

HP Route Analytics Management System

Software Version: 5.20

Traffic Analysis User's Guide

Manufacturing Part Number: BA129-90025

Document Release Date: May 2007

Software Release Date: May 2007



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

© Copyright 1999-2007 Hewlett-Packard Development Company, L.P.

Contains software from Packet Design, Inc.

©Copyright 2006 Packet Design, Inc.

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.

Support

You can visit the HP software support web site at:

<http://support.openview.hp.com/support.jsp>

HP Software online support provides an efficient way to access interactive technical support tools. As a valued support customer, you can benefit by using the support 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 and HP Passport, go to:

http://support.openview.hp.com/new_access_levels.jsp

Contents

1	Introduction	17
	How RAMS Works	18
	Key Components of RAMS	19
	Using RAMS	22
	The RAMS Web Interface	22
	The RAMS Application Interface	23
2	Configuration and Management	25
	Configuration Overview	26
	Connecting to the RAMS Home Page	27
	Logging In to RAMS	29
	Applying License Keys	31
	Designating the Master and Clients	33
	Assigning the Master Role to a RAMS Appliance	34
	Adding Clients to the Configuration	35
	Relinquishing the Master Role of a RAMS Appliance	36
	Setting the Time and Date	38
	Configuring the Network Interfaces	40
	Selecting the Administration Interface	42
	Configuring an Alias Interface	42
	Configuring a Static Route	42
	Configuring Multiple Port Options	44
	Viewing System Settings	46
	Configuring RAMS for Recording	47
	Creating a Configuration Hierarchy	47
	Configuring the Route Recorder	50
	Adding Protocol Instances	51
	Interconnection of BGP and IGP Protocol Instances	54

Configuring a BGP Confederation	55
Configuring Multiple EIGRP Autonomous Systems	55
Configure an IGP Instance	56
Configure a BGP Instance.	60
Starting and Stopping the Route Recorder	64
Viewing and Modifying Route Recorder Settings	65
Changing the Area ID Format of an OSPF Protocol Instance	65
Deleting an Existing Domain or Protocol Instance	66
Removing an Interface from a Protocol Instance.	66
Changing an OSPF Authentication Password	67
Deleting an OSPF Authentication Password	67
Configuring the Flow Collector	68
Starting and Stopping the Flow Collectors	70
Viewing Flow Collector Settings	71
Configuring the Flow Analyzer	75
Starting and Stopping the Flow Analyzer.	77
Viewing and Modifying Flow Analyzer Settings.	77
Additional Configuration Tasks	79
Configuring GRE Tunnels	79
Configuring a Loopback Interface for Cisco Routers	82
Enabling Technical Support Access	83
3 Administration	85
Creating a Traffic Class.	87
Editing a Traffic Class.	91
Deleting a Traffic Class	93
Creating Traffic Groups.	94
Editing Traffic Groups.	98
Configuring the VNC Server	99
Connecting to the VNC Persistent Session	100
Using On-Demand VNC Sessions	101
Managing Databases	102
Using Offline Databases.	104
Deleting a Database	104
Renaming Databases.	105
Using Online Databases.	106

Backing Up and Restoring Data	107
Creating a Backup File	107
Saving a Backup File	110
Uploading a Backup File	111
Restoring a Backup File	111
Deleting a Backup File	112
Enabling SMB and Adding a Remote Server	113
Restoring a Backup File from a Remote Server	114
Transferring Backup Files Using FTP	115
Choosing a Data Source	116
Archiving Data	119
Configuring Automatic Archival Settings	119
Manually Archiving Data	121
Restoring Archived Data	122
Updating Software	124
Updating with Internet Access	125
Updating without Internet Access	126
Returning to a Previous Version of RAMS	128
Creating Daily Reports	129
Configuring the Mail System	130
Scheduling Daily Reports	131
Understanding Daily Report Contents	132
Viewing Saved Daily Reports	132
Configuring the FTP Server	134
Managing Users	136
Adding, Deleting, and Editing Users	137
Setting Session Login Timers	139
Viewing and Exporting RAMS Log Pages	140
Uploading Layout Backgrounds	142
Using Diagnostic Functions	143
Pinging a Network Device	143
Running a Traceroute	143
Shutting Down the RAMS Appliance	144

4	RAMS Viewers	147
	The X Window Server	148
	Using X Window System Software for MS Windows™	148
	Using X Window Server for Unix Platforms	150
	VNC Viewer	152
	Downloading VNC	152
	Downloading and Installing VNC on Windows	152
	Downloading and Installing VNC on Linux	153
	Opening the RAMS Application in VNC	155
	Installing SVG Plug-In	155
	View System Information	155
5	The Routing Topology Map	157
	Opening a Routing Topology	158
	Managing Topology Map Layouts	159
	Creating a Layout	159
	Applying Layout Backgrounds	159
	Saving Topology Map Layouts	160
	Understanding Symbols and Colors on the Map	161
	Understanding Links and Peerings on the Map	162
	Using the Toolbar	163
	Using the Status Bar	165
	Using the Menu Bar	167
	Topology Menu	167
	Edit Menu	168
	View Menu	168
	Tools Menu	170
	Topology Diagnostics Submenu	172
	Traffic Menu	172
	Help Menu	173
	Legend Panel	173
	Understanding the Topology Hierarchy	175
	Viewing Node and Link Details	177
	Node Information Panel: Protocol Tab	177
	Node Information Panel: Traffic Tab	179

Link Information Panel: Protocol Tab	180
Link Information Panel: Traffic Tab	182
Viewing Complex Routing Hierarchies	184
Selecting Nodes in a Specific Area	184
Trimming the Displayed Nodes	186
Trimming Leaves	187
Finding a Router	189
Viewing Network Anomalies At A Glance	190
Viewing a Complete Up-to-the-Minute Network Inventory	191
Router List	191
Links List	192
Prefix List	193
OSI Prefixes List	195
Viewing the Exit Routers from the Network for any Internet Destination	196
Using the Network Summary Option	197
Assigning Router Names using Router Name Repository	198
Assigning AS Names using AS Name Repository	202
Verify and Manually Assign BGP AS Assignments to Routers	205
Highlighting the IP Route Between Two Points in the Network	207
Finding a Route By Prefix	209
Viewing the Highlighted Path Cost for EIGRP	210
Using the Configuration Options Dialog Box	212
Analysis Options Panel	212
Node Labels Option Panel	213
Colors Option Panel	214
Miscellaneous Options	214
History Navigator Options	215
Auto-Hide Options	216
Diagnosing EIGRP Topology Errors	219
List Topology Errors	219
List Inaccessible Routers	221
List Mismatched Distances	222
Find Invisible Links	224

6 The History Navigator	227
About the History Navigator Window	228
History Navigator Controls	230
Status Bar	230
Cursor	231
Buttons	231
Playback Controls	232
Zoom Buttons	233
Zooming the Time Line	233
Selecting History Graphs	234
Analyzing Historical Data	235
Root Cause Analysis	235
Animation Window	238
Playing an Animation	240
Saving an Animation	241
RIB Visualization	242
Generating a Visualization	242
Changing RIB Visualization Thresholds	244
Saving a Visualization	246
RIB Browser	247
IGP Protocols	247
BGP Protocol	248
VPN Protocol	250
OSI IS-IS Protocol	251
RIB Browser Comparison	252
IGP Protocols	253
BGP Protocols	254
VPN Protocol	255
OSI IS-IS Protocol	255
Event Analysis	256
IGP Protocols	256
BGP Protocol	257
VPN Protocol	258
OSI IS-IS Protocol	258
Traffic Reports	259
VPN Reports	260

Understanding the Events List	261
Events List Controls	262
Event Details	263
Highlighting Associated Nodes	267
Filtering the Events List	267
Adjusting the Time Range	268
Moving Time and Executing Events	269
Using the History Navigator as a Forensic Tool	271
Correlating Time Series Data	278
Using Filters	280
Expression Syntax	282
Examples	283
Expression Definitions	283
7 Network Planning	291
Network Planning and Analysis Tools	292
The Planning Toolbar	292
The Edit Menu	294
The Planning Reports Window	294
Editing Network Topology	297
Add a Node	297
Create a Peering	298
BGP Routers	299
IGP Routers	300
Add a Prefix	301
BGP Routers	302
IGP Routers	304
OSI IS-IS Router	306
Change a Prefix	308
Add Traffic Flow	308
Edit Traffic Flows	310
Change the Bitrate of a Flow	311
Move a Traffic Flow from One Router to Another	312
Delete a Flow from a Router	313
Add a Flow Server	313
Add a Traffic Flow	314

Bring Down a Node	314
Bring Down Peerings	315
Bring Down a Prefix	316
Bring Down an OSI Prefix	318
Set or Change Link Metrics and Bandwidth.	319
View and Remove Edits	320
Trending and Estimation	321
Link Traffic Estimation	321
Total Traffic Estimation	322
Link Utilization Prediction	324
Analyzing Topology Edits	328
Changes in Link Utilization.	329
Changes in Link Traffic	330
Changes in Available Capacity	331
Changes in Neighbor AS Traffic	332
Changes in Destination AS Traffic	333
Changes in Exit Router Traffic	333
Changes in Community Traffic	334
Using the Set Interface Capacities Table	336
Customizing Interface Capacity Values.	336
Clearing Manual Interface Capacity Changes	338
Reverting to Saved Interface Capacity Values	339
Importing and Exporting Planning Data	340
Export Topology Edits	340
Import Topology Edits	343
8 BGP Reports	345
Accessing the BGP Report Pages	347
Creating BGP Activity Reports.	348
BGP Activity Summary Report	348
BGP Activity by AS Report	349
BGP Activity by Peer Report	350
Route Flap Report.	350
Prefix Event Detail	351
Creating BGP Logical Topology Reports.	353
Route Distribution Detail Report.	353

Redundancy by Prefix Report	354
Baseline Redundancy by Prefix Report	355
AS Reachability Report	355
Baseline AS Reachability Report	355
Prefix Reachability Report	356
9 IGP Reports	357
Preparing to Create IGP Reports	359
Accessing the IGP Report Pages	360
Creating IGP Reports	361
Changed Metrics	363
Flapping Links	364
Network Churn	366
New Prefixes	368
New Routers and Links	369
Prefix List	371
Prefix Origination Changes	374
Prefix Origination from Multiple Sources	376
Prefixes Withdrawn	378
10 VPN Configuration and Reports	381
Understanding the Reachability and Participation Index	383
Using VPN Reports	384
Creating Customer and Route Target Associations	386
Configuring VPN Alarms	389
Configuring VPN Reports	392
VPN Summary Report	392
VPN Alarms Report	392
VPN Reachability Reports	394
VPN Participation Reports	395
Summary Reports for More Detailed Information	396
Details by PE	396
List Routes	397
Highlight PEs or Highlight VPN	397
PE Participation over Time or Reachability Over Time	397
History Navigator	399

VPN Prefixes Report	399
IPv4 Prefixes Report	399
Links Report	400
Routers Report	400
Using the VPN History Navigator	401
11 Path Reports	403
Selecting a Topology for Path Reports Analysis	404
Using the Path Reports Window	406
Path Statistics Report	407
Path Statistics by Source	411
Path Statistics by Destination	412
ECMP Paths Analysis	413
Single (Non-ECMP) Path Analysis	414
Asymmetric Paths Analysis	415
Asymmetric Paths by Path	416
Asymmetric Paths by Metric	416
Asymmetric Paths by Source	417
Asymmetric Paths by Destination	417
Network Element Analysis	418
Router Hot Spots	418
Link Hot Spots	419
Unused Links	420
Down Nodes	420
Down Links	420
Asymmetric Link Metrics	421
Failure Analysis	422
Path Failure Analysis	422
Link Failure Analysis	423
Failure-Induced ECMP Analysis	424
12 Traffic Reports	425
Viewing Traffic Reports	426
Using the Traffic Reports Summary Window	427
Using Filters	428
Select Traffic Groups Tab	429

Display Modes	430
Window Icons	430
Creating Time Range Reports	432
Using the Summary Report Window	434
Traffic and Routing Events Correlation Report	434
Link Utilization Report	435
Neighbor AS Traffic Report	435
Exit Router Traffic Report	436
Destination AS Traffic Report	436
Using the IGP Link Report Window	438
History Tab	439
Traffic Groups Tab	440
Flow Tab	441
Details by Flow Report	441
Show Prefixes for Selected Flow Report	442
Show Paths for Selected Flow Report	442
Using the BGP Peering Report Window	443
Neighbor AS Traffic Report	443
Destination AS Traffic Report	444
Transit AS Traffic Report	445
Exit Router Report	446
Distribution by BGP Community Report	446
Using the Flow Collectors Report Window	448
Traffic By Protocol Report Window	450
Exporting Routers Report Window	451
Using the View Traffic Groups Window	452
Using the Traffic Menu	453
13 Alerts	455
Configuring an SNMP Server	456
Configuring the Remote Syslog and RAMS	459
Setting an Alert	460
Understanding Alert Formats	463
SNMP Trap Alert Format	463
Syslog Alert Format	463

Configuring IGP Alerts	465
Adjacency Lost Alert	465
Established Peer Alert	466
Adjacency Flap Alert	466
Prefix Change Alert	468
Prefix Origination Change Alert	469
Prefix Flap Alert	470
Route Change Alert	470
Routing Event Alert	471
Excess Churn Alert	472
Peer Change Alert	473
Configuring BGP Alerts	475
BGP Route Flap Alert	475
BGP Lost Redundancy Alert	476
BGP Lost Peer Alert	476
BGP Prefix Flood and BGP Prefix Drought Alerts	477
BGP Acquired Redundancy Alert	478
BGP Established Peer Alert	478
BGP AS Path Longer Alert	479
BGP Down to One Path and BGP Down to Zero Paths Alerts	479
Configuring OSI IS-IS Alerts	480
ES Neighbor Flap Alert	480
ES Neighbor Change Alert	481
ES Neighbor Origination Change Alert	481
Prefix Neighbor Flap Alert	482
Prefix Neighbor Change Alert	483
Prefix Neighbor Origination Change Alert	484
OSI Route Change Alert	485
Configuring Traffic Alerts	486
Traffic Link Utilization Alert	486
Traffic Route Correlation Alert	488
A License Key Details	491
B Protocol Compliance	495
Index	497

1 Introduction

The Route Analytics Management System (RAMS) is an IP Route Analytics tool that listens to routing protocols and builds a real-time routing topology map. This map helps you visualize and understand the dynamic operation of the network. RAMS also collects and aggregates traffic data, so you can 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), or IGP area. The *History Navigator* window lets you play back a history of your routing topology changes.
- **Monitoring and alerts.** Monitor vital service parameters (network churn, prefix flaps, traffic link utilization, and so on), watch for changes in specific end-to-end service paths and prefixes, and look for degrading redundancy. RAMS can also raise alerts on all watched parameters to head off costly outages.
- **Interactive routing and traffic analysis.** Perform “before and after” comparisons and detailed event analysis using a comprehensive routing base and complete event history to help you rapidly establish the cause of the problem.
- **Planning support.** Display network activity patterns to help you 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 will respond to specific failures or upgrades. You can also import and export these simulated changes, so you can manage multiple routing scenarios using external editors.
- **Reports.** View trends and identify emerging issues before they become problems. You can generate web-based reports for any recorded time period, which show you key information about network health.

How RAMS Works

RAMS appliances physically connect to the network in either of two ways: 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 RAMS to listen to link-state routing protocols (OSPF or IS-IS) in one location as each router knows of all adjacencies in the network. Link-state routers send periodic update messages, which communicate network information to each other, and to RAMS.

Unlike links between OSPF and IS-IS routers, BGP peerings may not follow physical paths. RAMS must discover BGP routers and their peerings indirectly, by receiving routes whose next hop attribute contains the address of a BGP router. Therefore, beyond the physical connection between a BGP router and a peer, RAMS can only infer the existence of a BGP peering if it is advertising prefixes.

When you first connect RAMS to the network, the appliance acquires the topology in a matter of minutes in most cases. The process can take up to an hour for an EIGRP network. As noted in the *RAMS Appliance Setup Guide*, you should connect the unit to the core routers. When connected to a core router, the RAMS appliance becomes more resilient with respect to loss of edge connectivity and remains useful for recovery purposes even during a widespread outage.

RAMS then maintains a real-time topological view of the entire network, which you can view and manage from your desktop computer through the RAMS graphical user interface.

Key Components of RAMS

A typical multi-unit deployment of RAMS appliances includes the following components:

HP RAMS Route Recorder(s) — A RAMS 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 autonomous systems from a single appliance.

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

HP RAMS Flow Analyzer — A Traffic Explorer appliance that correlates traffic and routing data, then uses this combined data to produce reports.

HP RAMS Modeling Engine — A RAMS appliance that creates a synthesized view of routing and traffic data collected across the network by one or more RAMS units. The Modeling Engine presents this data in a graphical user interface accessible from your desktop, providing a single, cohesive view of network activity.

These components may exist on a single RAMS appliance or on separate units, depending upon the size of the network and number of concurrent users desired.

With RAMS, 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, rather than remotely through generic routing encapsulation (GRE) tunnels. A centralized Modeling Engine retrieves the recorded data from each recorder. Users can then monitor network-wide routing information 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 deployment, you can configure each appliance to record data per protocol or per multiple protocols; per area or per area within a protocol; or any combination thereof. Recorder configuration is described in [Chapter 2, “Configuration and Management.”](#)

Figure 1 shows how data flows through these components:

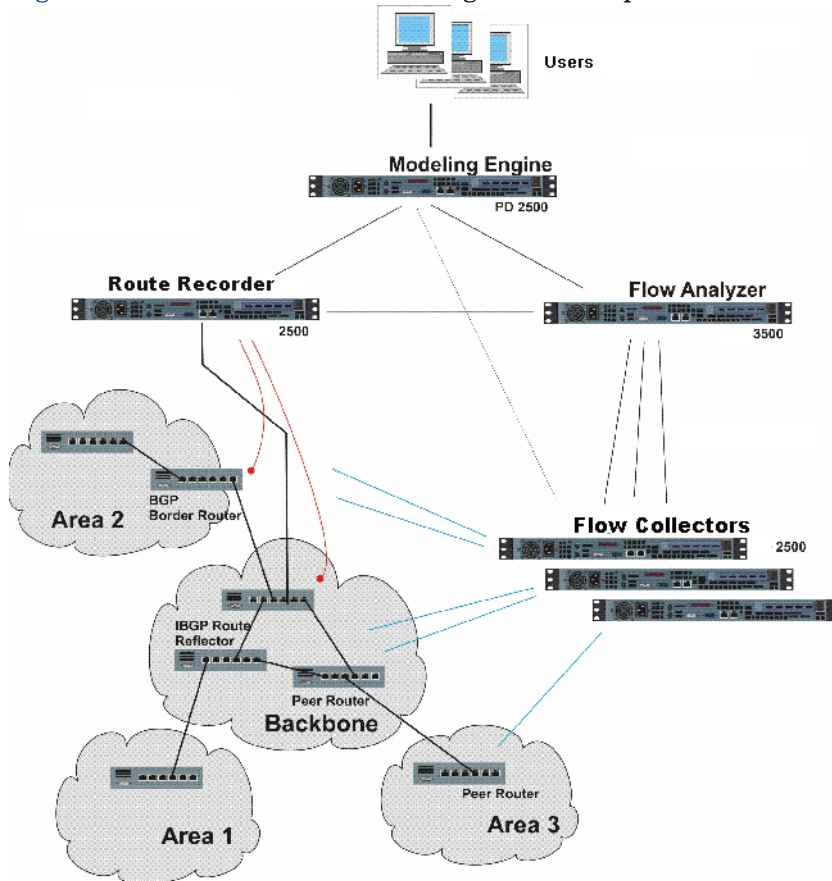


Figure 1 Connection to a Network

- 1 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. Meanwhile, Route Recorders collect network routing data.
- 2 The Flow Analyzer receives routing data from Route Recorder and traffic data from the Flow Collectors. The Flow Analyzer then correlates the traffic and routing data.

- 3 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, providing an accurate, real-time view of how the network is directing traffic.

In a distributed environment where multiple RAMS appliances are installed on the network, one unit in the deployment will act as the master appliance, with other appliances in the deployment acting as clients of the master unit. For example, the Modeling Engine may be designated the master appliance, while the Route Recorder, Flow Collectors, and Flow Analyzer are client units. From the master appliance, you view and configure the clients for recording. You also manage license keys for the entire configuration from the master appliance.

Using RAMS

You connect to the RAMS appliance using two types of viewers:

- A web browser, which primarily displays the RAMS *Administration* pages. The system administrator uses this web-based interface to perform tasks like managing the database, updating software, and configuring the recorders.
- A VNC or X Window System client, which displays the RAMS application. Network engineers and operators use the VNC or X client interface to perform tasks like viewing the routing topology map, creating traffic reports, and analyzing network activity.

Both types of viewers allow remote access to the RAMS appliance, 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.

The RAMS Web Interface

After you log into a RAMS appliance through a web browser, the *Home* page of the appliance appears. The *Home* page provides a central launching point for the product.

At the top of the *Home* page are the following navigational links, which provide access to each area of the web interface:

- **Administration** — Connects you to the appliance *Administration* pages, where you perform administrative tasks such as user management, create traffic classes and groups, and software updates.
- **Recorder Configuration** — Connects you to the *Recorder Configuration* page. In a distributed system, you configure the Route Recorder, Flow Analyzer, and Flow Collector 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.

The RAMS Application Interface

After you launch the RAMS 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. RAMS offers three modes for viewing and manipulating the topology map:

- **Online Mode** — The topology is currently being recorded and updates to the routing database are shown on the topology map as they occur.
- **Design Mode** — Planning features are enabled for the topology map.
- **History Mode** — Only previously recorded information in the routing database is shown on the topology map.

In both Design mode and History mode, you can focus on most any snapshot 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.

By default, topologies open in Online mode. You can toggle from mode to mode with a click of a button.

In addition to the main topology map window, RAMS provides the following powerful tools:

- **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 you've made to the topology map in Design mode. This tool also provides analysis of how the edits theoretically affect network traffic.
- **Traffic Reports** — Allows you to view reports based on traffic collected by Flow Collectors, then correlated and analyzed by the Flow Analyzer.

- **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 Recorder(s). 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 Configuration and Management

After the hardware required to connect one or more RAMS appliances to the network is installed, the administrator can configure each unit using the RAMS web-based *Administration* pages.

In a distributed configuration, where multiple appliances are deployed, one of the units, a Route Recorder or a Modeling Engine, must be designated as the “master unit” during the configuration process. The other units in the deployment become clients of the master.

The following descriptions apply to the master and client(s):

- **Master:** A RAMS that allows centralized configuration of all client units in a multi-appliance deployment. You can monitor and manage these clients using the web-based *Administration* pages on the master unit.
- **Client(s):** One or more RAMS appliances configured and monitored by the master unit in a distributed configuration. A client can be a Modeling Engine, Route Recorder, Flow Collector or Flow Analyzer.

Follow the procedures in this chapter to configure the Route Recorder to listen to particular protocols, as well as configure the Flow Collectors to collect traffic data, and the Flow Analyzer to correlate routing and traffic information.

Configuration Overview

Proceed through the tasks described in this section to initially configure one or more RAMS systems. You perform the majority of these tasks on the unit designated as the master. Tasks that require you to access client machines individually are clearly noted.

- 1 [Connecting to the RAMS Home Page](#) on page 27.
- 2 [Logging In to RAMS](#) on page 29.
- 3 [Applying License Keys](#) on page 31.
- 4 [Designating the Master and Clients](#) on page 33.
- 5 [Setting the Time and Date](#) on page 38.
- 6 [Configuring the Network Interfaces](#) on page 40.
- 7 [Configuring RAMS for Recording](#) on page 47.
- 8 [Additional Configuration Tasks](#) on page 79.
- 9 [Enabling Technical Support Access](#) on page 83.

