Verify BGP AS Assignments to Routers

The system can show the path resolved between two points in a network. For network configurations that include BGP, correct BGP AS assignments to routers are required to accurately resolve the path. For a BGP confederation topology, it is not always possible to automatically determine the correct assignment for all routers. For network configurations that include BGP without confederations, the system can create BGP AS assignments for all routers automatically, however, some routers may not be running BGP. For these cases, you can change the BGP AS assignment manually in the BGP AS for Routers window.

If one or more routers do not belong in a BGP AS, you can select **No BGP** in The selected routers are in AS drop-down list. By specifying routers as not belonging to a BGP AS, the system can calculate IP routes across the topology more accurately. This may be needed in topologies without BGP confederations when only some routers run BGP and others follow a static default route.

To verify and manually assign AS assignments to routers, perform the following steps:

1 Choose **Administration > Assign BGP ASs to Routers** to open the BGP AS for Routers window, as shown in Figure 67. Some routers may have AS numbers already assigned to them as detected by BGP peering or computed from network topology.

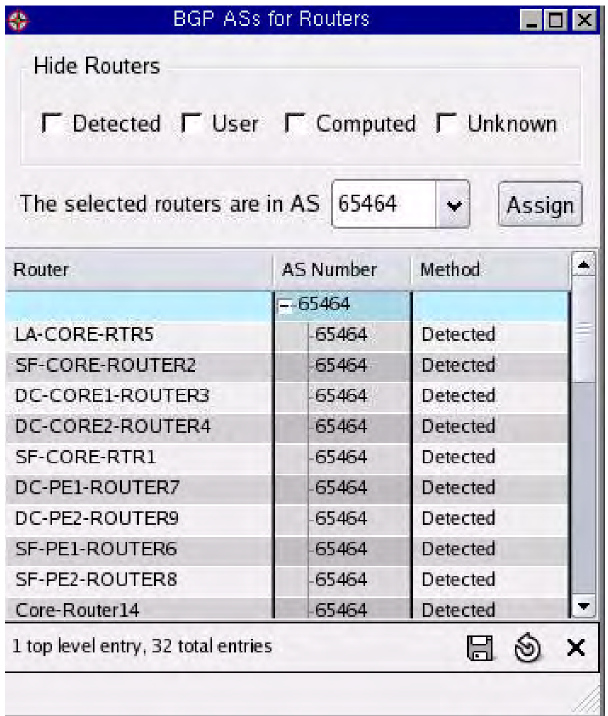

Figure 67 BGP AS for Routers

2 Click a router in the Router list. You can also select multiple routers or a range of routers by holding down the Ctrl or Shift key when you click another router in the list. Routers for which an assignment was detected cannot be reassigned.

3 Select the appropriate AS or No BGP in the Selected routers are in AS drop-down list. Select No BGP if, for example, the router is not running BGP and is following a default route.

4 Click **Assign**.

5 Repeat Steps 2-4 to manually assign other routers.

6 Click **Save User Input**.

7 Click **Close**.

4 The History Navigator

This chapter describes how to use the History Navigator to analyze the detailed routing history of the network.

Chapter contents:

Understanding the History Navigator

The History Navigator allows you to display the detailed routing history of a network. The Route Recorder obtains the routing history by monitoring the routing protocols and recording all protocol events in a database.

Every ten minutes the Route Recorder saves a complete, time-stamped snapshot of the routing topology. During each interval between snapshots, the system records all routing announcements with timestamps. This data enables the system to display a precise routing map of the network at any point in time.

The History Navigator includes the following options to view and analyze recorded data:

- View summaries of recorded data in graphical format to understand changes in vital network statistics over time.
- View the number of events between snapshots.
- For each snapshot, view the number of routers, routing adjacencies, and prefixes.
- Display detailed lists of routing events for a specified time period to aid in diagnosing a network outage or performing forensic analysis after an outage.
- Move back in time to a specific event and see that event replayed in the topology map.
- Distill large quantities of data related to an event down to the essentials.
- View a real-time graph of events as they occur.
- Perform root cause analysis on event data.
- Display the contents of the RIB or visual representations of the RIB at any point in time.
- Perform a before-and-after comparison of the state of the network at different times.

Accessing the History Navigator

You can open the History Navigator from the client application.

To access the History Navigator, perform the following steps:

1. Choose **Reports > History Navigator.**

   The History Navigator window opens to show a graph of recorded network routing events (Figure 1).

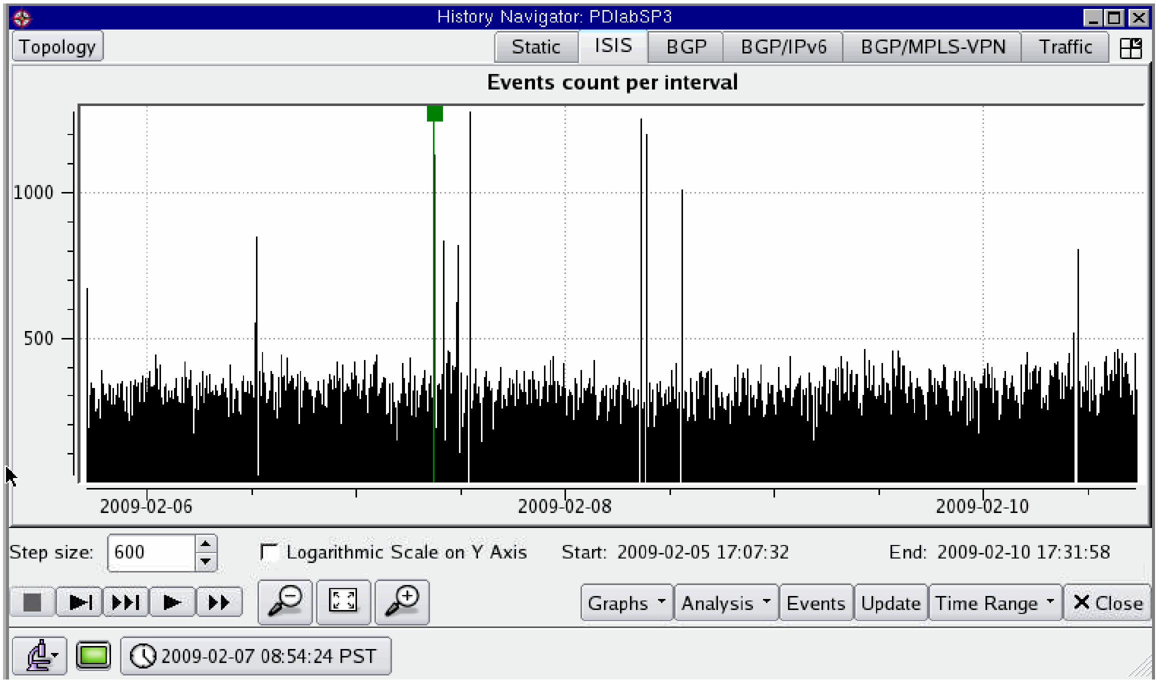

Figure 1 History Navigator Window (Analysis Mode)

2 Perform operations as described in the next section, Working With the History Navigator on page 157.

Working With the History Navigator

This section describes the following History Navigator capabilities:

- History Navigator Controls on page 158
- History Graphs on page 165

If a database contains multiple protocols, the History Navigator window displays multiple tabs, one for each protocol. For example, Figure 1 shows the tabs for a database that contains BGP, EIGRP, and OSPF protocols.

History Navigator Controls

The History Navigator window controls allow you to navigate through the routing database and customize the presentation of data. The controls are available in the main History Navigator window (see Accessing the History Navigator on page 156).

Modes

The elements display in History Navigator depend on the current topology map mode. The following shown in Table 1 are available in the History Navigator window:

Table 1 Modes

	**Monitoring mode**	Indicates that the topology is currently being recorded and updates to the routing database are shown in the graphs as they occur. In this mode, the playback controls are disabled. Just above the playback controls is a text box that specifies the interval in minutes between updates. In addition, the Graphs button is disabled in this mode if the Time Range option is set to Online, traffic data is not displayed. The History Navigator window displays routing events as they occur.
	**Analysis Mode**	Indicates that only previously recorded information in the routing database is shown on the topology map. In this mode, the playback buttons are enabled and just above them is a text box that specifies the step size in seconds that is used during playback.
	**Planning Mode**	A network icon indicates that planning features are enabled for the topology map.

The controls for the window in Monitoring mode are the same as those present when the window is in Analysis mode or Planning mode. However, options that do not apply in Monitoring mode and Planning mode are disabled, such as playback controls.

To switch modes in the History Navigator, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Click the mode icon in the lower left corner of the window and choose the desired mode.

The time range shown on the graph changes from Online to One Week . See Buttons on page 162 for information about the values that can be set with the Time Range button.

When switching from Analysis to Monitoring mode, a pop-up box informs you that if the amount of events exceeds a certain threshold, the analysis of the events could take a few minutes. You can restart Monitoring mode without the analysis, however the event history graph will show no data for that period.

Topology and Protocol Selection

The History Navigator includes a Topology button at the top of the window and buttons for each available protocol.

To select the topology and protocol, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Click the mode icon in the lower left corner of the window and choose the desired mode.

3 Click **Topology** and select the topologies from the drop-down list.

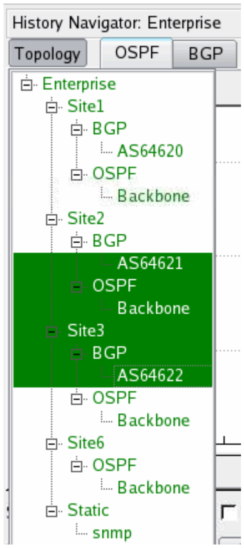

Figure 2 Topology Drop-Down List

4 Click the tab for the desired protocol.

Status Bar

The tool bar at the bottom of the History Navigator window is the same as the status bar on the Topology Map window. See Main Window Toolbar on page 52 for information about the icons and indicators on the status bar.

Cursor

The cursor is the green vertical hairline with green squares at the top and the bottom (see Figure 3).

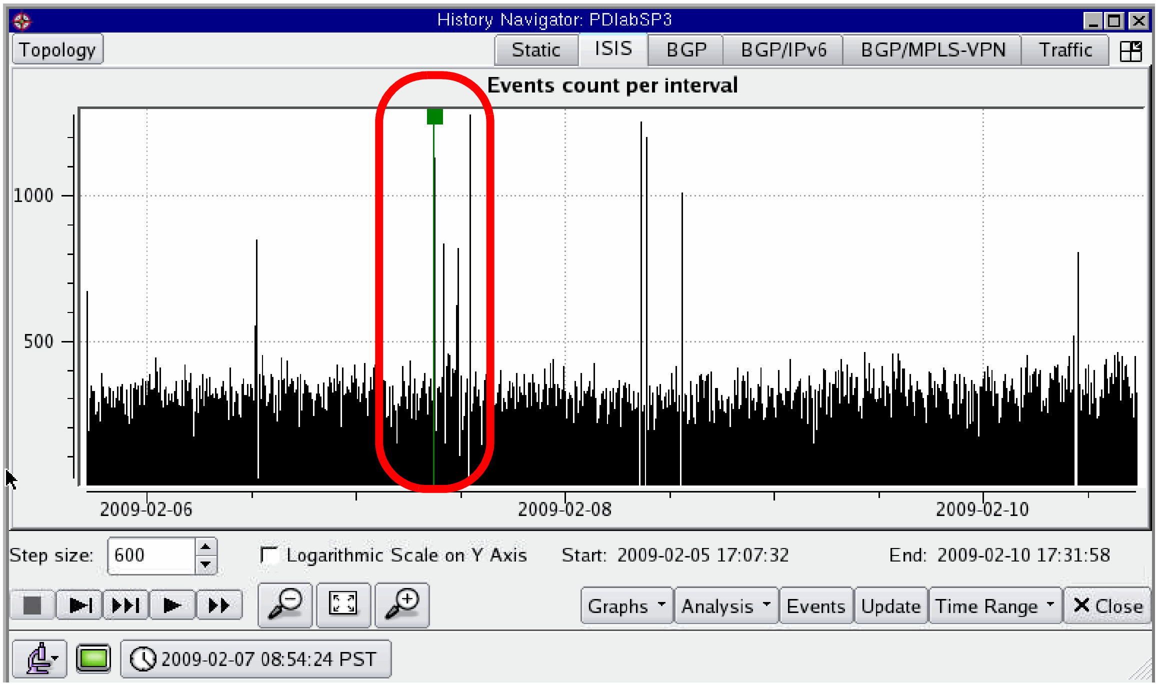

Figure 3 Cursor in History Navigator Window

The cursor indicates the currently displayed point in time within the routing topology history. There are several ways to move the cursor:

- Use the mouse to drag the cursor to a different point on the time line. The routing topology map immediately displays the routing topology as it existed at that time.

  In RAMS Traffic, traffic data may not be available if the selected time is within 30 minutes of the current time.

- Right-click a point in the time line. A pop-up prompts you to confirm whether you want to move time to that point. Click **Yes**. The topology map immediately displays the routing topology as it existed at that time.
- Move the cursor through time by stepping or animating (automated stepping) using the playback controls as described in Playback Controls on page 163. Any paths highlighted on the routing topology map are recomputed and redisplayed at each step of the replay.

The routing topology database does not store nodes and links that are down; therefore, objects that are down when the topology is first opened are not shown. If the cursor is moved back to a time when a down node or link was up, and then the cursor is moved to the current time again, the down node or link may remain on the map in red to indicate that it is down, if the time traversed by the cursor movement is less than the failed node or link timeout interval. See Auto-Hide on page 79 for information about node/link timeout intervals.

Buttons

Use the panel of buttons in the lower right-hand corner of the window to access graphs, data analysis tools, and events lists, and to specify the time range for display:

- **Graphs**—Choose the graphs to display. See History Graphs on page 165 for a description of the graphs. The Graphs button is disabled if you are working with an actively recording database and the Time Range option is set to **Online**.
- **Analysis—**Choose from a list of data analysis tools. See Analyzing Historical Data on page 167 for information about these tools.
- **Events**—Display a detailed list of routing events. See Understanding the Events List on page 200 for information about the Events list.
- **Update**—Update the display. This button is enabled only if the current window represents an actively recording database. If you display the window for more than 15 minutes, you can add newly recorded data to the current graph by clicking the **Update** button.
- **Time Range**—Choose the data range to include in the History Navigator window. By default, the time range is set to **Online** when in Monitoring mode and to **One Week** when in Analysis mode. You can set the time range to one hour, one day, one week, one month, or a custom range. If you choose **Custom**, a pop-up window opens to allow you to set the range. If you choose **Recent Custom**, a pop-up window opens to allow you to choose from recently-specified custom ranges.

    In RAMS Traffic, traffic data may not be available if the selected time is within 30 minutes of the current time.

- **Close** button—Close the History Navigator window.

Playback Controls

The panel of buttons in the lower left-hand corner of the window control playback (see Table 1).

Table 2 Playback Controls

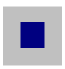	**Stop**	Stops animated playback. This button is enabled only in animated playback mode.
	**Step**	Advances the cursor by the number of seconds specified in the Step size text box. The topology map is updated with the recorded data from the new point in time.
	**Fast Step**	Advances the cursor by 10 times the number of seconds specified in the Step size text box. The topology map is updated with data from the new point in time.
	**Animate**	Automatically steps through routing history by executing a continuous sequence of cursor advances, with the network map being updated at each step. If any paths are highlighted, the routes will also be recomputed and redisplayed at each step. Click **Stop** to stop the animated playback.
	**Fast Animate**	Starts animated playback in fast mode, automatically advancing the cursor by 10 times the number of seconds specified in the Step size text box. Click **Stop** to stop the animated playback.

Because stepping and animation advance time in steps of the specified interval, a routing change will not be shown if it occurs and then changes back within one time interval. The Events List window, described on Understanding the Events List on page 200, includes all routing changes.

Logarithmic Units

To display the Y axis of the graph in logarithmic units, select the **Logarithmic Scale on Y Axis** check box. The graph is redisplayed with logarithmic Y values.

Zooming the Time Line

You can zoom into a subsection of the recorded history shown on the History Navigator graph. Zooming in the time dimension may help you to see more detail within a cluster of event spikes when the graph covers a long period of time.

The zoom buttons zoom only the values on the X axis.

The panel of buttons toward the bottom center of the screen control zooming functions (see Table 3).

Table 3 Zoom Controls

	**Zoom In**	Allows you to zoom into a subsection of the History Navigator graph.
	**Zoom Out**	Allows you to zoom out of a subsection of the History Navigator graph.
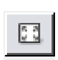	**Reset**	Resets the view back to the standard viewer setting.

To zoom the time line, perform the following steps:

1. While holding the Ctrl key down, right-click and drag a rectangle with the mouse to select the area to be expanded to fill the graph.
2. While holding the Ctrl key down, release the left mouse button to set the zoom area.
3. Repeat steps 1 and 2 to increase the level of zoom.

4 To broaden (unzoom) the view one level, hold the Ctrl key down, and then click the right mouse button. Repeat until the graph returns to the original zoom level.

History Graphs

The primary feature of the History Navigator window is the Events graph that is shown in the default window. Several additional graphs can be selected to show a wide range of statistics about the state of the network.

In RAMS Traffic, open the Traffic tab to display additional traffic-related graphs.

To access history graphs, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Click **Graphs**, and then choose the desired graph.

The Graphs button is disabled when you are working with an actively recording database and the Time Range option is set to Online.

The following graphs may be available, depending upon the selected protocol:

- **Routers**—Displays the number of physical entities in the network. For OSPF, this includes AS Border Routers in other areas that are visible from the viewed area.
- **Routes**—Displays the number of routes advertised in the network. This graph does not appear if the topology currently selected in the History Navigator window is an IGP area or AS.
- **Links**—Displays the sum of router-to-router links plus the number of router-to-prefix links. This graph does not appear if the topology currently selected is a BGP AS.
- **IPv4 Prefixes/IPv6 Prefixes**—Displays the number of IPv4 or IPv6 prefixes available in the entire network.

- **BGP Updates**—Displays the number of how many announced and withdrawn packets received in the preceding 10 minutes. Announced packets are represented by blue lines and re-announced packets are represented in dark yellow lines.
- **ISIS Activity**—Displays LSP activity for IS-IS domains. New activity displays in blue, refreshed activity displays in dark yellow. (option is shown only for IS-IS)
- **ES Neighbors—**Displays activity between the router and end systems and routers.
- **Prefix Neighbors—**Displays activity between neighboring routers.
- **OSPF Activity**—Displays LSA activity for OSPF domains. New activity displays in blue, refreshed activity displays in dark yellow.
- **OSPFv3 Activity**—Displays LSA activity for OSPFv3 domains. New activity displays in blue, refreshed activity displays in dark yellow.
- **EIGRP Activity**—Displays updated and inferred activity for EIGRP domains. New activity displays in blue, refreshed activity displays in dark yellow.
- **Events—**Displays the number of routing protocol changes that occurred in the network between recorded snapshots. Example routing events include a neighbor adjacency going down or a new prefix being announced. For EIGRP, both distance-vector events and derived link-state events are included.
- **Interfaces**—Displays the number of interfaces discovered by static protocol.

You can configure the system to include any or all of the graphs in the default window. See Applying Configuration Options on page 71.

In Monitoring mode, the default interval for updating the events graph is 10 seconds. If any events are generated during this time frame, a spike corresponding to the number of events is drawn on the graph.

Analyzing Historical Data

The History Navigator supports the following tools for detailed analysis of historical data:

- Root Cause Analysis on page 167
- RIB Visualization on page 175
- RIB Comparison on page 186
- Trending on page 192
- Event Analysis on page 193
- Flow Record Browser on page 199

Root Cause Analysis

The Root Cause Analysis function analyzes the huge amounts of data generated by BGP-related routing events and distills the data down to the essential information that helps you to pinpoint the cause and location of the event.

Traffic data in RAMS Traffic is not relevant for this type of analysis, and that loading traffic information from the database may slow the analysis process. When you open a topology for BGP-specific analysis, it is recommended that you deselect any traffic databases from the tree in the Open Topology dialog box as described in Chapter 3, "The Routing Topology Map"

To perform a root cause analysis, perform the following steps:

1. Use one of these options to open the root cause analysis window:
   - Choose **Reports > History Navigator** to open the main History Navigator window. Choose **Analysis > Root Cause Analysis**.
   - If your topology has BGP data, choose **Reports > Root Cause Analysis**.

   If you have multiple BGP topologies, the system prompts you to select a single topology.

2 Left-click in the graph just before the events occurred.

3 Left-click in the graph just after the events occurred.

4 (Optional) If you have more than 500k events, a window prompts you to **Continue**, **Abort**, or **Prefilter** the events.

5 (Optional) If you select **Prefilter,** the Event Prefiltering window opens. From here, you can select from a list of filters from the drop-down menu, decreasing the time it takes to generate the event list. For more information on using filters, see Using Filters on page 221.

If no significant incidents occurred during the time you specify, a message indicates that the Root Cause Analysis algorithm did not find any major BGP problems. Adjusting the analysis options as described in Chapter 3, "The Routing Topology Map," may increase the number of incidents the algorithm will find.

If incidents are found, the Root Cause Analysis Results window opens. Figure 4 shows an example.

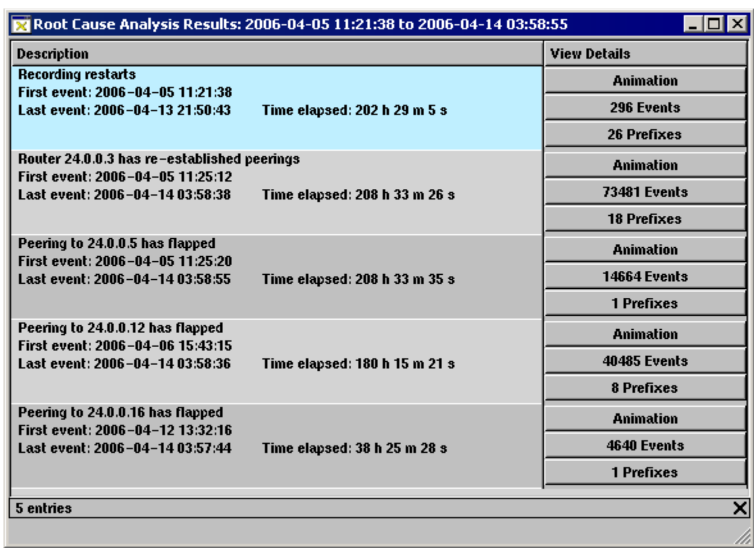

Figure 4 Root Cause Analysis Results

The events that occurred during the time period you specified in Steps 2 and 3 are analyzed and correlated into groups. All of the BGP routing messages that apply to related events are summarized. For each group, the inferred root cause will be one of the following:

- **Prefixes shifted—**When the event messages of a group indicate that prefixes have shifted, the number of prefixes that have shifted on a specified link is listed. "Shifted" refers to the total number of prefixes that have either left or joined the link. The count is approximated because the same prefixes may join and leave the link more than once, or different prefixes may be joining and leaving. Therefore, the system calculates a maximum shift size or change in count.
For example, if 2 prefixes left and 3 prefixes joined, the maximum shift size is 4. The total number of prefixes to traverse the link is at least 4, but can be any number between 4 and 9, hence the approximation.

- **Prefixes are flapping—**When the event message of a group indicates that prefixes are flapping, the prefixes are being announced and withdrawn, possibly over and over. The corresponding message takes one of two forms:

  — 192.168.103.0/24 is flapping on 128.32.0.66 > 11423.calre_2

    In this case, one prefix (192.168.103.0/24) is flapping on the specified link
    (128.32.0.66 > ...).

  — 3435 prefixes are flapping on 128.32.0.66 > 11423.calre_2

    In this case, 3435 prefixes are flapping on the specified link (128.32.0.66 > ...).

- **Peering established —** Each time the system establishes a BGP peering, the peer router sends all of its BGP routes. At the end of this sequence, the system writes a synchronized event.

- **Peering lost—**When a peering is lost, the connection between the system and the peer router is closed. The Animation window will continue to show the effect of associated prefix withdrawals, but the withdrawn prefixes are not listed in the Prefixes table.

- **Peering has flapped—**When a peering is both established and lost, the peering has flapped. If the peer router has several established and lost events during the selected interval, all of these events are combined into a single incident.

- **Router has lost peerings—**When the system loses its connection to a peer router, it may detect that other routers have also lost contact with the peer. As a result, the system infers that the router has lost peerings. This message might indicate, for example, that the router is rebooting.

- **Router has established peerings**, or **re-established peerings—**When a peering is established, all prefixes are added at once, which creates an artificial spike in activity.
- **Recording stopped**, **started**, or **restarted—**Recorders write to the database when they start recording and when they shut down. Any peerings that are established within five minutes of the start of recording are included in this event message. Multiple stops, starts, and restarts within the selected interval are combined into a single event.

To the right of each group of event messages are three buttons:

- The **Animation** button generates an animated visualization of the routing topology during the related events. The Animation window is described below.
- The **Events** button displays a detailed list of the events that constitute the group.
- The **Prefixes** button displays a list of all prefixes affected by the group of events.

Animation Window

The routes of a BGP router form a virtual tree rooted at the router. The visualization function creates a graphical representation of this tree from the viewpoint of each BGP edge router (or core route reflector) and merges them into a single tree. The system appears at the left side of the tree.
The time range of the animation is indicated in the rectangle. To the right of the rectangle are its BGP peers; to their right are its BGP next hops; to their right are the AS the next hops serve; to the right of each AS is the downstream AS, if any; to the right of the downstream AS are the prefixes it advertises.

The visualization function assigns a weight to each edge (that is, each trunk or branch or twig) of the tree that is equivalent to the number of unique prefixes carried by the edge, and uses this weight to determine the thickness of the line representing that edge. The thickness of an edge displayed on the Animation window is based solely on the number of prefixes that are routed over that edge, not how much traffic is flowing over the edge. (The visualization function is a routing diagnostic tool, not a traffic diagnostic tool.)

[Figure 5](#) shows an example of the Animation window.

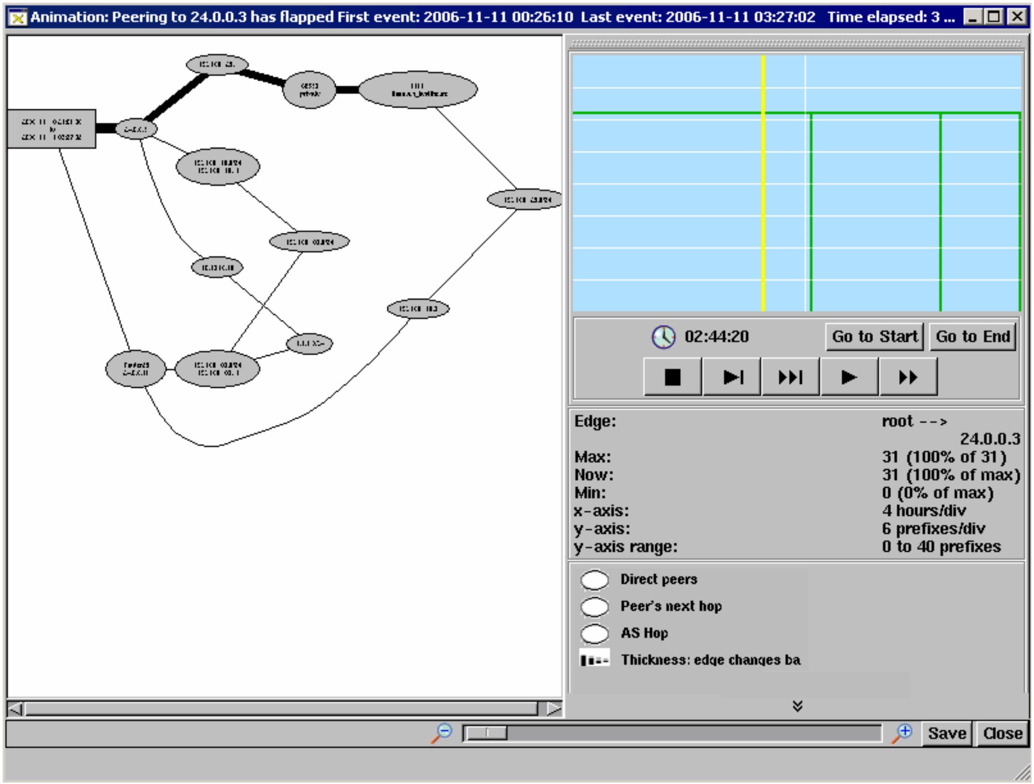

Figure 5 Animation Window

Animations can help you to identify, isolate, and resolve problems that are difficult to diagnose, for example, continuous customer route flaps, persistent MED oscillations, and “leaky” routes.

The upper pane of the window displays the visualization. The lower pane of the window contains the following elements:

- **Clock—**Indicates elapsed time during animation.
- **Playback** controls—Control the animation. These are identical to the playback controls on the History Navigator window (see Playback Controls on page 163).

- **Go to Start** and **Go to End** buttons—Reposition the yellow cursor in the graph. In addition, you can move the cursor by clicking the position in the graph to which you want to move.
- **Graph—**Represents the change in number of prefixes carried by an edge. By default, the edge on which the graph is based is the most active edge, that is, the edge that lost or gained the most prefixes. You can change the perspective of the graph by clicking on another edge in the visualization.

  To the left of the graph is a list of details about the graph, including the nodes at either side of the edge on which the graph is based, the maximum, current, and minimum number of prefixes carried by the edge, and the scale of the x and y axes of the graph.

Playing an Animation

When you animate the visualization using one of the playback controls, the group of related events you selected is replayed in both the visualization pane and the graph pane.

- In the visualization pane of the window, the thickness and color of an edge indicates the level of activity on the edge.
  - The thickness of the line representing an edge changes based on the number of unique prefixes that are routed over the edge. The thickness of a gray shadow surrounding a line indicates the maximum number of prefixes the edge ever carried, while the thickness of the colored portion of the line indicates the current number of prefixes the edge carries.
  - The color of the edge changes as the edge gains or loses prefixes. A black line indicates that no changes are occurring. Green indicates that the edge is gaining prefixes, while blue indicates that the edge is losing prefixes.
- In the graph pane of the window, a yellow line indicating the current position in the animation moves from left to right, while the clock indicates elapsed time.

To play an animation, perform the following steps:

1. Choose **Reports > History Navigator** to open the main History Navigator window.
2. Choose a time period (see Cursor on page 160).

3 Choose a playback option (see Figure 6):

- Step mode advances the cursor. The step interval depends upon the actual duration covered by the visualization.

  The formula for calculating the interval in milliseconds is $I = (A/P) \times 25$, where A is the actual duration in seconds of the period covered by the animation, and P is 30 for step mode or 15 for fast step mode.

- Fast step mode advances the cursor by 10 steps.

- Animate mode advances the cursor in a continuous series of steps through the time range covered by the visualization. The animation completes in 30 seconds, regardless of the actual interval covered by the visualization.

- Fast Animate mode advances the cursor through the time range covered by the visualization. The animation completes in 15 seconds, regardless of the actual interval covered by the visualization.

- Stop halts the playback of an animation.

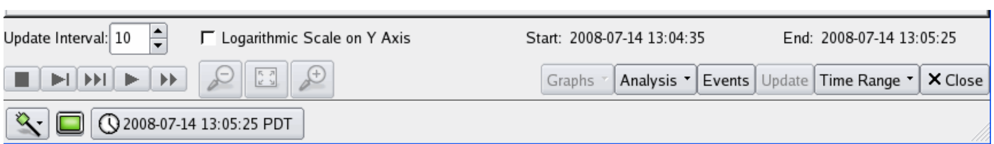

Figure 6 Playback Controls

4 Click the corresponding playback button to begin the animation.

If the adjacency between the appliance and its peers was down at the specified start point, the Animation window may initially be blank except for the rectangle representing the appliance. However, when you click a playback button, the tree is filled in as adjacencies stabilize.

Saving an Animation

You can save an animation for later viewing by clicking Save in the lower right corner of the Animation window. The animation is saved in SVG format (file extension .svg) on the appliance hard disk. SVG is a W3C standard for producing high quality graphics. SVG support is not yet standard in most browsers, so you may need to download a plug-in to view saved animations.

Adobe offers a free SVG plug-in that you can download from **http://www.adobe.com/svg/viewer/install/main.html**). The Adobe plug-in is compatible with a variety of browsers on Linux, Mac OS X and Microsoft Windows platforms.

Alternatively, the Apache Batik project offers a standalone SVG viewer, squiggle, that you can download from **http://xml.apache.org/batik/install.html#distributions**. Because squiggle is written in Java, it runs on almost any platform, but the current version may require more CPU usage than the Adobe viewer.

To view a saved animation, perform the following steps:

1 Open the web application on a Route Recorder and choose **Home**.

> Administrator privileges are not required to view a saved animation.

2 Choose **Reports Portal** on the top navigation bar.

3 Click **BGP Animations** on the left navigation bar.

The BGP Animations page displays a list of all saved SVG files, and contains a link to the installation page for the Adobe plug-in.

4 Click the title of the animation you wish to view to open an Animation window for that animation. All of the information and controls present in the original animation are available in the saved animation.

RIB Visualization

The RIB Visualization function provides you with a still image that represents the BGP Routing RIB at the time indicated by the current History Navigator window cursor position. Visualizations can help you identify difficult-to-diagnose problems such as prefix load-balancing issues.

Generating a Visualization

To generate a RIB visualization, perform the following steps:

1 Use one of these options to open the RIB visualization window:

- Choose **Reports > History Navigator** to open the main History Navigator window. Move the History Navigator cursor to the time desired, and choose **Analysis > RIB Visualization**.
- If your topology has BGP data, choose **Reports > RIB Visualization**.

2 If you have multiple BGP topologies, the system prompts you to select a single topology.

The RIB Visualization window opens, as shown in Figure 9.

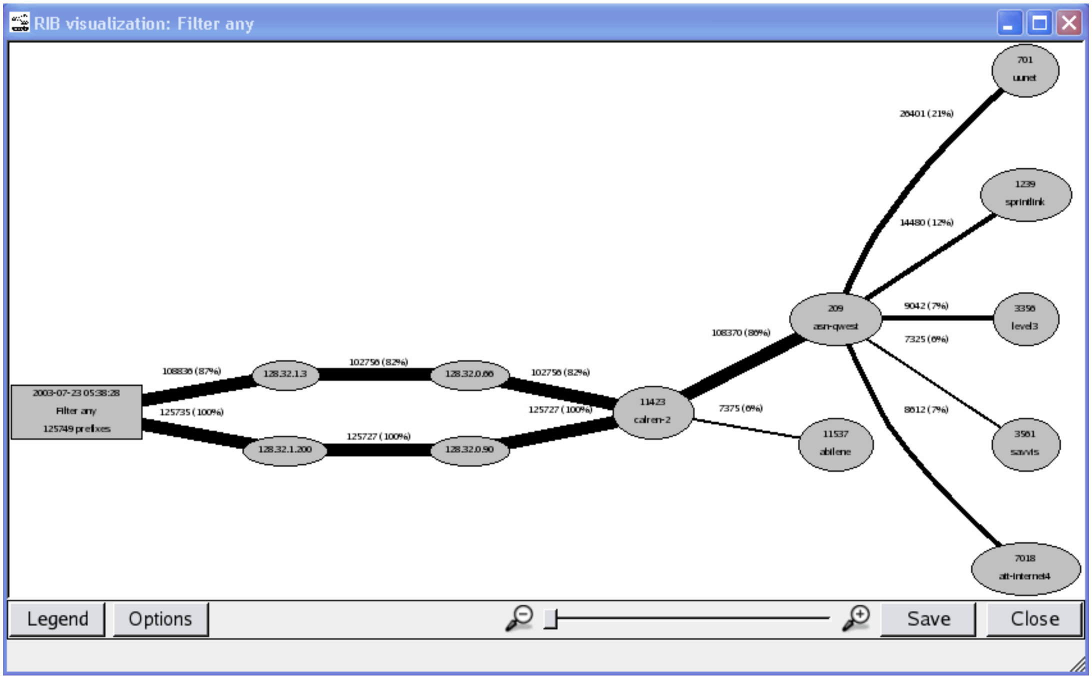

Figure 7 RIB Visualization Window

RIB visualization is enabled only if you are viewing one BGP AS.

The routes of a BGP router form a virtual tree rooted at the router. The Visualization function creates a graphical representation of this tree from the viewpoint of each BGP edge router (or core route reflector) and merges these trees into a single tree. The appliance appears at the left side of the tree. The appliance rectangle indicates the date and time of the RIB snapshot and the total number of prefixes. In addition, the rectangle indicates how the routes were filtered before the picture was generated. You can create visualizations filtered to include only a subset of the routes from the RIB Browser. See RIB Browser on page 180.

To the right of the rectangle are its BGP Peers, followed by its BGP next hops; to their right are the ASs that the next hops serve; to the right of the ASs are any downstream ASs; to the right of the downstream ASs are the prefixes they advertise.

The RIB Visualization function assigns a weight to each edge (that is, each trunk or branch or twig) of the tree that is equivalent to the number of unique prefixes carried by the edge, and uses this weight to determine the thickness of the line representing that edge. The thickness of an edge displayed on the RIB Visualization window is based solely on the number of prefixes that are routed over that edge, not how much traffic is flowing over the edge. The visualization function is a routing diagnostic tool, not a traffic diagnostic tool.

Each entity in the visualization is identified, and each edge is labeled with the number of unique prefixes advertised on that edge and the percentage of the total number of prefixes in the network.

The bottom of this screen has a zoom slider which allows you to zoom in on any portion of the window by moving the zoom slider to the right. You can also pan across the screen by holding down the space bar while left-clicking on the mouse.

Changing RIB Visualization Thresholds

Visualization options control whether a network entity appears in a RIB Visualization window or a root cause analysis Animation window. For each type of entity, you can choose to always include it, include it if it announces more than a specified percentage of prefixes to any of its peers, or to never include it. The **Always** option is disabled if choosing it could create a visualization too big or crowded to read.

The following entities are included on the Options panel of the RIB Visualization window:

- **Show a Peer**. The default is to include a peer if it announces 5% or more of the total number of prefixes.
- **Show a Nexthop**. The default is to include a nexthop if it announces 5% or more of the total number of prefixes.
- **Show a Neighbor AS**. The default is to include the neighbor AS if it announces 5% or more of the total number of prefixes.
- **Show a Non-Neighbor AS**. The default is to include the non-neighbor AS if it announces 5% or more of the total number of prefixes.

- **Show an Edge to a Prefix**. Select **Always** to show an ellipse for the prefix on the visualization and connect that ellipse to the other elements. For example, Figure 8 shows a visualization without the **Never** Show an Edge to a Prefix option selected. Figure 9 shows the same visualization with **Always** option selected.

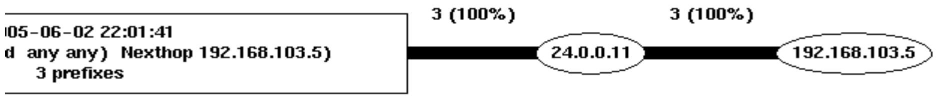

Figure 8 RIB Visualization without Show an Edge to a Prefix Option

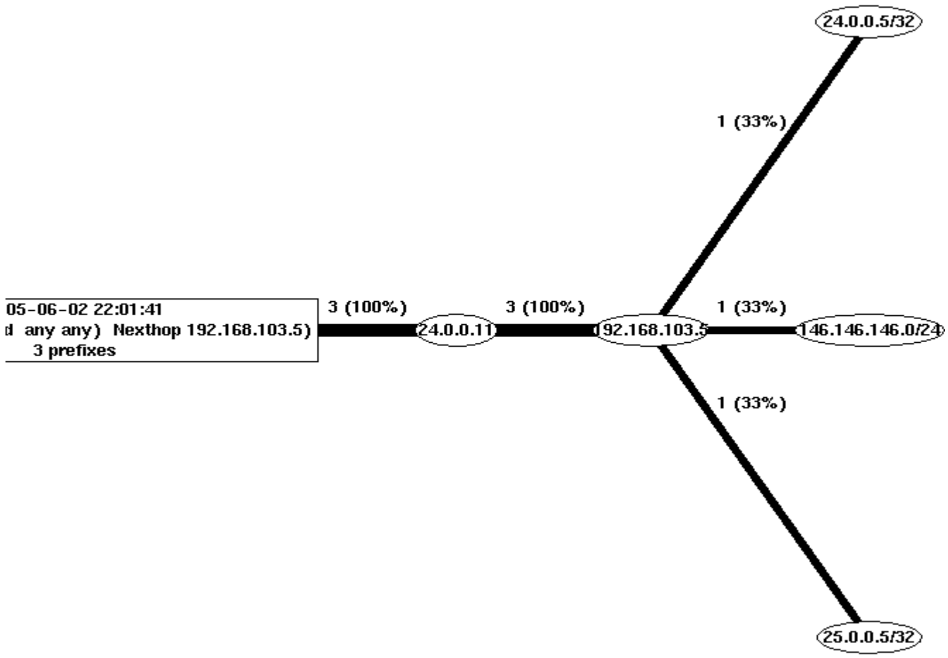

Figure 9 RIB Visualization with Show an Edge to a Prefix Option

To change RIB visualization thresholds, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Move the History Navigator cursor to the time desired.

3 Choose **Analysis > RIB Visualization**.

4 Click **Options**.

The Visualization options are displayed in the left pane of the window.

5 Change thresholds as desired.

Lowering a threshold increases the number of entities that are included in the visualization, giving you a more detailed picture. Conversely, raising a threshold decreases the number of entities and level of detail.

If choosing one of the Always options would create a visualization too big or crowded to read, that option is disabled.

6 Click **Redraw**, and then select **In Place** or **In New Window**.

If you select In Place, your changes are applied only to the current window. If you select In New Window, your changes are applied only to the new window, not to the original window.

Saving a Visualization

You can save a visualization for later viewing by clicking **Save** in the Visualization window.
The visualization is saved in SVG format (file extension `.svg`) on the hard disk. SVG is a W3C standard for producing high quality graphics. SVG support is not yet standard in most browsers, so you may need to download a plug-in to view saved visualizations. See Saving an Animation on page 174 for available plug-ins.

To view a saved visualization, perform the following steps:

1 Open the web application on a Route Recorder and choose **Home**.

Administrator privileges are not required to view a saved animation.

2 Choose **Reports Portal** on the top navigation bar.

3 Click **BGP Animations** on the left navigation bar.

The BGP Animations page displays a list of all saved SVG files, and contains a link to the installation page for the Adobe plug-in.

4 Click the title of the visualization you wish to view to open a window for that visualization. All of the information and controls present in the original animation are available in the saved animation.

RIB Browser

Use the RIB Browser to display the following types of information:

- For IGP, distribution of the number of down links and prefixes per router.
- For BGP, distribution of routes based on attributes such as Peer, Nexthop, MED, and so on.
- For OSI IS-IS, in addition to the fields displayed for IGP, distribution of the number of down ES Neighbors and Prefix Neighbors.

To open the RIB browser, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Specify a time range (see Adjusting the Time Range on page 210).

3 Choose **Analysis > RIB Browser** to open the RIB browser (Figure 4-9).

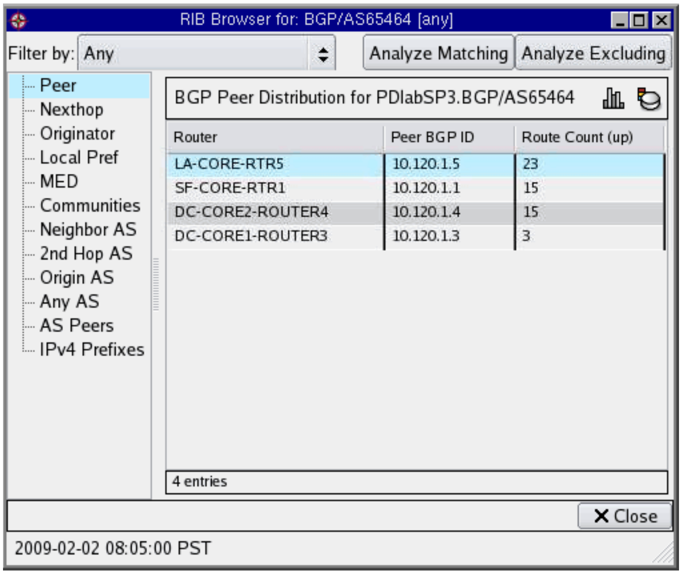

Figure 10 RIB Browser

4 Use the **Filter By** drop-down list to select filter parameters to restrict the input data included in the analysis. Click **Analyze Matching** to reanalyze the input data including only the items that match the filter criteria, or click **Analyze Excluding** to reanalyze the input data including only the items that do not match the filter criteria.

For example, the IGP RIB Browser includes a report on distribution of down prefixes per router. If you apply an OSPF Prefix Type filter to select just Internal prefixes and then click **Analyze Matching**, the report shows a distribution of down internal prefixes per router.

5 If multiple topology domains are included in the analysis, there are two levels of menus of the left. Select a domain from the first menu.

6 Select the type of information to view from the second (or only) menu on the left.

7 The table on the right redraws to show the selected information.

IGP Protocols

For IGP protocols, the RIB Browser conditionally shows two types of information, depending upon state of objects in the network:

- If there are any down links or down prefixes at the time selected for display, then the report shows the distribution of the number of these down objects per router.
- If there are any isolated routers (those with all adjacencies down) or routers indicated as “overloaded” in an IS-IS network, then the RIB browser shows a list of those routers.

Figure 4-10 shows an example of the RIB browser with IS-IS data.

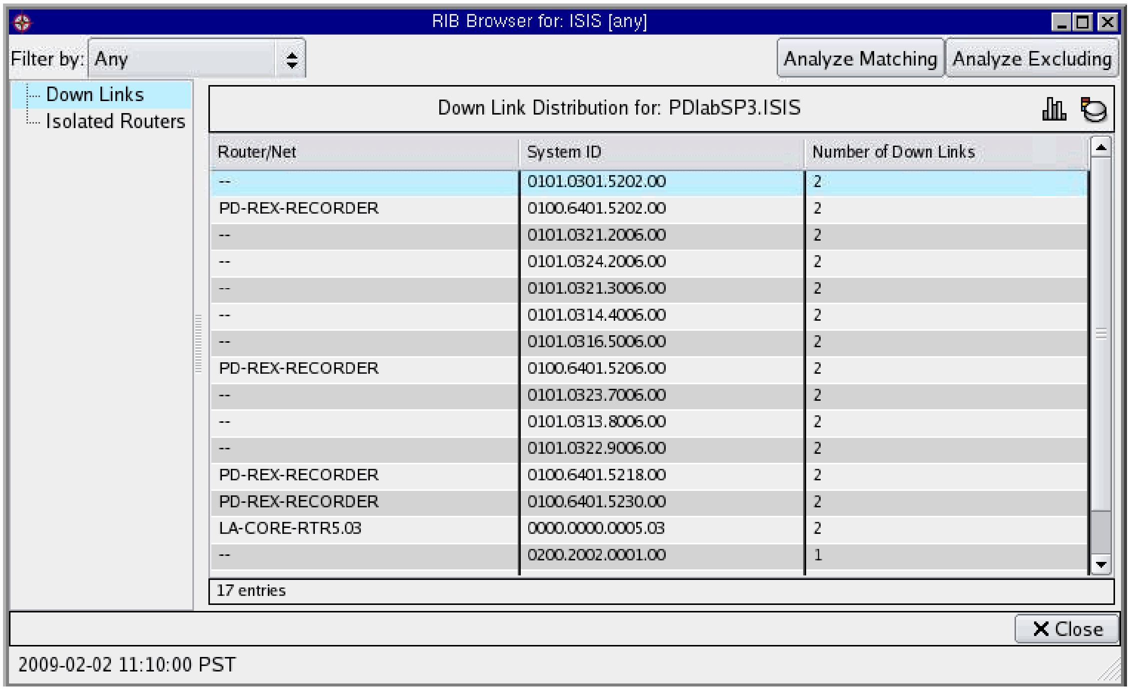

Figure 11 RIB Browser for IGP

On the left side of the RIB browser, click **Down Links** or **Down Prefixes** to view the list of routers that have links or prefixes that are currently down. For each router listed in the Router column, the Number of Down Links (or Prefixes) column displays the number of links or prefixes on that router that are down. To view the list of routers as a bar chart, click **View as Bar Chart** in the upper right corner of the window.

If the network has any isolated or overloaded routers, click **Isolated Routers** or **Overloaded Routers** on the left side to display a list of routers in that state. The Overloaded condition indicates routers in an IS-IS network that have the “overloaded” bit set. These options appear only if there are isolated or overloaded routers in the network.

To locate one of the identified routers on the routing topology map, click the list entry for that router. That entry will be highlighted in the list and the router will flash yellow on the routing topology map. Alternatively, click Color Routers to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of links or prefixes that are down.

To view the details for a particular router, right-click the corresponding list entry, and then select one of the following choices from the pop-up menu:

- **Show Links/Show Prefixes**—Displays a list of the links or prefixes associated with the selected router that are down.
- **Filter Analysis**—Displays a new RIB Browser window with data for the selected router only.

BGP Protocol

Figure 12 shows an example of the RIB Browser with BGP peer data.

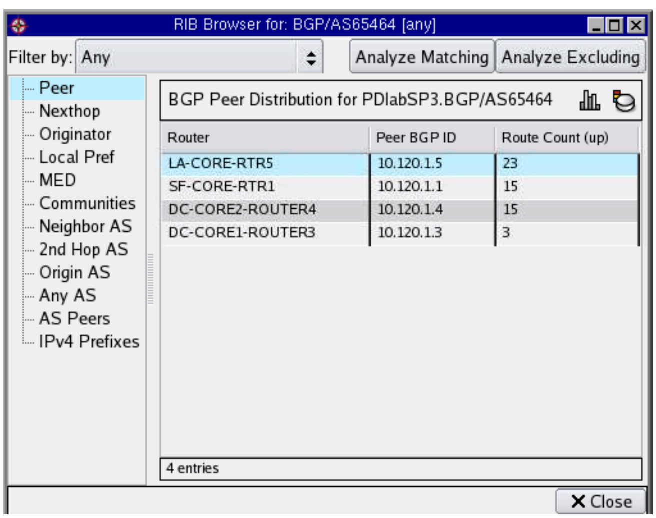

Figure 12 RIB Browser for BGP

For BGP protocol, the RIB Browser displays distributions of the advertised prefixes according to their attributes. A tree structure on the left side of the window presents attribute options that may include the following:

- **Peer**—Displays the number of routes advertised by the peer.
- **Nexthop**—Displays the number of routes that list that Nexthop router among their attributes.
- **MP Nexthop**—Displays the number of multi-protocol (VPN or IPv6) routes that list that MP Nexthop router among their attributes.
- **Originator—**Displays the number of routes that list that Originator among their attributes.
- **Local Pref**—Displays the number of routes that list that Local Pref among their attributes.
- **MED**—Displays the number of routes that list that MED among their attributes.
- **Communities**—Displays the number of routes that list that community among their attributes.
- **Neighbor AS**—Displays the number of routes that list that neighbor AS among their attributes.
- **2nd Hop AS**—Displays the number of routes that list that 2ndHopAS among their attributes.
- **Origin AS**—Displays the number of routes that list that origin AS among their attributes.
- **Any AS**—Displays the number of routes that list that AS among their attributes.
- **AS Peers**—Displays the number of routes that list that peer-pairing among their attributes.
- **IPv4 Prefixes/IPv6 Prefixes**—Displays a count of routes for that prefix among their attributes.

For the Peer, Nexthop, and Originator options, click a router in the list and the corresponding node flashes yellow on the routing topology map. Alternatively, click **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of times the link or prefix that are down. To view any of the lists as a bar chart, click **View as Bar Chart**.

To view additional information for a particular entry, right-click the entry, and then select one of the following choices on the pop-up menu:

- **Show Routes**—Displays a list of the routes that include that entry among their attributes.
- **Visualize**—Displays a visualization of the BGP tree as seen by the selected entity (see RIB Visualization on page 175).
- **Filter Analysis**—Displays a new RIB Browser window with data for the selected entity only.

VPN Protocol

For VPN protocols, the RIB Browser displays distributions of the advertised prefixes their attributes.

A tree structure on the left side of the window presents the following attribute options:

- **Peer**—For each peer, displays the number of routes advertised by that peer.
- **MP Nexthop—**For each MP NextHop, displays a count of routes that list that MP NextHop router among their attributes.
- **Local Pref**—For each Local Pref, displays the number of routes that list that Local Pref among their attributes.
- **MED**—For each MED, displays the number of routes that list that MED among their attributes.
- **Ext. Communities**—For each Extended Community, displays a count of routes that list that Extended Community among their attributes.
- **VPN Customer**—For each VPN Customer, displays a count of routes that list the RTs associated with that VPN Customer among their attributes.
- **Prefixes**—For each prefix, displays a count of routes for that prefix.
- **Route Distinguishers**—For each route distinguisher, displays a count of routes that list that Route Distinguisher among their attributes.
- **VPN Prefixes**—For each VPN prefix, displays a count of routes for that VPN prefix.

OSI IS-IS Protocol

For OSI IS-IS protocol, the RIB browser displays a tree structure on the left side of the window and presents the following attribute options:

If OSI IS-IS is not detected by the system, this window will not open.

- **Down Links**—Provides details of the number of links between the routers that are down.
- **Down Prefixes**—Displays the number of prefixes that are down in the topologies that are loaded.
- **Overloaded Routers**—Displays the routers that have their overload bit set.
- **Down ES Neighbors**—Displays the number of ES Neighbors that are down.
- **Down Prefix Neighbors**—Displays the number of Prefix Neighbors that are down.

For more information on using filters, see Using Filters on page 221.

RIB Comparison

Use this option to compare the state of the network at two points in time. This option is useful for analyzing the before and after state of the network when an unusually large number of events occur within a given period of time. Figure 13 provides an example of one such instance.

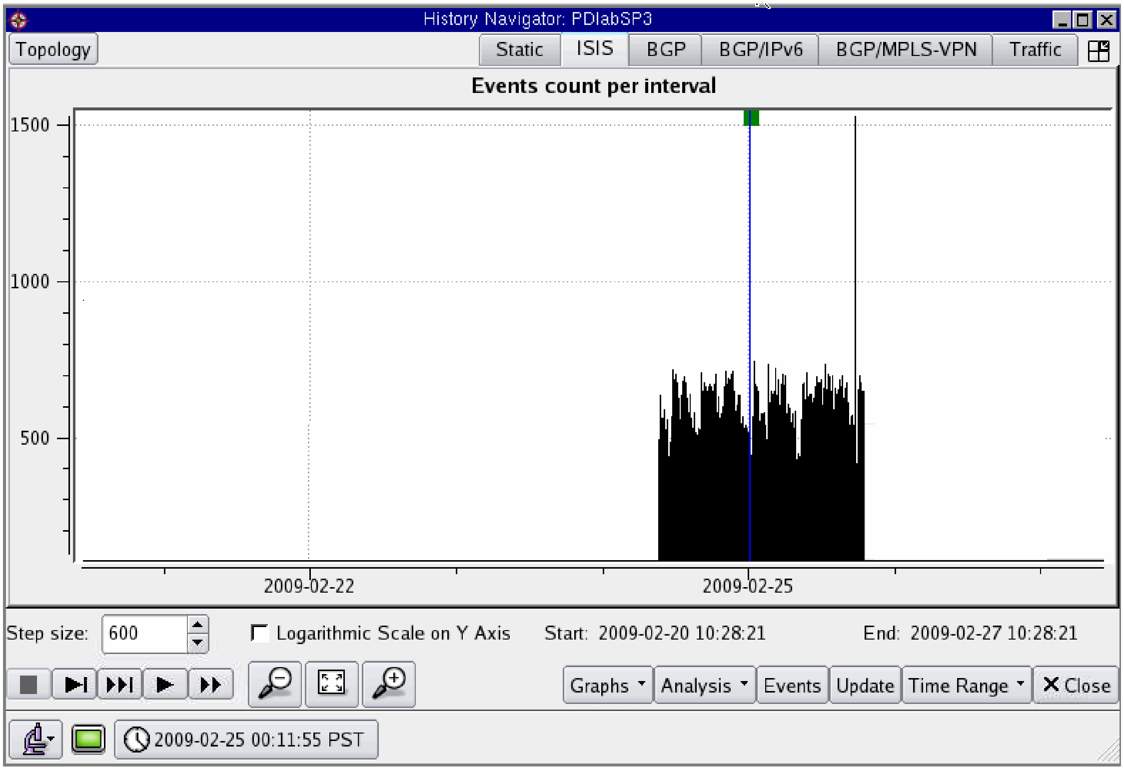

Figure 13 Large Number of Events in Short Space of Time

To analyze the state of the network just before these events occurred against the state of the network just after, use the RIB Comparison function.

To use the RIB Comparison, perform the following steps:

1. Choose **Reports > History Navigator** to open the main History Navigator window.
2. Choose **Analysis > RIB Comparison**.
3. Click in the graph just before the events occurred.

4 Click in the graph just after the events occurred.

The RIB Before-N-After Comparison window opens. It is similar to the RIB browser window.

Use the **Filter By** drop-down list to select filter parameters to restrict the input data included in the analysis. Click **Analyze Matching** to reanalyze the input data including only the items that match the filter criteria, or click **Analyze Excluding** to reanalyze the input data including only the items that do not match the filter criteria.

For example, the BGP RIB Comparison includes a report showing the difference in the number of routes advertised by each router at the two selected times. If you apply a Neighbor AS filter for a particular AS and then click **Analyze Matching**, the report shows the difference in the number of routes with that neighbor AS advertised by each router. Routers with no such routes are omitted from the list.

5 If multiple topology domains are included in the analysis, there are two levels of menus of the left. Select a domain from the first menu.

6 Select the type of information to view from the second (or only) menu on the left.

7 The table on the right redraws to show the selected information.

IGP Protocols

For IGP protocols, the RIB Before-N-After Comparison window includes columns for the link and prefix counts before and after the events, and also a column for the difference between the two. Figure 14 show an example of the RIB Comparison window with IGP link data.

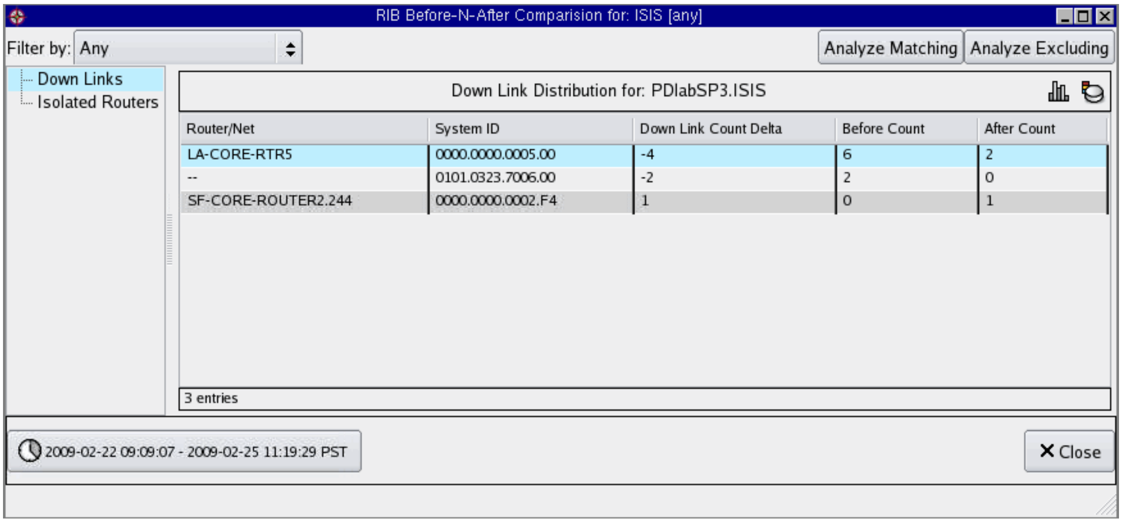

Figure 14 RIB Comparison for IGP

Click the **Down Link** and **Down Prefix** options to view information about the down links and prefixes, respectively. **Overloaded Routers** shows all the routers that have had their overload bit change between the two set time frames, along with the bit state at the beginning and end. **Changed Metrics** shows all the links whose metric values have changed between the time frames, along with the values at the beginning and end. To view the list of links or prefixes as a bar chart, click **View as Bar Chart**.

Click a router in the list and the corresponding node flashes yellow on the routing topology map. Alternatively, click **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of times the link or prefix that are down.

To view additional information for a particular entry, right-click the corresponding list entry and select one of the following choices on the pop-up menu:

- **Show Events**—Displays a list of the events reported by that entry. This window is similar to the Events window described in Understanding the Events List on page 200.
- **Filter Analysis**—Displays a new window with data for the selected router only.

BGP Protocols

For BGP protocols, the same options appear in the RIB Before-N-After Comparison window as those that appear in the RIB browser. In addition, **Before**, **After**, and **Delta** columns display the route count before and after the specified events, and the difference between the two counts.

Figure 15 displays the RIB Comparison window.

For the Peer, Nexthop and Originator options, click any entry in the list, and the corresponding node flashes yellow on the map. Alternatively, click **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of times the link or prefix went down. To view any of the lists as a bar chart, click **View as Bar Chart**.

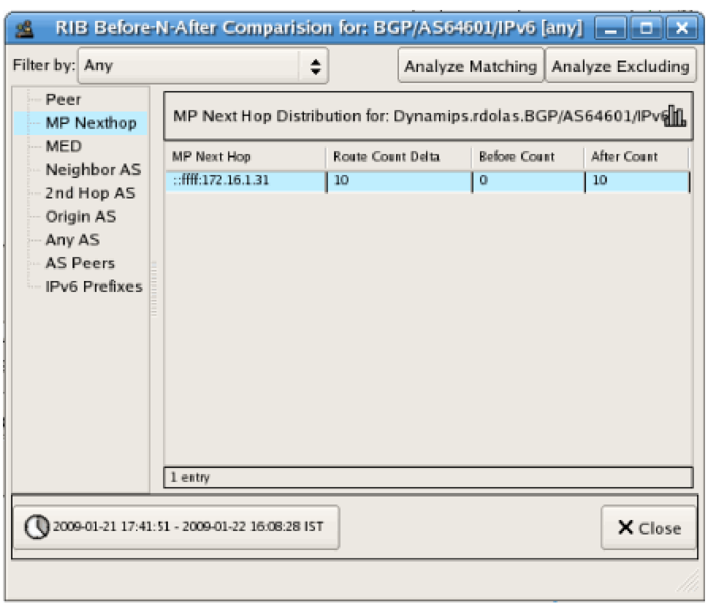

Figure 15 RIB Comparison for BGP

To view additional information for a particular entry, select and right-click the entry, and then select one of the following choices from the pop-up menu:

- **Show Differences**—Displays detailed information about the delta between the before state and the after state, based on the attribute selected in the RIB Comparison window. See Figure 16 for an example of the display.
- **Filter Analysis**—Displays a new RIB Browser window with data for the selected entity only.

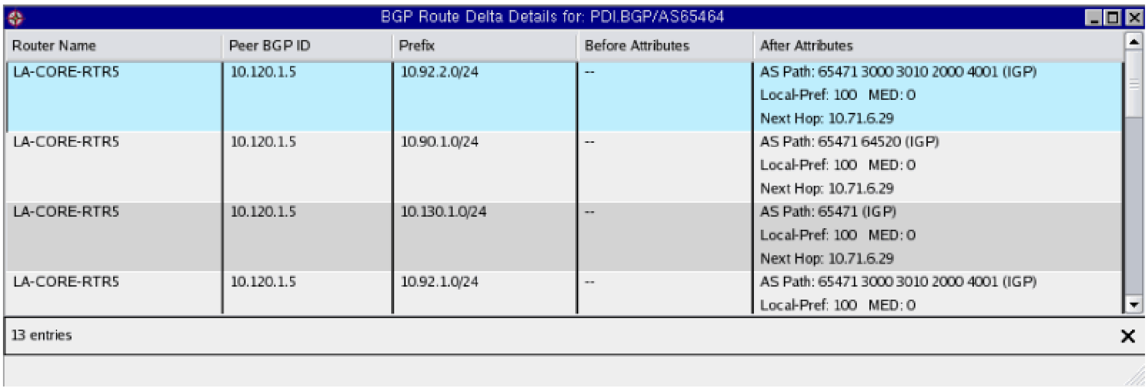

Figure 16 Show Differences Display

VPN Protocol

In addition to the options found for BGP, RIB Browser Before-N-After Comparison for VPN includes the following options: MP Nexthop, Ext. Communities, VPN Customers, Route Distinguishers, and VPN Prefixes. These options are described in the RIB Browser section for VPN Protocol on page 185.

The functions for RIB Browser Before-N-After Comparison for VPN protocol are the same as for BGP protocol (described on page 190) with the exception of the Animation feature, which is not included for VPN.

OSI IS-IS Protocol

As in IGP protocols, similar columns appear in the RIB Before-N-After Comparison window for OSI IS-IS.

If OSI IS-IS is not detected by the system, this window will not display.

Trending

Use this option to view aggregate event counts over an extended period of time. Figure 17 provides an example of one such instance.

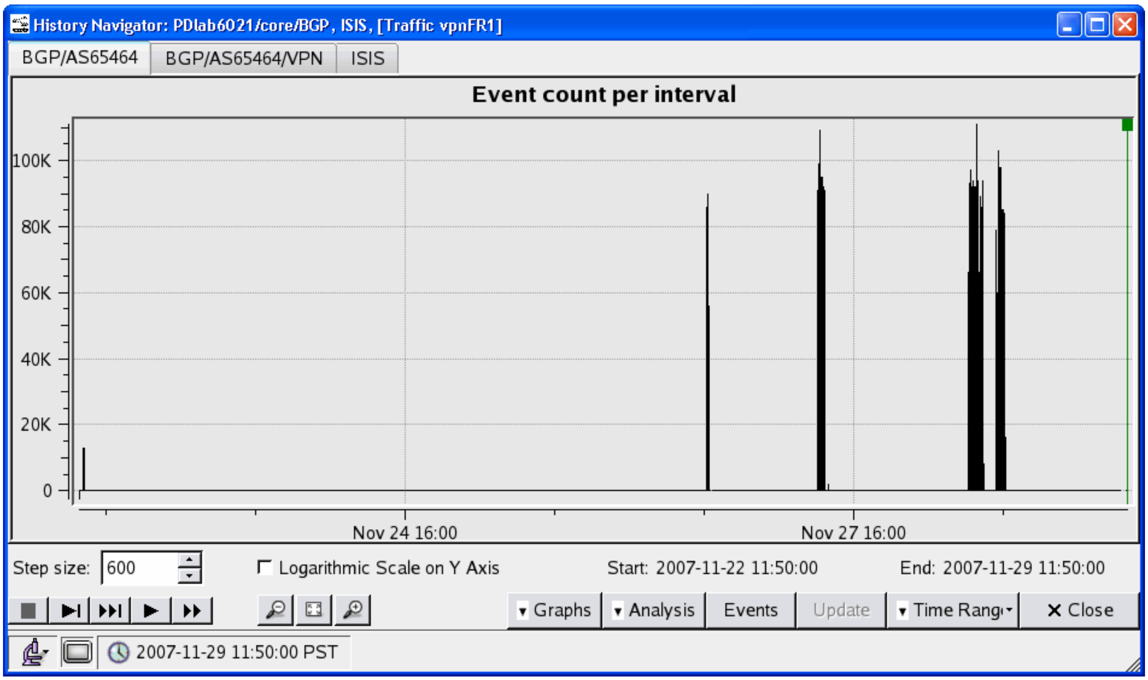

Figure 17 Event Trending

To display a trending graph, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose **Analysis > Trending**.

3 Select whether to display the trending graph in linear or exponential format, and choose the desired end date. You can enter values directly or click a component of the date and use the up and down arrows (Figure 18).

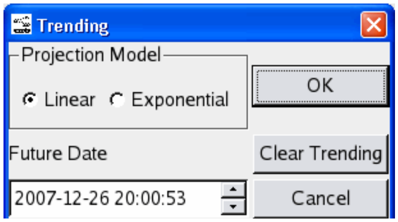

Figure 18 Trending

4 Click **OK** to display the graph.

Event Analysis

When a large cluster of routing events occurs, it may be difficult to grasp the nature of the problem by looking at individual events. The system can help you analyze the series of routing events to determine the distribution of events according to which routers, links, prefixes, and BGP attributes were involved. These distributions are presented as tables or bar charts.

In Monitoring mode, the default interval for updating the events graph is 10 seconds. If any events are generated during this time frame, a spike corresponding to the number of events is drawn on the graph.

In Analysis mode, each data point on the events graph shows the number of events that occurred in the previous 600 seconds.

To analyze a series of events, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose **Analysis > Event Analysis**.

3 Click in the Events graph just before the events occurred.

4 Click in the Events graph just after the events occurred.

5 (Optional) If you have more than 500k events, a window prompts you to **Continue**, **Abort**, or **Prefilter** the events.

6 (Optional) If you select **Prefilter,** the Event Prefiltering window opens. From here, you can select from a list of filters from the drop-down menu, decreasing the time it takes to generate the event list. For more information on using filters, see Using Filters on page 221.