These tasks should be completed in the order in which they appear in this chapter. Task 4, Designating the Master and Clients, does not apply to single-unit configurations.

Many of the following configuration tasks are performed using the RAMS *Administration* pages. Access the *Administration* pages by clicking **Administration** on the RAMS *Home* page. The *View Configuration* page appears by default, with a blue navigation bar on the left side of the screen.

Note The links and buttons on the *Administration* pages might differ from those shown in the examples in this guide, depending on the functions for which your system is licensed. If your system is not licensed for a particular function, the links or buttons related to that function do not appear on the *Administration* pages.

Connecting to the RAMS Home Page

RAMS has a built-in web server, and this provides the primary administrative access. The first configuration step is connecting to the RAMS Home page by using a supported web browser.

RAMS supports the following browsers:

- Netscape 6.2 and 7.1.
- Internet Explorer 5.5, 6.0, 6.22, and 7.0.
- Firefox 2.0
- Macintosh Internet Explorer 5.2.1.
- Mozilla 1.4, 1.5, and 2.0 (Linux).

To connect to the RAMS Home page, perform the following steps:

- 1 Open a standard web browser and type the initially configured address or hostname of the appliance.

In a distributed configuration, type the address or hostname of the master unit.

The *Home* page is password-protected and user input is encrypted using SSL. A confirmation security dialog box appears the first time you access the website. The secure web pages have self-signed certificates from the Route Analytics Management System.

Note If you are using Internet Explorer 7.0 to connect to the RAMS, a window with the warning “There is a problem with this websites security certificate” opens after you enter the IP address in Step 1. It is safe to ignore this warning. Select “Continue to this website (not recommended)” to connect to the appliances’ *Home* page.

- 2 Click **Yes** to accept the certificate and allow the secure web pages to open. The *Login* page appears as shown in [Figure 2](#).

Login

Username

Password:

Your web browser must accept cookies to login.

Figure 2 Log-In Page

Logging In to RAMS

Before you begin to configure RAMS, you must log into the Administration pages and change the default administration password.

To log into the RAMS Administration pages, perform the following steps:

- 1 Connect to the RAMS *Home* page as described in the previous section, [Connecting to the RAMS Home Page](#) on page 27.

The *Login* page appears.

- 2 Type the default administrator user name (`admin`) and the default password (`admin`).
- 3 Click **Login**.

After a period of inactivity, you must repeat this step to have continued access to the *Administration* pages.

Note To log into the system, your browser must accept cookies.

After a successful login, the *Home* page appears as shown in [Figure 3](#).

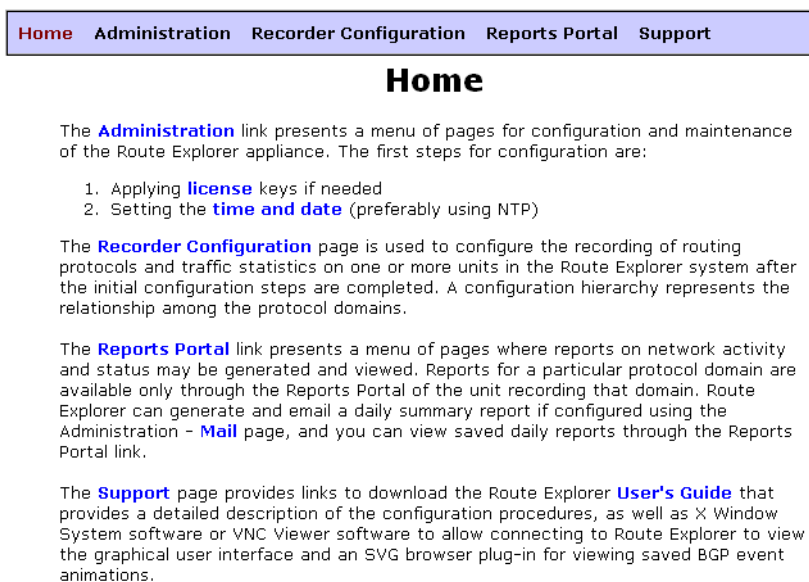


Figure 3 Home Page

The Home page provides links to RAMS support and documentation downloads for the X Window system or VNC, either of which is required to run the RAMS software. The X Window system is recommended for users with high-speed Internet connections, while VNC is more appropriate for users with dial-up or DSL connections.

Note The third-party X Window system software provided with RAMS comes with a 30-day evaluation license. After this initial 30-day period, you must purchase a license from StarNet Communications, Corp to continue using the software.

For security reasons, it is recommended that you change the administrator password when you first log in, and on a regular schedule after that.

To change the administration password, perform the following steps:

- 1 Click **Administration** at the top of the *Home* page.
- 2 On the left navigation bar, click **Users**.
- 3 On the *User Administration* page, select the administrator's user name from the **Current Users** list, then type the new password.
- 4 Click **Update**.

For more information about user administration, see [Chapter 3, "Administration."](#)

Applying License Keys

A license key determines what functions a RAMS unit can perform, as well as how many users and routers it can support, if it has master capability, and in the case of Route Recorder, what protocols it participates in.

Each license key is tied to an identification number, or Unit ID, which corresponds to a RAMS appliance with the same Unit ID. In a distributed configuration, license keys for all units in the configuration are applied to the master unit, which in turn assigns the appropriate license keys to its clients based on the Unit ID numbers. RAMS rejects licenses without a Unit ID.

To view the license key information, start on the *Administration* page. Locate the **Maintenance** link in the left frame, then click **License** to launch the *License Update* page. This page displays the current license information for the system, and lets you activate your temporary license or install new license keys. This page displays three buttons:

- **RAMS** — Provides GUI functionality and is enabled for viewing traffic information
- **Traffic** — Provides traffic functionality
- **Modeling Engine** — Provides functions for testing configurations with replication enabled

RAMS comes with a temporary license that unlocks all protocols for up to three users and 1000 routers. The temporary license key expires after 60 days. To activate a temporary license, click the corresponding button on the *License Update* page, shown in [Figure 4: RAMS, Traffic, or Modeling Engine](#). For example, to activate the RAMS license, click **RAMS**, and the temporary license is applied.

If the temporary license expires before you have installed a permanent license key, data recording is disabled, protocol configuration is disabled, and a warning message is displayed on the RAMS user interface. For uninterrupted use of the RAMS, you must obtain and install your permanent license key before the temporary license expires.

Note If you are setting up a system of multiple units, you will need to activate the temporary license on each unit separately by logging into the *Administration* page, navigating to the *License Update* page, and installing the appropriate temporary license key. From the master unit, you can then proceed to add clients, configure recording, etc.

To install a new license key, perform the following steps:

- 1 Enter the license key text-string in the space provided, exactly as it was given to you, including punctuation. If you received your license key electronically, you can use the cut-and-paste feature on your computer. Otherwise, you can manually type it in. Take care to avoid typing mistakes.
- 2 Click **Update**. This installs the license key. The functionality authorized by the license key is immediately available. Each licensing component is described in [Appendix A, “License Key Details.”](#)

License Update

The screenshot shows a software interface for license management. At the top, a blue header reads "License Information". Below it, the "Unit ID" is "00118550FF79". There are three tabs: "RAMS", "Traffic", and "Modeling Engine". The "Traffic" tab is selected, and the text below it says "Feature Value Expiration Date" and "No license information found". Below this is a link "Activate Temporary License".

Below the license information is another blue header "License Update". Underneath it is a text area with the instruction "Copy-Paste the License Update File contents here:". The text area is empty and has a scroll bar on the right. Below the text area are three buttons: "Update", "Cancel", and "Clear".

Figure 4 License Update Page

The *License Update* page displays the applied license key components and the corresponding expiration dates. The master later distributes license components to client units as described in [Designating the Master and Clients](#) on page 33.

In some cases, you may need to reapply licenses from the master to its clients. For example, if a client is unreachable when a new license key is applied to the master, that client will not receive the new key when it is initially distributed among the appliances in the configuration. When the client reestablishes contact with its master, click the **Re-Apply All** button to re-deploy the new license key.

Designating the Master and Clients

Distributed configuration and management allows an administrator to configure and monitor a number of RAMS components from a central location. If you are working with a single-unit configuration, skip this section and proceed to [Configuring the Network Interfaces](#) on page 40.

The following component definitions apply to the units in a distributed configuration:

- **Route Recorder(s):** Collects routing information and stores it in the database.
- **Modeling Engine:** Provides a synthesized view of all network data collected by the Route Recorders and the Flow Collectors in a configuration. In a deployment with multiple Route Recorders, the centralized Modeling Engine replicates the data collected by each Route Recorder, providing a locally accessible database of network-wide information. The centralized Modeling Engine also feeds prefix information to the Flow Collectors and supplies data to the Flow Analyzer to create reports. The Modeling Engine displays this data in a VNC or X viewer as described in [Chapter 4, “RAMS Viewers.”](#)
- **Flow Collector:** Collects traffic flow information from NetFlow recorders and stores it in the database. The Modeling Engine queries the database to create a synthesized view of traffic flow through the network.
- **Flow Analyzer:** Receives traffic data from the Flow Collectors. The Flow Analyzer uses this information to generate traffic flow reports and alerts.

Depending on the type of multi-unit deployment you are configuring, a Modeling Engine or Route Recorder will typically act as the master appliance in your environment. Additional Route Recorders can be configured to listen to different protocols. For example, in new RAMS deployments, the Modeling Engine will usually become the master unit. On the other hand, if you have migrated an existing single-unit RAMS deployment to a multi-unit, distributed configuration, the Route Recorder will be the master unit.

Assigning the Master Role to a RAMS Appliance

The master unit in a distributed configuration allows you to view and manage multiple RAMS appliances through its administration interface. You must designate which appliance (a Route Recorder or a Modeling Engine) in the configuration will act as the master.

Note The unit designated as master will associate itself with client units through an HTTP POST. To authenticate that initial POST, you must use the Master Access password. A default password is already set. If you want to choose a different password, use the Configure Master Password command on the master unit and on each client unit to set the new password. The Master Access password **must** be the same on all units.

To designate the master unit, perform the following steps on the system that will be the master:

- 1 On the left navigation bar, click **Units**.
- 2 On the *Units* page, the IP address of the administrative interface on the master unit is listed in the text box. Click **Make Master**.

The master unit is designated and the Client Configuration section appears on the *Units* page as shown in [Figure 5](#).

The screenshot shows the 'Units' page in the RAMS administration interface. The top navigation bar includes 'Home', 'Administration', 'Recorder Configuration', 'Reports Portal', and 'Support'. The left navigation menu lists various sections: Application, Diagnostics, Maintenance, License, System, Units (highlighted), and Users. The main content area is titled 'Units' and contains a 'Client Configuration' section. This section includes a 'Client List' table with one entry showing the IP address 192.168.0.165 and a 'Delete' button. To the right is an 'Add New Client' form with an 'IP Address' field and an 'Add' button. Below the 'Client Configuration' section is a 'Client Status' table with columns for Unit, Type, Status, and Version. At the bottom of the page are two buttons: 'Relinquish Master Role' and 'Refresh Client Status'.

Figure 5 Units Page

The address of the master unit appears in the *Client List* box.

After you have assigned the configuration master, you must now designate each of the clients the master unit will manage.

Adding Clients to the Configuration

Before you assign client units to the master, be sure the appliances in the configuration are running and reachable. When you add a client, the master unit automatically applies the appropriate license key(s) to that appliance. Once a client is bound to the master, that client cannot be bound to another master until the master-client relationship is dissolved as described in [Client Status Section](#) on page 36. License details are returned when a client is added successfully. For instance, a message indicating that two license keys were applied will appear on the master's *Units* page after you have added a client.

To add clients to the configuration, perform the following steps:

- 1 On the *Units* page of the master appliance, type the IP address of a client in the *Add New Client* text box.
- 2 Click **Add**.

A success message appears. The IP address of the client is now listed in the *Client List* box. If the client does not appear in the *Client List* box, make sure the appliance is properly connected to the network and reachable by the master.

As you add each client, the *Client Status* table on the *Units* page is refreshed to display an updated list of clients. Corresponding to the IP address of each client is its type. For example, if an appliance is licensed to act as a Modeling Engine, the **Type** column lists this description as shown in [Figure 6](#). The **Status** column indicates whether the client is running (Up) or not (Down).

Use the **Refresh Client Status** button at the bottom of the *Units* page to update the displayed version, type, and status of the clients listed on the page.

Units

The screenshot shows a web interface for managing units. It is divided into two main sections: 'Client Configuration' and 'Client Status'.

Client Configuration

- Client List:** A list of IP addresses: (192.168.0.165), (192.168.0.148), and (192.168.0.160). A 'Delete' button is next to the list.
- Add New Client:** A form with an 'IP Address:' label and an input field. Below it is an 'Add' button and the text 'This may take several seconds'.

Client Status

Unit	Type	Status	Version
192.168.0.165 00304871DEEE	Modeling Engine	Up	Build 4.0.72-E
192.168.0.148 00E0188A5F3C	Route Recorder	Up	Build 4.0.72-E
192.168.0.160 00E0188A5880	Flow Collector	Up	Build 4.0.72-E
192.168.0.166 003048850D56	Flow Analyzer	Up	Build 4.0.71-E

At the bottom of the interface, there are two buttons: 'Relinquish Master Role' and 'Refresh Client Status'.

Figure 6 Client Status Section

Relinquishing the Master Role of a RAMS Appliance

If necessary, you can remove the master status from a Route Analytics Management System, and if desired, assign this function to another unit in the configuration. Keep in mind that relinquishing the master role from an appliance requires you to delete all client configurations.

To remove the master status from an appliance, perform the following steps in the order listed.

- 1 Delete each of the client configurations on the master appliance from the *Recorder Configuration* page. To delete client configurations on the master unit, perform the following steps:
 - a On the top navigation bar, click **Recorder Configuration**.
 - b On the *Recorder Configuration* page, click **Stop All Recording**.
 - c Click the client to remove from the tree structure.
A blue pop-up menu appears.
 - d Click **Delete**.
 - e On the confirmation screen that opens, click **Yes**.

- f Repeat Steps c-e to remove additional clients from the configuration.



Caution If a client is inaccessible for any reason, you will receive a warning message when you attempt to delete its configuration. RAMS then allows you to forcibly remove the configuration of the inaccessible client. If you choose to forcibly remove the client configuration of an inaccessible appliance, the client will be unusable until it is reset to factory defaults.

- 2 On the *Units* page of the master unit, remove each of the clients from the *Client List* box. To remove clients from the client list, perform the following steps:
 - a Return to the *Units* page by clicking **Administration** on the *Recorder Configuration* page, then clicking **Units** on the left navigation bar.
 - b On the *Units* page, select the client IP address from the *Client List* box.
 - c Click **Delete**.

A success message appears. The IP address of the client is now removed from the *Client List* box and from the *Client Status* table on the *Units* page.

Deleting an appliance from the *Client List* box means the client unit no longer has a relationship with the master and is free to become the client of another master appliance.

- 3 On the *Units* page, click **Relinquish Master Role**.

The appliance no longer functions as the master of the distributed configuration. You can now designate a second unit the configuration master and, if desired, add the first unit as a client of the new master. For more information, see [Designating the Master and Clients](#) on page 33.

Setting the Time and Date

You must set the time and date for every unit in your configuration before continuing with the procedures described in this chapter. To access the *Time and Date* page, locate **System** on the left navigation bar, then click **Time and Date**.

Note In a multi-unit configuration, access the *Time and Date* page of each client unit individually to configure time and date settings for that appliance.

It is strongly recommended that the time and date for each unit be synchronized with an NTP server to avoid ever having to make manual time adjustments, potentially backwards in time. (For multi-unit configurations, NTP is required.) Because the recorded routing topology database requires that time progress monotonically, the NTP daemon will not adjust the time if the discrepancy is large enough to require a step adjustment rather than slowly slewing the clock. Therefore, you should first set the time manually to the correct time (within a few minutes), and then select **Get time from server**. After the time is set using the NTP server option, verify the results on the *View Configuration* page.



Caution If you have to manually set the clock backwards, you must first rename all currently recording databases and start a new database.

To set the time and date for each Route Analytics Management System, perform the following steps:

- 1 On the *Time and Date* page, shown in [Figure 7](#), select the time zone from the **Time Zone** drop-down list.
- 2 Choose whether to set the time and date from an NTP server or set it manually.
 - If time is to be obtained through NTP, type the primary and secondary NTP server (if necessary) in the **Primary NTP Server** and **Secondary NTP Server** text boxes, respectively.
 - If time is to be maintained manually, and needs to be corrected now, click the **Set time now** check box and adjust the time fields as necessary.

3 Click **Update**.

Note Modifications to the *Time and Date* page are not permitted if recording is in progress.

Time and Date

Time Zone:

◦ **Get time from server**

Primary NTP Server:

Secondary NTP Server:

NTP is intended to keep the system clock accurate. To make large time changes, manually set the time then configure the NTP servers.

◦ **Set time manually**

Set time now:

/ / :

Year Month Day Hour Minute
yyyy mm dd h24 mm

Warning: Route Recorder must be stopped before altering time configuration. It is strongly recommended to use the Database Administration page to delete all previously recorded databases, then change time configuration before resuming recording.

Figure 7 **Time and Date Page**

Configuring the Network Interfaces

Use the *Network and Interface Configuration* page, shown in [Figure 8](#) to set the RAMS IP address, the netmask, the default router, and primary and secondary DNS servers. To access this page, locate **System** on the left navigation bar, then click **Network and Interface**.

Network & Interface Configuration

Any operation on this page may cause the http server to go away for up to 2 minutes.

Hostname:

Primary DNS: Secondary DNS:

Name: Region
 Input: Region
 Output: None Selected
 Mode: Image Capture

Interface	Name	Use DHCP	IP Address	Netmask	Allow Admin
Slot 0/Port 1 (00:30:48:70:B6:B0)	<input type="text"/>	<input type="checkbox"/>	<input type="text" value="65.192.41.45"/>	<input type="text" value="255.255.255.0"/>	<input checked="" type="radio"/>
Slot 0/Port 2 (00:30:48:70:B6:B1)	<input type="text"/>	<input checked="" type="checkbox"/>	<input type="text" value="192.168.3.45"/>	<input type="text" value="255.255.252.0"/>	<input type="radio"/>

Interface	Auto Negotiate	Speed			Duplex	
Slot 0/Port 1	<input checked="" type="checkbox"/>	<input checked="" type="radio"/> 10 Mbps	<input type="radio"/> 100 Mbps	<input type="radio"/> 1000 Mbps	<input checked="" type="radio"/> Half	<input type="radio"/> Full
Slot 0/Port 2	<input checked="" type="checkbox"/>	<input checked="" type="radio"/> 10 Mbps	<input type="radio"/> 100 Mbps	<input type="radio"/> 1000 Mbps	<input checked="" type="radio"/> Half	<input type="radio"/> Full

Alias Interfaces		
Interface	Alias Name	IP Address
Slot 0/Port 2	ConfedAlias1	192.168.122.45
Slot 0/Port 2	ConfedAlias2	192.168.122.90
Slot 0/Port 1	<input type="text"/>	<input type="text"/>

Static Routes		
Destination	Default Router	Netmask
192.168.0.0	192.168.0.2	255.255.0.0
25.0.0.0	192.168.0.2	255.255.255.0
0.0.0.0	65.192.41.1	0.0.0.0
<input type="text"/>	<input type="text"/>	<input type="text"/>

Figure 8 Network & Interface Configuration Page



Caution Before you can change the IP address on a RAMS appliance, you must delete all recorder configurations, delete all associations between master and client, and relinquish the master role. Failure to do so may require reset to factory defaults on all units to recover.

Note In a multi-unit configuration, access the *Network & Interface Configuration* page of each client unit individually to configure network and interface settings for that appliance.

Before you configure the network manually or using DHCP, type the hostname of the Route Analytics Management System in the *Hostname* text box.

Note It is strongly recommended that you use static IP addresses rather than DHCP when you configure a stand-alone RAMS unit. In a multi-unit deployment, static IP addresses are required.

To configure the network manually (recommended), perform the following steps:

- 1 On the *Network & Interfaces Configuration* page, type the following information into the appropriate text boxes:
 - Primary DNS
 - Secondary DNS (optional)
 - Properties for each interface (name, IP address, and netmask)
 - Default router (Gateway Address)
- 2 Click **Update**.

To configure the network using DHCP (single units only), perform the following steps:

- 1 On the *Network & Interface Configuration* page, select the **Use DHCP** check box.
- 2 In the *Name* text box, type a name to identify the interface.
- 3 Select the **Auto Negotiate** check box to use automatic speed and duplex settings.
- 4 Click **Update**. DHCP automatically configures the IP address, netmask, default router, and primary and secondary DNS servers.

Selecting the Administration Interface

Administrative access to RAMS is only available through one of the configured interfaces. For configurations without an option card, this interface is one of the two RJ-45 jacks labeled *Port 1* and *Port 2* on the RAMS (Port 1 by default). If you have a multiple port option card, any of the interfaces can act as the Administration Interface. Click **Allow Admin** beside the interface that acts as the Administration Interface.

Use the IP address associated with the interface to administer RAMS. The IP address is also used in the following ways:

- As an FTP address for file transfer to RAMS.
- As an address where XML queries are sent.
- As an address required by technical support when any request for assistance is made.
- As an address for the X Window System or VNC client software connections.

Configuring an Alias Interface

An Alias Interface is used to add an additional IP address to an interface inside the netmask configured for that interface.

To configure an alias interface, perform the following steps:

- 1 On the *Network & Interfaces Configuration* page in the Alias Interfaces section, select the desired interface from the **Interface** drop-down list.
- 2 Type a name for the alias in the *Alias Name* text box.
- 3 Type the IP address in the *IP Address* text box.
- 4 Click **Update**.

Configuring a Static Route

Use a static route to route packets directly to a specific router in a particular network area. If the address of the Administration Interface is manually configured (recommended), then you must also configure a default static route.

For EIGRP topologies, the default route configured on an interface must be suitable for a Telnet or SSH transmission to all routers in the autonomous systems monitored on that interface. When multiple physical interfaces or tunnels are configured, a separate default route may be configured on each interface by specifying a target router on the subnet of the interface.

If no default route is set on a particular interface, the EIGRP Route Recorder uses policy routing to direct Telnet or SSH packets through one of the EIGRP peer routers. If not, all the EIGRP peer routers on an interface are suitable for this purpose, then you must install a default route on that interface to avoid the possible selection of an unsuitable peer. If more than one interface is connected to the same autonomous system (for improved visibility), the Route Recorder prefers a broadcast interface over a tunnel interface.

Note RAMS does not monitor the Internet Control Message Protocol (ICMP) redirect messages used by some enterprises to route around failures. RAMS does not accept ICMP redirects to update its routing table.

To add a static route, perform the following steps:

- 1 On the *Network & Interfaces Configuration* page, in the Static Routes section, type the destination, default router, and netmask details in the appropriate text boxes.
- 2 Click **Update**.

Note Only one static route may be added at a time. After you click **Update**, a new row of blank fields is provided so you can add another route.

Tip Override the default gateway inserted by DHCP by adding two static routes: 0.0.0.0/128.0.0.0 and 128.0.0.0/128.0.0.0.



Caution It can take up to 30 seconds before the static route becomes visible while the new information is written to the system asynchronously. Be aware that if you click **Update** again in this time period, the new static route information is erased. If the page returns earlier and does not display the new routes, click **Reload** on the browser to refresh the page.

To modify a static route, perform the following steps:

- 1 On the *Network & Interfaces Configuration* page, in the Static Routes section, make any required changes to the route details.
- 2 Click **Update**.

To delete a static route, perform the following steps:

- 1 On the *Network & Interfaces Configuration* page, in the Static Routes section, manually erase the route details.
- 2 Click **Update**.

Configuring Multiple Port Options

You can monitor a different OSPF or IS-IS area or EIGRP autonomous system with each of the ports on the RAMS appliance, including those on a multiple port option card if installed. Multiple area monitoring is also possible using tunnels. See [Configuring GRE Tunnels](#) on page 79.

To set up multiple port monitoring, keep the following factors in mind:

- The default router must be on the same network as the Administration Interface.
- One of the interfaces must be the Administration Interface. RAMS defaults to slot 0 port 1.
- You can use DHCP to set the IP address on the Administration Interface.
- You must assign an IP address to each of the ports or interfaces.

To configure multiple ports from the *Network & Interfaces Configuration* page, perform the following steps:

- 1 To use an expansion port as the Administrative Interface, click **Allow Admin** to switch the administrative interface to the desired port.
- 2 To use DHCP on the Administrative Interface, select the **Use DHCP** check box.

Note It is strongly recommended that you use a static IP address rather than DHCP. In a distributed configuration, a static IP address is required.

- 3 For each of the other interfaces on the card, type the following information in the appropriate text boxes:

- Name (optional)
 - IP address
 - Netmask
- 4 Click **Update**. It may take some time for all of the interfaces to become active.

After the network is configured, check that the network settings are correct using the *View Configuration* page.

Viewing System Settings

The settings configured in this chapter are listed on the *View Configuration* page for each unit. To access the *View Configuration* page from the **RAMS Home** page, click **Administration**, then **View Configuration** on the left navigation bar.

Configuring RAMS for Recording

If you are working with a single-unit configuration, proceed to [Configuring the Route Recorder](#) on page 50, where you'll configure the Route Recorder to listen to particular protocols.

In a multi-unit deployment, RAMS components are configured and managed from a central location, the unit designated as the master. All configuration tasks described in this section are performed on the Recorder Configuration page of the master unit. You cannot configure or manage RAMS components from the Recorder Configuration pages of client units. The Recorder Configuration page at the client level displays only that client's branch of the configuration hierarchy tree, where settings are view-only.

The following RAMS component management tasks must be performed in the order listed:

- 1 [Creating a Configuration Hierarchy](#) on page 47
- 2 [Configuring the Route Recorder](#) on page 50
- 3 [Adding Protocol Instances](#) on page 51
- 4 [Configuring the Flow Collector](#) on page 68
- 5 [Configuring the Flow Analyzer](#) on page 75

If the Flow Analyzer is already running when you begin recording on the Route Recorder or Flow Collectors, you must restart the Flow Analyzer as described on [page 69](#).

Additional configuration tasks include the following:

- [Configuring GRE Tunnels](#) on page 79
- [Configuring a Loopback Interface for Cisco Routers](#) on page 82

Click **Recorder Configuration** on the top navigation bar to access the *Recorder Configuration* page.

Creating a Configuration Hierarchy

This section describes how to create a configuration hierarchy, the first step in managing RAMS components from the *Recorder Configuration* page of the master unit. Use a configuration hierarchy to organize RAMS components into a tree structure, with each “branch” representing a different part of the

configuration. For example, in [Figure 9](#), branches exist for different protocols and the recorders that listen to those protocols, divided into groups of BGP-protocol routers and IGP-protocol routers. You can also organize the tree according to the area being monitored within each protocol.

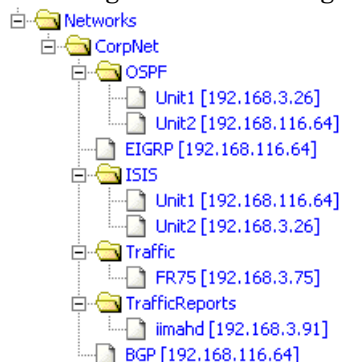


Figure 9 Configuration Hierarchy

The branches in the tree are grouped under a top-level administrative domain, represented by a folder icon on the web interface. You should first establish one top-level administrative domain, like *CorpNet* in the following examples, so you can easily rename the complete hierarchy of recorded databases for backup purposes. You can add multiple administrative domains underneath the top-level administrative domain as needed to organize the structure of the network, for example, to reflect geographical regions, or management divisions running separate protocol instances, or to designate traffic and reports instances. This tree structure is useful when opening a topology in a RAMS viewer as you can load the entire tree or focus only on particular branches, as described in [Chapter 5, “The Routing Topology Map.”](#)

The organization of the tree is reflected by the Modeling Engine when you open a topology in the VNC or X viewer.

To begin creating a configuration hierarchy, perform the following steps:

- 1 On the master unit, click **Recorder Configuration** on the top navigation bar.

The *Recorder Configuration* page appears.

- 2 Click **Networks**.

A blue pop-up menu appears.

- 3 On the blue pop-up menu, move the cursor to **Add**.

The **Administrative Domain** menu item appears as shown in [Figure 10](#).



Figure 10 Administrative Domain

- 4 Click **Administrative Domain**.

The **Name of new Administrative Domain** box appears as shown in [Figure 11](#).

Name of new Administrative Domain:

IP Address

Domain is a BGP AS Confederation

BGP AS Confederation Id:

Figure 11 Adding an Administrative Domain

- 5 Type a name for the administrative domain. The name must consist solely of alphanumeric characters, with an alphabetic character first.
Keep in mind that this should be the top-level domain under which you will create the rest of your hierarchy as recommended.
- 6 For distributed configurations only: In the *IP Address* field, choose **None** from the drop-down list if you are establishing a top-level domain. Otherwise, choose the IP address of the unit to add to the hierarchy.
- 7 If the administrative domain represents a BGP confederation, perform the following steps:
 - a Select the **Domain is a BGP AS Confederation** check box.
A BGP confederation is a domain that contains multiple member autonomous systems, but appears to outside autonomous systems to have a single AS identifier. For more information, see [Configuring a BGP Confederation](#) on page 55.
 - b Type the confederation ID number in the *BGP AS Confederation ID* text box.
- 8 Click **Add Domain**.

- 9 Click the **+** symbol to the left of **Networks** to open the folder that shows the domain name you just entered.
- 10 Proceed to [Configuring the Route Recorder](#) on page 50 to configure one or more Route Recorders to listen to different protocols across the network.

Configuring the Route Recorder

This section introduces the main functions of the Route Recorder and describes how to configure it to start recording. You must configure the Route Recorder before configuring other RAMS components.



Caution Set the time and date on RAMS before configuring the recorder and starting to record data, according to the instructions in [Setting the Time and Date](#) on page 38. The date and time must be set before data recording begins because RAMS relies on accurate time stamps for generating report information. It is strongly recommended that NTP is used to set the time and date.

In a deployment with multiple Route Recorders, you can configure different network segments to be recorded by different Route Recorders. Each Route Recorder can be configured to listen to one protocol or multiple protocols per area or Autonomous System (AS). Only one Route Recorder can record per area or AS.

To add a Route Recorder to the configuration hierarchy, perform the following steps:

- 1 On the *Recorder Configuration* page of the master unit, click the top-level domain name you added in [Creating a Configuration Hierarchy](#) on page 47 (for example, *CorpNet*).
- 2 Move the cursor to **Add**.

The option to add another level of domain name hierarchy appears, as shown in [Figure 10](#).

- 3 Click **Administrative Domain**.

The **Name of new Administrative Domain** box appears as shown in [Figure 11](#).

- 4 Type a name for the administrative domain. The name must consist solely of alphanumeric characters, with an alphabetic character first (for example, *IGPRecorders*).
- 5 For distributed configurations only: From the **IP Address** drop-down list, choose the IP address of a Route Recorder, or choose **None**. When you choose **None**, you create a branch in the configuration tree that can be used to organize the tree by protocol or by area. For example, imagine the *IGPRecorders* branch was added to the tree. In subsequent steps, you can add *OSPF* and *IS-IS* branches under *IGPRecorders*. Then, under *OSPF* and *IS-IS*, you can add one or more Route Recorders.
- 6 If the administrative domain represents a BGP confederation, perform the following steps:
 - a Select the **Domain is a BGP AS Confederation** check box.

A BGP confederation is a domain that contains multiple member autonomous systems, but appears to outside autonomous systems to have a single AS identifier. For more information, see [Configuring a BGP Confederation](#) on page 55.
 - b Type the confederation ID number in the *BGP AS Confederation ID* text box.
- 7 Click **Add Domain**.

The domain name you just entered (for example, *IGPRecorders*) appears in the hierarchy.

The Route Recorder listens to routing protocol packets and records that data in a database. To set up the database(s), proceed to [Adding Protocol Instances](#) on page 51.

Adding Protocol Instances

RAMS uses a hierarchical tree to represent the collection of IGP and BGP routing protocols to be recorded and the relationships among them.

Each instance of an IGP or BGP routing protocol is represented by a page icon as a leaf in the tree. A protocol instance includes the set of routers communicating directly with each other using that particular routing protocol. For example, the routers within a set of interconnected OSPF areas form one protocol instance.

Note that multiple instances of the same protocol cannot be configured within a single administrative domain. If a network contains two instances of the same protocol, then you must create two administrative domains to contain them. Also, a BGP instance cannot be configured in an administrative domain whose ancestor directly contains a BGP instance.

Before starting to record routing data using the Route Recorder, you must assign a name to the database where the data will be stored. You must also specify the routing protocol (IS-IS, OSPF, EIGRP, or BGP) and the network interface used to listen to the routing protocol packets.

Note For each interface, the RAMS supports an interface password and an area or domain password. If that interface is recording Level 1, the area password should be entered when configuring that interface in the recorder configuration. If the interface is recording Level 2, the domain password should be used. If the interface is recording both Level 1 and Level 2, then the area and domain passwords configured on the router must be the same.

When configuring RAMS for the first time, at least one database must be specified for storing routing data.

To add a protocol instance and start recording routing data, perform the following steps:

- 1 In the tree on the *Recorder Configuration* page of the master unit, click the label you added in [Configuring the Route Recorder](#) on page 50 (for example, *IGPRecorders*).
- 2 On the blue pop-up menu, click the desired type of protocol instance (BGP, OSPF, IS-IS, or EIGRP).

One of the following options appears to the right of the configuration tree:

- If the label you added in [Configuring the Route Recorder](#) on page 50 is bound to an IP address, the configuration section for the selected protocol appears. Proceed to [Configure an IGP Instance](#) on page 56 or [Configure a BGP Instance](#) on page 60 for detailed instructions about how to configure the protocol instance.
 - If the label you added in [Configuring the Route Recorder](#) on page 50 is unbound, a drop-down list of Route Recorders appears with the **Select a Unit** button. Proceed to Step 3.
- 3 Choose one of the following options and then click **Select a Unit**:

- If you will be adding only one Route Recorder under this protocol, choose the Route Recorder's IP address from the drop-down list. The configuration section for the selected protocol appears. Proceed to [Configure an IGP Instance](#) on page 56 or [Configure a BGP Instance](#) on page 60 for detailed instructions about how to configure the protocol instance.
 - If you will be adding more than one Route Recorder under this protocol, choose **Multiple** from the drop-down list. Proceed to Step 4.
- 4 Click the protocol label you just added and select **Add Route Recorder** from the pop-up menu. Then choose the IP address of the Route Recorder from the pop-up menu.

The configuration section for the selected protocol appears on the right side of the *Recorder Configuration* page. For detailed instructions about how to configure the protocol instance, see [Configure an IGP Instance](#) on page 56 or see [Configure a BGP Instance](#) on page 60.

- 5 You can configure additional protocol instances by repeating the preceding steps.

Note Multiple different protocol instances can be added under one administrative domain, but only one instance of a particular protocol is allowed.

The structure of the administrative domain hierarchy also affects how RAMS associates the protocol instances to connect them in the routing topology map and to calculate routes across them. The next sections explain three requirements to consider when you configure the hierarchy:

- Correct interconnection of BGP and IGP protocol instances depends on their proximity in the hierarchy.
- If the network is a BGP confederation, this must be indicated in the administrative domain configuration.
- Multiple EIGRP autonomous systems can be configured in one administrative domain or separate ones.

After you determine the administrative domain hierarchy that is appropriate for your network, see the specific instructions in [Configure an IGP Instance](#) on page 56.

The RAMS supports up to 100 areas. This limitation applies to the total of OSPF areas, IS-IS areas (levels), EIGRP autonomous systems, and BGP autonomous systems.

Interconnection of BGP and IGP Protocol Instances

A single physical router can run multiple routing protocols and be a member of multiple protocol instances. RAMS will attempt to consolidate all the instances of that router as a single node on the routing topology map so that the protocol instances will be connected. Because the various protocols identify routers in different ways, RAMS employs a heuristic algorithm to match routers.

A BGP node that peers with RAMS is identified by the peering address, while a BGP node created as a next-hop from a BGP peer is identified by the address in the BGP NextHop attribute of some of the routes learned from that BGP peer. RAMS searches for the nearest IGP protocol instance in the hierarchy containing a router with matching router ID or interface address, or, failing those, a router that advertises a prefix containing the BGP address. If the hierarchy is configured with an administrative domain for each AS containing the BGP and IGP protocol instances of that AS, the intended IGP protocol instance will be nearest. However, if the hierarchy is configured inappropriately, the closest IGP with a match might not be correct.

For example, consider a network with two BGP and two IGP instances. Here, the two BGP instances represent two BGP autonomous systems that are connected with an external BGP peering across a link. In each AS, RAMS would peer with the BGP router on the end of the link in that AS, and from that peer it would learn routes of the other AS. Along with these routes, it would also learn an interface address of the BGP router on the other end of the link using the BGP NextHop attribute of the routes. In order to join the two autonomous systems by a link on the map, RAMS must consolidate this interface address with an IGP router in the other AS. To do this, RAMS first finds the BGP instance of the next-hop router using the AS number from the BGP AS path attribute, then it searches from that point in the hierarchy for the closest IGP instance containing a matching router as described above. If each of the BGP and IGP instances was in its own administrative domain under the *Separate* domain rather than being paired in the *East* and *West* domains as shown, then both IGP's would be equally distant and either might have matched since both are likely to advertise a prefix covering that interface address. As a result, the wrong IGP might be found first, causing the next-hop router to be consolidated with the router at the wrong end of the link.

Configuring a BGP Confederation

A BGP confederation is a domain that contains multiple member autonomous systems but appears to outside autonomous systems to have a single AS identifier. If your network is configured as a BGP confederation, you must create an administrative domain to represent the confederation and configure it with the confederation AS identifier. Under that administrative domain, you then configure an administrative domain for each member AS.

There are two types of BGP confederations. The member BGP autonomous systems in the confederation may all be contained within one IGP domain, or each member BGP AS may be running a separate IGP instance. The administrative domains *Common* and *Separate* are the ones representing the confederation in each case. Underneath these are the *West* and *East* administrative domains, each of which contains a BGP instance for the member AS. For the common IGP case, a single IGP instance is configured under the confederation domain. For the separate IGP case, an IGP instance is configured along with the BGP instance in each member AS domain. You can adapt the approaches in these examples for monitoring your BGP confederation. See [Configure a BGP Instance](#) on page 60 for guidelines on configuring BGP.

Configuring Multiple EIGRP Autonomous Systems

A network with multiple EIGRP autonomous systems can be configured in either of two ways:

- Configure a single protocol instance with multiple network interfaces connected to the different EIGRP autonomous systems, as long as no two autonomous systems have the same number.

Advantages: Fewer configuration steps are required, and the autonomous system configuration need only be entered once. In the table listing protocol events, the events from all autonomous systems within the instance are merged so you can view their relative timing.

- Create separate administrative domains for each AS, each with its own protocol instance. However, you can only configure a particular network interface under a single protocol instance, so you must configure all of the autonomous systems that may be heard on a single interface under the same domain and protocol instance.

Advantages: Each AS can be given a name rather than being identified only by number.

Configure an IGP Instance

After you create an OSPF, IS-IS, or EIGRP protocol instance, the configuration section for the selected protocol instance appears on the right side of the *Recorder Configuration* page, as shown in [Figure 12](#).

Recorder Configuration

Protocol OSPF	
Name : CorpNet	
<input type="checkbox"/> Save raw protocol packets	
Interfaces	
Active	Not Active
Slot 0/Port 2	Slot 0/Port 1 rightsidetunnel router16

Save Start Recording

Starting Route Recorder may take up to 2 minutes.

Disk space used: 12%

Start All Recording

Figure 12 Configuring an IGP Protocol Instance

To configure an IGP instance, perform the following steps:

- 1 To receive diagnostic trace information for your network, select **Save raw protocol packets**. The information is saved in a log file within the ftp directory.
- 2 Beneath the Interfaces section, the configured network interfaces are displayed in the **Not Active** column. Select the appropriate interfaces that are connected to network areas or autonomous systems for this protocol instance, and move them from the **Not Active** column to the **Active** column using the < button.
- 3 If a generic routing encapsulation (GRE) tunnel is required to connect to a remote router, click **New Tunnel**, and then follow the instructions in [Configuring GRE Tunnels](#) on page 79 to create the tunnel interface.

For EIGRP, there can be multiple interfaces connected to a single autonomous system in order to get complete coverage. You may need to configure a default route for each interface as described in [Configuring a Static Route](#) on page 42.

- 4 If OSPF is selected, you can enable OSPF authentication separately for each interface. Select the desired interface and then click **Configure**.

There are two types of OSPF authentication:

- simple: requires only a password (up to eight characters)
- MD5: requires a password (up to sixteen characters) and a Key-Id (a number between 1 and 255, inclusive).

Note To monitor two OSPF areas that are linked to a single Cisco router, you must configure a loopback interface on the router to monitor both areas with RAMS. See [Configuring a Loopback Interface for Cisco Routers](#) on page 82.

- 5 For an EIGRP instance, one or more autonomous systems must also be configured. Click **New AS** to display the autonomous system configuration section, as shown in [Figure 13](#).

For a network with more than one EIGRP AS, there are two configuration choices as explained in [Configuring Multiple EIGRP Autonomous Systems](#) on page 55.

Configure Autonomous System	
AS Number (0 for any):	<input type="text"/>
Periodic Explore Interval* (hours):	<input type="text" value="8"/>
Periodic Explore Start Time (00:00 - 23:59):	<input type="text" value="4:00"/>
Full Explore Interval* (hours):	<input type="text" value="48"/>
Full Explore Start Time (00:00 - 23:59):	<input type="text" value="4:00"/>
Max. Outstanding Queries:	<input type="text" value="10"/>
Telnet Line Password:	<input type="text"/>
TACACS/SSH Username:	<input type="text"/>
TACACS/SSH Password:	<input type="text"/>
Login Method :	<input checked="" type="checkbox"/> telnet <input type="checkbox"/> ssh

Configure Blocked Interfaces	
Interface Address	Remove
<input type="text"/>	<input type="text"/>

Configure Interface Passwords			
Interface Address	Username	Password	Remove
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

Figure 13 Configuring an EIGRP Autonomous System

- 6 In the *Configure Autonomous System* table, fill in the following fields:
- AS Number: Type either an explicit AS number or zero to allow RAMS to record all autonomous systems that are heard on the selected interfaces. To restrict the recording to a subset of the autonomous systems that may be heard, configure each desired AS number explicitly, one at a time.
 - Topology Exploration parameters:
 - Periodic Explore Interval determines the frequency of periodic exploration, and may be configured or left at its default value of 8 hours. Entering zero disables periodic exploration.
 - Periodic Explore Start Time determines when periodic exploration begins, and is expressed in 24-hour clock mode in the time zone set on RAMS.
 - Full Explore Interval determines the frequency of full topology exploration, and may be configured or left at its default value of 48 hours. Entering zero disables exploration.

- Full Explore Start Time determines when full exploration begins, and is expressed in 24-hour clock mode in the time zone set on RAMS.
 - You can change the setting for Max. Outstanding Queries, but the default value is recommended. This setting controls the maximum number of routers to which simultaneous Telnet or SSH connections will be issued to query topology information using the command line interface.
 - Router Password, TACACS User name, and TACACS Password. If all routers use the same password or TACACS user name and password combination, type either the simple router login password or a TACACS user name and password combination, or both if some routers will use one and some routers will use the other. The TACACS fields can also be used for user name and password authentication configured through RADIUS or configured locally on the routers, although it must be the same for all routers in the AS.
- Note** See Step 9 if each router in the AS has a unique simple password or TACACS user name and password combination.
- In the **Login Methods** area, select the method that RAMS will use to log into routers to collect EIGRP topology information by selecting **telnet**, **SSH**, or both if some routers use Telnet and some use SSH. If both are selected, SSH will be tried first.
- 7 In the *Configure Blocked Interfaces* table, specify any router interfaces that you do not want included in the topology map.
You can use this feature to specify routers that RAMS should not attempt to log into, or to limit the scope of the RAMS topology exploration.
 - 8 Type an interface address, then click **Update AS**.
 - 9 Use the *Configure Interface Passwords* table to override the generic passwords specified in Step 6 for any router that has a password different from those specified in the *Configure AS* table. Type an interface address and its password or TACACS user name and password combination, then click **Update AS**.
Note You must configure the password for *all* interfaces on the router.
 - 10 When you have configured all of the autonomous systems, Blocked Interfaces, and Interface Passwords, click **Save** to complete the recorder configuration process.

You can now begin recording routing topology information for this protocol instance, or you can complete the configuration of all other protocol instances in the topology, if any, before you begin recording.

Configure a BGP Instance

After you create a BGP protocol instance as described in [Adding Protocol Instances](#) on page 51, the BGP configuration section appears on the *Recorder Configuration* page as shown in [Figure 14](#). Before getting to the detailed steps for configuring the BGP instance, this section presents some guidelines for choosing what peerings to establish between RAMS and the BGP routers.

For every BGP AS configured, the AS number and list of IP addresses of peers in the AS are required during configuration. Each peer should be configured to an IBGP peer with RAMS. It is important that no policies are applied to the routes sent to RAMS. RAMS does not send any routes to its peers. Nevertheless, you should install filters on the peer routers to prevent the acceptance of any routes from RAMS.

If multiple autonomous systems are monitored, an alias must be assigned to the RAMS appliance for each additional AS. Aliases are logical interfaces created by the OS that are assigned their own IP Address and behave the same as any other physical interface. These aliases facilitate the multiple personalities required for participating in multiple autonomous systems. These additional addresses can be the tunnel end-point addresses or they can be configured in the Administration Interface as IP alias addresses as described in [Configuring an Alias Interface](#) on page 42. Any tunnels should be into the corresponding autonomous systems.

If you use IP alias addresses, all the addresses should be from the AS to which the main/physical routing interface connects. The router at the other end of this interface connection should ensure that each of these addresses is routable. The easiest way to ensure this is to assign all of the alias addresses from the same address block as the interface, meaning, from the subnet of the interface. In this case, no additional configuration is required in the AS routers. However, when the additional IP addresses are not from the same subnet block, the static routes to the AS routers must be configured and injected into the IGP/BGP so that the RAMS appliance is reachable from each BGP peer.

There are three ways to configure a BGP instance in relation to the IBGP full mesh. In order of preference, these are the following:

- Ensure that the RAMS appliance participates as part of the IBGP full mesh with a neighbor relationship (peer relationship) with all of the IBGP routers. This allows RAMS to see all of the IBGP updates sent by all of the routers in the mesh. That is, RAMS sees the topology in exactly the same way that the other IBGP routers see it.
- If your network has route reflectors, you can set up a peer (neighbor) relationship among RAMS and all of the Route Reflectors in the network. Note that you are configuring the Route Reflectors to treat RAMS as a *peer*, not as a *client*. In this scenario, RAMS receives all of the updates from Route Reflector clusters. This provides RAMS with information about all routes being advertised, but the information is more limited than if RAMS were part of the full IBGP mesh. In particular, if some of the Route Reflector clusters have multiple exit points for the same route, RAMS might be able to see only a few (if there are multiple route reflectors present in the cluster), or only one of these exit points.
- The least preferred method is to configure RAMS as a Route Reflector client. This option gives RAMS only one view of the network, that of the Route Reflector. This limits the ability of RAMS to do some types of analysis, but route tracing should work correctly.

To increase the amount of routing information available to RAMS, you can configure RAMS to be a client of several key Route Reflectors. For best coverage, you should select Route Reflectors in geographically distant major PoPs.