The Event Analysis window appears.

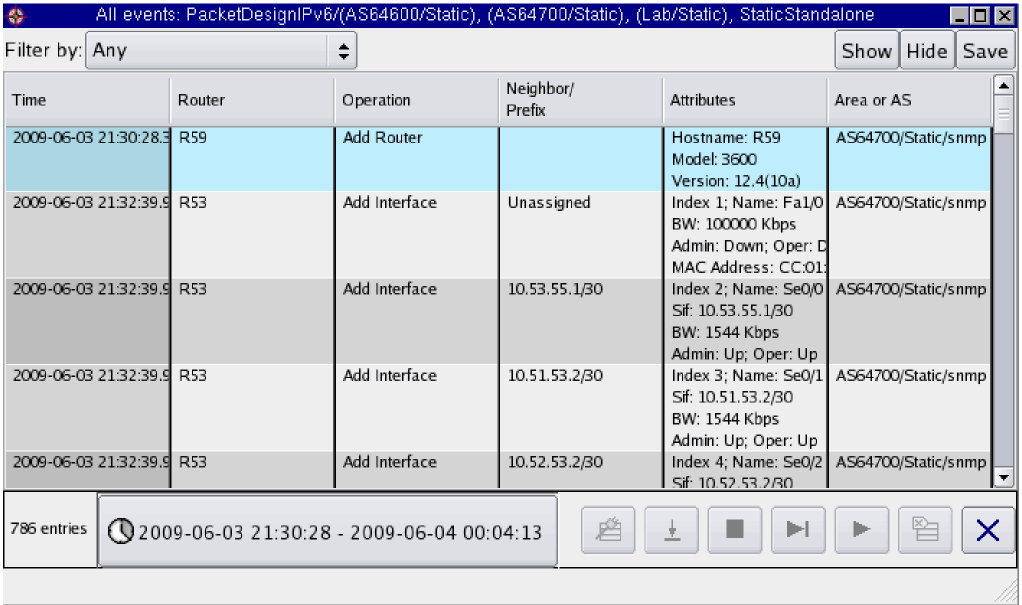

Figure 19 Event Analysis

Use the **Filter By** drop-down list to select filter parameters to restrict the set of events included in the analysis. Click **Analyze Matching** to reanalyze the events including only those that match the filter criteria, or click **Analyze Excluding** to reanalyze the events including only those that do not match the filter criteria.

For example, if you select a Prefix filter, specify a prefix that was flapping, and then click **Analyze Matching**, the analysis will include only the events that affected that prefix (Add, Drop or Change Prefix). Then, if you select the Initiator item from the tree at the left, you can see list of the routers that were sending those events and the number of events for each router.

7 If multiple topology domains are included in the analysis, there are two levels of menus of the left. Select a domain from the first menu.

8 Select the type of information to view from the second (or only) menu on the left.

9 The table on the right redraws to show the selected information.

IGP Protocols

For IGP protocols, this window displays the number of routing events that occurred in the specified time interval for each involved initiator, router, link, or prefix. Figure 20 shows an example Event Analysis window. The following options may be available, depending on the specific protocol:

- **Initiator**—Distribution of events according to the router that initiated the event.
- **Router**—Events that were initiated by that router plus other events that refer to that router (such as a neighboring router in an Add Neighbor event). If you specify a Router filter and click **Analyze Excluding**, that filter applies only to the Router column of the event list, which is where the initiator of the event is shown. The excluded router may still show up in the distribution of the Routers tab, however, because the excluded router was referred to as a neighbor.
- **Link**—Link-related events.
- **IPv4 Prefix/IPv6 Prefix**—Events involving the specified prefix type.
- **Messages**—Events involving the selected message type. The available types depend on the specific protocol.

To locate a router on the routing topology map, click the entry for that router. The selected entry is highlighted in the list and the router flashes yellow on the routing topology map. Alternatively, click  **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of events per router.

To view additional information for a particular entry, right-click the corresponding row in the table, and then select one of the following choices on the pop-up menu:

- **Show Events**—Displays a list of events reported by the selected entity. See Events List Controls on page 203 for information on the controls that allow you to replay the listed events.
- **Filter Analysis**—Displays a window with data for the selected entity only.

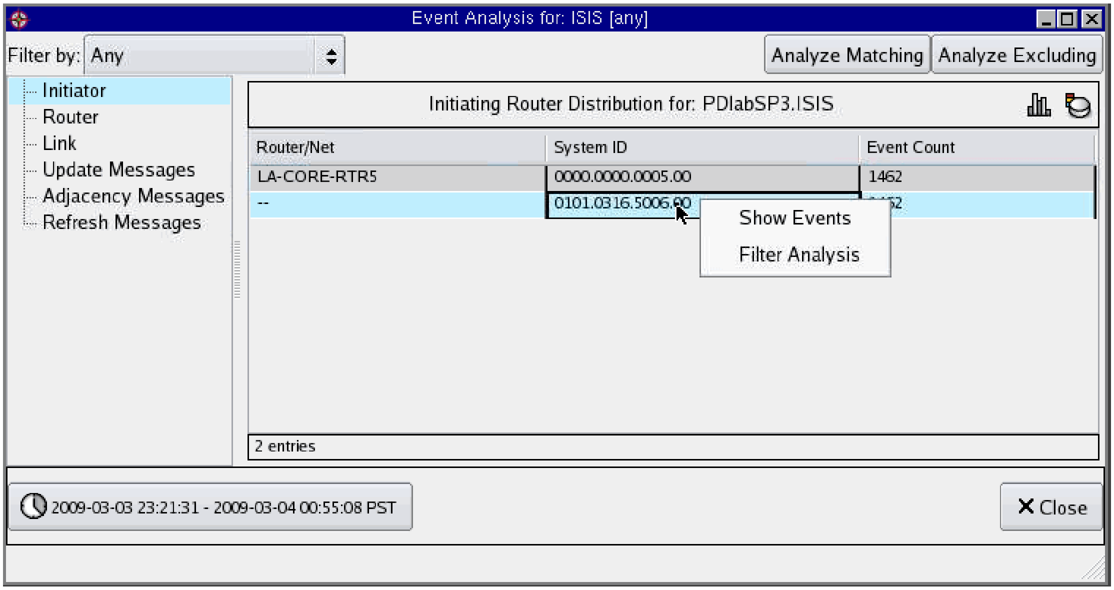

Figure 20 Event Analysis for IGP

BGP Protocol

For BGP, the Event Analysis window displays the number of routing events that occurred in the specified time interval with particular values for the Peer, Nexthop, Originator, Local Pref, MED, Communities, Neighbor AS, 2nd Hop AS, Origin AS, Any AS, AS Peers, or Prefix attributes.

To locate a router on the routing topology map, select the list entry for that router. The selected entry is highlighted in the list and the router flashes yellow on the routing topology map. Alternatively, click **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of events per router.

To view additional information for a particular entry, right-click the corresponding row in the table, and then select one of the following choices on the pop-up menu:

- **Show Events**—Displays a list of events reported by the selected entity. See Understanding the Events List on page 200 for information about this window.

- **Animate**—Displays an Animation window that animates the events reported by the selected entity. No Root Cause Analysis is performed. See Root Cause Analysis on page 167 for information about the controls that appear in this window.
- **Filter Analysis**—Displays a window with event data for the selected entity only.

Figure 21 shows an example of the Event Analysis window for BGP, and includes the pop-up menu that appears when you right-click an entry in the list.

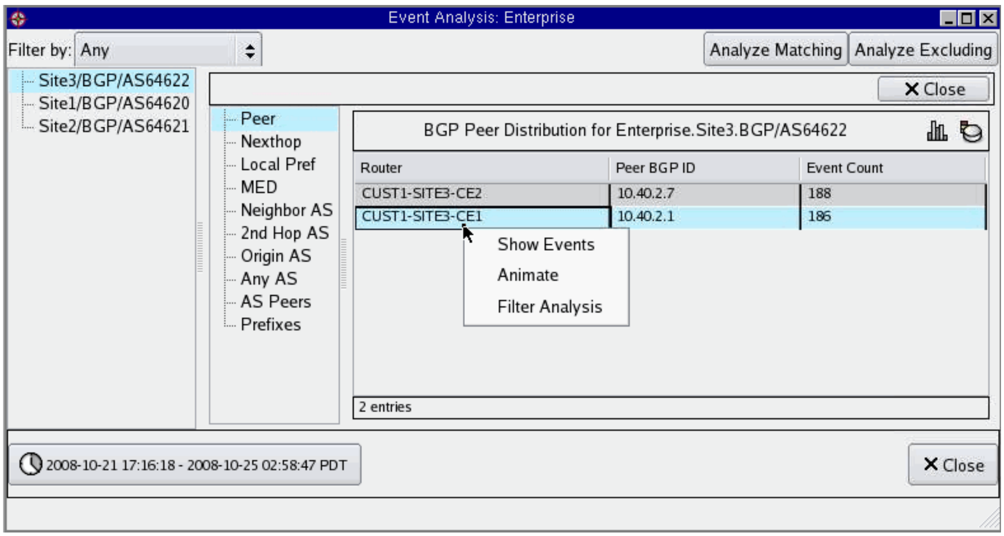

Figure 21 Event Analysis for BGP

VPN Protocol

In addition to the options found for BGP, Event Analysis window for VPN includes the following options: MP Nexthop, Ext. Communities, VPN Customers, Route Distinguishers, and VPN Prefixes. These options are described in the RIB Browser section for VPN Protocol on page 185.

The functionality for Event Analysis for VPN is the same as the Event Analysis for BGP, described on page 196.

OSI IS-IS Protocol

For OSI IS-IS protocol, the Events Analysis window displays the number of routing events that occurred in a specified time interval with particular values for the Initiator, Router, Link, Prefix, ES Neighbor, and Prefix Neighbor.

If OSI IS-IS is not detected by the system, this window will not display.

To locate a router on the routing topology map, select the list entry for that router. The selected entry is highlighted in the list and the router flashes yellow on the routing topology map. Alternatively, click **Color Routers** to color all the listed routers on the map at the same time using a spectrum of red to green for highest to lowest number of events per router.

To view additional information for a particular entry, right-click the corresponding row in the table, and then select one of the following choices on the pop-up menu:

- **Show Events**—Displays a list of events reported by the selected entity. See Understanding the Events List on page 200 for information about this window.
- **Animate**—Displays an Animation window that animates the events reported by the selected entity. No Root Cause Analysis is performed. See Root Cause Analysis on page 167 for information about the controls that appear in this window.
- **Filter Analysis**—Displays a window with event data for the selected entity only.

Flow Record Browser

This section applies to RAMS Traffic only.

Use the Flow Record Browser to display aggregated flow information.

The columns of information vary according to the type of aggregation chosen on the left side of the window. Choose a type from the side to display the flow information.

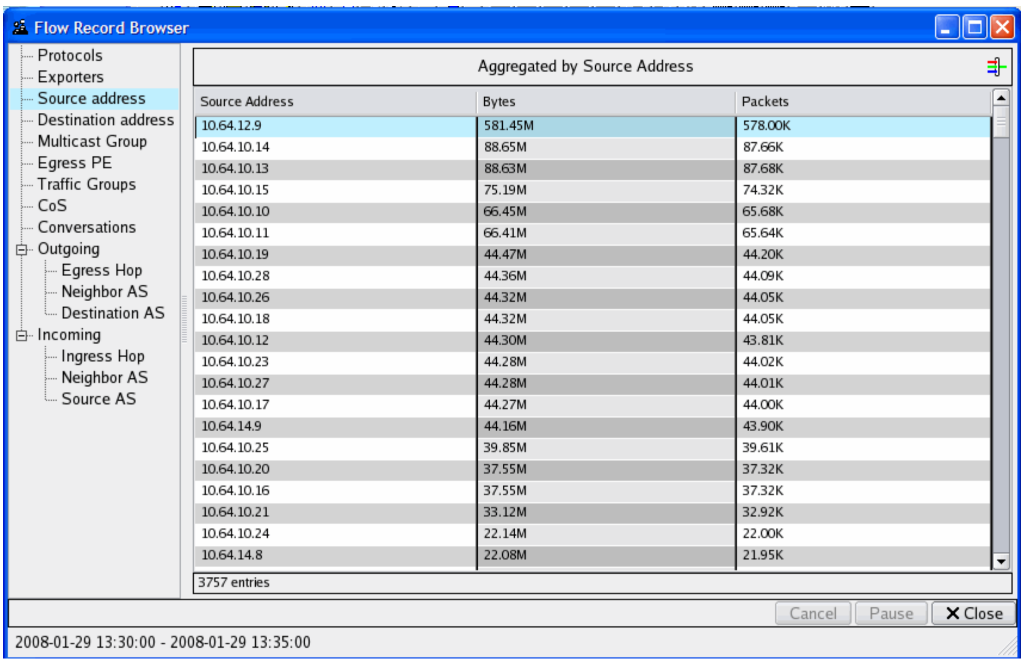

Figure 22 Flow Record Browser with Source Address Aggregation

To use the flow record browser, perform the following step:

1 Choose **Reports > Flow Record Browser** to open the Flow Record Browser window..

The following aggregation choices are available:

- **Protocols**– Aggregation by protocols involved in the flow. Select **Protocols** and then expand or contract the listing as needed.
- **Exporters**– Aggregation by IP addresses of exporters
- **Source Address**– Aggregation by source IP addresses
- **Destination Address**– Aggregation by destination IP addresses
- **Multicast Group**– Aggregation by groups configured as multicast groups
- **Egress PE**– Aggregation by IP addresses of egress PEs
- **Traffic Groups**– Aggregation by groups configured as traffic groups
- **CoS**– Aggregation by class of service (CoS) levels
- **Conversations**– Aggregation by active connections, indicated by source > destination IP address pairs
- **Outgoing**– Choices include:
  - Egress Hop– Aggregation by IP address for the egress hop
  - Neighbor AS– Aggregation by IP address of the neighboring AS
  - Destination AS– Aggregation by IP address of the destination AS
- **Incoming**–
  - Ingress Hop– Aggregation by IP address for the ingress hop
  - Neighbor AS– Aggregation by IP address of the neighboring AS
  - Destination AS– Aggregation by IP address of the destination AS

Understanding the Events List

After the general nature of a routing problem is identified, you may want to look at individual routing events to determine what caused the problem. The All Events list shows a sequential list of all routing events recorded in the database for a specified time interval. For each event, the list shows several columns of details, such as the router that initiated the event.

To view a list of individual events, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Click **Events**.

3 Move the mouse cursor, displayed as blue crosshairs, to the desired starting time in the graph and left-click to leave a blue line marking that time.

4 Move the mouse cursor to the ending time and left-click again to mark that time.

5 (Optional) If you have more than 500k events, a window prompts you to **Continue**, **Abort**, or **Prefilter** the events.

6 (Optional) If you select **Prefilter**, the Event Prefiltering window opens. From here, you can select from a list of filters from the drop-down menu, decreasing the time it takes to generate the event list. For more information on using filters, see Using Filters on page 221.

   The All Events window opens displaying details of the events that occurred within the selected time period, as shown in Figure 23. Detailed information, including information about OSPFv3 attributes for OSPFv3 networks, is shown in the Attributes column.

If the time period selected has a large number of events associated with it, a warning appears stating that the table will take time to load and may exceed memory capacity.

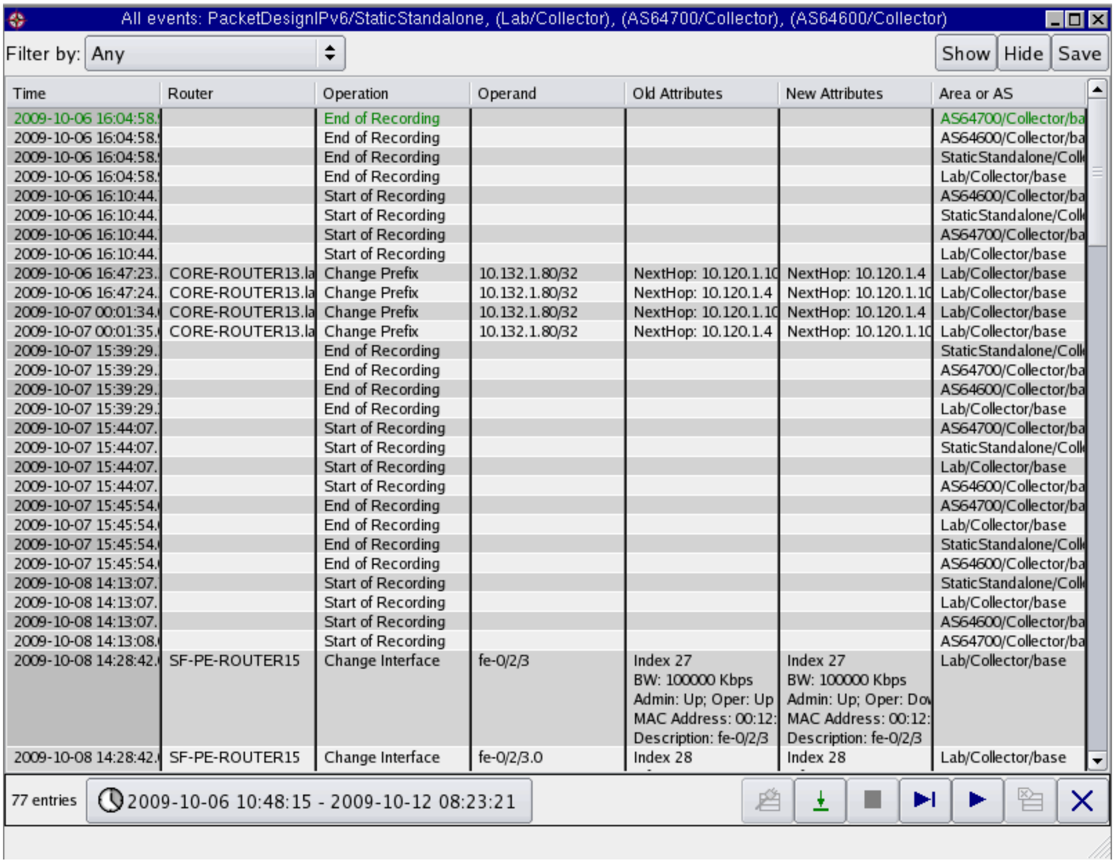

Figure 23 Events List

Events List Controls

Use the **Filter By** drop-down list and the **Show** and **Hide** buttons at the top of the Events window to filter the results displayed in the events list (see Filtering the Events List on page 209).

The panel of buttons toward the bottom center of the screen control zooming functions (see Table 4).

Table 4 Events List Controls

	**Time Range**	Opens the Select Time Range window, which allows you to change the time range covered by the Events list. To the right of the icon is a box that indicates the start and end of the current time range.
	**Online Update**	Refreshes the Events list with events that occurred within the past 10 minutes. This button is disabled when the History Navigator window is in Analysis mode.
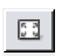	**Reset**	Resets the view back to the standard viewer setting.
	**Show Current Event**	See Moving Time and Executing Events on page 211.
	**Stop Execution**	See Moving Time and Executing Events on page 211.
	**Execute One event**	See Moving Time and Executing Events on page 211.
	**Start Execution**	See Moving Time and Executing Events on page 211.
	**Clear Events List**	Clears all events from the Events window. This button is only functional in Monitoring mode.
	**Close**	Closes the Events window.

Event Details

The entries shown in the events list are a generalized representation of the state changes communicated in the routing protocol.

For link-state protocols, these are adjacency changes for neighbors or prefixes that are carried in OSPF LSA packets or OSI IS-IS LSP. EIGRP is a distance-vector protocol, so it does not communicate link-state changes directly. However, the system determines what link-state changes caused the distance change, and inserts those link-state changes into the events list.

For BGP, peers communicate a stream of prefix (route) announcements, reannoucements, and withdrawals. In addition to the events indicating network state changes, entries are inserted in to the events list when the peering between the appliance and a neighbor router is established or lost.

Several state changes may be communicated at once within the routing protocol; these are displayed as separate events in the list, but all having the same timestamp. The timestamp is the first of several columns of details that are displayed for each event in the events list as described in Table 5 .

Table 5 Events List Columns

Name	Description
Time	Date and time of the event.
Router	The router to which the event is related. In OSPF and EIGRP, the router ID is displayed in dotted decimal. For OSI IS-IS, the router is identified by SysID, the unique value that is programmed into the router. The router name is shown for protocols that provide it.
Operation	The operation can be Add, Drop, or Change a Router, a Neighbor, a Prefix, or a RexPeering. For EIGRP, the operation can also be an EIGRP Update or an Unresolved EIGRP Change. For BGP, the operation may be Open or Close of the peering, or Announce, Reannounce or Withdraw of a prefix.

Table 5 Events List Columns (cont'd)

Name	Description
Neighbor/Prefix	Displays either the neighbor router for neighbor operations or the prefix for prefix operations.
Attributes	Displays the affected attributes of the router or prefix. This includes node isolation for IS-IS domains. A corresponding alert will occur once this isolation is detected.
Area or AS	The OSPF area, OSI IS-IS level, and EIGRP and BGP AS where the event took place. The areas and ASs are interleaved in this column.

The format of the **Attributes** column will vary depending on the protocol and the event type, but generally includes details such as the type of a router or the metric to a prefix or neighbor.

For example, starting at the first event in the list shown in Figure 23, router 24.0.0.11 changes its metric for prefix 192.168.116.0/24 to 1, and then, in subsequent Change Prefix events, changes its metric to 2 and back to 1. In the **Attributes** column, the type of the prefix, Area External, indicates that this prefix is being redistributed by router 24.0.0.11 in its role as an ABR. The highlighted Add Router event in the middle of the list indicates that a new router 192.168.0.72 of type Internal (meaning, not a border router) is being added to the routing topology. This event was implicitly generated as a result of the next event in which router 192.168.0.2, acting as DR for its subnet, added router 192.168.0.127 as a neighbor with metric 0. (The metric from a pseudonode to a router is always 0). About 20 minutes later, the adjacency was dropped and the router along with it.

If the Router Isolated alert is enabled, an alert is sent when all of the adjacencies in a specific area attached to a router are down. This action isolates the node from the rest of its area, and the router is no longer accessible from the connected area where the peering resides.

The following information is included in the isolated router alert:

- Time and date the router was isolated.
- Specific area in which the router isolation occurs.
- Information on the router that was isolated

When a Router Isolated alert is received, open the routing topology map to view the effect of the isolation. The History Navigator window is also useful for bringing up the Events List for the network, which will display information about adjacency losses that cause router isolations in the **Attributes** column.

To configure router isolated alerts, see Creating New Alerts on page 508.

Event Operations and Attributes

There are many different possible combinations of event operations and attributes. While the event list format is generalized to allow a consistent representation in multiprotocol networks, there are some protocol-specific characteristics due to the differences in nomenclature and behavior of the protocols:

- OSPF: In the **Router** column, the letters DR following a router address or DNS name indicate events pertaining to the pseudonode that represents a LAN subnet. These letters may appear in the **Router** column for events originated by the Designated Router in its role as the DR for the LAN (versus its role as an individual router). The letters may also appear in the **Neighbor/Prefix** column for events for which that column lists the neighbor router, such as an Add Neighbor event, which indicates that the router sending the event has added an adjacency to the pseudonode represented by the DR.

  The router types are Internal, ABR, Autonomous System Border Router (ASBR), a combination of these, or LAN Pseudo-Node. The prefix metric type is Internal if not explicitly identified as one of Area External, AS External Comparable (Type 1), or AS External (Type 2). The **Attributes** column for a Drop Prefix or Drop Neighbor event may indicate “Cause: Expired” if the prefix advertisement of the router has timed out without a refresh, or “Cause: Premature” when the router advertises a graceful withdrawal (for example, on shutdown). For protocol details, see Appendix A, “Protocol Compliance.”

- OSI IS-IS: Since OSI IS-IS is a link-state protocol such as OSPF, the event list details are similar. Routers are identified by a 7-byte hexadecimal SystemID in the form C0A8.00E0.0000.00, or by a name communicated within the protocol. The 7th byte of the SystemID is non-zero when a router is acting as the Designated Router for a LAN. Different values of this byte distinguish different subnets. In the **Router** column the Designated Router is indicated by the SystemID followed by “DR” or by the router name followed by a period, the hexadecimal subnet byte, and “DR.” The system labels an OSI IS-IS router that is just in level 1 or level 2 as “Internal,” while a router that participates in both level 1 and level 2 as “Area BR.” Nodes representing subnets are labeled “LAN Pseudo-Node.” The prefix metric type is Internal unless explicitly identified as External or TE. For protocol details, see Appendix A, “Protocol Compliance.”

- EIGRP: Since EIGRP is a distance-vector protocol, the only routing events recorded directly from the protocol are EIGRP Update events, which tell the distance from one of the peer routers to a prefix. These events are obtained from EIGRP Update and EIGRP Query packets.

  The distance is measured in the EIGRP metric with two components:

  — Inverse bandwidth (bw)

  — Delay (dly).

  The prefix metric type is Internal, if it is not specified. For External prefixes, the originating router is identified. Several special-case prefix types are identified:

  — Loopback, the prefix of a loopback interface

  — Dialup, a /32 prefix that is contained within a less-specific prefix advertised by the same router

  — Auto-Summary and Manual Summary

  — Static prefixes that are redistributed in EIGRP

  The system analyzes the EIGRP Update events to determine what link-state changes caused the EIGRP distances to change, then issues CLI queries to the affected routers to verify the change. One or more link-state events are then synthesized and recorded in the routing topology database.

  The basic link-state events have the same format as OSPF and OSI IS-IS events: Add/Drop of Router/Neighbor/Prefix. In addition, for the EIGRP protocol the database records other changes in the routing configuration that are learned through CLI queries to the routers: Add/Drop of Route Filter, Route ACL, or Static Route. These events are interspersed with the EIGRP Update events. Since the analysis may take tens of seconds, the link-state events will appear later in the events list than the EIGRP Update events.

In case the analysis of EIGRP Updates cannot determine what link-state change was the cause, an Unresolved EIGRP Change event will be written to the database. The **Attribute** column gives the reason:

— Unknown path—Due to the nature of the EIGRP metric, it is possible, although rare, that the changed state of a link not on the shortest path between two end points will affect the choice of that path. The system cannot infer the link change in this case.

— Not on old path; new path broken—If the internal routing topology model provided by the system has become inaccurate, perhaps due to a previous Unresolved EIGRP Change, the analysis algorithm may not be accurately locate the shortest path between two nodes. If this happens, The system will not be able to infer a link failure that partitions some nodes from the viewpoint of the appliance.

— Query failed—A query by the system to the problematic node failed, perhaps because the node itself became unreachable or was busy, so the link state is unknown.

— Unexpected value—The analysis algorithm lost track while inferring a topology change. This may happen for various reasons; for example, while tracing a changed route, the route may change again.

— Fast route flap—The link state changed back before the change could be verified to have existed.

When the system detects a route that appears to be stuck in active state, it follows the stuck route until it gets to the last responding router before the nonresponding router. The table entry for this event (EIGRP Stuck in Active) identifies the router that is waiting for the nonresponding router to communicate. The Attribute column reports the statistics about the nonresponding router that are obtained using the **show ip eigrp neighbor** command issued on the last responding router on the route. The cause of this event could be that the nonresponding router is down but has not yet been reported down by its neighbor.

- BGP: The set of event operations for BGP is small: Open or Close of a peering, or announce, reannounce, or withdraw of a prefix. However, the number of different attributes is much larger than for IGP events: AS Path, Local-Pref, MED, Communities, Next Hop, Originator ID, Cluster List, and Aggregator. For protocol details, see Appendix A, “Protocol Compliance.”

In addition to the protocol-specific events outlined above, there are Add Peering and Drop Peering events that indicate when the system established or lost peering with its neighbor router. Routing topology changes cannot be recorded when the peering is lost.

Highlighting Associated Nodes

Selecting an event in the list highlights that entry in reverse color, as shown for the Add Router event at 09:43:25 in Figure 23, and also causes any associated nodes to flash on the routing topology map. (If the map is displaying the routing topology at a different time than the time of the event, it is possible that no nodes will flash, because the associated nodes are not present.)

Filtering the Events List

When many events occur during a period of interest, it may be difficult to isolate the events relevant to a particular problem. To make finding the desired events easier, the displayed list of events may be filtered by a wide range of criteria, which differ depending on the protocol represented by the current tab of the History Navigator window from which you generated the Events list.

Use the **Filter By** drop-down list to select the filter parameters and the **Show** or **Hide** buttons to list only those events that match the filter criteria, or exclude those events that match the filter criteria, respectively.

If you have more than 500k events, the system will display a window prompting you to prefilter. If you select Prefilter, you can select from an array of filters to help cut down processing time. See Step 5 in Understanding the Events List on page 200 for more information.

Some parameters require that you enter a value in a text box (for example, if you filter by router, you must enter the name or IP address of the router in the text box to the right of the Filter By list). Other parameters require that you choose one or more items from a list (for example, if you filter by event operation, you are presented with a list of event types from which to choose).

Using Filters on page 221 explains how to combine filter parameters using the **Expressions** option on the filter drop-down list.

See Using Filters on page 221 for information about how to compose complex filters.

Alternatively, you can focus in on events related to a particular node or link on the routing topology map. Right-click the object of interest to display the node information panel or link information panel (Node Information Panel on page 82 and Link Information Panel on page 84, respectively), and then click **Events** on the information panel. A new Events List window is displayed showing only the events originated by the selected router or related to adjacency changes on the selected link.

Adjusting the Time Range

The initial time range for the All Events list is selected by setting the blue lines on the History Navigator Events graph, as described in Steps 3 and 4 in Understanding the Events List on page 200. You can adjust the time range as needed.

In Analysis mode, you can double-click the date/time area to display the playback controls. Choose and date and time and a step size in sections for the playback, and click **OK**. Then click the right-facing arrow to begin the playback.

Figure 24 Main Window Status Bar

To adjust the start and end of the time range, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Adjust the time range in any of the following ways:

- Click **Time Range > Online** to view current data in a moving 1-hour window.
- Click **Time Range > Custom**. Enter new values into the **From** and **To** fields, or adjust the values with the up and down triangle buttons. Click **OK**.
- Choose **Time Range** and select a pre-set interval: one hour, day, week or month. The time range will be centered around the currently displayed point in time.

- Choose **Time Range > Recent** to display a drop-down list of recently used time ranges from which you can select.
- Choose **Time Range > All** to include all events recorded in the database.

In RAMS Traffic, traffic data may not be available if the selected time is within 30 minutes of the current time.

3 Click **OK** to accept the adjusted time range.

If the time period selected has numerous events associated with it, a warning appears stating that the table will take time to load and may exceed memory capacity.

Moving Time and Executing Events

The current time for the routing topology map may be moved to the time of any event in the list so that the map shows the state of the network at the time just before the event occurred.

If you right-click an event in the list, its text temporarily changes to blue, and a pop-up dialog box asks if you want to move time to before or after that event. When you choose an option, the event text changes to green to indicate that it is the next event to be executed. In the example events list shown in Figure 23, the next event is the Add Router event at 09:43:25.

Click **Start Execution** to execute events one after another starting with the next event and continuing to the last event in the Events list, and observe their effect on the network. During execution, the routing topology map marks nodes or links that go DOWN as a result of event execution with a red cross (✕), while nodes and links that change state to UP are marked with a green dot (•). When an EIGRP Update event is executed, indicating a change in the distance to a prefix, the routers to which that prefix is attached are marked with a blue dot (•). As each event is executed, the text for the next event in the list turns green and the current time for the routing topology map advances as shown by the green time cursor moving to the right on the Events graph in the History Navigator window. To stop the execution, click **Stop Execution**.

Click **Execute One Event** to execute events one at a time and observe their effect on the network.

Conversely, you can drag the time cursor to any point of interest on the time line. This displays the state of the network corresponding to that point in time in the routing topology map. There are three possible situations:

- If the time cursor is within the time range covered by the events list (between the blue lines), you can click **Show Current Event** to quickly find the next event to be executed. The next event, highlighted by green text, scrolls to the top of the list.
- If the time cursor is earlier than the start of the time range of the events list, the next event to be executed is the first event in the list. The time cursor jumps to the time of the first event if it is executed.
- If the time cursor is later than the end of the time range, the **Show Current** and **Execute** buttons are disabled and no event is highlighted by green text.

Using the History Navigator as a Forensic Tool

When diagnosing a network outage or performing forensic analysis after an outage, having complete historical data and analysis capability is invaluable. The RIB Comparison on page 186 showed how the History Navigator window displays event churn in a time line and analyzes the state of the RIB before and after network churn. This section provides an example of the steps you can take to narrow down the event churn to its root cause.

In the example shown in Figure 25, a period of instability (a high level of churn) lasts for more than an hour. Using the History Navigator Event Analysis tool, you can focus on a small part of the total churn period.

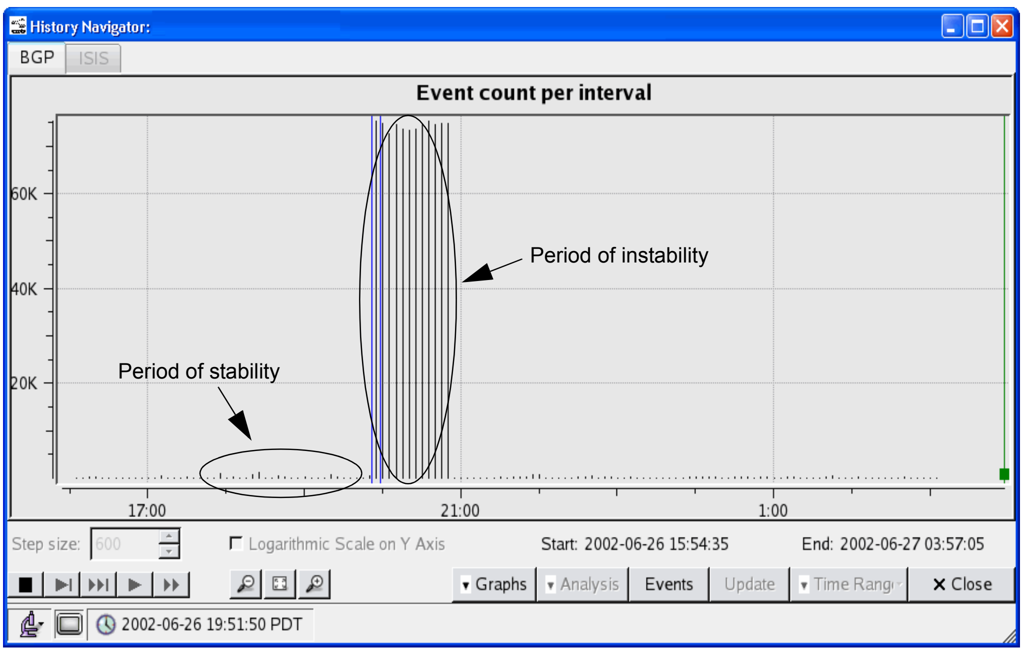

Figure 25 Stability and Instability

To perform an events analysis, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose **Analysis > Event Analysis**.

3 Mark the start and end time for the analysis using the blue cross-hairs.

The Event Analysis window appears, as shown in Figure 26.

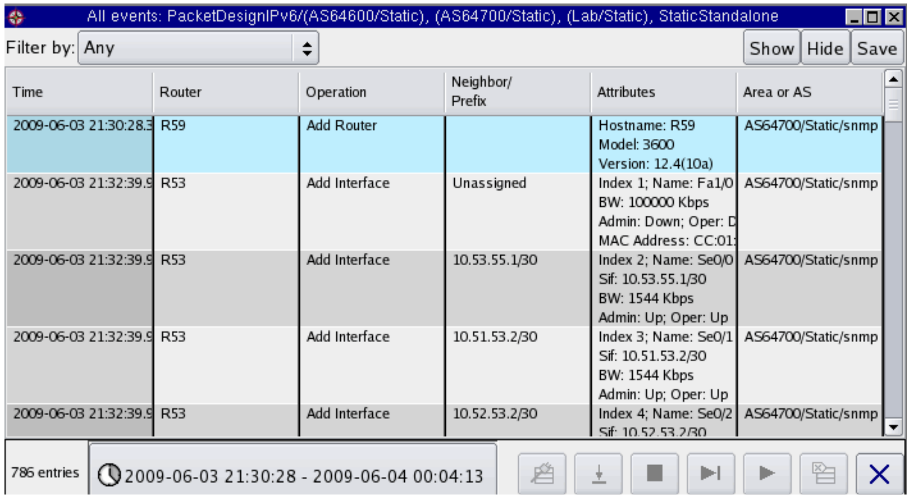

Figure 26 Event Analysis Window

The Event Analysis table can be filtered, sorted by column heading, or viewed as a bar chart.

When the **MED** option is selected, as shown in Figure 27, a small number of MEDs may have a large number of events associated with them. This could represent a MED oscillation.

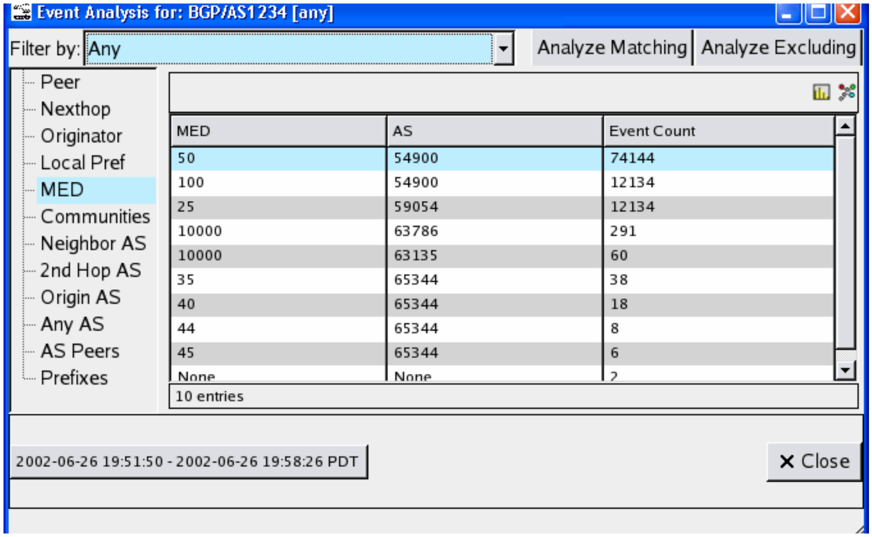

Figure 27 MEDs

To identify the prefixes affected by this possible MED oscillation, select the **Prefixes** tab as shown in Figure 28.

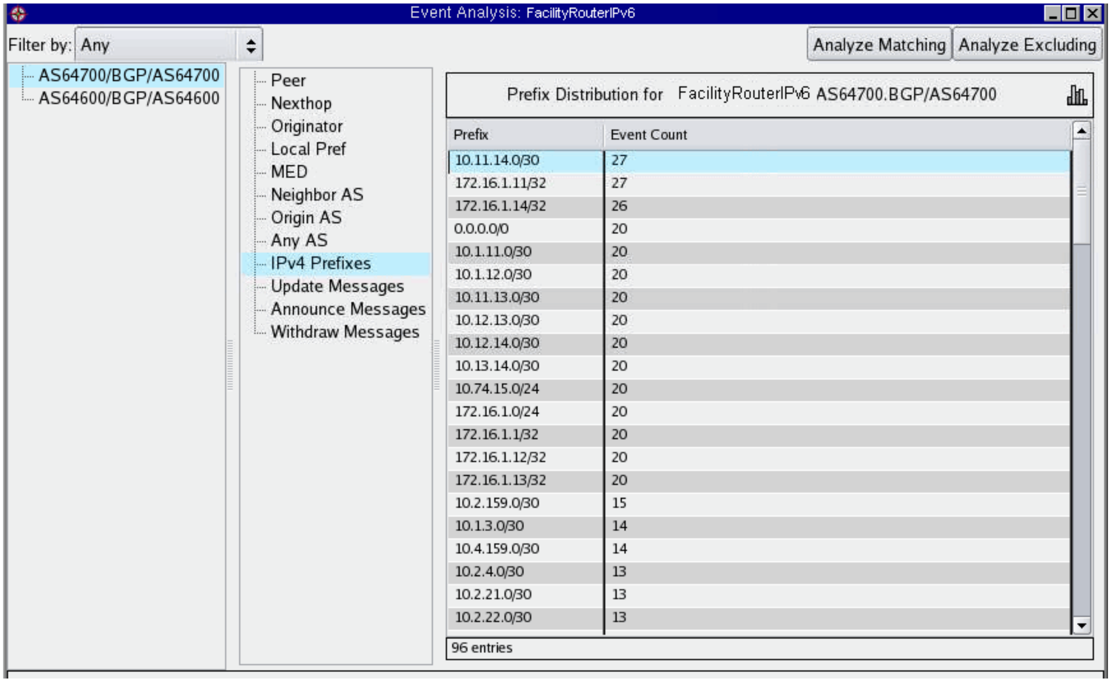

Figure 28 Prefix Details

A single prefix has a numerous events associated with it. You can drill down to determine which BGP peers have generated these events by filtering the analysis to include just these events, and then observing the peers involved.

To drill down and view details, perform the following:

1. Choose **Reports > History Navigator** to open the main History Navigator window.
2. Choose **Analysis > Event Analysis**.
3. Specify a time range (see Adjusting the Time Range on page 210).
4. Right-click the desired prefix.
5. Click **Filter Analysis** in the pop-up window that appears.

Figure 29 displays the results of the drill-down filter analysis.

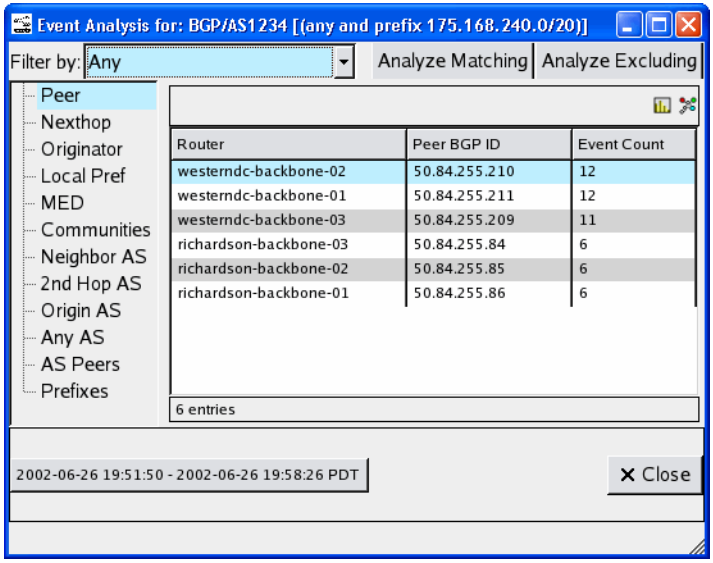

Figure 29 Filtered Event Distribution

It appears that five peers have generated the majority of events. This increases the suspicion of a MED oscillation. To confirm this suspicion, you should look at the actual events in question in
more detail.

Many routing instabilities are caused by interactions between multiple routers and are very difficult to isolate because routers do not keep an event history. Diagnosis of the outage can require login to multiple routers and the execution of show ip bgp...commands, a very tedious and time-consuming task.

The RIB analysis identified the possibility of a MED oscillation, and the Event Analysis identified the exact prefix and the peers involved in the oscillation. The following procedures show how you can confirm the exact cause of the problem by looking at the events list.

To view the events associated with a particular problem, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose **Analysis > Event Analysis**.

3 Select the desired start and end times with the blue cross-hairs.

4 To view the events associated with an individual table entry, right-click the entry, and then select **Show Events** in the pop-up menu.

The Filtered Events window appears, as shown in Figure 30. This window lists all of the events associated with the selected table entry.

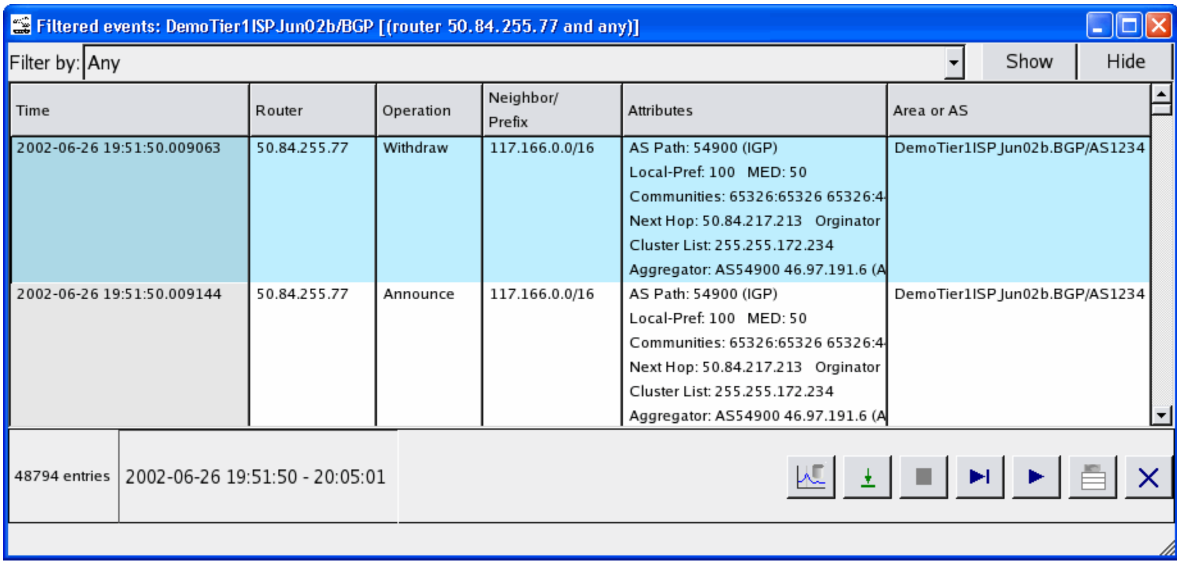

Figure 30 Filtered Events Window

The final step is to use the Root Cause Analysis function to distill this event information down to its root cause and display an animation of the events, so you can visualize the events as they occurred.

To perform the root cause analysis, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose **Analysis > Root Cause Analysis**.

3 Select the desired start and end times with the blue cross-hairs. If you have multiple BGP topologies, the system prompts you to select a single topology.

The Root Cause Analysis Results window opens, as shown in Figure 31. For more information on using root cause analysis, see Root Cause Analysis on page 167.

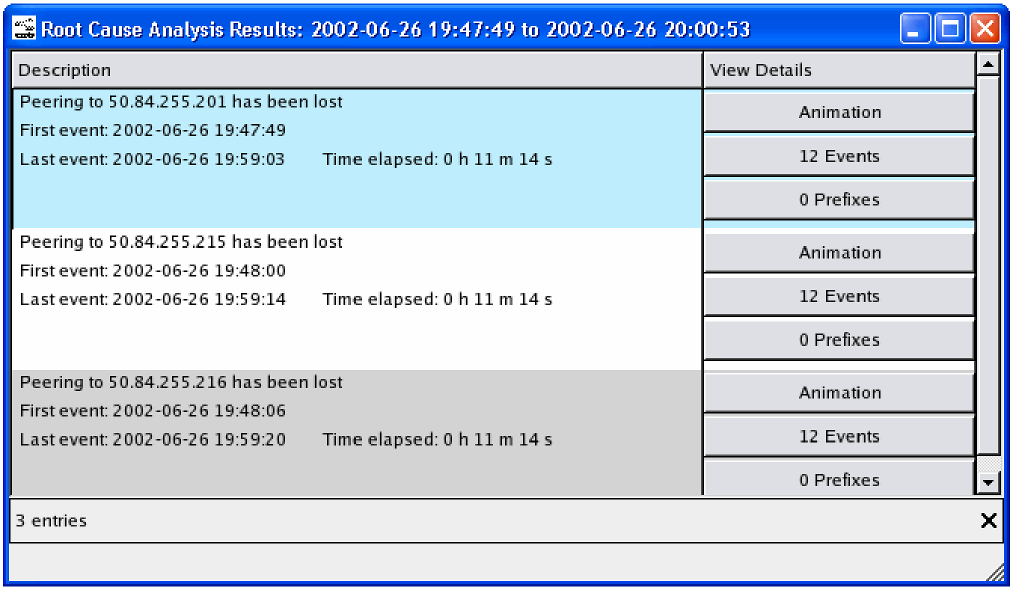

Figure 31 Root Cause Analysis Window

Correlating Time Series Data

The power of routing navigation is extended by letting you import and display external time series data such as link utilization or jitter measurements in correlation with the state of the network at the routing layer. This makes it easier to identify the cause and effect of events visually by having all of the data on one screen. As routing data is played back from the database to visualize changes in routing, a time cursor simultaneously steps along the time line graph of the external time series data. The time series data correlation feature provides an unprecedented visualization and analysis capability. Intermittent and intractable problems can be approached from a new perspective and analyzed within seconds or minutes, rather than hours.

One popular source of time series data is the Multi Router Traffic Grapher (MRTG), a free tool that monitors the traffic load on network links. MRTG generates HTML pages that contain PNG graph images providing a live visual representation of this traffic. You can use MRTG graphs as a way of monitoring the health and status of your network. When an anomaly appears, importing the graph data and performing a time correlation with routing events can help diagnose the root cause of the anomaly.

You can import data in MRTG ASCII `.log` file format, Round Robin Database (rrd) `.rrd` binary format version 1 or 3 little-endian architecture, RRD xml dump format, and a simple, generic ASCII time series format called graf file format. The graf format consists of two floating-point numbers on each line; the first is the time coordinate and the second is the data value corresponding to that time. The first line of the graf file can optionally be a '#' character followed by a title for the graph.

Any external event data (jitter, packet loss, traffic statistics, server statistics, and so on) can be viewed if it can be transformed with text processing tools into graf format. Data exported from a two-column spreadsheet in tab-separated `.csv` format is one suitable source.

To view a time series data file, it must be uploaded to the appliance. The administrator must first enable the File Transfer Protocol (FTP) server. See "Administration" in the *HP Route Analytics Management Software Administrator Guide*.

To upload files to the appliance, perform the following steps:

1. Use FTP or secure FTP (SFTP) to the appliance using the IP address of the appliance.
2. Log in with your user name and password.

   The administrator must enable FTP access for your account.

3. Change the directory to `pub`.
4. Transfer one or more files to the appliance.

To display a time series file (all formats), perform the following steps:

1. Choose **Reports > Correlate Time Series**.

2 The available files are listed on the left side of the window. Select a file to display the time series data.

The green cursor in the time series graph is time-aligned with the cursor in the History Navigator window and any other time series graphs already displayed. Moving the cursor in any displayed time series graph moves it in the others. If you move the cursor, the routing topology map updates to display the state of the network at the time indicated by the cursors.

If the time interval covered by a graph does not include the point in time chosen by moving a cursor in another window, the cursor for the first graph will be positioned either at the beginning or end of its timeline, whichever is closer to the chosen time. In this case, the cursors will not all be positioned at the same time.

When displaying MRTG data, note that MRTG files contain four datasets for the average and maximum bytes/second input and output on a network interface. The four datasets are displayed together in one graph window.

Using Filters

A filter option is provided on several tables, including the RIB Browser and Events List, to allow you to focus in on items of interest. Figure 32 shows an example of the **Filter by** drop-down list on the Events window for BGP. Note that the items in the list differ depending on the current routing protocol.

The following filter levels are available:

- Simple filters let you choose a single operator (for example, “router”) from a list and specify one or more parameters (for example, router IP addresses or names) to be matched or excluded.
  See Expression Definitions on page 228 for examples of the parameter syntax as illustrated using filter expressions.

  With a simple filter, you enter only the parameter; the operator is selected from a list.

The filter is translated internally into a filter expression combining the filter operator with the parameters. Figure 33 shows an example of a filter specifying a router address.

- Advanced filters let you choose two or more different operators from a list and specify their corresponding parameters to be matched or excluded.
- Filter expressions let you manually enter a filter expression that is too complex to be set up with either simple or advanced filter menus.
- You can also pre-filter events for the following features:
  - Root Cause Analysis
  - Event Lists
  - Event Analysis

  This option is prompted if you have more than 500k events for the system. Using this pre-filter will cut down on the time it takes to generate the information you are looking for.

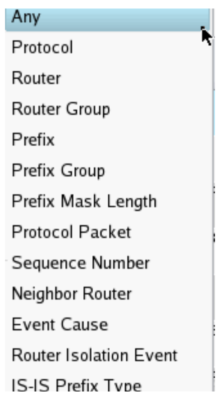

Figure 32 Filter By Drop-Down List

In many cases, the built-in **Filter by** selections, such as the ones shown in Figure 32, provide sufficient flexibility in filtering.

The Router filter accepts a space-separated list of router addresses or names when several specific routers are of interest. The Community filter accepts a space-separated list of community strings when prefixes with different community strings are of interest. The filter matches if any of the community strings matches (OR relationship). The Sequence Number filter accepts LSP packet sequence numbers.

The protocol related filters, such as BGP Ingress Router, allow you to specify the IP address of the node.

To set up a simple filter, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose the report of interest.

3 Select an operator from the Filter by drop-down list.

A text box is shown on the right of the Filter by parameter.

4 Enter the address of the node in the text box.

For example, if you want to see only the events that are reported by the node with IP address 192.168.167.166, choose **Router** in the Filter by drop-down list and then enter the IP address in the text box (Figure 33).

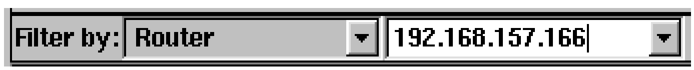

Figure 33 Parameter Text Box

5 Click **Show** to list only items that match the address, or click **Hide** to list only the items that do not match the address.

To set up an advanced filter, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose the report of interest.

3 Select **Advanced** from the Filter by drop-down list.

The Composing Advanced Filter window opens, as shown in Figure 34.

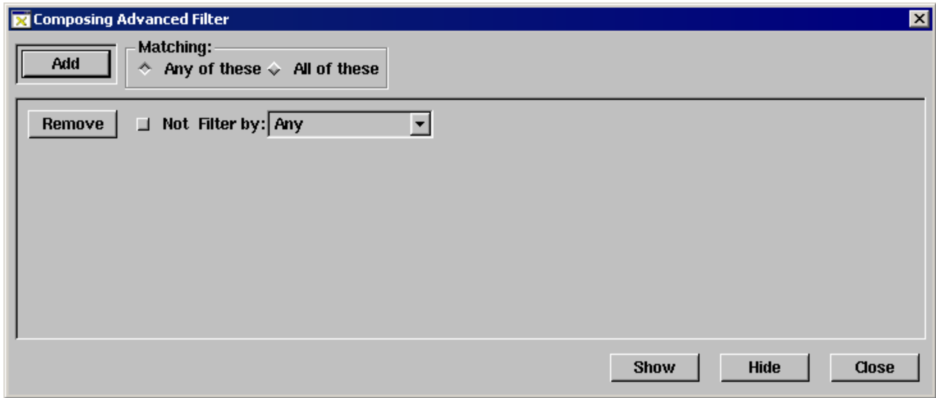

Figure 34 Advanced Filter Window

4 Choose a filter operator from the Not Filter by drop-down list.

The Remove button removes the Not Filter By field.

Some operators require that you enter the parameter in a text box; others let you choose among items in a list.

5 Specify the appropriate parameter value.

6 To exclude matching items, click **Not** for that operator.

7 To add another operator to the filter, click **Add** in the upper left corner of the window and define the parameters for the new operator.

8 After you add the desired filter operators (and their corresponding parameters), choose an option from the Matching box:

- **Any of these** includes an item if it matches any one of the filter criteria.
- **All of these** includes an item only if it matches every filter criteria.

9 Click **Show** to list all events that match the filter, or click **Hide** to list only the events that do not match the filter.

The filter is translated internally into a filter expression that combines the filter operators with the parameters that you specified.

You can combine multiple levels of advanced filters to construct any logical AND-OR expression.

To enter a filter expression manually, perform the following steps:

1 Choose **Reports > History Navigator** to open the main History Navigator window.

2 Choose the report of interest.

3 Choose **Expression** from the Filter by drop-down list.

4 If desired, select a filter from the Custom Filters drop-down list, which will then populate the adjacent text box as shown in Figure 35. Otherwise, go to Step 3. You can enter and save filter expressions in the Custom Filter Repository as described in Creating Custom Filters on page 101.

5 If necessary, modify the selected expression or enter a new expression in the text box.

See Expression Syntax on page 225 for information about the syntax used to enter expressions, and Expression Definitions on page 228 for a complete list of operators and examples showing their use.

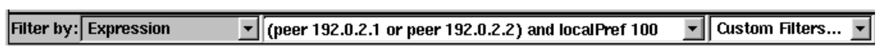

Figure 35 Expression Text Box

6 Click **Show** to display only those items that match the filter, or **Hide** to display only those items that do not match.

Expression Syntax

Filter expressions are specified in prefix notation, which means that the filter operator must be placed first with the parameter coming after the operator. An expression may include multiple terms (operators and parameters).

The following syntax rules apply:

- Operators are case-insensitive. Mixed capitalization is used in the examples for clarity.
- Operators and parameters are separated by whitespace.
- Operator `not` has higher precedence than operator `and`, which in turn has higher precedence than operator `or`.
- Parentheses may be used when needed to group subexpressions and override the precedence of operators.

Examples

The following expression is equivalent to selecting the **Router** option of the Filter by menu and supplying the three addresses 10.1.1.1, 10.2.2.2 and 10.3.3.3:

```
router 10.1.1.1 or router 10.2.2.2 or router 10.3.3.3
```

The following expression on the RIB browser would restrict the display to just the portion with BGP peers 192.0.2.1 and 192.0.2.2 that also has LocalPref 100:

```
(peer 192.0.2.1 or peer 192.0.2.2) and localPref 100
```

The previous example demonstrates the use of parentheses. Without them, the display would include the portion of the RIB with BGP peer 192.0.2.1 independent of LocalPref plus the portion of the RIB with both BGP peer 192.0.2.2 and LocalPref 100.

If the entries with LocalPref 100 are not interesting but other values are, then the expression could be modified as follows:

```
(peer 192.0.2.1 or peer 192.0.2.2) and not localPref 100
```

Regular Expressions

Some of the filter operators accept a regular expression as the parameter. A regular expression, also referred to as regex or regexp, is a pattern string that is used to describe or match a set of strings, according to certain syntax rules. Regular expressions are used in text editors and programming languages to search and manipulate bodies of text based on certain patterns. Several variants of the specification exist; the one used here is called POSIX Extended Regular Expressions.

Regular expressions can be arbitrarily complex, but even simple ones can prove useful as part of a filter expression. A regex pattern consists of literal characters and metacharacters. The literal characters in the pattern, such as letters and digits, match case-insensitively with the same characters in the target string such as a router name. Thirteen metacharacters have special meanings, but they can be matched literally by preceding them with a backslash:

```
. [ ^ $ ( ) { } \ ? + * |
```

The references below provide a full explanation of the pattern syntax, but a few examples will illustrate the power of regular expressions:

```
^NYC
```

The caret at the start of the pattern restricts it to match only at the beginning of the string, for example a set of routers in New York if their names begin with NYC.

```
-core$
```

The dollar sign at the end of the pattern restricts it to match only at the beginning of the string, for example a set of core routers if their names end with a hyphen and that word.

```
^NYC.*-core$
```

This pattern will match any name that begins with NYC and ends with -core, including NYC-core, but won't match Backup-NYC-central. The period matches any character and the asterisk indicates that the period should be repeated zero or more times, so the pattern `.*` will match zero or more arbitrary characters. Replacing the period with `[0-9]` would restrict the match to names with zero or more digits in the middle, while `[^-_.z]` would match any number of characters in the middle that were *not* hyphen, underscore, a literal period, or the letter z. Replacing the period with `(abc)` would restrict the match to names with zero or more repeats of the sequence abc, such as NYCabcabc-core. Note that the meaning of asterisk is different in regular expressions than it is in shell or file manager filename matching, so `*-core` is not a correct regular expression.

The plus sign is similar to asterisk but indicates a repeat of one or more times, while the question mark indicates a repeat of zero or one times, `{n}` indicates exactly n times, and `{m,n}` indicates at least m but no more than n times.

Two of the filter operators, **asPath regexp** and **mplsLabels regex**, take a Cisco extension to the regular expression syntax. That extension is the addition of underscore as a metacharacter. The underscore will match the beginning or end of the target string or a space in the middle of the target string.

For further information about regular expressions, see any of the following web pages:

http://www.zytrax.com/tech/web/regex.htm

http://analyser.oli.tudelft.nl/regex/

http://www.regular-expressions.info/posix.html

http://en.wikipedia.org/wiki/Regular_expression

Expression Definitions

This section defines the filter operators and also describes the function of each. The three conjunctive operators are listed first, followed by the others in alphabetical order.

not

Used to negate the next operator or parenthesized subexpression in the expression. For example,

```
not router 10.2.2.2
not (router 10.2.2.2 or router 10.3.3.3)
```

and

Requires that both the preceding and following operators in the expression be matched. For example,

```
router 10.1.1.1 and prefix 192.168.5.0/24
```

or

Matches if either the preceding or following operator in the expression matches. For example,

```
peer 10.1.1.1 or peer 10.2.2.2 or peer 10.3.3.3
```

asEdge <from-asn> <to-asn>

Matches an AS edge, meaning, a hop from one AS number to another, anywhere in the AS path. For example,

```
asEdge 1234 5678
```

asPath <asn> [atHead] [atTail]

Matches an AS number anywhere in the AS path, or optionally selects the AS at the head and/or tail. This example matches a singleton path containing only 1234:

```
aspath 1234 atHead atTail
```

asPathLen <n>
asPathLength <n>

Matches an AS path of length n. For example,

```
asPathLen 5
```

asPath regexp <regular expression>

Matches any AS path matching the Cisco-extended regular expression (see Regular Expressions on page 226). The first example will match all AS paths ending with the number 655, so AS path 124 444 1655 matches, but AS path 123 655 111 does not. The second example matches any path that includes AS 655 anywhere, so it matches 123 655 111 but not 124 444 1655:

```
AsPath regexp 655$
AsPath regexp _655_
```

availableBandwidthAfter <n> <relop>

Matches the available bandwidth after the planned changes. For example,

```
availableBandwidthAfter 100 lt
```

availableBandwidthBefore <n>

Matches the available bandwidth before the planned changes. For example,

```
availableBandwidthBefore 1000 lt
```

availableBandwidthChange <n> <relop>

Matches the difference of the available bandwidth before and after the planned changes. For example,

```
availableBandwidthChange 100 eq
```

bgpState <state>

Matches BGP routes in the specified state with respect to the baseline. The states are as follows:

`bgpState Dead`	(not in baseline and dead)
`bgpState Down`	(not in baseline and down)
`bgpState Up`	(not in baseline but up)
`bgpState Down/B`	(in baseline but down)
`bgpState Up/B`	(in baseline and up)

capacity <n> <relop>

In RAMS Traffic, matches the traffic link with capacity equal, less than or greater than the given capacity. For example,

```
capacity 100 lt
```

community <x:y>
community <x>

Matches a complete community attribute; it cannot match just the AS or just the value.

```
community 208:888
```

or

```
community 13632376
```

In the first form of notation, x is the first two octets (the AS number) and y is the second two octets (a value) of the community attribute. In the second form of notation, x is a four-octet quantity representing the complete community attribute.

destination / MPFilterFlowDst

In RAMS Traffic, matches traffic flow with the specified destination prefix. For example,

```
destination 182.168.0.1/24
```

Matches traffic flow destinations with the prefix of 192.168.0.1/24

destinationTrafficAfter <n> <relop>

In RAMS Traffic, matches the destination traffic after the planned changes. For example,

```
destinationTrafficAfter 1000 lt
```

destinationTrafficBefore <n> <relop>

In RAMS Traffic, matches the destination traffic before the planned changes. For example,

```
destinationTrafficBefore 100 lt
```

egressCapacity <value> <gt/eq/lt>

Matches with the specified egress capacity value in bps according to greater than, equal to, or less than comparisons. For example,

```
egressCapacity 100 gt eq
```

Matches egress capacity greater than 100 bps

egressTraffic <value> <gt/eq/lt>

In RAMS Traffic, matches with the specified egress traffic value in bps according to greater than, equal to, or less than comparisons. For example,

```
egressTraffic 100 gt eq
```

Matches egress capacity greater than 100 bps

egressUtilization <value> <gt/eq/lt>

Matches with the specified egress utilization value in bps according to greater than, equal to, or less than comparisons. For example,

```
egressUtilization 100 gt eq
```

Matches egress capacity greater than 100 bps.

eventCause <cause>

Matches events with the specified cause of a neighbor or prefix going down. Causes include:

```
eventCause premature
eventCause expired
```

eventType <operation>

Matches an event operation, meaning, a value in the *Operation* column of an events list, one of the following (where "*" is a wildcard to match any value):

```
eventType drop router
eventType add router
eventType change router
eventType drop neighbor
eventType add neighbor
eventType change neighbor
eventType drop prefix
eventType add prefix
eventType change prefix
eventType add rexpeering
eventType drop rexpeering
eventType change rexpeering
eventType drop *
eventType add *
eventType change *
eventType * router
eventType * neighbor
eventType * prefix
eventType * rexpeering
```

The following event types apply to EIGRP only:

```
eventType EIGRP Update
eventType Unresolved EIGRP Change
eventType EIGRP stuck-in-active
eventType start of exploration
eventType end of exploration
eventType drop static
eventType add static
eventType change static
eventType add route filter
eventType drop route filter
eventType change route filter
eventType add route acl
eventType drop route acl
```

```
eventType change route acl
eventType * static
eventType * route filter
eventType * route acl
```

The following event types apply to BGP only:

`eventType open` (open connection)
`eventType close` (close connection)
`eventType announce` (route announcement restoring a withdrawn route)
`eventType reannounce` (route announcement with same attributes as before)
`eventType new announce` (route announcement not previously existing and withdrawn)
`eventType withdraw` (route withdrawal)

exitRouter <ip address>

Matches exit router with a given IP address. For example,

```
exitRouter 192.168.0.1
```

Matches the exit router with the IP address of 192.168.0.1.

exportingInterface <interface address>

Matches flow export router with given router interface IP address. For example,

```
exportingRouter 192.168.0.1
```

Matches the flow export router interface with the IP address of 192.168.0.1

exportingRouter <router name/ip>

Matches flow export router with the specified router name or IP address. When a router name is given, it matches all routers whose names begin with that string, or the string can be a regular expression (see Regular Expressions on page 226) to match router names with any desired pattern. For example,

```
exportingRouter 192.168.0.1
exportingRouter .*core.*
```

Matches the flow export router with the IP address of 192.168.01

extCommunity RT:<route_target>
extCommunity SoO:<source_of_origin>

Matches a complete extended community attribute, including the type and all bits of the value. For either route_target or source_of_origin, the value can be expressed as a 16-bit global administrator value (AS number) followed by a 32-bit assigned value, or as a 32-bit value global administrator value, in the form of an IPv4 address or decimal number, followed by a 16-bit assigned value:

```
extCommunity RT:208:888
extCommunity RT:192.0.2.55:7
extCommunity SoO:13632376:123
```

externalOriginator <router>

Matches a router that is the originator of an external prefix in an EIGRP update event, using any of the forms of router identification described below for `router`. For example,

```
externaloriginator 192.168.0.36
```

igpPrefixType <type>

Matches the specified IGP prefix type. The available prefix types include the following:

- EIGRP: Internal, ASExt_Type2, ASExt_Type1, Loopback, Dialup, AutoSum, ManualSum, StaticInt, StaticExt, NotFoundInt, NotFoundExt
- IS-IS: Internal, ASExt_Type2, ASExt_Type1, TE, TEL2L1, InternalL2L1, AreaExtL2L1, ASExt_Type1L2L1, or ASEXT_Type2L2L1
- OSI: ESNeighbor, PrfxNeighbor, PrfxNeighborComparable
- OSPF: Internal, AreaExt, ASExt_Type2, or ASExt_Type1

For example, in an OSPF network:

```
igpPrefixType AreaExt
```

igpSeqNum <number> <gt/eq/lt>

Matches LSP packet sequence numbers according to greater than, equal to, or less than comparisons. For example:

```
igpSeqNum 58723 eq
```

igpState <state>

Matches IGP prefixes with the specified area. States include:

```
igpState up
igpState down
```

interface <IP Address>

Matches the interface with the given IP address. For example,

```
interface 192.168.0.1
```

Matches the interface 192.168.0.1

interfaceIdex <Interface Index>

Matches the interfaces with the specified interface index. For example,

```
interfaceIdx 1
```

Matches interfaces using index 1

inTraffic <value> <gt/eq/lt>

In RAMS Traffic, this matches communities with the specified traffic (bps) flowing out of a community according to greater than, equal to, or less than comparisons. For example:

```
inTraffic 100 gt eq
```

Matches total traffic received by a community greater than or equal to 100 bps.

linkBandwidth <value> <gt/eq/lt>

Matches EIGRP links with the specified bandwidth value according to greater than, equal to, or less than comparisons. The following example shows the matches for EIGRP links with bandwidth ≥ 1:

```
linkBandwidth 1 gt eq
```

linkDelay <value> <gt/eq/lt>

Matches EIGRP links with the specified delay value according to greater than, equal to, or less than comparisons. The following example matches EIGRP links with delay ≥ 1:

```
linkDelay 1 gt eq
```

linkState <state>

Matches links with the specified state. States include the following:

```
linkState up
linkState down
```

localPref <value>

Matches the BGP LocalPref value. For example,

```
localPref 888
```

med <value>

Matches the MED attribute only, and not the neighboring AS:

```
med 987
```

To match both the MED attribute and the neighboring AS, combine both operators:

```
med 987 and neighAS 208
```

metricBandwidth <value> <gt/eq/lt>

Matches EIGRP links with the specified EIGRP inverse bandwidth value according to greater than, equal, or less than comparisons. The metric value is $(10^7 / \mathrm{bw}) * 256$, where bw is in units of kilobits per second. For example,

```
metricBandwidth 25600 gt eq
```

matches EIGRP links with inverse bandwidth ≥ 25600 which is bandwidth ≤ 100 Mb/s

metricDelay <value> <gt/eq/lt>

Matches EIGRP links with the specified EIGRP delay metric according to greater than, equal to, or less than comparisons. The delay metric is in units of 10 μs, multiplied by 256. For example,

```
metricDelay 2560 gt eq
```

matches EIGRP links with delay ≥ 100 μs

metric <value> <gt/eq/lt>

Matches links with the specified metric value according to greater than, equal to, or less than comparisons. The following example matches links with metric ≥ 1:

```
metric 1 gt eq
```

mplsLabels <n> [atHead] [atTail]

Matches MPLS label anywhere in the label stack, or optionally selects the label at the head or tail. The following example matches MPLS labels with the first label 123:

```
mplsLabels 123 atHead
```

mplsLabels regex <regular expression>

Matches labels matching the Cisco-extended regular expression (see Regular Expressions on page 226). The following example shows MPLS labels starting with 111:

```
mplsLabels regexp ^111
```

neighbor <router>

Matches a neighbor router in an events list using any of the forms of router identification described above for `router`.

```
neighbor labnet-gw
```

neighAS <asn>

Matches the neighbor (nexthop) AS, meaning, the first AS in an AS path. For example,

```
neighAS 288
```

nexthop <addr>

Matches a BGP Nexthop. For example,

```
nexthop 192.0.2.67
```

noCommunity

Matches a BGP route or event that has no community attribute.

noExtCommunity

Matches a BGP route or event that has no extended community attribute.

noLocalPref

Matches a BGP route or event that has no LocalPref attribute.

noMed

Matches a BGP route or event that has no MED attribute, which is different than a MED value of zero.

noOrig
noOrigin
noOriginator

Matches a BGP route or event that has no Originator ID.

orig <addr>
origin <addr>
originator <addr>

Match a BGP originator ID. For example,

```
originator 192.0.2.4
```

originAS <asn>

Matches the origin AS, meaning, the last AS in an AS path. For example,

```
originAS 289
```

outTraffic <value> <gt/eq/lt>

In RAMS Traffic, matches communities with the specified traffic (bps) flowing out of a community according to greater than, equal to, or less than comparisons. For example,

```
outTraffic 100 gt eq
```

Matches total traffic flowing out of a community greater than or equal to 100 bps.

peer <router>

Matches a specific BGP peer address using any of the applicable forms of router identification described below for `router`. For example,

```
peer 192.0.2.3
```

peeringDestination <value> <gt/eq/lt>

In RAMS Traffic, matches ASs with the specified traffic whose final destination is the AS according to a greater than, equal to, or less than comparison. For example,

```
peeringDestination 100 gt eq
```

Matches destination traffic greater than or equal to 100 bps.

peeringNextHop <value> <gt/eq/lt>

In RAMS Traffic, matches ASs with the specified traffic flow transiting across the AS according to greater than, equal to, or less than comparisons. For example,

```
peeringNextHop 100 gt eq
```

Matches destination traffic greater than or equal to 100 bps.

peeringTotal <value> <gt/eq/lt>

In RAMS Traffic, matches ASs with the specified traffic flow from the AS according to greater than, equal to, or less than comparisons. For example,

```
peeringTotal 100 gt eq
```

Matches egress capacity greater than or equal to 100 bps.

percent <number>

In RAMS Traffic, matches traffic groups with the percentage of the total traffic flowing for the traffic group which is greater than or equal to the specified percentage. For example,

```
percent 0.10
```

Matches the traffic group whose traffic flow is greater than or equal to 10% of the total traffic.

prefix <addr/masklen> [moreSpecifics][lessSpecifics][ge <masklen>][le <masklen>]

Matches a prefix; optionally followed by either or both of the `moreSpecifics` and `lessSpecifics` operators to also display prefixes more or less specific than the given prefix. Alternatively, the prefix can be specified by an address followed by either or both of the operators `ge` or `le` with a mask length to include prefixes with mask lengths greater-than-or-equal or less-than-or-equal to the given integer parameter. For example,

```
prefix 10.2.0.0/16
prefix 10.2.0.0/16 moreSpecifics
prefix 10.2.0.0 ge 16
prefix 10.2.0.0 ge 16 le 24
```

proto <proto>
protocol <proto>

Selects a particular protocol. The available protocols are IS-IS, OSPF, EIGRP, BGP and Static (the last currently only available in an EIGRP topology):

```
proto isis
```

routeTarget RT:<route_target>

routeTarget So0<source_of_origin>

Matches VPN routers with the specified VPN route target (see page 234 for input format). For example,

```
routeTarget RT:59300:460210
RouteTarget So0: 12632376:123
```

router <router>

Matches a specific router using any of the forms of identification that are shown in a table: name, address, prefix (for a LAN pseudonode), or SystemID (for IS-IS). An address or name may be followed by "DR" to select the role of a router as the designated router for a subnet. When a router name is given, it matches all routers whose names begin with that string, or the string can be a regular expression (see Regular Expressions on page 226) to match router names with any desired pattern. For example,

```
router labnet-gw
router 192.168.0.36
router 192.168.0.36/24
router 1921.6800.0036:00
router 192.168.0.36 dr
router labnet-gw:01 DR
router .*core.*
```

routerState <state>

Matches routers with the specified state. States include the following:

```
routerState up
routerState down
```

routerType <type>

Matches routers with the specified router type. Router types include the following:

```
routerType Internal
routerType LANPseudonode
routerType AreaBR
routerType AreaBR_ASBR
routerType ASBorderRouter
routerType VirtualRouter
routerType RouteRecorder
routerType IBGPpeer
routerType EBGPpeer
routerType RouteReflector
routerType Originator
```

```
routerType EBGPNextHop
routerType NeighborAS
routerType IBGPPeerOriginator
routerType EBGPPeerOriginator
routerType Implicit
routerType IBGP
routerType Static
routerType StaticNextHop
routerType L1Internal
routerType L2Internal
routerType L1L2Router
routerType L1L2RouterASBR
routerType ASBRProxyOutsideArea
```

secondHopAS <asn>

Matches the second hop AS, meaning, the one after the neighbor AS in an AS path. For example,

```
secondHopAS 288
```

source <ip prefix>

In RAMS Traffic, matches traffic flows with the specified source prefix. For example,

```
source 182.168.0.1/24
```

Matches traffic flow with the source prefix of 182.168.0.1/24

staticNexthopType <type>

Matches the specified nexthop type for a static route. The available types are: Network, Interface, Gateway, Default. For example,

```
staticNexthopType Network
```

totaltrafficAfter <n> <relop>

In RAMS Traffic, matches the total traffic after the planned changes. For example,

```
totaltrafficAfter 100 eq
```

totaltrafficBefore <n> <relop>

In RAMS Traffic, matches the total traffic before the planned changes. For example,

```
totaltrafficBefore 100 eq
```

totalTrafficChange <n> <relop>

In RAMS Traffic, matches the difference of total traffic before and after the planned changes. For example,

```
totalTrafficChange 100 eq
```

trafficAfter <n> <relop>

In RAMS Traffic, matches the specified traffic after the planned change. For example,

```
trafficAfter 1000 eq
```

trafficBefore <n> <relop>

In RAMS Traffic, matches the specified traffic before the planned changes. For example,

```
trafficBefore 200 eq
```

trafficChange <n> <relop>

In RAMS Traffic, matches the specified traffic changes. For example,

```
trafficChange 1000 eq
```

traffic <value> <gt/eq/lt>

In RAMS Traffic, matches total egress traffic greater than or equal to 100 bps. For example,

```
traffic 100 gt eq
```

Matches traffic greater than or equal to 100 bps.

transitBandwidthAfter <n> <relop>

Matches the transit bandwidth after the planned changes. For example,

```
transitBandwidthAfter 1000 lt
```

transitBandwidthBefore <n> <relop

Matches the transit bandwidth before the planned changes. For example,

```
transitBandwidthBefore 1000 gt
```

utilizationAfter <n> <relop>

In RAMS Traffic, matches the specified utilization of the traffic link after the planned change. For example,

```
utlizationAfter 100 eq
```

utilizationBefore <n> <relop>

In RAMS Traffic, matches the specified utilization of the traffic link before the planned change. For example,

```
utilizationBefore 100 eq
```

utilizationChange <n> <relop>

In RAMS Traffic, matches the specified utilization of the traffic link with the specified utilization changes. For example,

```
utilizationChange 200 gt
```

vpnCustomer <name>

Matches VPN routes with the specified VPN customer. For example,

```
vpnCustomer Customer1
```

vpnPrefix <target:addr/masklen> [moreSpecifics][lessSpecifics][ge <masklen>][le <masklen>]

Matches a VPN prefix, which is composed of a route distinguisher (RD) plus a prefix; see Chapter 8, "VPN Routing" for a description of RD formats. The prefix is optionally followed by either or both of the `moreSpecifics` and `lessSpecifics` operators to also display prefixes more or less specific than the given prefix. Alternatively, the prefix can be specified by an address followed by either or both of the operators `ge` or `le` with a mask length to include prefixes with mask lengths greater-than-or-equal or less-than-or-equal to the given integer parameter. For example,

```
vpnPrefix 192.168.0.36:65522:101:10.2.0.0/16
vpnPrefix 192.168.0.36:65522:101:10.2.0.0/16 moreSpecifics
vpnPrefix 192.168.0.36:65522:101:10.2.0.0 ge 16
vpnPrefix 192.168.0.36:65522:101:10.2.0.0 ge 16 le 24
```

5 Network Planning

This chapter describes how to use Route Analytics Management Software to plan for network growth and change.

Chapter contents:

About the Network Planning Tools

The Route Analytics Management Software network planning tools help you identify and eliminate hot spots and avoid potential service failures. You can perform failure analysis and move prefixes among border routers. By combining routing and traffic data, you can view traffic trends and their impact on the available capacity and reliability of the network.

In contrast with other tools that are based only on synthetic models of network activity or limited link utilization measurements, the planning tools in RAMS and RAMS with traffic analysis add-on are based on actual measurements.

To use the Route Analytics Management Software planning tools, you enter Planning mode on the routing topology map. In Planning mode, you can view edits, import and export data, or undo changes. In RAMS Traffic, you can also compare network activity before and after applied changes take effect to analyze the differences in traffic measurements that result from simulated network modifications. Using Planning Reports, you can analyze how traffic changes across the entire network and on specified nodes, interfaces, exit routers, next-hop ASs, or destination ASs.

To enter Planning mode, click the mode icon in the lower left corner of the window and choose the Planning mode icon .

When you enter Planning mode, the options listed in the Planning menu are activated, except for Capacity planning, which is activated when you enter Analysis mode. To enter Analysis mode, click the mode icon in the lower left corner of the client application window and choose the Analysis mode icon . When you enter Analysis mode, the Capacity Planning item in the Planning menu is activated.

If you are recording a BGP VPN topology when you switch to Planning mode, VRF configurations are automatically discovered based on a heuristic algorithm.

Planning Menu

In Planning mode, use the Planning menu at the top of the routing topology map window (Figure 1) to perform the network planning tasks described in this chapter. Working in Planning mode, you can simulate changes to the network by editing the topology map. For instance, simulating the addition of a node allows you to see realistic effects of the new router on network activity.

When you enter Planning mode, the planning toolbar on the right side of the window is displayed.

As noted in this section, some of the options found in the Planning menu are for RAMS Traffic only.

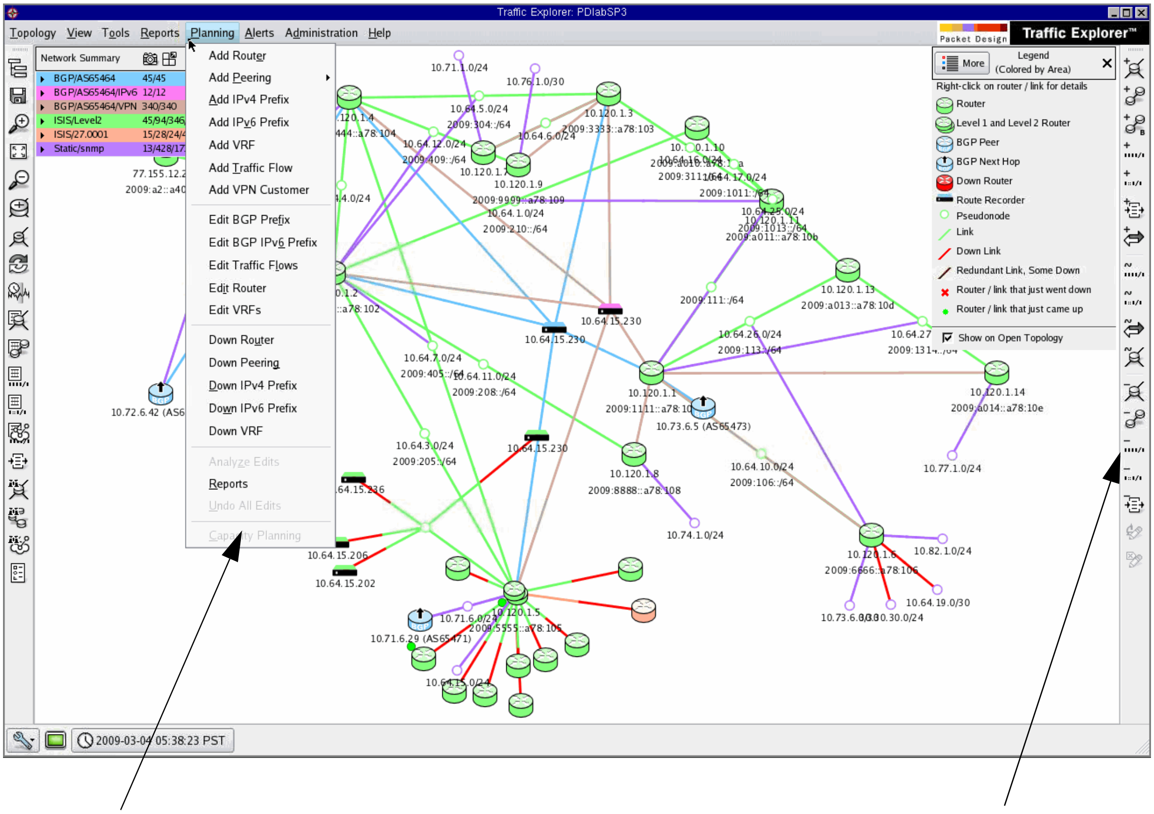

Planning menu

Planning toolbar

Figure 1 Planning Menu and Toolbar

The Planning menu includes the following items:

- **Add Router**—Places a new node on the topology map.
- **Add Peering**—Creates a peering relationship between two nodes on the topology map.
  - **Add BGP Peering**—Creates an eBGP peering relationship between two nodes on the topology map.
  - **Add IGP Peering**—Creates an IGP peering relationship between two nodes on the topology map.

- **Add IPv4 Prefix**—Applies IPv4 prefixes to a router on the topology map. For BGP routers, you can add prefixes manually or select a filtering method for the prefix.
- **Add IPv6 Prefix**—Applies IPv6 prefixes to a router on the topology map. For BGP routers, you can add prefixes manually or select a filtering method for the prefix.
- **Add OSI Prefix**—Applies OSI prefixes to an OSI IS-IS router on the routing topology map.

  The Add OSI Prefix option is available only if OSI IS-IS is included on the routing topology map.

- **Add VRF**—Adds Virtual Routing & Forwarding, which allows multiple instances of a routing table to coexist within the same router at the same time. The routing instances are independent, thus allowing the same or overlapping IP addresses to be used without conflict.

  The Add VRF option is available only when VPN is licensed on the appliance and is included in the opened topology.

- **Add Traffic Flow**—Creates one or more flows from one node to another on the routing topology map. (RAMS Traffic only)
- **Add VPN Customer**—Opens the Add Customer Wizard, which guides you through the process of adding a full mesh, or hub and spoke VPN customer to the network, including VRF routes.

  The Add VPN Customer option is available only when VPN is licensed on the appliance and is included in the opened topology.

- **Edit BGP Prefix**—Removes or changes the attributes of an IPv4 prefix on a particular node on the topology map. You can change multiple prefixes at once by selecting more than one prefix from the table.
- **Edit BGP IPv6 Prefix**—Removes or changes the attributes of an IPv6 prefix on a particular node on the topology map. You can change multiple prefixes at once by selecting more than one prefix from the table.
- **Edit Traffic Flows**—Launches the Edit Flows window, where you can add, move, or delete IPv4 VPN flows. (RAMS Traffic only)

- **Edit Router**—Edits the overload bit for an IS-IS node.
- **Edit VRFs**—Enables editing of the Virtual Routing and Forwarding router table instances.

  The Edit VRFs option is available only when VPN is licensed on the appliance and is included in the opened topology.

- **Down Router**—Changes the state of a node from Up to Down, simulating what would happen if the selected router should fail.
- **Down Peering**—Changes the state of a peer relationship from Up to Down, simulating what would happen if the selected peering should fail. Bring down all the peerings in the table or only selected relationships.
- **Down IPv4 Prefix**—Changes the state of one or more IPv4 prefixes from Up to Down on a particular router, simulating what would happen if the selected prefixes are withdrawn.
- **Down IPv6 Prefix**—Changes the state of one or more IPv6 prefixes from Up to Down on a particular router, simulating what would happen if the selected prefixes are withdrawn.
- **Down VRF**—Changes the state of one or more VPN nodes from Up to Down on a particular router, simulating what would happen if the selected item should fail.

  The Down VRF option is available only when VPN is licensed on the appliance and is included in the opened topology.

- **Reports—**Launches the Planning Reports window. For more information, see Chapter 9, "Traffic Flows and Reports" (RAMS Traffic only)
- **Analyze Edits**—Updates all traffic and routing edits simultaneously.
- **Capacity Planning**—(Analysis mode only) Launches the Capacity Reports window, which enables you to view an estimate of potential future traffic demands based on past data collection. This allows you to plan network expansion to meet future demands. See Working with Capacity Planning Tools on page 304 for more information. (RAMS Traffic only)
- **Show Edits**—Displays edits that have been made to the topology (present only if the opened topology does not include traffic).

The Planning Toolbar

By default, the Planning toolbar is docked on the right side of the Topology window (Figure 1). You can access planning functions from the toolbar, or use the Planning menu.

Move the toolbar to the left side or the top or bottom of the window by dragging the dimpled strip at the top of the toolbar. Add and edit elements on the topology map using the buttons shown in Table 1 .

Table 1 Planning Toolbar

	**Add Router**	See Add Router on page 253.
	**Add IGP Peering**	See Add IGP Peering on page 256.
	**Add BGP Peering**	See Add BGP Peering on page 254.
	**Add IPv4 Prefix**	See Add a Prefix on page 257.
	**Add IPv6 Prefix**	See Add a Prefix on page 257.
	**Add VRF**	See Add VRF on page 262.
	**Add Traffic Flow**	See Add a Traffic Flow on page 275. (RAMS Traffic only)
	**Edit BGP Prefix**	See Add VRF on page 262.
	**Edit BGP IPv6 Prefix**	See Add VRF on page 262.
	**Edit Traffic Flow**	See Edit Traffic Flows on page 272. (RAMS Traffic only)
	**Edit Router**	See Edit Node Properties on page 278.
	**Down Router**	See Bring Down a Router on page 280.
	**Down Peering**	See Bring Down Peerings on page 280.
	**Down IPv4 Prefix**	See Bring Down a Prefix on page 282.

Table 1 Planning Toolbar (cont'd)

	**Down IPv6 Prefix**	See Bring Down a Prefix on page 282.
	**Down VRF**	See Bring Down VRF on page 283.
	**Analyze Edits**	Select to update all traffic and routing edits simultaneously.

Add Router

Use the Add Router function in Planning mode to simulate the effect of adding one or more nodes to the network. When you add a node, you must specify its protocol and properties. You can then create peering relationships with other nodes on the network, as described in Add Peering on page 254. In RAMS Traffic with BGP routers, the node is automatically peered.

You can also add a protocol instance to a node, and this procedure is described in Adding a Protocol Instance to an Existing Node on page 254.

To add a router, perform the following steps:

1 In Planning mode, choose **Add Router** from the Planning menu.

2 Click the desired protocol tab (Figure 2).

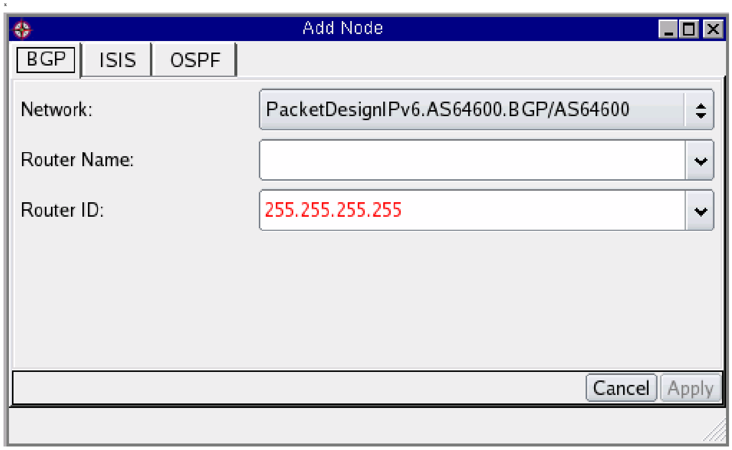

Figure 2 Add Router (BGP Options)

3 The current network is displayed in Network field. If necessary, click the down arrow to choose a different network from the drop-down list. The node will be added to the topology of the specified network.

4 Enter a name for the node in the Router Name text box.

5 Specify ID and address information:

- For BGP OSF, specify the router ID.
- For IS-IS, select the type of address from the Address Family drop-down list. If you choose one of the IPv4 and/or OSI options, enter the IPv4 address. If you choose the IPv4 + IPv6 option, enter the IPv4 and IPv6 addresses.

6 Click **Apply** to save the node.

For OSPFv3, the Add Router operation sets the R and V6 bits. To change those settings, see Edit Node Properties on page 278.

Adding a Protocol Instance to an Existing Node

To add a protocol instance to an existing node, perform the following steps:

1 In Planning mode, choose **Add Router** from the Planning menu.

2 Click a node on the topology map.

3 Choose a protocol to add to the node.

4 Click **Apply**.

Add Peering

You can create peerings between two nodes to simulate the effect that the specified relationship would have on the network.

Add BGP Peering

Create a BGP peering between two nodes to simulate the effect this relationship would have on the network.

To Add a BGP peering, perform the following steps:

1 In Planning mode, choose **Add Peering > Add BGP Peering** from the Planning menu.

2 Click a node on the topology map to specify the source node for the new peering (Figure 3).

If a IPv4/IPv6 node is selected, the Add BGP peering window includes a BGP IPv6 tab in addition to the BGP tab.

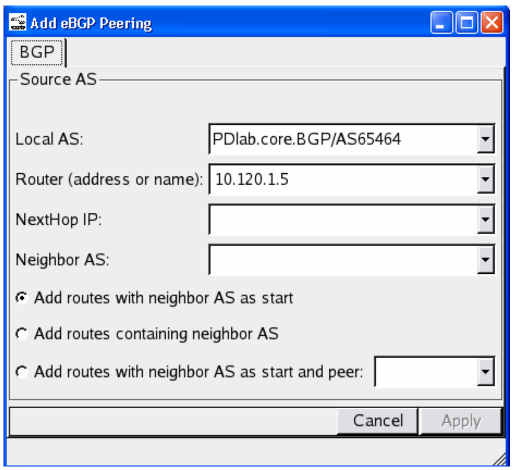

Figure 3 Add eBGP Peering

3 The current AS is shown in the Local AS text box. If necessary, click the down arrow to choose a different AS from the drop-down list. The peering will be added to the topology of the AS specified in the text box.

4 In the Router text box, specify the IP address or name of the source node in the peering.

5 In the NextHop IP text box, specify the next hop IP address of the source node. If IPv6 is supported, you can specify an IPv6 address.

6 The Neighbor AS text box is populated with the neighbor AS number you are creating a peering for. You can enter other neighboring AS numbers or choose an AS number from the drop-down menu.

7 Select radio button to determine the criterion to use in creating routes for this peering.

8 Click **Apply**.

Add IGP Peering

Create an IGP peering between two nodes to simulate the effect this relationship would have on the network.

To create an IGP peering, perform the following steps:

1 In Planning mode, choose **Add Peering > Add IGP Peering** from the Planning menu.

2 Click a node on the topology map to specify the source node for the new peering.

3 Click a node on the topology map to specify the destination node for the new peering (Figure 4).

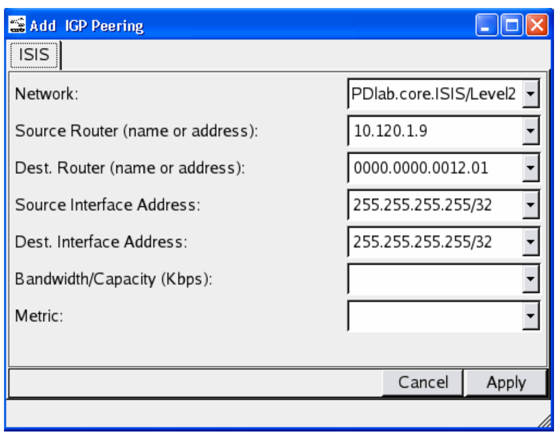

Figure 4 Add IGP Peering

4 The network you are currently editing appears in the Network text box. If necessary, click the down arrow to choose a different network from the drop-down list. The peering will be added to the network specified in the text box.

5 The Source Router text box is populated with the IP address or name of the source node in the peering. Edit this text box if necessary.

6 The Dest Router text box is populated with the IP address or name of the destination node in the peering. Edit this text box if necessary.

7 The Source Interface Address text box is populated with the default IP address. Replace the default value with the address of the source interface and optional mask (for example, 192.168.1.101 or 192.168.1.101/24.

8 The Dest Interface Address text box is populated with the default IP address. Replace the default value with the address of the destination interface and optional mask (for example, 192.168.1.101 or 192.168.1.101/24.

9 For OSPFv3, specify the source and destination interface IDs.

10 Specify the available bandwidth allocated for this peering in the Bandwidth/Capacity text box.

11 Specify the metric value for the peering. Metric values help traffic determine the best path to take through the network and typically take bandwidth, communication cost, delay, hop count, load, and reliability into consideration.

12 Click **Apply**.

Add a Prefix

Simulate the effect of adding one or more IPv4 or IPv6 prefixes to the topology by using the Add Prefix functions.

Adding a Prefix for BGP or BGP/MPLS-VPN Routers

You can specify prefix attributes manually or by using a filter.

To manually add a prefix, perform the following steps:

1 In Planning mode, choose **Add IPv4 Prefix** or **Add IPv6 Prefix** from the Planning menu and click a node on the map to specify the router that will advertise the prefix.

2 The topology you are currently editing appears in the Network text box. If necessary, click the down arrow to choose a different network from the drop-down list.

3 In the Router text box, enter the IP address or name of the router that advertises the prefix.

4 In the **Metric Type** text box, enter the peer type of the router.

5 Enter a valid metric according to the metric type.

6 For OSPFv3, select whether to include the local address (LA) bit and the no unicast (NU) bit.

7 Click **Apply**. (This radio button is selected by default when you click **Add Prefix**).

8 In the RD (route distinguisher) field, select the VRF label for the prefix. (BGP/MPLS-VPN only)

9 In the MPLS Label field, select the label for the prefix. (BGP/MPLS-VPN only)

10 In the Prefix text box, enter the address of the new prefix.

11 Set the attributes of the prefix by specifying the origin.

12 Click **Apply**.

To add a prefix using a filter, perform the following steps:

1 In Planning mode, choose an **Add Prefix** item from the Planning menu and click the desired node to open the Add Prefix window.

2 Click **Using Filter** at the top of the window.

3 In the RD field, select the VRF label for the prefix. (BGP/MPLS-VPN only)

4 In the MPLS Label field, select the label for the prefix. (BGP/MPLS-VPN only)

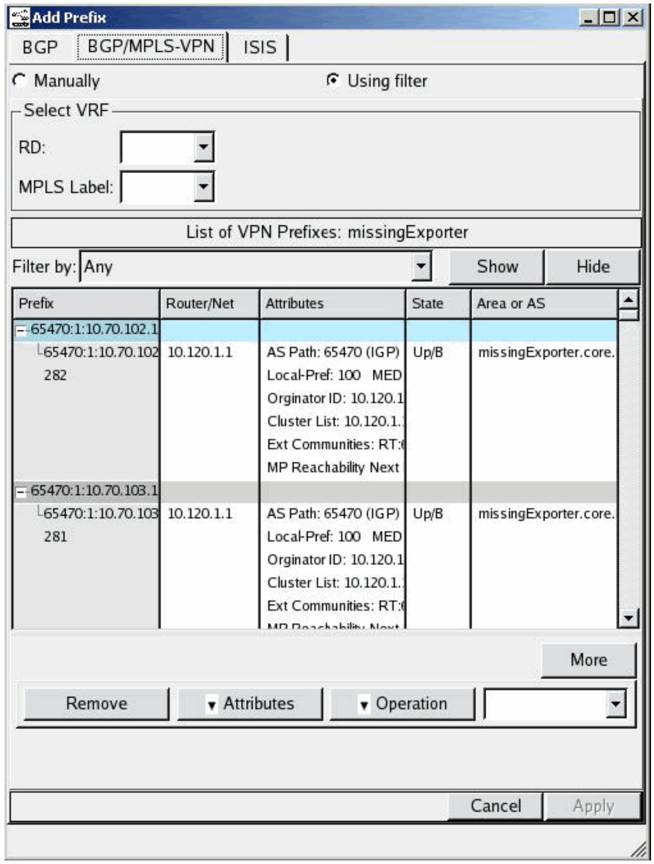

Figure 5 Using Filters Radio Button

5 Use the **Filter by** drop-down list to choose which prefixes to show or hide in the table.

Simple filters let you choose a single operator (for example, Router) from a list and specify one or more parameters (for example, router IP addresses or names) to be matched or excluded. See Using Filters on page 221 in Chapter 4, "The History Navigator" for examples of the parameter syntax using filter expressions described in Expression Syntax on page 225.

With a simple filter, you enter only the parameter; the operator is selected from a list.

The filter is translated internally into a filter expression that combines the filter operator with the parameters.

You can choose two or more different operators from a list and specify their corresponding parameters to be matched or excluded. The Composing Advanced Filter window opens when **Advanced** is selected.

Filter expressions let you manually enter a filter expression that is too complex to be set up with either simple or advanced filter menus.

See Using Filters on page 221 in Chapter 4, "The History Navigator" for more detailed information.

6 Click **Show** to list only items that match the parameters of the filter, or click **Hide** to list only items that do not match the parameters of the filter.

7 Click a prefix in the table to highlight it.

8 Choose a filtering method for the highlighted prefix from the Attributes drop-down list and enter a corresponding value in the text box to the right.

9 Use the Operation drop-down list to set, append, or prepend the value to the corresponding attribute of the prefix.

For example, to set the local preference value of the highlighted prefix to 99, choose **Local Pref** from the Attributes list, enter **99** in the text box and choose **Set** from the Operation list.

10 Click **More** to apply additional filtering methods to the prefix or **Remove** to remove the last filtering method applied to the prefix.

11 Click **Apply**.

12 Click **Close (X)** to close the window.

Add prefix for IS-IS Routers

If you chose the **ISIS** tab in the Add Prefix window, the Add Prefix window opens.

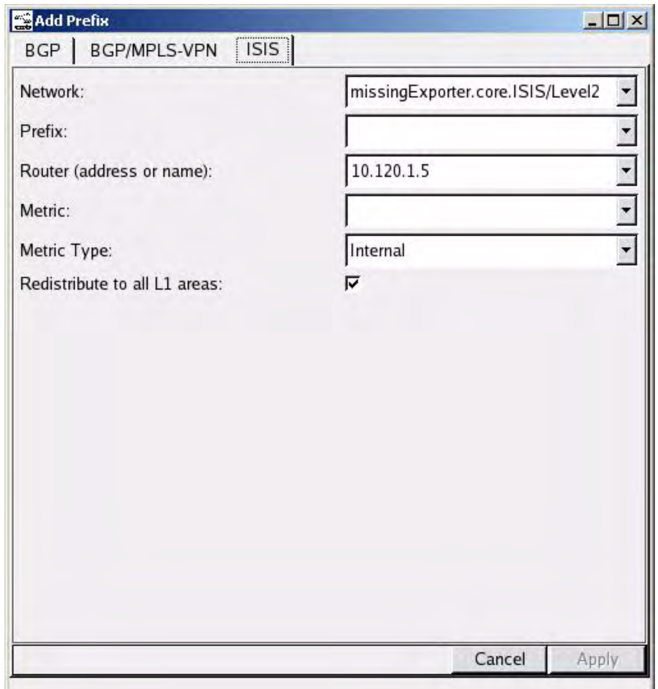

Figure 6 Add Prefix (IS-IS)

To add a prefix to an IS-IS router, perform the following steps:

1. In Planning mode, choose **Add IPv4 Prefix** or **Add IPv6 Prefix** from the Planning menu.
2. Click an IS-IS router on the topology map to specify the router that will advertise the prefix. You can move the prefix later, if necessary. If you choose **Add IPv6 Prefix**, the IPv6 address is shown by default.
3. The Add IPv4 Prefix or Add PIv6 Prefix window opens.

   Click the ISIS tab.

4 The network you are currently editing appears in the Network text box. If necessary, click the down arrow to choose a different network from the drop-down list.

5 Enter the address of the new prefix in the Prefix text box.

6 The Router text box is populated with the system ID or name of the node that advertises the prefix. Edit this text box if necessary.

7 Enter a valid metric according to the metric type.

8 Select the check box to redistribute to all L1 areas.

9 Click **Apply**.

Add VRF

Adding a VRF to a PE allows multiple instances of a routing table to coexist within the same router at the same time. The routing instances are independent, allowing the same or overlapping IP addresses to be used without conflicting with each other.

This option is enabled if your appliance is licensed for VPN, and if VPN protocol is present in the opened topology.

To add a VRF to a PE, perform the following steps:

1 In Planning mode, choose **Add VRF** from the Planning menu.

2 On the topology map, click a VPN node that is to be the PE for the VRF.

The network you are currently editing appears in the Network text box. If necessary, click the down arrow to choose a different network from the drop-down list.

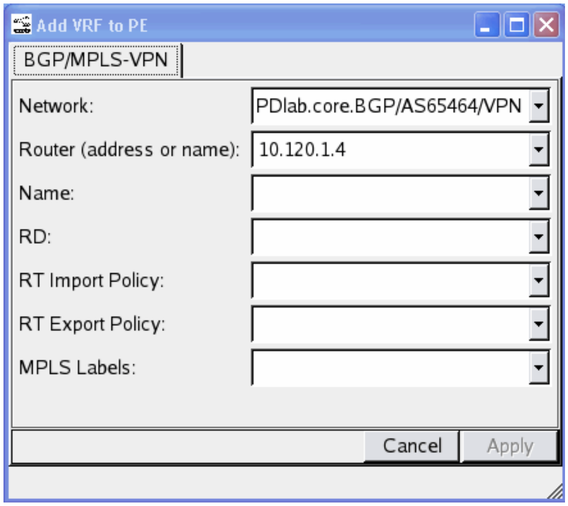

Figure 7 Add VRF to PE

3 Enter the router that identifies the VRF in the Router field.

4 Enter a name for the VRF in the Name field.

5 Enter a list of RTs in the RT Import Policy field.

6 Enter the RTs that attaches to router when in the RT Export Policy field.

7 Enter the MPLS label to attach to the exporting router in the MPLS Labels field.

8 Click **Apply** to save the information.

Add Traffic Flow

This section applies to RAMS Traffic and IPv4 only.

You can add of IPv4 traffic flows to the topology map using the Add Traffic Flow function. When adding a traffic flow, specify the source and destination prefixes and the bandwidth for the flow. To add more than one flow at a time, use the Multiple Flows tab. Edit existing flows using the Edit Flows window as described in Edit Traffic Flows on page 272.

To add an IPv4 traffic flow, perform the following steps:

1 In Planning mode, choose **Add Traffic Flow** from the Planning menu.

> You can also add traffic flow in the Edit Traffic Flows window using the **Add Flow** button described in Edit Traffic Flows on page 272.

2 Click a node on the topology map to open the Add Traffic Flow window (Figure 8).

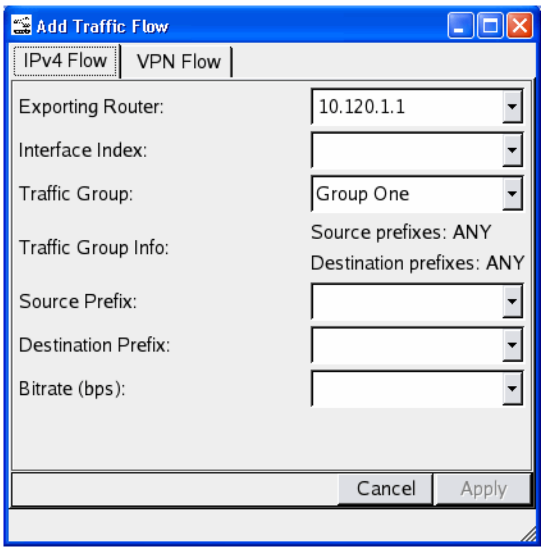

Figure 8 Add Traffic Flow-IPv4 Tab

3 Select the **IPv4 Flow** tab to add an IPv4 traffic flow to the node you previously selected.

4 In the Exporting Router text box, the address of the router chosen on the topology map is displayed. Choose another router by entering its address in the text box or by clicking the down arrow to select one from the drop-down list.

5 In the Interface Index text box, enter the value of the interface index to associate with the traffic flow.

6 Select the traffic group you want the IPv4 traffic flow to go to from the **Traffic Group** drop-down menu.

In the Traffic Group Info section, the source and destination information reflect the configuration of the traffic group that was specified in the previous step. You can view the configuration of these groups from the Traffic Groups web page or in the Traffic Reports window under the Traffic Group Definition column.

7 Enter the address of the prefix where the new traffic flow originates in the Source Prefix text box.

8 Enter the address of the destination prefix for the new traffic flow in the Destination Prefix text box.

9 Enter the bitrate of the new traffic flow in the Bitrate text box.

10 Click **Apply**.

Add VPN Traffic Flows

This section applies to RAMS Traffic only. The appliance must be licensed for VPN, and the open topology must contain VPN protocol for this option to be enabled.

The appliance supports the following methods to identify the source and destination of a VPN traffic flow:

- **PE-PE Flow—**Allows you to specify the ingress and egress PEs for the flow, without the flow belonging to a specific customer. An unspecified customer is created in the background for such flows, along with any necessary VRFs and PEs for that customer on its PEs, and a flow is created from the ingress PE to the egress PE with source and destination prefixes of 0/0. You can also specify the CoS and bitrate for the flow.
- **Customer Flow**—Allows you to specify the information necessary to determine the ingress VRF of a flow along with its destination prefix, at which point the existing VPN topology is used to determine its egress PE and VRF. Enter the necessary information to determine the ingress VRF of the flow. The destination prefix is then combined (the import RTs of the ingress VRF) to determine the egress PE and egress VRF of the flow. If there is no VPN prefix for the ingress VRF and destination prefix combination, you are notified and no flow is entered.

- **VRF Flow**—Allows you to specify the egress PE and egress VRF, which are automatically determined in the Customer flow. If there is no VPN prefix for the ingress VRF and destination prefix combination, you are notified and no flow is entered.

To add a VPN PE-PE traffic flow, perform the following steps:

1. In Planning mode, choose **Edit Traffic Flows** from the Planning menu.
2. Select the VPN Flows tab and click **Add Flow**. If there is only one type of flow, tabs are not displayed.

The PE-PE flow is the default selection.

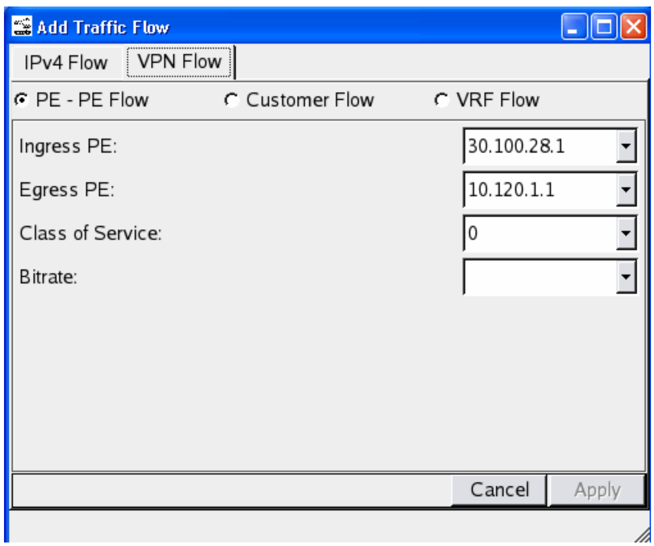

Figure 9 Add PE-PE Selection for VPN Traffic Flow

3. Determine the method of specifying the flow by choosing Customer Flow, VRF Flow, or PE-PE flow.
4. Select the IP addresses for the ingress PE and egress PE from the drop-down menus.
5. Select the class of service from the drop-down menu.
6. Select a the bitrate from the Bitrate drop-down menu.

7 Select **Apply**.

To add a VPN customer flow, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Select the VPN Flows tab and click **Add Flow**.

3 Select the **Customer Flow** radio button.

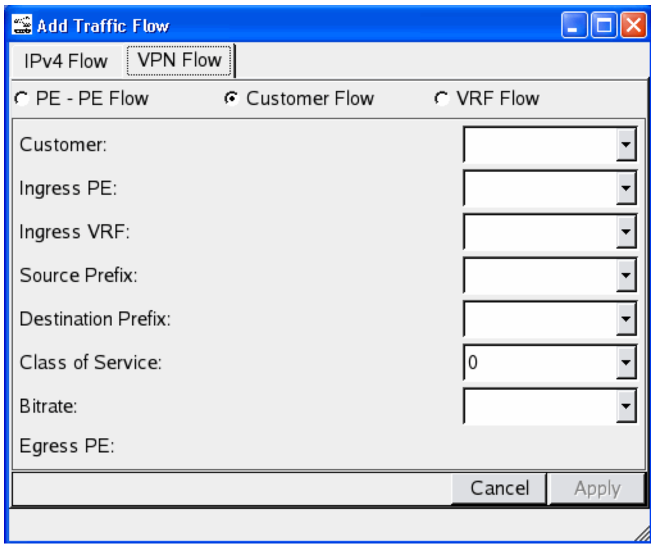

Figure 10 Add Customer Flow Selection for VPN Traffic Flow

4 Select the customer from the drop-down menu.

5 Select IP addresses for the ingress PE and egress PE from the drop-down menus.

6 Enter the source and destination IP prefixes for the traffic that you want to identify drop-down menus.

7 Select the class of service from the drop-down menu.

8 Select a the bitrate from the Bitrate drop-down menu.

9 Select **Apply**.

To Add a VRF Flow, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Select the VPN Flows tab and click **Add Flow**.

3 Select the **VRF Flow** radio button.

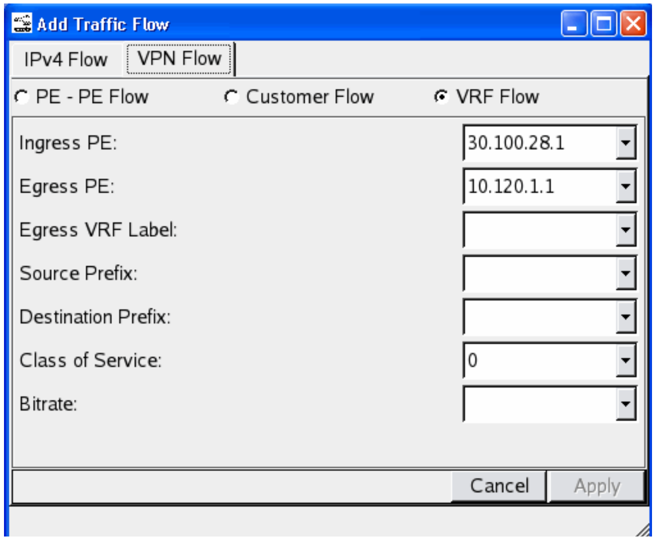

Figure 11 Add VRF Flow Selection for VPN Traffic Flow

4 Select IP addresses for the ingress PE and egress PE from the drop-down menus.

5 Enter or select a label to identify the egress VRF.

6 Enter the source and destination IP prefixes for the traffic that you want to identify drop-down menus.

7 Select the class of service from the drop-down menu.

8 Select a the bitrate from the Bitrate drop-down menu.

9 Select **Apply**.

Add VPN Customer

To add VPN customers, use the Add VPN Customer wizard.

This option is enabled if your appliance is licensed for VPN, and if VPN protocol is present in the opened topology.

To add a VPN customer to the topology, perform the following steps:

1 In Planning mode, choose **Add VPN Customer** from the Planning menu.

2 Enter the customer name in the Customer field.

3 Select the VPN topology (Full Mesh or Hub-and-Spoke).

Two columns are shown and identified below each column: Available PEs is on the left, and Selected PEs is shown on the right. Both have Select (RegEx) and a drop-down menu at the top of each column.

The syntax of extended regular expressions is explained in Regular Expressions on page 226. The syntax is not the same as shell or file manager pattern patching, so a pattern like *-core-gw is not correct.

4 Select the available PEs for this customer by double-clicking the PE or by clicking on the PE, and then clicking on the right-arrow (->).

If you selected Full Mesh in Step 3, the Select Customer RTs window opens. Continue to Step 5.

If you selected Hub-and-Spoke in Step 3, continue to Step 6 to select the Hub-and-Spoke PEs in the Define Hub-and Spoke Configuration window.

5 If you selected Full Mesh in Step 3, enter the customer's RTs in the Full Mesh RTs text box in the Select Customer RTs window, click **Next** and continue to Step 7.

6 If you selected Hub-and-Spoke in Step 3, the Define Hub-and Spoke Configuration window opens. Two columns display: the column on the left displays Hub PEs at the bottom of the column, and the column on the right displays Spoke PEs at the bottom of its column. Select the Hub and Spoke PE's, and click **Next**.

The Add Customer Routes window opens.

7 Enter the prefix range used by the routes in this VPN topology in the Prefix Range text field, and click **Next**.

The Select VPN Traffic Sources window opens.

8 Choose **All PEs** if you want all customer PEs selected as VPN traffic sources, or choose **Specify PEs** and select the sources. Click **Next**.

   The Select VPN Traffic Destinations window opens.

9 Choose **All PEs** if you want all customer PEs selected as VPN traffic sources, or choose **Specify PEs** and select the destinations. Click **Next**.

   The Specify Traffic Flow Distribution window opens.

10 Enter the class of service from the corresponding drop-down menu.

11 Enter the average bitrate per flow in the corresponding text box.

12 Select the distribution type from the following choices:

   - **Equal**—Each flow has the average bitrate entered in the previous step.
   - **Uniform**—Each flow has an equal probability of having any bitrate from 0 bps to twice the bitrate entered in the previous step.
   - **Pareto**—The average bitrate of the generated flows is set to the bitrate that was entered in the previous step. This setting is based on the fact that in a flow distribution, a small number of flows has a high bitrate, while most of the flows have a lower bitrate.

13 Click **Next.**

   The Review Traffic Flow Distribution window opens.

14 Review the information you entered in the previous screen. You can go back to previous screens to edit what you entered in previous steps. You can also edit the CoS and bitrate of the flows shown on the Review Traffic Flow Distribution window. When you are satisfied with the information shown and have added all the flow distributions that you want, click **Next.**

   The Review Customer Topology window opens.

15 If desired, edit the information on following tabs:

   VRFs: Name

   VPN Prefixes: AS Path, Local Pref, MED

   VPN Flows: Cos, Bitrate.

16 Click **Finish** when you are done editing the Review Customer Topology window.

Edit BGP Prefixes

You can modify BGP prefixes from the Planning menu.

To change a prefix, perform the following steps:

1. In Planning mode, choose **Edit BGP Prefix** or **Edit BGP IPv6 Prefixes** from the Planning menu.
2. Click a router on the topology map that advertises the prefix you want to change.
3. The Change Prefixes window opens to show the following columns of information:
   - The **Prefix**—Displays prefixes on the topology map.
   - The **Router/Net**—Lists the router that advertises the prefix and the current topology name.
   - The **Attributes**—Displays prefix attributes, such as AS path, local preference, and next hop IP addresses.
   - The **State**—Displays the current state of the prefix (Up or Down).
   - The **Area or AS**—Displays the AS of the network.
4. Use the **Filter By** drop-down list to select the prefix you want to modify. There are three levels of filtering:
   - Simple filters let you choose a single operator (such as “router”) from the drop-down list and specify one or more parameters (such as router IP addresses or names) to be matched or excluded.
   - Advanced filters let you choose two or more different operators from a list and specify their corresponding parameters to be matched or excluded.
   - Filter expressions let you manually enter a filter expression that is too complex to be set up with either simple or advanced filter menus.

     Click **Show** to display only the items that match the parameters of the filter.

     Click **Hide** to display only the items that do not match the parameters of the filter.

For more information about using filters, see Using Filters on page 221 in Chapter 4, "The History Navigator"

5 Click a prefix in the table to highlight it.

6 Choose a filtering method for the highlighted prefix from the Attributes drop-down list and enter a corresponding value in the text box at the right of the screen.

7 Use the Operation drop-down list to set, append, or prepend the value to the corresponding attribute of the prefix.

For example, to set the local preference value of a prefix to 99, choose **Local Pref** from the Attributes list, enter **99**, and choose **Set** from the Operation list.

8 Click **Apply**.

Edit Traffic Flows

This section applies to RAMS Traffic only.

You can manipulate IPv4 and VPN traffic flows in Planning mode to help plan the most effective routing scenario for your network.

To edit IPv4 traffic flows, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

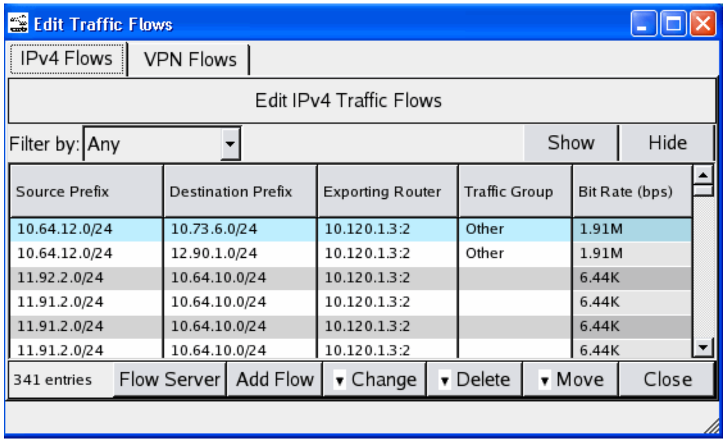

Figure 12 Edit Traffic Flows

2 Limit the list of traffic flows shown using the **Filter by** drop-down list. Filter by the following choices: Any, Source, Destination, Flow Exporter, or use the Advanced Filtering window.

You can perform the following functions using the Edit IPv4 Traffic Flows window. At the bottom of this window, select from the following buttons:

- **Flow Server**—Open the Add Flow Distribution window. See Add a Flow Server on page 273 for more information.
- **Add Flow**—Open the Add IPv4 Traffic Flow window. See Add a Flow Server on page 273 for more information.
- **Change**—Choose **All** to open the Change flow bitrate to window to edit all flows, or choose **Selected** and then make changes in the Bitrate column of the Edit IPv4 Traffic flow window. See Change the Bitrate of a Flow on page 275 for more information.
- **Delete**—Remove one or more flows. See Delete a Flow from a Router on page 276.
- **Mov—**Change the exporting/ingress router for traffic flows. See Move a Traffic Flow from One Router to Another on page 276 for more information.

Add a Flow Server

You can add flow servers from the Planning menu.

To add a flow server, perform the following steps:

1. In Planning mode, choose **Edit Traffic Flows** from the Planning menu.
2. Click **Flow Server** at the bottom of the Edit IPv4 Traffic Flows window to add a flow server to the simulated network.

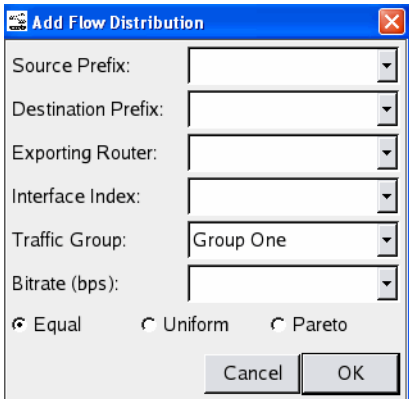

Figure 13 Add Flow Distribution

3. Enter the source prefix, destination prefix, exporting router, interface index, traffic group, and bitrate (bps) of the new server in the appropriate text boxes.
4. Choose how traffic will flow from the new server to its clients by clicking **Equal**, **Uniform** or **Pareto**.
   - **Equal**—Each flow has the average bitrate entered in the previous step.
   - **Uniform**—Each flow has an equal probability of having any bitrate from 0 bps to twice the bitrate entered in the previous step.
   - **Pareto**—Set the average bitrate of the generated flows to the bitrate that was entered in the previous step. This setting is based on the fact that in a flow distribution, a small number of flows has a high bitrate, while most of the flows will have a lower bitrate.
5. Click **OK**.

Add a Traffic Flow

See Add Traffic Flow on page 263 for instructions.

Change the Bitrate of a Flow

Edit the bitrate of a flow to increase or decrease the amount of traffic flowing through the network.

To change the bit rate of a flow, perform the following steps:

1. In Planning mode, choose **Edit Traffic Flows** from the Planning menu.
2. Click a flow in the table to highlight it. Hold down the Ctrl key and click to highlight more than one flow at a time.
3. Right-click the highlighted flows and choose **Change Flow**.

   You can also perform this function using the Change button at the bottom of the Edit IPv4 Traffic Flows window. Choose **Selected** to change the bit rate in only highlighted flow or **All** to change every flow shown in the table.
4. Configure the following values:
   - **Set (bps)**—Assigns a single bitrate value to specified traffic flows. For example, to set the bitrate of specified flows to 5, click **Set** and type 5 in the text box.
   - **Scale (decimal)**—Multiplies the bitrate by N, where N is the value you supply. For example, to triple the bitrate of all highlighted traffic flows, click **Scale** and type 3 in the text box. A bitrate of 2 bps is now 6 bps. A bitrate of 10 bps is now 30 bps.
   - **Add (bps)**—Increases the bitrate of highlighted traffic flows by adding N to the current value, where N is the value you supply. For example, to increment all selected flows by 5, click **Add** and enter 5 in the text box. A bitrate of 2 bps is now 7 bps. A bitrate of 10 bps is now 15 bps.
   - **Add Proportional (bps)**—Increases the bitrate of highlighted flows so the total bitrate of the set of flows is equal to N (where N is the value you supply). For example, to increase the total bitrate of a set of flows from an exporting router to equal 200, click **Add Proportional** and enter 200 in the text box. Each highlighted flow is proportionally increased based on its original value. A set of three flows with bitrates of 103 bps, 23 bps, and 7 bps are increased to 155 bps, 35 bps, and 10 bps.

5 Click **OK**.

Delete a Flow from a Router

You can remove traffic flows from the Edit IPv4 Traffic Flows window, thereby removing them from the topology.

To delete a flow from a router, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Click a flow in the table to highlight the flow. Hold down the Ctrl key and click to highlight more than one flow at a time.

3 Right-click the highlighted flows and choose **Delete Flow** from the pop-up menu.

You can also perform this function using the **Delete** button at the bottom of the Edit Flows window. Choose **Selected** to delete only highlighted flows or **All** to delete every flow in the table.

4 Click **Yes** to complete deletion or click **No** to cancel.

Move a Traffic Flow from One Router to Another

This procedure allows you to change the exporting/ingress router for one or more traffic flows.

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Click a flow in the table to highlight the flow. Hold down the Ctrl key and click to highlight more than one flow at a time. Use the **Filter by** drop-down menu to limit the list of results.

3 Right-click the highlighted flows and choose **Move flow** from the pop-up menu.

You can also perform this function using the **Move** button at the bottom of the Edit Flows window. Choose **Selected** to move only highlighted flows or **All** to move every flow in the table.

4 Enter the address of the router where you'll move the traffic flow.

5 Enter the interface index for the traffic flow exporter.

6 Click **OK** to save your changes.

Edit VPN Traffic Flows

This section applies to RAMS Traffic only. This option is enabled if your appliance is licensed for VPN, and if VPN protocol is present in the opened topology.

To edit VPN traffic flows, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Select the **VPN Flows** tab to open the Edit VPN Traffic Flows window.

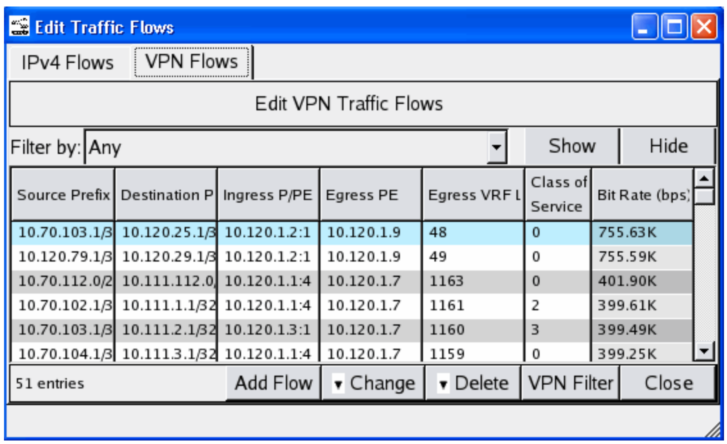

Figure 14 Edit VPN Traffic Flows

This window displays all of the VPN traffic flows currently running in the system. At the bottom of this window are the following buttons:

Add Flow: Select this to invoke the Add VPN Traffic Flow window. See Add a Flow Server on page 273 for more information.

Change: Select this to change the bitrate of all flows in the table, or all flows currently selected in the table. See Change the Bitrate of a Flow on page 275 for more information.

Delete: Select this to delete the bitrate of all flows in the table, or all flows currently selected in the table. See Delete a Flow from a Router on page 276 for more information.

VPN Filter: Select this to filter all of the flows in the window.

Close: Select this to close the window.

Using the VPN Filter

Use this window to display a subset of VPN Traffic Flows.

To use the VPN Filter, perform the following steps:

1 In Planning mode, choose **Edit Traffic Flows** from the Planning menu.

2 Click **VPN Filter** to edit the VPN filters.

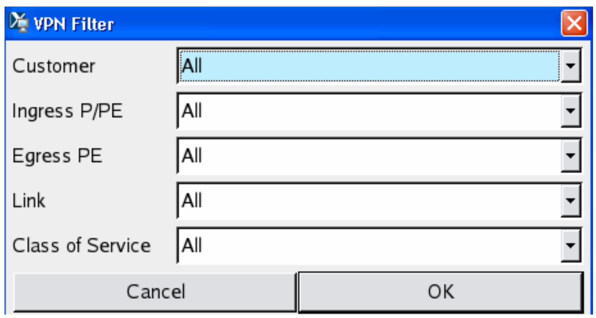

Figure 15 VPN Filter

3 Select values from the drop-down menus, or select **All**.

4 Click **OK**.

Edit Node Properties

Use this option to set the overload bit for IS-IS routers.

To edit a node, perform the following steps:

1 In Planning mode, choose **Edit Router** from the Planning menu**.**

2 On the topology map, select the node you wish to edit.

The network information appears in the Network field, and the router's address or name will display in the Router field.

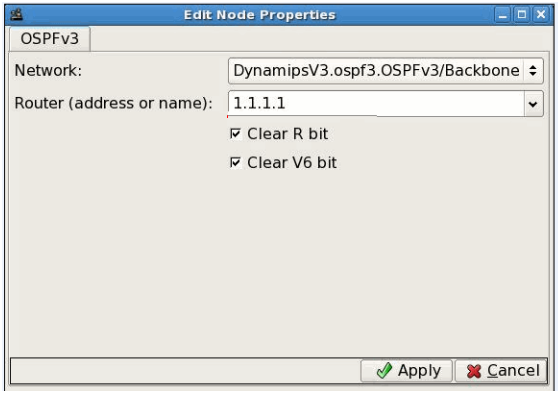

Figure 16 Edit Node Properties

3 (OSPF 3 only) Select check boxes to clear or set the R and/or V6 bits.

4 Select **Set overloaded bit** to disable the transit traffic from being routed.

5 Click **Apply**.

Edit VRFs

This option is enabled if your appliance is licensed for VPN, and if VPN protocol is present in the opened topology.

To edit a VRF, perform the following steps:

1 In Planning mode, choose **Edit VRFs** from the Planning menu.

The Edit VRFs window opens to show all of the VRFs in the current VPN topology.

2 To edit an entry, click the box in which the entry appears.

Edited values are shown in purple. PE, RD, and AS cannot be edited.

3 Click **Apply** to apply the VRF edits.

VRF parameters are cleared when you exit Planning mode.

Bring Down a Router

The Down Router function changes the state of a node from up to down, simulating the impact on network activity should the selected router fail.

To bring down a router, perform the following steps:

1. In Planning mode, choose **Down Router** from the Planning menu.
2. On the topology map, click a node to bring down.

   As an alternative to steps 1 and 2, you can right-click a node to bring up its information panel, then click **Down** on the panel.

   The node turns red, indicating that its state is Down.

   To bring down a single IGP protocol instance on a node, right-click the node to bring up its information panel, select the appropriate tab, and then click **Down**.

3. To bring a node back up, right-click it.

   The node information panel for the node appears.

4. Click **Up** for each protocol tab shown on the node information panel.

   The node is no longer red, indicating that its state is Up.

Bring Down Peerings

The Down Peering function changes the state of a peer relationship from up to down, simulating the impact on network activity should the selected relationship fail.

For BGP protocols, the peerings from and to the NextHop are taken down. For OSPF and IS-IS protocols, both halves of the duplex link are taken down. For the EIGRP protocol, only one half is taken down, but note that all peers of the selected interface are brought down as well.

To bring down a peering, perform the following steps:

1 In Planning mode, choose **Down Peering** from the Planning menu.

2 On the topology map, click the node whose peering you want down.

3 Click the tab for the appropriate protocol.

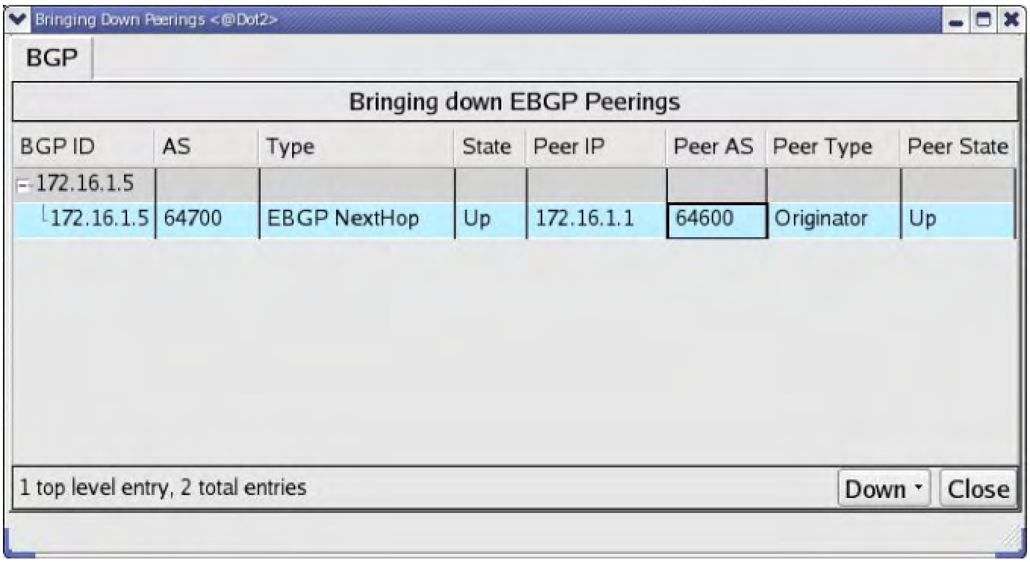

Figure 17 Bringing Down Peerings

4 click a peering to highlight it in the table.

5 Click **Down** or right-click the peerings.

6 Choose **All** to bring down all peerings in the table or **Selected** to bring down only highlighted peerings.

When you click a node to bring down an IPv6 peering, the BGP IPv6 peerings are listed in that table along with BGP IPv4 peerings; however, the BGP ID and Peer IP show only the IPv4 address.

7 Click **Yes** to complete the process of bringing down the peerings. Click **No** to return to the Bringing Down Peerings window without bringing down the peerings.

8 Click **Close** to close the window.

Alternatively, you can bring down a peering by right-clicking a link to bring up its information panel, then clicking **Down** on the panel.

Bring Down a Prefix

The Down Prefix function changes the state of a prefix from up to down, simulating the impact on network activity should the selected prefix be withdrawn.

To bring down a prefix, perform the following steps:

9 In Planning mode, choose **Down IPv4 Prefix** or **Down IPv6 Prefix** from the Planning menu.

10 On the routing topology map, click the node that advertises the prefix you want to bring down.

11 Click the tab for the protocol.

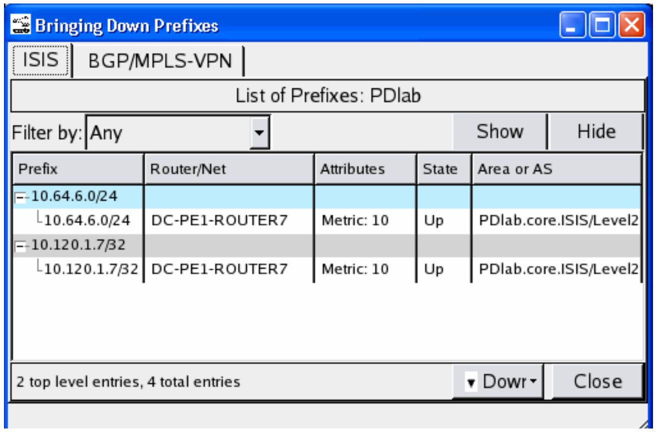

Figure 18 Bringing Down Prefixes

12 Use the **Filter By** drop-down list to choose the attributes you'll use to narrow the list of prefixes shown in the table. For more information about choosing a filter, see Add a Prefix on page 257.

13 Click **Show** to list only items that match the parameters of the filter, or click **Hide** to list only items that do not match the filter parameters.

14 Click a line in the table to highlight the prefix to bring down. Hold down the Ctrl key and click the mouse button to highlight more than one prefix at a time.

15 Right-click the highlighted prefixes and choose **All** or **Selected** from the pop-up menu.

You can also click **Down** at the bottom of the Bringing Down Prefixes window. **All** deletes every prefix listed in the table. **Selected** deletes only highlighted prefixes.

16 Click **Yes** to complete the process of bringing down the prefixes or click **No** to cancel the process.

Bring Down VRF

This option is enabled if your appliance is licensed for VPN and if the VPN protocol is present in the opened topology.

The Down VRF function changes the state of one or more VPN nodes to Down on a particular router, simulating what would happen if the selected item should fail.

To bring down a VPN customer site, perform the following steps:

1 Choose **Down VRF** from the Planning menu.

2 On the topology map, click a VPN node to select the customer site PE router.

3 In the Network field, select the network of the customer you site you want to bring down.

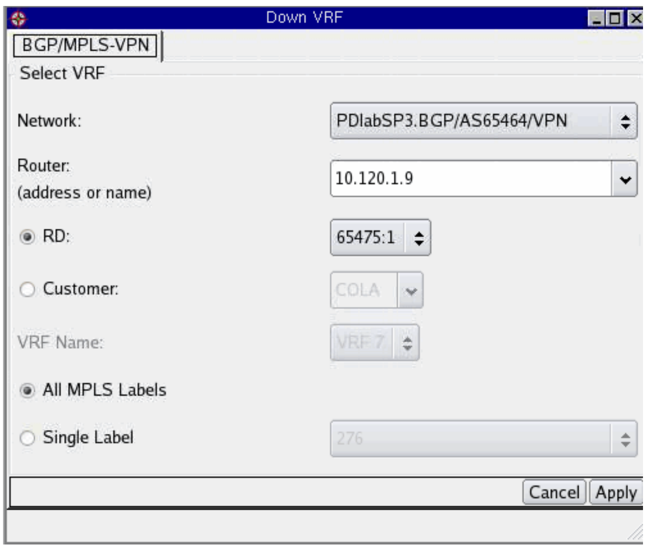

Figure 19 Down Customer Site

4 Select the router IP address or host name in the Router field.

5 Choose **RD** or **Customer**.

- If you choose **RD**, select the RD that identifies the VRF for the customer, and continue to Step 6.
- If you choose the **Customer**, select the VRF name for the customer.