Note When you set up peer relationships with Route Reflectors as a client, the total number of received routes is the product of the number of Route Reflectors multiplied by the number of prefixes. Because the total number of advertised routes can be very high, it is recommended that RAMS support no more than 10 Route Reflectors when RAMS is configured as a client.

You can choose the third option even if you do not currently have any Route Reflectors in your network. Any router can be configured to become a Route Reflector for the purpose of peering with RAMS. This does not affect the other BGP neighbors of the router, because being a Route Reflector is a per-neighbor setting. You might consider the third option if the first two options could entail too much configuration effort. Note that if you choose the third option, and later decide to change to the first or second option, you can simply reconfigure the Route Reflectors to treat RAMS as a peer instead of a client.

Configuring RAMS as an external BGP peer is not recommended because there are several drawbacks associated with EBGP peer relationships:

- EBGP routing information does not include certain BGP attributes such as the NextHop, Local-Pref, and MED attributes. These attributes are essential to determine the best BGP routes inside an AS. Without them, RAMS is unable to determine the correct exit routers or find correct paths for BGP prefixes.
- RAMS will learn from each of the EBGP peers one route for each prefix known in the autonomous system, rather than just those routes for which the peer is responsible, unless some policy filter is used to limit this information. If RAMS peers with many BGP routers in the AS, the total number of routes may exceed its capacity.
- Even though RAMS will be maintaining many more routes than with IBGP peering, the fidelity of the routing information is much less because of the missing attributes. Without those attributes, the routes passed to RAMS by different routers will be almost identical.

After you have configured a BGP instance for IBGP peering with the routers in your own AS, you may want to monitor what routes are being advertised to other autonomous systems. For this purpose, you can configure RAMS with a separate BGP instance to EBGP peer with one or a few of the border routers in your AS.

In a BGP confederation, RAMS should be configured to peer with each member AS using one of the three approaches described earlier in this section for a non-confederated BGP AS. That is, separately for each member AS, RAMS should be configured as an IBGP peer participating in the full mesh of IBGP routers in the member AS, or as a Route Reflector peer or client. A different alias must be assigned to RAMS for each member AS.

If your RAMS unit has a license for both BGP and VPN protocols, and you want to monitor VPN routing, then in order to collect complete routing you should configure peering from RAMS either to all of the provider's edge (PE) routers or as a peer to all of the Route Reflectors serving the VPN routes in the AS. As each peering is configured, you must enable BGP extensions for MPLS VPNs, as indicated in the detailed instructions that follow. See [Chapter 10, "VPN Configuration and Reports"](#) for more information about configuring customer/Route Target (RT) associations for VPN reports.

Once you have decided what BGP peerings you want to configure, you can type the appropriate information in the BGP configuration section as shown in [Figure 14](#).

Protocol BGP	Peers
BGP Id: <input type="text" value="192.168.122.90"/> AS: <input type="text" value="65533"/> Name: CorpTwo Interface: <input type="text" value="Slot 0/Port 1"/> <input type="button" value="v"/> <input type="checkbox"/> Save raw protocol packets <input type="button" value="Save"/>	<div style="border: 1px solid gray; width: 100px; height: 100px; margin-bottom: 5px;"></div> <input type="button" value="Add"/> <input type="button" value="Edit"/> <input type="button" value="Delete"/>

Figure 14 Configuring a BGP Instance

To configure a BGP instance, perform the following steps:

- 1 Type the BGP IP address in the **BGP ID** box.
- 2 Type the autonomous system number in the **AS** box. If the BGP protocol instance is within a confederation, type the AS number of the member in the **Member AS** box. This is *not* the confederation ID, which is shown separately.
- 3 From the **Interface** drop-down list, select the physical interface slot and port or select the logical interface alias.
- 4 To receive diagnostic trace information for your network, select **Save raw protocol packets**. The information is saved in a log file within the ftp directory.

- 5 In the **Peers** column, click **Add** to add peers as shown in [Figure 15](#).

The peer configuration table appears at the right side of the screen.

Peer Information	Options
IP Address(es) or Prefix(es): <input type="text"/> <input type="button" value="v"/> AS: <input type="text" value="65510"/> MD5 Password: <input type="text"/>	<input checked="" type="radio"/> Internal <input type="radio"/> External <input type="checkbox"/> BGP ext for MPLS VPNs
<input type="button" value="Save"/> <input type="button" value="Cancel"/>	

Figure 15 BGP Peer Configuration

- 6 In the **Peer Information** column, type the IP address(es) of the peer(s), and the MD5 password of the peer, if applicable.

You can add multiple peers at once by typing each IP address on a separate line in the *IP Address* text box.

- 7 In the **Options** column, specify whether the peer(s) is internal or external. If you enter multiple peers at once, all of the peer addresses must share the same internal or external setting.
- 8 If this BGP topology is a BGP MPLS VPN, select **BGP ext for MPLS VPNs**. Otherwise, leave this check box empty.
- 9 When you have entered all information for the peer(s), click **Update Peer** to return to the BGP instance configuration.
- 10 Click **Save**.

You can now begin recording routing topology information for this protocol instance, or you can complete the configuration of all other protocol instances in the topology, if any, before you begin recording.

Starting and Stopping the Route Recorder

To start recording routing topology data for any of the protocols you added in [Configuring the Route Recorder](#) on page 50, perform the following steps on the *Recorder Configuration* page of the master unit:

- 1 If a Flow Analyzer is running in the configuration, you must stop the Flow Analyzer before you can start recording.
- 2 In the configuration hierarchy tree, click the protocol instance where you'll start recording routing data.
A blue pop-up menu appears.
- 3 Click **View**.
If the protocol instance is not already recording, the **Start Recording** button is displayed on the *Recorder Configuration* page.
- 4 Click **Start Recording**.
When recording has begun, the **Stop Recording** button appears.
- 5 Restart the Flow Analyzer, if you stopped it in Step 1.

Route Recorder will write routing data for the selected protocol to the database until you click **Stop Recording**.

Use the **Start All Recording** and **Stop All Recording** buttons to start and stop recording for the entire configuration tree.

Viewing and Modifying Route Recorder Settings

After the Route Recorder is started, as described in [Starting and Stopping the Route Recorder](#) on page 64, RAMS is ready to receive routing announcements from the peer routers. You can check the status of the recorder using the *Status* table on the *Recorder Configuration* page.

To access the *Status* table, perform the following steps:

- 1 In the configuration hierarchy tree, click a protocol instance.

A blue pop-up menu appears.

- 2 Click **View**.

Depending on the protocol instance, the following information is displayed:

- For non-BGP protocols, the table shows the status of the protocol adjacency, the time the last “hello” packet was received as part of the routing protocol, the time the last event was written to the database by the Route Recorder. The table also identifies the interface, AS, and neighbor AS.
- For BGP protocols, the table shows the peer status for the AS, including IP address of the peer, the BGP Id, the state of the peering, and how long that state has existed (“Since”).

From the *Status* table for each protocol instance, you can also change recorder settings, as described in the following sections.

Changing the Area ID Format of an OSPF Protocol Instance

For an OSPF instance, viewing the configuration also displays the Area ID Format selection. An OSPF Area ID can be displayed either as a single decimal number or in dotted decimal format. The format selection controls both the administration pages and the RAMS client, but you must restart the client for a change in the format to take effect. If you use VNC, you must also stop the VNC server and restart it to see the change.

To change the Area ID format, perform the following steps:

- 1 Follow Steps 1 and 2 in the previous section to navigate to the *Recorder Configuration* page *Status* table.
- 2 From the **Area ID Format** drop-down list, choose Decimal or Dotted Decimal.
- 3 Click **Submit**.
- 4 Click the **Administration** link at the top of the *Recorder Configuration* page.
- 5 Click **VNC Configuration** in the left navigation bar.
- 6 On the *VNC Configuration* page, click **Stop**. This stops the VNC server and closes any active VNC sessions.
- 7 Click **Start** to restart the VNC server.
- 8 Restart any VNC client sessions.

Deleting an Existing Domain or Protocol Instance

To delete an existing administrative domain or any instance, perform the following steps:

- 1 Navigate to the *Recorder Configuration* page *Status* table.
- 2 Click **Stop Recording**.
- 3 Click an existing instance from the tree structure on the left side of the screen.

A blue pop-up menu appears.

- 4 On the blue pop-up menu, click **Delete**.
- 5 Click **Yes** on the confirmation screen that opens.

This deletes the instance, but the database remains in RAMS as a historical record.

Removing an Interface from a Protocol Instance

To remove an interface from a protocol instance, perform the following steps:

- 1 On the *Recorder Configuration* page, select the protocol instance name on the tree.
- 2 From the blue pop-up menu, click **View** to see the protocol configuration.

- 3 Click **Stop Recording**.
- 4 Choose an existing interface name. If the chosen interface is in the **Active** column, move it to the **Not Active** column.
- 5 If the interface is a tunnel that you want to delete, highlight the interface name and click **Configure**.
The Configure Interface section opens on the *Recorder Configuration* page.
- 6 Click **Delete**. A confirmation message page opens.
- 7 Click **Yes**. The tunnel is deleted and the *Recorder Configuration* status page returns.

Changing an OSPF Authentication Password

To change an OSPF authentication password, perform the following steps:

- 1 On the *Recorder Configuration* page, select the protocol instance name on the tree.
- 2 From the blue pop-up menu, click **View** to see the protocol configuration.
- 3 Click **Stop Recording**.
- 4 Choose an existing interface name. If the chosen interface is in the **Active** column, move it to the **Not Active** column.
- 5 Select the interface name, and then click **Configure**.
- 6 Highlight the password and type another password.
- 7 Click **Update**. This changes the password.

Deleting an OSPF Authentication Password

To delete an OSPF authentication password, perform the following steps:

- 1 On the *Recorder Configuration* page, select the protocol instance name on the tree.
- 2 From the blue pop-up menu, click **View** to see the protocol configuration.
- 3 Click **Stop Recording**.
- 4 Choose the existing interface name with the affected password. If the interface is Active, it is first necessary to move it to the **Not Active** column.
- 5 Select the interface name and then click **Configure**.

- 6 Deselect **Enable**.
- 7 Delete the password.
- 8 Click **Update**. This deletes the OSPF authentication password.

Configuring the Flow Collector

A Flow Collector receives and stores NetFlow traffic data from various routers on the network. Traffic information is stored in a database, then processed by the Flow Analyzer, which correlates routing and traffic data, and generates corresponding traffic reports and alerts.

You must configure the Route Recorder as described in [Configuring the Route Recorder](#) on page 50 before you begin configuration of the Flow Collectors or the Flow Analyzer.

To configure a Flow Collector, perform the following steps:

- 1 On the *Recorder Configuration* page of the master unit, click the top-level domain name you added in [Creating a Configuration Hierarchy](#) on page 47 (for example, *CorpNet*).

- 2 Move the cursor to **Add**.

A list of options appears on the blue pop-up menu.

- 3 Click **Traffic** near the bottom of the list.

A status message appears in the browser window indicating that a traffic label was added to the configuration hierarchy. The traffic label is used to organize one or more Flow Collectors, which you will add in Steps 4-11.

- 4 Click the **Traffic** label in the configuration hierarchy and move the cursor to **Add Flow Collector**.

A choice of IP addresses appears. Each IP address corresponds to a RAMS client licensed to function as a Flow Collector.

- 5 Click an IP address.

The *Traffic Configuration* page appears as shown in [Figure 16](#).

Flow Collector Configuration			
Domain	CorpNet		
Flow Collector Id	<input type="text" value="FC1"/>		
UDP Port	<input type="text" value="9991"/>		
Sampling Rate	<input type="text" value="1000"/>		
Prefix Source*	<input type="text" value="(192.168.1.85)"/>		
Routers sending NetFlow data			
Exporter IP	Router IP**	Sampling Rate***	Delete
<input type="text" value="192.168.0.254"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>
<input type="text" value="192.168.3.123"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/>

* Route Recorder or Report Server
 ** corresponding router ip if different from exporter ip.
 *** sampling rate of individual exporter ip if different from global setting.

Figure 16 Traffic Configuration Page

Domain is pre-populated with the name of the administrative domain (*CorpNet*).

- 6 In the *Flow Collector Id* text box, enter a label for the Flow Collector. RAMS uses this label to identify the Flow Collector in the configuration hierarchy tree, and as the name of the Flow Collector database file. This text box is editable until a valid ID is entered. After the ID is validated, the text box is no longer editable, and you cannot change the Flow Collector ID.
- 7 In the *UDP Port* text box, edit the UDP port number if necessary. The number in this text box must match the port number to which the routers will send NetFlow data. The text box is pre-populated with the port from which flow exports are received.
- 8 In the *Sampling Rate* text box, type the NetFlow sampling rate configured on the routers, if any. All routers must use the same sampling rate. If no sampling is done on the routers, type 1.

- 9 In the *Exporter IP* text box, type the IP address of the router from which the Flow Collector will accept NetFlow data exports. The number of routers specified in this section is not limited.
- 10 In the *Router IP*** text box, enter the IP address of the corresponding router if the exporter is not known to the Route Recorder.
- 11 In the *Sampling Rate**** text box, type the Sampling Rate configured on the routers if the sampling rate from the exporter is different from the sampling rate shown in Step 8.
- 12 Repeat steps 9 through 11 to choose additional routers.

Note If you add or delete NetFlow exporting routers using the **Add** or **Delete** buttons, the change will not be preserved until you click **Save** or **Start Recording**.
- 13 Review the information in the *Traffic Configuration* box and click **Save** to apply the configuration. Clicking **Start Recording** also saves the configuration.

The Flow Collector appears in the domain hierarchy.

Starting and Stopping the Flow Collectors

After configuring the Flow Collectors as described in the previous section, you can start or stop recording NetFlow data using the *Recorder Configuration* page of the master unit.

To start recording traffic flow data, perform the following steps:

- 1 If a Flow Analyzer is running in the configuration, you must stop the Flow Analyzer before you can start recording. See on page 69.
- 2 In the configuration hierarchy tree, click the Flow Collector instance where you'll start recording traffic data.

A blue pop-up menu appears.

- 3 Click **View**.

If the Flow Collector is not already recording, the **Start Recording** button is displayed on the *Recorder Configuration* page.

- 4 Click **Start Recording**.

When recording has begun, the **Stop Recording** button appears.

- 5 Restart the Flow Analyzer, if you stopped it in Step 1.
The Flow Collector will write traffic data to the database until you click **Stop Recording**. Changes to Flow Collector configuration are not allowed after recording has started.

Viewing Flow Collector Settings

You can view settings for a configured Flow Collector using the *Recorder Configuration* page on the master unit.

To view Flow Collector settings, perform the following steps:

- 1 On the *Recorder Configuration* page, click the Flow Collector to view in the tree structure on the left side of the page.

A blue pop-up menu appears.

- 2 Click **View**.

A *Status* table appears at the bottom of the *Recorder Configuration* page. If this table is not empty, then a database is being created for RAMS. If the table is empty, verify that the connection to the router or the tunnel is properly configured.

The Flow Collector *Status* table displays each Flow Collector Id and gives its status (*Recording*, *Not Recording*, or *Down*). The table also shows the time the status was last updated, and when the last NetFlow record was received.

You can view additional, more detailed status messages about the traffic database by clicking **Show Detailed Status**. The *Detailed Status* table displays information such as the number of NetFlow packets received by the Flow Collector, the number of NetFlow packets lost, if any, and the number of traffic flows known to the Flow Collector. These details are useful in troubleshooting and are described in [Table 1](#).

Table 1 Detailed Status Table

Time elapsed since last start in seconds	Amount of time in seconds since the last Flow Recorder start.
--	---

NetFlow packets	Number of NetFlow packets processed from all exporting routers. This count is cumulative since the most recent time the recorder was started (or restarted).
Netflow packets from invalid exporters	This counts the number of packets that arrived at the recorder but were dropped because they came from an invalid exporter. If this number is non-zero, it could indicate misconfiguration of the Flow Recorder (or the router), or a malicious outsider attempting to flood the Flow Recorder.
NetFlow records	Number of NetFlow records processed from all exporting routers. This count is cumulative since the most recent time the recorder was started (or restarted).
NetFlow records lost	Number of NetFlow records estimated to be lost since the recorder started (or was restarted). This is estimated by examining the sequence numbers in the NetFlow packets. When recording a new traffic database, a burst of loss at the start of recording is normal (a number less than .01% of the total NetFlow records). Low traffic deployments may see zero lost packets.
Current/Mean/Maximum five-minute count of NetFlow Records	Current/Mean/Maximum number of NetFlow records received from all exporting routers in the past five minutes since the last Flow Recorder start.
NetFlow records in past 5 minutes	Number of records sampled since the recorder started. For low traffic level deployments, this number should be the same as the number of NetFlow records received.
Mean NetFlow records in past 5 minute windows	Mean number of NetFlow records received by the Flow Collector per five-minute interval since the recorder was started (or restarted).
Maximum NetFlow records in past 5 minute windows	Maximum number of NetFlow records received by the Flow Collector per five-minute interval since the recorder was started (or restarted).
NetFlow sequence resets	Number of times RAMS attempted to estimate NetFlow packet loss, usually because a gap in packet sequence numbers was too large.

NetFlow records sampled	<p>Number of records processed through algorithms since the recorder started.</p> <p>For low traffic level deployments, this number should be the same as the number of NetFlow records.</p> <p>For high traffic environments, sampling is done at the level of NetFlow records and this number counts only the records that were selected by the sampling algorithm. (Currently, a cap of 500k records per five minutes).</p>
NetFlow records successfully processed	<p>Number of Netflow records processes since the last Flow Recorder start to generate traffic statistics. In the absence of errors in NetFlow records, this number should be the same as the number of NetFlow records received.</p>
Netflow records dropped due to bad time stamps	<p>Number of Netflow records dropped due to late export (greater than 20 minutes) by the exporting router.</p>
Netflow records dropped due to source or destination not in prefix table	<p>Number of Netflow records dropped due to source or destination address of the reported flow not in the set of routing prefixes heard by the Route Recorders.</p>
Netflow records dropped with multicast information:	<p>Number of Netflow records with multicast information that are received by the Recorder.</p> <p>Multicast Netflow records are currently not processed by the Flow Recorder.</p>
Packets dropped by access list	<p>Number of packets that arrived at the recorder, but were dropped because they came from an invalid source address. If this number is non-zero, it could indicate misconfiguration of the Flow Collector or the router. Alternatively, a malicious intruder could be attempting to flood the Flow Collector.</p>
Input queue length	<p>Instantaneous measure of the amount of data to be processed by RAMS. This number should generally be either 0 or 1.</p>
Maximum input queue length	<p>The maximum value of the input queue length.</p>

Known flows	Instantaneous measure of the number of flows (each flow represents a source or destination prefix/exporter combination) that the Flow Recorder is tracking. This influences the amount of memory being used by the appliance.
Prefixes	Instantaneous measure of the number of prefixes known to the Flow Collector. It uses these prefixes for aggregation purposes. RAMS uses these prefixes when aggregating data. This number should be the count of prefixes known to all of the IGP and BGP areas.
Exporters	Instantaneous measure of the number of exporters on a router (An exporter refers to an interface on a router). For example, in a deployment with two routers each exporting from ten interfaces, this value would be 20.
Next snap_version ID	The next record number to be written to the <i>flowvec</i> table. A new record is written to this table every five minutes, so this number can generally be used to determine how long the current traffic database has been recording.
Writer work queue size	Instantaneous measure of the amount of database traffic that remains to be written to the RAMS disk. A value of 2 or larger indicates a backlog of data waiting to be written to the disk.

To delete a Flow Collector from the configuration hierarchy, perform the following steps:

- 1 Stop the Flow Collector, as described in [Starting and Stopping the Flow Collectors](#) on page 70.
- 2 On the *Recorder Configuration* page of the master unit, click the Flow Collector to delete in the tree.
A blue pop-up menu appears.
- 3 On the pop-up menu, click **Delete**.
- 4 On the confirmation screen that appears, click **Yes**.
This deletes the instance, but the database remains in RAMS as a historical record.

Configuring the Flow Analyzer

A Flow Analyzer processes traffic data received from the Flow Collector(s) and routing data from the Route Recorder. It uses the combined information to generate traffic reports and alerts.

You must configure the Route Recorder and the Flow Collectors as described on [Configuring the Route Recorder](#) on page 50 and [Configuring the Flow Collector](#) on page 68 before you begin configuration of a Flow Analyzer.

To add a Flow Analyzer to a domain hierarchy and configure its settings, perform the following steps:

- 1 On the *Recorder Configuration* page of the master unit, click the top-level domain name you added in [Creating a Configuration Hierarchy](#) on page 47 (for example, *CorpNet*).

- 2 Move the cursor to **Add**.

A list of options appears on the blue pop-up menu.

- 3 Click **Traffic Reports**.

A status message appears in the browser window indicating that a traffic label was added to the configuration hierarchy. The traffic label is used to organize one or more Flow Analyzers, which you will add in the following steps.

- 4 Click the **Traffic Reports** label in the configuration hierarchy and move the cursor to **Add Flow Analyzer**.

The Flow Analyzer's IP address appears.

Note If you have more than one Flow Analyzer configured, then multiple IP addresses will appear. Select the IP address for the Flow Analyzer you are configuring for.

- 5 Click an IP address.

The *Flow Analyzer Configuration* page appears as shown in [Figure 17](#).

The Flow Analyzer Id is the name of the Flow Analyzer (for example, *Reports*). The *Flow Analyzer Id* text box remains editable until a valid label is supplied for the Flow Analyzer. After the label is validated, the Flow Analyzer ID cannot be changed.

- 6 From the **Unselected Databases** list box, choose the databases that will provide the Flow Analyzer with information to generate reports (for example, *Traffic* and *RouteRecorder*).
- 7 Click the “>” button to move the highlighted database to the **Selected Databases** column.

Note You must select a database before you can save the Flow Analyzer configuration.
- 8 In the Data Source section of this page, select from the following two choices:
 - **Get data from Route Recorders**, or
 - **Connect to a Replicating Modeling Engine**, selecting an IP address from its corresponding drop-down menu.
- 9 Click **Save** or **Start Recording** to complete the Flow Analyzer configuration process.

Changes to Flow Analyzer configuration are not allowed after recording has started.

Recorder Configuration

Flow Analyzer Configuration

Flow Analyzer Id

Selected Databases	Unselected Databases
	<div style="font-size: 0.8em;"> A2CorpNet ConfedsTest.ConfedTestBottom ConfedsTest.ConfedTestTop ConfedsTest CorpNet </div>

Data Source

Get data from Route Recorders
 Connect to a Replicating Modeling Engine RME not found

Flow Analyzer Status

Id	Status	Last Status Update
-----------	---------------	---------------------------

Figure 17 Flow Analyzer Configuration Page

Starting and Stopping the Flow Analyzer

In a distributed configuration, you should start the Flow Analyzer last, after the Route Recorder and Flow Collectors are configured and have started recording. If the Flow Analyzer is already running when you begin to record data on the Route Recorder or Flow Collectors, you must restart the Flow Analyzer.

To start and stop the Flow Analyzer, perform the following steps:

- 1 In the configuration hierarchy tree, click the Flow Analyzer instance.

A blue pop-up menu appears.

- 2 Click **View**.

If the Flow Analyzer has not already been started, the **Start** button is displayed on the *Recorder Configuration* page.

- 3 Click **Start**.

When the Flow Analyzer is running, the **Stop** button appears.

- 4 Click **Stop** to stop the Flow Analyzer.

Viewing and Modifying Flow Analyzer Settings

You can view settings for a configured Flow Analyzer using the *Recorder Configuration* page of the master unit.

To view Flow Analyzer settings and modify them, perform the following steps:

- 1 On the *Recorder Configuration* page, click the Flow Analyzer in the tree structure on the left side of the page.

A blue pop-up menu appears.

- 2 Click **View**.