6 Click **Apply**.

Working with Planning Reports

Planning Reports apply to RAMS without traffic only.

RAMS Traffic uses present-time statistics derived from traffic data to help plan changes to the network. Use the resulting information to create network configuration that will maximize available resources and provide consistent quality service to users.

After using Planning mode to make changes to the topology, as described in the previous sections, you can analyze the effects of those simulated changes on such variables as link utilization, bandwidth and available capacity using the Planning Reports window. For example, to see how the addition of a node influences link traffic, use Planning reports to view how much traffic was present on each link before the node was added and after it was added, and the amount of traffic that changed as a result of the edit.

Planning Report Access

You can access planning reports from the Planning menu in RAMS Traffic.

To access planning reports in RAMS Traffic, perform the following steps:

1 Choose **Reports** from the Planning menu.

The Show Topology Edits window opens (Figure 21) to display a list of all edits made to the network (see Planning Report Access on page 285).

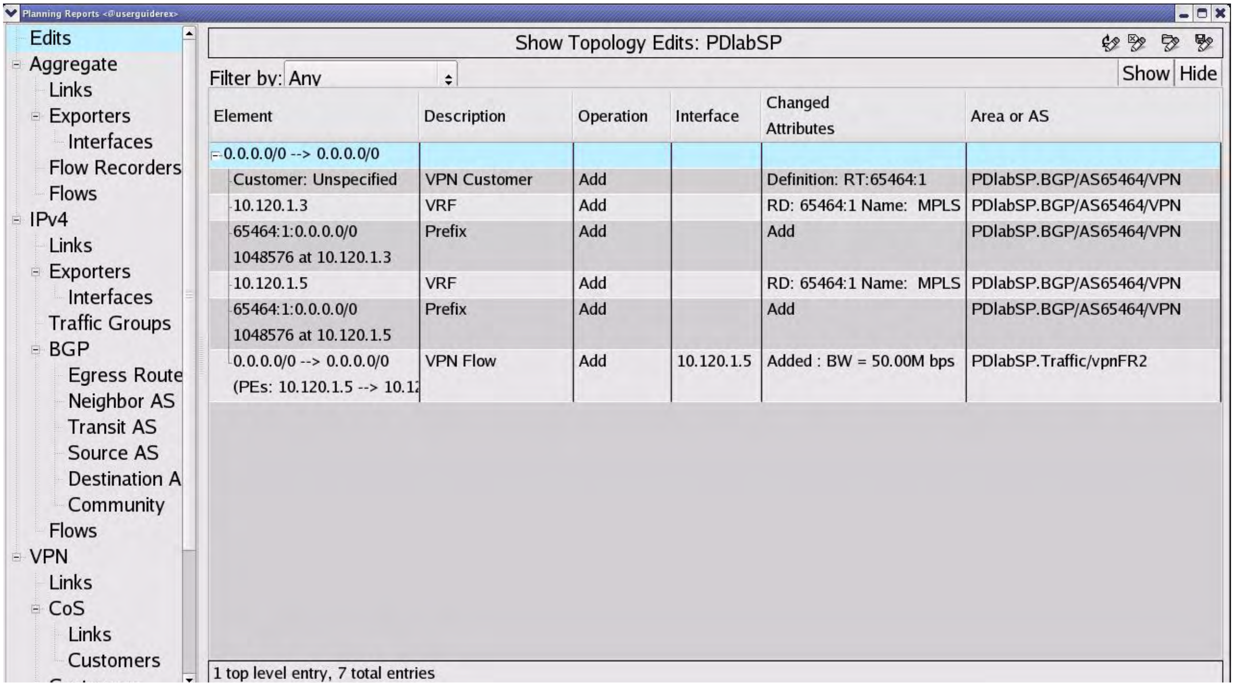

Figure 20 Show Topology Edits Window

2 See these sections for information on working with Planning reports:

- Navigation Tree on page 287
- Drill-down Function on page 287
- Planning Report Icons on page 288
- Edits on page 288
- Advanced Filtering on page 290

Navigation Tree

The navigation tree, shown in the left-hand pane in Figure 21, defines the protocols found in your network. If the network supports multiple protocols, the tree includes an aggregate category for IPv4 and VPN traffic.

You can expand or collapse any category that is preceded by a plus (+) or minus (-) symbol. Within each category are reports that contain columns of generally requested information.

Drill-down Function

A drill-down menu is available at the top-right of each report window (except for the Edit Reports option). Drill-down allows you break down data into finer detail. If drill-down is not available, it is because finer detail is not stored or you have already reached the finest level of detail. For example, an IPv4-related report does not display VPN-related drill-downs. The following drill-down options may be available, depending upon the current report and level:

- BGP Community
- Destination AS
- Egress PE
- Exit Router
- Exporter
- Flows
- Exporters
- Interfaces
- Flow Collector
- IPv4 Flows
- Ingress P
- Ingress PE
- Link
- Neighbor AS

- Traffic Groups
- Transit AS
- VPN Flows

Planning Report Icons

The following icons are found at the top-right of the planning Reports window (see Planning Report Access on page 285).

Table 2 Planning Report Icons

	**Analyze Edits**	Update all traffic and routing edits simultaneously.
	**Undo All Edits**	Reverse the edits made.
	**Import Edits**	Import edits from another database or clipboard
	**Export Edits**	Export edits to another database or clipboard.

Edits

The Show Topology Edits window (Figure 21) displays a list of all edits made to the network (see Planning Report Access on page 285).

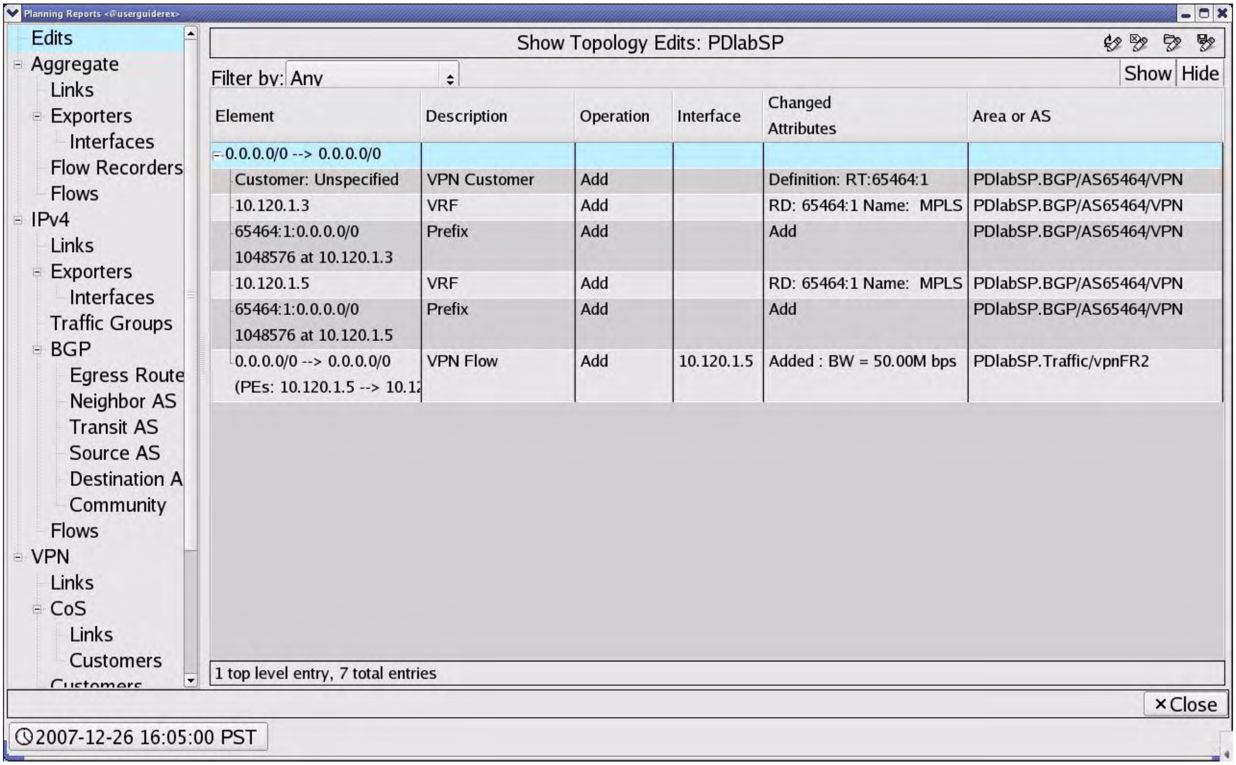

Figure 21 Show Topology Edits Window

The table in this window includes the following columns:

- **Element**—Displays the object (or element). This could be a router name, a link (specified by two end points), or a prefix.
- **Description**—Describes the item in the Element column.
- **Operation**—Displays the type of edit operation that occurred for the element.
- **Interface**—Displays the interface address of the element.
- **Changed Attributes**—Displays what changed for the edited element. For example, if you changed the bitrate of a flow it will display what the bitrate was before and after the edit.

- **Area or AS**—Displays the topology name of the network, which includes information on the administrative name you specified when you set up recording, protocol name, and an AS number or area.

Advanced Filtering

Simple filters let you choose a single operator from a list and specify one or more parameters to be matched or excluded.

Advanced filters let you choose two or more different operators from a list and specify their corresponding parameters to be matched or excluded. From a Filter workspace, select the drop-down list in the Filter by field and choose **Advanced**.

For more information about filtering, see Using Filters on page 221 in the Chapter 4, "The History Navigator"

Understanding Planning Reports

This section describes the contents of the specific Planning reports:

- Aggregate Reports on page 290
- IPv4 Planning Reports on page 294
- BGP Traffic Reports on page 297
- VPN Traffic Reports on page 300

Aggregate Reports

The Aggregate Reports hierarchy is displayed only if both IPv4 and VPN traffic are present. These reports present the total load of both IPv4 and VPN traffic on the network.

The following Aggregate reports may be available:

- Links Report
- CoS Report
- CoS Links Report

- Exporters Report
- Interfaces Report
- Flow Collectors Report
- Flows Report

The Aggregate Report table includes the following columns of information for the protocols shown in the left navigational pane:

- **Element**—Displays the object (or element). This could be a router name, a link (specified by two end points), or a prefix.
- **Description**—Describes the item in the Element column.
- **Operation**—Displays the type of edit operation that occurred for the element.
- **Interface**—Displays the interface address of the element.
- **Changed Attributes**—Displays what changed for the edited elements. For example, if you changed the bitrate of a flow it will display what the bitrate was before and after the edit.
- **Area or AS**—Displays the topology name of the network, which includes information on the administrative name you specified when you set up recording, protocol name, and an AS number or area.

Links Report

Selecting this option invokes the Link Utilization Total Traffic report. Each row in this report represents a link on the topology map. The report contains the following columns:

- **Source**—Identifies the source address of the link.
- **Destination**—Identifies the destination address of the link.
- **Capacity**—Lists the capacity of the link, if known.
- **Traffic Before Edit**—Displays the amount of traffic flowing across the link before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing across the link after editing
- **Traffic Change**—Displays the traffic delta for the link.

CoS Report

This report displays the total amount of VPN traffic seen with each CoS. The report contains the following columns:

- **CoS**—Displays the class of service defined for each router.
- **Traffic Before Edit**—Displays the amount of traffic flowing for each CoS before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for each CoS after editing.
- **Traffic Change**—Displays the traffic delta for the CoS.

CoS Links Report

This report displays information on the CoS links. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **CoS—**Displays the CoS attributed for the link.
- **Traffic Before Edit**—Displays the amount of traffic flowing though the links before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the links after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Exporters Report

Each row shown in this table represents a router that is sending NetFlow data throughout the topology. The report contains the following columns:

- **Exporter**—Displays the address of the exporting router.
- **Traffic Before Edit**—Displays the amount of traffic flowing before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Interfaces Report

Each row represents a router interface where traffic was seen and exported. The report contains the following columns:

- **Exporter: Interface Index**—Identifies a network interface on the exporting router by its Simple Network Management Protocol (SNMP) interface index.
- **Traffic Before Edit**—Displays the amount of traffic flowing before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Flow Collectors Report

Each row in this report represents a Flow Collector on the network. The report contains the following columns:

- Flow Collector—Identifies the Flow Collectors found on the network.
- **Traffic Before Edit—**Displays the amount of traffic flowing across the link before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing across the link after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Flows Report

Each row in this report represents a prefix-to-prefix flow found in the network. The report contains the following columns:

- **Flow Source—**Identifies the source address for the flow.
- **Flow Destination**—Identifies the destination address for the flow.
- **Exporter—**Identifies exporting routers found in the network.
- **Traffic Group—**Identifies all traffic groups for the flow found on the network (IPv4 flows only).
- **ToS/CoS—**Displays Type of Service (ToS) and CoS for the flow.

- **Egress PE**—Displays each PE router found at the edge of an Internet service provider (ISP) (VPN flows only).
- **Traffic Before Edit**—Displays the amount of traffic flowing across the link before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing across the link after editing.
- **Traffic Change**—Displays the traffic delta for the link.

IPv4 Planning Reports

This section describes the following IPv4 Planning reports:

- Links Report
- CoS Report
- CoS Links Report
- Exporters Report
- Interfaces Report
- Traffic Groups Report
- Flows Report

Links Report

This report displays all IPv4 links found in the network. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **Traffic Before Edit**—Displays the amount of traffic flowing though the links before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the links after editing.
- **Traffic Change**—Displays the traffic delta for the link.

CoS Report

This report displays the total amount of VPN traffic seen with each CoS. The report contains the following columns:

- **CoS**—Displays the class of service defined for each router.
- **Traffic Before Edit**—Displays the amount of traffic flowing for each CoS before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for each CoS after editing.
- **Traffic Change**—Displays the traffic delta for the CoS.

CoS Links Report

This report displays information on the CoS links for each CoS level. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **CoS—**Displays the class of service attributed for the link.
- **Traffic Before Edit**—Displays the amount of traffic flowing though the links before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the links after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Exporters Report

Each row shown in this table represents a router that is sending NetFlow data throughout the topology. The report contains the following columns:

- **Exporter**—Identifies the exporting router.
- **Traffic Before Edit**—Displays the amount of traffic flowing through the exporters before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the exporters after editing.

- **Traffic Change**—Displays the traffic delta for the exporters.

Interfaces Report

Each row represents a router interface where traffic was seen and exported. The report contains the following columns:

- **Exporter, Interface Index**—Identifies a network interface on the exporting router by its SNMP interface index.
- **Traffic Before Edit**—Displays the amount of traffic flowing before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Traffic Groups Report

This report displays the IPv4 traffic groups found in the network. For more information about traffic groups, see “Administration” in the *HP Route Analytics Management Software Administrator Guide*. The report contains the following columns:

- **Traffic Group**—Identifies all traffic groups found in the network.
- **Traffic Before Edit**—Displays the amount of traffic flowing for the group before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for the group after editing.
- **Traffic Change**—Displays the traffic delta for the groups.

Flows Report

Each row in this report represents a prefix-to-prefix IPv4 flow found in the network. The report contains the following columns:

- **Flow Source—**Identifies the source address for the flow.
- **Flow Destination**—Identifies the destination address for the flow.
- **Exporter**—Identifies exporting routers found in the network.
- **Traffic Group**—Identifies all traffic groups for the flow found on the network (IPv4 flows only).

- **ToS/CoS**—Displays ToS and CoS for the flow.
- **Egress PE**—Displays each PE router found at the edge of an ISP (VPN flows only).
- **Traffic Before Edit**—Displays the amount of traffic flowing across the link before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing across the link after editing.
- **Traffic Change**—Displays the traffic delta for the link.

BGP Traffic Reports

This section describes the following available BGP Traffic reports:

- Egress Router Report
- Neighbor AS Report
- Transit AS Report
- Source AS Report
- Destination AS Report
- Community Report

Egress Router Report

Each row in this report lists all the neighboring ASs (the immediate neighbor for the Exit Router) found in the network. The report contains the following columns:

- **Egress Router—**Displays each exiting router in the network.
- **Next Hop**—Identifies the next router along the exit router's path.
- **Traffic Before Edit**—Displays the amount of traffic flowing through the neighboring ASs before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for the neighboring ASs after editing.
- **Traffic Change**—Displays the traffic delta for the neighbor ASs.

Neighbor AS Report

Each row in this report lists all the neighboring ASs (the immediate neighbor for the Exit Router) found in the network. The report contains the following columns:

- **Neighbor AS—**Displays every neighboring AS in the network.
- **Traffic Before Edit**—Displays the amount of traffic flowing through the neighboring ASs before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for the neighboring ASs after editing.
- **Traffic Change**—Displays the traffic delta for the neighbor ASs.

Transit AS Report

This report displays all of the traffic passing through a given AS (but not starting or ending with this particular AS) along the path to its destination. The report contains the following columns:

- **Transit AS**—Identifies the transit AS in the network.
- **Traffic Before Edit**—Displays the amount of traffic flowing through transit ASs before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through transit ASs after editing.
- **Traffic Change**—Displays the traffic delta for the ASs.

Source AS Report

This report displays the inferred source AS of all the traffic flows passing through the network.

For every flow, the appliance finds the most-specific BGP route to the source of the flow, and uses the originating AS of this route to determine the breakdown of traffic for this report. This enables the appliance to find the AS that originated the route, which is the most likely source AS. The report contains the following columns:

- **Source AS**—Identifies the source AS in the network.
- **Traffic Before Edit**—Displays the amount of traffic originating from the source ASs before editing occurred.

- **Traffic After Edit**—Displays the amount of traffic flowing through source ASs after editing.
- **Traffic Change**—Displays the traffic delta for the ASs.

Destination AS Report

Each row lists the amount of traffic categorized by the AS of its ultimate destination.

The report contains the following columns:

- **Destination AS—**Identifies each destination AS in the network.
- **Traffic Before Edit**—Displays the amount of traffic flowing through the destination ASs before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the destination ASs after editing.
- **Traffic Change**—Displays the traffic delta for the destination ASs.

Community Report

This report shows the amount of traffic that is destined for routes that belong to different BGP communities. Community attributes are a way for ISPs to apply routing policies to the routes that are received by a particular router. Each route can belong to zero or more communities, and the communities to which a route belongs are carried as attributes of the route.

The report contains the following columns:

- **First 2 Octets—**Displays the first two AS numbers for the given community attribute.
- **Second 2 Octets**—Displays the second two AS numbers for the given community.
- **Traffic Before Edit**—Displays the amount of traffic flowing through exit routers before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic using routes with a given community attribute.
- **Traffic Change**—Displays the traffic delta for the exit routers.

VPN Traffic Reports

The following VPN traffic-related planning reports are available:

- VPN Links Report
- VPN CoS Report
- VPN CoS Links Report
- VPN CoS Customers Report
- VPN Customers Report
- Ingress PE Report
- Exporters Report
- Interfaces Report
- Egress PE Report
- VPN Flows Report

VPN Links Report

This report displays each VPN link that carries traffic in the network. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **Traffic Before Edit**—Displays the amount of traffic flowing though the links before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the links after editing.
- **Traffic Change**—Displays the traffic delta for the link.

VPN CoS Report

This report displays the total amount of VPN traffic seen with each CoS. The report contains the following columns:

- **CoS**—Displays the class of service defined for each router.

- **Traffic Before Edit**—Displays the amount of traffic flowing for each CoS before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing for each CoS after editing.
- **Traffic Change**—Displays the traffic delta for the CoS.

VPN CoS Links Report

This report displays the CoS links for each service class. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **CoS—**Displays the service class for the link.
- **Traffic Before Edit**—Displays the amount of traffic flowing though the links before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through the links after editing.
- **Traffic Change**—Displays the traffic delta for the link.

VPN CoS Customers Report

This report displays the VPN customers seen for each CoS. The report contains the following columns:

- **Customer—**Displays every VPN customer found in the network.
- **CoS—**Defines the CoS for each VPN customer found on the network.
- **Traffic Before Edit**—Displays the amount of traffic flowing through each customer before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through each customer after editing.
- **Traffic Change**—Displays the traffic delta for the customer.

VPN Customers Report

This report displays the traffic seen by each VPN customer. The report contains the following columns:

- **Customer—**Displays every VPN customer found in the network.
- **Traffic Before Edit—**Displays the amount of traffic flowing through each customer before editing occurred.
- **Traffic After Edit—**Displays the amount of traffic flowing through each customer after editing.
- **Traffic Change—**Displays the traffic delta for the customer.

Ingress PE Report

This report displays each ingress PE router found on the network. The report contains the following columns:

- **Ingress PE—**Identifies the ingress PE.
- **Traffic Before Edit—**Displays the amount of traffic flowing through each ingress PE before editing occurred.
- **Traffic After Edit—**Displays the amount of traffic flowing through each ingress PE after editing.
- **Traffic Change—**Displays the traffic delta for the ingress PE.

Exporters Report

This report shows every exporting router found in the network. The report contains the following columns:

- **Exporter—**Displays the address of the exporting router.
- **Traffic Before Edit—**Displays the amount of traffic flowing before editing occurred.
- **Traffic After Edit—**Displays the amount of traffic flowing after editing.
- **Traffic Change—**Displays the traffic delta for the link.

Interfaces Report

Each row represents a router interface where VPN traffic was seen and exported. The report contains the following columns:

- **Exporter: Interface Index**—Identifies a network interface on the exporting router by its SNMP interface index.
- **Traffic Before Edit**—Displays the amount of traffic flowing before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Egress PE Report

This report shows each PE router found at the edge of an ISP. The report contains the following columns:

- **Egress PE**—Identifies the egress PE.
- **Traffic Before Edit**—Displays the amount of traffic flowing through each ingress PE before editing occurred.
- **Traffic After Edit**—Displays the amount of traffic flowing through each ingress PE after editing.
- **Traffic Change**—Displays the traffic delta for the ingress PE.

VPN Flows Report

This report lists every prefix-to-prefix flow for VPN traffic. The report contains the following columns:

- **Flow Source—**Identifies the source address for the flow.
- **Flow Destination**—Identifies the destination address for the flow.
- **CoS—**Identifies the service class for the flow.
- **Exporter**—Identifies exporting routers found in the network.
- **Egress PE**—Displays each PE router found at the edge of an ISP (VPN flows only).
- **Traffic Before Edit**—Displays the amount of traffic flowing across the link before editing occurred.

- **Traffic After Edit**—Displays the amount of traffic flowing across the link after editing.
- **Traffic Change**—Displays the traffic delta for the link.

Working with Capacity Planning Tools

Capacity Planning applies to RAMS Traffic only. To enable capacity planning, you must be in Analysis mode.

Capacity planning enables you to view an estimate of future traffic demands based on past data collection. A linear regression model is used to extrapolate traffic demands into the future. This allows you to plan potential expansion of the network in order to meet future demands. This feature enables you to be sure that the network will have the capacity to carry future traffic loads.

To access the capacity planning reports, perform the following steps:

1 If you are not already in Analysis mode, click the icon in the lower right corner of the window and choose the Analysis mode icon.

2 Choose **Capacity Planning** from the Planning menu. The Capacity Planning window opens.

Navigation Tree and Report Window Settings

The navigation tree in the left pane of the Capacity Planning window defines the protocols of your network. If your network supports multiple protocols, you will see an Aggregate category for IPv4 and VPN traffic.

You can expand or hide any category that is preceded by a plus (+) or minus (-) symbol. Within each category are reports that contain columns of generally requested information.

Use the Trending menu (Figure 22) to set parameters to control the trending period.

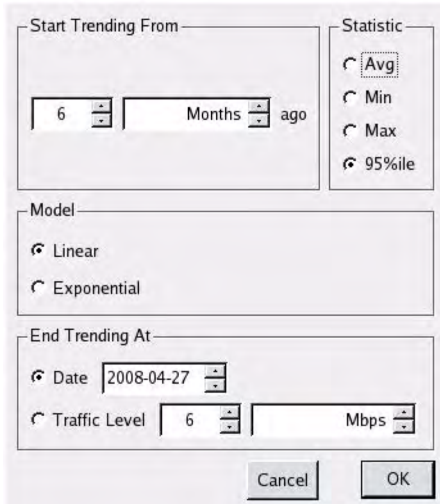

Figure 22 Trending Options

Aggregate Capacity Planning Reports

This section describes the following aggregate capacity planning reports:

- Links Report
- CoS Report
- CoS Links Report
- Exporters Report
- Interfaces Report
- Flow Collectors Report

Links Report

Selecting this option invokes the Links report. Every row displaying in this report represents a link on the topology map. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the link.
- **Daily %-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

CoS Report

This report displays the total traffic amount of all exporting routers with a particular class of service assigned to the routers. The report contains the following columns:

- **CoS**—Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

CoS Links Report

Each line of the displays the total amount of traffic per link at each class of service. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.

- **CoS—**Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Exporters Report

Each row shown in this report represents a router sending NetFlow data throughout the topology. The report contains the following columns:

- **Exporter—**Displays the address of the exporting router.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Interfaces Report

This report displays how much traffic each exporting router detects, broken down by network interface. The interfaces are identified by the exporting router and the SNMP interface index. This report shows the amount of traffic that can be expected to arrive on each router interface at some point in the future. The report contains the following columns:

- **Exporter: Interface Index—**Identifies a network interface on the exporting router by its SNMP interface index.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Flow Collectors Report

Each row in this report represents a Flow Collector on the network. The report contains the following columns:

- **Flow Collector**—Identifies the Flow Collectors found on the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

IPv4 Capacity Planning Reports

This section describes the following IPv4 Capacity Planning reports:

- Links Report
- CoS Report
- CoS Links Report
- Exporters Report
- Interfaces Report
- Traffic Groups Report

Links Report

This report displays every IPv4 link found in the network. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the link.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that

95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

CoS Report

This report displays the total traffic amount of all exporting routers with a particular class of service assigned to the routers. The report contains the following columns:

- **CoS**—Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

CoS Links Report

Each line of the displays the total amount of traffic per link at each class of service.

The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **CoS—**Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Exporters Report

Each row shown in this report represents a router sending NetFlow data throughout the topology. The report contains the following columns:

- **Exporter—**Displays the address of the exporting router.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Interfaces Report

This report displays how much traffic each exporting router detects, broken down by network interface. The interfaces are identified by the exporting router and the SNMP interface index. This report shows the amount of traffic that can be expected to arrive on each router interface at some point in the future. The report contains the following columns:

- **Exporter: Interface Index—**Identifies a network interface on the exporting router by its SNMP interface index.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Traffic Groups Report

This report displays the traffic groups found in the network. For more information about traffic groups, see “Administration” in the *HP Route Analytics Management Software Administrator Guide*. The report contains the following columns:

- **Traffic Group**—Identifies all traffic groups found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that

95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

BGP Capacity Planning Reports

This section describes the following BGP Capacity Planning reports:

- Egress Router Report
- Neighbor AS Report
- Transit AS Report
- Source AS Report
- Destination AS Report

Egress Router Report

This report shows traffic segregated by egress PE. The report contains the following columns:

- **Egress Router**—Identifies each egress router in the network.
- **Next Hop**—Identifies each router along the exit router's path.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Neighbor AS Report

Each row in this report lists all the neighboring ASs (the immediate neighbor for the Exit Router) found in the network. The report contains the following columns:

- **Neighbor AS—**Displays every neighboring AS in the network.
- **Daily 95%-tile—**Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that

95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Transit AS Report

This report displays all of the AS traffic found along the path to its destination. The report contains the following columns:

- **Transit AS**—Identifies the transit AS in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Source AS Report

This report shows the amount of traffic from each source AS at some point in the future, calculated using trending models. The report contains the following columns:

- **Source AS**—Identifies the source AS in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Destination AS Report

Each row lists the amount of traffic categorized by the AS of its ultimate destination. The report contains the following columns:

- **Destination AS—**Identifies each destination AS in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that

95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

VPN Capacity Planning Traffic Reports

The following are VPN traffic-related Capacity Planning reports:

- VPN Links Report
- VPN CoS Report
- VPN CoS Links Report
- VPN CoS Customers Report
- VPN Customers Report
- Ingress PE Report
- Exporters Report
- Interfaces Report
- Egress PE Report

VPN Links Report

Selecting this option invokes the Links report. Every row displaying in this report represents a VPN link on the topology map. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the link.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

VPN CoS Report

This report displays the total traffic amount of all exporting routers with a particular class of service assigned to the routers. The report contains the following columns:

- **CoS**—Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

VPN CoS Links Report

Each line of the displays the total amount of traffic per link at each class of service. The report contains the following columns:

- **Source**—Displays the source address for the link.
- **Destination**—Displays the destination address for the link.
- **Capacity**—Displays the capacity value for the links.
- **CoS—**Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

VPN CoS Customers Report

Each line of this report displays how much traffic is being sent to this customer at each class of service. The report contains the following columns:

- **Customers—**Displays every VPN customer found in the network.
- **CoS—**Defines the service class for each VPN customer found in the network.

- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

VPN Customers Report

This report displays the total number of user-defined VPN Customers found on the network. The report contains the following columns:

- **Customers—**Displays every VPN customer found in the network.
- **CoS—**Defines the service class for each VPN customer found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Ingress PE Report

This report displays each ingress PE router found on the network. The report contains the following columns:

- **Ingress PE**—Identifies the ingress PE.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Exporters Report

Each row shown in this report represents a router sending NetFlow data throughout the topology. The report contains the following columns:

- **Exporter—**Displays the address of the exporting router.

- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Interfaces Report

This report shows the amount of traffic that can be expected to arrive on each router interface at some point in the future. The report contains the following columns:

- **Exporter: Interface Index—**Identifies a network interface on the exporting router by its SNMP interface index.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and five % of the five-minute intervals have traffic demands above this amount.

Egress PE Report

This report shows each PE router found at the edge of an ISP. The report contains the following columns:

- **Egress PE**—Identifies every egress PE router found in the network.
- **Daily 95%-tile**—Displays the 95th percentile of actual traffic for five-minute intervals throughout the day. Of those five-minute intervals, you can compute the 95th percentile. This is the amount of traffic such that 95% of the five-minute intervals have traffic demands below this amount, and 5% of the five-minute intervals have traffic demands above this amount.

Show Edits

This option displays only if the open topology does not include traffic. To see a list of edits made for a topology that includes traffic, see Edits on page 288.

Every edit made to the topology map is displayed in the Show Edits window. To view the edits, choose **Show Edits** from the Planning menu.

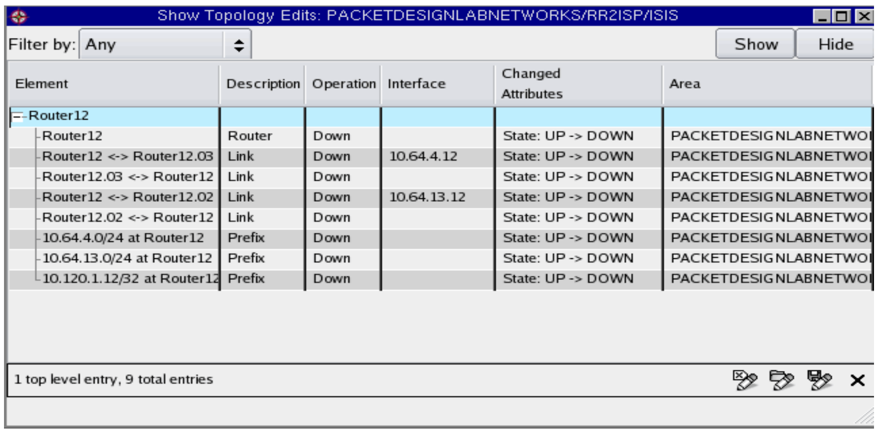

Figure 23 Show Edits Window

This window shows edits that have been made. The following columns are included:

- **Element—**Item that has been changed.
- **Description—**Description of change.
- **Operation—**Operation performed to effect the change.
- **Interface—**Interface involved in the change.
- **Changed Attributes—**Attributes that were modified.
- **Area or AS—**Affected area or AS.

You can import edits, export edits, or undo edits that have been made. Use the buttons shown in Table 3).

Table 3 Events List Controls

	**Undo All Edits**	Clears the list of changes in the Show Topology Edits table and removes the corresponding edits from the topology map. The original layout of the topology map is restored.
	**Import**	Allows you to import edits from the clipboard or a database.
	**Export**	Allows you to send the edits listed in the Show Topology Edits table to either the clipboard or to a database. Exported edits can then be manipulated in external programs, such as Emacs or Excel.

6 BGP Reports

This chapter describes how to use BGP reports to display information about BGP routing events in the network. The BGP reports are accessible from the web interface. For information on the web interface, see the *HP Route Analytics Management Software Administrator Guide*.

Chapter contents:

Understanding BGP Reports

BGP reports allow you to check BGP routing status and view the state of the routing tables to help identify problems. BGP reports are available in HTML format on Route Recorders and on the centralized Modeling Engine in deployments with multiple Route Recorders.

Following initial installation, we recommend that you generate and print all of the BGP reports to obtain a baseline view of network status. When establishing a baseline, the system looks at routes that have been up for more than 80% of the time over the course of seven days. Before a route reaches the seven day mark, its baseline is determined on a day-to-day basis (for example, 80% of 24 hours or 80% of 48 hours).

To identify network problems, begin with the summary report to identify the general problem area and then run additional reports to better identify the problem.

The following types of BGP reports are available:

- Activity Reports—Check BGP routing status on a day-to-day or shift-to-shift basis and quickly identify potential problems.
  - BGP Activity Summary
  - BGP Activity by AS
  - BGP Activity by Peer
  - Route Flap Report
  - Prefix Event Detail
- Logical Topology—View routing tables at specified times to aid in problem identification. These reports also provide a multiple router summary.
  - Route Distribution Detail by RRC, Next Hop, Peer Router, or Next Hop AS
  - Redundancy by Prefix
  - Baseline Redundancy by Prefix
  - AS Reachability
  - Baseline AS Reachability
  - Prefix Reachability

Accessing the BGP Report Pages

If you have a deployment with multiple Route Recorders, we recommend that you access the IGP and BGP reports pages from the centralized Modeling Engine, for the following reasons:

- The Modeling Engine is physically closer to the user.
- Requesting data from the Modeling Engine reduces overhead on the Route Recorder.

When you obtain reports directly from a Route Recorder, information local to the area or protocol being monitored is returned. When you obtain reports from the centralized Modeling Engine, information collected from recorders across entire network is returned.

To access the BGP report pages, perform the following steps:

1. Open a web browser, enter the appliance IP address, and log in as prompted to open the web interface.
2. Choose **Reports Portal**.
3. On the Daily Reports page that opens by default, click **BGP Reports** on the left navigation bar.

   The BGP Reports Activity Summary page opens (Figure 1).

Report data is not available until 15 minutes after recording begins. If you attempt to run a report sooner, a “Report data not available” message is displayed, and you may see an additional message that explains why the report was not generated. If the database is not recording data, ask your administrator to start the recording process.

Figure 1 BGP Reports Page

Generating BGP Activity Reports

This section describes the following BGP activity reports:

- BGP Activity Summary Report on page 322
- BGP Activity by AS Report on page 323
- BGP Activity by Peer Report on page 325
- Route Flap Report on page 326
- Prefix Event Detail on page 327

BGP Activity Summary Report

The BGP Activity Summary report provides a high level overview of BGP network activity over a specified period of time, including any changes or problems with the network. Changes may include new peering sessions or routers appearing in the network.

We recommend that you run the report daily or on a per-shift basis to quickly determine if there is a problem within the BGP network. Problems can include instabilities caused by convergence failures, oscillations, or unstable links or routers.

If an unusual amount of activity is spotted, you can run the History Navigator to obtain more information about the events occurring during this time period (see Chapter 4, "The History Navigator"). If BGP activity is high and stays high, it may indicate a configuration error during scheduled maintenance. You can also use the data in this report to obtain a high-level view of the scaling characteristics of the network as new routes, routers, and peers are added to the network.

To generate the BGP Activity Summary report, perform the following steps:

1. From the web interface, choose **Reports Portal**.
2. Choose **BGP Reports > Activity Summary**.
3. Choose a time period and database from the drop-down lists.

   ➤ If you specify a time period longer than the number of days of data in the database, the generated report begins with the first recorded data. This also causes the generated report to display an ending time that is in the future.

4. Click **Create Report**.

If data is available as specified, the report is displayed with the following information:

- Total churn – sum of all the types of churn.
- Internal update churn – number of internal prefixes.
- External update churn – number of updates from external peers.
- Internal withdrawals – number of withdrawn internal prefixes.

External withdrawals – number of withdrawn prefixes from external peers.

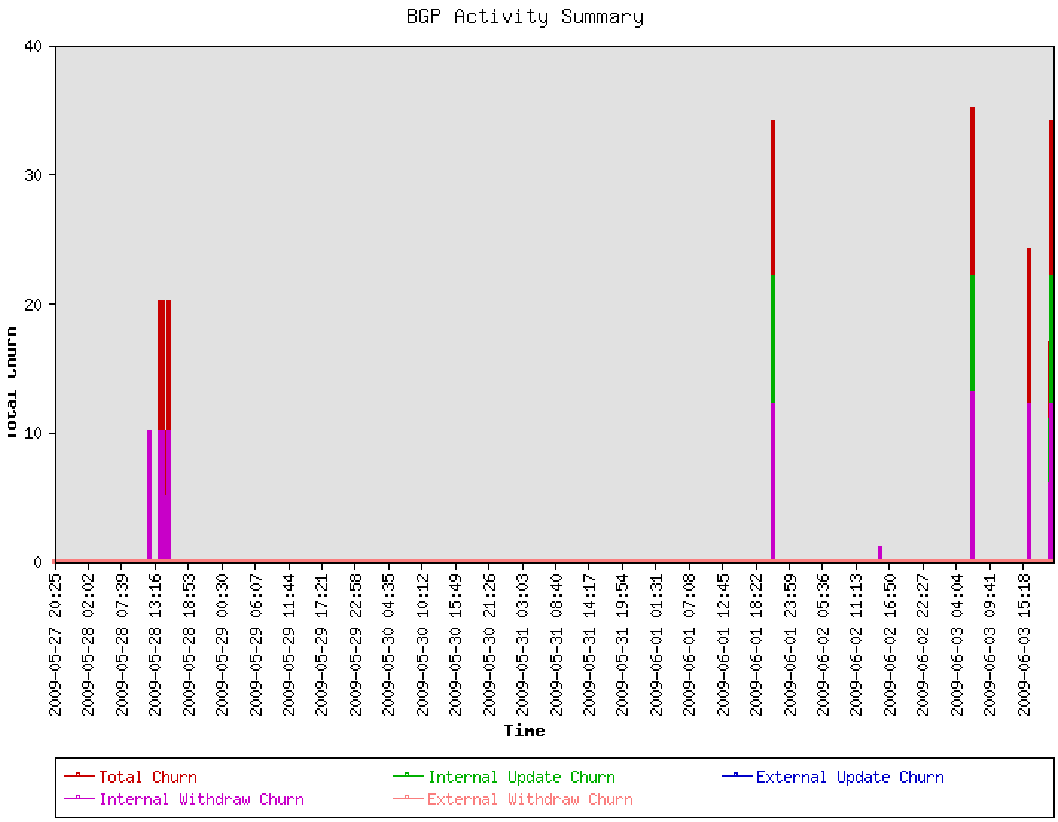

Figure 2 BGP Activity Summary

BGP Activity by AS Report

The BGP Activity by AS report (Figure 3) provides a filtered view of BGP activity for the individual ASs that comprise the entire network. For each AS, you can identify sources of instability or excessive activity. This report is useful for private enterprise networks in which BGP connects multiple ASs.

To generate the BGP Activity by AS report, perform the following steps:

1 From the web interface, choose **Reports Portal**.

2 Choose **BGP Reports > Activity by AS**.

3 Select the desired database from the drop-down list.

4 Click **List AS**.

If data is available as specified, the report opens to display a graph of updates and withdrawals.

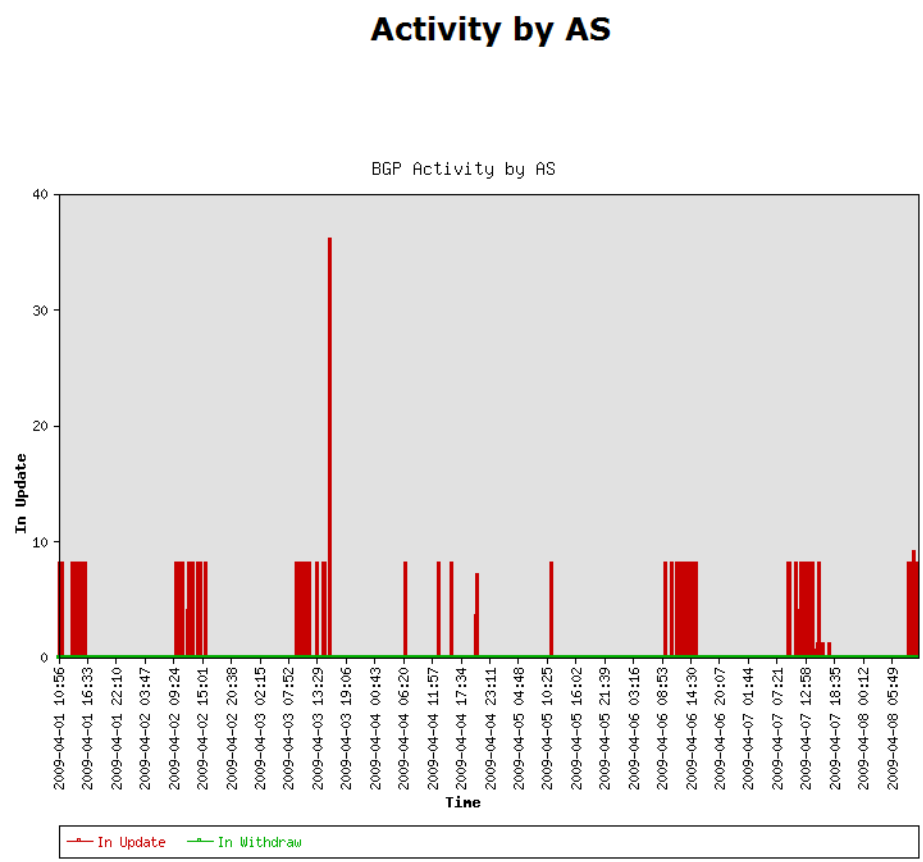

Figure 3 BGP Activity by AS Report

After viewing this report, refer to the BGP Activity Summary report to view specific sources of instability, or open the History Navigator window and perform an event analysis for the time period (see Event Analysis on page 193).

BGP Activity by Peer Report

The BGP Activity by Peer report (Figure 4) allows you to identify the BGP peers that are most active and to diagnose internal churn. Start with the BGP Activity Summary report and then run this report if you notice an unusual amount of activity.

To generate the BGP Activity by Peer report, perform the following steps:

1 From the web interface, choose **Reports Portal**.

1 Choose **BGP Reports > Activity by Peer**.

2 Select the desired database from the drop-down list.

3 Click **List Peer**.

4 Select the desired peer from the Peer drop-down list.

5 Select the desired time period from the Time drop-down list.

6 Click **Create Report**.

If data is available as specified, a line chart is displayed showing update churn (updated prefixes) and withdrawn churn (prefixes withdrawn).

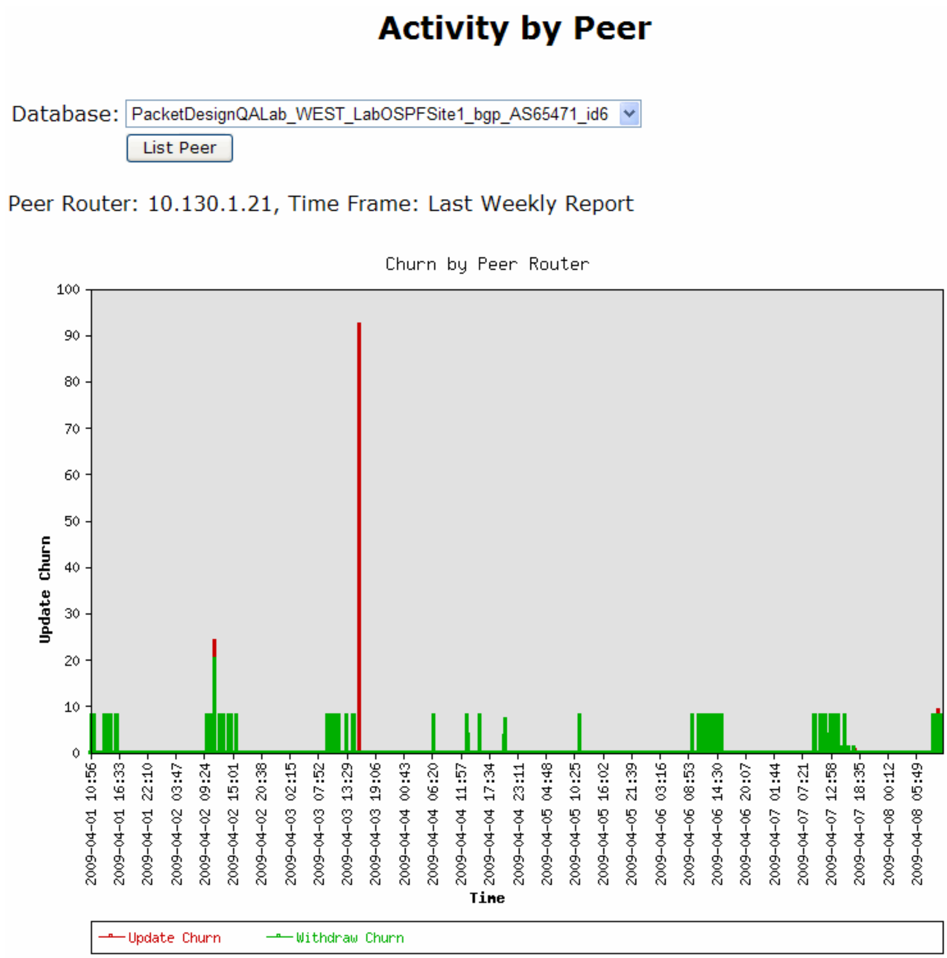

Figure 4 BGP Activity by Peer Report

Route Flap Report

The Route Flap report (Figure 5) provides a list of prefixes that have oscillated ("flapped") between announced and withdrawn from the BGP protocol. It provides a way to quickly identify where service has been lost or degraded due to flapping links and which routers are responsible for activity spikes. We recommend that you run this report as a periodic status check to determine if a problem requires further investigation.

To generate the Route Flap report, perform the following steps:

1 From the web interface, choose **Reports Portal**.

1 Choose **BGP Reports > Route Flap Report**.

2 Select the desired time period and database from the drop-down lists.

3 Click **Create Report** to display the report.

If data is available as specified, the report is displayed. To sort on a column, click the column header. Click again to reverse the sort order.

Route Flap Report (Last 15 Mins)

Database: PacketDesignQALab_WEST_LabOSPFSite1_bgp_AS65471_id6

Create Report

Report creation Time: 2009-04-08 11:00
(Considering 1000 most flapping routes)

Showing entries 1 - 9 of 9

Prefix	Peer ID	▼ Flaps	Last Update	Current State
10.132.1.48/32	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.132.1.47/32	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.132.1.42/32	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.72.15.0/24	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.72.11.0/24	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.72.10.0/24	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.72.9.0/24	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.72.6.0/24	10.130.1.21	2	Apr 8 10:57:44 2009	Announce
10.70.211.0/24	10.130.1.21	1	Apr 8 10:46:01 2009	Announce

Figure 5 Route Flap Report

Prefix Event Detail

The Prefix Event Detail report shows how long a problem has existed and identifies the affected prefixes. For example, you can run this report if the Route Flap Report report identifies a flapping prefix. You can see how long a prefix has experienced intermittent service and identify the customers that have been affected by the associated service degradation. With this report, it is not necessary to log into each individual router and view the routing tables to identify the problem.

To generate the Prefix Event Detail report, perform the following steps:

1 From the web interface, choose **Reports Portal**.

1 Choose **BGP Reports > Prefix Event Detail**.

2 Select the database from the drop-down lists.

3 Select a date and time from the Time drop-down lists.

4 Enter the in the Prefix text box.

5 Click **Create Report** to display the report.

If data is available as specified, the report is displayed. To sort on a column, click the column header. Click again to reverse the sort order.

Prefix Event Detail

Event Detail for Prefix: 167.15.30.16/28

Showing entries 1 - 16 of 16

Time	Direct Peer	Op	Attributes
May 27 08:55:24 2009	10.120.1.5	Announce	ORIGIN: IGP ASPATH: NEXT_HOP: 10.64.15.167 MED: 300 LOCAL_PREF: 300 ORIGINATOR_ID: 167.167.167.167 CLUSTER_LIST: 10.120.1.5 EXT_COMMUNITIES: RT:213:213
May 28 12:32:52 2009	10.120.1.5	Withdraw	ORIGIN: IGP ASPATH: NEXT_HOP: 10.64.15.167 MED: 300 LOCAL_PREF: 300 ORIGINATOR_ID: 167.167.167.167 CLUSTER_LIST: 10.120.1.5 EXT_COMMUNITIES: RT:213:213
May 28 12:34:54 2009	10.120.1.5	Announce	ORIGIN: IGP ASPATH: NEXT_HOP: 10.64.15.167 MED: 300 LOCAL_PREF: 300 ORIGINATOR_ID: 167.167.167.167 CLUSTER_LIST: 10.120.1.5 EXT_COMMUNITIES: RT:213:213

Figure 6 Route Flap Report

Creating BGP Logical Topology Reports

This section describes the following BGP logical topology reports:

- Route Distribution Detail Report on page 329
- Redundancy by Prefix Report on page 334
- Baseline Redundancy by Prefix Report on page 335
- AS Reachability Report on page 335
- Baseline AS Reachability Report on page 336
- Prefix Reachability Report on page 337

To configure a BGP logical topology report, perform the following steps:

1. Open the web application and choose **Reports Portal**.
2. Click the name of the report to configure in the Logical Topology section.
3. On each report page, select the desired database from the **Database** drop-down list.
4. Select a date and time from the Time drop-down lists.
5. Click **Create Report**.

If data is available as specified, the report is displayed. For tabular reports, click a column header to sort on that column. Click again to reverse the sort order.

Route Distribution Detail Report

The Route Distribution Detail report (Figure 7) provides information on the distribution of BGP routes as determined by the BGP path selection algorithm. The distributions are important in traffic engineering, capacity planning, and maintenance. During troubleshooting, you can refer to this report if there is a problem with traffic getting to a particular AS from a particular AS over a BGP link.

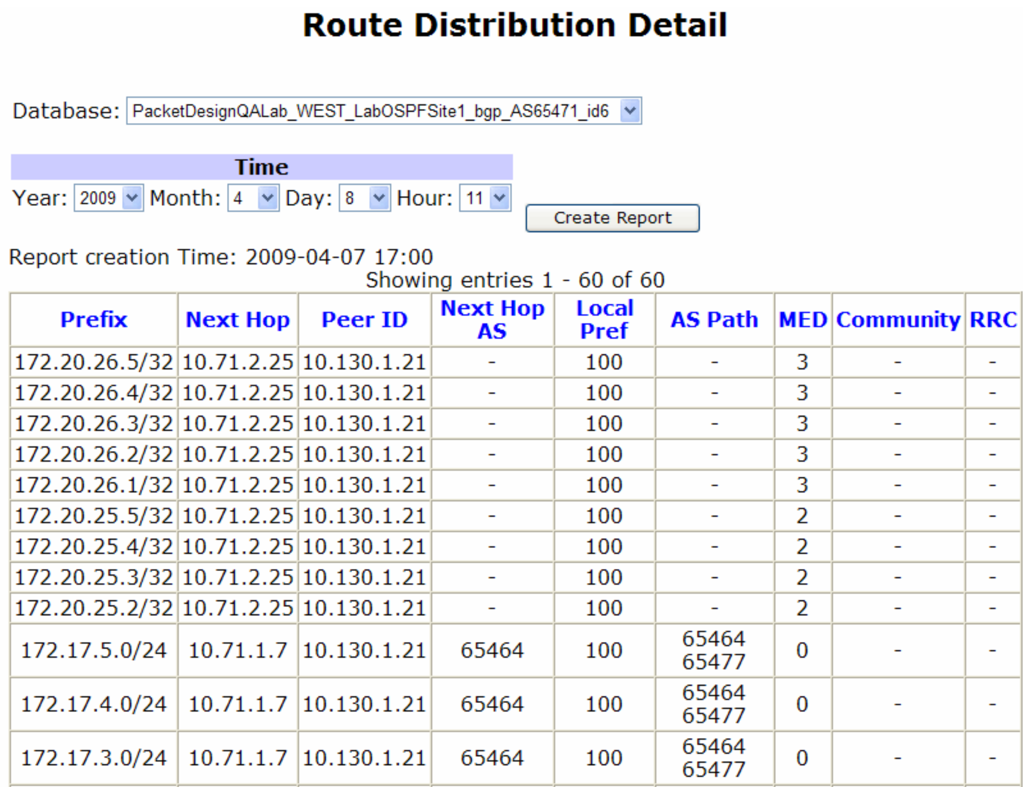

Route Distribution Detail

Database: PacketDesignQALab_WEST_LabOSPFSite1_bgp_AS65471_id6

Time

Year: 2009 Month: 4 Day: 8 Hour: 11 Create Report

Report creation Time: 2009-04-07 17:00

Showing entries 1 - 60 of 60

Prefix	Next Hop	Peer ID	Next Hop AS	Local Pref	AS Path	MED	Community	RRC
172.20.26.5/32	10.71.2.25	10.130.1.21	-	100	-	3	-	-
172.20.26.4/32	10.71.2.25	10.130.1.21	-	100	-	3	-	-
172.20.26.3/32	10.71.2.25	10.130.1.21	-	100	-	3	-	-
172.20.26.2/32	10.71.2.25	10.130.1.21	-	100	-	3	-	-
172.20.26.1/32	10.71.2.25	10.130.1.21	-	100	-	3	-	-
172.20.25.5/32	10.71.2.25	10.130.1.21	-	100	-	2	-	-
172.20.25.4/32	10.71.2.25	10.130.1.21	-	100	-	2	-	-
172.20.25.3/32	10.71.2.25	10.130.1.21	-	100	-	2	-	-
172.20.25.2/32	10.71.2.25	10.130.1.21	-	100	-	2	-	-
172.17.5.0/24	10.71.1.7	10.130.1.21	65464	100	65464 65477	0	-	-
172.17.4.0/24	10.71.1.7	10.130.1.21	65464	100	65464 65477	0	-	-
172.17.3.0/24	10.71.1.7	10.130.1.21	65464	100	65464 65477	0	-	-

Figure 7 Route Distribution Detail Report

Route Distribution By RRC Report

The By Route Reflector Client (By RRC) report shows the number of routes from each route reflector client in the topology. You can view the route distribution information for a particular route reflector client by entering its address into the RRC Detail text box, and clicking **Create Report**. This produces a report showing a table of the routes from the selected route reflector client.

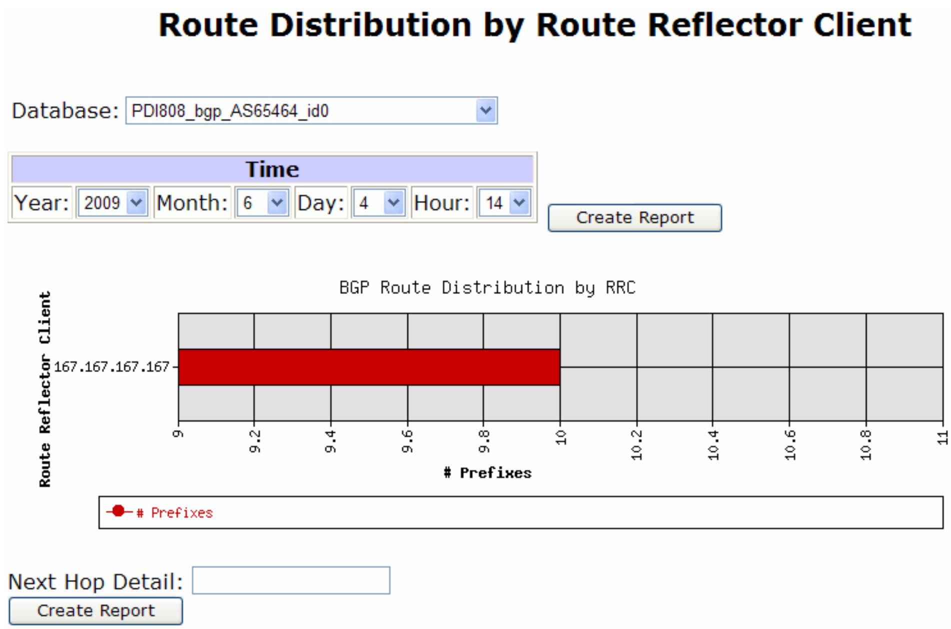

Figure 8 Route Distribution Detail by RRC

Route Distribution By Next Hop Report

The By Next Hop report (Figure 9) shows the number routes with each next hop in the topology. You can view the route distribution information on the routes that contain a particular next hop by entering its address into the Next Hop Detail text box, and clicking **Create Report**. This produces a report showing a table of the routes containing the specified next hop information.

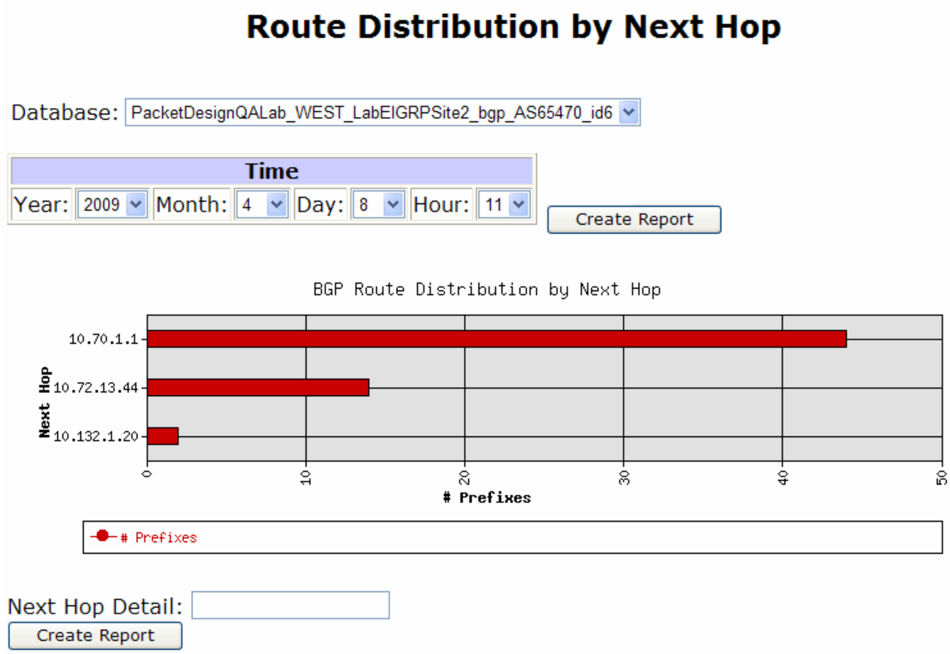

Figure 9 Route Distribution Detail By Next Hop Report

Route Distribution By Next Hop AS Report

The By Next Hop AS report (Figure 10) shows the number of routes with next hop AS in the topology. You can view the route distribution information on the routes that contain a particular next hop AS by entering its AS number into the Next Hop AS Detail text box, and clicking **Create Report**. This produces a report showing a table of the routes containing specified next hop AS information.

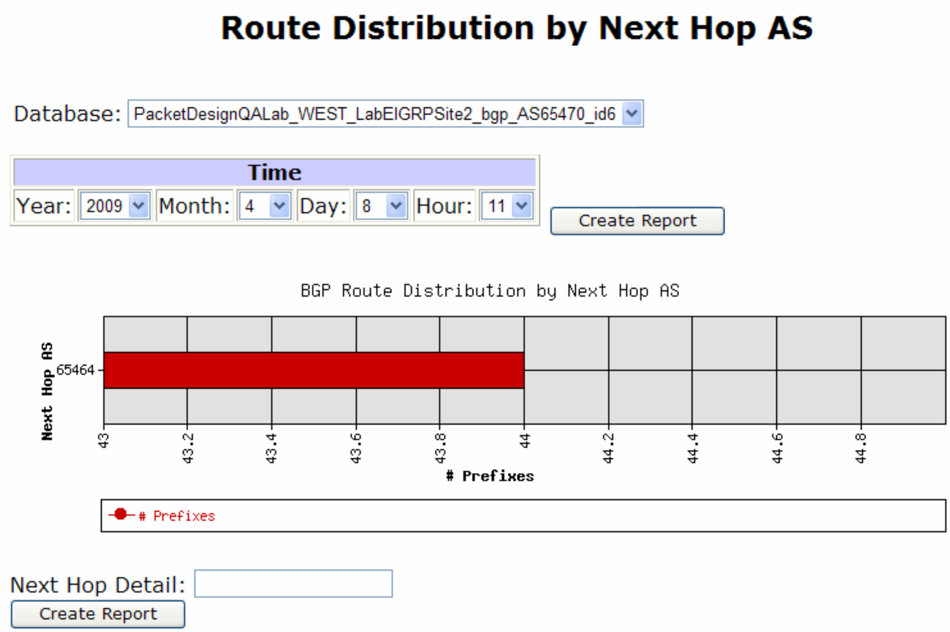

Figure 10 Route Distribution By Next Hop AS Report

Route Distribution Detail By Peer Router

The By Peer Router report (Figure 11) shows the number of routes from each peer router in the topology. You can then view the route distribution information for a particular peer by entering its address into the Peer Router Detail text box, and clicking **Create Report**. This will produce a report showing a table of the specified peer routes.

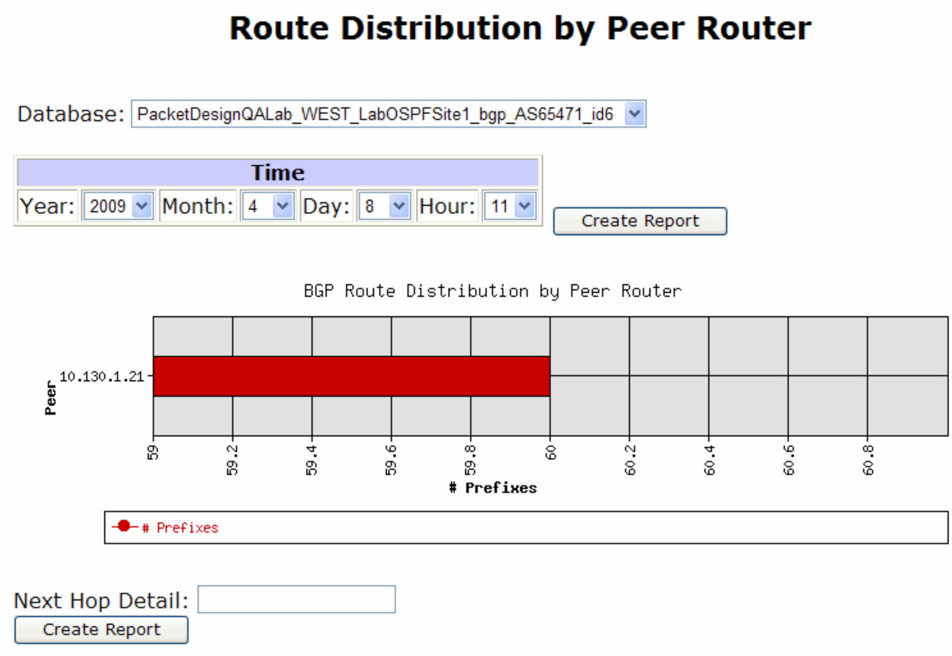

Figure 11 Route Distribution Detail By Peer Router

Redundancy by Prefix Report

The Redundancy by Prefix report displays the degree of redundancy available for each routed prefix on the network and the number of available hops for each route where the number of available hops differs from the number of baseline hops (see Baseline Redundancy by Prefix Report on page 335). You can use this report along with the Baseline Redundancy by Prefix report to see the comprehensive prefix redundancy of the topology. Use this report periodically to review network redundancy and check that redundant paths are available as planned. If you are involved in network planning and design, you will also find this report useful as you plan network updates and expansion.

To identify prefixes that are available only through a single point path, sort the report by the number of available next hops.

The baseline number of next hops is calculated by looking at routes that have been up for more than 80% of the time over a course of seven days. For example, 80% of the first 24 hours, 80% of 48 hours, etc.

Baseline Redundancy by Prefix Report

The Baseline Redundancy by Prefix report (Figure 12) identifies whether or not each routed prefix on the network is available. Run this report to ensure that all network prefixes are available.

Baseline Redundancy by Prefix

Database: PacketDesignQALab_WEST_LabOSPFSite1_bgp_AS65471_id6

Time

Year: 2009 Month: 4 Day: 8 Create Report

Report creation Time: 2009-04-07 17:00

Showing entries 1 - 60 of 60

Prefix	Baseline # Next Hops	Next Hop(s)	Next Hop AS(es)
172.20.26.5/32	1	10.71.2.25	-
172.20.26.4/32	1	10.71.2.25	-
172.20.26.3/32	1	10.71.2.25	-
172.20.26.2/32	1	10.71.2.25	-
172.20.26.1/32	1	10.71.2.25	-
172.20.25.5/32	1	10.71.2.25	-
172.20.25.4/32	1	10.71.2.25	-
172.20.25.3/32	1	10.71.2.25	-
172.20.25.2/32	1	10.71.2.25	-
172.17.5.0/24	1	10.71.1.7	65464

Figure 12 Baseline Redundancy by Prefix Report

AS Reachability Report

The AS Reachability report shows the connectivity, next hops, and AS paths toward all reachable ASs. You can sort through the list of reachable ASs, where the number of available hops to the AS differs from the number of baseline hops to the AS, and the paths taken to reach them. Use this report with the Baseline AS Reachability report to see the comprehensive AS reachability for the topology.

This report can help validate security and routing policies to ensure that there are no single points of failure between the network and key external ASs. You can refer to this report during planning to see whether there is adequate network redundancy.

Baseline AS Reachability Report

The Baseline AS Reachability report (Figure 13) shows when an entire AS is reachable displays the AS paths. This report can help ensure that all ASs are monitored as planned (baseline) and to assist in network planning.

The baseline number of next hops is calculated by looking at routes that have been up for more than 80% of the time over a course of seven days. For example, 80% of the first 24 hours, 80% of 48 hours, etc

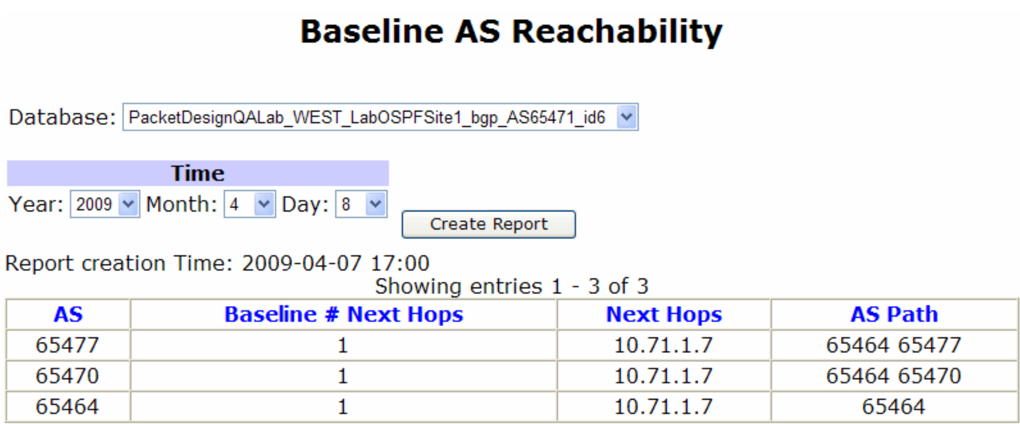

AS	Baseline # Next Hops	Next Hops	AS Path
65477	1	10.71.1.7	65464 65477
65470	1	10.71.1.7	65464 65470
65464	1	10.71.1.7	65464

Figure 13 Baseline AS Reachability Report

Prefix Reachability Report

The Prefix Reachability report (Figure 14) indicates the degree of connectivity and BGP attributes for prefixes that are routable by BGP across the entire network. It can help validate routing policies and identify configuration errors. During network design, it provides a simple way to identify the paths chosen by BGP for a given prefix or set of prefixes.