The *Flow Analyzer Status* table displays each Flow Analyzer Id and gives its status (*Up* or *Down*). The table also shows the time the status was last updated.

You must stop the Flow Analyzer as described in on page 69 before you can make changes to the configuration.

To delete a Flow Analyzer from the configuration, perform the following steps:

- 1 Stop the Flow Analyzer, as described in on page 69.
- 2 On the *Recorder Configuration* page of the master unit, click the Flow Analyzer to delete in the tree.

A blue pop-up menu appears.

- 3 On the pop-up menu, click **Delete**.
- 4 On the confirmation screen that appears, click **Yes**.

This deletes the instance, but the database remains in RAMS as a historical record.

Additional Configuration Tasks

This section introduces generic routing encapsulation (GRE) tunnels and describes how they are configured. GRE is used to listen to routing traffic on a network other than the local network.

This section also describes how to configure a loopback interface for Cisco routers. Use a loopback interface to monitor two areas that are linked to a single router.

Configuring GRE Tunnels

To listen to routing traffic on a network other than the local network, there are two options:

- Remotely connect that network to a secondary network interface on RAMS.

You can use a GRE tunnel to form the adjacency.

The GRE tunnel follows standard routing across the network to the destination, so that only the source router and RAMS need to be configured to bring up the tunnel.

In general, a small address block is assigned to the tunnel (usually a /30 address block with four addresses in it). Of these four addresses, the first and last address are the network and broadcast addresses, respectively. The second address is assigned to the router end of the tunnel, and the third to the RAMS end of the tunnel. For example, assume that the address block 10.0.1.0/30 is assigned to the GRE tunnel. The four addresses are allocated as follows:

- 10.0.1.0/30 is the network address.
- 10.0.1.1/30 is the router end of the tunnel.
- 10.0.1.2/30 is the RAMS end of the tunnel.
- 10.0.1.3/30 is the broadcast address.

The tunnel extends from the router to RAMS, and the tunnel must be configured both on the router and on RAMS.

When you configure the router, keep the following points in mind:

- From the router perspective, the tunnel *source* is the IP address of either a loopback or physical address on that router.

- The tunnel *destination* is the IP address of the physical interface on RAMS.
- The IP address of the tunnel has to be assigned to the monitored protocol on the router. For example, there must be a network statement in OSPF for the network IP address of the tunnel.

When you configure RAMS, keep the following points in mind:

- From the perspective of RAMS, the tunnel *source* is the IP address for the physical interface on RAMS.
- The tunnel *destination* is the IP address of the physical interface of the remote router.

To create a GRE tunnel on Cisco destination router, the information in [Table 2](#) must be configured into the router. (This description applies to Cisco routers and may vary for others. Refer to the router documentation for detailed router configuration instructions.)

Table 2 Required Tunnel Information on Remote Router

Item	Example of Command
Tunnel Interface	<code>int tunnel <n></code>
Tunnel IP	<code>ip address 10.0.1.1 netmask 255.255.255.252</code>
Tunnel Source	<code>tunnel source loopback 0</code>
Tunnel Destination	<code>tunnel dest <RAMS IP address></code>
Network Statement	The routing protocol configuration must include a network statement with the assigned IP address of the tunnel and the inverse of the network mask. For example: <code>router ospf <n> network 10.0.1.0 0.0.0.3 area <area to be monitored></code>

To configure a GRE tunnel on RAMS, perform the following steps:

- 1 On the *Recorder Configuration* page, select a protocol instance name on the tree.
- 2 From the blue pop-up menu, click **View** to see the protocol configuration.
- 3 Click **New Tunnel** on the *Recorder Configuration* page.

The Configure Interface section opens, as shown in [Figure 18](#).

Configure Interface		
Interface	OSPF Authentication	GRE Tunnel
Tunnel2OSPF	<input type="checkbox"/> Enable Password: <input type="text"/> MD5 Key-Id: <input type="text"/>	Remote IP Address: (Tunnel Destination) <input type="text" value="192.123.4.5"/> Local IP Address: <input type="text" value="10.1.2"/> Netmask: <input type="text" value="/"/> Interface: <input type="text" value="Slot0/Port1"/>

Figure 18 Configuring a GRE Tunnel Interface

- 4 In the *Interface* text box, type a descriptive name for the tunnel.
- 5 If configuring a tunnel for an OSPF instance, and OSPF Authentication is configured for the area to be monitored, then you must configure authentication by selecting the **Enable** check box.
 - For simple authentication, type a password in the Password text box that matches the password in the remote area.
 - For MD5 authentication, type both a password and an MD5 Key-ID that match the password and key used in the remote area.

If an MD5 key is entered, it is assumed that there is MD5 authentication for this OSPF area. If no MD5 key is entered, then simple authentication is presumed.

- 6 In the *Remote IP Address* text box, type the IP Address of the physical interface or loopback on the remote router.

This IP address should be the same as the **Tunnel Source** address you configured on the router.
- 7 In the *Local IP address* text box, type the IP Address assigned to the tunnel on RAMS.

Using the example addresses listed above, this would be 10.0.1.2.

- 8 In the *Netmask* text box, type the network mask of the **Tunnel IP** address that you configured on the router.

The format for the *Netmask* field can be one of the following:

- CIDR notation (/x)
- Netmask notation (x.x.x.x)

Here is an example based on the addresses used in [Figure 18](#):

- Remote IP Address: 192 . 123 . 4 . 5 (IP address of the physical interface or loopback on the router that is the tunnel destination)
- Local IP Address: 10 . 0 . 1 . 2 (IP address assigned to the tunnel on RAMS)
- Netmask: /30 (mask length for the Tunnel IP address of the tunnel)

9 Click Update.

The protocol instance configuration box returns on the *Recorder Configuration* page with the new tunnel name appearing in the *Not Active* interface list.

10 To begin recording data for the tunnel, move the tunnel name into the Active column.

11 Click Update.

Configuring a Loopback Interface for Cisco Routers

To monitor two areas that are linked to a single router, you must configure a loopback interface on the Cisco router to monitor both areas with RAMS. If you do not do this, only one adjacency will be formed with the remote router. Both areas will come up initially with full adjacency, but the first one to come up will fail soon after. Use the command line interface to the router to configure the loopback interface.

To configure a loopback, perform the following steps:

1 Open a session with the router.

2 Type the following commands:

```
int loopback <n>
ip address <ip address> <netmask>
int tunnel <n>
ip address <tunnel ip address> <netmask>
tunnel source <loopback ip address>
tunnel destination <RAMS ip address>
router ospf <process number>
network <loopback ip address> <inverse netmask> area <a>
network <tunnel ip address> <inverse netmask> area <a>
```

Enabling Technical Support Access

To enable or disable technical support options, locate **System** on the left navigation bar, then click the **Support Access** link. The *Technical Support Configuration* page displays as shown in [Figure 19](#).

Note In a multi-unit configuration, access the *Technical Support Configuration* page of each appliance individually to enable tech support access for that appliance.

The *Technical Support Configuration* page has two buttons that control access to the RAMS by technical support personnel. The first, **Technical support access**, is enabled by default. It allows HP technical support personnel to connect to the RAMS using an SSH connection and a specially encrypted key. Access is restricted to HP technical support personnel and is initiated only after obtaining your permission as part of your requested assistance. To disable technical support access, click **Disable Access**.



Caution Disabling technical support access makes it impossible for HP to reset the password or perform diagnostic services. In this case, you must return the appliance to receive support.

If the RAMS is connected to a network where direct remote access is not possible due to firewall restrictions, the “Technical support callback” feature can be enabled by clicking **Enable Callback**. This feature is disabled by default and your explicit action is required to enable it. Doing so initiates an SSH connection from the RAMS to a dedicated and tightly secured server at HP. Firewall rules usually allow such outbound SSH connections. The connection is configured in such a way that new login sessions can be tunneled from the server at HP through the SSH connection back to the RAMS. As in the case of direct remote access, these login sessions use SSH and require a specially encrypted key.

Note When you enable technical support access, technical support callback is also enabled. When you disable technical support access, technical support callback is disabled as well.

Technical Support Configuration

Technical support access is enabled.

Technical support callback is disabled.

Enabling technical support callback enables technical support access.
Disabling technical support access disables technical support callback.

Figure 19 Technical Support Configuration Page

3 Administration

RAMS administration includes ongoing maintenance tasks such as managing users, updating software, and maintaining databases. You perform these tasks per client, rather than using the *Administration* pages of the master unit. To access the *Administration* pages of an appliance, click **Administration** at the top of the RAMS *Home* page for each unit.

Before performing the administrative tasks described in this chapter, see [Chapter 2, “Configuration and Management”](#) for information about how to access the RAMS Administration pages, log in, and configure RAMS components.

The following topics are included in this chapter:

- [Creating a Traffic Class](#) on page 87
- [Creating Traffic Groups](#) on page 94
- [Configuring the VNC Server](#) on page 99
- [Managing Databases](#) on page 102
- [Backing Up and Restoring Data](#) on page 107
- [Choosing a Data Source](#) on page 116
- [Archiving Data](#) on page 119
- [Updating Software](#) on page 124
- [Creating Daily Reports](#) on page 129
- [Configuring the FTP Server](#) on page 134
- [Managing Users](#) on page 136
- [Viewing and Exporting RAMS Log Pages](#) on page 140
- [Uploading Layout Backgrounds](#) on page 142
- [Using Diagnostic Functions](#) on page 143

- [Shutting Down the RAMS Appliance](#) on page 144

Note The links and buttons on the *Administration* pages might differ from those shown in the examples in this guide, depending on the functions for which your system is licensed. If your system is not licensed for a particular function, the links or buttons related to that function do not appear on the *Administration* pages.

Creating a Traffic Class

Traffic class is a set of traffic belonging to one or more Class of Service (CoS), as defined by Type of Service (ToS) byte or Diff-Serv Code Point (DSCP) settings in the IP header. Traffic classes can be a component of a traffic group definition.

If a traffic class is specified, it will match a subset of traffic associated with a particular Class of Service defined by DSCP or TOS bits in the IP header. A traffic explorer deployment can only have DSCP or TOS specified, but not both.

Note You must create traffic classes on the Master or Standalone appliance.

To access *Traffic Classes*, locate *Traffic* on the left Navigation bar, then click on **Traffic Classes**.

The *View DSCP Traffic Classes* page opens as shown in [Figure 20](#).

View DSCP Traffic Classes

DSCP ▼
 TOS

DSCP Value ▼	Name ▼
12	AF12
14	AF13
18	AF21
20	AF22
22	AF23
26	AF31
28	AF32
30	AF33
34	AF41
36	AF42
38	AF43
0	CS0
8	CS1
16	CS2
24	CS3
32	CS4
40	CS5
48	CS6
56	CS7
46	EF
10	Lomond

Figure 20 View DSCP Traffic Classes Page

There are two traffic class modes:

- **DSCP**—The default mode for traffic classes. You can edit traffic class names in DSCP mode, but you cannot create them.
- **TOS**—You can create traffic classes in this mode to monitor valid TOS byte values in the IP header.

To create traffic classes in TOS mode:

- 1 Toggle to **TOS** mode located at the top of the *Traffic Classes* page.

The *Traffic Classes Warning Dialog* page opens. This page cautions that if you change modes, all previous information input for related Traffic Classes and Traffic Groups will be deleted.

- 2 Click **Submit** to change to TOS mode.

The *View TOS Traffic Classes* page opens.

- 3 Select **New** to create a TOS traffic class.

The *Create TOS Traffic Classes* page opens, as shown in [Figure 21](#).

Traffic Classes




Figure 21 shows the 'Create TOS Traffic Classes' form. The form is titled 'Create TOS Traffic Classes' and contains the following fields:

- Name: Text input field.
- Precedence (Bit 0-2): Dropdown menu with '000 Routine' selected.
- Delay (Bit 3): Dropdown menu with '0 Normal' selected.
- Throughput (Bit 4): Dropdown menu with '0 Normal' selected.
- Reliability (Bit 5): Dropdown menu with '0 Normal' selected.
- Monetary Cost (Bit 6): Dropdown menu with '0 Normal' selected.

At the bottom of the form, there are two buttons: 'Submit' and 'Cancel'.

Figure 21 Create TOS Traffic Classes Page

- 4 Choose from a list of select values from the corresponding drop-down menus for the following fields:
 - Name
 - Precedence
 - Delay
 - Throughput
 - Reliability
 - Monetary Cost

Note You can only use alpha-numeric (upper and/or lower case) and underscore (“_”) to name the class.

Note Be careful not to replicate bit values for the classes your are creating. Traffic classes should have distinct sets of values.

- 5 Click **Submit** to save the information.

After you create a traffic class, the *View TOS Traffic Classes* will display a column for deleting the traffic class.

Editing a Traffic Class

Note The only values you can edit are the names for the traffic classes.

To edit traffic classes in DSCP mode:

- 1 Select the **Edit** button at the bottom of the *Traffic Classes* page.

The *Edit DSCP Traffic Classes* page opens, as shown in [Figure 22](#).

Traffic Classes

Edit DSCP Traffic Classes

DSCP Value	Name
12	AF12
14	AF13
18	AF21
20	AF22
22	AF23
26	AF31
28	AF32
30	AF33
34	AF41
36	AF42
38	AF43
0	CS0
8	CS1
16	CS2
24	CS3
32	CS4
40	CS5
48	CS6
56	CS7
46	EF
10	Lomond

Figure 22 Edit DSCP Traffic Classes Page

- 2 Under the *Name* field, change the name of class or classes you wish to rename.
- 3 Click **Submit** to save the information.

Deleting a Traffic Class

Caution Deleting a traffic class that is part of a traffic group(s) will also delete the traffic group(s) that are associated with it.

To Delete a Traffic Class

- 1 Select **Traffic Classes** from the Navigation Bar.
The *View DSCP Traffic Class* page opens.
- 2 Toggle to TOS mode.
- 3 Find the traffic class you want to delete and click on **delete**.
The *Warning* dialog box opens, advising that the traffic class will be deleted.
- 4 Click **Submit** to delete the class. Clicking **Cancel** will return you to the *View Tos Traffic Classes* page.

Creating Traffic Groups

Traffic Explorer 2.0 increases the visibility into the traffic flowing through the network core by classifying traffic into user defined Groups.

The network admin can create groups in a flexible manner. The rules of the group can be a combination of Source/Destination prefixes, TCP/UDP ports, an IP protocol, and Traffic Classes.

The traffic group then matches a subset of network traffic from or to specific locations per application(s) and/or class(es) of service.

Further, if a traffic class is specified, it matches a subset of traffic associated with a particular Class of Service defined by TOS or DSCP bits in the IP header. A TEX deployment can only have TOS or DSCP at one time, not both.

Once the appropriate groups are defined, network administrators can track traffic:

- To or from specific application servers (based on prefixes)
- For specific services like "Voice Premium" or "Voice Gold" (based on ToS and DSCP: AF or EF classes)
- Determine (via reports and GUI) per traffic flow, the application or CoS that it belongs to
- Alert (SNMP Trap) when %utilization (or bps) exceeds (or falls below) a specified threshold per traffic group.

To access *Traffic Groups*, locate *Traffic Groups* on the left Navigation bar, then click on **Traffic Groups**.

The *View Group* page opens, as shown in [Figure 23](#).

Traffic Groups

The screenshot shows the 'View Group' page. On the left, there is a list of traffic groups: '1 Davinded' (selected) and '2 Bitterenz'. Below the list is a 'New' button. On the right, the configuration details for the selected group 'Davinded' are displayed in a table:

Traffic Group Name	Davinded
Source prefix	192.168.0.0/16
Destination prefix	any
Protocol	EIGRP
Classes	ANY

Below the table are several buttons: 'Edit', 'Change Order *', 'Configure Alerts', 'Disable', 'Delete', and 'Copy'. At the bottom, a note states: '* Traffic is classified to the first matching group in the ordered list.'

Figure 23 View Group Page

The following is a list of the buttons you can select from this page:

New—Select this to create a new traffic group.

Edit—Select this to edit a previously created traffic group.

Change Order*—Select this to set the priority of the traffic group the flow recorder will track traffic for.

Configure Alerts—Select this to configure traffic group alerts. See [Chapter 13, “Alerts”](#) for more information.

Disable—Select this to disable a group.

Delete—Select this to delete a group.

Note The Traffic Explorer assigns the lowest priority to the most recently created traffic group.

To Create Traffic Groups:

- 1 In the *View Traffic Groups* page, click **New**.

The *New Group* page appears.

Traffic Groups

New Group

Traffic Group Name

Source prefix

Destination prefix
192.168.0.0/16,
10.0.0.1/32

Source ports

and or

Destination ports
1-100,200,300

Protocol

Classes

All Selected

Bittern
Bittor

Submit Reset Cancel

Figure 24 New Group Page

The following is a list of fields you will need to enter information to create the traffic group:

- **Traffic Group Name:** Field where you name the traffic group.

Note When naming traffic groups, use upper and lower case alpha-numeric, ampersand (“&”) and underscore (“_”) to name the values.

- **Source Prefix:** Enter one or more source IP prefixes.

- **Destination Prefix:** Enter one or more destination prefixes, separated by commas.
- **Source Ports:** Enter the source port range (e.g., 1-200, 300)
- Choose either the **and** or the **or** radio button
- **Destination Ports:** Enter the destination port range.
- **Protocol:** Select **Any** or a protocol from the drop-down menu.

In the *Classes* area of this page, you can select the traffic class you want to be a part of the traffic group you are creating by clicking once on the class, and then clicking on the arrow button to move the class to the *Selected* category.

Below this field are the following buttons:

Submit: Select this to save the information you entered.

Reset: Select this to bring back the values you had before creating this group.

Cancel: Select this to cancel this operation.

- 2 Enter the name, source and destination prefix, source and destination ports, and the protocol for the traffic group you are creating.
- 3 In the *Classes* section, select the traffic classes you want to be part of the traffic group you are creating. You can do this by clicking once on the class, and then clicking the arrow button to the *Selected* category.
- 4 Click **Submit** to save the information.

Editing Traffic Groups

To Edit Traffic Groups:

- 1 In the *View Traffic Group* page, and in the left navigational pane, select the traffic group you wish to edit by clicking once on the group, and select **Edit**.

The *Edit Group* page opens.

Traffic Groups

The screenshot shows the 'Edit Group' page. On the left, a list of traffic groups is displayed: '1 Davinded' with a green checkmark and '2 Bitterenz' with a blue checkmark. Below the list is a 'New' button. The main area is titled 'Edit Group' and contains the following fields:

- Traffic Group Name: Bitterenz
- Source prefix: any
- Destination prefix: any
- Source ports: any
- Destination ports: 2000
- Protocol: ICMP 1

Below these fields is a 'Classes' section with two columns: 'All' and 'Selected'. The 'All' column contains 'Bittern' and 'Bitter'. There are 'Add' and 'Remove' buttons between the columns. At the bottom of the page are 'Submit', 'Reset', and 'Cancel' buttons.

Figure 25 Edit Group Page

Note The history for the edited traffic group may not match the new definition

- 2 Edit the fields you need to revise.
- 3 Select **Submit** to save the information.

Configuring the VNC Server

If you plan to use a VNC viewer to access RAMS (see [Chapter 4, “RAMS Viewers”](#)), you need to configure the VNC server before users can access RAMS using the VNC viewer on a desktop machine. Alternatively, if all RAMS users will use the X Window System rather than VNC, you might want to stop the VNC server.

Note VNC server configuration applies only to RAMSs operating with a GUI license key.

You can log into a persistent VNC session that is shareable by multiple operators, or you can log into a session created on demand if both session types are enabled.

To access the *VNC Configuration* page, locate **Application** on the left navigation bar, then click **VNC Configuration** to display the page shown in [Figure 26](#).

Window Size	Colors	Share Session
1024x768	True Color (24 bit)	Enable

VNC Authentication	
Password: <input type="text"/>	Confirm password: <input type="text"/>

VNC Displays 2 and Higher

Figure 26 VNC Server Configuration Page

Note The VNC server consumes system resources, even when no sessions are active and VNC is not configured for sharing. Start the VNC server only when necessary.

You can configure the following on the *VNC Configuration* page:

- The window size, which is the size of the virtual screen of the VNC viewer.
- The number of colors displayed.
- The settings to enable or disable session sharing.

- The VNC authentication password for the persistent session (VNC display 1).
- The availability of on-demand sessions using VNC displays 2 and higher.

All users who access RAMS using VNC share the settings configured on this page.

After configuring the VNC settings, set a VNC authentication password. Type the password information in the appropriate text boxes in the *VNC Authentication* table and then click **Update**.

After you click **Start** to start the VNC server, the button toggles to **Stop**. If VNC server changes are required in the future, stop the VNC server before making the changes.

Note If you change the VNC authentication password after the VNC server is started, you must stop and restart the VNC server for the new setting to take effect.

Connecting to the VNC Persistent Session

After you start the VNC viewer, type *hostname:1*, and then log in using the password you entered in the *VNC Configuration* page. The VNC session displays using the window size, colors, and sharing properties specified in the *VNC Configuration* page for VNC display 1.

If you enable sharing of the persistent session, anyone with the VNC password can access the currently active session and issue commands from their desktop. Multiple users can connect at the same time and take turns controlling the user interface to jointly work on a problem. If you disable sharing, only one person at a time can connect to VNC display 1 and all other users are locked out.

When you disconnect from the persistent VNC session, the RAMS user interface continues to run. If you connect to the session again later, it resumes as you left it unless someone else has connected and made changes in the meantime.

Note If the network connection drops while someone is connected with sharing disabled, the VNC server might refuse to allow new connections. To restore VNC access to RAMS, use the **Stop/Start** button on the *VNC Configuration* page to stop and restart the VNC Server.

Using On-Demand VNC Sessions

RAMS allows multiple operators to log into their own sessions that are initiated on demand and not shared.

After you start the VNC viewer, specify the VNC display window size by typing the hostname, a colon, and the VNC display number that corresponds with the window size you want, as shown in [Table 3](#). For example, type *hostname:8* to connect to the VNC session with a 1280 × 1024 pixel window size. All of the window sizes will display in 24-bit true color. The VNC viewer will display a login window where you should log in using a user name and password you configured on the *User Administration* page (see [page 136](#)).

Table 3 VNC Configuration Window Settings

VNC Display Number	Window Size (Pixels)
1	Configure on VNC Configuration Page
2	1016 × 700
3	1024 × 768
4	1152 × 768
5	1152 × 864
6	1152 × 900
7	1280 × 864
8	1280 × 1024
9	1400 × 1050
10	1600 × 1200

If you and another user connect to the same display number, such as *hostname:8* as in the example above, RAMS creates separate VNC sessions so that both users can operate independently. RAMS allows multiple users to connect to on-demand VNC sessions up to the user count limit displayed on the *License* page.

Note When you disconnect from an on-demand VNC session, RAMS terminates the session immediately.

Managing Databases

Use the *Database Administration* page to delete, rename, or trim an existing database on a RAMS appliance. To access this page, locate **Maintenance** on the left navigation bar of an individual unit, then click **Databases**. The *Database Administration* page opens as shown in [Figure 27](#).