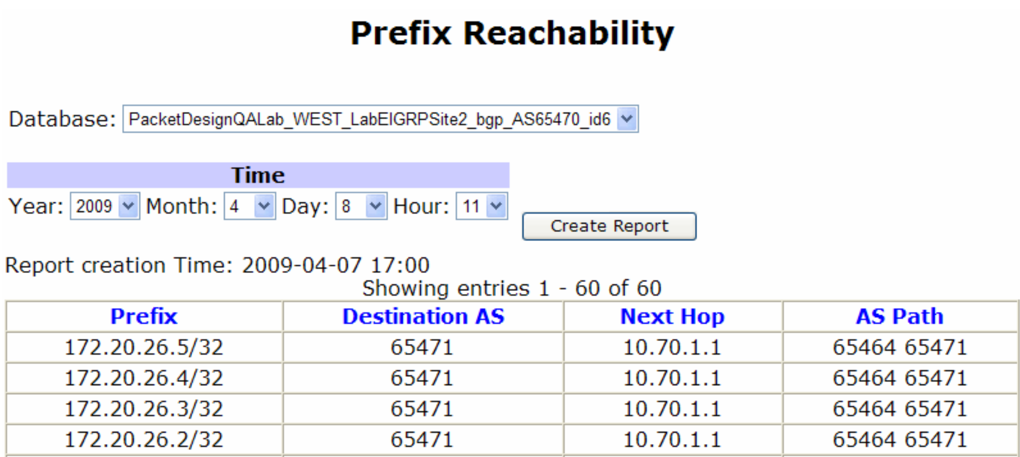

Prefix	Destination AS	Next Hop	AS Path
172.20.26.5/32	65471	10.70.1.1	65464 65471
172.20.26.4/32	65471	10.70.1.1	65464 65471
172.20.26.3/32	65471	10.70.1.1	65464 65471
172.20.26.2/32	65471	10.70.1.1	65464 65471

Figure 14 Prefix Reachability Report

7 IGP Reports

This chapter describes how to use IGP reports to display information about IGP routing events in the network.

Chapter contents:

Understanding IGP Reports

IGP reports allow you to view activity and resources, display network and configuration changes, and identify potential problems. Following initial installation, we recommend that you generate and print all of the IGP reports to obtain a baseline view of network status.

To identify network problems, begin with a summary report to identify the general problem area. You can then run additional reports to better identify the problem. For example, if the summary report displays a large number of flapping links, run a Flapping Links report to further analyze the problem.

The following types of IGP reports are available:

- Summary Reports—View high-level network activity over a specified time period, display day-to-day changes, and quickly flag potential problems.
  - Network Events Summary
  - Network Churn
- Drill-Down Reports—View detailed information to verify configuration changes after troubleshooting problems.

- Changed Metrics
- Flapping Links
- Prefix Origination Changes
- New Prefixes
- New Routers and Links
- Prefixes Withdrawn

• Inventory Reports—View available network resources.
  - Prefix List
  - Prefix Origination from Multiple Sources

Configuring IGP Report Pages

If you have a deployment with multiple Route Recorders, we recommend that you access the IGP and BGP reports pages from the centralized Modeling Engine, for the following reasons:

- The Modeling Engine is physically closer to the user.
- Requesting data from the Modeling Engine reduces overhead on the Route Recorder.

When you obtain reports directly from a Route Recorder, information local to the area or protocol being monitored is returned. When you obtain reports from the centralized Modeling Engine, information collected from recorders across entire network is returned.

Figure 1 IGP Reports Page

To configure IGP reports, perform the following steps:

1 Open a web browser, enter the appliance IP address, and log in as prompted to open the web interface.

2 Choose **Reports Portal > IGP Reports** to open the IGP Reports page.

Figure 2 IGP Reports Page

3 Choose a report from the drop-down list and click **Configure Report**.

4 Configure the following report options:

- Choose the desired database from the Administrative Domain drop-down list.
- For reports that present interval options, choose the top button to specify a time period up till the current time or the bottom button to specify an exact start and end time (Figure 3).

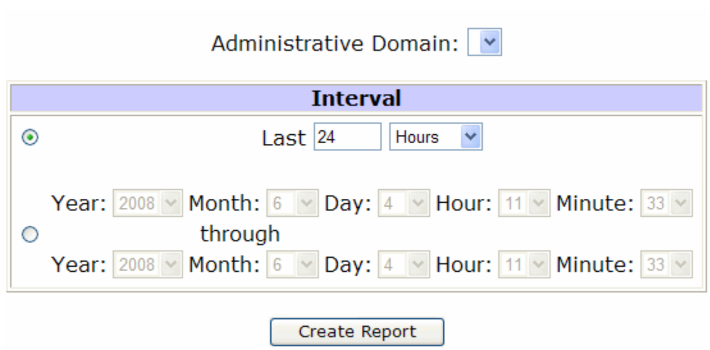

Figure 3 Configuring IGP Reports - Interval Options

- For reports that present a time option, choose a time period from the drop-down lists. For the Prefix Origination from Multiple Sources report, you can also specify the minimum number of originating routers (Figure 4).

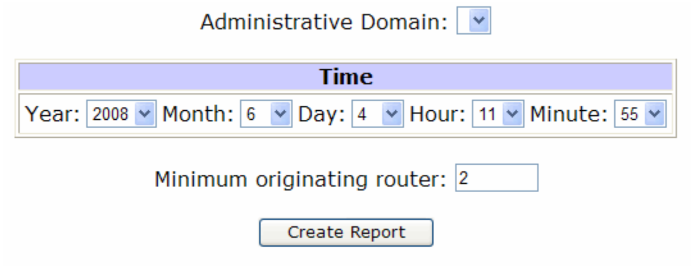

Figure 4 Configuring IGP Reports - Time Options

5 Click **Create Report**.

The time it takes to generate a report depends on the input parameters and size of the database. Reports from large databases normally take longer to generate than those from small databases.

The report is generated from the selected database and displayed in a new page. You can use the print command on your browser to print a copy of the report or click **Reconfigure Report** to change the time period.

Understanding IGP Report Contents

This section describes the IGP report contents:

Network Events Summary

The Network Events Summary report summarizes network changes for the specified time period and is a good place to start when diagnosing network problems or checking general network status. Table 1 describes the fields in the report.

Table 1 Network Events Summary Fields

Field	Description
Number of Link Flaps	Number of times the interface goes up and down.
Number of Links with Changed Metrics	Number of links that have had a metric change between the two time frames, along with the values at the beginning and end.
Number of Events	Number of events recorded in the database.
Number of Prefix Origination Changes	Number of prefixes advertised on a different router.
Number of Prefixes Withdrawn	Number of prefixes that have been withdrawn from the network.
Number of New Prefixes Advertised	Number of prefix advertisements recorded.

Changed Metrics

The Changed Metrics report provides a summary of all link metrics that have changed in the network. It identifies the router that is advertising the changed metric, the original metric, the new metric, and the time when it changed. Run this report after scheduled maintenance to confirm that planned metric changes have occurred. Table 2 describes the fields in the report.

Table 2 Changed Metric Report Fields

Field	Description
Link	Link from source router to the destination router.
Source Interface	Source of the link.
Destination Interface	Destination of the link.
Original Metric	Metric originally assigned to the interface.
New Metric	New metric assigned to the interface.
Time	Date and time the change occurred.

Flapping Links

The Flapping Links report lists all routing links that have gone down and come up recently, including the source router and interface that is flapping, how many times it has changed its status during the period, and when the last change occurred. You can sort the report to list the links with the highest flap count at the top.

Run this report if you suspect a problem with links in the network. For example, the Network Events Summary report may display a high incidence of flapping link events, and this report will indicate which router links are flapping. This report is also useful in situations where the real-time topology map displays links that are going up and down.

In cases where a route is flapping an abnormally high number of times, you can configure an alert on the link to monitor it more closely for a period of time to ensure that an outage is avoided. See Chapter 12, “Alerts” Table 3 describes the fields in the Flapping Links report.

Table 3 Flapping Links Report Fields

Field	Description
Link	Link from the source to destination router
Source Interface	The source of the link
Destination Interface	The destination of the link
Count	The number of times the link has changed state
Time	The date and time the last change occurred

Network Churn

The Network Churn report displays a summary of routing events that took place over the selected time period and identifies all of the sources (routers) and number of events attributed to each source. The outing events tabulation excludes “hello” packets, which are exchanged periodically. Run this report if the Network Events Summary report displays an unusually high level of network events.

The report has two columns. Table 4 describes the fields in the Router column and Table 5 describes the fields in the Number of Events column.

Table 4 Router Column Fields

Field	Description
Name	Router name provided by the routing protocol, if available.
IP Address	Address of the router originating the events.

Table 5 Number of Events Column Fields

Field	Description
Total	Cumulative total of all events.
Router	Number of router events: •Router dynamic hostname change (for IS-IS only) •IS-IS overload bit change •Router type change, for example between Internal and ABR for OSPF
Prefix	Number of prefix events for this router: •Addition and dropping of prefix adjacencies •Changes in the metric to a prefix
Link	Number of link events for this router: •Addition and dropping of neighbor router adjacencies, including peering •Changes in the metric on the link to a neighbor

New Prefixes

The New Prefixes report lists all the newly advertised prefixes for the report period and the advertising frequency. Sources of new prefixes include new links, networks, tunnels, and routers. Run this report after scheduled maintenance to verify the changes. Table 6 describes the fields in the report.

Table 6 New Prefixes Fields

Field	Description
Prefix	IP address of the new prefix.
System ID	System Identifier for the router.
Name	Name of the router.

Table 6 New Prefixes Fields

IP Address	IP address of the router. This column may not display if the routers do not have an IP address attributed to them.
Time	Time the new prefix first appeared on the network.
Count	Number of times the prefix was advertised.

New Routers and Links

The New Routers and Links report lists newly advertised source and destination routers and links in the network. Run this report after new routers are inserted into the network or new links are set up to verify the changes. Table 7 describes the router columns in the report and Table 8 describes the links columns.

Table 7 New Routers Table Fields

Field	Description
IPv4 Address	IPv4 address of the newly advertised router.
IPv6 Address	IPv6 address of the newly advertised router.
Type	Router category: •unknown •internal •area-external •AS-external •both internal and area-external •both internal and AS-external •internal, area-external and AS-external
Name	Router name provided by the routing protocol, if available.

Table 8 New Links Table Fields

Field	Description
Link	Link from the source to destination router
Source Interface	The source of the link
Destination Interface	The destination of the link

Prefix List

The Prefix List report lists all of the currently or previously advertised prefixes and indicates reachability into and out of the network. We recommend that you run this report on a weekly basis to verify and assess the reachable networks inventory. It provides a way to quickly check that new networks have been added as planned or obsolete paths removed. Table 9 describes the fields in for an IGP domain. Separate tables are presented for IPv4 and IPv6 prefixes.

Table 9 Prefix List Fields for an IGP Domain

Field	Description
Prefix	Address of the prefix.
Type	Prefix category: •unknown •internal •area-external •AS-external •both internal and area-external •both internal and AS-external •internal, area-external and AS-external

Table 9 Prefix List Fields for an IGP Domain (cont'd)

Area or AS	Area or AS where the prefix is located.
Name	Router name provided by the routing protocol, if available.
IP Address	IP address of the router advertising the prefix.

If you choose OSI as the Administrative Domain, Table 10   describes the fields in the report.

Table 10 Prefix List Fields for an OSI Domain

Field	**Description**
Prefix Neighbor/ES Neighbor	Address of the prefix neighbor or ES neighbor.
Type	Prefix category: •unknown •internal •area-external •AS-external •both internal and area-external •both internal and AS-external •internal, area-external and AS-external
Area or AS	Area or AS where the prefix is located.
System ID	System identifier for the router.
Name	Router name provided by the routing protocol, if available.

Prefix Origination Changes

The Prefix Origination Changes report lists all of the prefixes that have changed their source router over a specified time period. This is a summary of any changes to the entry points of routes into a network. External routes are not visible. We recommend that you run this report every few days to identify potential problems, such as a router losing an interface or a flapping link. Table 11 describes the columns in the report for current routers and Table 12 describes the columns for the changes version of the report.

Table 11 Prefix Origination Changes Current Routers Fields for IGP Domain

Field	Description
Prefix	Prefix address of the network.
Name	Name of the router (if available in the routing protocol).
IP Address	IP address (router ID) in dotted decimal notation.

Table 12 Prefix Origination Router Changes Report Fields for IGP Domain

Field	Description
Name	Name of the router (if available in the routing protocol).
IP Address	IP address (router ID) in dotted decimal notation.
Advertised	Number of times the router advertised the prefix. *indicates whether the route was recently advertised.
Withdrawn	Number of times the router withdrew the prefix. *indicates whether the route was recently withdrawn.
Time	Time when the route was most recently announced or withdrawn.

If you choose OSI as the Administrative Domain, Table 13 describes the fields in the report for current routers and Table 14 for the changes version of the report.

Table 13 Prefix Origination Changes Current Routers Fields for OSI Domain

Field	Description
Prefix	Prefix neighbor and ES neighbor.
System ID	System identifier for the current router.
Name	Router name provided by the routing protocol, if available.

Table 14 Prefix Origination Router Fields for OSI

Field	Description
Prefix	Prefix neighbor and ES neighbor.
System ID	System identifier for the changed router.
Name	Router name provided by the routing protocol, if available.
Advertised	Number of times the router advertised the prefix. *indicates whether the route was recently advertised.
Withdrawn	Number of times the router withdrew the prefix. *indicates whether the route was recently withdrawn.
Time	Time when the route was most recently announced or withdrawn.

Prefix Origination from Multiple Sources

The Prefix Origination from Multiple Sources report lists all of the prefixes advertised by multiple routers. Run this report to determine if redundant links or hosts (such as redundant DNS servers and “Anycast” IP multicast rendezvous points) are operating normally. The absence of a redundant link or host from the list indicates that a redundant link or host is down, possibly resulting in reduced service levels or other problems within the network. This report can also detect configuration errors.

Run this report if the Prefix Origination Changes report identifies a problem or when you receive an alert that a redundant link has failed. Table 15 describes the fields in the report. Separate tables are presented for IPv4 and IPv6 prefixes.

Table 15 Prefix Origination from Multiple Sources Report Fields

Field	Description
Prefix	IP address of the network prefix.
Type	Prefix category: •unknown •internal •area-external •AS-external •both internal and area-external •both internal and AS-external •internal, area-external and AS-external
Area or AS	Area or AS where the prefix is located.
Name	Name of the router (if available in the routing protocol).
IP Address	IP address of the router.

If you selected OSI for the Administrative Domain, Table 16 describes the fields in the report.

Table 16 Prefix Origination from Multiple Sources Router Report Fields for OSI

Field	Description
Prefix	IP address of the network prefix.
Type	Prefix category: •unknown •internal •area-external •AS-external •both internal and area-external •both internal and AS-external •internal, area-external and AS-external
Area or AS	Area or AS where the prefix is located.
System ID	System identifier of the router.
Name	Name of the router (if available in the routing protocol).

Prefixes Withdrawn

The Prefixes Withdrawn report lists all of the prefixes that have been withdrawn from the network during the specified period. Run this report to identify clients that can no longer access the network or after scheduled maintenance to verify that no prefixes have been unintentionally dropped. Table 17 describes the fields in the report.

Table 17 Prefixes Withdrawn Report Fields

Field	Description
Prefix	IP address of the network prefix.
Name	Name of the router (if available in the routing protocol).
IP Address	IP Address of the router, if available. This field may not be displayed if you choose OSI as the administrative domain.
Time	Time the prefix was withdrawn.
Count	Number of times the prefix was withdrawn.

8 VPN Routing

This chapter describes how to configure the MPLS VPN protocol module and display VPN routing reports.

Chapter contents:

➤ The information in this chapter applies only to units that have licenses for both the BGP protocol and the VPN protocol.

About VPN Routing

The BGP/MPLS VPN, described in RFC 4364, is the most common form of service provider VPN. The edge routers of a VPN customer (CE routers) announce their routes to the service provider's edge routers (PE routers). The service provider then uses BGP to exchange the routes of the VPN among the PE routers associated with that VPN in a way that ensures that the routes from different VPNs are distinct and separate, even if the VPNs' address space overlaps. Because the CE routers do not peer with one another, there is no VPN overlay visible to the VPN routing algorithm.

Each route within a VPN has an MPLS label. When BGP advertises a VPN route, it also announces the MPLS label for the route. Before a VPN data packet is sent across the service provider backbone, the packet is encapsulated with the MPLS label that corresponds to the VPN route to the packet destination. The resulting packet is re-encapsulated so that it can be tunneled appropriately over the backbone to the destination PE router. In this way, the backbone routers do not need to know the details of the VPN route, thus protecting the privacy and security of the VPN.

In RFC 4364, a single mesh of tunnels is required between the PE routers. Although routes are stored in separate forwarding tables, the routes are still passed between PE routers using the same instance of BGP that exchanges Internet routes in the provider network. This means that problems with BGP routes can affect normal Internet connectivity.

To manage a large-scale deployment of VPNs in a robust manner, it is important to integrate protocol diagnostics with VPN service metrics such as reachability and participation. To accomplish this, the VPN protocol module includes the following features:

- A VPN topology overlay that lets you visualize the VPN topology on a per-customer basis.
- Summary and detailed views of reachability for advertised prefixes and status of the PE routers.
- Reports that signal problems in the VPN.
- Integrated VPN and BGP routing diagnostics to isolate reachability issues down to a single prefix, determine the routers participating in any VPN, and isolate and debug complex routing problems.

Understanding the Reachability and Participation Index

The system provides dynamic tracking for every prefix that is advertised from each customer on every VPN and tracks the PE routers that participate in each VPN.

Reachability and participation metrics normally remain stable in customer networks; however, these metrics change continually in service provider networks. Changes can be caused by periodic addition of new customer sites or VPN prefixes added to existing sites, addition of new PE routers to the network, or reallocation of prefixes between PE routers for load balancing or other reasons. Changes to the VPN overlay can also be introduced inadvertently due to BGP misconfiguration.

To provide a visual picture of VPN stability, a baseline is established with the number of prefixes seen at each PE router per VPN and the number of PE routers that participate in each VPN in a steady-state condition. When establishing the baseline, the system looks at routes that have been up for more than 80% of the time over the course of seven days. (Before a route reaches the seven-day mark, its baseline is determined on a day-to-day basis, such as 80% of 24 hours, 80% of 48 hours.) This number is assigned a stability index of 100. As the number of prefixes or routers change, the corresponding index number changes accordingly.

If you begin recording new routes for an existing database, the system continues to consider the baseline history of routes stored in that database. In the case of PEs, however, historical data is not considered; each time recording is started, the process of determining a baseline begins again.

The system displays the change in the reachability and participation index in graphical and text-based reports. You can use these reports to prioritize the allocation of technical resources to resolve customer VPN problems. See Viewing VPN Routing Reports on page 364 for information about the available reports.

Creating Customer and RT Associations

Because the MPLS labels of the VPN are carried in the MP-BGP protocol messages, the VPN protocol module does not require any additional configuration other than establishing peering with the PE routers when you install the appliance in the network. See the "Configuration and Management" chapter in the *HP Route Analytics Management Software Administrator Guide* for information about how to enable VPN when establishing peering.

However, if you want to display summary reports on a per-customer or per-PE-router basis, you can associate a customer identifier with one or more RTs by entering the association information manually in the Routing Reports window.

To configure multiple customer/RT associations at once, it is easier to copy and paste using a .csv file.

Specify one of the following RT formats to match the conventions used in your network:

- `RT:<AS number>:<VRF ID>`

  This format consists of the letters RT, followed by the 16-bit AS number, followed by the unique 32-bit VPN routing and forwarding (VRF) ID. Separate each of the three elements with a colon; for example, `RT:65522:101`.

- `RT:<IPv4address>:<ID>`

  This format consists of the letters RT, followed by the 32-bit IPv4 address of the appliance announcing the routes, followed by a unique 16-bit ID number. Separate each of the three elements with a colon; for example, `RT:192.168.0.1:5`.

Before creating customer and RT associations, you must enable queries on the VPN client. In a deployment with Route Recorders and a centralized Modeling Engine, consider the following recommendations when you enable queries:

- For network-wide information, enable queries on the centralized Modeling Engine.
- For information local to a recorder's area or protocol, enable queries on the Route Recorder.

To enable queries, perform the following steps:

1. From the web interface, choose **Administration**.
2. Choose **Queries**.
3. Select **XML-RPC Query Server** and **Enable remote access**.
4. Enter a password and confirm it. The password can be from one to eight alphanumeric characters in length, is case-sensitive, and must not contain nulls, blanks or underscores.
5. Click **Update**.

To set up customer/RT associations manually, perform the following steps:

1 From the client application, choose **Administration > VPN > VPN Customers** to open the VPN Customer Configuration window (Figure 2).

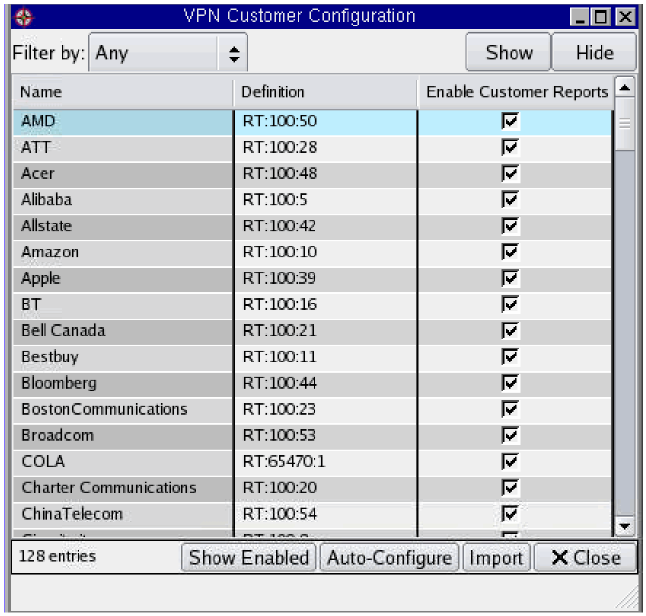

Figure 1 VPN Customer Configuration Table

2 Click **Import** and enter customer data using the format

```
cust_id,rt
```

where `cust_id` is a customer identifier, and `rt` is the route target you want to associate with that customer.

To associate multiple route targets for a given customer, separate the route targets with white space. For example:

```
customer1, RT:65522:101 RT:192.168.0.1:1
```

To set up more than one association, enter each association on a separate line.

3 After entering all of the required customer/RT associations, click **Import**.

The new associations appear in the table. If a customer is associated with more RTs than can be displayed in the Definition column, placing the mouse pointer over the definition opens a pop-up window containing the complete list.

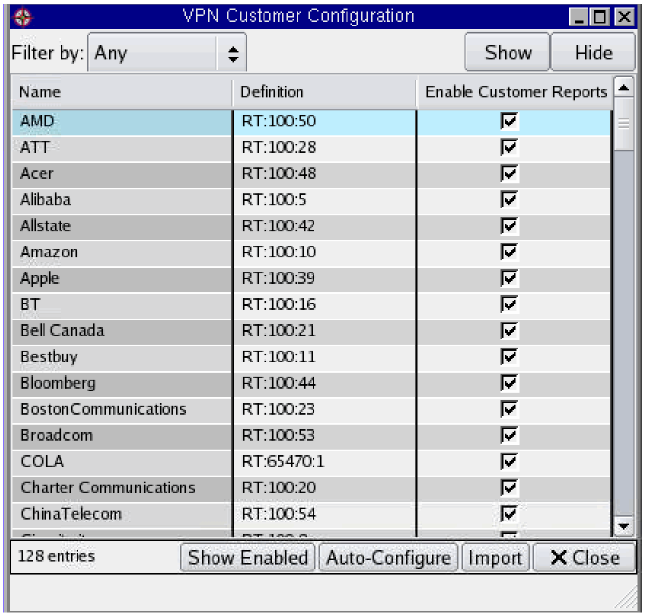

Figure 2 VPN Customer Configuration Table

If the appliance does not have a VPN Customer Reports license, the Enable Customer Reports column is hidden.

To set up customer/RT associations automatically, perform the following steps:

1. From the client application, choose **Administration > VPN > VPN Customers** to open the VPN Customer Configuration window.

2. Click **AutoConfigure** (you may need to scroll to see this button).

   The AutoConfigure feature uses heuristics to identify a collection of route targets that could plausibly be associated with a single customer. The route targets are selected from the received VPN routes, and the identified customers are named with a “Customer_RT:” prefix.

After the Autoconfig process finishes, the edit option allows customer names to be edited, and the set of associated route targets to be adjusted as needed.

To set up associations by copy and pasting a file, perform the following steps:

1 Create a text file with customer data in the following format:

```
cust_id,rt
```

where cust_id is a customer identifier, and rt is the route target you want to associate with that customer.

To enter multiple route targets for a given customer, separate the route targets with white space. Place each customer/RT association on a separate line. For example:

```
customer1, RT:65522:101 RT:192.168.0.1:1
customer2, RT:65511:102
customer3, RT:192.168.245.3:22
```

2 Copy the contents of the file.

3 Choose **Administration > VPN > VPN Customers** to open the VPN Customer Configuration window.

4 Paste the text into the text box near the bottom of the window.

5 Click **Submit**.

The new associations are added to the table in the lower portion of the pane.

Generating VPN Customer Traffic Reports

This feature is available only if the appliance has a VPN Customer Reports license.

The system can generate a set of reports through an XML-RPC interface, which you can make available through a web reports portal for your customers. By default, the ingress PE and egress PE of a particular VPN customer flow are specified as the source and destination customer site for computing site-related statistics.

You can change this definition from the web portal and communicate it to the Flow Analyzer and the Modeling Engine through the XML-RPC interface.

You can generate reports for up to 100 customers using this mechanism. The reports are available on the Flow Analyzer or Modeling Engine.

For more information about the XML RPC queries used for the customer reports, see the "VPN Customer Traffic API" chapter in the *HP Route Analytics Management Software Administrator Guide*.

To make the XML reports available, perform the following steps:

1 From the client application, choose **Administration > VPN > VPN Customers** to open the VPN Customer_RT Mapping Configuration table.

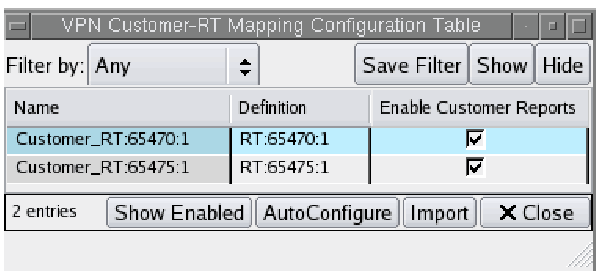

Figure 3 VPN Customer_RT Mapping Configuration Table

2 Perform any of these tasks:

- Select check boxes for the reports that you want to generate.
- Click **Show Enabled** to list only the customers for which reports are enabled.
- Click **Show All** to list all entries in the table.

3 Click **Close**.

You can also add names to selected routers to make the customer reports more meaningful. The names are included in the XML file along with the IP addresses.

To add router location names for XML customer reports, perform the following steps:

1 From the client application, choose **Admin > VPN > Router Location Names**.

2 Choose a router from the list and click in the Location Name column.

3 Enter the location name.

4 Click **Save**.

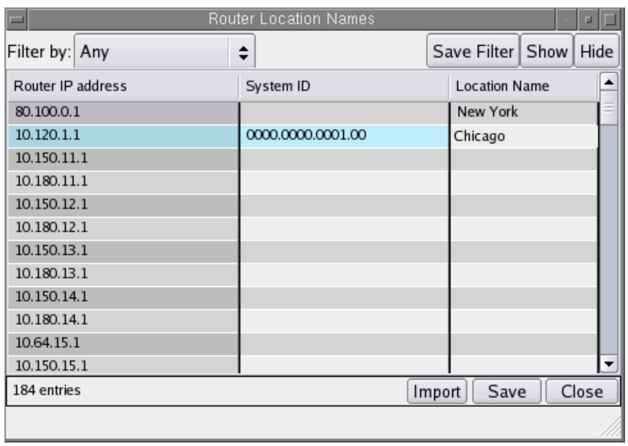

Figure 4 Router Location Names

Viewing VPN Routing Reports

In a deployment with multiple Route Recorders and a centralized Modeling Engine, we recommend that you obtain reports from the Modeling Engine. Reports from the Modeling Engine return network-wide information and are faster to obtain than reports obtained directly from Route Recorders.

To access the VPN report pages, perform the following steps:

1 From the client application, choose **Reports > Routing Reports**.

The Routing Reports window opens (Figure 5).

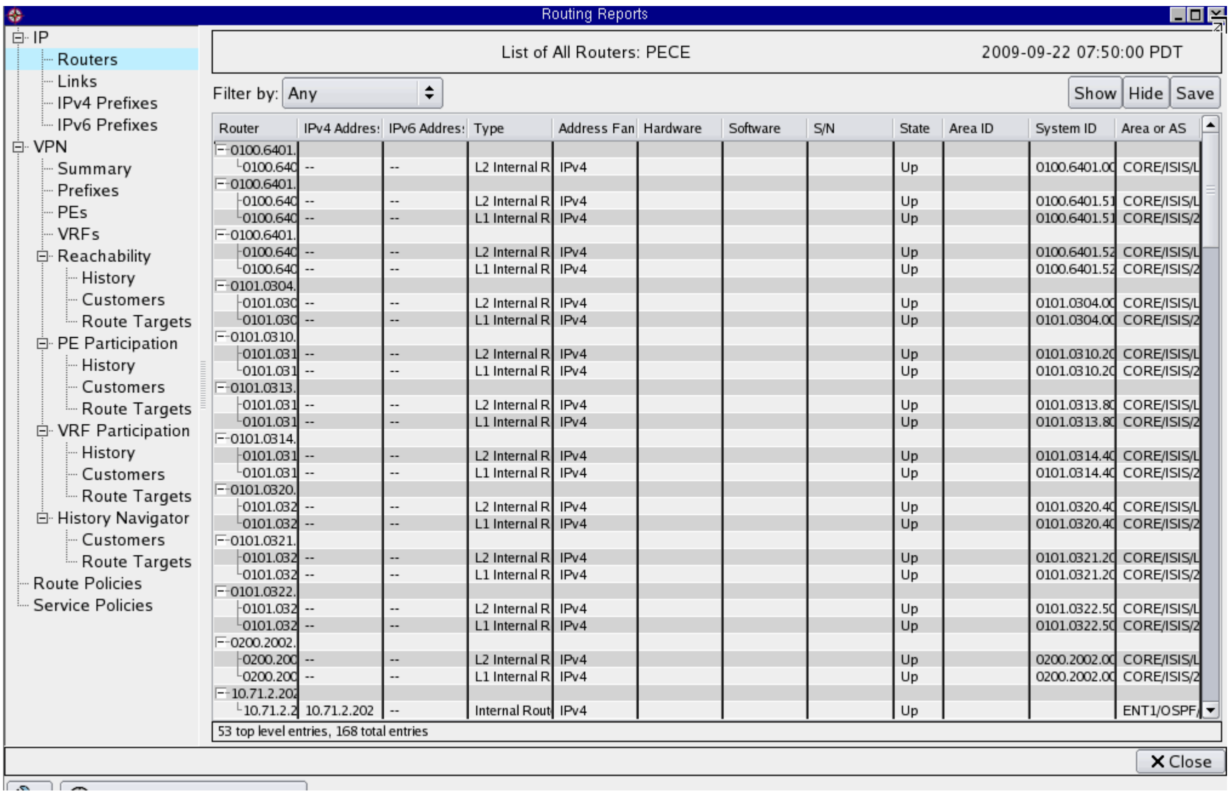

Figure 5 Routing Reports Window

2 Choose individual reports from the tree menu on the left side of the window. The remaining sections in this chapter describe the reports.

3 The Filter By menu lets you filter the report to include only the information you want to see. See Using Filters on page 221. You can resort data by clicking any column heading in the report. Click again to change the sort order (descending/ascending).

Modifying the Report Table

You can modify the report table by right-clicking in the column header and selecting from the following options:

- **Sort**—Sort on the selected column. Click the column header to change the sort order.
- **Group or UnGroup**—Combine elements of the same values in the selected column into a single group, or remove a grouping.
- **Collapse All/Expand All**—Hide all the elements in the groups created in the column, or show all the elements.
- **Hide**—Hide the selected column.
- **Show**—Show a column that was previously hidden.

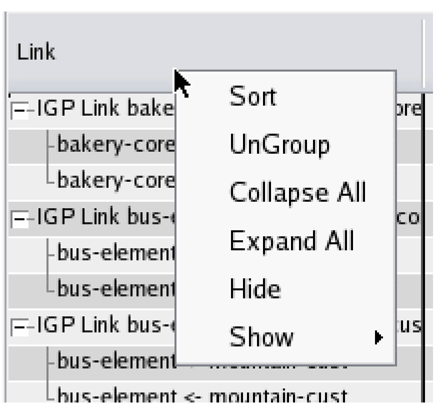

Figure 6 Modifying Report Tables

VPN Summary Report

Choose **Reports > Routing Reports > VPN > Summary** in the client application to open the VPN Summary report (Figure 7). The report displays the following information:

- At the top of the pane, the report contains pie charts that indicate reachability by RT and by customer and participation by RT and by customer, respectively.
- Below each pie chart, the report lists customers and RTs that experienced the greatest deviation from the baseline index in the same categories (reachability and participation).

See Chapter 12, "Alerts," for information about VPN alerts.

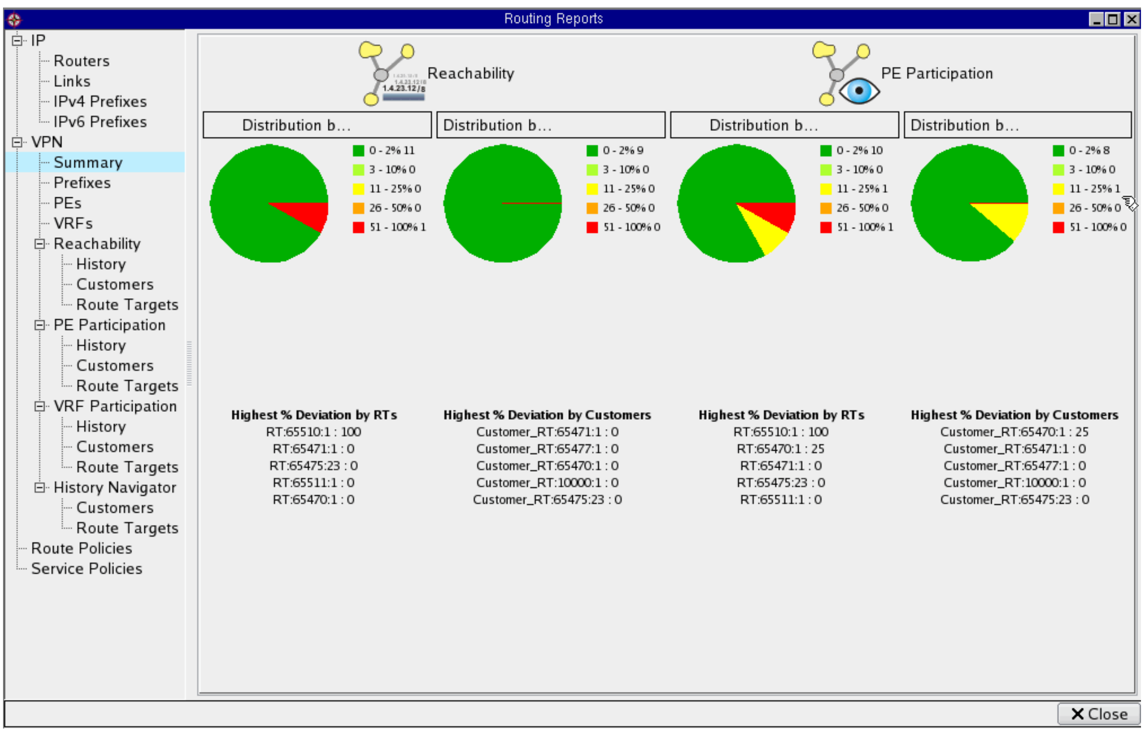

Figure 7 VPN Summary Report

VPN Prefixes Report

Choose **Reports > Routing Reports > VPN > Prefixes** in the client application to open the VPN Prefixes report (Figure 8). The report lists all prefixes advertised by VPNs in the network with the associated router, attributes, state, and area or AS.

Each entry is identified in the first column of the table, as follows:

- The first part of the entry (for example, 65453:1) is the route distinguisher (RD). It can take one of two formats, each of which has two numbers separated by a colon.
  - — The format shown in the example 65453:1 consists of a 16-bit number (typically a BGP AS number) followed by the unique 32-bit VPN routing and forwarding (VRF) ID.
  - — The second format is a 32-bit number represented in the format of an IPv4 address followed by the unique 16-bit VPMN routing and forwarding (VRF) ID. An example RD in the second format is 198.51.100.1:1001.
- The second part is the IPv4 address prefix (for example, 10.85.1.0/24).
- The second line is the MPLS label (for example, 131065).

These three items (RD, prefix, label) are carried together in BGP to convey a VPN-IPv4 route, and the Routing Report > VPN > Prefix table displays the routes as received in BGP.

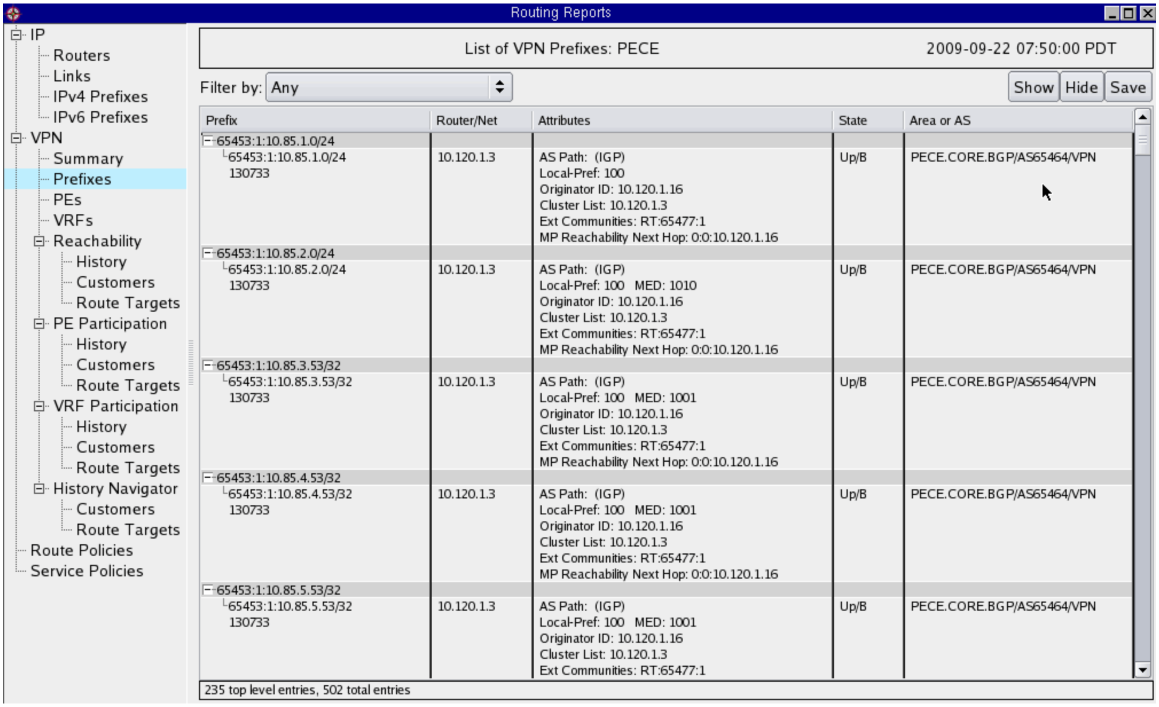

Figure 8 VPN Prefixes Reports

VPN PEs Report

Choose **Reports > Routing Reports > VPN > PEs** in the client application to open the VPN Prefixes report (Figure 9). The report lists all PEs advertised by VPNs in the network with the information about the associated VRFs and VPN routers.

To display information involving VRFs, you must first configure the Collector. See the instructions in the "Configuration" chapter of the *HP Route Analytics Management Software Administrator Guide*.

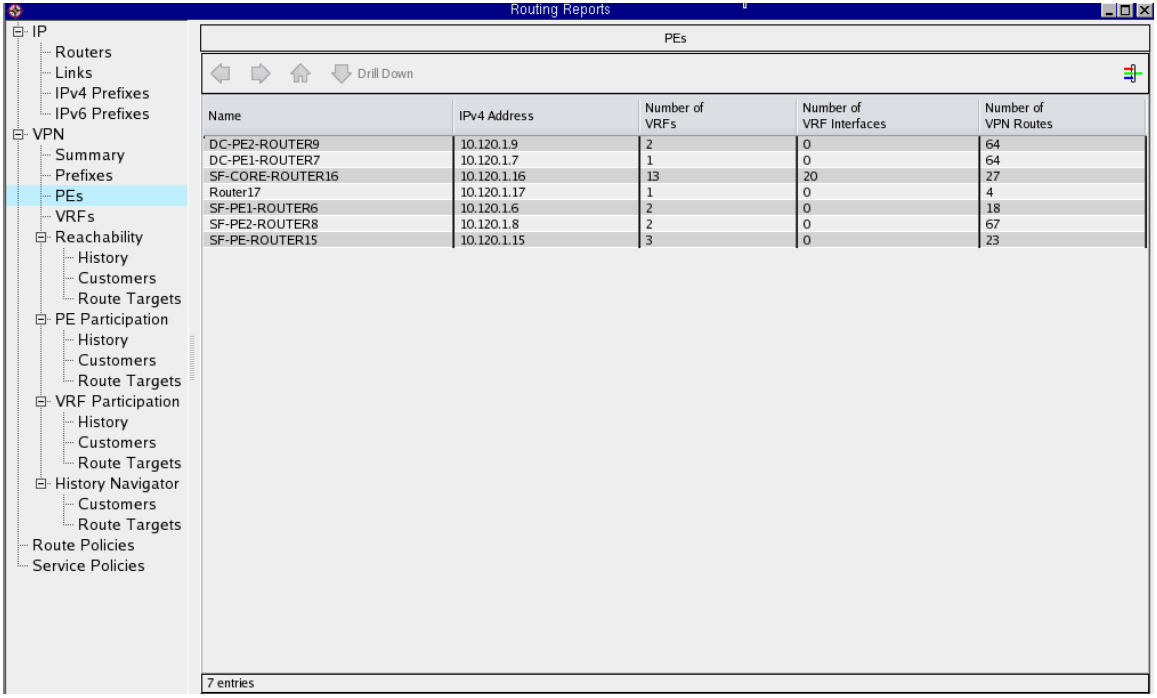

Figure 9 VPN PEs Reports

VPN VRFs Report

Choose **Reports > Routing Reports > VPN > VRFs** in the client application to open the VPN Prefixes report (Figure 10). The report lists all the VRFs for the VPN, along with information about each VRF entry.

To display information involving VRFs, you must first configure the Collector. See the instructions in the “Configuration” chapter of the *HP Route Analytics Management Software Administrator Guide*.

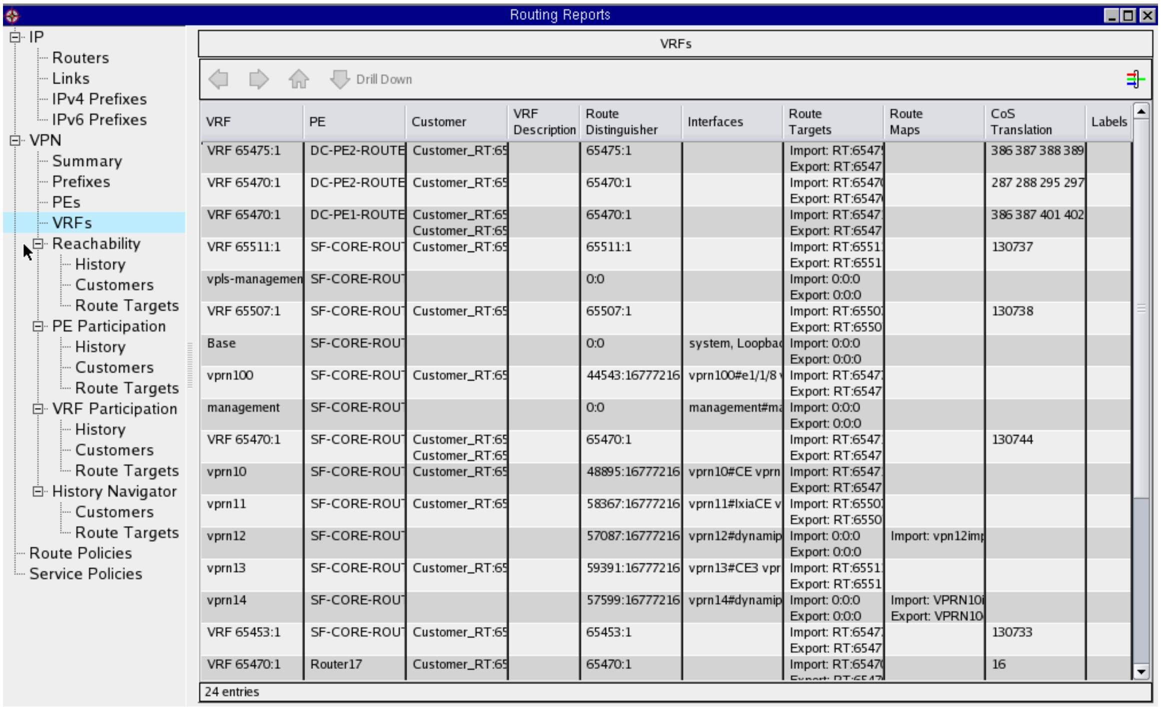

Figure 10 VPN VRFs Reports

VPN Reachability History Report

Choose **Reports > Routing Reports > VPN > Reachability > History** in the client application to open the VPN Reachability History report (Figure 11).

The graphs in this report show deviation from the baseline by RT and by customer. The x axis is time and the y axis is percentage of deviation. A limited set of History Navigator functions is available. For a description of the functions, see Chapter 4, "The History Navigator"

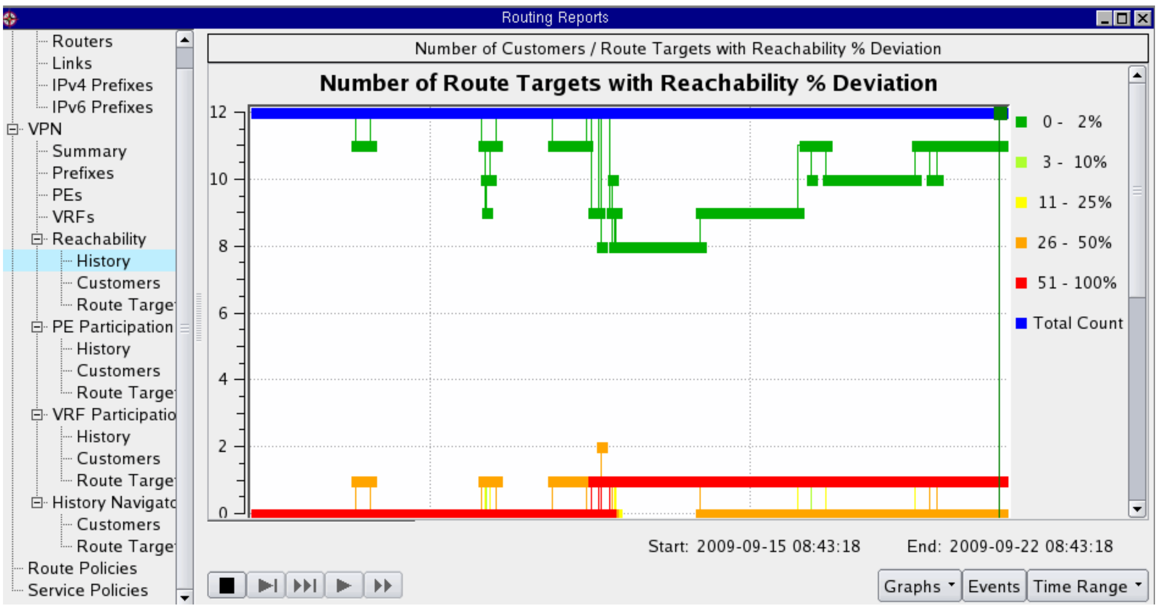

Figure 11 VPN Reachability History Report

VPN Reachability Customers Report

Choose **Reports > Routing Reports > VPN > Reachability > Customers** in the client application to open the VPN Reachability Customers report (Figure 12).

The report includes the customer identifier or RT identifier and the numbers of active PE participants, active routes, baseline routes, down routes, new routes, and the deviation from the baseline.

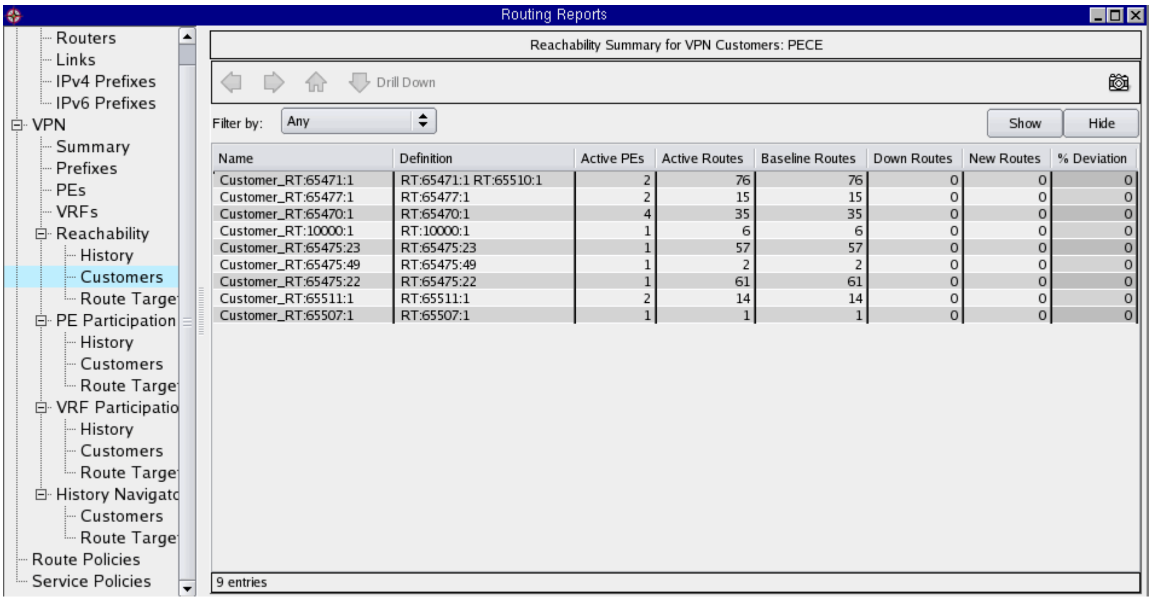

Figure 12 VPN Reachability Customers Report

VPN Reachability Route Targets Report

Choose **Reports > Routing Reports > VPN > Reachability > Route Targets** in the client application to open the VPN Reachability Route Targets report (Figure 13).

The report includes the RT identifier and the numbers of active PE participants, active routes, baseline routes, down routes, new routes, and the deviation from the baseline.

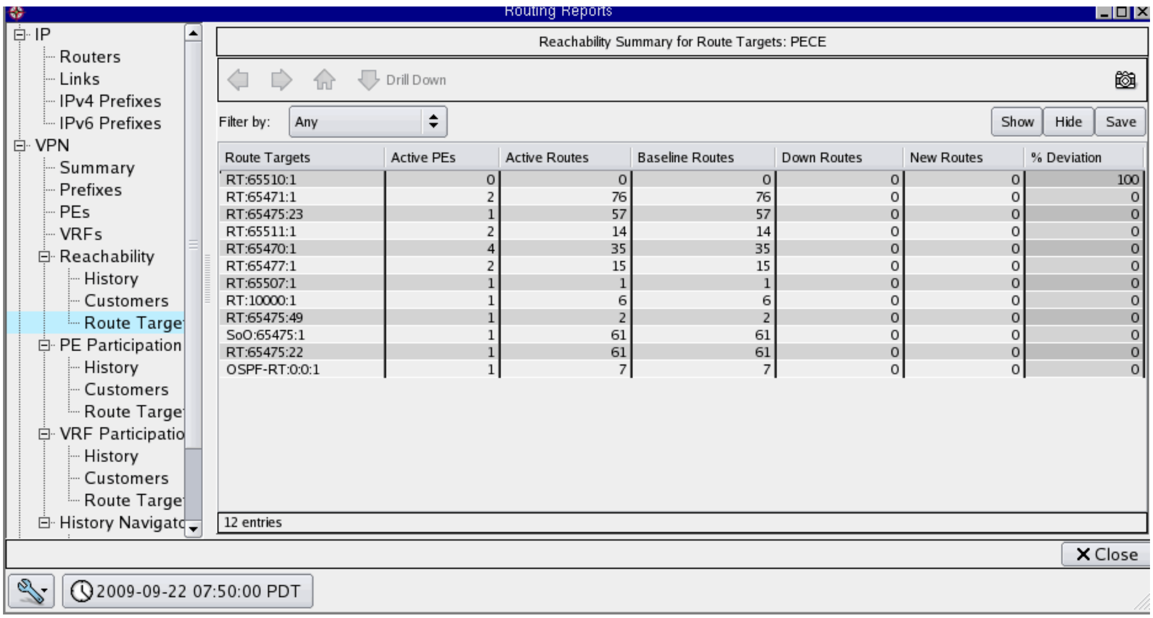

Figure 13 VPN Reachability Route Targets Report

VPN PE Participation History Report

Choose **Reports > Routing Reports > VPN > PE Participation > History** in the client application to open the VPN PE Participation History report (Figure 14).

The graphs in this report show deviation from the baseline by RT and by PE. The x axis is time and the y axis is percentage of deviation. A limited set of History Navigator functions is available. For a description of the functions, see Chapter 4, "The History Navigator"

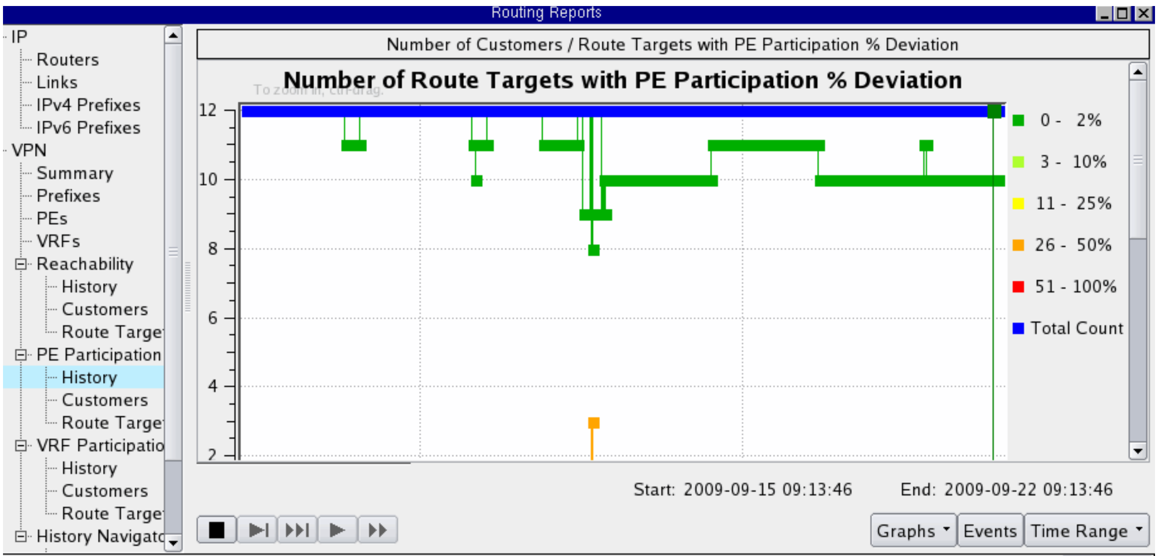

Figure 14 VPN PE Participation History Report

VPN PE Participation Customers Report

Choose **Reports > Routing Reports > VPN > PE Participation > Customer** in the client application to open the VPN PE Participation Customers report (Figure 15).

The report includes the customer identifier or RT identifier and the numbers of active PE participants, baseline PE participants, down PEs, new PEs, and the deviation from the baseline

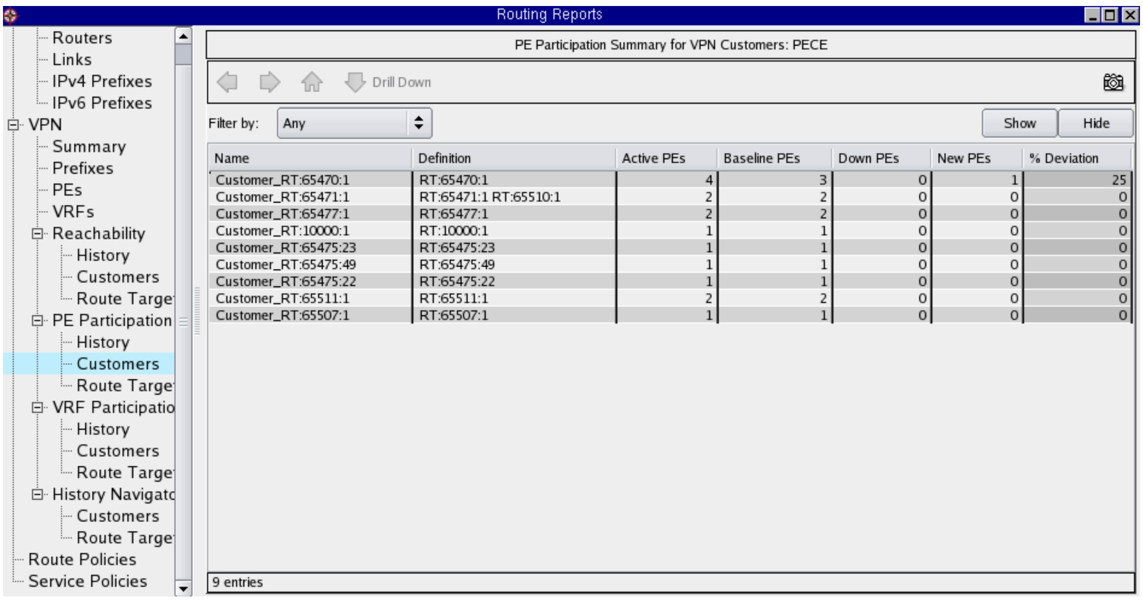

Figure 15 VPN PE Participation Customers Report

.

VPN PE Participation Route Targets Report

Choose **Reports > Routing Reports > VPN > PE Participation > Route Targets** in the client application to open the VPN PE Participation Route Targets report (Figure 16).

The report includes the RT identifier and the numbers of active PE participants, active PEs, baseline down PEs, new PEs, and the deviation from the baseline.

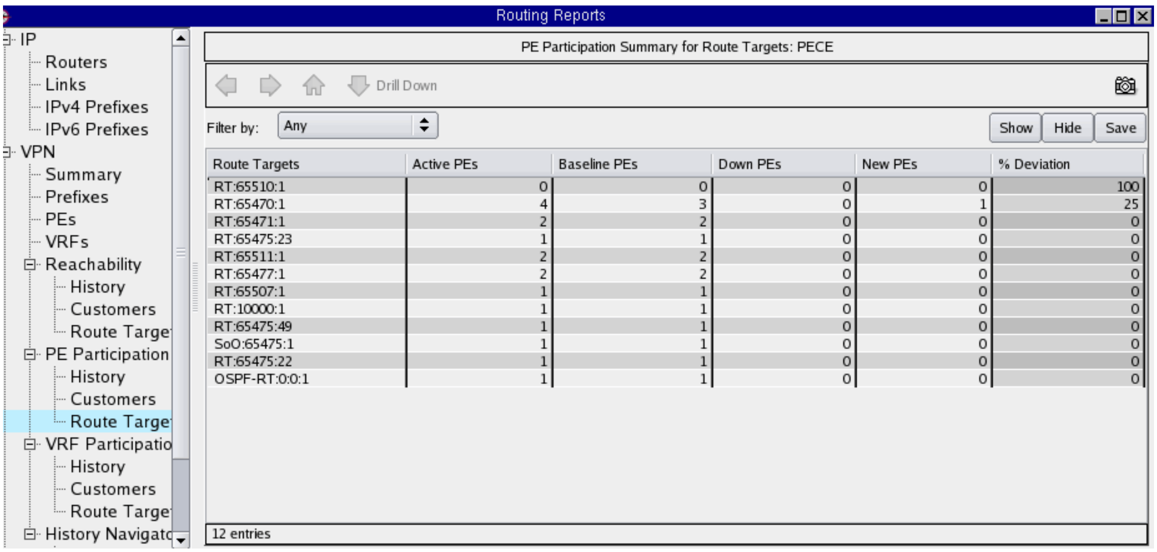

Figure 16 VPN PD Participation Route Targets Report

VPN VRF Participation History Report

Choose **Reports > Routing Reports > VPN > VRF Participation > History** in the client application to open the VPN VRF Participation History report (Figure 17).

The graphs in this report show deviation from the baseline. The x axis is time and the y axis is percentage of deviation. A limited set of History Navigator functions is available. For a description of the functions, see Chapter 4, "The History Navigator"

To display information involving VRFs, you must first configure the Collector. See the instructions in the "Configuration" chapter of the *HP Route Analytics Management Software Administrator Guide*.

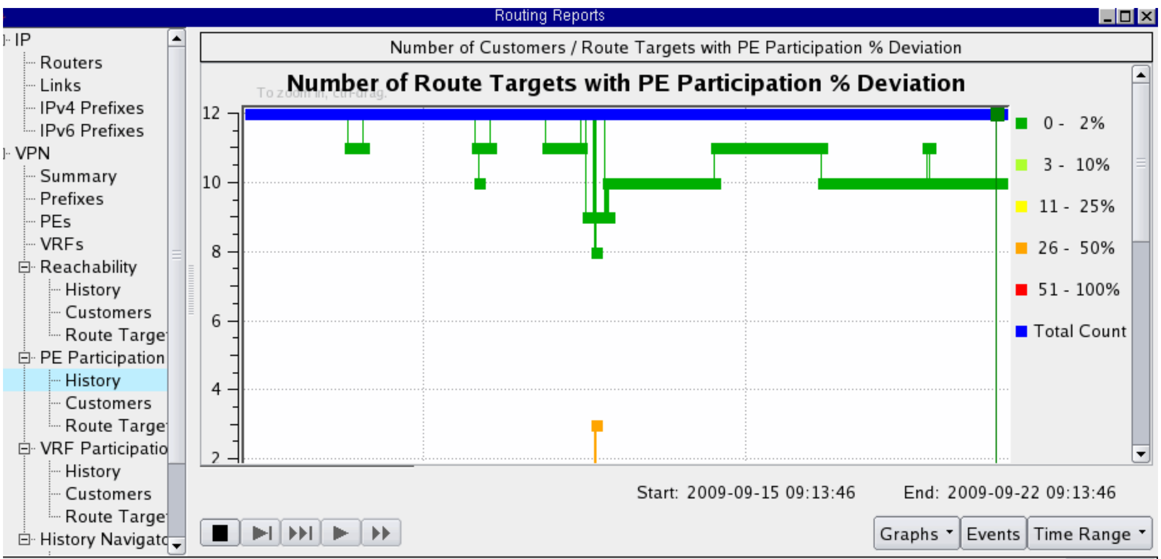

Figure 17 VPN VRF Participation History Report

VPN VRF Participation Customers Report

Choose **Reports > Routing Reports > VPN > VRF Participation > Customer** in the client application to open the VPN VRF Participation Customers report (Figure 18).

The report includes the customer identifier or RT identifier and the numbers of active VRF participants, baseline VRF participants, down VRFs, new VRFs, and the deviation from the baseline.

To display information involving VRFs, you must first configure the Collector. See the instructions in the “Configuration” chapter of the *HP Route Analytics Management Software Administrator Guide*.

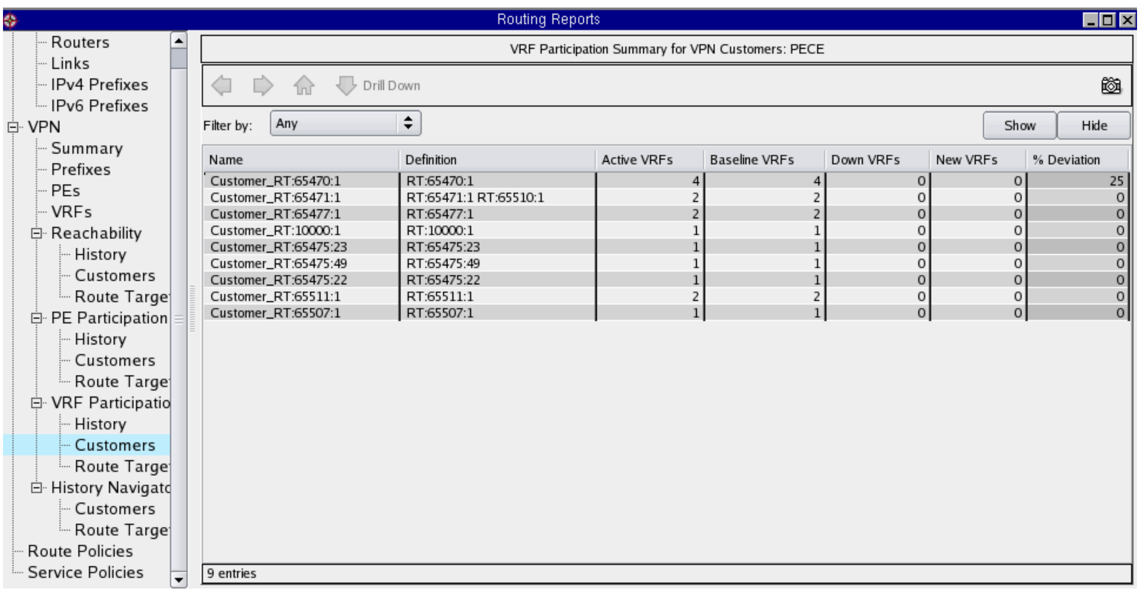

Figure 18 VPN VRF Participation Customers Report

VPN VRF Participation Route Targets Report

Choose **Reports > Routing Reports > VPN > VRF Participation > Route Targets** in the client application to open the VPN VRF Participation Route Targets report (Figure 19).

This report includes the RT identifier and the numbers of active PE participants, active PEs, baseline down PEs, new PEs, and the deviation from the baseline.

To display information involving VRFs, you must first configure the Collector. See the instructions in the “Configuration” chapter of the *HP Route Analytics Management Software Administrator Guide*.

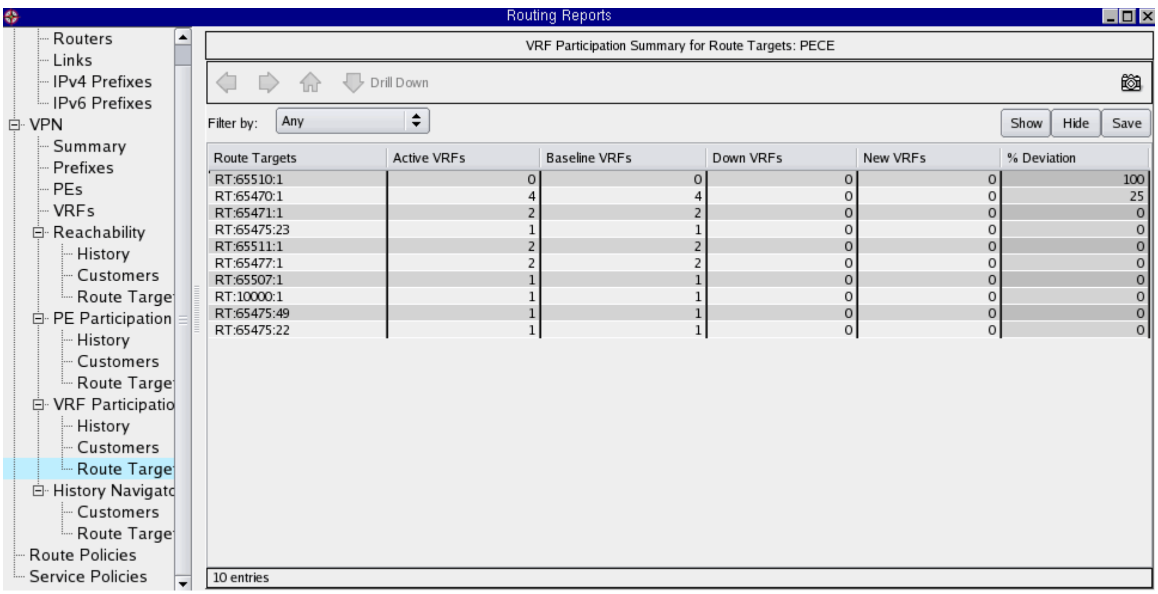

Figure 19 VPN VRF Participation Route Targets Report

VPN History Navigator Customers Report

Choose **Reports > Routing Reports > VPN > History Navigator > Customers** in the client application to open the VPN Customers History Navigator report(Figure 20).

Select an RT from the list to display a graph of associated events. A limited set of History Navigator functions is available. For a description of these functions, see Chapter 4, "The History Navigator"

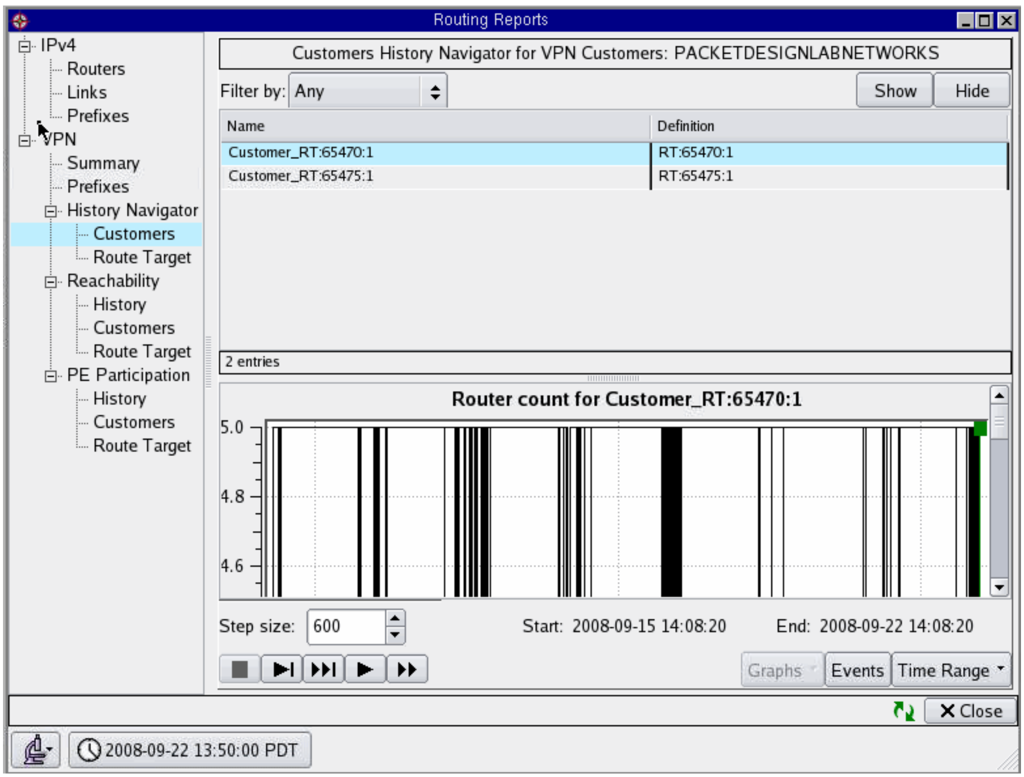

Figure 20 VPN Customers History

VPN History Navigator Route Targets Report

Choose **Reports > Routing Reports > VPN > History Navigator > Route Targets** in the client application to open the VPN History Navigator Route Targets report (Figure 21).

Select an RT from the list to display a graph of associated events. A limited set of History Navigator functions is available. For a description of these functions, see Chapter 4, "The History Navigator"

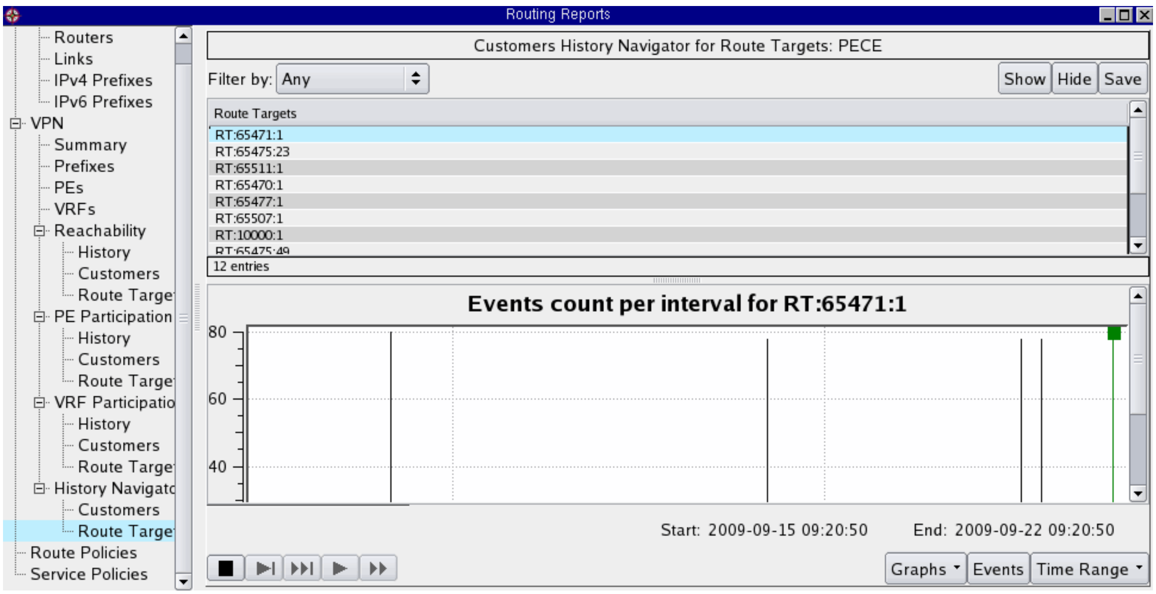

Figure 21 VPN History Navigator Route Targets Report

Route Policies and Service Policies Reports

The Route Policies report lists the policies that control route modifications, for example, when routes are moved or copied from one protocol to another or routes are imported for VRF. Each policy specifies how route elements are moved or copied.

For each interface and interface direction (input or output), the Service Policies report determines the forwarding class and priority and may rewrite some fields in the packets that pass through the interfaces.

For route policies and service policies, you can click the "i" icon in the upper right corner of the report window to open the Inspector panel. This panel presents policy information in the format that is appropriate to the router vendor. The inspector information includes information about the policy itself and all of the objects that the policy uses. Tabs in the Inspector panel allow you to drill down to policy and object details.

For example, the next figure shows the Inspector panel for the "routemap22" policy for Router17 (highlighted in the figure). The Inspector panel has two tabs. The routemap22 tab provides a summary of the policy in vendor-appropriate format, showing that the policy involves prefix list matching. The Prefix Lists tab provides additional information about the prefix list matching criteria.

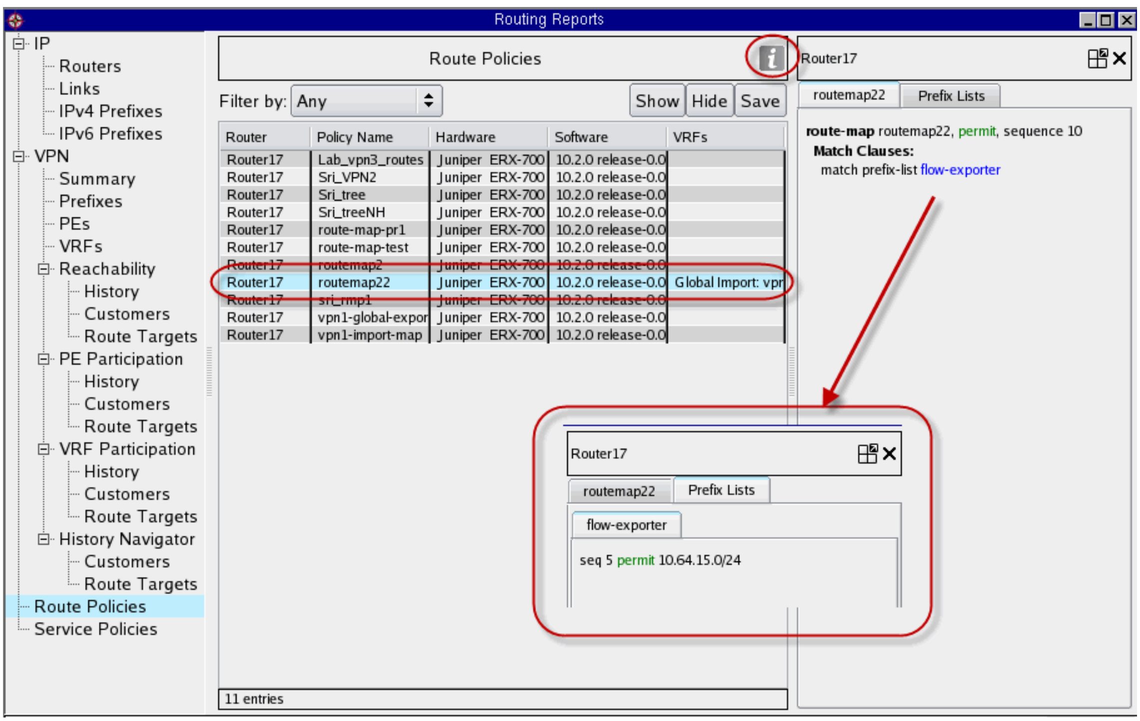

Figure 22 Inspector Panel for Route Policies Report

Obtaining Detailed Information

From the Reachability and PE Participation reports, you can display detailed information or focus your attention on specific RTs or customers.

To display the details of a summary report, perform the following steps:

1 Choose **Reports > Routing Reports**.

2 Choose **Reachability** > **Customers** or **PE Participation > Customers** open the Reachability Summary report or Participation Summary report.

3 Right-click an entry in the report to open and choose one of the detail options.

9 Traffic Flows and Reports

This chapter applies to RAMS Traffic only.

This chapter describes the traffic collection and analysis process and the traffic reports that are available to monitor behavior and anticipate usage trends in the network.

Chapter contents:

Understanding Traffic Flows

RAMS Traffic adds traffic flow analysis to route analytics to provide an integrated, real-time view of network-wide routing and traffic behavior, allowing you to view complex IP networks as integrated systems and maximize IT efficiency and productivity.

With RAMS Traffic, you can understand the dynamic impact of routing changes or failures on traffic flow, determine root cause of problems, optimize network operation, and effectively analyze and plan for network change and growth. Network-wide, end-to-end visibility is provided, without requiring broad deployment of probes or the overhead associated with polling-based techniques. You can interact with an “as-running” model of the network, in which actual traffic flow information is dynamically overlaid on the routing topology map.

Because RAMS Traffic is able to determine the actual routed path through the network for every flow, you can quickly focus attention on suspect devices or links and pinpoint the cause of poorly performing applications. You can determine whether failures are due to new traffic loads on the network or to traffic that was rerouted due to router failure or congestion, and set response priorities or generate alerts accordingly.

Routing is done according to prefixes (each destination is a prefix). To keep the magnitude of the traffic data workable, all data to the same prefix is aggregated in a single aggregate data flow. All sources that pass through a specific source router (exporter) to the same destination prefix are aggregated together. This reduces the number of flows for which a projection across the network is required for reporting purposes.

By simulating network changes, such as failing network elements, modifying prefixes, or adjusting metrics, you can understand the impact of proposed changes before implementing them. Traffic paths can be engineered to avoid performance problems or SLA violations during peak traffic loads. You can analyze and manipulate a network-wide traffic matrix that shows traffic volumes between every source/destination pair in the network.

VPN RAMS Traffic

VPN RAMS Traffic adds comprehensive VPN monitoring to RAMS Traffic. Customer traffic flows are mapped across their individual VPN topology, providing traffic visibility throughout the provider network, from the PE router where customer traffic enters the network, through the routers and links in the MPLS core that forward that traffic, to the PE router connected to the customer's destination site.

Service providers can view individual customer VPN topologies, visualize the complete end-to-end path between any two sites, and analyze a customer's site-to-site service by prefix reachability, traffic utilization, and CoS breakdown on all links connecting the sites.

Because it maintains a complete history of traffic and routing events, VPN RAMS Traffic can rewind the network's traffic and routing state to a previous point in time when an intermittent problem may have been occurring, and compare current traffic loads against historical baselines.

An aggregate report is available to aggregate VPN and IPv4 non-VPN traffic. The standard aggregation gives a feel for the overall traffic flow in the network. To address specific traffic issues, use the Flow report options within traffic reports to see the details on the individual flows.

For VPN RAMS Traffic requires Netflow data to be collected from either of the following collection points:

- At the ingress to the P router from the PE router. This option is preferable, however, it requires Netflow version 9 with MPLS label information.
- At the ingress to the PE router from the CE router. To map this traffic onto the provider network using this option, it is necessary to collect additional information from the PE routers. To do so, you must configure the Collector. See the "Configuring the Route Recorder for the Collector" section in the *HP Route Analytics Management Software Administrator Guide*.

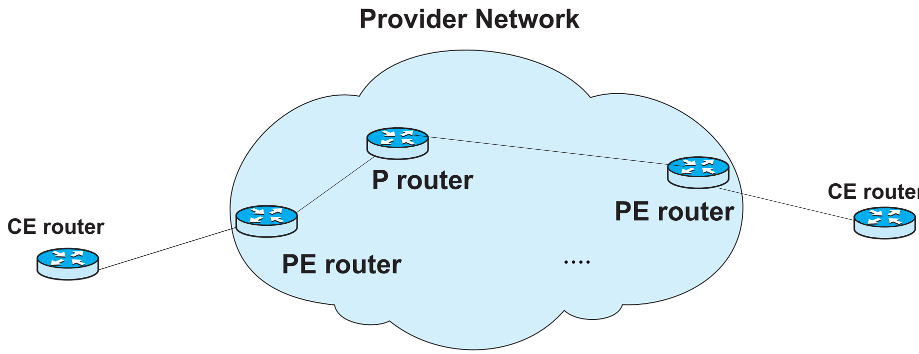

Figure 1 Collection Points for VPN Traffic

How Traffic Flow Collection Works

Traffic flow collection, analysis, and reporting follows this process:

1 The Flow Collector collects Netflow data that is exported from routers at important traffic sources, such as data centers, Internet gateways, and WAN links. In Netflow, each flow has the following characteristics in common:

— Source IP address

— Destination IP address

— Source port for UDP or TCP, 0 for other protocols

— Destination port for UDP or TCP, type and code for ICMP, or 0 for other protocols

— IP protocol

— Ingress interface

— IP Type of Service

2 The Netflow data is projected using the routing model across the network. The information is provided to the Flow Analyzer, which performs the analysis and generates reports.

3 The Flow Analyzer aggregates reports as needed and issues alerts based on network-wide routing and traffic data, including application and CoS details.

See the “Configuration” chapter in the *HP Route Analytics Management Software Administrator Guide* for information on configuring the Flow Collector, including Netflow settings.

Understanding Traffic Reports

Traffic Reports allow you to monitor behavior and anticipate trends of network elements, such as routers, links, customers, and ASs. You can study network element usage, optimize network resource allocation, produce realistic network planning strategies, and troubleshoot and identify network problems.

RAMS Traffic constructs traffic reports from real time and historic data generated every few minutes (default is 5 minutes). Incremental data is used to produce hourly, daily, weekly, and monthly data and to determine averages, minimums, maximums, and 95 percentiles. The resulting statistical and trend reports show bandwidth utilization, packet and protocol distribution, and error and outage information. Many of the reports have drill-down options. In some cases you can drill down and see the raw flows that make up the aggregate flow.

The report data is collected and correlated from the Flow Collector and the Route Recorder and processed by the Flow Analyzer. Some data is concealed by prefix-level aggregation so that only the Flow Collector (not the Flow Analyzer) can generate reports for this concealed data.

Examples of this data include:

- Top source addresses
- Top destination addresses
- Traffic distribution per protocol
- Traffic distribution per Flow Collector

The Flow Collector generates data for the 5-minute reports. The Flow Analyzer obtains the data from multiple Flow Collectors and aggregates the data into hourly, daily, weekly, and monthly data sets. For example, you can compare utilization values of the network elements listed in the table to find over- and under-utilized links and make changes to help route traffic more evenly among the links.

For calculation of link utilization, the measurement of traffic exported from an upstream router will stop when the traffic reaches another router that is exporting traffic on the interface into which the traffic exported by the upstream router is arriving.

For ECMP, if the first exporter is at the branch point or upstream from the branch point, and the second exporter is on one of the legs of the ECMP path, it is not possible to know the accurate distribution of the traffic after the branch point until the traffic reaches the second exporter. Therefore, the reported utilization for the link going out from the second exporter may not correspond to the reported utilization for the link coming into it.

Traffic data is based on the time that is displayed in the status bar at the bottom of the routing topology map window in the client application. For current data, there is typically a 30-minute delay, although data acquisition time varies based on the complexity of your network. See Main Window Status Bar on page 50 for instructions on changing the date, time, and modes.

You can view and analyze traffic distribution over the network for the following protocol families:

- IPv4
- VPN (if your appliance is licensed for VPN as well as traffic)

Aggregate reports are available only when your network uses more than one protocol. All aggregate reports include data from all Flow Collectors known to the Flow Analyzer.

Right-click options are available with most data fields within the workspace. These options
can duplicate or extend drill-down functionality, as described in Understanding Report Types
on page 406.

Traffic data is based on the time appearing in the status bar, which you can change to review any previous data. For current data, there is typically a 20- to 30-minute delay, although data acquisition times varies based on the complexity of your network.

Accessing Traffic Reports

The following conditions are required to access traffic reports in the client application:

- You must have a license for RAMS Traffic SPI.
- The topology that you open must include a Traffic Reports database (which has a corresponding routing database).
- The Flow Collector database must be open.

To access Traffic Reports, perform the following steps:

1 Choose **Reports > Traffic Reports** to open the Traffic Reports window (Figure 2).

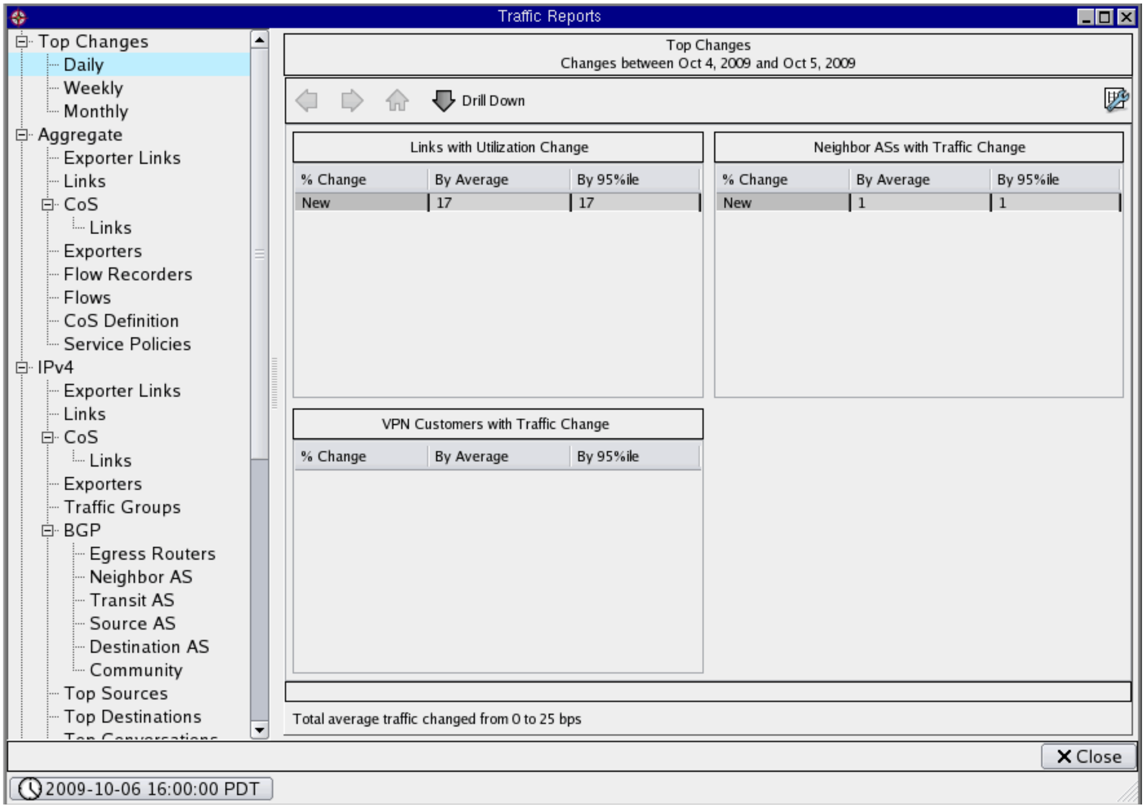

Figure 2 Traffic Reports Window

2 Choose individual reports from the tree menu on the left side of the window.

Navigation Tree

The navigation tree in the left pane is grouped by protocol families in your network. If your network supports multiple protocols, you will see an Aggregate category and a category for IPv4 and VPN traffic. If your network supports only one protocol there will not be an Aggregate category.

You can expand or compress any category preceded by a plus or minus symbol. Within each category are reports that contain diagnostic information.

Traffic reports for protocol families appear only if the protocol is configured and only if that specific license has been purchased.

Traffic Reports Buttons

The buttons listed in Table 1   are available in Traffic Reports, depending on the table and conditions.

Table 1   Traffic Report Buttons

	**Go back one drill-down**	During a drill-down, goes back one drill-down level.
	**Go forward one drill-down**	During a drill-down, goes forward one-drill down level.
	**Go back to top; undo all drill-downs**	Goes to the highest point in the drill-down hierarchy, and “unrolls” the drill-down view from the window.
	**Configure**	Allows you to add columns or modify the data within a column using conditions and parameters. For Top Changes reports, allows you to set up ranges.
	**Color Links By Traffic Volume**	Each link is colored according to according to the aggregate traffic volume of all flows traversing the link. The legend indicates the correspondence between colors and levels of traffic volume. Note: This icon is present for any report or drill-down based on links.

Table 1 Traffic Report Buttons (cont'd)

	**Advanced Filter**	Allows you to define advanced filters. See Understanding Report Types on page 406 and Using Filters on page 221 for more information.
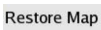	**Restore Map**	Displays after the Show Map feature is evoked. This is done by right-clicking a row in a VPN traffic report that is broken down by customer.
	**Snapshot**	Opens a new window containing a snapshot of the current window.
	**Table**	Changes the workspace data to columns and rows. (This button is not always available.) This icon is present for Top Sources, Top Destinations, and Protocol reports.
	**Pie Chart**	Changes the workspace data to a pie chart. This icon is present for Top Sources, Top Destinations, and Protocol reports.
	**Bar Graph**	Changes the workspace data to a bar graph. This icon is present for Top Sources, Top Destinations, and Protocol reports.
	**Drill-down**	If available, this button allows to see finer detail within a set of data.
	**Close**	Closes the Traffic Reports workspace.

Working with Traffic Reports

Most reports allow you to edit, move, or add columns appearing in the report.

Columns Settings

Most traffic reports allow you to manipulate columns or customize the data in existing columns by clicking the configuration icon , which opens the Traffic Reports Columns window (Figure 3).

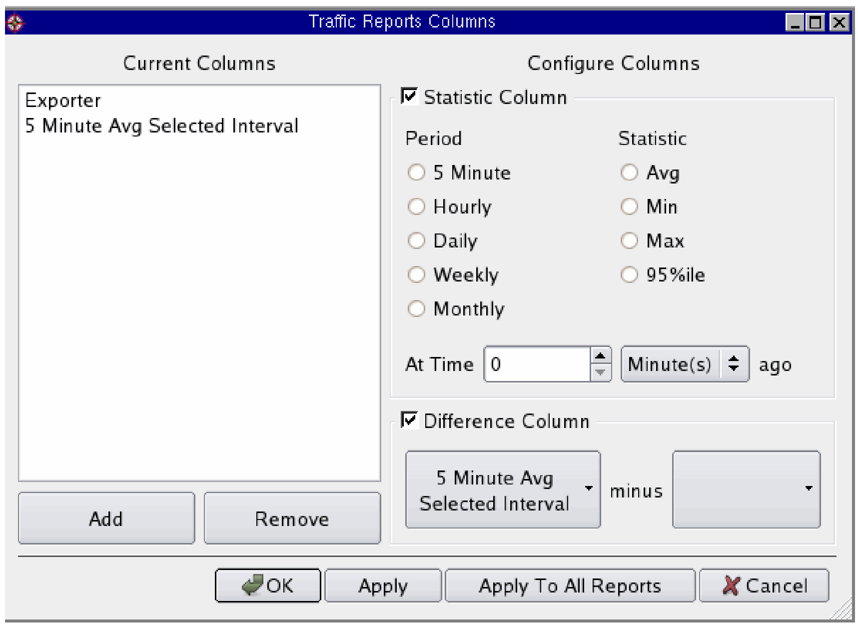

Figure 3 Traffic Reports Columns

The available Configure columns vary. You can modify the Statistic (including the default 5 Minute Average column) or Difference Column. Do not delete or change the default 5 Minute Avg Selected Interval column if you want to retain drill-down capability.

If you click **Apply to All Reports**, all statistic and difference columns in all other reports are replaced by those appearing the current Traffic Reports Columns window. This process cannot be undone.

Statistics Column Area

The Statistics Column area allows you to create additional columns defining new data parameters. Periods are based on the time appearing in the status bar. Within the available periods, you can select from the following options:

- **5 Minute** – Last 5 minutes of available data based on the standard 5-minute increments of a clock. For example, if you choose this option and the time is set between 1:30 and 1:34, data is retrieved for the period 1:25 to 1:30.
- **Hourly** – Last 1 hour of available data from the time appearing in the status bar based on the standard twelve 1-hour increments of a clock. For example, if you choose this option and the time is set for 1:30, data is retrieved for the period 12:00 to 1:00.
- **Daily** – Last full day (12:00:01 am to 11:59:59 pm) of available data from the time appearing in the status bar. For example, if you choose this option and it is Tuesday, data is retrieved for the previous Monday.
- **Weekly** – Last full week (Monday to Sunday) of available data from the time appearing in the status bar. For example, if you choose this option and it is Friday, data is retrieved for the period Monday to Sunday of the previous week.
- **Monthly** – Last calendar month of available data from the time appearing in the status bar.
- **At Time** – Specific time intervals for obtaining report data.

Within the available statistics options are:

- **Average** – Sum of all selected values divided by the total number of elements.
- **Minimum** – Lower bounds within a group of values.
- **Maximum** – Upper bounds within a group of data.
- **95 percentile** – Value equal to or greater than 95 percent of the values.

Selecting reports on links allows you to create additional columns related to the Average statistic:

- **Bit Rate (bps)** – Bit rate level that a link is reaching.
- **Utilization (%)** – Percentage of use that a link is achieving if capacity is configured.

Difference Column Area

The Difference Column area allows you to create a new column that generates data based on subtracting one statistics column from another. The applicable statistic columns appear in the drop-down lists within the Difference Column area.

Advanced Filtering

Simple filters let you choose a single operator from a list and specify one or more parameters to be matched or excluded.

Advanced filters let you choose two or more different operators from a list and specify their corresponding parameters to be matched or excluded. From a Filter workspace, select the drop-down list in the Filter by field and choose **Advanced**. The Composing Advanced Filter window opens.

Figure 4 Composing Advanced Filter

Advanced filtering may not be enabled for all reports.

See Using Filters on page 221 for more information.

Minute Data Granularity

You can reset the clock to retrieve data based on minute granularity. This can be useful information if you need to determine the environment at a specific moment. However, depending on the amount of data involved, retrieving granular data can be a resource-intensive process and require additional time in retrieving and displaying the data. This type of information automatically creates a 5 Min Avg Current Time column in the report.

Drill-Down Capabilities

A measure is a metric or a performance indicator used to determine how well routing and traffic are operating. The aggregated data you need to examine are called measures — numeric values that are measurable and additive. You can also examine measures under certain conditions called “dimensions.” Dimensions are often time-based, such as by day, week, or month. Dimensions can have a hierarchy that allows you to perform drill down functions on the data.

Traffic Reports contains measures (data) organized by different dimensions to provide faster retrieval and drill-down capability. Drill-down allows you breakdown data into finer detail. Many report measurements offer the ability to take summary information and drill-down through a hierarchy to show the detailed data used to develop the summary data. Drill-down capability is not available for all measures, either because finer detail is not stored or you have already reached the finest detail. The values you see when you drill-down can vary, depending on the drill-down selections you have made. In some cases you can drill down to see the raw flows that make up the aggregate flow.

➤ You can select multiple rows to drill-down by selecting Ctrl + Shift keys, and then selecting the rows you want to drill-down.

When the **Drill-Down** button appears, you can gain additional detail for any of the following:

- **Community—**Displays how much traffic each community receives in bits.
- **CoS**—Displays the CoS for the traffic flow.
- **CoS**—Displays a user-defined CoS name.

- **Customer—**Displays the user-defined customer ID.
- **Destination AS**—Identifies the name or AS number of the final AS.
- **Egress PE**—Identifies the router from which flow exits the current AS.
- **Egress Router**—Identifies the router from which flow exits the current AS in a VPN topology.
- **Exit Router**—Identifies the name or IP address of the exit router used to reach a peer. This information is obtained from the router name repository prioritized by router names, DNS names, IP address, and system IDs. For more information regarding the router name repository, see Assigning Router Names on page 104.
- **Exit Router-Next Hop**—Identifies the IP address of next hop.
- **Exporter**—Identifies the name or IP address of the peer from which flow information was received.
- **Exporter Links**—Identifies the interface of the peer from which flow information was received.
- **Interfaces**—Identifies the name or IP address of peer from which flow information on the interface it exited from.
- **Flow Collector**—Identifies each Flow Collector on the network.
- **Flows**—Displays the details of individual flows within the aggregate and allows you to see the exact source and distribution.
- **History**—Displays detailed routing history for the network using graphs and trending. For more information about this feature, see Drill-Down History Option on page 401.
- **Ingress PE**—Identifies the name or IP address of peer from which flow information was received.
- **IPv4 Flows**—Displays details of each IPv4 traffic flow
- **Link**—Displays how much traffic each link is carrying.
- **Neighbor AS**—Identifies the name or AS number of directly connected AS to which traffic is being sent.
- **Source AS**—Identifies the name or AS number of the AS from which the flows originated.

- **Traffic Group**—Displays a user-defined group name. For more information about creating traffic groups, see Creating Groups Using the Menu on page 122.
- **Transit AS**—Displays the name or AS number of each transit AS used in delivering data to a destination.
- **VPN Flows**—Displays the details of each VPN traffic flow.

Viewing the summary data helps show the general utilization of your network elements, while dividing the data into more specific segments can help you analyze and troubleshoot your network elements more accurately.

Depending on the amount of data involved, drilling down can be a resource-intensive process and may

The following example displays the drill-down from the BGP Flows report to show the links for specified flows.

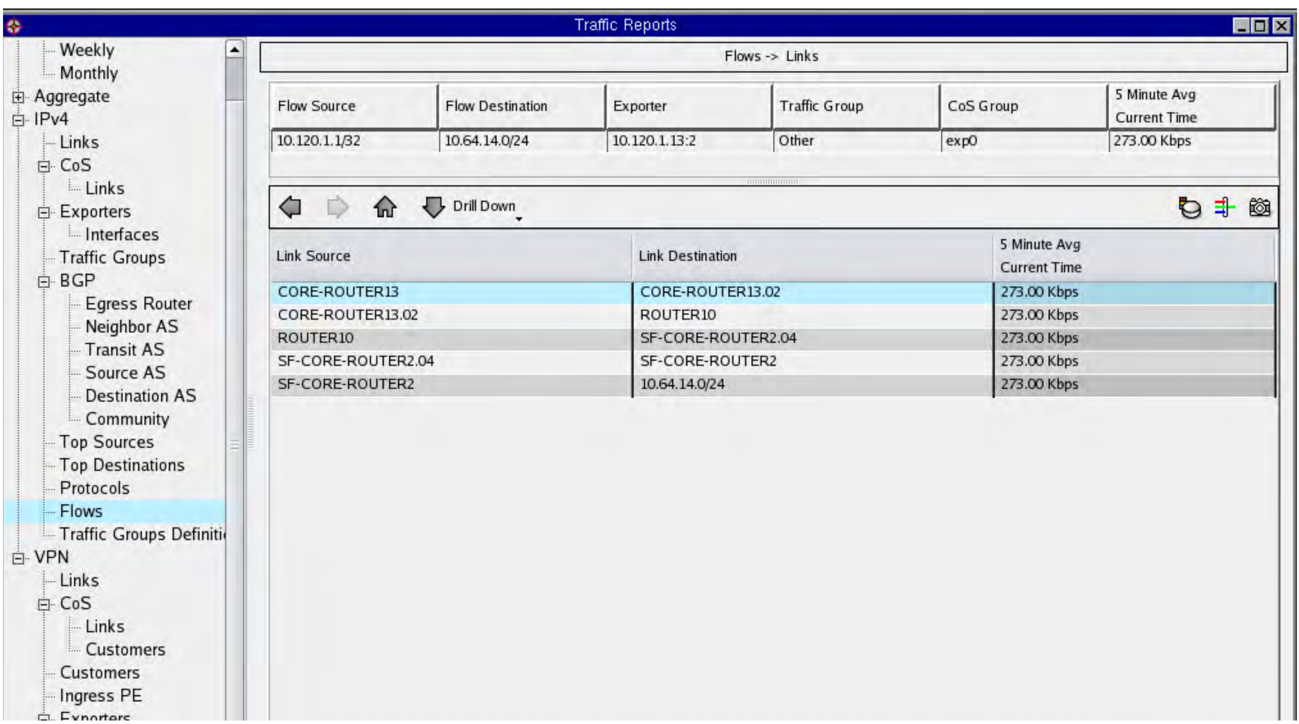

Figure 5 Example of Drill-Down to Identify BGP Flows on a Link

Drill-Down History Option

The History option delivers ways for you to obtain graphical representation of past traffic trends, enabling you to anticipate and plan for future traffic needs. To access this option, select a report type from the left navigation pane, and then select the element you wish to view traffic data for. Next, select **Drill-Down>History.**

The History window opens.

Several features within this option provide you with more tools to examine traffic data:

The **Trend** drop-down menu opens the Trending window (shown below). This feature provides a way for you to project and estimate traffic statistics at the selected future date and time. Either a Linear or Exponential model can be used for this estimation. Your organization can use this information to allocate your future traffic needs.

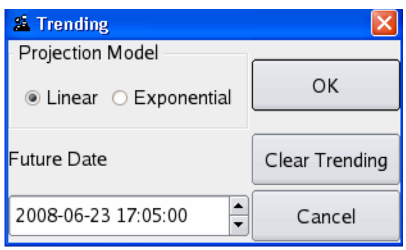

Figure 6 Trending

- The **Graphs** drop-down menu displays, by default a “best fit” (in other words, the best traffic data available) graphical representation of past traffic data for the selected network elements. The **Graphs** drop-down menu provides the option to show/hide other graphs, including the default graph. There is one graph for each of the reports that displays in the top level report.

The default time range for the graphs is seven days.

- The **Time Range** drop-down menu provides a selection of time periods (including a custom time frame) to confine the graph to.

Columns in the History window enable you to select the reports to be shown simultaneously. For more information about Columns in Traffic Reports, see Columns Settings on page 395.

Setting Interface Capacities

To compute percentage utilization across the links, RAMS Traffic must be able to determine the capacity of link interfaces. Some protocols (IS-IS with TE enabled or EIGRP) allow the system to determine the capacity directly, and capacity can also be determined through static data collection.

To accommodate other situations and also permit additional capacity adjustments, RAMS Traffic allows you to configure interface capacities manually. This is useful for situations in which the capacity is not automatically discovered or there are conflicts in the capacities that are reported for different protocols or physical links.

The following use cases may arise:

- Multiple physical links between the same two routers—The system automatically sums the capacities for the different physical links.
- Multiple protocols over a single physical link. If the individual protocols are reporting capacities, then the following rules are automatically applied to determine default capacity:
  - If the same capacity is reported for different multiple protocols, that capacity is used by default.
  - If different capacities are reported for the multiple protocols (or only some capacities are reported), then the largest reported capacity is used by default.
  - If the capacities are not known, then it is necessary to assign capacities manually.

In all of these case, you can adjust capacity on a per-interface basis by directly setting the capacity. If you configure a capacity for a protocol instance on a physical link, that configured capacity supersedes all of the discovered capacities for that physical link.

For OSPF and IS-IS interfaces, the system provides the option of configuring a reference bandwidth, which is used with the advertised link metric to determine the capacity according to the formula (capacity = reference bandwidth / metric).

To set interface capacities, perform the following steps:

1 Enter Planning mode or Analysis mode. You cannot set interface capacities in Monitoring mode. If you open the Set Interface Capacities window and then change to Monitoring mode, a warning message indicates that the window will be closed and prompts for confirmation that you want to continue.

2 Choose **Administration > Traffic > Set Interface Capacities** to open the configuration options window.

   The window presents information in a hierarchy, with the top level entry identifying a router-router link and each sub-entry representing an individual protocol. A parent entry has multiple sub-entries if there are multiple protocols running over a single physical link or multiple physical links running in parallel between routers.

   On the parent (group) rows, the effective capacity for the multi-protocol link is shown in the Discovered Capacity or Configured Capacity column, depending upon whether it is the sum of discovered capacities only or it includes at least one configured capacity.

   After you modify capacities at the sub-entry level, the system will consolidate the results to generate a capacity for the overall link (top level entry).

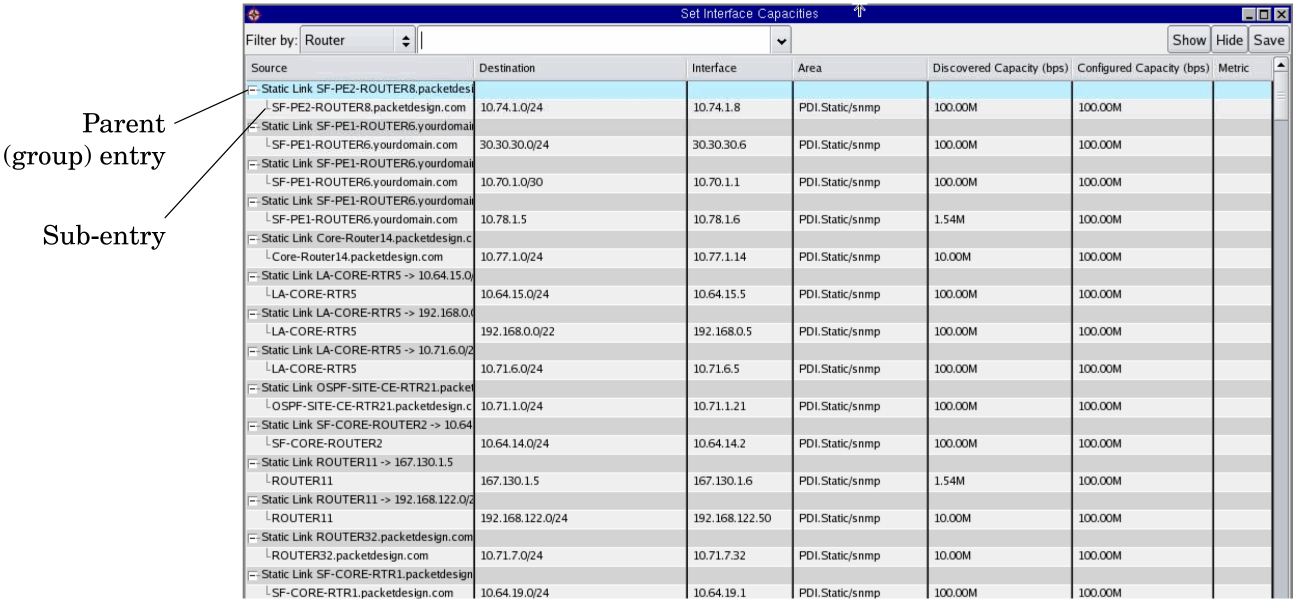

Source	Destination	Interface	Area	Discovered Capacity (bps)	Configured Capacity (bps)	Metric
Static Link SF-PE2-ROUTER8.packetdesi						
SF-PE2-ROUTER8.packetdesign.com	10.74.1.0/24	10.74.1.8	PDI.Static/snmp	100.00M	100.00M	
Static Link SF-PE1-ROUTER6.yourdomai						
SF-PE1-ROUTER6.yourdomain.com	30.30.30.0/24	30.30.30.6	PDI.Static/snmp	100.00M	100.00M	
Static Link SF-PE1-ROUTER6.yourdomai						
SF-PE1-ROUTER6.yourdomain.com	10.70.1.0/30	10.70.1.1	PDI.Static/snmp	100.00M	100.00M	
Static Link SF-PE1-ROUTER6.yourdomai						
SF-PE1-ROUTER6.yourdomain.com	10.78.1.5	10.78.1.6	PDI.Static/snmp	1.54M	100.00M	
Static Link Core-Router14.packetdesign.c						
Core-Router14.packetdesign.com	10.77.1.0/24	10.77.1.14	PDI.Static/snmp	10.00M	100.00M	
Static Link LA-CORE-RTR5 -> 10.64.15.0/						
LA-CORE-RTR5	10.64.15.0/24	10.64.15.5	PDI.Static/snmp	100.00M	100.00M	
Static Link LA-CORE-RTR5 -> 192.168.0.(						
LA-CORE-RTR5	192.168.0.0/22	192.168.0.5	PDI.Static/snmp	100.00M	100.00M	
Static Link LA-CORE-RTR5 -> 10.71.6.0/2						
LA-CORE-RTR5	10.71.6.0/24	10.71.6.5	PDI.Static/snmp	100.00M	100.00M	
Static Link OSPF-SITE-CE-RTR21.packet						
OSPF-SITE-CE-RTR21.packetdesign.c	10.71.1.0/24	10.71.1.21	PDI.Static/snmp	100.00M	100.00M	
Static Link SF-CORE-ROUTER2 -> 10.64						
SF-CORE-ROUTER2	10.64.14.0/24	10.64.14.2	PDI.Static/snmp	100.00M	100.00M	
Static Link ROUTER11 -> 167.130.1.5						
ROUTER11	167.130.1.5	167.130.1.6	PDI.Static/snmp	1.54M	100.00M	
Static Link ROUTER11 -> 192.168.122.0/2						
ROUTER11	192.168.122.0/24	192.168.122.50	PDI.Static/snmp	10.00M	100.00M	
Static Link ROUTER32.packetdesign.com						
ROUTER32.packetdesign.com	10.71.7.0/24	10.71.7.32	PDI.Static/snmp	10.00M	100.00M	
Static Link SF-CORE-RTR1.packetdesign						
SF-CORE-RTR1.packetdesign.com	10.64.19.0/24	10.64.19.1	PDI.Static/snmp	100.00M	100.00M	

Figure 7 Set Interface Capacities Window

3 Choose a desired filtering option, if needed, to display the entries of interest.

4 Choose one of the following options:

- To change the capacity for a single sub-entry, click the **Configured Capacity** column and enter the value. The default units are bits per second (bps). You must enter K, M, or G to specify Kbps, Mbps, or Gbps.
- To modify the capacities for the full table, choose **Change > All**.
- To configure the capacity for multiple rows, select the sub-entry rows and choose **Change > Selected**.

5 If you choose one of the **Change** options, the Change Interface Capacity window opens.

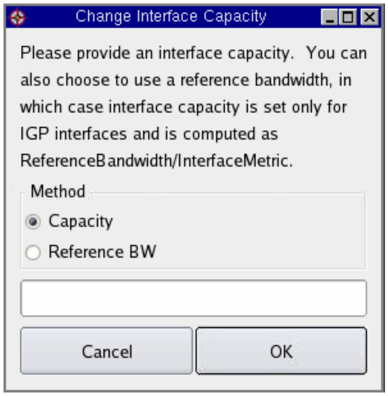

Figure 8 Setting Interface Capacity

6 Choose one of the following options:

- Select the **Capacity** button and enter the desired capacity. If you choose this option, no additional metric is calculated. The default units are bits per second (bps). You must enter K, M, or G to specify Kbps, Mbps, or Gbps.
- (IS-IS and OSPF only) Choose the **Reference BW** button and enter the bandwidth. IS-IS and OSPF have a mechanism that allows the system to compute the capacity for each link according to the link's metric (capacity = reference bandwidth/metric).

7 Click **OK**.

The window closes, and the Set Interface Capacities window displays the changes.

8 Choose one of the following actions:

- Click **Apply** to apply the values in the Configured Capacity column to the working topology model. (**Apply** is automatically performed when you click **OK** in the Change Interface Capacity window, as described in step 7.) If you enter a value in the Configured Capacity column and then close the window without saving or applying, the change is lost.
- Click **Save** to perform an **Apply** operation and then save new values to the database. You must choose this option to keep the applied changes after closing the topology.

- Click **Cancel** to remove any unsaved changes made to the table since it was opened, whether or not those changes have been applied. If you make a change, click **Apply** or **Save**, and then close the table and reopen it, the change will be visible and Cancel will not remove it.
- Click **Retrieve** to reload all the configured capacities that were saved in the database.

Understanding Report Types

There are four categories of report types available within Traffic Reports:

- Flow Analyzer reports contain history, period (5 minute, hourly, weekly, monthly), statistical data (average, minimum, maximum, and 95th-percentile). These are reports that include the **Columns** button.
- Flow Collector reports contain data from the last 5 minutes. These are the IPv4 Top Source Address, IPv4 Top Destination Address, and IPv4 Protocols reports.
- Reports with no history that are created on-the-fly. These are the Aggregate Flows, BGP Community, IPv4 Flows, and VPN Flows reports.
- General information reports that contain configuration information and do not change based on any time change. This type of report is represented by the Traffic Groups Definition report.

The navigation tree contains only those protocols identified on your network, have been loaded and opened, and for which you have purchased a license. This section provides a list of all default measures and dimensions within each report. This list is presented alphabetically and does not match the flow within the navigation tree. Your configuration may differ. Following the descriptions of each default report are suggestions for other useful measures.

The following diagram is intended to help clarify the content that appears within Traffic Reports by defining terms. All reports are based on your administrative domain being the central location.

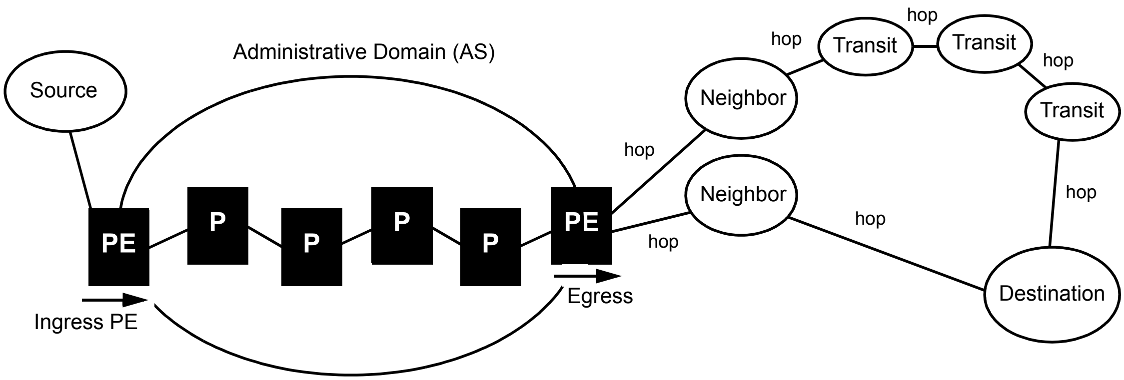

Figure 9 Graphical Representation of Networking Concepts

Note the following:

- Each source, neighbor, transit, and destination AS is an administrative domain. Administrative domains are based on perspective; for example, a transit is also a destination, neighbor, and source in relation to other administrative domains.
- Each administrative domain is connected to multiple administrative domains. Typically there are multiple entry and exit points for data flow within an administrative domain.
- Multiple Provider Edge (PE) routers can be connected to the Provider (P) routers. Only two PE routers are shown in Figure 9 for the ingress (data coming in) and egress (data going out) examples.

Top Changes Report

Top Changes Report highlights the top traffic changes across the entire network. Choose the daily, weekly, or monthly time frame, as shown in Figure 10.

If CoS is not configured, a message is shown in red at the bottom of the window when you open any of the Top Changes reports (Table 1). To include reports on CoS, click the configuration icon in the Top Changes window and select the desired CoS definitions under the Link Utilization by CoS menu option (see next procedure). When you apply your changes, a new table is added to the Top Changes report for each of the CoS definitions. For information on creating CoS definitions, see the “Administration” chapter in the *HP Route Analytics Management Software Administrator Guide*.

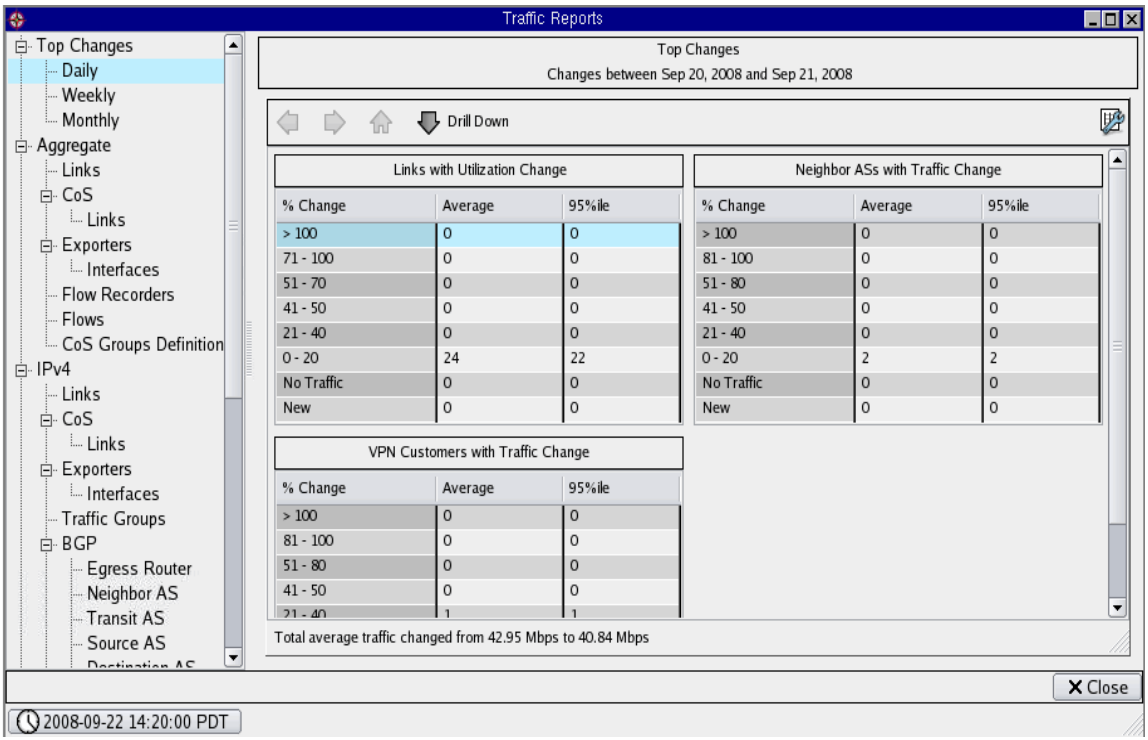

Figure 10 Top Changes Report Window

To modify the information presented in the display, perform the following steps:

1 Choose **Reports > Traffic Reports** to open the Traffic Reports window (Figure 9-1).

2 Click the **Configure** button to open the Configure Traffic Change Reports window (Figure 12).

3 Choose one of the following categories to modify:

- Link Utilization—Select to compute the traffic changes for each link found in the network.
- Neighbor AS bitrate—Select to compute change in traffic traveling to each Neighbor AS.
- VPN Customer bitrate—Select to compute traffic changes per VPN customer.
- Link Utilization by CoS—Select to compute the traffic changes for the selected CoS links found in the network.

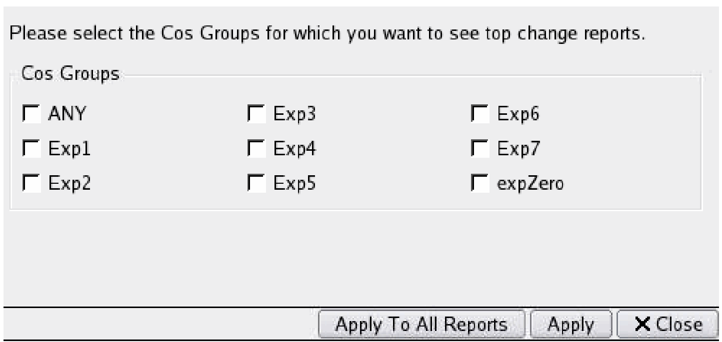

Figure 11 CoS Selection

4 Enter ranges for the top changes, where each entry represents the upper bound of a range. For example, the ranges shown in Figure 12 correspond to ranges 0-20, 20-40, 40-60, 60-80, and 80-100.

You can also open the Traffic Change Report by choosing **Administration > Traffic > Traffic Change Report**.

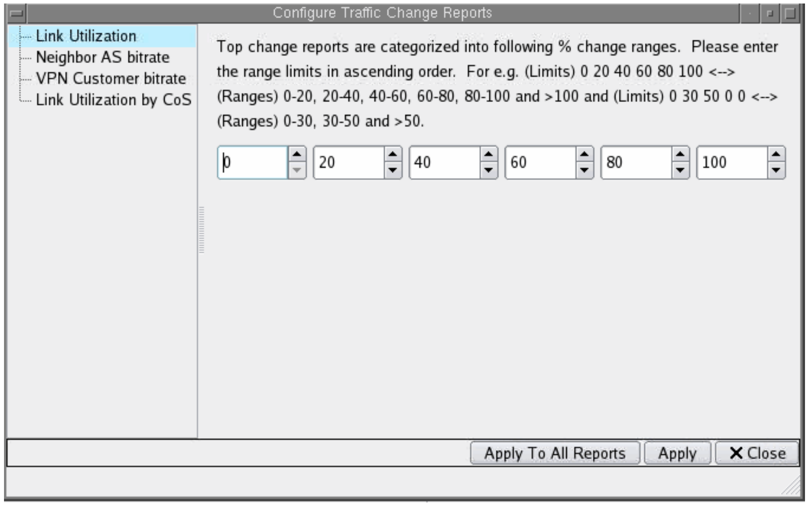

Figure 12 Configure Traffic Change Report Window

5 Click **Apply**.

Aggregate Reports

Aggregate Reports are available only when your network uses more than one protocol family. All aggregate reports include data from all Flow Collectors known to the Flow Analyzer. You cannot dynamically change the set of Flow Collectors. The Aggregate reports always shows a network view that is as complete as possible. By default, these reports cover the most recently recorded five minutes of traffic activity.

Aggregate – Exporter Links

This report shows the combined total of IPv4 and VPN traffic that the links attached to the exporting interfaces carry. The default fields for the report are shown in the following table.

Table 2 Aggregate - Exporter Links

Measure	Description
Link Source	Source router ID of an identified pair.
Link Destination	Destination router ID of an identified pair.
Index	Input SNMP interface index on the exporting router.
Name	Interface name.
Address	Interface IP address.
Capacity	The amount of traffic the link is capable of handling, in bits per second (bps).
Description	Description set on the router for the exporting interface.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Interfaces report include:

• Link	• IPv4 Flows
• Flow Collectors	• VPN Flows
• Flows	• History

Aggregate – Links

The Links report shows the combined total of all IPv4 and VPN traffic links in the network, and allows you to color links on the topology map based on traffic volume. The links that are identified with an asterisk (*) are attached to the exporting interface and their traffic volume information comes directly from the Netflow.

The default fields for the Links report are shown in the following table.

Table 3 Aggregate – Links

Measure	Description
Source	Source router ID of an identified pair.
Destination	Destination router ID of an identified pair.
Capacity	The amount of Egress traffic the link is capable of handling, in bits per second (bps).
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Color by Links is enabled only if the **5 Minute Avg Selected Interval** column is present in the window.

Options for the Links report include the following:

• Links	• IPv4 Flows
• Exporters	• VPN Flows
• Exporter Links	• CoS
• Flow Collectors	• History
• Flows	

Aggregate—CoS

This report displays the total of all IPv4 and VPN traffic seen with the specified CoS.

Table 4 Aggregate - CoS

Measure	Description
CoS	User-defined CoS name.
5-Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS report include the following:

•Links	•Customers	•Flow Collectors	•History
•Egress Routers	•Ingress PE	•Exporters	
•Source AS	•Ingress VRF	•Exporter Links	
•Neighbor AS	•Egress PE	•Flows	
•Transit AS	•Egress VRF	•IPv4 Flows	
•Destination AS		•VPN Flows	
•Community			

Aggregate—CoS - Links

The Links report shows the combined total of all IPv4 and VPN traffic links associated with a particular CoS in the network.

Table 5 Aggregate - CoS Links

Measure	Description
Source	Source router ID of an identified pair.
Destination	Destination router ID of an identified pair.
Capacity	The amount of traffic the link is capable of handling, in bits per second (bps).
CoS	CoS associated with this link.
5-Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS Links report include the following:

•Links	•Customers	•Flow Collectors	•History
•Egress Routers	•Ingress PE	•Exporters	
•Source AS	•Ingress VRF	•Exporter Links	
•Neighbor AS	•Egress PE	•Flows	
•Transit AS	•Egress VRF	•IPv4 Flows	
•Destination AS		•VPN Flows	
•Community			

Aggregate – Exporters

This report lists the total of all IPv4 and VPN Exporters traffic on the network. The default fields for the Exporters report are shown in the following table.

Table 6 Aggregate - Exporters

Measure	Description
Exporter	Name or IP address of peer from which flow information was received.
5-Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Exporters report include the following:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporter Links	•Traffic Groups
•Source AS	•Ingress VRF	•Flows	•History
•Neighbor AS	•Egress PE	•IPv4 Flows	
•Transit AS	•Egress VRF	•VPN Flows	
•Destination AS			
•Community			

If an Exporter is not present in the topology, the **Exporter** column is listed as unknown.

Aggregate – Flow Collectors

This report lists each Flow Collector on the network. The default fields for the Flow Collectors report are shown in the following table.

Table 7 Aggregate - Flow Collector

Measure	Description
Flow Collector	The ID of each Flow Collector on the network.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Flow Collectors report include:

•Links	•Ingress PE	•Exporters	•CoS
•Source AS	•Egress PE	•Exporter Links	•Traffic Groups
•Neighbor AS		•Flows	•History
•Transit AS		•IPv4 Flows	
•Destination AS		•VPN Flows	
•Community			

Aggregate – Flows

This report provides details for the all IPv4 and VPN flows on the network. The default fields for the Flows report are shown in the following table.

Table 8 Aggregate - Flows

Measure	Description
Flow Source	Source router ID of an identified pair.
Flow Destination	Destination router ID of an identified pair.
Exporter	Name or IP address of peer from which flow information was received.
Traffic Group	User-defined group name. See Creating Groups Using the Menu on page 122 for more information.
CoS	User-defined group name for the CoS.
Egress PE	Router from which flow exits the current AS.
VRF Label	Label assigned to VRF.
5 Minute Average Current Time	Average amount of traffic during the most recent 5 minutes of time.

Options for the Flows report include:

•Links
•Flow Collectors
•Show flow records
•Show Prefixes
•Show Paths

Aggregate – CoS Definition

This report lists the currently defined CoS. See the “Administration” chapter in the *HP Route Analytics Management Software Administrator Guide* for information on creating CoS definitions.

Table 9 Aggregate - COS Definition

Measure	Description
Name	CoS name.
EXP	EXP value or CoS value. This column is available only if you have a valid MPLS VPN license.
DSCP or TOS	DSCP or TOS value for the CoS. This is based on the mode used for CoS definitions.

Aggregate – Service Policies

For each interface and interface direction (input or output), this report lists the forwarding class, priority, and direction for services policies assigned to routers.

you can click the “i” icon in the upper right corner of the report window to open the Inspector panel. This panel presents policy information in the format that is appropriate to the router vendor. The inspector information includes information about the policy itself and all of the objects that the policy uses. Tabs in the Inspector panel allow you to drill down to policy and object details.

For more information on service policies, see “Route Policies and Service Policies Reports” on page 381.

Table 10 Aggregate - Service Policies

Measure	Description
Router	Router name.
Interface	Router interface.
Direction	Direction that the policy applies to (input or output).
Policy	Summary of the policy parameters.

IPv4 Traffic Reports

IPv4 is the dominant network layer protocol on the Internet. It is a best effort protocol in that it does not guarantee delivery, does not make any guarantees on the correctness of the data, and it can result in duplicated packets and/or packets out-of-order.

IPv4 Traffic reports include data from all Flow Collectors specific to IPv4 traffic. The IPv4 Traffic reports always shows a view that is as complete as possible. By default, these reports cover the most recently recorded five minutes of traffic activity.

These reports are available only for IPv4.

IPv4 Traffic – Exporter Links

Shows the IPv4 traffic that the links attached to the exporting interfaces carry. The default fields for the Interfaces report are shown in the following table.

Table 11 IPV4 Traffic - Exporter Links

Measure	Description
Link Source	Source router ID of an identified pair.
Link Destination	Destination router ID of an identified pair.
Index	Input SNMP interface index on the exporting router.
Name	Interface name.
Address	Interface IP address.
Capacity	The amount of traffic the link is capable of handling, in bits per second (bps).
Description	Description set on the router for the exporting interface.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for this report include:

• Links	• Flow Collectors	• Traffic Groups	• Top Sources
• Egress Routers	• IPv4 Flows	• History	• Top Destinations
• Source AS			• Top Protocols
• Neighbor AS			
• Transit AS			
• Destination AS			
• Community			

IPv4 Traffic – Links

This report shows how much IPv4 traffic each link is carrying, and allows you to color links on the topology map based on traffic volume. Links that are identified with an asterisk (*) are attached to the exporting interface and their traffic volume information comes directly from the Netflow.

The default fields for the Links report are shown in the following table.

Table 12 IPV4 Traffic - Links

Measure	Description
Source	Source router ID of an identified pair.
Destination	Destination router ID of an identified pair.
Capacity	User-defined capacity level.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the IPv4 Links report include:

•Links	•Flow Collectors	•CoS
•Egress Routers	•Exporters	•History
•Neighbor AS	•Exporter Links	•VPN Flows
•Transit AS	•IPv4 Flows	
•Destination AS		
•Community		

IPv4 Traffic – CoS

This report lists the total amount of IPv4 traffic seen with each CoS.

Table 13 IPV4 Traffic - CoS

Measure	Description
CoS	User-defined CoS name.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS report include the following:

•Links	•Egress Routers	•Flow Collectors	•History
•Ingress to Sites	•Source AS	•Exporters	
•Egress from Sites	•Neighbor AS	•Exporter Links	
•Source Sites	•Transit AS	•IPv4 Flows	
•Transit Sites	•Destination AS		
•Destination Sites	•Community		

IPv4 Traffic - CoS - Links

This report shows how much IPv4 traffic each CoS link is carrying.

Table 14 IPV4 Traffic - CoS Links

Measure	Description
Link Source	Source router ID of an identified pair.
Link Destination	Destination router ID of an identified pair.

Table 14 IPV4 Traffic - CoS Links

Measure	Description
Capacity	The amount of traffic the link is capable of handling, in bits per second (bps).
CoS	User-defined name of the CoS.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS Links report include the following:

•Links	•Flow Collectors	•History
•Egress Routers	•Exporters	
•Source AS	•Exporter Links	
•Neighbor AS	•IPv4 Flows	
•Transit AS		
•Destination AS		
•Community		

IPv4 Traffic – Exporters

This report lists each router used for obtaining traffic data. The default fields for the Exporters report are shown in the following table.

Table 15 IPV4 Traffic - Exporters

Measure	Description
Exporter	Name or IP address of Flow Collector peer that is exporting Net Flow data.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Exporters report include:

•Links	•Flow Collectors	•CoS
•Egress Routers	•Exporter Links	•Traffic Groups
•Source AS	•IPv4 Flows	•History
•Neighbor AS		
•Transit AS		
•Destination AS		
•Community		

IPv4 Traffic – Traffic Groups

This report shows the details of each user-defined traffic group. The default fields for the Traffic Groups report are shown in the following table.

Table 16 IPV4 Traffic - Traffic Groups

Measure	**Description**
Traffic Group	User-defined group name. See Creating Groups on the Routing Topology Map on page 115 for more information.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Traffic Groups report include:

•Links	•Exit Router-Next Hop	•Flow Collectors	•History
•Ingress to Sites	•Source AS	•Exporters	
•Egress from Sites	•Neighbor AS	•Exporter Links	
•Source Sites	•Transit AS	•IPv4 Flows	
•Transit Sites	•Destination AS		
•Destination Sites	•Community		

IPv4 Traffic – Top Sources

This report shows the top 100 source addresses per Flow Collector generating the most traffic. Advanced filtering is enabled. The default fields for the Top Sources report are shown in the following table.

Table 17 IPV4 Traffic - Top Sources

Measure	Description
Prefix	Top 100 source IP addresses with the highest amount of traffic.
Average Traffic (bps)	Average amount of traffic for each IP address.

There is no drill-down option for this report. Graphing capabilities are provided.

Right-click options for the Top Sources report include:

- Show Flow Records
- Show flow record browser

IPv4 Traffic – Top Destinations

This report shows the top 100 destination address prefixes that receive the most traffic. Advanced filtering is enabled. The default fields for the Top Destinations report are shown in the following table.

Table 18 IPV4 Traffic - Top Destinations

Measure	Description
Prefix	Top 100 destination IP addresses with the highest amount of traffic.
Average Traffic (bps)	Average amount of traffic for each IP address.

There is no drill-down option for this report. Graphing capabilities are provided.

Right-click options for the Top Destinations report include:

- Show Flow Records
- Show flow record browser

IPv4 Traffic – Top Conversations

This report shows the connections with the highest average amount of traffic (Mb/s) for the most recent five minutes of the selected time. Advanced filtering is enabled. The default fields for this report are shown in the following table.

Table 19 IPV4 Traffic - Top Conversations

Measure	Description
Source Address	Source address of an identified pair.
Destination Address	Destination address of an identified pair.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

There is no drill-down option for this report. Graphing capabilities are provided.

Right-click options for the Top Destinations report include:

- Show Flow Records
- Show flow record browser

IPv4 Traffic – Protocols

This report shows the Flow Collector traffic breakdown by IP protocol. The default fields for the Protocols report are shown in the following table.

Table 20 IPV4 Traffic - Top Protocols

Measure	Description
Protocols	All IP protocols. UDP and TCP protocols include the source and destination ports.
Port	Port for the protocol traffic.
Average Traffic (bps)	Average amount of traffic for each protocol.

There is no drill-down option for this report. Graphing capabilities are provided.

Right-click options for the Protocols report include:

- Show Flow Records
- Show flow record browser

IPv4 Traffic – Protocols – Source Ports

This report shows traffic flows for source ports for each protocol instance.

Table 21 IPV4 Traffic - Top Sources

Measure	Description
Protocol	Type of protocol.
Port	Source port.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

There is no drill-down option for this report.

Right-click options for the report include:

- Show Flow Records
- Show Flow Record browser

IPv4 Traffic – Protocols – Destination Ports

This report shows traffic flows for destination ports for each protocol instance.

Table 22 IPV4 Traffic - Top Sources

Measure	Description
Protocol	Type of protocol.
Port	Destination port.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.

There is no drill-down option for this report.

Right-click options for the report include:

- Show Flow Records
- Show Flow Record browser

IPv4 Traffic – Flows

This report shows the details of each traffic flow. The default fields for the Flows report are shown in the following table.

Table 23 IPV4 Traffic - Flows

Measure	Description
Flow Source	Within a flow, the prefix of the starting point.
Flow Destination	Within a flow, the prefix of the ending point.
Exporter	Name or IP address of peer from which flow information was received.
Traffic Group	User-defined group name. See Creating Groups on the Routing Topology Map on page 115 for more information.
CoS	User-defined group name.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Flows report include:

•Link	•Flow Collector
•Egress Router	•Show Flow Records
•Source AS	•Show Prefixes
•Neighbor AS	•Show Paths
•Transit AS	
•Destination AS	

Right-click options for the Flows report include the drill-down options and the following:

- Show Flow Records

- Show Prefixes
- Show Paths

IPv4 Traffic – Traffic Groups Definition

This report allows you to see the current definition of each traffic group.The default fields for the Traffic Groups Definition report are shown in the following table.

Table 24 IPV4 Traffic - Traffic Groups Definition

Measure	Description
Priority	User-defined order of groups.
Name	User-defined group name. See Creating Groups on the Routing Topology Map on page 115 for more information.
Source Prefixes	Source address of an identified pair.
Destination Prefixes	Destination address of an identified pair.
Source and/or Destination Ports	UDP and TCP protocols include the source and destination ports.
Protocols	All IP protocols.
CoS	User-defined identification assigned to a CoS.

There is no drill-down option for this report.

BGP Traffic Reports

The Border Gateway Protocol (BGP) maintains a table of IP networks or prefixes that designate network reachability among ASs. An AS is a collection of IP networks and routers frequently under the control of one company that presents a common routing policy to the Internet. BGP makes routing decisions based on path, network policies, and rulesets. BGP traffic is a subset of IPv4 traffic.

BGP Traffic reports include data from all Flow Collectors specific to BGP traffic. You cannot dynamically change the set of Flow Collectors. The BGP Traffic reports always shows a view that is as complete as possible. By default, these reports cover the most recently recorded five minutes of traffic activity.

BGP Traffic – Egress Router

This report shows the amount of traffic exiting the customer network through a given router at a single point in time. Traffic exiting at this router may go to many different BGP next hops. A next hop address usually corresponds to a peering link. The default fields for the Egress Router report are shown in the following table.