Database Administration

Offline Databases

	Administrative Domain	Protocol	Area ID	Size
<input type="checkbox"/>	A2CorpNet	EIGRP	1	3.6M
<input type="checkbox"/>	A2CorpNet	OSPF	0.0.0.1	50M
<input type="checkbox"/>	A2CorpNet	ISIS	490001	4.3M
<input type="checkbox"/>	A2CorpNet	ISIS	Backbone	4.3M
<input type="checkbox"/>	ConfedsTest	BGP	AS6003	236K
<input type="checkbox"/>	ConfedsTest.ConfedTestBottom	BGP	AS6003	1.9M
<input type="checkbox"/>	ConfedsTest.ConfedTestTop	BGP	AS6001	2.0M
<input checked="" type="checkbox"/>	CorpNet	EIGRP	1	23M
<input checked="" type="checkbox"/>	CorpNet	ISIS	4700230001000000000020001	2.7M
<input type="checkbox"/>	CorpNet	BGP	AS65522	112M
<input type="checkbox"/>	CorpNet	BGP	AS65522/VPN	535M
<input type="checkbox"/>	ISPtraf1	ISIS	498888096018253000	35M
<input type="checkbox"/>	ISPtraf1	ISIS	Backbone	952M
<input type="checkbox"/>	ISPtraf1	BGP	AS8888	1.5G
<input type="checkbox"/>	ISPtraf1	Traffic	fr1	4.6G
<input type="checkbox"/>	ISPtraf1	TRS	fa1	476M
<input type="checkbox"/>	UCBJul03a	OSPF	169.229.128.128	3.9M
<input type="checkbox"/>	UCBJul03a	OSPF	169.229.128.168	3.8M
<input type="checkbox"/>	UCBJul03a	OSPF	169.229.128.176	3.9M
<input type="checkbox"/>	UCBJul03a	OSPF	Backbone	5.0M
<input type="checkbox"/>	UCBJul03a	BGP	AS25	379M

New Top-Level Administrative Domain:

12148MB of **106528MB** used on disk (**11%**)

Rename Selected Databases

Delete Selected Databases

Archive Selected Databases

(*) In use databases – cannot be archived, deleted or renamed. Possible solutions - quit all instances of gui client, including stopping the VNC server

Figure 27 Database Administration Page

Database administration is usually performed for housekeeping reasons — for example, when the RAMS disk is becoming full. Deleting unneeded databases and trimming active databases can help you to gain disk space. Since each database is stored and managed per-unit, you must access RAMS appliances individually to perform these administrative tasks on the database.

It is recommended that you run an FTP application to connect to the RAMS FTP file storage area, and then use the delete command to delete any unneeded files. (See [Configuring the FTP Server](#) on page 134 for information about the FTP file storage area.)

The *Database Administration* page consists of two sections:

- The Offline Databases section lists administrative domains that are not currently being recorded.
- The Online Databases section lists administrative domains that are currently being recorded. See [Figure 28](#) to view this portion of the window.

Online Databases

	Administrative Domain	Protocol	Area ID	Size
*	CorpNet	OSPF	0.0.0.1	38M
*	CorpNet	ISIS	Backbone	6.5M
*	LabRight.ConfedsTest	OSPF	Backbone	6.5M
*	LabRight.ConfedsTest.ConfedTestBottom	BGP	AS65520	12M
*	LabRight.ConfedsTest.ConfedTestBottom	BGP	AS65520/VPN	9.3M
*	LabRight.ConfedsTest.ConfedTestTop	BGP	AS65510	12M
*	LabRight.ConfedsTest.ConfedTestTop	BGP	AS65510/VPN	9.3M

9 weeks recorded; more than one year of recording capacity at current rate;	Number of weeks to trim: <input type="text" value="1"/>	<input type="button" value="Trim Online Databases"/> Trimming databases may take some time, depending on the amount of data to be removed.
---	---	---

(*) Online database – cannot be archived, deleted or renamed.

Figure 28 Online Databases Page

Note Before performing any database operations, you must quit all running instances of the RAMS client, including stopping the VNC server.

Using Offline Databases

Offline databases contain data from administrative domains that are not currently being recorded. These databases can be deleted to save disk space, or renamed if the administrative domain name changes.

Deleting a Database

To permanently delete a database, perform the following steps:

- 1 Stop all instances of the RAMS client.
- 2 On the top navigation bar, click **Recorder Configuration** to display the *Recorder Configuration* page.
Note In a distributed configuration, use the *Recorder Configuration* page of the master unit, which displays a tree representing all units in the configuration.
- 3 If the database you want to delete is currently recording, stop recording on that database.
- 4 Click the node on the tree that corresponds to the database to be deleted and choose **Delete** from the blue pop-up menu.
- 5 If you are not working with a distributed configuration, proceed to Step 8. Otherwise, click **Units** on the left navigation bar to access the *Units* page of the master unit and a list of clients in the configuration.
- 6 In the Client Status section of the *Units* page, click the client where the database was recorded to access the *Home* page of that client.
- 7 On the *Home* page of the client, click **Administration** on the top navigation bar.
- 8 Click **Databases** on the left navigation bar.
- 9 Select the check box(es) to the left of the database(s) to be deleted.
- 10 Click **Delete Selected Databases**.
A confirmation page opens displaying the names of the selected databases.
- 11 Click **Yes** if you want to delete the selected databases.
The selected databases are deleted.

To delete an existing database and start recording to a new database, perform the following steps:

- 1 Stop all instances of the RAMS client.
- 2 Navigate to the *Home* page of the client where the database is being recorded and click **Administration** on the top navigation bar.
- 3 Click the **Recorder Configuration** link on the top navigation bar to display the *Recorder Configuration* page.
- 4 If the database you want to delete is currently recording, stop recording on that database
- 5 Click **Databases** on the left navigation bar.
- 6 Select the check box(es) to the left of the database(s) to be deleted.
- 7 Click **Delete Selected Databases**.
A confirmation page opens displaying the names of the selected databases.
- 8 Click **Yes** if you want to delete the selected databases.
The selected databases are deleted.
- 9 Navigate back to the *Recorder Configuration* page and start recording on the configuration whose database you deleted.
A new database is created for that configuration and recording begins.

Renaming Databases

To rename one or more databases, perform the following steps:

- 1 Stop all instances of the RAMS client.
- 2 Navigate to the *Home* page of the client where the database is being recorded and click **Administration**.
- 3 Click the **Recorder Configuration** link on the top navigation bar to display the *Recorder Configuration* page.
- 4 If the databases you want to rename are currently recording, stop recording on these databases.
- 5 Click **Databases** on the left navigation bar.
- 6 Select the check box(es) to the left of the database(s) to be renamed to a single, new name (typically, only databases that have a common first-level administrative domain name already).

- 7 Type the new name in the *New Top-Level Administrative Domain* text box. Only the first-level administrative domain name can be changed, meaning, no period is allowed in the new name. Names must begin with an alphabetic character and can contain only alphanumeric characters.
- 8 Click **Rename Selected Databases**.
A confirmation page opens displaying the names of the selected databases.
- 9 Click **Yes** to confirm that you want to rename the selected databases. The top-level name of the selected databases is renamed to the new top-level administrative domain.

Using Online Databases

The Online Databases section of the *Database Administration* page lists all of the Administrative Domains that are currently being recorded. Below the list is a field that indicates how many days of recording capacity remain, based on the current growth rate of the most active of the databases, and a field that indicates the number of days remaining until trimming is needed. The default value of the *Number of Weeks to Trim* field is 7. You can specify a different interval.

When you trim the online databases, RAMS deletes data from the earliest point in the databases until the aggregate database size allows for a number of days' growth specified in the *Number of Weeks to Trim* text box. All databases are trimmed so that the data begins at approximately the same date.

To trim a database, perform the following steps:

- 1 Navigate to the *Home* page of the client where the database is being recorded and click **Administration**.
- 2 Locate **Maintenance** on the left navigation bar and click **Database Administration**.
- 3 On the *Database Administration* page, click **Trim On-line Databases**.

If the databases contain data older than the date implied by the *Number of Weeks to Trim* field, the databases are trimmed, and a confirmation message displays on the *Database Administration* page.

If the databases do not contain data older than the date implied by the *Number of Weeks to Trim* field, the message "No action required" displays.

Backing Up and Restoring Data

RAMS stores recorded route topology in databases, and also stores system configuration files. The *Backup and Restore* page allows you to selectively save any or all data files, including databases and system configuration files, to one of two storage systems:

- The RAMS appliance hard disk.
- A remote storage system, such as your desktop computer. If you back up to a remote storage system, you must configure Server Message Block (SMB) information for that system. See [Enabling SMB and Adding a Remote Server](#) on page 113.

You can save the backup file through the RAMS administration interface or by using File Transfer Protocol (FTP). Since most web browsers do not allow file transfers larger than 2-4 GB, it is recommended that you use FTP to manage large backup files. See [Transferring Backup Files Using FTP](#) on page 115.

RAMS saves the data files into a single backup file named `backup.dat`. The system saves the backup file with the date and time in your local time zone.

Note If you back up only to the RAMS, the appliance contains only one backup file. After you create a new backup file, this new file automatically overwrites the old backup file.

Creating a Backup File

During the backup file creation process, RAMS archives the selected databases and/or system configurations, and then attaches a metadata header into the backup file. This metadata header contains database and/or system configuration information as well as general information about the backup file, including the following items:

- The RAMS OS version number.
- The RAMS unit ID.
- The schema version number.
- The backup creation date and time.

This information opens in the *Restore from Backup on This Unit's Disk* section at the bottom of the *Backup and Restore* page. See [Restoring a Backup File](#) on page 111.

The system configuration information contains the following items:

- The recorder configuration information.
- System licenses.
- All system logs.
- All system layouts.
- User account information.

To open the *Backup and Restore* page, locate **Maintenance** on the left navigation bar, then click **Backup and Restore** (shown in [Figure 29](#)).

Backup and Restore

Create Backup on This Unit's Disk

	Administrative Domain	Protocol	Area ID
<input type="checkbox"/>	System Configuration		
	Data		
<input type="checkbox"/>	A2CorpNet	eigrp	1
<input type="checkbox"/>	A2CorpNet	isis	490001
<input type="checkbox"/>	A2CorpNet	isis	Backbone
<input type="checkbox"/>	A2CorpNet	ospf	0.0.0.1
<input type="checkbox"/>	ConfedsTest	bgp	AS6003
<input type="checkbox"/>	ConfedsTest.ConfedTestBottom	bgp	AS6003
<input type="checkbox"/>	ConfedsTest.ConfedTestTop	bgp	AS6001
<input type="checkbox"/>	CorpNet	bgp	AS65522
<input type="checkbox"/>	CorpNet	bgp	AS65522/VPN
<input type="checkbox"/>	CorpNet	ospf	0.0.0.1
<input type="checkbox"/>	CorpNet	eigrp	1
<input type="checkbox"/>	CorpNet	isis	4700230001000000000020001
<input type="checkbox"/>	CorpNet	isis	Backbone
<input type="checkbox"/>	ISPtraf1	bgp	AS8888
<input type="checkbox"/>	ISPtraf1	isis	498888096018253000
<input type="checkbox"/>	ISPtraf1	isis	Backbone
<input type="checkbox"/>	ISPtraf1	traffic	fr1
<input type="checkbox"/>	ISPtraf1	trs	fa1
<input type="checkbox"/>	LabRight.ConfedsTest	ospf	Backbone
<input type="checkbox"/>	LabRight.ConfedsTest.ConfedTestBottom	bgp	AS65520
<input type="checkbox"/>	LabRight.ConfedsTest.ConfedTestBottom	bgp	AS65520/VPN
<input type="checkbox"/>	LabRight.ConfedsTest.ConfedTestTop	bgp	AS65510
<input type="checkbox"/>	LabRight.ConfedsTest.ConfedTestTop	bgp	AS65510/VPN
<input type="checkbox"/>	UCBJul03a	bgp	AS25
<input type="checkbox"/>	UCBJul03a	ospf	169.229.128.128
<input type="checkbox"/>	UCBJul03a	ospf	169.229.128.168
<input type="checkbox"/>	UCBJul03a	ospf	169.229.128.176
<input type="checkbox"/>	UCBJul03a	ospf	Backbone

Create Backup

Figure 29 Backup and Restore Page

This page contains two sections:

- The **Create Backup on This Unit's Disk** section, which contains a list of all system configurations and databases in the system. The table orders information in three columns:
 - **Administrative Domain:** The system configuration and/or database administrative domain name.
 - **Protocol:** The protocol the file uses.
 - **Area ID:** The file area ID.
- The **Restore from Backup on This Unit's Disk** section, shown in [Figure 30](#), contains the backup file information and the database(s) and system configuration contained in that file.

To create a backup file, perform the following steps:

- 1 On the *Backup and Restore* page, select one or more check boxes to the left of the database and/or system configuration domain name that you want to back up.

Note You cannot back up or restore an actively recording database. An actively recording database opens in the list in green text and an asterisk in place of the check box. Stop database recording in the *Recorder Configuration* page. See [Updating Software](#) on page 124.

- 2 Click **Create Backup**.

The backup process begins. Depending on the size of the database or configuration, the backup process can take several minutes. The *Backup and Restore* page opens and periodically updates with the file size as it progresses.

- 3 When RAMS finishes creating the backup file, the **Finished** button opens on the *Backup and Restore* page. Click **Finished** to continue.

The *Backup and Restore* page displays the new backup file(s) in the *Restore from Backup on This Unit's Disk* table at the bottom of the page, as shown in [Figure 30](#).

Restore from Backup on This Unit's Disk

Backup Info	
Backup Date:	Tue Jul 4 13:22:39 CEST 2006
Software Version	4.5.8-C
OS Version	1.2.2
Unit ID:	003048884326
Size (in bytes)	3051520

	Administrative Domain	Protocol	Area ID
<input type="checkbox"/>	IssyLabRiv	isis	490751
<input type="checkbox"/>	IssyLabRiv	bgp	AS3215

Overwrite layouts and other properties

Restore Selection

Upload Backup File

Figure 30 Table for Restore from Backup on This Unit's Disk

Note Depending on the size of the database and on the Login Linger Timer settings, the *Administration* login might time out before the backup completes. If the **Finished** button does not appear after approximately 15 minutes, log in again.

The *Backup Info* table appears above the backup file table. This table contains the following backup file information:

- **Backup Date:** The backup date and time.
 - **Software Version:** The RAMS software version used to create the backup file.
 - **OS Version:** The operating system (OS) version used to create the backup file.
 - **Unit ID:** The unit ID.
 - **Size (in bytes):** The size of the backup file in bytes.
- 4 If you stopped recording before you backed up, start recording data again in the *Recorder Configuration* page.

Saving a Backup File

After creating a backup file, you can save the file to another location.

To save a backup file, perform the following steps:

- 1 After creating a backup file, click **Backup and Restore**.
- 2 On the *Backup and Restore* page, click the **Download Backup File** button.
A *File Download* dialog box appears.
- 3 Click **Save**.
- 4 Type a filename for the backup or accept the default filename (`backup.dat`), and then select a destination for the backup file.
- 5 Click **OK**. The system saves the backup file to the destination you specified in Step 3.

Uploading a Backup File

If you downloaded a backup file that you want to restore, you must upload the file before you restore it.

To upload a backup file, perform the following steps:

- 1 Go to the *Backup and Restore* page.
- 2 Click **Upload Backup File**, which appears below the Restore from Backup on This Unit's Disk section.

The *Restore Backup from File* window appears.

- 3 Type the location of the backup file in the *Backup Filename* text box, or browse for the file by clicking **Browse**.
- 4 Click **Upload File**.

The upload process begins. Depending on the size of the database or configuration, the upload process can take several minutes.

The database(s) and/or system configuration in the uploaded backup file appear in the *Restore from Backup on This Unit's Disk* table.

Restoring a Backup File

The Restore from Backup on This Unit's Disk section of the *Backup and Restore* page contains the most recent backup file information in the table.

To restore a backup file, perform the following steps:

- 1 Upload the backup file if necessary as described in [Uploading a Backup File](#) on page 111.
- 2 On the *Backup and Restore* page, select one or more checkboxes to the left of the database and/or system configuration domain name that you want to back up.

Note You cannot back up or restore an actively recording database. Stop database recording in the *Recorder Configuration* page. See [Updating Software](#) on page 124

- 3 Select the **Overwrite layouts and other properties** check box if you want to restore a backup set of layouts and properties.
- 4 Click **Restore Selection**.

The restore process begins. Depending on the size of the database or configuration, the restore process can take several minutes. The *Backup and Restore* page appears and periodically updates with the file size as it progresses.

Note If you are restoring the system configuration, a warning appears saying the system will reboot after the restore is complete. The configuration is restored and a message appears indicating that the system is rebooting. If RAMS does not respond to the browser, wait approximately three minutes for the boot process to complete, and log in again.

The *Backup and Restore* page displays the restored backup file in the *Restore from Backup on This Unit's Disk* table at the bottom of the page. If you have more than one backup table in this section, the restored backup file appears in the top table.

- 5 If you want to append new data to the restored database, start recording again in the *Recorder Configuration* page.

Deleting a Backup File

You can delete the backup file from the RAMS hard disk to free up space.

To delete the backup file from the RAMS, perform the following steps:

- 1 Go to the *Backup and Restore* page.

- 2 Click **Delete Backup File**, which appears below the *Create Backup on This Unit's Disk* table.
- 3 In the dialog box that appears, click **YES**.

The backup file information no longer appears in the Restore from Backup on This Unit's Disk section.

Enabling SMB and Adding a Remote Server

RAMS 5.20 uses the Common Internet File System (CIFS) implementation of the Server Message Block (SMB) to create backup files on, and restore backup files from, a remote storage server. You can enable SMB in RAMS 5.20 to create more than one backup database and/or system configuration file on a remote storage server.

Note RAMS checks remote server reachability every time you open the *Backup and Restore* page. If the system cannot open the remote server, then the *Backup and Restore* page displays the backup and restore information on the RAMS appliance.

To enable SMB and add a remote server, perform the following steps:

- 1 On the left navigation bar, locate **System** and click **Archival Configuration**. The *Archival Configuration and Remote Storage* page appears as shown in [Figure 32](#).
- 2 Select the **Enable SMB** check box to create or restore backup files from a remote storage server.
- 3 Type the remote server login user name in the *Remote user name* text box.
Note If the login is for a guest only, then leave the user name blank.
- 4 Type the remote server password in the *Remote password* text box.
- 5 Type the remote server IP address in the *Remote server IP* text box.
- 6 Type the remote server share name in the *Share name* text box.
- 7 Click **Update**.

The remote storage server information appears in the Restore from Backup on This Unit's Disk section of the *Backup and Restore* page.

Restoring a Backup File from a Remote Server

RAMS saves the data files into a single backup file named `backupYYMMDDHHSS_UnitID.dat`, where `YYMMDDHHSS` is the date and time the file was saved and `UnitID` is the server unit ID. For example, `0506251345_00304870B6B0.dat` specifies the file was saved on June 25, 2005 at 1:45 p.m. with the unit ID of `00304870B6B0`.

Note If you back up to a remote storage server, you only have access to the backup files on the remote server, not on the RAMS appliance.

If you enabled SMB to save multiple backup files, all backup files appear in the Backup to SMB Storage section of the file backup table. The Restore from SMB Storage section appears underneath the file backup table

This section displays the information of each backup file in a separate table. The file information of the table that was backed up most recently appears at the bottom of the section. The table contains the following information about the file:

- **Backup Date:** The date and time the file was backed up.
- **Software Version:** The RAMS software version used to create the backup file.
- **OS Version:** The operating system (OS) version used to create the backup file.
- **Unit ID:** The unit ID.
- **Size (in bytes):** The size of the backup file in bytes.
- **Contents:** Displays the database name(s) and/or indicates that the backup file includes the system configuration.

To restore a backup file from a remote server, perform the following steps:

- 1 On the *Backup and Restore* page, select the option button in the left column of the backup file.
- 2 Click **Select Backup File**.

The backup file database(s) and/or system configuration appears in the file backup table at the bottom of the Restore from SMB Storage section.

- 3 Select one or more check boxes to the left of the database and/or system configuration domain name that you want to back up.

Note You cannot back up or restore an actively recording database. Stop database recording in the *Recorder Configuration* page. See [Updating Software](#) on page 124.

- 4 Select the **Overwrite layouts and other properties** check box if you want to restore a backup set of layouts and properties.
- 5 Click **Restore Selection**.

The restore process begins. Depending on the size of the database or configuration, the restore process can take several minutes. The *Backup and Restore* page appears and periodically updates with the file size as it progresses.

Note If you are restoring the system configuration, a warning appears saying the system will reboot after the restore is complete. The configuration is restored and a message appears indicating that the system is rebooting. If RAMS does not respond to the browser, wait approximately three minutes for the boot process to complete, and log in again.

The *Backup and Restore* page displays the restored backup file in the *Restore from Backup on This Unit's Disk* table at the bottom of the page. If you have more than one backup table in this section, the restored backup file appears in the top table.

- 6 If you want to append new data to the restored database, start recording again in the *Recorder Configuration* page.

Transferring Backup Files Using FTP

Most web browsers cannot manage transfers of files that are 4 GB and larger. Therefore, RAMS gives you the option of saving the `backup.dat` file to the RAMS appliance hard disk using File Transfer Protocol (FTP). You must enable the FTP server to transfer backup files using FTP. See [Configuring the FTP Server](#) on page 134.

Choosing a Data Source

RAMS 5.20 provides you with the capability to have multiple recorders geographically and topologically distributed in remote parts of the network. This provides you with a global view of the entire network topology from a local copy of the database, which will minimize delay. Data replication is enabled by default on the Modeling Engine.

Using the *Data Source Configuration* page, you can have the RAMS unit that is licensed to act as a Modeling Engine to replicate data from the Route Recorders in your distributed configuration. This centralized Modeling Engine consolidates the data that has been collected across the network to a centralized location. You can retrieve network-wide information from a local source.

The *Data Source Configuration* page allows you to specify where RAMS obtains routing data. In previous versions of RAMS, users obtained routing data by querying each Route Recorder individually. Users could configure Generic Routing Encapsulation (GRE) tunnels to connect Route Recorders deployed to remote areas of the network. Each Route Recorder provided reports specific to the local areas and protocols where it was listening.

You can still obtain reports from individual Route Recorders with RAMS 5.20. However, using a centralized Modeling Engine to generate reports not only provides a global view of the network, but also reduces the amount of time required to obtain that view. In addition, querying the Modeling Engine reduces the amount of work required by the Route Recorders, whose primary function is to record data.

Choose a data source using the *Data Source Configuration* page. To access this page, locate **Maintenance** on the left navigation bar, then click **Data Source Configuration**. The *Data Source Configuration* page appears as shown in [Figure 31](#).

Data Source Configuration

Data Source Configuration

Get data from Route Recorders

Connect to a Replicating Modeling Engine

Enable Replication

Bring back data for the past

Enable Centralized Report Server

Report Server is required to run Alerts and IGP Reports

Replication Status		
IP Address	Databases	Status

Figure 31 Data Source Configuration Page

The Data Source Configuration page provides the following three options:

- **Get data from Route Recorders** — Use this option to disable replication of data on the Modeling Engine. To obtain network information, you must query individual Route Recorders.
- **Connect to a Replicating Modeling Engine** — Use this option in a deployment where multiple Modeling Engines are co-located. You can choose one unit to act as the centralized Modeling Engine, then point additional Modeling Engines to this central unit to obtain network-wide data. When you select this option, you must provide the IP address of a Modeling Engine who has replication enabled as described in the following paragraph.
- **Enable Replication** — Use this option to begin replicating data from all Route Recorders in the deployment.

When you enable replication, you choose the amount of data to be copied from the network's Route Recorders to the centralized Modeling Engine. For example, choose **One Week** from the drop-down list to obtain data recorded over the last 7 days. You can bring back one week, two weeks, one month, or a maximum of 2 months of data.

While data is being replicated, the Replication Status table lists the IP addresses of each Route Recorder in the deployment, which databases are being copied, and the status of the replication. Only databases that are currently recording are replicated.

Note When a Route Recorder is added to the network, the centralized Modeling Engine automatically begins replicating data from that unit. If you opt to bring back 2 weeks of data, for example, the Modeling Engine copies data recorded up to 14 days before the Route Recorder was added to the configuration.

If replication is enabled on the Modeling Engine, you can select the **Enable Centralized Report Server** check box. This option allows you to produce network-wide reports from the local Modeling Engine.

After a database is replicated on the Modeling Engine, the Route Recorder retains ownership of that data. Any operation - such as delete or rename - that is performed on the Route Recorder will be replicated on the Modeling Engine.