Table 25 BGP Traffic - Egress Router

Measure	Description
Egress Router	Name or IP address of the exit router used to reach a peer. This information is obtained from the Router Name repository prioritized by router names, DNS names, IP address, and system IDs.
Next Hop	IP address of next hop.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Egress Router report include:

•Links	•Customers	•Flow Collectors	•CoS
•Source AS	•Ingress PE	•Exporters	•Traffic Groups
•Neighbor AS	•Ingress VRF	•Exporter Links	•History

•Transit AS	•Egress PE	•Flows	
•Destination AS	•Egress VRF	•IPv4 Flows	
•Community		•VPN Flows	

BGP Traffic – Neighbor AS

This report shows how much traffic a neighbor AS is receiving. This information can help ensure a peering relationship is meeting expected levels. The default fields for the Neighbor AS report are shown in the following table.

Table 26 BGP Traffic - Egress Router

Measure	Description
Neighbor AS	Name or AS number of directly connected AS to which traffic is being sent.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Neighbor AS report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporters	•Traffic Groups
•Source AS	•Ingress VRF	•Exporter Links	•History
•Transit AS	•Egress PE	•Flows	
•Destination AS	•Egress VRF	•IPv4 Flows	
•Community		•VPN Flows	

BGP Traffic – Transit AS

This report lists each transit AS used in delivering data and how much data each transit was transmitting. The default fields for the Transit AS report are shown in the following table.

Table 27 BGP Traffic - Transit AS

Measure	Description
Transit AS	Name or AS number of each transit AS used in delivering data to a destination.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Transit AS report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporters	•Traffic Groups
•Source AS	•Ingress VRF	•Exporter Links	•History
•Neighbor AS	•Egress PE	•Flows	
•Destination AS	•Egress VRF	•IPv4 Flows	
•Community		•VPN Flows	

BGP Traffic – Source AS

This report shows how much traffic is being received from each source AS. The default fields for the Source AS report are shown in the following table.

Table 28 BGP Traffic - Source AS

Measure	Description
Source AS	Name or AS number of the source AS.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Source AS report include:

•Links	•Flow Collectors	•Traffic Groups
•Egress Routers	•Exporters	•History
•Neighbor AS	•Exporter Links	
•Transit AS	•IPv4 Flows	
•Destination AS		
•Community		

BGP Traffic – Destination AS

This report shows how much traffic each destination AS is receiving. The default fields for the Destination AS report are shown in the following table.

Table 29 BGP Traffic - Destination AS

Measure	Description
Destination AS	Name or AS number of the final AS.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Destination AS report include:

•Links	•Flow Collectors	•Traffic Groups
•Egress Routers	•Exporters	•History
•Neighbor AS	•Exporter Links	
•Transit AS	•IPv4 Flows	
•Destination AS		
•Community		

BGP Traffic – Community

This report shows how much traffic each community receives in bits per second (bps) by the first and second 2 octets (groupings of 8 bits). The default fields for the Community report are shown in the following table.

Table 30 BGP Traffic - Community

Measure	Description
First 2 Octets	The first two AS numbers for the given community.
Second 2 Octets	The second two AS numbers for the given community.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Community report include:

•Links	•Exporters	•Traffic Groups
•Egress Routers	•Exporter Links	
•Source AS	•IPv4 Flows	
•Neighbor AS		
•Transit AS		
•Destination AS		

VPN Traffic Reports

A virtual private network (VPN) is a communications network tunnelled through another network and dedicated for a specific network. VPNs can be used to separate the traffic of different user communities over an underlying network with strong security features.

All VPN Traffic reports include data from all Flow Collectors specific to VPN traffic. The VPN Traffic reports always shows a view that is as complete as possible. By default, these reports cover the most recently recorded five minutes of traffic activity.

VPN Traffic – Exporter Links

Each row in this report represents a router interface where VPN traffic was seen and exported over the network and identifies the links attached to the exporting interfaces.

The default fields for the Interfaces report are shown in the following table.

Table 31 VPN Traffic - Exporter Links

Measure	Description
Link Source	Source router ID of an identified pair.
Link Destination	Destination router ID of an identified pair.
Index	Input SNMP interface index on the exporting router.
Name	Interface name.
Address	Interface IP address.
Capacity	The amount of traffic the link is capable of handling, in bits per second (bps).
Description	Description set on the router for the exporting interface.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Interfaces report include:

•Links	•Flow Collectors	•History
•Customers	•VPN Flows	
•Ingress PE		
•Ingress VRF		
•Egress PE		
•Egress VRF		

VPN Traffic – Links

This report shows how much VPN traffic each link is carrying, and allows you to color links on the topology map based on traffic volume. Links that are identified with an asterisk (*) are attached to the exporting interface and their traffic volume information comes directly from the Netflow. The default fields for the Links report are shown in the following table.

Table 32 VPN Traffic - Links

Measure	Description
Source	Source router ID of an identified pair.
Destination	Destination router ID of an identified pair.
Capacity	User-defined capacity level.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Links report include:

•Links	•Flow Collectors	•CoS
•Customers	•Exporters	•History
•Ingress PE	•Exporter Links	
•Ingress VRF	•VPN Flows	
•Egress PE		
•Egress VRF		

VPN Traffic – CoS

This report displays the total amount of VPN traffic seen with each CoS. The fields for the CoS report are shown in the following table.

Table 33 VPN Traffic - CoS

Measure	Description
CoS	User-defined name CoS name.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS report include:

•Links	•Customers	•Flow Collectors	•History
•Egress Routers	•Ingress PE	•Exporters	
•Neighbor AS	•Ingress VRF	•Exporter Links	
•Transit AS	•Egress PE	•VPN Flows	
•Destination AS	•Egress VRF		
•Community			

VPN Traffic – CoS Links

This report displays the total amount of VPN traffic per link at each CoS. The default fields for the CoS Links report are shown in the following table.

Table 34 VPN Traffic - CoS Links

Measure	Description
Link Source	Displays the source dress for the link.
Link Destination	Displays the destination address for the link.
Capacity	Displays the capacity value for the link.
CoS	User-defined CoS name.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS Links report include:

•Links	•Customers	•Flow Collectors	•History
•Egress Routers	•Ingress PE	•Exporters	
•Neighbor AS	•Ingress VRF	•Exporter Links	
•Transit AS	•Egress PE	•VPN Flows	
•Destination AS	•Egress VRF		
•Community			

VPN Traffic – CoS Customers

This report displays the total amount of traffic per customer for each CoS. The default fields for the CoS Customers report are shown in the following table.

Table 35 VPN Traffic - CoS Customers

Measure	Description
Customer	User-defined customer ID. See Creating Customer and RT Associations on page 357 for more information.
CoS	User-defined CoS.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the CoS Customers report include:

•Links	•Ingress PE	•Flow Collectors	•History
•Egress Routers	•Ingress VRF	•Exporters	
•Neighbor AS	•Egress PE	•Exporter Links	
•Transit AS	•Egress VRF	•VPN Flows	
•Destination AS			
•Community			

VPN Traffic – Customers

This report displays traffic for each VPN customer and allows you to view the customer on the topology map. The default fields for the Customers report are shown in the following table. Options for the Customers report include the following.

Table 36 VPN Traffic - Customers

Measure	Description
Customer	User-defined customer ID. See Creating Customer and RT Associations on page 357 for more information.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Customers report include:

•Links	•Ingress PE	•Flow Collectors	•CoS
•Egress Routers	•Ingress VRF	•Exporters	•History
•Neighbor AS	•Egress PE	•Exporter Links	•Show Map
•Transit AS	•Egress VRF	•VPN Flows	
•Destination AS			
•Community			

The Show Map highlights the portion of the topology map that is utilized for a particular VPN by fading all the other nodes and links. The highlighted links can then be viewed in any of the color modes available for the full map. To return the topology map to its original setting, go to the Traffic Reports window and select **Restore Map** button.

VPN Traffic – Ingress PE

The Ingress PE report shows how much VPN traffic each link is carrying, and allows you to color links on the topology map based on traffic volume. The default fields for the Ingress PE report are shown in the following table.

Table 37 VPN Traffic - Ingress PEs

Measure	Description
Ingress PE	Displays each Ingress PE router found on the network.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Ingress PE report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Egress PE	•Exporters	•History
•Neighbor AS	•Egress VRF	•Exporter Links	•VPN Flows
•Transit AS		•VPN Flows	
•Destination AS			
•Community			

VPN Traffic – Ingress PE – Ingress VRF

The Ingress VRF report shows how much VPN traffic is entering each VRF at the ingress PE. It is populated only if Netflow is collected at the PE router. See "VPN RAMS Traffic" on page 387.

Table 38 VPN Traffic - Ingress VRFs

Measure	Description
Ingress PE	Displays each Ingress PE router found on the network.
VRF	Displays the VRF for the ingress traffic.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Ingress VRF report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Egress PE	•Exporters	•History
•Neighbor AS	•Egress VRF	•Exporter Links	
•Transit AS		•VPN Flows	
•Destination AS			
•Community			

VPN Traffic – Exporters

This report lists each router used for obtaining traffic data. The default fields for the Exporters report are shown in the following table.

Table 39 VPN Traffic - Exporters

Measure	Description
Exporter	Name or IP address of peer from which flow information was received.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Exporters report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporter Links	•History
•Neighbor AS	•Ingress VRF	•VPN Flows	
•Transit AS	•Egress PE		
•Destination AS	•Egress VRF		
•Community			

VPN Traffic – Egress PE

This report shows the amount of traffic leaving the PE routers. The default fields for the Egress PE report are shown in the following table.

Table 40 VPN Traffic - Egress PEs

Measure	Description
Egress PE	Router from which flow exits the current AS.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Egress PE report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporters	•History
•Neighbor AS	•Ingress VRF	•Exporter Links	
•Transit AS	•Egress VRF	•VPN Flows	
•Destination AS			
•Community			

VPN Traffic – Egress PE – Egress VRF

The Egress VRF report shows how much VPN traffic is destined for each VRF at the egress PE.

Table 41 VPN Traffic - Egress VRFs

Measure	Description
Ingress PE	Displays each egress PE router found on the network.
VRF	Displays the VRF for the egress traffic.
5 Minute Avg Selected Interval	Average amount of traffic during the most recent 5 minutes of the selected time.
5 Minute Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Ingress PE report include:

•Links	•Customers	•Flow Collectors	•CoS
•Egress Routers	•Ingress PE	•Exporters	•History
•Neighbor AS	•Ingress VRF	•Exporter Links	
•Transit AS		•VPN Flows	
•Destination AS			
•Community			

VPN Traffic – Flows

This report shows the details of each traffic flow. The default fields for the Flows report are shown in the following table.

Table 42 VPN Traffic - Flows

Measure	Description
Flow Source	Source router ID of an identified pair.
Flow Destination	Destination router ID of an identified pair.
CoS	User-defined CoS name.
Exporter	Address of the exporting router.
Egress PE	Router from which flow exits the current AS.
VRF Label	Label assigned to VRF.
5 Min Avg Current Time	Average amount of traffic during the most recent 5 minutes of the selected time.

Options for the Flows report include:

•Links	•Flow Collectors	•Show Flow Records
•Customers	•Exporters	•Show Prefixes
•Ingress PE	•Exporter Links	•Show Paths
•Ingress VRF	•VPN Flows	
•Egress VRF		

10 Path and Routing Analysis Reports

This chapter describes how to use path reports and routing analysis reports to analyze network connectivity and optimize routing performance.

Chapter contents:

Using Path Reports

Network connectivity is the ability of a router to reach all other routers in the network by sending packets of data along paths between source and destination routers. Each path consists of one or more links. Links are associated with a metric value, which is used to calculate the cost of the path.

The Path Reports tool computes and provides a summary of paths between routers in a selected topology, and allows you to break down the summary by area of interest, such as asymmetric links, unused links, and source routers. The reports are organized by analysis type. For example, you can view an analysis of all paths in the selected topology, or you can choose to view only asymmetric paths.

Path Reports are disabled in Monitoring mode.

Accessing the Path Reports Window

You can access path reports from the client application.

To access the Path Reports window, perform the following steps:

1 From the client application, choose **Reports > Path Reports**. to open the Select Topologies and Routers for Path Analysis window.

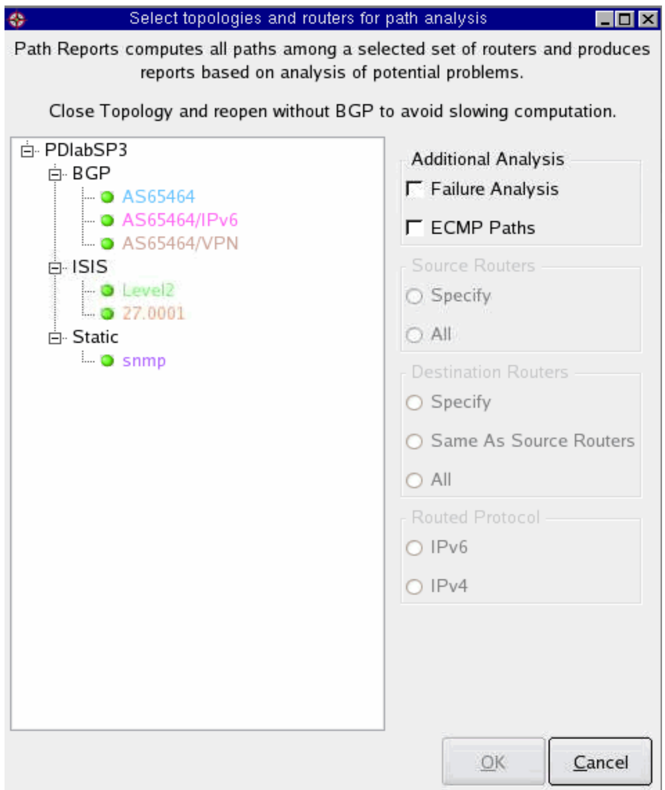

Figure 1 Path Reports - Select Topology

2 Choose one or more databases from the list.

You can choose any combination of databases using the Shift key to extend the range of selected items, or the Ctrl key to add or remove selected items. Selecting a higher-level folder implicitly selects all the folders contained within it.

When you open a topology for path analysis, avoid selecting BGP data when only IGP paths are of interest. Selecting BGP data may significantly increase the amount of time required to generate reports, as compared to non-BGP data. In general, computation of path reports for a 300-node, non-BGP topology can take up to 2 minutes to process, while a 300-node BGP topology can take 30-60 minutes to process.

3 If desired, select one or both of the following check boxes:

- **ECMP Paths**—Select this check box to enable the ECMP analysis option, which finds and lists multiple paths of the same cost. See ECMP Paths Analysis on page 460 for more information. If you do not select this check box, a single path is computed for each router pair, rather than multiple paths.
- **Failure Analysis**—Select this check box to cause the links in the selected database to fail, one link at a time. This can help determine which link failures are the most costly. See Failure Analysis on page 468 for more information.

  Enabling failure analysis, increases computation time significantly compared to reports that are generated without failure analysis. For example, when you enable failure analysis for a 300-node, non-BGP topology, you can add up to eight minutes to the computation time.

4 Choose the source routers to include in the path analysis:

- **All**—Include all source routers in the selected database.
- **Specify**—Choose the routers to include.

  When you click **Specify**, the system presents a list of available source routers. You can filter the list of routers by entering a regular expression in the RegEx text box at the top of the window. Select one or more routers from the Available Routers column, then click the arrow to move the specified routers to the Selected Routers column. Click **OK**.

The syntax of extended regular expressions is explained in Regular Expressions on page 226. The syntax is not the same as shell or file manager pattern patching, so a pattern like *-core-gw is not correct.

5 Select the destination routers to include in the path analysis:

- **All**—Include all destination routers in the selected database.
- **Specify**—Choose the routers to include.

  When you click **Specify**, the system presents a list of available source routers. You can filter the list of routers by entering a regular expression in the RegEx text box at the top of the window. Select one or more routers from the Available Routers column, then click the arrow to move the specified routers to the Selected Routers column. Click **OK**. For more information on regular expressions, see Regular Expressions on page 226 and Expression Definitions on page 228.

- **Same as Source Routers**—Destination routers included in the analysis will be the same as the source routers you chose in Step 4.

6 Select the routed protocols to include in the path analysis:

- **IP4** or **IPv6**—Route resolution for IP prefixes.
- **ISO OSI**—Route resolution for OSI prefixes.

  The options in Step 6 are displayed only if OSI IS-IS or IPv6 topologies are present.

7 Click **OK** to begin generating reports.

You can cancel the report generation at any time. Partial results are displayed in the Path Reports window.

Using the Path Reports Window

The Path Reports window displays reports in the following categories:

- Path Analysis Reports on page 455
- Network Element Analysis on page 464
- Failure Analysis on page 468 (included only if you select the **Failure Analysis** check box when opening the report)

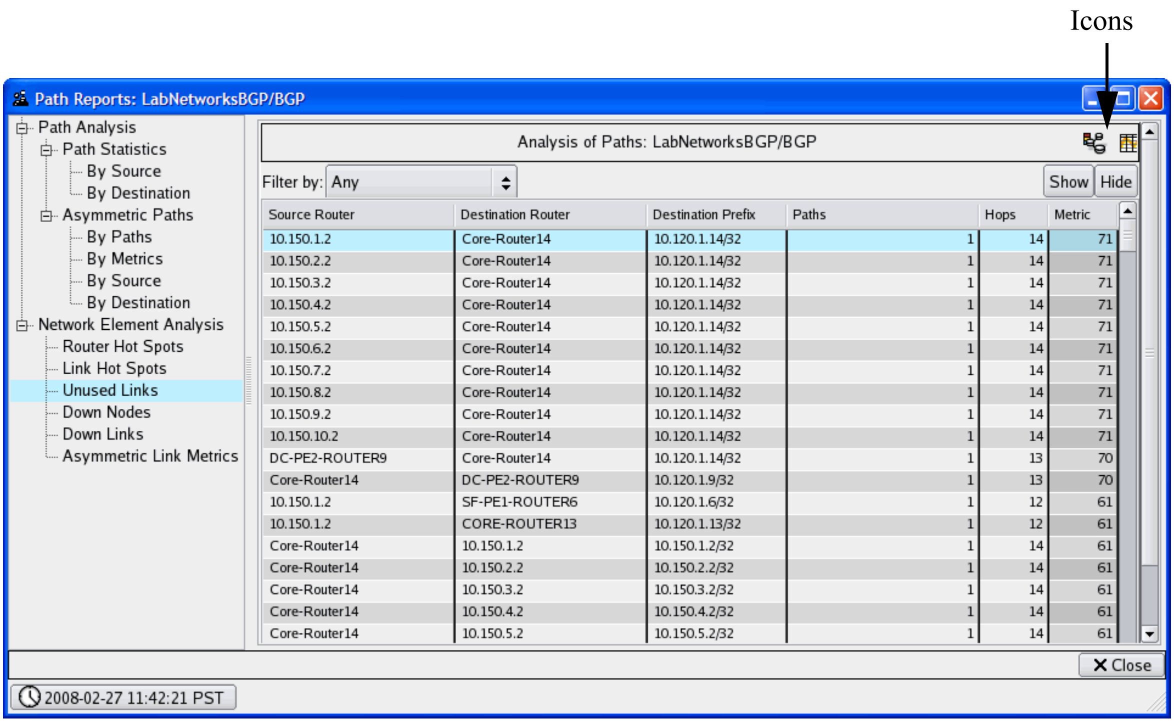

Figure 2 Path Reports Window

The Path Reports window may show the icons in Table 1 , depending on the selected report.

Table 1 Path Reports Icons

	**Color paths**	Highlight paths on the routing topology map. Click a particular row in the Path Reports table and the routing topology map highlights the paths between the selected the source and destination routers (see Figure 3).
	**Show paths**	Display a table that lists all paths between the selected source and destination router pair. The paths are broken down by hop or link.
	**Color by**	Color elements on the routing topology map. For example, click the **Color by** ... icon in the Hot Nodes table to color all hot nodes on the routing topology map. A second Legend panel is displayed on the topology map to describe each colored element. Click the icon again to uncolor the highlighted elements.
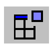	**Tear Off**	Open the table in a new window.
	**Put Back**	Place the table that you opened in a new window back in the original window.
	**Show Detailed Path Statistics**	View more details about the paths originating or ending with a selected router. For example, clicking this icon opens an Analysis of Paths table in the lower half of the Path Reports window (see Figure 4). You can also right-click a router in the table and choose **Show Detailed Path Statistics** from the pop-up menu to achieve the same result.
	**Show Effect of Link Failures**	View more details about the failed links.
	**View as Bar Chart**	Display the information in bar chart format.
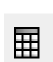	**View as Table**	Display the information in table format.
	**Color Hot Routers**	Highlight the routers that have routing issues on the routing topology map.

Path Analysis Reports

This section describes the following path analysis options, which are available from the Path Reports window (see Accessing the Path Reports Window on page 450):

- Path Statistics Report on page 455
- ECMP Paths Analysis on page 460
- Single (Non-ECMP) Path Analysis on page 461
- Asymmetric Paths Analysis on page 462

Path Statistics Report

The Path Statistics report lists all paths between router pairs in the selected database. If you choose specific source and destination routers to analyze, as described in Accessing the Path Reports Window on page 450, only those routers appear in the Path Statistics table.

The Path Statistics table allows you to identify where connectivity has been lost. For example, if Router A cannot reach Router D, then Path Not Found is listed in the **Paths** column of the table. In addition, the Path Statistics table lists paths from highest metric value to lowest, making the costliest paths immediately evident.

The Path Statistics table includes the following columns:

- **Source Router** — Router where the path originates.
- **Destination Router** — Router where the path ends. Because a router may advertise multiple prefixes, a destination prefix is used as the destination IP address for the path.
- **Destination Prefix** — Unique prefix that identifies the destination router. For example, if a node advertises 20 prefixes, the system isolates one prefix that is advertised by no other router, and uses this prefix as the address of the destination router. If multiple unique prefixes are found, the following rules determine the destination:

  a. Select a unique prefix with lowest metric value.

  b. If the metrics of all unique prefixes are equal, select the prefix with the longest length.

c. If the metrics of all unique prefixes are equal, and all are of the same length, select the prefix with the smallest IP address.

If none of the prefixes advertised by a router is unique, then the system considers non-unique prefixes. However, if a destination prefix cannot be determined, then no paths to the destination router are computed. All such routers are included in the Down Nodes report as described in Down Nodes on page 467.

- **Paths** — Number of paths found between the source router and the destination router. Unless you have chosen to compute ECMP paths, this number is 1.

  If the system cannot compute a path between the source and destination router, one of the following messages is displayed:

  — **Destination Not Reached** — The final hop of the path is not the destination router. This can occur if the destination prefix used to reach the destination router (Router B) is also being advertised by another router (Router C). If the final hop is Router C rather than Router B, the path cannot be computed.

  — **Path Not Found** — No path is found between the source and destination router.

  — **Loop Detected** — The number of hops between the source and destination router exceeds 30. A loop is preventing the system from computing the path properly.

- **Hops** — Number of hops between the source router and destination router. If a range of values is displayed in this column, the range represents the minimum and maximum number of hops for the paths between the source and destination router. For example, if there are three paths between Router A and Router B, with a range of 4-9 hops, Path 1 is made up of four hops, Path 2 is made up of five hops, and Path 3 is made up of nine hops.

- **Metric** — The metric value of a path is calculated by adding up the link metric values along the path. For example, a path between Router A and Router B consists of two hops. Hop 1 has a link metric value of 10, and Hop 2 has a link metric value of 20. The total metric value of the path is 30. If a range of metrics is displayed in the column, the range represents the minimum and maximum metric for paths found between the source and destination routers.

To drill down further and view path details for a specific row in the table, highlight the row and then click the **Show Paths** in the upper right corner of the Path Reports window.

Each path between the highlighted source and destination routers is listed in a separate row, broken down by link (Hop 1, Hop 2 ...). For example, if four paths are found between Router A and Router B, each of the four paths and the associated hops are listed along with the following information:

- **Path** — Paths corresponding to the row you selected in the upper half of the window. If more than one path is found between the source and destination router, the paths are labeled Path 1, Path 2 ... and so on. A collapsible list of the hops for each path is also shown in this column.
- **Source Router** — Router where the link originates.
- **Destination Router** — Router where the link ends.
- **Metric** — Metric value for the link.
- **Protocol** — Routing protocol associated with the link.

To view path information on the routing topology map, click the corresponding row in the table. The path, link, or node is highlighted on the routing topology map as shown in Figure 3.

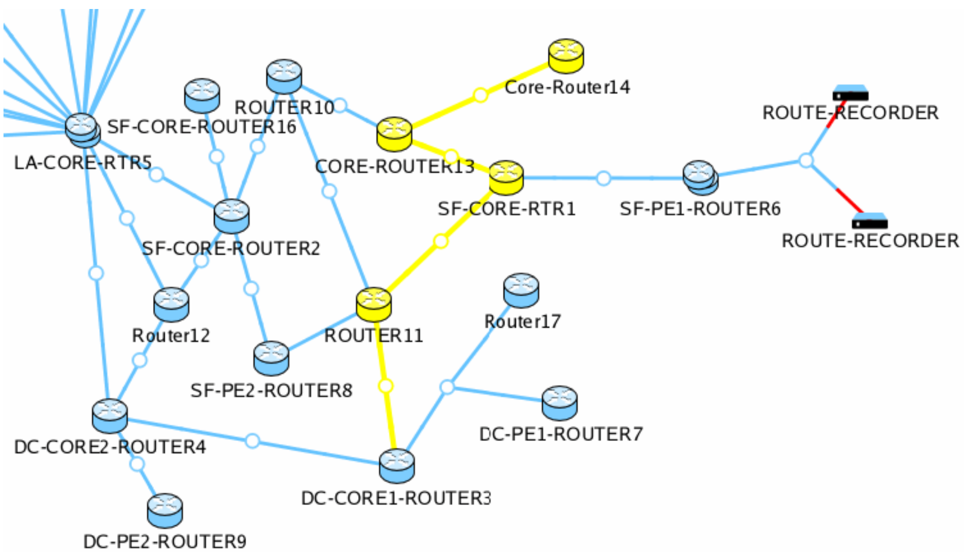

Figure 3 Path Highlighted on the Routing Topology Map

Path Statistics by Source

To view the number of paths that are reachable by each source router, click **By Source** in the left panel. The following information is displayed:

- **Source Router** — Router where a path originates.
- **Reachable Destinations** — Number of routers that the source router can reach.
- **Paths** — Range of paths to all reachable destinations. For example, 1-5.
- **Hops** — Number of hops along the path between the selected node and a corresponding destination router. If a range of numbers is displayed, this column reflects the minimum and maximum number of hops for the set of paths originating with the source router.
- **Metric** — The metric value for the paths originating with the selected node. If a range of values is displayed, this column reflects the minimum and maximum metric value for the set of paths originating with the source router.

To drill down further and view detailed path statistics for a source router in the Paths by Source table, highlight the source router, then click the **Show Paths** icon.

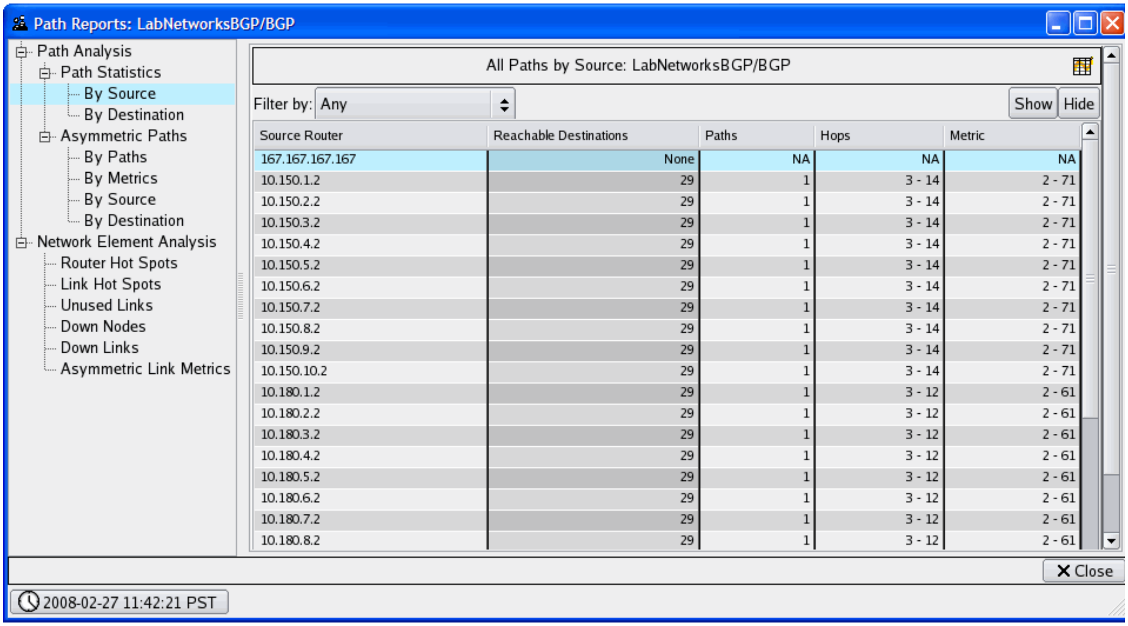

Figure 4 All Paths by Source with Analysis of Paths Table

To view details for each of the paths listed in the Analysis of Paths table, click the Show Paths icon, as described in Path Statistics Report on page 455.

Path Statistics by Destination

To view the number of paths reachable by each destination router, click **By Destination** in the left panel. The following information is displayed:

- **Destination Router** — Node where a path ends.
- **Reachable by** — Number of routers that can reach the specified destination router.
- **Paths** — Range of paths whose destination is the specified node.
- **Hops** — Number of hops along the path between the selected node and a corresponding source router. If a range of numbers is displayed, this column reflects the minimum and maximum number of hops for the set of paths ending with the source router.

- **Metric** — The metric value for the paths ending with the selected node. If a range of values is displayed, this column reflects the minimum and maximum metric value for the set of paths ending with the source router.

To drill down further and view detailed path statistics for a destination router in the Paths by Destination table, highlight the destination router, then click the **Show Paths** icon.

ECMP Paths Analysis

The Analysis of ECMP Paths table lists the paths between source and destination router pairs that are of equal cost. This information helps identify the amount of redundancy in the network and allows you to make adjustments accordingly. For example, if there are two equal-cost paths between Router A and Router B, you might determine that two paths do not provide enough redundancy and decide to increase the number of equal-cost paths between Router A and Router B. The opposite is also true: if you do not want equal-cost paths in your network, you can use the Analysis of ECMP Paths table to find and eliminate such paths.

- **Source Router —** Router where the path originates.
- **Destination Router —** Name or ID of the router where the path ends. Since a router may advertise multiple prefixes, a destination prefix is used as the destination IP address for the path.
- **Destination Prefix** — A prefix unique to each router is used as the destination for the destination router. For more information about destination prefixes, see Path Statistics Report on page 455.
- **Paths** — Number of equal-cost paths between the source and destination routers.
- **Hops** — Number of hops making up each equal-cost path. If a range of values is displayed, this column reflects the minimum and maximum number of hops for the set of paths.
- **Metric** — Total metric value of the path.

To view the Analysis of ECMP Paths table by source or destination router, click **By Source** or **By Destination** in the left pane to organize data accordingly, then see Path Statistics by Source on page 458 and Path Statistics by Destination on page 459 for more information.

To drill down further and view path details for a specific row in the By Source or By Destination table, highlight the row, then click the **Show Paths** icon.

Single (Non-ECMP) Path Analysis

The Analysis of Single Paths table lists paths between source and destination router pairs for which there is not another path of equal cost. This table can help identify a lack of redundancy between vital routers. The Analysis of Single Paths table is displayed only when the ECMP Paths check box is selected on the Select Topologies dialog box, as described in Accessing the Path Reports Window on page 450.

The Analysis of Single Paths table has the following columns:

- **Source Router** — Router where the path originates.
- **Destination Router** — Name or ID of the router where the path ends. Since a router may advertise multiple prefixes, a destination prefix is used as the destination IP address for the path.
- **Destination Prefix** — A prefix unique to each router is used as the destination for the destination router. For more information about destination prefixes, see Path Statistics Report on page 455.
- **Paths** — Number of paths between the source and destination routers whose cost is not equal to any other path between the source and destination router.
- **Hops** — Number of hops making up each path.
- **Metric** — Total metric value of the path.

To view the Analysis of Single Paths table by source or destination router, click **By Source** or **By Destination** in the left pane to organize data accordingly, then see Path Statistics by Source on page 458 and Path Statistics by Destination on page 459 for more information.

To drill down further and view path details for a specific row in the By Source or By Destination table, highlight the row, then click the **Show Paths** icon..

Asymmetric Paths Analysis

An asymmetric path can exist when the forward cost of a path between two routers differs from the reverse cost of a path between the same two routers. In other words, if the cost of a path from Router A to Router B is 20, and the cost of a path from Router B to Router A is 40, the paths are asymmetric. In addition, a path may be asymmetric if the number of forward hops differs from the number of reverse hops; if the forward and reverse hops themselves differ; or if the number of forward paths differs from the number of reverse paths. Identifying asymmetric paths can help isolate network misconfigurations.

To drill down further and view path details for a specific row, highlight the row, then click the **Show Paths** icon .

The Analysis of Asymmetric Paths window is shown in Figure 5. The Analysis of Asymmetric Paths table has the following columns:

- **Source Router** — Node where the path originates.
- **Destination Router** — Node where the path ends.
- **Forward Paths** — Number of paths from the source router to the destination router.
- **Reverse Paths** — Number of paths from the destination router to the source router.
- **Forward Hops** — Number of hops taken along the path from the source router to the destination router.
- **Reverse Hops** — Number of hops taken along the path from the destination router back to the source router.
- **Forward Metric** — Metric value, or cost, of the path from the source router to the destination router.
- **Reverse Metric** — Metric value, or cost, of the path from the destination router back to the source router.
- **Metric Difference** — Difference between the metric values of the forward path and the reverse path between the source and destination routers. For EIGRP protocol routers, this value represents the difference in bandwidth plus the difference in delay.

Asymmetric path reports do not include ECMP paths. For more information about ECMP paths, see ECMP Paths Analysis on page 460.

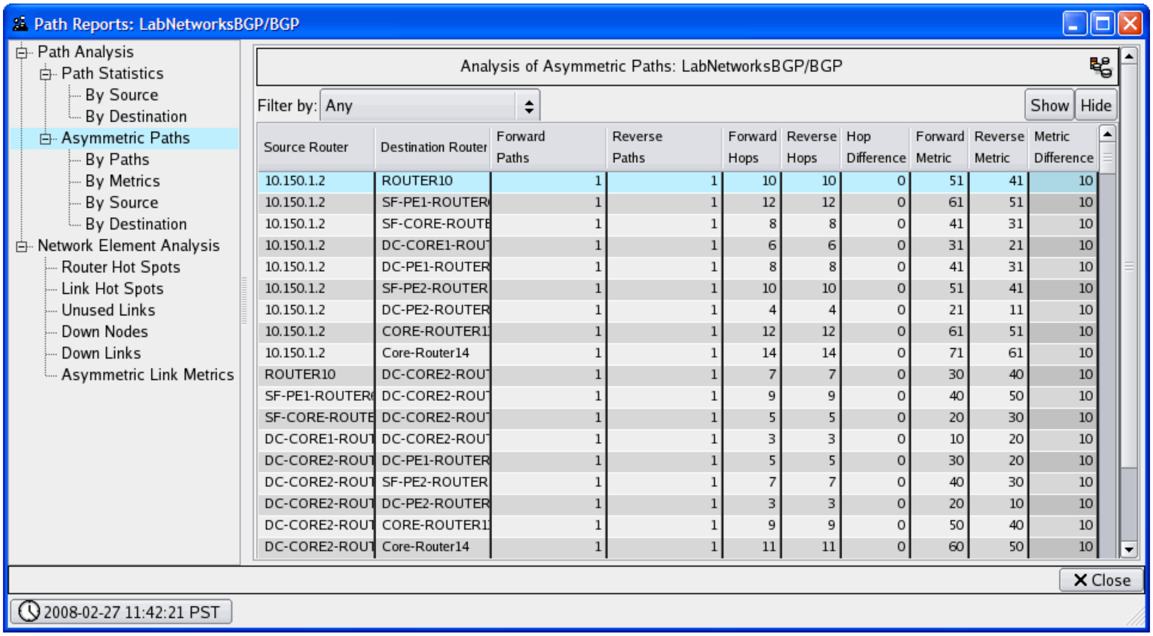

Source Router	Destination Router	Forward Paths	Reverse Paths	Forward Hops	Reverse Hops	Hop Difference	Forward Metric	Reverse Metric	Metric Difference
10.150.1.2	ROUTER10	1	1	10	10	0	51	41	10
10.150.1.2	SF-PE1-ROUTER(	1	1	12	12	0	61	51	10
10.150.1.2	SF-CORE-ROUTE	1	1	8	8	0	41	31	10
10.150.1.2	DC-CORE1-ROUT	1	1	6	6	0	31	21	10
10.150.1.2	DC-PE1-ROUTER	1	1	8	8	0	41	31	10
10.150.1.2	SF-PE2-ROUTER	1	1	10	10	0	51	41	10
10.150.1.2	DC-PE2-ROUTER	1	1	4	4	0	21	11	10
10.150.1.2	CORE-ROUTER1	1	1	12	12	0	61	51	10
10.150.1.2	Core-Router14	1	1	14	14	0	71	61	10
ROUTER10	DC-CORE2-ROUT	1	1	7	7	0	30	40	10
SF-PE1-ROUTER(	DC-CORE2-ROUT	1	1	9	9	0	40	50	10
SF-CORE-ROUTE	DC-CORE2-ROUT	1	1	5	5	0	20	30	10
DC-CORE1-ROUT	DC-CORE2-ROUT	1	1	3	3	0	10	20	10
DC-CORE2-ROUT	DC-PE1-ROUTER	1	1	5	5	0	30	20	10
DC-CORE2-ROUT	SF-PE2-ROUTER	1	1	7	7	0	40	30	10
DC-CORE2-ROUT	DC-PE2-ROUTER	1	1	3	3	0	20	10	10
DC-CORE2-ROUT	CORE-ROUTER1	1	1	9	9	0	50	40	10
DC-CORE2-ROUT	Core-Router14	1	1	11	11	0	60	50	10

Figure 5 Analysis of Asymmetric Paths Table

Asymmetric Paths by Path

The Analysis of Asymmetric Paths: By Paths table lists paths that are asymmetric due to a mismatch in forward and reverse hops. Click **By Path** in the left pane to display the data.

Asymmetric Paths by Metric

The Analysis of Asymmetric Paths: By Metric table lists paths that are asymmetric due to a mismatch in forward and reverse metrics. Click **By Metric** in the left pane to display the data.

Asymmetric Paths by Source

To view the Analysis of Asymmetric Paths table by source router, click **By Source** in the left pane to organize data accordingly. See Path Statistics by Source on page 458 for a description of this table.

Asymmetric Paths by Destination

To view the Analysis of Asymmetric Paths table by destination router, click **By Destination** in the left pane to display the data.

Network Element Analysis

Determining which network elements play too large a part in network routing and which are playing no part at all is required for network performance optimization. Hotspots are routers or links that are used more frequently than other elements in the network. For example, if Router A is used in 20 paths, all 20 of those paths will be affected should the router fail. Conversely, coldspots are network elements that are under-utilized. If Router B is not used in any paths, you can optimize performance by making better use of Router B, and relieving Router A of some of the load.

This section describes the following network element analysis options, which are available from the Path Reports window (see Accessing the Path Reports Window on page 450):

- Router Hot Spots on page 464
- Link Hot Spots on page 465
- Unused Links on page 466
- Down Nodes on page 467
- Down Links on page 467
- Asymmetric Link Metrics on page 467

To drill down further and view path details for a specific row, highlight the row, then click the **Show Paths** icon.

Router Hot Spots

This table is sorted by the routers that are used most frequently in paths on the network (Figure 6).

The Hot Nodes table has the following columns:

- **Node** — Name or ID of the “hot” router.
- **Paths** — Number of source and destination router pairs that include the “hot” router.

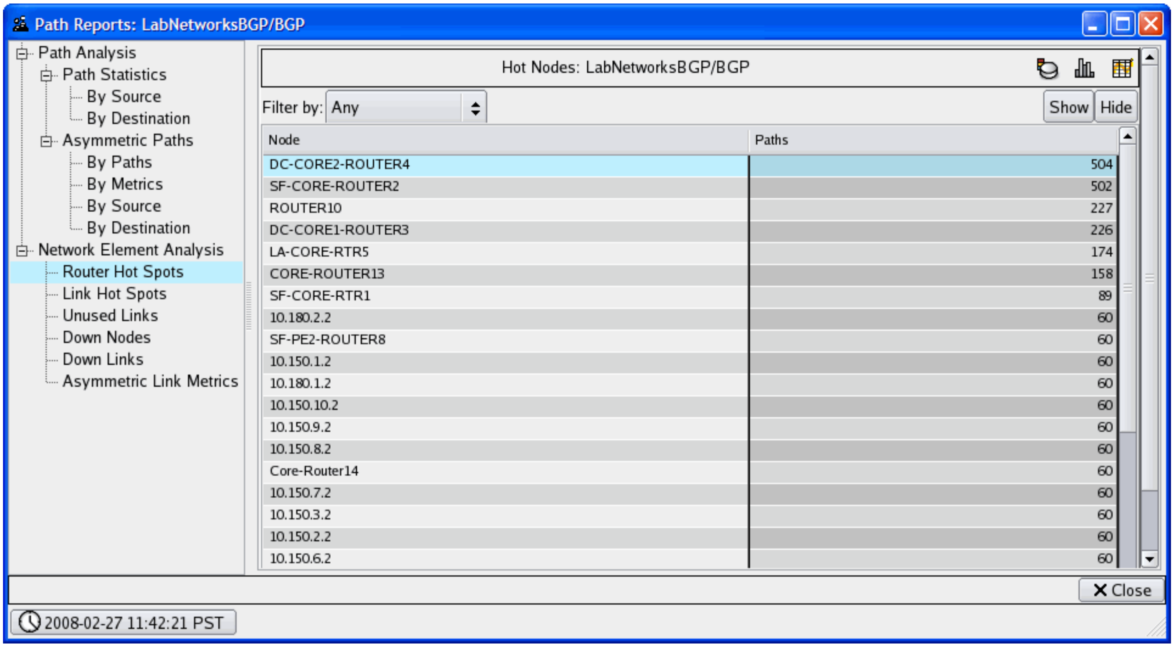

Figure 6 Router Hot Spots Table

Link Hot Spots

Similar to the Router Hot Spots table, the Link Hot Spots table is sorted by the links that are most frequently used in paths (Figure 7). The more frequently a link is used in a path, the more paths will be affected if that link fails.

The Hot Links table has the following columns:

- **Link** — The address of the "hot" link.
- **Paths** — The number of source and destination router pairs that use the "hot" link in their paths.

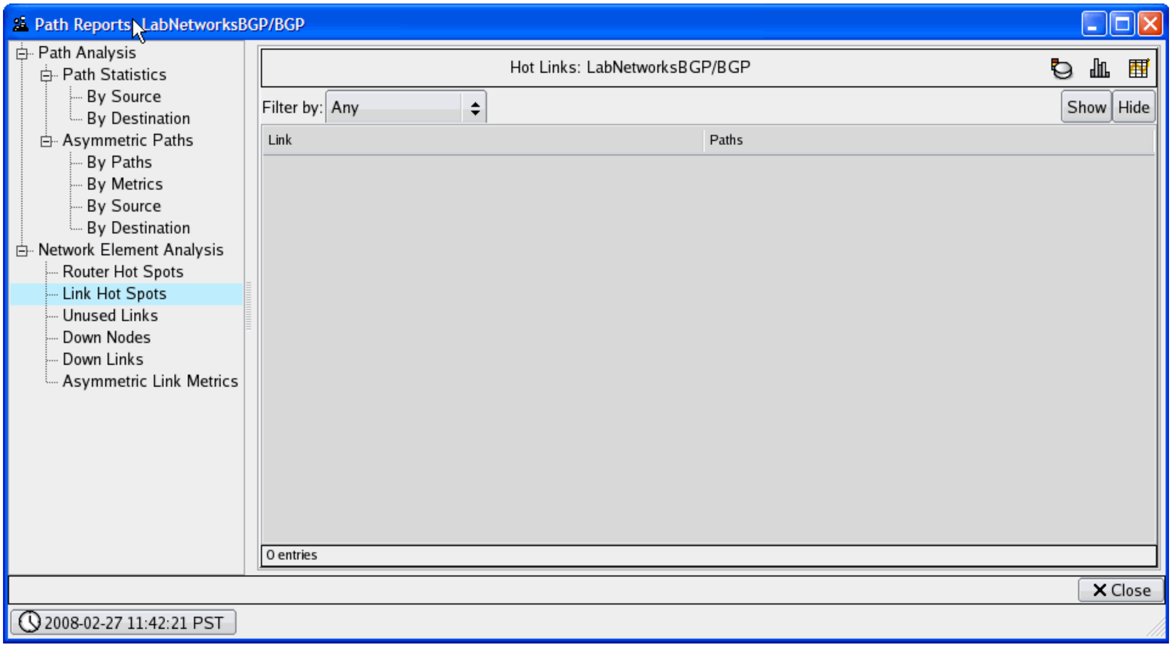

Figure 7 Link Hot Spots Table

Unused Links

This table lists all links in the network that are not being used in any paths. The Unused Links table has the following columns:

- **Link** — The unused link.
- **Source Interface** — The interface where the link originates.
- **Destination Interface** — The interface where the link ends.
- **Metric** — Metric value of the unused link.
- **State** — Indicates whether the link is up or down.
- **Area** — Location of the link.

Down Nodes

This table lists routers that are currently down, have failed, or do not have a destination prefix. The Down Nodes table has the following columns:

- **Router** — The node that has gone down.
- **IP Address** — Address of the node that has gone down.
- **Type** — Indicates the type of router (for example, internal, area border router, AS external).
- **Protocol** — Indicates the protocol of the router.
- **State** — Indicates whether the node is down or does not have a destination prefix.
- **Area** — Location of down node.

Down Links

This table lists links that are currently down, or have failed. The Down Links table has the following columns:

- **Link** — The down link.
- **Source Interface** — The interface where the down link originates.
- **Destination Interface** — The interface where the down link ends.
- **Metric** — Total metric value of the down link.
- **State** — Indicates that the link is down.
- **Area** — Location of down link.

Asymmetric Link Metrics

This table lists links whose forward and reverse metrics are different. The Asymmetric Link Metrics table is similar to the Asymmetric Paths table, described (see Asymmetric Paths Analysis on page 462). The Asymmetric Links table has the following columns:

- **Link** — The address of the asymmetric link.
- **Source Interface** — The interface where the link originates.

- **Destination Interface** — The interface where the link ends.
- **Forward Metric** — The value of the metric from the source interface to the destination interface.
- **Reverse Metric** — The value of the metric from the destination interface to the source interface.
- **Metric Difference** — The difference between the forward and reverse metrics.
- **Area** — Location of the link.

Failure Analysis

When you select the Failure Analysis check box on the Select Topology dialog box, as described in Accessing the Path Reports Window on page 450, each link is systematically failed along every path in the selected database. The results of the simulated link failure appear in the Failure Analysis tables, which you can use to isolate the links that would be most damaging in the event of a failure.

This section describes the following failure analysis options, which are available from the Path Reports window (see Accessing the Path Reports Window on page 450):

- Path Failure Analysis on page 468
- Link Failure Analysis on page 469
- Failure-induced ECMP Analysis on page 470

To drill down further and view the effect of the link failure for a specific path in the table, highlight the row corresponding to the path, then click the **Show Effect of Link Failures** icon.

Path Failure Analysis

This table lists each path that failed as part of the failure analysis and includes the following columns:

- **Source Router** — Node where the path originates.
- **Destination Router** — Name or ID of the node where the path ends. Since a router may advertise multiple prefixes, a destination prefix is used as the destination IP address for the path.

- **Destination Prefix** — A prefix unique to each router is used as the destination for the destination router. For more information about destination prefixes, see Path Statistics Report on page 455.
- **Original Paths** — Number of paths discovered between the source and destination router before failure analysis.
- **Original Hops** — Range of hops discovered between the source and destination router before failure analysis.
- **Original Metric** — Metric of the path between the source and destination routers before failure analysis.
- **Worst Link Failure**— Address of the link whose failure caused path cost to increase most of all failed links.
- **Worst Metric Change**— The difference between the original metric of the path and largest metric in the List of Paths table resulting from the link failure.

The Link Failure Analysis for Path table has the following columns:

- **Link** — This column lists all links whose failure had an effect on the path (for example, caused a change in the number of hops or in the metric value).
- **New Paths** — Number of paths after the link failure. Compare this value to **Original Paths** in the Path Failure Analysis table.
- **New Hops** — Range of hops after the link failure. Compare this value to **Original Hops** in the Path Failure Analysis table.
- **New Metric** — Metric value after the link failure. Compare this value to **Original Metric** in the Path Failure Analysis table.
- **Damage Metric** — The difference between the original metric of he path and the largest metric resulting from link failure.

Link Failure Analysis

This table lists each failed link and the number of paths that were damaged as a result of the link failure and includes the following columns:

- **Link** — Address of the failed link.
- **Damaged Paths** — Number of paths with a cost that increased after the link failed.

View the effect of the link failure for a specific path in the Analysis of Paths table, highlight the row, then click the **Show Effect of Link Failures** icon.

Failure-induced ECMP Analysis

This table lists paths that have become equal-cost as the result of link failure. If equal-cost paths are not desired in the network, use this table to pinpoint the links whose failure would cause equal-cost paths to occur and reconfigure the network accordingly. See ECMP Paths Analysis on page 460 for more information about equal-cost paths.

The Failure Induced ECMP table contains the same information as that in the Path Statistics table (see Path Statistics Report on page 455). To drill down further and view the effect of the link failure for a specific path in the table, highlight the path row, then click the **Show Effect of Link Failures** icon.

Using Routing Analysis Reports

Routing Analysis Reports provide detailed routing information on IGP topologies.

You must be in Analysis mode to view these reports. IPv6 prefixes are included if IPv6 is licensed and current topologies have IPv6 data.

To configure and view the Routing Analysis reports, perform the following steps:

1 In Analysis mode, choose **Reports > Routing Analysis Reports**.

The Routing Analysis Reports configuration window opens.

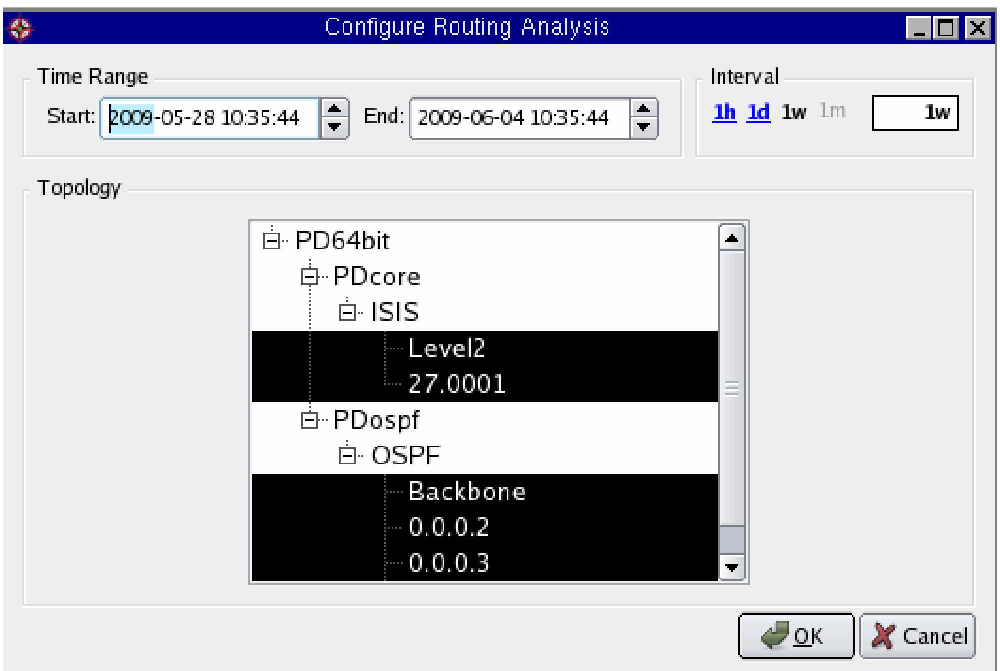

Figure 8 Routing Analysis Reports Configuration

2 Configure a start and end time for the analysis by specifying dates and times in the Start and End fields. The scale units, 1h (one hour), 1d (one day), 1w (one week), and 1m (one month) adjust the start time such that the interval between the start and end times is the selected period (such as one hour or one day).

3 To restrict the report topology to a subset of the current topology, click the button with the topology name just above the **OK** button. This option allows you to focus on the items of interest, while reducing the size of the overall reports.

4 Click **OK** to accept the settings and open the report window.

Routing Analysis Reports Buttons

The following buttons are available in Routing Analysis reports, depending on the specific report selection.

Table 2 Routing Analysis Report Buttons

	**Drill-down**	If available, this button allows to see finer detail within a set of data.
	**Go back one drill-down**	During a drill-down, goes back one drill-down level.
	**Go forward one drill-down**	During a drill-down, goes forward one-drill down level.
	**Go back to top; undo all drill-downs**	Goes to the highest point in the drill-down hierarchy, and "unrolls" the drill-down view from the window.
	**Configure**	Opens the configuration window. See Using Routing Analysis Reports on page 471.
	**Advanced Filter**	Allows you to define advanced filters. See Advanced Filtering on page 397 and Using Filters on page 221 for more information.
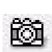	**Snapshot**	Opens a new window containing a snapshot of the current window.

Routing Analysis Reports

The Routing Analysis Reports are grouped into Routing Stability Reports and Change Reports.

All reports are based on the configured time interval and topology.

The following Routing Stability Reports present the numbers of events of different types in the selected topology during the specified time interval:

- Router Churn on page 473
- Link Flaps on page 474
- IPv4 Prefix Flaps on page 474
- IPv6 Prefix Flaps on page 475

Change Reports compare the state of the network elements in the selected topology at the start and end times. A row appears in the report only if an attribute has changed, such as a change to the router or to a prefix that the router is announcing. The Change reports also have special tooltip options. If you move your cursor over an entry in the first column of the report, you will see the list of attribute values that did not change.

The following reports are available:

- Routers on page 475
- Links on page 476
- IPv4 Prefixes on page 477
- IPv6 Prefixes on page 477

Router Churn

This report categorizes all of the events that are generated by the routers in the selected topology. Similar data is also presented in the Network Churn report that is available in the web interface.

This report includes the following columns:

- Name—Name of IP address of the router.

- Type—Router type.
- Total—Number of total events in the specified interval.
- Router—Number of events that were specific the router (such as going up or down).
- Link—Number of link events.
- IPv4 Prefixes—Number of IPv4 prefix events.
- IPv6 Prefixes—Number of IPv6 prefix events.
- Area—Router's area.

Link Flaps

For each link, this report lists the number of times that the link went from up to down.

This report includes the following columns:

- Link—Specific link (as included in the List of all Links window, see Links List on page 109).
- Interface—Link interface (as included in the List of all Links window, see Links List on page 109).
- Count—Total number of flaps (up to down).
- Current State—State when the last add or drop event occurred for that link in the interval.
- Last Change Time—Last time that the link had an add or drop event.
- Area—Area in which the link belongs.

IPv4 Prefix Flaps

For each IPv4 prefix, this report lists the number of times that the prefix went from up to down.

This report includes the following columns:

- Prefix—Specific prefix.
- Router—Name or IP address of the router announcing the prefix.
- Count—Total number of flaps (up to down).

- Current State—State when the last add or drop event occurred for that prefix in the interval.
- Last Change Time—Last time that the prefix had an add or drop event.
- Area—Area in which the router belongs.

IPv6 Prefix Flaps

For each IPv6 prefix, the this report lists the number of times that the prefix went from up to down.

This report includes the following columns:

- Prefix—Address of the prefix.
- Router—Individual router involved in the flap event, if there is a specific router.
- Count—Total number of flaps (up to down).
- Current State—State when the last add or drop event occurred for that prefix in the interval.
- Last Change Time—Last time that the prefix had an add or drop event.
- Area—Area in which the router belongs.

Routers

The Routers change report covers changes that occurred between the start and end times. It tracks the following attributes:

- Name
- Type
- Address Family
- IPv4 Address (IS-IS only)
- IPv6 Address (if applicable)
- Model (EIGRP only)
- Software Version (EIGRP only)
- Serial Num (EIGRP only)
- State

This report includes the following columns:

- Router—Name or IP address of the router.
- Attributes—Attributes, if they have changed between the start and end times:
- IPv4 Prefixes—Number of IPv4 prefixes that the router was announcing at the start and end times. The attributes considered here are the same as described for the IPv4 Prefixes reports (see IPv4 Prefixes on page 477).
- IPv6 Prefixes—Number of IPv6 prefixes that the router was announcing at the start and end times. The attributes considered here are the same as described for the IPv6 Prefixes reports (see IPv6 Prefixes on page 477).
- Neighbors—Links to the neighbors at the start and end times. The attributes considered here are the same as described for the Links reports, below.

Links

For OSPF and EIGRP, this report compares interfaces (a single link can have multiple interfaces). For IS-IS there are no interfaces unless TE is enabled, so comparing interfaces is usually equivalent to comparing links. The following attributes are tracked:

- Interface Name
- Addresses
- Bandwidth
- State
- Delay (EIGRP only)
- Borrowed Address (EIGRP only)
- Metric (IS-IS only and OSPF)

This report includes the following columns:

- Link—Specified link.
- Interface—Link interface (as included in the List of all Links window, see Links List on page 109).
- Before—State of the attributes at the start time.
- After—State of the attributes at the end time.

- Area—Link area.

IPv4 Prefixes

The IPv4 Prefixes change report covers IPv4 prefix changes that occurred between the start and end times. The following attributes are tracked:

- Metric (for EIGRP, this is derived using inverse bandwidth and delay)
- Metric Type (for EIGRP, this is derived using the metric type, name, and whether it the metric is an SPF metric)
- Forwarding Address (only for OSPF)
- State

This report includes the following columns:

- Prefix—IPv4 prefix.
- Router—Router involved in the event.
- Before—State at the start time.
- After—State at the end time.
- Area—Prefix area.

IPv6 Prefixes

The IPv6 Prefixes change report covers IPv6 prefix changes that occurred between the start and end times. The following attributes are tracked:

- Metric (for EIGRP, this is derived using inverse bandwidth and delay)
- Metric Type (for EIGRP, this is derived using the metric type, name, and whether it the metric is an SPF metric)
- Forwarding Address (only for OSPF)
- State

This report includes the following columns:

- Prefix—IPv6 prefix.
- Router—Router involved in the event.
- Before—State at the start time.

- After—State at the end time.
- Area—Prefix area.

11 MPLS WAN

The MPLS WAN feature allows Route Analytics Management Software to support enterprises that have multiple sites that are connected by a WAN through ISPs that use MPLS within their own networks. This chapter describes how to configure the MPLS WAN feature.

Chapter contents:

The information in this chapter applies only to units that have licenses for both the BGP protocol and the MPLS WAN protocol.

Understanding MPLS WAN

MPLS allows ISPs to support large numbers of VPNs. Although Route Analytics Management Software appliances do not have visibility into the routing structure within the ISP network, it is still possible to display and analyze routing topologies that extend across the WAN.

The MPLS WAN feature is important if your company has multiple sites that are connected by a Layer 3 VPN. Each of your sites will typically have one or more CE routers that are connected to the ISP's PE routers. The ISP handles all the routing (including BGP), as well as the VPN tunneling through its own network.

With the MPLS WAN feature, you can use Route Analytics Management Software to monitor all of the sites and provide reachability and enterprise connectivity information. The complete topology view shows each site with complete information up to the provider's PE routers.

Although detailed routing through the ISP is not available, the appliance can indicate whether there is connectivity between the ISP's PE routers. When one of your sites advertises prefixes, you can determine whether the ISP is correctly passing all the routing prefixes (not dropping any or sending additional prefixes).

Routing Between Sites

The appliance can support BGP or static routing between the CE and PE routers. If static routing is used, the CE routers inject a default route into the rest of the site, and the PE routers are preconfigured with the prefixes for the connected sites.

The following guidelines apply when you set up sites for the MPLS WAN feature:

- Sites are specified as part of the recording configuration when the recording hierarchy is set up. Sites cannot overlap or be nested.
- When BGP is used across the PE/CE link, it is necessary to have a BGP recorder that is configured with internal BGP (IBGP) to each CE router in the site. This approach allows the system to learn which routes are being announced and received at the CE router.
- The PE/CE links that use static routes are supported by configuring a Static recorder instance, usually at the top level of the hierarchy (not associated with an individual site). The Static recorder collects static routes from all the CE routers that use static routes to the VPN. You can specify each router individually or specify the appropriate prefix. In addition to recording Static/Collector on the CE, you must also record IGP on the CE router (and other routers in the sites, if there are any).

It is possible to discover static routes from all routers by entering 0.0.0.0/0; however, this is not recommended for large networks) due to the high resource cost of doing so.

Satellite Sites

A satellite site is a location, such as a small retail store, with a single prefix and with no topology recording (no IGP). Satellite sites usually use static routing into the VPN. Although there is no visibility into satellite sites, it is important to know whether they are reachable. (Because there may be too many of these sites to record topologies, the system determines only basic reachability.)

At satellite sites, routing information is not recorded and the system can infer reachability only by routing advertisements that come from other sites.

Satellite sites do not show up on the map, but are shown in reachability reports. For example, consider a retail chain that has many small stores, each with one assigned prefix. Recording takes place at major sites (corporate headquarters, regional headquarters) but not at the individual stores. The system can infer reachability of an individual store based on information about the store prefix that is obtained from other sites. The system can do this even though the individual store is outside the monitored network (is not on the routing topology map and has no recorded information).

➤ We recommend that you add all of the prefixes for your satellite site to the prefix group MPLS WAN Satellite Sites.

MPLS WAN and the Routing Topology Map

Figure 9 shows an example network for which MPLS WAN can provide an comprehensive network view. The network includes multiple sites, each running OSPF and BGP. Prior to configuring MPLS WAN, the routing topology map displays the sites as disconnected islands, as shown in the figure.

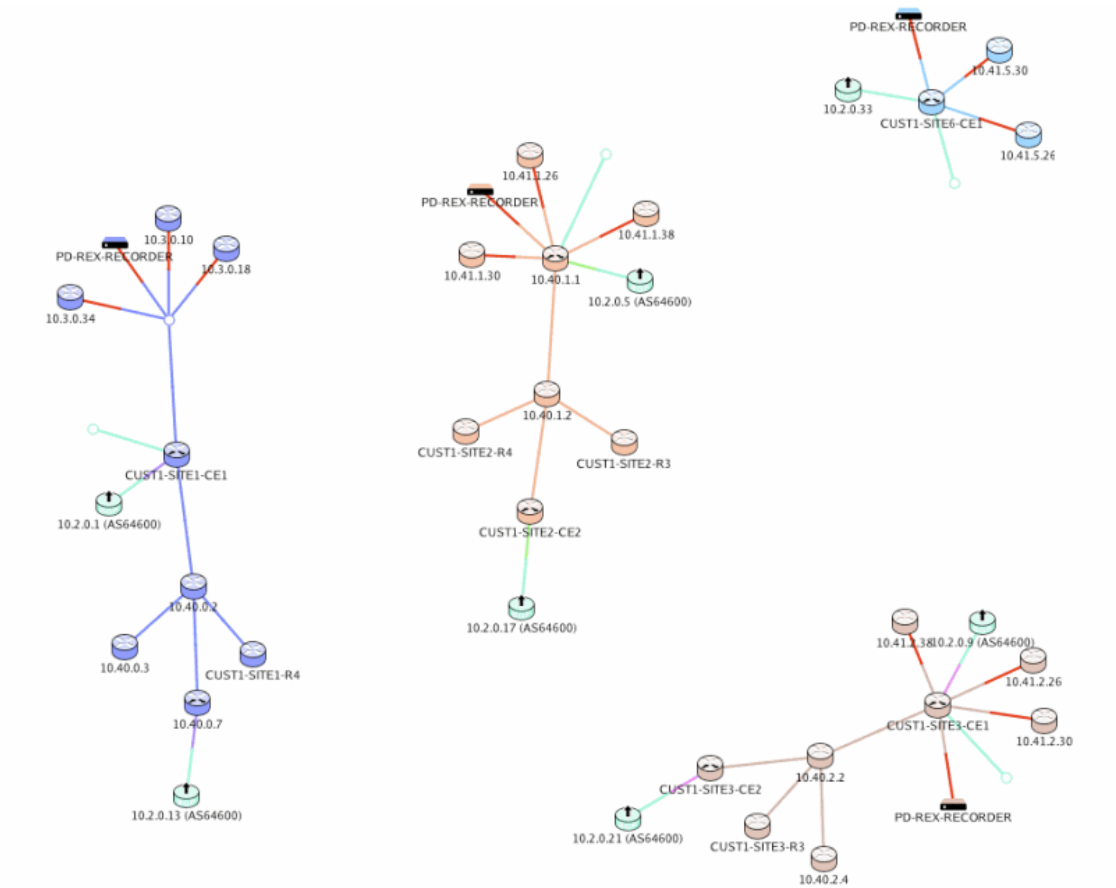

Figure 9 Multi-Site Network

The VPN Connection Configuration table allows you to specify connections between sites and VPNs. This table shows the candidate connections between CE and PE routers and indicates which connections are currently configured as VPN connections (Figure 10). These connections are learned through the routing topologies, specifically the next hop routers seen in BGP or static routes.

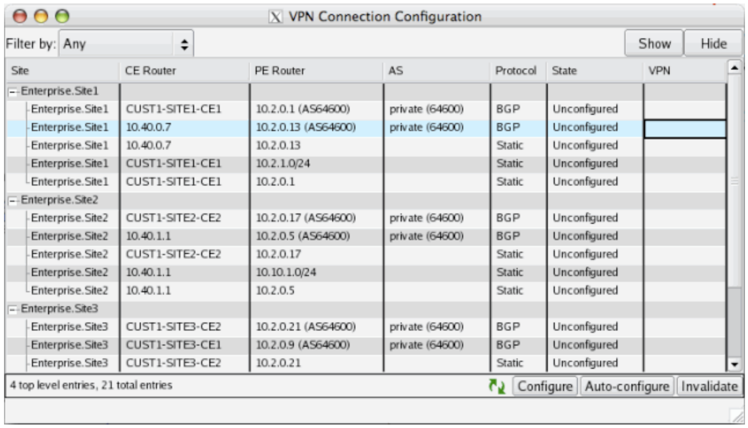

Figure 10 VPN Connection Configuration Before Auto-Configuration

If BGP is used between the CE and PE routers, an auto-configuration option is available. If you choose this option, the system creates a VPN for each set of connections that have a common AS number. Each VPN name is of the form

AS_*number*. Following auto-configuration, the VPN for the connected sites is shown in the table, and the VPN is represented by a cloud on the routing topology map (Figure 11).

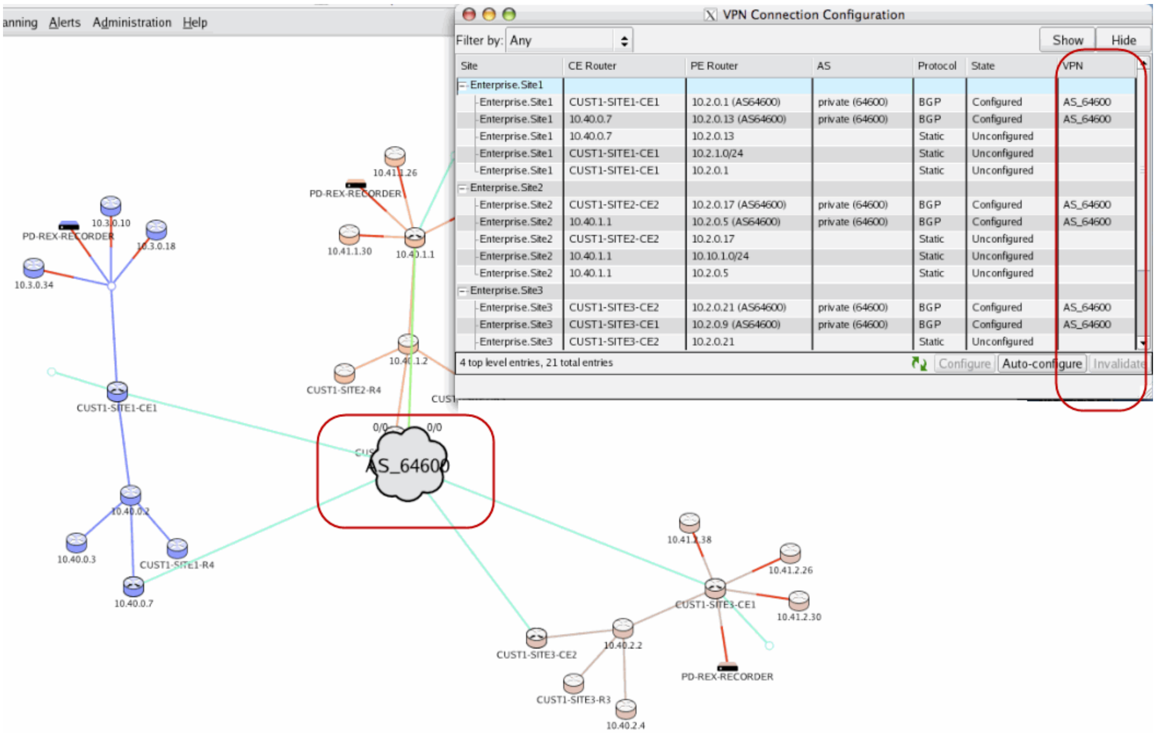

Figure 11 VPN Connection Configuration After Auto-Configuration

After the VPN is set up, you can find paths between selected routers at different sites. The complete path is shown at the site extending through the CE router to the PE routers. If you double-click the VPN cloud, it opens to show dotted line paths between the ISP's PE routers (Figure 11). The dotted lines indicate that a path exists between the PE routers, although the exact details (intermediate hops through the ISP's network) are not known.