Archiving Data

RAMS allows you to archive data using its automated archival tool. The automatic archival of data differs from creating backup files, which is described in [Backing Up and Restoring Data](#) on page 107. First, you can archive data without stopping the recording process or requiring users to log off of the appliance. Second, RAMS archives segments of data on an incremental basis, rather than storing an entire database at once. Segments of data are created each week, on Sunday, when RAMS divides the database into manageable files, labeled by date and time. These files are archived at regular intervals (every 7 days, on Monday at 0:00). Only unarchived data is stored during this process. Archived data remains accessible for analysis through the History Navigator (see [Chapter 6, “The History Navigator”](#)).

Note that unlike using Backup and Restore, automatic archiving data does not allow up-to-the-minute data storage and retrieval. For example, if data is automatically segmented on Sunday, June 25 at 10:55 PM, those segments are archived Monday, June 26, at 0:00. Any data written to the database beginning on Sunday, June 25 at 10:56 PM will not be archived until the following Monday, July 3 at 0:00. However, RAMS provides an “archive now” feature to capture data between intervals if necessary. This manual option requires you to stop recording on the database before archiving, and is described in [Manually Archiving Data](#) on page 121.

Note The procedures for archiving data on the Modeling Engine are the same for a distributed configuration as they are for stand-alone configurations. The only difference is the archived data for the distributed configuration will be replicated.

Note If more than one recorder is recording and archiving the same network area, or if archiving is configured on both the recorder and Modeling Engine that is replicating the data, and if both are pointing to the same storage unit, then the recorder or Modeling Engine that was invoked first will actually store the data.

Configuring Automatic Archival Settings

Manage the archival tool using the *Archival Configuration and Remote Storage* page. To access this page, locate **System** on the left navigation bar, then click **Archival Configuration**. The *Archival Configuration and Remote Storage* page appears as shown in [Figure 32](#).

Home Administration Recorder Configuration Reports Portal Support Unit: 19

Alerts

- Destinations
- SNMP Test
- IGP
- BGP

Application

- VNC Configuration
- Layout Backgrounds
- Queries

Diagnostics

- System Diagnostics
- View Log
- View Configuration

Maintenance

- Backup and Restore
- Databases
- License
- Software Update
- Restore Archives
- Shutdown

System

- Support Access
- Network and Interface
- Time and Date
- Mail
- Archival Configuration**
- FTP Server
- Users

Archival and Remote Storage Configuration

Configure Archival and SMB share information for creating/restoring archives/backups to and from a remote storage

Enable SMB Client

* Remote username:

Remote password:

Remote server:

Share name:

Enable Archival

Figure 32 Archival and Remote Storage Configuration Page

The RAMS archives data at a frequency of every seven days. Archiving of data occurs per-unit; you must configure archival settings on each unit in a distributed system.

To configure the archiving tool, perform the following steps:

- 1 Enable remote storage as described in [Enabling SMB and Adding a Remote Server](#) on page 113.
- 2 On the *Archival Configuration and Remote Storage* page, select the **Enable Archival** check box.
- 3 Specify how frequently RAMS will archive data. The frequency must be a multiple of 7. For example, type 7 to archive weekly. Type 14 to archive every two weeks.
- 4 Click **Update**.

Any unarchived databases that exist before you enable archival functionality will be archived at the next scheduled archiving interval.

Manually Archiving Data

To archive data between regularly scheduled intervals, which occur weekly on Monday at 0:00 in the time zone set on the unit, you can use the **Archive Selected Databases** button on the *Database Administration* page. Only offline databases can be archived. Therefore, before manually archiving data, you must stop recording on the selected database if necessary. RAMS then archives the most recent unarchived data. When recording is restarted, RAMS creates and begins writing to a fresh segment of the database.

Note Databases created with RAMS 3.7x and earlier cannot be archived.

To manually archive data, perform the following steps:

- 1 Locate **Maintenance** on the left navigation bar, then click **Databases**.

The *Database Administration* page is displayed as shown in [Figure 27](#).

- 2 From the Offline Databases section, select one or more check boxes corresponding to the databases to archive.

If the database to archive is currently online (recording), it appears in green. You must stop recording to the database before you can archive data. See [Chapter 2, “Configuration and Management”](#) for more information about starting and stopping the recorders.

- 3 Click **Archive Selected Databases**.

Segments of the selected databases are archived. Filenames of the archived segments use the following format:

```
DatabaseName_<segment time in epoch>
```

For example:

```
CorpNet_Common_West_bgp_AS65520_1149663600
```

When you restart recording, RAMS creates and begins writing to a fresh segment of the database.

The web-based configuration page includes a **Re-archive** button as well. Using **Re-archive** forces RAMS to archive all available data, whether or not the data was previously archived.

Restoring Archived Data

Retrieve archived data using the *Restore from Archive* page. To access this page, locate **Maintenance** on the left navigation bar, then click **Restore Archives**. The *Restore from Archive* page is displayed.

On the *Restore from Archive* page, the following text boxes appear:

- *Start Time*: The start time of the data to restore, in the following format:

YYYY-MM-DD HH:MM:SS

followed by the time zone (for example, PDT).

- *End Time*: The end time of the data to restore, using the format shown above.
- *Databases to Restore*: Contains the names of all databases that have segments available for restoration. The *Databases to Restore* box lists database names in the following format:

CorpNet.Common/ospf/Backbone

This format corresponds to the literal database name:

corpnet_common_ospf_backbone

Labels are created based on the name of the database. For example, note that if you have renamed a database as described in [Renaming Databases](#) on page 105, the original database name is used in the *Databases to Restore* list box.

For example, three labels are listed in the *Labels to Restore* box. The first label, SegmentA, contains data that began recording on June 4 at 10:31:55 PDT. The third label, SegmentC, contains data that stopped recording on June 24 at 10:31:55 PDT. The data contained in all three listed segments was recorded between the start and end times indicated.

Restoration of data occurs on a per-unit basis; you must access each unit in a multi-appliance deployment individually to enable and configure archival settings.

To restore archived data, perform the following steps:

- 1 On the *Restore from Archive* page, supply a start and end time corresponding to the data to restore.

For example, to restore data that was recorded between Wednesday, June 7, 2006 and Wednesday, June 14, 2006, type the following:

2006-06-07 0:00:00 in the *Start Time* text box

2006-06-14 0:00:00 in the *End Time* text box.

Since each segment of data begins on Sunday and ends on Sunday, the restored archives will include more than the requested Wednesday-to-Wednesday time range. In other words, restored data will include Sunday, June 4 through Sunday, June 18.

- 2 Click to select one or more databases in the *Databases to Restore* list box that corresponds to the database to restore.
- 3 Click **Restore Now**.

When restoring the archived data, RAMS retrieves archived database segments for the time range you specified. RAMS then creates a new database in which to store the retrieved segments. The new database is named *DatabaseName<Time1>to<Time2>*, where *DatabaseName* is the name of the database whose segments have been restored, *<Time1>* is the start time of the first restored segment, and *<Time2>* is the end time of the last restored segment.

Following the example used in Step 1, the restored database name would be:

```
CorpNet_Common_West20060607to20060614_bgp_AS65520  
[or in epoch: 1149663600to1150268400]
```

Updating Software

HP provides software updates for RAMS. If the appliance has a Software Update license key installed, you can download software updates directly from the HP FTP site. Contact HP customer service for a license key if necessary.

To download a new software update, locate **Maintenance** on the left navigation bar and click **Software Update**. The *Software Update* page appears, as shown in [Figure 33](#).



Caution In a distributed configuration, where more than one RAMS unit is installed, you must update the software of the master unit before updating the software on each client unit.

Software Update

Version Information			
	Software Version	OS Version	
Installed Software and OS:	5.0.35-E	2.0.28	Installed: Wed Feb 14 14:45:40 2007
Alternate Software and OS:	5.0.33-E	2.0.26	<input type="button" value="Install Alternate Software and OS"/>

Note: Installing an alternate Software and OS will cause the system to be rebooted.

Operators currently connected: 0

Download Software Update

URL:

Key:

Use Proxy

Host: Port:

Username: Password:

Figure 33 Software Update Page

RAMS has its own operating system software and application software. How you update the RAMS software depends on how it is connected to the Internet:

- The download process is easiest when RAMS is connected to the Internet, either directly or through a proxy server. See [Updating with Internet Access](#) on page 125.
- If the RAMS cannot get access to the Internet, you can download an update, move it to a local FTP server that the appliance can access, then update the appliance from there. See [Updating without Internet Access](#) on page 126.

Updating with Internet Access

If the RAMS appliance can access the Internet directly or through a proxy server, follow the steps in this section. Otherwise, proceed to [Updating without Internet Access](#) on page 126.



Caution Before updating RAMS software, stop recording on Route Recorders and the Flow Collectors. Restart recording after the update is complete. This ensures that traffic databases are renamed correctly.

To download updates when connected to the Internet through a proxy server, first you must set up the proxy configuration. You can then proceed to the second set of steps to download the update.

To set up the proxy configuration, perform the following steps:

- 1 On the *Software Update* page of the unit where you are updating software, select the **Use Proxy** check box.
- 2 Type the host and port details for the proxy server.
- 3 If the proxy server is password protected, type the user name and password.
- 4 Click **Save Proxy Settings** to preserve these settings for future downloads.

In a distributed configuration, where more than one RAMS unit is installed, you can update software on multiple machines at once *after* the software on the master unit has been updated. You must leave each the *Software Update* page open in a browser window until the software download for the machine is complete.

To download updates when connected directly to the Internet, perform the following steps:

- 1 On the *Software Update* page, click **Check for Update**.

An Update Available message tells you if an update is available. If so, the URL of the update appears automatically in the *URL* text box. Otherwise, a message tells you no update is available.

- 2 In the *Key* text box, type the Update key provided by HP customer support.
- 3 Click **Update** to begin download and installation of the update.

This can take some time, depending on your connection speed, since well over 100 MB of data is transferred.

- 4 If the download includes an operating system update, a message appears stating that you must reboot RAMS to complete the download. Click **Reboot Now** and wait for the system to reboot. This should take approximately two or three minutes.

In a multi-unit deployment, when updating software update on a client machine, you will be automatically directed to the master unit's *Home* page 5 seconds after you click **Reboot Now**.

- 5 Log in again.

Note If at any time during the update process a 404 page error appears, click **Back** on the browser and then click **Refresh**.

Updating without Internet Access

If the RAMS appliance cannot access the Internet directly or through a proxy server, you can download an update to a local FTP server accessible by the RAMS appliance, and then install the update from the local FTP server.

Use the full URL to download the update because the file may be hidden in an unreadable directory.

In a distributed configuration, where more than one RAMS unit is installed, you can update software on multiple machines at once *after* the software on the master unit has been updated. You must leave the *Software Update* page for each unit open in a browser window until the software download is complete on that machine.

To download updates when RAMS is behind a firewall, perform the following steps:

- 1 Go to the HP Route Analytics Management System product web site (<http://www.openview.hp.com/products/ovrams/index.html>), and follow the **Software Patches** link. This will take you to the HP software patches web site.
- 2 Use the Browse By Product version list to navigate to patches for the Route Analytics Management System product. Use the information there to determine if an update is available for you.

If an update is available, download it and save it to a convenient location. Make a note of the associated update key, which you will need to complete the update.

Note Instructions that accompany an upgrade may differ in some details from the steps given in this section. If so, use the instructions from the web site, as they are more recent.

- 3 Move the update package to a local server configured for anonymous FTP. The RAMS appliance itself is an acceptable server, providing you set it up as described in [Configuring the FTP Server](#) on page 134.
- 4 In the **URL** field, enter the URL for the local server you are using. For example, `ftp://anonftp.company.com/<dir>/<patch>`

If you use the RAMS appliance as your FTP server, the URL could easily be:
`file:///var/ftp<patch>/`

Note You must input the three slashes (“///”) in the URL.

- 5 On the Software Update page, click **Check for Updates**.

A message tells you if an update is available. If so, the **URL** field is automatically filled in for you.

- 6 Enter the Update key provided by HP customer support.
- 7 Click **Update**.

Downloading begins. If the download includes an operating system update, a message appears stating that you must reboot RAMS to complete the download.

- 8 Click **Reboot Now** and wait for the system to reboot. This should take approximately two or three minutes.
- 9 Click the **Home** link and log in again.

Returning to a Previous Version of RAMS

The previously installed version of the software is saved on the RAMS appliance. If you experience difficulty running a new version of RAMS software, you can return to the previous software version. The Software Update page displays the previously installed version number of RAMS.



Caution Reverting to the previous version of RAMS may require a reset to factory defaults, as described in [Shutting Down the RAMS Appliance](#) on page 144, because updates may not be completely reversible. Resetting to factory defaults will erase all data and configuration settings except the installed license. After reverting to the previous version and resetting to factory defaults, you can restore the data and configuration settings if you have a backup file created with the previous version.

To reinstall a previous version of the operating system, perform the following steps:

- 1 From the RAMS *Home* page, click **Administration**, then click **Software Update** on the left navigation bar.
- 2 On the *Software Update* page, click **Install Alternate Software and OS**.
To install only alternate software, click **Install Alternate Software**.
- 3 Click **Yes** to install the previously installed version of RAMS.
An informational window opens stating that an alternate version is installing and the system is rebooting.
- 4 (Optional) If re-set is required, refer to page [page 145](#) for re-set instructions.

Creating Daily Reports

You can schedule a time for Route Recorder to send a daily report summarizing BGP and IGP activity using the *Mail* page. The report arrives via email. To access the *Mail* page, locate **System** on the left navigation bar, then click **Mail**. The *Mail* page appears as shown in [Figure 34](#).

Note Daily reports are configured only on units that record routing data.

Mail

Mail System Configuration

Mail Server:

Sender:

Recipient(s):

Report Setup

Configure daily generation and emailing of Health and Routing Reports or generate one of these reports now.

Health Report

Summarizes the health of each unit in the network, including the status of the various recording processes and their databases, database replication (if applicable), SQL, and RAID (if applicable). Also displays a hierarchical view of the networks monitored by each unit and the license information present on each unit.

Enable Health Report

Email

Routing Report

Presents a variety of IGP reports (topology counters, flapping links, flapping prefixes, active routers, and withdrawn watch list prefixes) for each actively recording IGP domain and a variety of BGP reports (topology counters, BGP route flaps, prefix redundancy divergence, and AS reachability divergence) for each actively recording BGP topology. Only present on Master and standalone units with Route Analyzer enabled.

Enable Routing Report

Email

Report generation begins at:

Figure 34 Mail Page

Configuring the Mail System

Before you can schedule daily reports on a Route Recorder, you must specify the mail system information used to deliver the reports.

You can also enable health reports and routing reports. Health reports provide status of processes of each machine running on a network (for example, this report will show recording processes and status of databases that are active). Routing reports will show IGP and BGP reports for databases that are recording.

To configure the mail system, perform the following steps:

- 1 On the *Mail* page, type the outbound mail server or relay in the **Mail Server** box using either a DNS name or an IP address enclosed in square brackets (for example, [192.168.0.200]). If you do not specify a mail server, RAMS attempts to send directly to the mail server(s) of the recipient(s).
- 2 In the *Sender* text box, type the full e-mail address to be shown as the sender of mail sent from RAMS. Bounced e-mail messages may be sent to this address, so this address should be valid.
- 3 Type the recipient(s) e-mail address(es) in the **Recipient(s)** box. If you have more than one recipient, separate each recipient address with a comma.
- 4 Click **Update Mail Configuration**.
You can send a test message to the recipient(s) to verify receipt by clicking **Send Test Message**.
- 5 In the *Report Setup* portion of this screen, check the **Enable Health Report** to initiate the health report.
- 6 Click the **Email** box to have the report sent to the recipient.
- 7 Check the **Enable Routing Report** to initiate the routing report.
- 8 Check the **Email** box to have this report sent to the recipient.
- 9 Select the time you want the report to generate in the *Report generation begins at:* drop down menu.
- 10 Click **Update Report Configuration**.

Scheduling Daily Reports

To schedule daily reports, perform the following steps:

- 1 On the *Mail* page, select the **Enable daily reports** check box to send reports to the configured recipient(s) at a specific time each day.

- 2 Select the time to generate and send reports from the **Report generation begins at** drop-down list.
- 3 Click **Update Report Configuration**.

Understanding Daily Report Contents

The daily report contains several sections summarizing network activity, including the following topics:

- **Networks Monitored** — Lists all databases and their current status: online, offline, or offline in the last 24 hours.
- **IGP Summary** — If RAMS records IGP data, this section lists the following results in all monitored databases:
 - Counts of the number of routers, adjacencies and prefixes
 - Top 5 flapping links
 - Top 5 flapping prefixes
 - Top 5 active routers
 - Prefixes on the Prefix Flap watch list that are withdrawn at the time the report is generated (for example, at midnight). Note that this is not a summary of all prefixes withdrawn during the last 24 hours.
- **BGP Summary** — If RAMS records BGP data, this section lists the following results in all monitored databases:
 - Top 5 BGP route flaps
 - Top 5 prefix redundancy divergence
 - Top 5 autonomous system (AS) reachability divergence
- **Information Summary** — Lists RAMS system details, including the license status.

Viewing Saved Daily Reports

RAMS stores the last 30 days' reports on each unit that records routing data, so you can compare changes to an earlier report.

To view saved daily reports, perform the following steps:

- 1 On the RAMS *Home* page, click the **Reports Portal** button at the top of the *Home Page*.
- 2 Click **Daily Reports** on the navigation bar.
- 3 Click on a report filename to download or view that report.

Note If you click **Daily Reports** on a unit that does not have the Route Analyzer Reports license key installed, RAMS displays a message advising you to access a unit that records routing data.

Configuring the FTP Server

A portion of the hard disk on the RAMS is available for the storage of users' files in the FTP or SFTP server directory. Upload time series files and MRTG files onto the RAMS using FTP for correlation with routing events (see [Correlating Time Series Data](#) on page 278). Backed-up database files are also stored on the appliance.

To access the *FTP Server Configuration* page, locate **System** on the left navigation bar, then click **FTP Server**. The *FTP Server Configuration* page, shown in [Figure 35](#), lets you set the password for FTP file operations.

FTP Server Configuration



Enable FTP Server
 Enable SFTP Server
Password:
Confirm Password:

Figure 35 FTP Server Configuration Page

To enable FTP file uploads, perform the following steps:

- 1 On the *FTP Server Configuration* page, check either the **Enable FTP Server** or **Enable SFTP Server** check box.
- 2 Type a password in the *Password* text box.
- 3 Type the password again in the *Confirm Password* text box.
- 4 Click **Update** to complete configuration.

After the FTP or SFTP server is enabled, you can log in for FTP file transfer.

To log into the FTP or SFTP server, perform the following steps:

- 1 Start the FTP or SFTP client software.
- 2 Type the following values in the appropriate text boxes:
 - The IP address of the RAMS.
 - The login name: `rexftp` for FTP or `sftp` for SFTP.

- The password set on the *FTP Server Configuration* page.
- 3 Change the FTP directory to *pub* when uploading time series and MRTG files.

Managing Users

Before users can access a RAMS unit, the administrator must add each user name and password to the database and specify the type of access privilege, to apply to each user.

Every RAMS user is associated with one of the following privileges:

- Administrators, who have full access privileges, such as configuring other users and setting alerts. Administrators can load, modify, or delete any layout, as described in [Chapter 5, “The Routing Topology Map.”](#)
- Operators, who can access the RAMS application and the report pages, but do not have any administration privileges. Operators can load any layout, but if they save a layout that they do not own, a new copy of the layout is created with the same name, but with the new owner. Operators can also save a layout under a new layout name.
- Guests, who are view-only users, have the same access as operators, but cannot save or delete routing topology map layouts.

When administrators, operators, and guests access a RAMS unit using SSH as described in [Chapter 4, “RAMS Viewers”](#) for the X Window System or VNC client displays the RAMS graphical interface.

- CLI Access users, who use SSH to connect to RAMS’s TTY user interface rather than RAMS’s graphical user interface. Once connected, the CLI access user can run diagnostic commands, reboot the unit, or shut it down. For more information about using diagnostic and other command-line functions, see the *RAMS Appliance Setup Guide*.

You can add, remove, and edit users on the *User Administration* page. To access the *User Administration* page, click **Users** on the left navigation bar. The *User Administration* page appears as shown in [Figure 36](#).

User Administration

Create Users

New User:

Privilege:

Password:

Confirm Password:

Update Accounts

Current Users:

Privilege:

Password:

Confirm Password:

Update Login Attributes

Login Expire Time (in seconds): (May be 'Never'. Minimum value 60.)

Login Idle-Timeout (in seconds): (May be 'Forever'. Minimum value 60.)

Figure 36 User Administration Page

Adding, Deleting, and Editing Users

The *User Administration* page displays all users allowed access to a particular RAMS unit. Each RAMS unit has its own user database. In other words, in a distributed configuration, you must set access privileges for each unit in the configuration individually. When you allow administrative privileges to a user on the master unit, this does not allow that user the same privileges on the client units and vice versa.

To add users to the Current Users list, perform the following steps:

- 1 In the *New User* text box, type the name of the user (for example, *user1*).
- 2 From the **Privilege** drop-down list, choose a privilege for the user (for example, **Guest**, which would allow *user1* view-only access to routing topology map layouts).
- 3 In the *Password* text box, type the password the user will use to log in.
- 4 Type the password again to confirm.
- 5 Click **New** to save the new user. In the example, *user1* appears in the *Current Users* text box.

The *Current Users* box on the *User Administration* page displays any users who are already in the database. You can edit the privilege and password for these users, or remove them from the database.

To change the privilege or password of an existing user, perform the following steps:

- 1 In the *Current Users* text box, click the name of the user to edit.
- 2 In the *Password* text box, type the new password and/or from the **Class** drop-down list, choose a class for the user.
- 3 If necessary, confirm the password.
- 4 Click **Update**.

To remove users from the list, perform the following steps:

- 1 In the *Current Users* text box, click the name of the user to delete.
- 2 Click **Delete**.

The user is removed from the list and can no longer access this RAMS unit.

Note If you change the class for the administrator user name you entered to log in, or if you delete that user, you will be automatically logged out.

Users can change their passwords by using the **Change Password** link in the row at the top of the web page, after they are logged in.

Setting Session Login Timers

The administration interface allows two session login timers: a login expiration timer and a login linger timer. On the *User Administration* page, each can be enabled as a security precaution.

The login expiration timer restricts the time between the start of login activity and access to the user interface. After the expiration time is set, a user must re-login within the specified time period. The login expiration time applies to all users on a particular RAMS unit.

To set the login expiration time, perform the following steps:

- 1 In the *Login expire time (in seconds)* text box, type the number of seconds that may elapse between logins. The number of seconds must be 60 or more.
- 2 Click **Update Login Timeouts**.

The login linger time setting is the time period between keystrokes after a user logs in. If the specified time elapses without a keystroke from the user, the user is automatically logged out.

To set the login linger time, perform the following steps:

- 1 In the *Login linger time (in seconds)* text box, type the number of seconds that may elapse between keystrokes. The number of seconds must be 60 or more.
- 2 Click **Update Login Timeouts**.

Viewing and Exporting RAMS Log Pages

You can view and export log files to help diagnose problems on a particular appliance. To access the *View Log* page, locate **Diagnostics** on the left navigation bar, then click **View Log**.

To view a log page, perform the following steps:

- 1 On the *View Log* page, specify which component logs to view or choose **All** from the **Component** drop-down list.

The number of pages in the log displays automatically.

- 2 From the **Lines** drop-down list, select the number of lines to display per page.
- 3 (Optional) Type the page number to view in the *Page* text box. If you leave the text box empty, page 1 displays by default.

Time	Component	Message
Apr 5 11:04:55	Route Recorder	Starting rex: succeeded
Apr 5 11:04:35	Route Recorder	Startup REX 4.0.78-E OS 1.0.1 (REX1200)
Apr 5 10:55:58	Route Analyzer	* Igp-loader.cc:948 Trying to create a self-looping link (12,12) (24.0.0.5 24.0.0.5)

Figure 37 View Log Page

- 4 Select the **Show Most Recent First** check box to see the last recorded log entries.
- 5 Click **Apply Filter**.