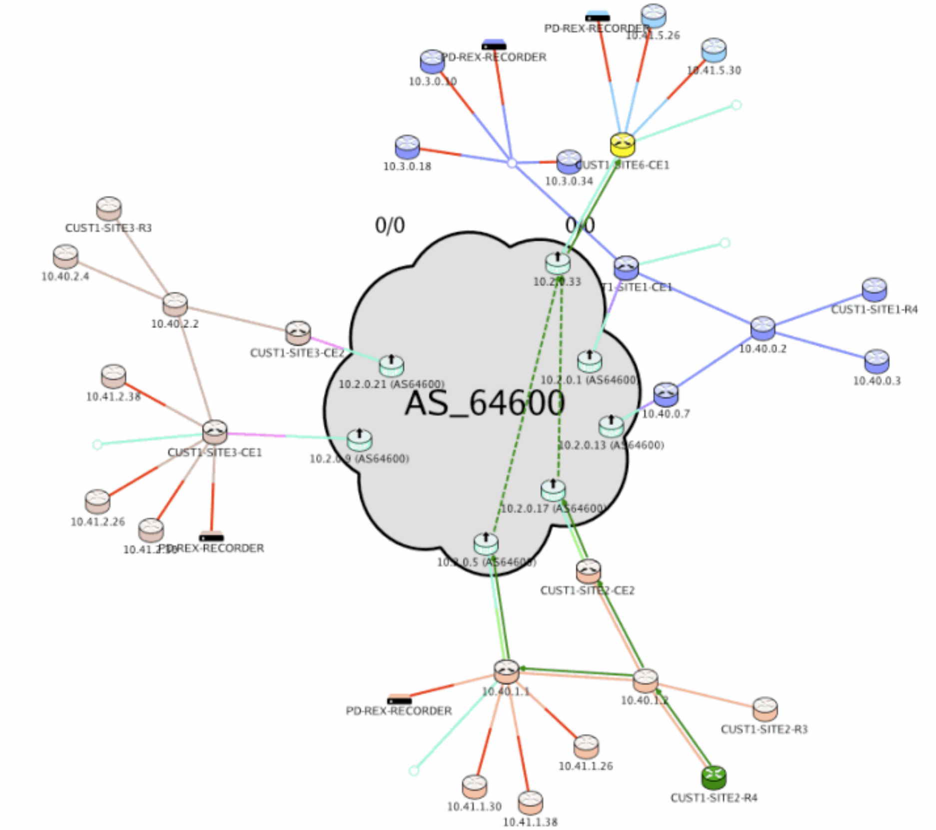

Figure 12 Paths Connecting Multiple Sites

Configuring MPLS WAN

Several tasks are required to set up the MPLS WAN feature:

1 Open the web interface and verify that your system is licensed for the MPLS WAN feature. See the "Configuration and Management" chapter in the *HP Route Analytics Management Software Administrator Guide*.

MPLS WAN functionality is visible in the web and client interfaces only if the system has a valid license for the feature.

If the system has a valid MPLS WAN license, then the Show BGP Next Hops and Show Static Next Hops fields, which are required to see the PE routers and PE/CE paths on the routing topology map, are automatically enabled and not displayed in the Options window (Administration menu). For more on these fields, see Map on page 76.

2 Use the web interface to set up sites. See Setting Up Sites on page 487.

3 Start recording, as described in the “Configuration and Management” chapter in the *HP Route Analytics Management Software Administrator Guide*.

4 For static links only, use the client interface to identify any unmatched interfaces. See Configuring Unmatched Interfaces on page 489.

5 Verify the desired reachability ranges for the sites. See Modifying Reachability Ranges on page 491.

6 Use the client interface to group the CE/PE connections into VPNs. See Setting Up the VPN Connection Configuration on page 491.

7 Use the client interface to set up static VPN connections. See Setting Up Static VPN Connections on page 493.

8 Use the client interface to specify expected announced and received prefixes for the sites. See Identifying Expected Prefixes on page 494.

Setting Up Sites

Each site within a multi-site topology is an administrative domain. When you configure the recorder, you can optionally specify that the administrative domain is a site. The site can be at any level of the recording hierarchy.

The following guidelines apply to site configuration:

- Sites cannot overlap or be nested.
- If you create an administrative domain inside an administrative domain that is already configured as a site, you cannot select the MPLS WAN option—it is not visible on the Recorder Configuration page.
- You can set up a single MPLS WAN site with multiple administrative domains; for example if you have a large site with different OSPF instances that you want to record individually. Because each administrative domain can use only a single protocol, it may be necessary or desirable to set up multiple domains within a site.

- You can set up an administrative domain that contains multiple sites. For example, you can create an organization of sites within different geographic regions. The system allows flexible set up, provided that the sites are not nested.
- You can configure recorder instances inside sites, as appropriate.
- Traffic recorder instances can be placed outside of the sites. Specifically, we recommend that you set up a Static recorder instance outside of all of the sites.
- In addition to recording Static/Collector on the CE, you must also record IGP on the CE router (and other routers in the sites, if there are any).

To specify the administrative domain for MPLS WAN, perform the following steps:

1. Open the web application.
2. Choose **Recorder Configuration**.
3. Click the top-level domain name.
4. Move the cursor to **Add**.

   The option to add another level of domain name hierarchy appears.
5. Click **Administrative Domain**.
6. Click **Add Domain**.
7. Select the check box for MPLS WAN. Selecting this check box identifies the administrative domain as a site.

   The domain name you just entered (for example, IGPRecorders) appears in the hierarchy.

It is not possible to edit an administrative domain; therefore, you cannot specify the domain as a site after saving the configuration.

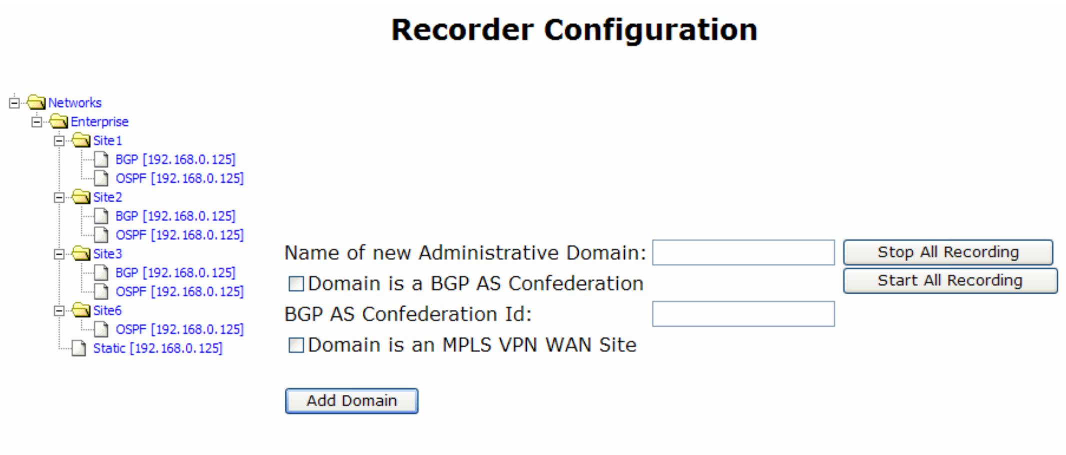

Figure 13 Setting Up the Administrative Domain for MPLS WAN

Configuring Unmatched Interfaces

For static links that start on a CE router, it may be necessary to identify the address of the corresponding PE router. If some interfaces are not identified or are incorrect, use the Unmatched Interfaces window to add or modify the PE addresses.

For point-to-point (serial) links for which the address at one end is known, the system can determine the other end automatically.

To set up unmatched interfaces, perform the following steps:

1 Choose **Administration > MPLS WAN > Unmatched Serial Interfaces**.

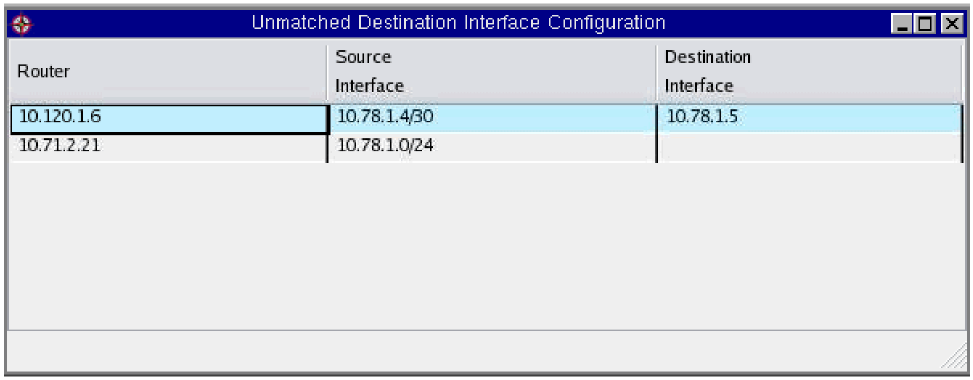

Figure 14 Unmatched Destination Interface Configuration

2 Double click an entry in the Destination Interface column. The entry changes to edit mode, allowing you to modify the addresses or prefixes of any connections that are not correctly configured.

Modifying Reachability Ranges

Reachability ranges are used in the MPLS WAN reachability reports to indicate how well the VPN is passing prefixes from one site to another. The system is pre-configured with a set of reachability ranges. Use the Configure Reachability Ranges window if you want to change the ranges for your company.

To modify reachability ranges, perform the following steps:

1 Choose **Administration > MPLS WAN > Reachability Ranges**.

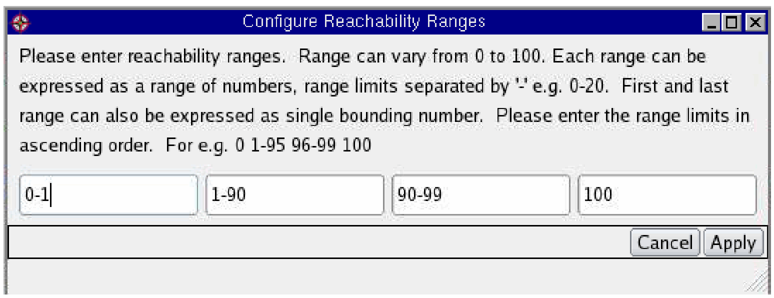

Figure 15 MPLS WAN Reachability Ranges

2 Modify the ranges as needed, specifying the ranges as *number-number* or as a single number.

3 Click **Apply**.

Setting Up the VPN Connection Configuration

Use the VPN Connection Configuration window to group the CE/PE connections into VPNs. Each line in the table is a potential connection into a VPN. It is recommended that you initially use the auto-configure option, and then modify the VPN assignments as needed. Auto-configuration takes all connections with the same AS number and places them in a VPN.

To set up the VPN connection configuration, perform the following steps:

1 Choose **Administration > MPLS WAN > VPN Connections**.

2 If you want to change only selected rows in the table, select the rows.

3 Perform any of the following actions:

- Select filtering options at the top of the window, if needed, to display items of interest or hide items not of interest, and click **Show** or **Hide**.
- Choose an auto-configuration option:
  - — To use auto-configuration on all the rows in the table, click **Auto-configure** and choose **All**.
  - — To use auto-configuration on selected rows in the table, select the rows, click Auto-configure and choose **Selected**.

  The system may warn you that auto-configuration may override the existing configuration. Click **OK** if that is acceptable.

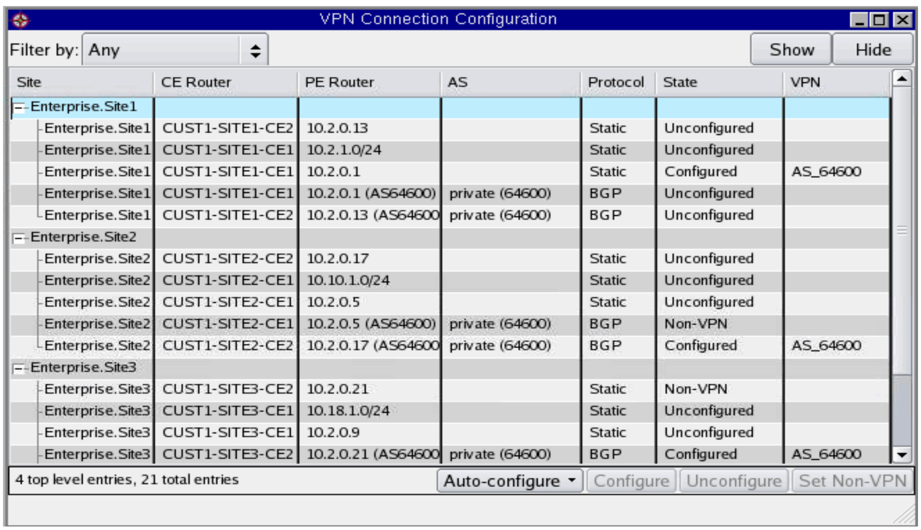

Figure 16 VPN Connection Configuration Window

- Make any manual adjustments to the VPN entries (including static links, which have no AS number), select the rows and click **Configure**. Enter the VPN name in the pop-up window and click **Apply**. For a single row, you enter or select the VPN name directly in the VPN column. The system creates the VPN and shows it on the routing topology map.
- To remove the VPN assignment, select the rows and click **Unconfigure**.

- To keep an entry in the table but remove it from a VPN, select the entry and click **Set Non-VPN**. For example, if there is currently a VPN connection for a CE/PE link, but you do not want that connection to be part of the VPN, use the Set Non-VPN option.

The unconfigure and set non-VPN options both remove the VPN configuration for a VPN connection. We recommend that you use the non-VPN option if the connection will definitely not be part of a VPN. If you filter the table to hide all of the non-VPN entries, you can focus on the VPN and potential VPN connections.

Setting Up Static VPN Connections

On static PE/CE links, the system polls the CE to determine the contents of the routing table. The system is aware of all the static routes that are used to reach other sites through the VPN. But in addition, the PE router must advertise routes to the rest of the VPN on behalf of the CE router.

Because the CE router may not have this information and the Route Analytics Management Software system does not typically have access to the PE router, you must specify routes manually using the Static VPN Connection Configuration window.

To set up the static connections, create a prefix group that includes all the prefixes at the site, and then select the prefix group on the Static VPN Connection Configuration window. This specifies the prefixes that the PE is injecting into the ISP network. This does not alter the routing topology map.

To set up the static VPN connections, perform the following steps:

1 Choose **Administration > MPLS WAN > Static VPN Connections**.

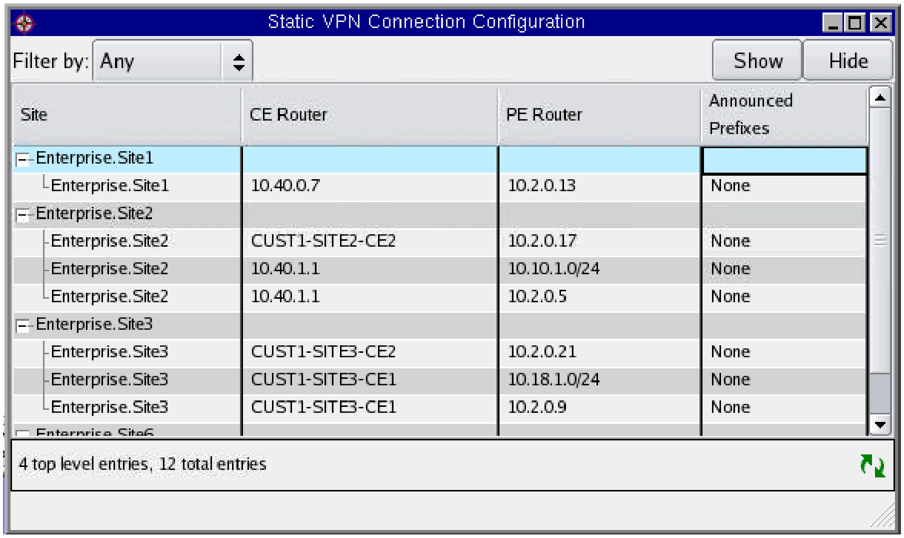

Figure 17 Static VPN Connection Configuration

2 Select filtering options at the top of the window, if needed, to display items of interest or hide items not of interest, and click **Show** or **Hide**.

3 Double-click an entry in the Announced Prefixes column to display select arrows, and then click the arrows to select from the available prefix groups.

Identifying Expected Prefixes

Many of the MPLS WAN reachability reports (MPLS WAN Reachability Reports on page 497) use sets of IPv4 prefixes called “expected prefixes.” These are the IPv4 prefixes that each site announces to, or receives from, each VPN to which that site is connected. Each site has at least two expected prefix sets, one for prefixes it is expected to announce to each attached VPN and another for prefixes it is expected to receive from each attached VPN.

The expected prefixes sets represent the normal, or expected, operating state of the network. Deviations from the expected state (such as unexpected prefixes appearing or expected prefixes withdrawn) may indicate a situation that requires attention and are among the items that are reported in the MPLS WAN reachability reports.

The appliance maintains these sets of expected prefixes, although they are manually configured. Several tools are available for configuring the expected prefix sets for the various sites.

Follow this process to determine expected prefixes:

- Take the current set of prefixes that are announced and received and use that as a snapshot to create the expected set.
- If you discover that some prefix assignments are not correct, make changes using the reachability reports (fine tuning).

To identify expected prefixes, perform the following steps:

1 Choose **Administration > MPLS WAN > Expected Prefixes**.

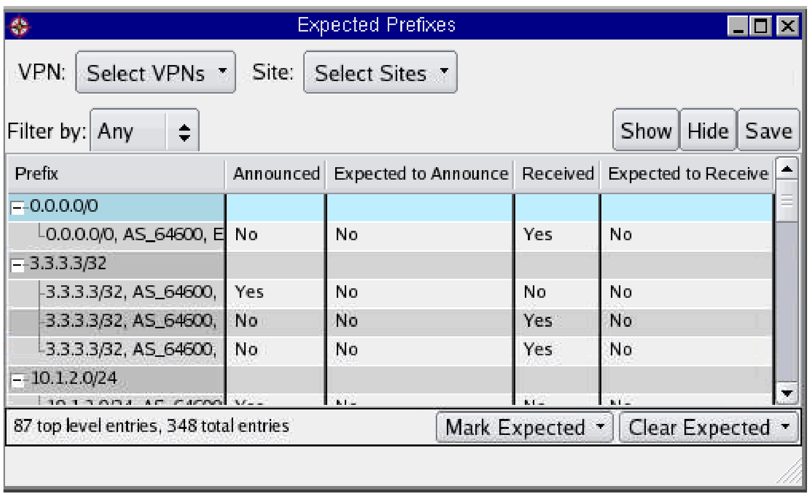

Figure 18 Expected Prefixes

2 Select filtering options at the top of the window, if needed, to display items of interest or hide items not of interest, and click **Show** or **Hide**.

3 To mark entries as expected, select the desired rows, click **Mark Expected,** and choose from the following options.

- **Active Announced**—Uses the actively announced prefixes as the expected announced prefixes.
- **Selected Announced**—Uses the selected prefixes as the expected announced prefixes.

- **Active Received**—Uses the actively received prefixes as the expected received prefixes.
- **Selected Received**—Uses the selected prefixes as the expected received prefixes.
- **All Active**—Takes a snapshot of the prefix state announced by each site to each connected VPN and the prefixes received by each site from each connected VPN, and uses those prefixes as the expected prefix sets. In new deployments this allows you to set up expected prefix sets quickly. You can then use the other Mark and Clear actions to perform fine-grain adjustments for cases in which the current state of the network is not the baseline or expected state. Some of the Reachability Reports also support the ability to mark and clear prefixes from expected prefix sets. See MPLS WAN Reachability Reports on page 497.

4. To clear expected entries, select the desired rows, click **Clear Expected**, and choose from the following options:
   - **All Announced**—Removes all of the announced prefixes from the expected announced prefixes list.
   - **Selected Announced**—Removes all of the selected prefixes from the expected announced prefixes list.
   - **All Received**—Removes all of the received prefixes from the expected received prefixes list.
   - **Selected Received**—Removes all of the selected prefixes from the expected received prefixes list.
   - **All**—Removes all prefixes from both the expected announced and expected received prefixes lists.

Understanding MPLS WAN Routing Reports

The MPLS WAN routing reports can help determine whether a site is announcing and receiving all the expected traffic.

Routing reports include drill-down options that provide additional information about reachability in the MPLS WAN context. For example, you can list the prefixes announced and received or how may prefixes are expected. To drill down for additional information, right-click an entry in the list and choose an option, or click the Drill-Down arrow above the table.

The settings in the Expected Prefixes window under the Administration menu determine the expected behavior of the prefixes.

MPLS WAN Reachability Reports

This section describes the following MPLS WAN reachability reports:

- Reachability from Other Sites on page 498
- Reachability to Other Sites on page 499
- Reachability by VPN on page 500
- Reachability to Satellite Sites on page 501

You can assign names to prefixes to more easily identify them in MPLS VAN WAN reachability reports. See Assigning IPv4 or IPv6 Prefix Names on page 107.

Reachability from Other Sites

Choose **Reports > Routing Reports > MPLS VPN WAN Reachability > Reachability from Other Sites** in the client application to open this report (Figure 19).

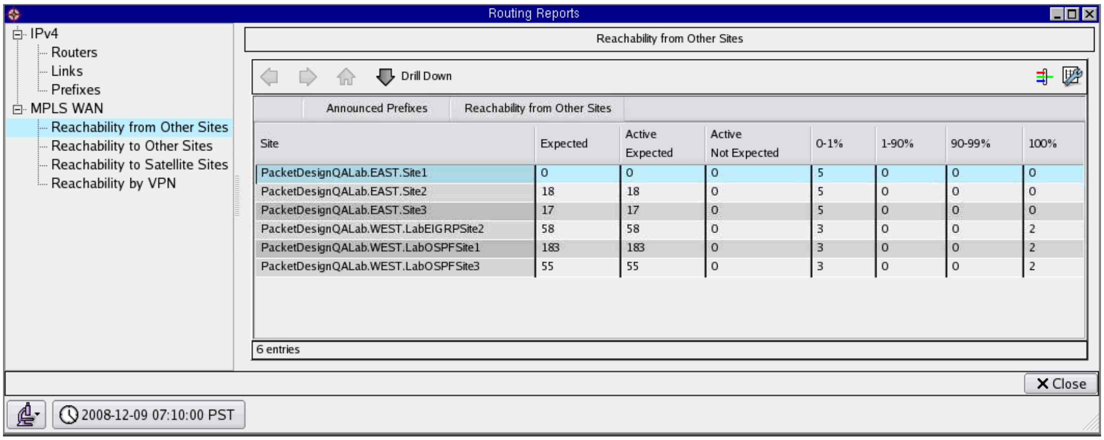

Figure 19 MPLS WAN - Reachability from Other Sites

The report table contains the following information:

- **Name**—Site name.
- **Expected**—Number of prefixes that the user has configured as expected for this site to announce (see Identifying Expected Prefixes on page 494).
- **Active Expected**—Subset of the expected announced prefixes that are actually being announced.
- **Active Not Expected**—Number of prefixes that are being announced but are not in the expected prefixes list.
- **Reachability from Other Sites**—Indication of how well the VPN is passing prefixes from this site to the other sites. Each site that receives announced prefixes from this site then has reachability to this site measured by the percentage of the announced prefixes that each site receives. Each table cell on the row shows the count of sites having a reachability percentage to this site that falls within the range indicated for the column. Some examples of reachability are as follows:

— The value 5 in the 0% column means that 5 sites are not receiving any of the expected prefixes that are announced by the specified site.

— If all values are in the 100% column, this indicates a mesh network in which all sites announce to all other sites.

— If only some sites are show all values in the 100% column, this may indicate a hub and spoke network, in which only the hubs are able to announce to all sites.

See Modifying Reachability Ranges on page 491 for information on modifying the reachability ranges on this report.

Reachability to Other Sites

Choose **Reports > Routing Reports > MPLS VPN WAN Reachability > Reachability to Other Sites** in the client application to open this report (Figure 20).

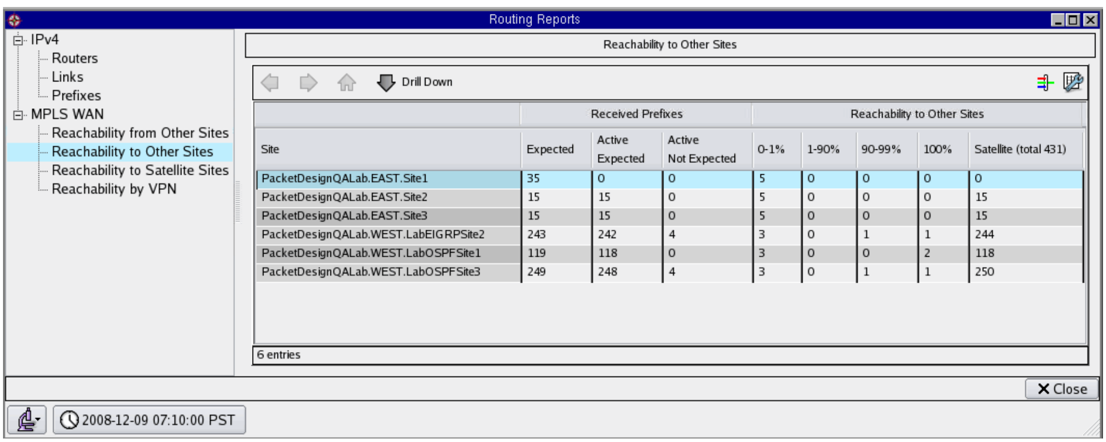

Figure 20 MPLS WAN - Reachability to Other Sites

The report table contains the following information:

- **Name**—Site name.
- **Expected**—Number of prefixes that the user has configured as expected for this site to receive (see “Identifying Expected Prefixes” on page 9-11).
- **Active Expected**—Subset of the expected received prefixes that are actually being received.

- **Active Not Expected**—Number of prefixes that are being received but are not in the expected prefixes list.
- **Reachability to Sites**—Indication of how well the VPN is passing prefixes from other sites to this site. This site has reachability to another site as measured by the percentage of the prefixes announced by the other site that it receives. Each table cell on the row shows the count of sites from which this site is receiving a percentage of the announced prefixes that falls within the range indicated for the column.

See Modifying Reachability Ranges on page 491 for information on modifying the reachability ranges on this report.

Reachability by VPN

Choose **Reports > Routing Reports > MPLS VPN WAN Reachability > Reachability by VPN** in the client application to open this report (Figure 21).

The report lists reachability information for each VPN.

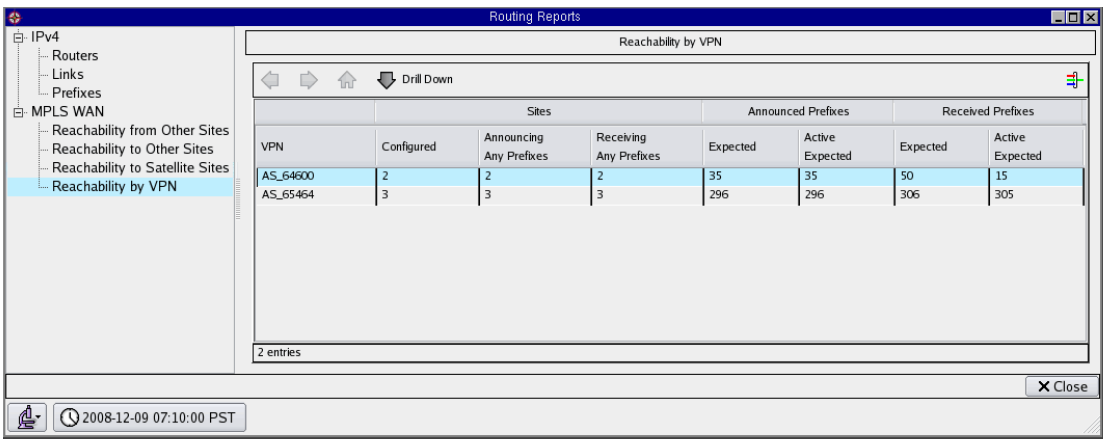

Figure 21 MPLS WAN - Reachability by VPN

The report table contains the following information:

- **Name**—VPN name.
- **Sites**—Number of sites that are included in the VPN.
- **Expected Prefixes Announced**—Total number of expected prefixes that are being announced for all sites in the VPN.

- **Expected Prefixes Received**—Total number of expected prefixes that are being received by at least one site in the VPN.
- **Sites Announcing 0% Expected Prefixes**—Number of sites in the VPN that are not announcing any expected prefixes.
- **Sites Receiving 0% Expected Prefixes**—Number of sites in the VPN that are not receiving any expected prefixes.

Reachability to Satellite Sites

Choose **Reports > Routing Reports > MPLS VPN WAN Reachability > Reachability to Satellite Sites** in the client application to open this report (Figure 22).

The report lists each site, the site prefix, and the number of other sites that is able to reach it.

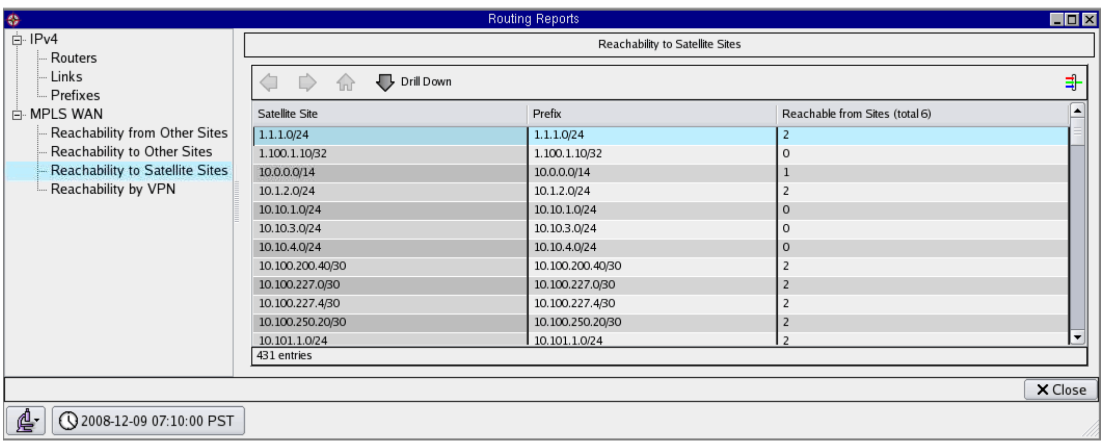

Figure 22 MPLS WAN - Reachability to Satellite Sites

The report table contains the following information:

- **Name**—Prefix name.
- **Prefix**—Associated prefix.
- **Reachable from Sites**—Number of sites from which this prefix is reachable.

12 Alerts

This chapter describes how to configure and view alerts.

Chapter contents:

Understanding Alerts

Alerts allow you to monitor network activity and obtain information about potential problems. You can obtain notification of any changes to network elements (such as routes and routers) based on configurable thresholds. The system can send alerts as SNMP traps to an SNMP-based network management system integrated with other network operations, logged to a syslog file, sent as a simple email notification, or logged to the database for display in the GUI.

You can configure alerts for selected protocols. You can also determine whether the alerts will operate globally or only in selected areas. For example, if you are monitoring a single or multi-area OSPF network, you can configure alerts that apply to the OSPF routing events for all the OSPF areas being monitored or that apply only to the OSPF routing events in selected areas. Alerts are disabled by default.

From the Alerts menu, select any of the following items:

- **View Alert Status**—List all alerts generated by the system. See Viewing Alert Types on page 505.
- **Configure Alerts—**Set up alerts and view the list of configured alerts. See Creating New Alerts on page 508.
- **Dispatch Specifications**—Specify the notification method for network and traffic status. See Creating Dispatch Specifications on page 525.
- **Suppression Specifications**—Identify periods when you do not want to generate any alerts, and set a rate limit on delivering alerts. See Creating Suppression Specifications on page 529.

All users can view alerts; however to configure alerts you must have administrator privileges.

The descriptions within this chapter are based on a full product license and a network configuration using all protocols.

Viewing Alert Types

You must have administrator privileges to configure alerts.

Because the privilege level of the GUI in VNC display 1 is **operator**, it is not possible to configure alerts using VNC display 1. Use one of the VNC displays 2-10 or an SSH/X Window connection to run the GUI when configuring alerts.

To view alert types, perform the following steps:

1 Choose **Alerts > Configure Alerts** to open the Configure Alerts window.

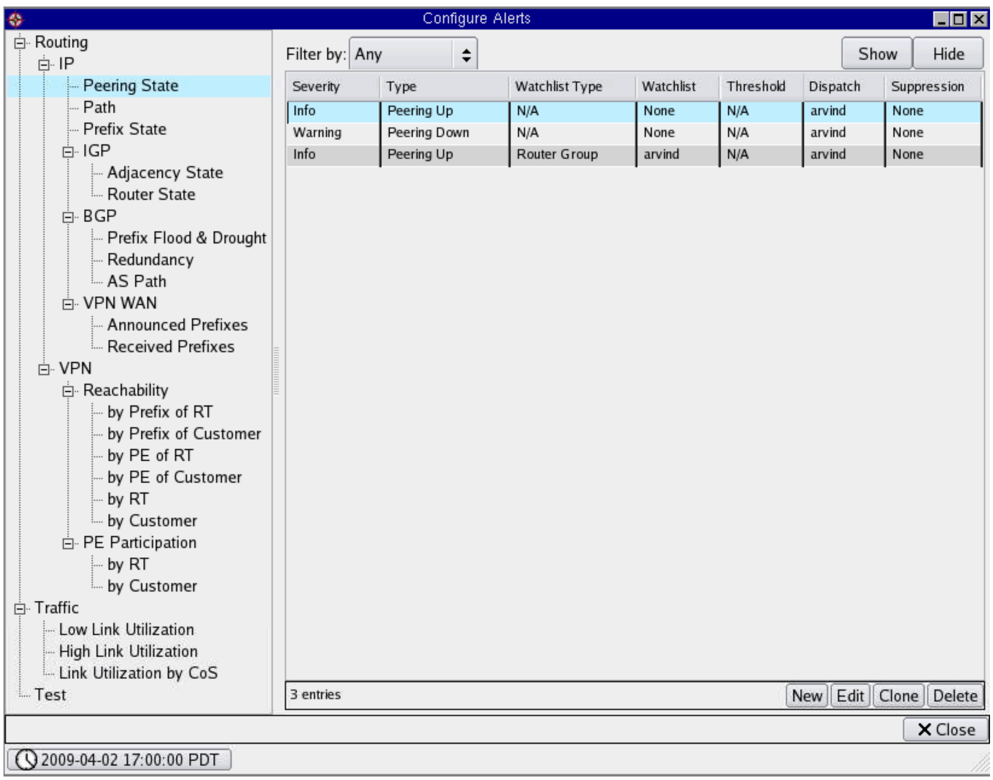

Figure 1 Configure Alerts Window

2 Choose any of the supported alert types from the side menu to display information. See Table 1 .

Table 1 Alert Types

Alert Type	**Description**
IP Alerts	
Peering State	Fired when a peering comes up, goes down or flaps,[a] depending upon the configuration selection. In this case, the term "peering" refers only to the connection between the route recorder and its neighbor routers. For BGP routers, it is only possible to infer whether a peering exists based on prefix announcements, so no alert is provided since the state is not known with sufficient accuracy. For IGP routers, the adjacency state change alert provides the needed information.
Path	Fired when a path between a router and a prefix changes.
Prefix State	Fired when a route comes up, goes down, or flaps, depending upon the configuration selection.
Adjacency State	Fired when a adjacency comes up, goes down or flaps, depending upon the configuration selection.
Router State	Fired when all incoming links to a router go down or are overloaded (IS-IS only), when a router that was previously isolated is no longer isolated, or when the router flaps.
Prefix Flood & Drought	Fired when the number of prefixes heard from a BGP peer changes significantly.
Redundancy	Fired when the redundancy (number of next hops) for a prefix in BGP goes above or below the configured threshold.
AS Path	Fired when the AS path of any route in BGP changes.
VPN WAN Alerts	
Announced Prefixes	Fired when prefixes are announced.
Received Prefixes	Fired when prefixes are received.

Table 1 Alert Types (cont'd)

Alert Type	Description
VPN Alerts	
Reachability by Prefix of RT	Fired when a prefix comes up or goes down in a VPN RT.
Reachability by Prefix of Customer	Fired when a prefix comes up or goes down in a customer.
Reachability by PE of RT	Fired when number or percentage of routes gained/lost in an RT from a PE exceeds a threshold.
Reachability by PE of Customer	Fired when number or percentage of routes gained/lost in a customer from a PE exceeds a threshold.
Reachability by RT	Fired when number or percentage of routes gained/lost in an RT exceeds a threshold.
Reachability by Customer	Fired when number or percentage of routes gained/lost in a customer exceeds a threshold.
PE Participation by RT	Fired when number or percentage of PEs participating in an RT exceeds a threshold[b].
PE Participation by Customer	Fired when number or percentage of PEs participating in a customer exceeds a threshold.
Traffic Alerts	
Low Link Utilization	Fired when link utilization is below a threshold.
High Link Utilization	Fired when link utilization is above a threshold.
Link Utilization by CoS	Fired when link utilization corresponding to a specified CoS is above a threshold.
Test Alerts	Used to verify that alerts are being sent and that all specified endpoints are able to receive the alerts.

a. Up, down, isolated, and connected events all count as flaps. For example, if a network element goes down and comes back up, the flap count increases by 2.

b. A PE is considered to be participating in a VPN if it announces one or more routes in that VPN (RT or Customer).

The Configure Alerts window lists existing alerts. and provides a snapshot of settings for each alert. Table 2   describes the columns on the page, which may vary according to the alert type.

Table 2 Alert Attributes

Item	Description
Severity	User-assigned importance levels are info, notice, warning, error, or critical.
Type	Type of alert (see Table 14-1).
Watchlist Type	Type of group to which this alert applies. See Creating Groups Using the Menu on page 122 for information on creating groups.
Watchlist	Specific group to which this alert applies. The alert is triggered only if the affected network elements are in this group.
Threshold	Parameter level that causes an alert to fire when it is reached. For flap events, the format is count:duration. Duration is presented in seconds (adjusted as needed if the alert was configured in minutes or hours).
Dispatch	User-defined name for the dispatch specification that occurs when an alert occurs.
Suppression	User-defined name for a set of conditions that allow alerts to ignore a condition until it reaches a specified level.

Creating New Alerts

You can configure custom alert types to match your notification needs if the following conditions are met:

- Only a single top-level database is opened (database are not opened from more than one top-level topology).
- The database that is opened is configured to record (indicated as green in the Open Topology dialog box).
- If multiple top-level database names are configured to record (green), only the first one is opened. (Usually there is only one topology configured to record.)

To create a new alert, perform the following steps:

1 Choose **Alerts > Configure Alerts**.

2 Click the alert type on the side menu and click **New** to display the configuration parameters.

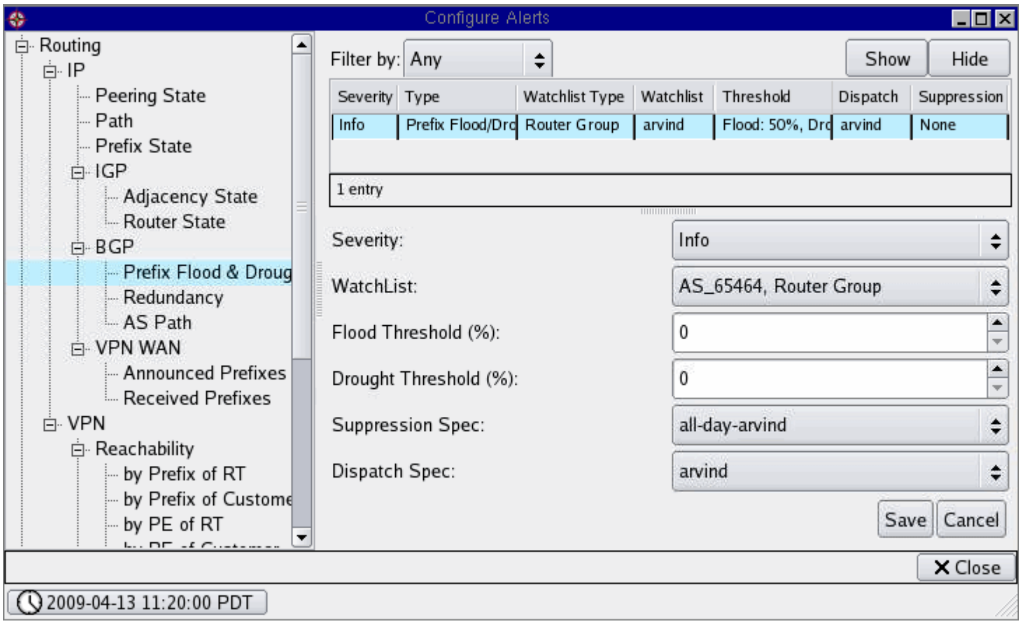

Figure 2 Creating a New Alert

3 Configure settings according to the descriptions in Table 3 and Table 4 .

Table 3 Alert Settings

Item	Description
Severity	User-assigned importance levels: info, notice, warning, or critical.
Watchlist	Group name representing the set of interest. The watchlist is a group, such as a router group or prefix group, depending upon the alert. If you choose an existing watchlist group or none, the field reflects that option. If you choose a new group, the New Group window opens. See Creating Groups Using the Menu on page 122 for information on creating groups. **Note:** The option None is included for some alert types to specify that all objects of the relevant type should be monitored. The option None is not included for alerts where the total number of objects is large enough that monitoring all of them is not practical due to processing or storage requirements. **Note:** Use the RegEx field in the Add Router Group window to restrict the displayed items in the Routers and Child Groups tabs. For example, if you only want IP addresses starting with 192, enter 192 in the Select (RegEx) field, click OK by the field, and then click the arrow to move the selected IP addresses from the Available Items field to the Selected Items field. The syntax of extended regular expressions is explained in Regular Expressions on page 226. The syntax is not the same as shell or file manager pattern patching, so a pattern like *-core-gw is not correct.

Table 3 Alert Settings (cont'd)

Item	Description
Parameters	Parameter that generates an alert if specified conditions are met. Availability of the following parameter types depends upon the type of alert, as listed in Table 4 . •Severity—Indication of how serious the alert is. •Watchlist—Set of routers associated with the alert. •Alert Condition—isolated/connected, gain/loss, up/down, flap. •Flap Count—Number of flaps within the specified duration. These options are available if the Alert Condition is set to **flap**. **Note:** Isolated, connected, up, and down are each counted as a flap. For example, a rule could be set to generate an alert if there is a flap 10 times (isolated and/or connected) within 100 seconds. •Duration—Interval for counting the number of flaps. •Suppression Spec. •Dispatch Spec. •Flood Threshold—Percent change from a baseline number of prefixes indicating a flood of prefixes. •Drought Threshold—Percent change from a baseline number of prefixes indicating a drought of prefixes. •Threshold Type—Absolute/percent. •Threshold—Level that must be reached to generate an alert. •Clear Threshold—Works in conjunction with the main alert threshold to implement hysteresis. An alert is issued when the alert threshold is reached or crossed, but if the parameter oscillates just above and just below this threshold, it is preferable not to generate an alarm each time the threshold is crossed. The level must drop below the clear threshold to enable triggering a new alert if the level subsequently rises above the alert threshold again. •RT—Selected RT to match. •Customer—Selected customer to match. •CoS—Class of service. •Number of next hops—Number of hops for the AS path.

Table 3 Alert Settings (cont'd)

Item	Description
Dispatch Specification	Name of the dispatch specification that controls how the alert should be delivered (email, SNMP, or syslog). See Creating Dispatch Specifications on page 525 for instructions on defining a dispatch specification, Configuring an SNMP Server on page 532 for instructions on setting up SNMP for alerts, and Configuring a Remote Syslog Server on page 533 for instructions on setting up a remote syslog server.
Suppression Specifications	Name of an optional suppression specification to limit the rate at which the alert may be generated or to specify a time interval when the alert should be ignored. See Creating Suppression Specifications on page 529 for instructions on defining a suppression specification.

Table 4 Available Parameters by Alert Type

Item	Description
Routing - IP	
Peering State	•Severity •Watchlist •Alert Condition •Flap Count •Duration •Suppression Spec. •Dispatch Spec.
Path	•Severity •Watchlist •Suppression Spec. •Dispatch Spec.

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
Prefix State	•Severity •Watchlist •Alert Condition •Flap Count •Duration •Suppression Spec. •Dispatch Spec.
Routing - IP - IGP	
Adjacency State	•Severity •Watchlist •Alert Condition •Flap Count •Duration •Suppression Spec. •Dispatch Spec.
Router State	•Severity •Watchlist •Alert Condition •Flap Count •Duration •Suppression Spec. •Dispatch Spec.
Routing - IP - BGP	

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
Prefix Flood and Drought	•Severity •Watchlist •Flood Threshold •Drought Threshold •Suppression Spec. •Dispatch Spec.
Redundancy	•Severity •Watchlist •Number of next hops •Suppression Spec. •Dispatch Spec.
AS Path	•Severity •Watchlist •Suppression Spec. •Dispatch Spec.
Routing - IP - VPN WAN	
Announced Prefixes	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (Pfxs) •Clear Threshold •Suppression Spec. •Dispatch Spec.

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
Received Prefixes	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (Pfxs) •Clear Threshold •Suppression Spec. •Dispatch Spec.
Routing - VPN - Reachability	
by Prefix of RT	•Severity •Watchlist •Alert Condition •RT •Suppression Spec. •Dispatch Spec.
by Prefix of Customer	•Severity •Watchlist •Alert Condition •Customer •Suppression Spec. •Dispatch Spec.

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
by PE of RT	•Severity •Watchlist •Alert Condition •RT •Threshold Type •Threshold (routes) •Clear Threshold •Suppression Spec. •Dispatch Spec.
by PE of Customer	•Severity •Watchlist •Alert Condition •Customer •Threshold Type •Threshold (routes) •Clear Threshold •Suppression Spec. •Dispatch Spec.

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
by RT	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (routes) •Clear Threshold •Suppression Spec. •Dispatch Spec.
by Customer	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (routes) •Clear Threshold •Suppression Spec. •Dispatch Spec.

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
Routing - VPN - PE Participation	
by RT	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (PEs) •Clear Threshold •Suppression Spec. •Dispatch Spec.
by Customer	•Severity •Watchlist •Alert Condition •Threshold Type •Threshold (PEs) •Clear Threshold •Suppression Spec. •Dispatch Spec.
Traffic	

Table 4 Available Parameters by Alert Type (cont'd)

Item	Description
Low Link Utilization	•Severity •Watchlist •Threshold Type •Threshold (kbps) •Clear Threshold •Suppression Spec. •Dispatch Spec.
High Link Utilization	•Severity •Watchlist •Threshold Type •Threshold (kbps) •Clear Threshold •Suppression Spec. •Dispatch Spec.
Link Utilization by CoS	•Severity •Watchlist •Threshold Type •Threshold (kbps) •Clear Threshold •CoS •Suppression Spec. •Dispatch Spec.

4 Click **Save** to create the new alert.

For example, you can define an adjacency flap alert that tracks links state transitions (up, down). Each up or down event counts as one flap, so a link that goes down and comes back up counts as two flaps of that link.

To configure the flap alert, specify the following parameters:

- Flap count
- Flap duration

If you define a flap count of 5 and a flap duration of 3 minutes, the alert is generated if five or more flaps are seen within any three minute interval. Because the alert may be watching multiple links, the flap count is maintained on a per link basis. You can apply the same types of settings to other flap alerts as well (such as prefix flap, peering flap, and router state).

Editing, Cloning, and Deleting Alerts

You can edit and delete alerts, or use the clone option to create a new alert from an existing one.

To edit alerts, perform the following steps:

1. Choose **Alerts > Configure Alerts**.
2. Choose the alert type from the side menu and click **Edit**.

   The alert parameters are shown in a panel at the bottom of the window.
3. Modify settings as described in Table 3 .
4. Click **Save**.

To create a new alert from an existing one, perform the following steps:

1. Choose **Alerts > Configure Alerts**.
2. Choose the alert type from the side menu and click **Clone**.

   The alert parameters for the original alert are displayed.
3. Modify settings as described in Table 3 .
4. Click **Save**.

To delete an alert, perform the following steps:

1 Choose **Alerts > Configure Alerts**.

2 Choose the alert type from the side menu and click **Delete**.

You cannot recover an alert after it is deleted.

Viewing Alert Status

Choose **Alerts > View Alert Status** to open the Alerts report. This report lists the time, severity, network, alert type, and description of the fired alert. You can sort the data in ascending or descending alphanumeric order by clicking any of the column headings. You can also show or hide selected alerts.

Only alerts that are logged to the database (as specified in the Dispatch Specifications window) are displayed in the Alerts report. Any records you delete cannot be recovered.

For networks that include the BGP/MPLS VPN feature, the Network column may occasionally contain an entry in the format Opaque:0xn:0xn rather than an RT number. This type of entry indicates an extended community attribute that is not Route Target or Source of Origin, and thus is not interpreted by the BGP/MPLS VPN protocol. Users who do not want to see alerts for the opaque community strings should set a watchlist (an RT Group) on the appropriate alerts to restrict to the desired RT values. Use the Show RT communities only configuration option to control whether opaque RTs are shown on the routing topology map. See Map on page 76 for a description of the option.

For more information about VPN reports, see Chapter 8, "VPN Routing"

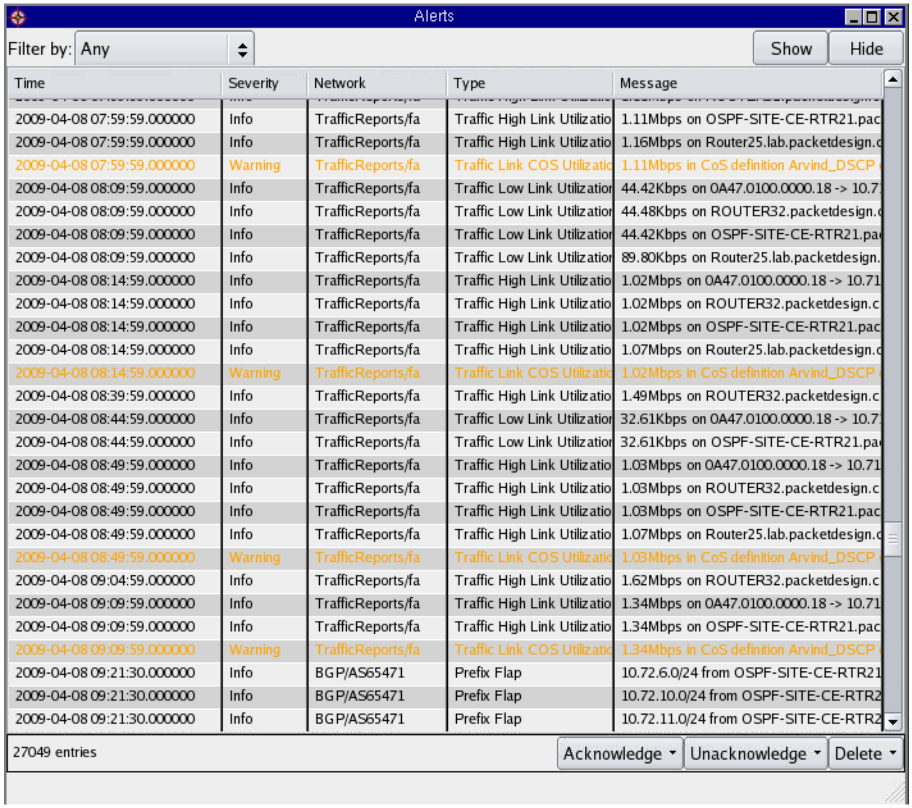

Figure 3 Viewing Alerts Status

Filtering Alerts

You can modify the alert list by using filtering options.

To view specific alerts using the Filter by option, perform the following steps:

1. Choose **Alerts > View Alert Status**.
2. Make a selection from the **Filter by** drop-down list.
   - If you choose **Severity**, select the severity levels you want to include from the check boxes in the **Options** drop-down list.

- If you choose **IP Alerts**, select the types of alerts you want to include from the check boxes in the **Options** drop-down list.
- If you choose **VPN Alerts**, select the types of alerts you want to include from the check boxes in the **Options** drop-down list.
- If you choose **Traffic Alerts**, select the types of alerts you want to include from the check boxes in the **Options** drop-down list.
- If you choose **Network**, enter the network to match, as indicated in the Network column. For the match, you can choose **Substring**, **Exact Match**, or **Begins With**.
- If you choose **Message**, enter the string to match the message name, as indicated in the Message column. For the match, you can choose **Substring**, **Exact Match**, or **Begins With**.
- If you choose **Acknowledgment**, you can choose **Acknowledged** or **Unacknowledged**.
- If you choose **Expression**, you can type a simple expression in the adjacent field. For information on using this type of filter, see Using Filters on page 221.
- If you choose **Advanced**, enter values in the window. For information on using this type of filter, see Using Filters on page 221.

3 Click **Show** to list only those records you want to view. Click **Hide** to have all hidden data reappear.

Acknowledging Alerts

You can acknowledge alerts by marking them in the alert list. When you acknowledge an alert, the entry turns grey. To acknowledge selected alerts, highlight the alerts and choose **Acknowledge > Selected**. To acknowledge all alerts in the list, choose **Acknowledge > All**.

If you inadvertently acknowledge an alert that you have not investigated, you can select it and click **Unacknowledge > Selected** to return the record to its original state.

A right-click menu is also available. Select one or more alerts and right-click to acknowledge, unacknowledge, or delete the alerts. For an individual alert, you can also choose **Move time to here** to change the historical time of the routing topology map to the time of the selected alert. Doing so allows you to view the network conditions at the time that the alert was generated.

Updating Alert Status

When the Alerts table opens, the page scrolls to the end (sorted by time) and starts updating every 30 seconds to add new alerts. A message indicates when the table is updating.

If you select any rows in the table, the update stops so that the current selection does not automatically change. If any new alerts occur while rows are selected, a reload icon appears on the left of the buttons, along with an indication of the last real time when alerts were loaded into the table. Click the reload icon to turn row selection off and continue displaying new alerts when they occur. If the table is sorted ascending by time (the default) and scrolled to the end, then it will stay at the end.

Creating Dispatch Specifications

If you want alerts to be sent using the SNMP or syslog methods, you must configure a server to receive those notifications. Dispatch specifications determine the notification method that is used to inform recipients about network and traffic status.

Before creating dispatch specifications, you must configure the an SNMP server and remote syslog server. See Configuring an SNMP Server on page 532 and Configuring a Remote Syslog Server on page 533.

To define dispatch specifications perform the following steps:

1 Choose **Alerts > Dispatch Specifications**.

2 Click **Add**.

3 Enter a name for the dispatch specification.

4 Select the **Log to DB** checkbox. This allows you to view the alerts in the GUI.

5 Click **OK**.

To create a new SNMP dispatch specification, perform the following steps:

1 Choose **Alerts > Dispatch Specifications**.

2 Click **Add**.

3 Enter a name for the dispatch specification.

4 Click the **SNMP** tab.

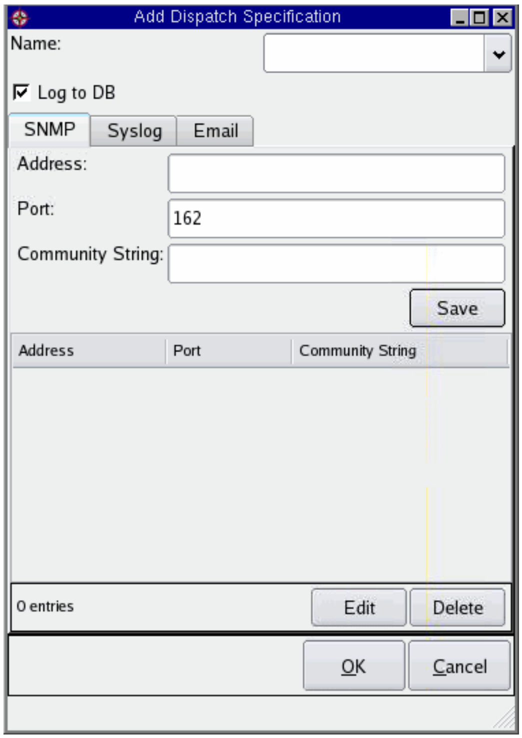

Figure 4 Defining an SNMP Dispatch Specification

5 Specify the address, port, and community string.

6 Click **Save** to add the SNMP information.

7 Click **OK** to save your settings or another tab to create additional specifications.

To create a new Syslog dispatch specification, perform the following steps:

1 Choose **Alerts > Dispatch Specifications**.

2 Click **Add**.

3 Enter a name for the dispatch specification.

4 Click the **SYSLOG** tab.

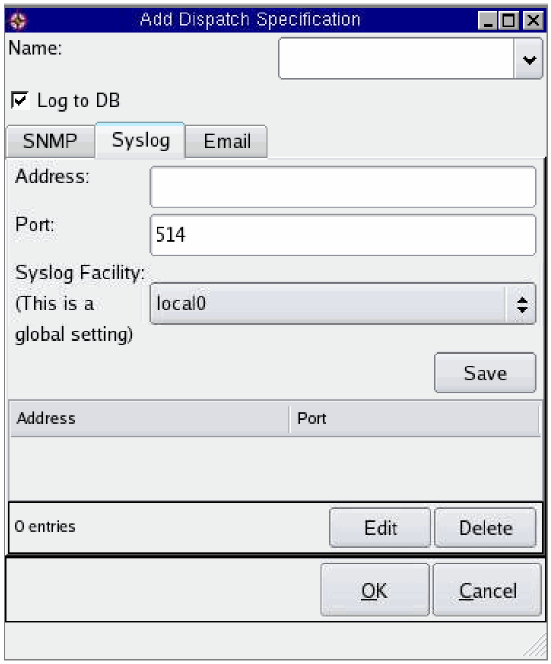

Figure 5 Defining a Syslog Dispatch Specification

5 Specify the address and port number for the syslog.

The SYSLOG Facility field is a global setting; that is, it is used by all alerts that specify a Syslog dispatch.

6 Click **Save** to add Syslog information.

7 You can modify existing Syslog information by selecting the address and clicking **Edit**.

8 Click **OK** to save your settings or another tab to create additional specifications.

To create a new Email dispatch specification, perform the following steps:

1 Choose **Alerts > Dispatch Specifications**.

2 Click **Add**.

3 Enter a name for the dispatch specification.

4 Click the **EMAIL** tab.

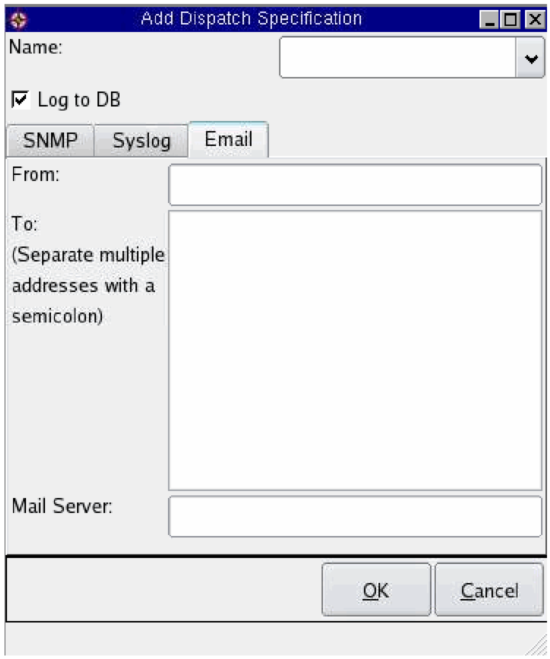

Figure 6 Defining an Email Dispatch Specification

5 Perform any of the following tasks on the Email tab.

- Set the email address that is the source for sending the dispatch notifications.
- Create a list of destination email addresses (separated with semicolons).
- Set the default mail server IP address or the DNS name that is used to send out the email. If left blank, the mail protocol sends email through the default mail server for that network as specified on the Mail page in the administration web pages. See "Administration" in the *HP Route Analytics Management Software Administrator Guide*.

6 Click **OK** to save your settings or another tab to create additional specifications.

Editing, Duplicating, and Deleting Dispatch Specifications

You can edit and delete dispatch specifications. You can also use the duplicate option to create a new dispatch specification from an existing one.

To a edit dispatch specification, perform the following steps:

1. Choose **Alerts > Dispatch Specifications**.
2. Choose the dispatch specification and click **Edit**.
3. Modify settings as described in Creating Dispatch Specifications on page 525.
4. Click **Save**.

To create a new dispatch specification from an existing one, perform the following steps:

1. Choose **Alerts > Dispatch Specifications**.
2. Choose the dispatch specification and click **Duplicate**.
3. Modify settings as described in Creating Dispatch Specifications on page 525.
4. Click **Save**.

To delete a dispatch specification, perform the following steps:

1. Choose **Alerts > Dispatch Specifications**.
2. Choose the dispatch specification and click **Delete**.

You cannot recover an alert after it is deleted.

Creating Suppression Specifications

This section describes how to create new suppression specifications. Suppression specifications determine the periods when no alerts are generated and sets rate limits on delivering alerts.

To define suppression specifications perform the following steps:

1 Choose **Alerts > Suppression Specifications** to open the Alerts Suppression Specifications window. This lists user-defined alert exclusions based on frequency and date/time.

2 Click **Add**.

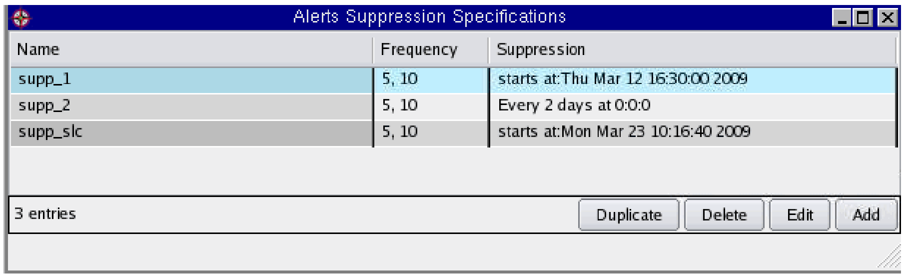

Figure 7 Suppression Specifications

3 Enter a name for the suppression specification.

4 Click the **Schedule** tab.

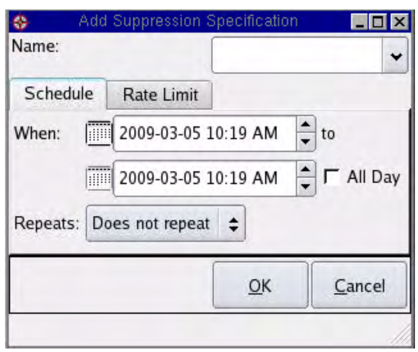

Figure 8 Suppression Specification - Schedule

5 Specify the schedule that determines when alerts are suppressed.

6 For the repeat frequency, you can choose never, daily, weekly, monthly or yearly. If this is a nonrepeating suppression, the remaining repeating frequency options do not appear.

7 Click the **Rate Limit** tab.

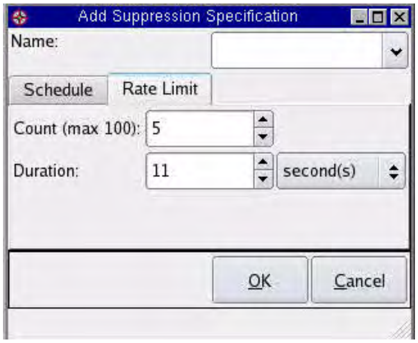

Figure 9 Suppression Specification - Schedule

8 From the Rate Limit tab you can determine the minimum acceptable count or duration.

You can specify a rate above which the alert notification is suppressed. For example, with a rate limit of five alerts in 10 minutes, alerts are suppressed starting with the sixth alert in that interval. The suppression applies per individual alert, not over an entire type of alert.

9 Click **OK** to save your settings.

Editing, Duplicating, and Deleting Suppression Specifications

You can edit and delete suppression specifications. You can also use the duplicate option to create a new suppression specification from an existing one.

To a edit suppression specification, perform the following steps:

1 Choose **Alerts > Suppression Specifications**.

2 Choose the suppression specification and click **Edit**.

3 Modify settings as described in Creating Suppression Specifications on page 529.

4 Click **Save**.

To create a new suppression specification from an existing one, perform the following steps:

1 Choose **Alerts > Suppression Specifications**.

2 Choose the dispatch specification and click **Duplicate**.

3 Modify settings as described in Creating Suppression Specifications on page 529.

4 Click **Save**.

To delete a suppression specifications, perform the following steps:

1 Choose **Alerts > Suppression Specifications**.

2 Choose the dispatch specification and click **Delete**.

3 You cannot recover an alert after it is deleted.

Configuring an SNMP Server

To use SNMP alert functionality, you must download the structure of the Management Information (SMI) and the Management Information Base (MIB) to the appropriate location for the SNMP agent software that you are using.

The RAMS SNMP dæmon does not support a full MIB walk. The value "integer 0" is returned for all queries to portions of the MIB that are not supported. This may be interpreted by the querying system as an error in the type of the returned value.

To download the Packet Design SMI and MIB, perform the following steps:

1 Open a web browser, enter the appliance IP address, and log in as prompted to open the web interface.

2 Choose **Support** to open the Support page.

3 In the Download MIBs section, click **Packet Design SMI** to display the Packet Design SMI in the browser window.

4. From the File menu on your browser, choose **Save As** and save the SMI file in the appropriate location for the SNMP agent software that you are using.
5. Click the **Back** button on your browser to return to the Support page, and then click **Packet Design Products MIB** to display the Packet Design MIB in the browser window.
6. From the File menu on your browser, choose **Save As** and save the file in the appropriate location for the SNMP agent software that you are using.
7. Click the **Back** button on your browser to return to the Support page, and then click the **PD Route Explorer MIB** link to display the system MIB.
8. From the File menu on your browser, choose **Save As** and save the MIB file in the appropriate location for the SNMP agent software that you are using.

Configuring a Remote Syslog Server

Before configuring RAMS to send alerts using the syslog method, you must set up a syslog server, such as syslogd, to accept remote logging of events.

Keep the following points in mind:

- The machine receiving syslog messages must have appropriate firewall settings to allow this. Some Linux systems come with Security Level set to High, which blocks the syslog port.
- syslogd may default to starting with no remote reception capability. To receive remote syslog messages, you may need to restart syslogd with the -r option. You can check the /var/log/messages file for a "`syslogd started < remote reception >`" log entry.

After you have configured the remote syslog server, you can configure syslog settings for alerts.

A Protocol Compliance

This appendix lists protocol compliance information.

- OSPF:
  - RFC 2328, OSPF Version 2
  - RFC 3101, The OSPF NSSA Option
- IS-IS
  - ISO 10589, or RFC 1142 (ISO 10589 draft), OSI IS-IS Intra-Domain Routing Protocol
  - RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
  - RFC 2763, Dynamic Hostname Exchange Mechanism for IS-IS
  - RFC 3784, IS-IS Extensions for Traffic Engineering
  - RFC 2966, Domain-wide Prefix Distribution with Two-Level IS-IS
  - RFC 3567, IS-IS Cryptographic Authentication
- BGP:
  - RFC 4271, BGP Version 4
  - RFC 2796, BGP Route Reflection
  - RFC 1997, BGP Communities Attribute
  - RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
  - RFC 2547, BGP/MPLS VPNs
  - RFC 2858, Multiprotocol Extensions for BGP-4
  - RFC 4364, BGP/MPLS IP VPNs

- EIGRP: Since EIGRP is a Cisco proprietary protocol, there is no RFC to specify the protocol. Documentation for EIGRP is available on the Cisco website (**http://www.cisco.com/warp/public/103/eigrp-toc.html)**.

The following implementations for multi-protocol route resolution are incomplete:

- RFC2328: 16.3 (summary LSAs in transit-area)
- ISO/IEC 10589: QoS metric, virtual link
- ECMP for BGP next hops

Index

B

C

D

E

F

G

H

I

L

M

R

S

T

U

V

W

X

Z