You can print a copy of this page using the print command on the web browser.

To export the log as plain text, perform the following steps:

- 1 Choose which section of the log to export by following the previous set of steps in this section.

- 2 Click **Export Log as Plain Text**.

The log page is redisplayed in plain text form in the browser window.

- 3 Use the *File* menu on your browser or right-click in the window to open the pop-up menu.

- 4 Choose **Save As**.

- 5 Type the directory where the log file will be stored, and then click **Save**.

- 6 Click the **Back** button on your browser to return to the *View Log* page.

Uploading Layout Backgrounds

Layout backgrounds are images you can apply to the routing topology map to provide additional visual cues to the layout. For example, you can upload a map depicting the geographic location of network routers, which would enable you to arrange nodes based on physical or logical groupings, such as per building or per lab.

Create or convert a desired background image using JPEG, PNG, BMP, the SVG format (Scalable Vector Graphic), or XPM (X PixMap, an ASCII image format used by the X Window System). Adobe Illustrator, Corel Draw, OpenOffice Draw, and a number of other graphics tools support SVG images. Import the image to RAMS using the *Layout Backgrounds* page. The image files are stored in a database and are accessible from any RAMS appliance on the network.

To import an image into RAMS, perform the following steps:

- 1 On the left navigation bar, locate **Application**, then click **Layout Backgrounds** to access the *Layout Backgrounds* page.
- 2 Click **Browse** to locate the image file to upload. Be sure the appropriate file extension is included in the file name.
- 3 Click **Upload**.

Note Note that binary images should not exceed 12 MB in size, and all other images should not exceed 16 MB.

The uploaded file, as well as any other layout background files that have been uploaded, appears in the **Image Name** column, along with the type of image file. You can preview the image by clicking **View**.

- 4 To delete any of the uploaded background image files, check the corresponding check box, and then click **Delete**.

If a background image is in use, a green star appears in the **Delete** column; the image cannot be deleted until it is removed from all routing topology map layouts.

- 5 To apply or remove a background image to or from the routing topology map layout, see [Finding a Router](#) on page 189.

Using Diagnostic Functions

RAMS supports the diagnostic functions *ping* and *traceroute*. Use these functions to investigate network failures or outages. To access the *System Diagnostics* page, locate Diagnostics on the left navigation bar, then click **System Diagnostics**.

Pinging a Network Device

Use *ping* to determine if a destination host is reachable on the network.

To ping another network device, perform the following steps:

- 1 On the *System Diagnostics* page, type the IP address or DNS name of the destination device.
- 2 Click **Ping**.

The *System Diagnostics* page displays the results of the ping function.

Running a Traceroute

Use *traceroute* to trace the path a packet takes through the network from the RAMS to the destination you specify.

To run the traceroute function, perform the following steps:

- 1 On the *System Diagnostics* page, type the IP address or DNS name of the destination device.
- 2 Click **Traceroute**.

The *System Diagnostics* page displays the results of the traceroute function.

Shutting Down the RAMS Appliance

The RAMS appliance can be shut down at any time if needed. The system shutdown options are displayed on the *Shutdown* page, shown in [Figure 38](#).

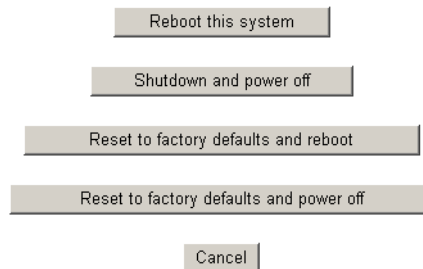
To access the *Shutdown* page, locate **Maintenance** on the left navigation bar, then click **Shutdown**.

The following options are available on the *Shutdown* page:

- **Reboot this system** – Click this button to reboot the system. A confirmation page appears. Click **Yes** to reboot the system. RAMS stops the VNC server and stops recording data, and reloads the operating system and recording software from the disk. Then, using the previous system settings, it automatically restarts the VNC server and the recorders. The message “Please wait for the system to reboot then click on Home” appears. If this message does not appear, wait three minutes, and then click **Home**. Log in to the *Administration* pages and verify that the VNC server and recorders are operating correctly.

Shutdown

Shutdown options



System uptime: 16 hours, 25 minutes

Figure 38 Shutdown Page

- **Shutdown and power off** – Click this button to shutdown the system and power off. A confirmation page appears. Click **Yes** to shutdown the system. RAMS stops the VNC server and stops recording data. To restart, press the power switch on the RAMS appliance.
- **Reset to factory defaults and reboot** – Click this button to restore the factory default settings and reboot the appliance. A confirmation page appears. Click **Yes** to reset the factory default settings and reboot the system. When the system reboots, use the serial console interface to reconfigure the network address and then connect to the *Home* page, and log-in as administrator. Restore the system configuration from a back-up file, or re-configure manually following the sequence of steps in [Applying License Keys](#) on page 31. If recording is enabled, verify on the *Recorder Configuration* page that Hellos and Events are being received from the areas or levels monitored by the RAMS.

Note If the factory settings are restored, the following information is lost:

- All configuration information, including Network, Route Recorder, user names and passwords
- Data files, including databases, user time-series data files, and log files

The current and alternate versions of the software and the installed licenses remain on the appliance.

- **Reset to factory defaults and power off** – Click this button to restore the factory default settings and power off the appliance. A confirmation page appears. Click **Yes** to reset the factory default settings and power off the system.

4 RAMS Viewers

The X Window System or AT&T Lab's Virtual Network Computing (VNC) Viewer is required to run the RAMS client application.

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

This chapter explains how to download and configure the X Window System and VNC on Microsoft Windows, UNIX, and Linux operating systems. The downloads are available from the web server on the RAMS appliance.

The X Window Server

To use the X Window System, the user's computer must run an X server to receive and display the output of the remote RAMS application. In addition, RAMS requires that the X session be connected using the SSH (secure shell) protocol for privacy and security.

Using X Window System Software for MS Windows™

Several X Window System products incorporating SSH are available for Microsoft Windows. The third-party X software included with RAMS 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.

Note The X-Win 32 evaluation from StarNet supports anti-ill-used fonts, which allows the BGP Root Cause Analysis and RIB visualizations to display correctly.

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

- 1 Using a web browser, connect to the RAMS Home page.
- 2 Click **Support** on the top navigation bar

The *Support* page appears as shown in [Figure 39](#).

Support

Download Documentation (PDF)

[Traffic User's Guide](#)

[Developer's Guide](#)

Download X Window System Software

The following product is recommended for its support of anti-aliased fonts:

[Link to StarNet Communications Corp. for X-Win32 Evaluation](#)

The following alternative product is available for direct download:

[Xmanager 30-Day Evaluation \(Windows 9x, NT, 2000, XP\)](#)

[Xmanager Setup \(ShockWave Demo\)](#)

Download VNC Viewer

[Windows 9x, NT, 2000, XP](#)

[Linux \(x86\)](#)

[Macintosh \(OS9\)](#)

[Macintosh \(OSX\)](#)

[Solaris \(sparc\)](#)

Install SVG Plug-in

<http://www.adobe.com/svg/viewer/install/main.html>

[View System Information](#)

For other technical support, visit:

<http://www.packetdesign.com/support.html>

Figure 39 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, you will receive in return a 30-day license key and download instructions for X-Win32.

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

- 1 Using a web browser, connect to the RAMS *Home* page.
- 2 Click **Support** on the top navigation bar

The *Support* page appears.

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

The *File download* window opens.

- 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 the desktop.
- 2 Open the window from the X-Win32 folder.
- 3 Type the connection details in the window that appears:
 - Name: Type a name for the session.
 - Hosts: Type either the hostname or IP address of the RAMS appliance.
 - Protocol: Select SSH; to ensure privacy and security, RAMS accepts only SSH connections.
 - User name and Password: Type your RAMS user ID and password.
- 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.

Note SSH is required to run RAMS through the X Window System.

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

- 1 Ensure a graphical user interface such as XDE or Gnome is running on the desktop.
- 2 From the shell in a terminal window, open an SSH connection to RAMS. Type the following command:

```
ash X userid@rex
```

For example:

```
ssh X op@10.0.0.24
```

- 3 Type your RAMSuser password when prompted.
The RAMS application opens on the desktop.

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

- 1 Ensure that a graphical user interface, such as OpenWindows or CDE, is running on the desktop.
- 2 Open an SSH connection to RAMS in a terminal window (CDE) or shelltool (OpenWindows). Type the following command at the shell prompt:

```
ssh X userid@rex
```

For example:

```
ssh X op@10.0.0.24
```

- 3 Type your RAMS user password when prompted.
The RAMS application opens on the desktop.

Note RAMS is expected to work with the X Window System on other platforms, but they may not be tested or fully supported. For more information on these platforms, go to <http://www.packetdesign.com>.

VNC Viewer

To use VNC, a VNC viewer (client) must be installed on the user's system. The VNC viewer then connects to the NVC server running on the RAMS appliance. Before starting the VNC viewer on the desktop, configure and start the VNC server.

Downloading VNC

The RAMS Support page offers five versions of the VNC viewer:

- 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 On the RAMS *Home* page, click **Support** on the top navigation bar.
- 2 On the *Support* page, click the link for the appropriate version of the VNC viewer.

The *File download* window opens.

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

This saves the VNC viewer to a local directory.

- 4 The downloaded VNC file is compressed. Before installing it, decompress it with an application such as WinZip.
- 5 Run the VNC viewer.exe file to install VNC.

Start VNC after installing it.

To start VNC, perform the following steps:

- 1 Double-click the VNC icon.

The *Connection Details* dialog box appears.

- 2 If this is a first-time installation, adjust the optional settings as desired. Click **Options** to display the *Connection Options* dialog box.
- 3 Check **Tight Encoding** to improve performance.
- 4 Check **Full-screen mode** to eliminate scroll bars on the VNC viewer and window frame. This prevents the taskbar and minimized icons on the RAMS desktop from being scrolled off-screen.

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

- 5 Click **OK** to close the *Options* dialog box.
- 6 Type the RAMS appliance IP Address or hostname in the VNC Server text box followed by “:1”.
- 7 Click **OK** to start the VNC viewer.

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

- 8 In the *VNC Password* dialog box, type the VNC authentication password.
- 9 Click **OK**.

This starts the VNC viewer.

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

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**.
The *Save As* dialog box opens.
- 3 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 was saved.
- 3 Type 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, type the following command:

```
rpm -qa | grep vnc
```

To start VNC, perform the following steps:

- 1 Type the following command at the command line, where a.b.c.d is the RAMS IP address or hostname:

```
vncviewer a.b.c.d:1
```

or

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

Note The warning “vncviewer: ConnectToTcpAddr: connect: Connection refused: will appear if you omit the “:1” at the end of the IP address or if the VNC server is not running. In the latter case, contact the RAMS administrator and request that he or she start the VNC Server from the *Administration* page.

- 2 At the password prompt, type the VNC authentication password.

This starts the VNC viewer.

Note If RAMS 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 RAMS administrator to resolve the problem. If shared access to the VNC desktop is appropriate, as the RAMS administrator to change the setting to Multiple Operators and restart the VNC server.

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

Opening the RAMS Application in VNC.

When the VNC server is initially started, the RAMS application is automatically started on the VNC desktop so it appears as soon as the VNC viewer connects to the server.

If the RAMS application is subsequently closed using the **Quit** menu command or by closing the main window, then the next time VNC is opened the desktop will appear without the RAMS main window. In that 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 RAMS.

The RAMS main window opens.

Installing SVG Plug-In

Adobe offers a free SVG plug-in which can be downloaded from the following url: <http://www.adobe.com/svg/viewer/install/main.html>).

HP has used the Adobe plug-in with a variety of 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 will display the currently configured settings for your appliance.

5 The Routing Topology Map

The routing topology map displays the status of the routers and links in the entire network at a glance, which lets you quickly spot outages. This chapter describes in detail the routing topology map and the various tools used to monitor the network.

The *Routing Topology Map* window provides a view of the network as it is currently running, including any tactical changes made during outage repairs that might not be reflected in the network design documents.

You can use the routing topology map to anticipate and head off user complaints by identifying any routing failures in the network. You can also identify potential configuration errors that might result in service outages following maintenance activities.

The routing topology map lets you view the routing events that led to failure, to aid forensic analysis. It also provides an accurate, vendor-independent view of the routing network that often points out potential implementation or interoperability issues that are not easily isolated using other tools.

Opening a Routing Topology

The routing topology is the overall view of the routing activity on the network that RAMS is monitoring. Each stored database has a corresponding routing topology. The routing topologies are displayed as a map in the main window of the RAMS application.

[Chapter 4, “RAMS Viewers”](#) explains how to start the RAMS application in the X Window System or in a VNC viewer. The next step is to select and open a routing topology.

To open a routing topology, perform the following steps:

- 1 Select **Open Topology** from the *Topology* menu on the main window.

The *Open Topology* dialog box appears. The names of databases that are configured for recording data appear in green letters, while the names of inactive databases appear in black letters.

- 2 Select the desired databases from the list provided. Selected items are highlighted, as shown by *CorpNet* in the figure. Any combination of databases can be selected 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.

Note If you plan to perform BGP-specific analysis operations, such as the Root Cause Analysis described in [Chapter 6, “The History Navigator,”](#) it is recommended that you deselect any traffic databases from the tree in the *Open Topology* dialog box. Traffic data is not relevant for BGP-specific analysis, and loading traffic information from the database may slow the analysis process.

- 3 Click **Open**.

After the database loads, the topology map appears in the main RAMS window. The larger the database, the more time is required to load the topology.

Managing Topology Map Layouts

When a topology is loaded initially, RAMS uses a randomizing process to place the nodes on the routing topology map. The placement is not geographical. After you save the layout, the database will load more quickly when it is reopened.

Creating a Layout

To create a custom layout, you can move nodes by dragging them with the mouse to a preferred location on the map.

Applying Layout Backgrounds

Layout backgrounds are images you can apply to the topology map to provide additional visual cues to the layout. For example, you can upload a map depicting the geographic location of network routers, which would enable you to arrange nodes based on physical or logical groupings, such as per building or per lab.

Create or convert a desired background image using the SVG format (Scalable Vector Graphic, a W3C standard for producing high-quality graphics), JPG, PNG, BMP, or XPM (X PixMap, an ASCII image format used by the X Window System). Import the image to RAMS using the *Layout Backgrounds* page as described in [Managing Databases](#) on page 102. The image files are stored in a database and are accessible from any RAMS appliance on the network.

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

- 1 Open RAMS and select a topology as described in [Opening a Routing Topology](#) on page 158.
- 2 From the *Topology* menu, choose **Select Layout Background**.

The *Select Layout Background* dialog box opens as shown in [Figure 40](#). Imported image files, including any images you imported in [Managing Databases](#) on page 102, appear in the dialog box.



Figure 40 Select Layout Background Window

- 3 Highlight the filename of the image to apply to the topology and click **OK**.
The image appears as a background for the topology in RAMS.
- 4 To remove the background image, repeat Step 2, then choose **<blank>** from the *Select Layout Background* dialog box.

When you save the layout, the background image is saved with it and will load each time you open the topology unless you change or remove the image as described in Step 4.

Saving Topology Map Layouts

Multiple different layouts can be named and saved, and each user can set a default named layout (see [Miscellaneous Options](#) on page 214). The default layout is then loaded whenever that user opens a topology in RAMS.

When you resize or hide nodes on a layout, those changes are preserved when the layout is saved. You can unhide nodes on the layout as described in [Trimming Leaves](#) on page 187.

Saved layouts are bound to your user name, and are view-only for other users. When another user loads a layout you have saved, RAMS creates a copy of the layout only if the other user makes changes to the original layout and then saves the revised layout.

Understanding Symbols and Colors on the Map

The routing topology map consists of a variety of symbols, representing nodes, interconnected by lines that represent links. The symbols have different shapes depending upon the nodes' functions:

- Square nodes are routers internal to a protocol domain (autonomous system or AS). Alternatively, these nodes represent routers on the border between domains.
- Hexagons denote Area Border Routers (ABR) located on the border of one or more OSPF areas and connecting those areas to the backbone OSPF area. In an EIGRP network, hexagons denote routers that are redistributing routes between EIGRP autonomous systems.
- Open circles indicate the pseudonodes that represent multiaccess networks (LANs) in OSPF and IS-IS.

Note IS-IS treats Non-broadcast Multiaccess (NBMA) networks, such as ATM, as a collection of point-to-point connections and will not show any pseudonodes on the topology map. However, OSPF treats NBMA networks as multiaccess, and will show pseudonodes between two OSPF routers connected to the NBMA network.

- Filled circles represent BGP peers.

The nodes and links in each routing area of the network are drawn in a different color. Nodes that are down in any area are colored red. You can change or turn off the coloring of elements on the map using the *Color Modes* submenu of the *View* menu. When coloring is turned off, all nodes are black and all links are gray.

Understanding Links and Peerings on the Map

Links on the topology map are divided into two halves representing the two directions of communication between a pair of nodes. The half adjacent to a node represents the direction outbound from that node. Each half can be separately up or down. Links that are down are bright red instead of the area color or gray. If a link represents multiple physical links between two nodes, and if only some but not all of the physical links are down, then the link will be dark red instead of bright red.

RAMS represents IGP links and BGP peerings similarly on the topology map, with a colored line connecting nodes. However, distinct differences exist between these two types of connections. Beyond the physical connection between a BGP router and a peer such as RAMS, peerings between BGP protocol routers and their clients are inferred, and may not follow physical paths. RAMS discovers and monitors these connections indirectly by receiving routes from BGP protocol routers whose next hop attribute contains the address of another BGP router.

Using the Toolbar

The RAMS window has a toolbar along the left side with buttons that control the map display or provide shortcuts to menu options. You can move the toolbar to the right side or the top or bottom of the window by dragging the dimpled strip at the top of the toolbar.

The following buttons appear on the toolbar:



Show Topology Hierarchy inserts a pane on the left side of the map window to display a hierarchical view (a tree structure) of the set of routing databases displayed in the map. See [Understanding the Topology Hierarchy](#) on page 175 for a description of the available operations.



Save Layout saves the current map layout (disabled if no changes). If this is the first time you have saved the layout, a dialog box allows you to name the layout and specify it as your default layout.



Zoom In and **Zoom Out** control the zoom level of the map display.



Reset View to 1:1 restores the zoom level to the default.



Node Size Up and **Node Size Down** increase or decrease the size of nodes and their accompanying text. Node size is saved when you save a layout.



Relayout generates a new randomized layout of the nodes.



History Navigator opens the *History Navigator* window for the current topology.



VPN Navigator opens the *VPN Navigator* window for the current topology. This icon is only present if the RAMS unit has a license for the VPN protocol module.



List Routers generates a list of all routers in the current topology map.



List Links generates a table that lists all links in the current topology map.



List Prefixes generates a list of all prefixes in the current topology map.



RIB Browser opens the *RIB Browser* window.



Find Router opens the *Find Router* dialog box.



List/Find Path opens the *List / Find Paths* dialog box.



Legend opens the *Legend* panel, which describes each symbol on the topology map.

Using the Status Bar

At the bottom of the *Routing Topology Map* window is a status bar. The icon on the left end of the status bar indicates the mode:



Online mode – A network icon indicates that the topology is currently being recorded and updates to the routing database are shown on the topology map as they occur. Note that traffic data is delayed by 30 minutes.



History mode – A graph icon indicates that only previously recorded information in the routing database is shown on the topology map. Note that traffic data is delayed by 30 minutes.



Design mode – A network icon indicates that planning features are enabled for the topology map.

You can switch between modes by clicking the mode icon.

A colored dot to the right of the mode icon indicates Route Recorder status with different colors as listed in [Table 4](#). If you click on the status LED icon, additional information on the recorder status and the status of RAMS peers is displayed.

Table 4 Status Indicator States

Color	Description
Green	Indicates Route Recorder is running and recording to the database, and adjacencies between Route Recorder and peer routers in all areas are up.
Blue	This color has two meanings: 1. Indicates data is being recorded, but EIGRP topology exploration is in progress so changes in the topology will not be shown until completion. 2. For all protocols, this indicates that the time on the recorder is out of sync with the time on the GUI.
Yellow	Indicates data is being recorded, but adjacencies to peer routers in some areas are down.
Red	Indicates no data is being recorded, either because Route Recorder is not running or all adjacencies with peer routers are down.

Gray	Indicates a historical database to which no new data is being recorded.
Purple	Indicates replication-related error exists.

Next to the status indicator is the date and time of the network state currently being shown on the routing topology map. Status messages appear to the right of the date and time when applicable.

Using the Menu Bar

The menu bar at the top of the RAMS window contains the following drop-down menus:

- *Topology* menu
- *Edit* menu
- *View* menu
- *Tools* menu
- *Traffic* menu
- *Help* menu

Use these menus to open routing topology maps and monitor and analyze routing data. The remainder of this section describes the options on these menus in more detail.

Topology Menu

The following options are available from the *Topology* menu:

- **Open Topology** – Opens a routing topology database, displaying the topology map in the RAMS window.
- **Close Topology** – Closes a routing topology database so that another can be opened.
- **History Mode** – Changes the RAMS window from Online or Design mode to History mode.
- **Online Mode** – Changes the RAMS window from History or Design mode to Online mode.
- **Design Mode** – Changes the RAMS window from Online or History mode to Design mode.
- **Go to Time** – In History mode, allows you to “rewind” and view the topology map as it appeared at a particular time in history.
- **Reload Layout** – Restores the node positions according to the previously loaded layout (disabled if no layout is loaded or no nodes have been moved).

- **Save Layout** – Saves the current node positions into the loaded layout (disabled if no changes). Same as **Save Layout As** if no layout is loaded.
- **Save Layout As** – Opens the *Save Layout* dialog box to allow saving the current node positions as a new layout with a new name. You can also set the new layout as the default.
- **Load Layout** – Loads a saved layout to reposition the nodes. Multiple layouts can be saved under different names for different purposes.
- **Delete Layout** – Deletes a named layout.
- **Relayout** – Creates a new randomized placement of the nodes on the topology map, which can be helpful to find a preferred orientation that can then be named and saved.
- **Resolve DNS** – Resolves all of the node addresses on the topology map into DNS names. In a large network, this can take some time.
- **Quit** – Closes the RAMS client (the RAMS application itself continues to run).

Edit Menu

The *Edit* menu is enabled in Design mode. See [Chapter 7, “Network Planning”](#) for a complete description of the *Edit* menu and Design mode.


View Menu

The following options are available on the *View* menu:

- **Show Topology Hierarchy** – Inserts a pane on the left side of the map window to display a hierarchical view (a tree structure) of the set of routing databases displayed in the map (see [Understanding the Topology Hierarchy](#) on page 175). Toggles to **Hide Topology Hierarchy** when the pane is shown
- **Show Network Summary** – Provides up-to-date counts of various network elements. Toggles to **Hide Network Summary**. See [Using the Network Summary Option](#) on page 197 for more information about this feature.
- **Hide Toolbar** – Turns off the toolbar on the left edge of the main RAMS window, and toggles to **Show Toolbar**.

- **Color Modes** – Selecting **Color by Area** colors each area of the topology map with a distinct color. Selecting **Color by Bitrate** colors each link according to the speed of traffic flowing across the link. Selecting **Color by Utilization** colors each link according to the amount of traffic flowing across the link. Selecting **Uncolor** removes colors from the topology map, so that failed links and nodes can be seen more easily. See [Colors Option Panel](#) on page 214 for more information.

Note If the color modes mentions color by bitrate and color by utilization, these two selections will be disabled in online mode.

- **Node Label Modes** – Opens the *Node Label Modes* submenu. See [Node Labels Option Panel](#) on page 213 for more information about these options.
- **Trim Nodes** – Opens the *Trim Nodes* dialog box where you can specify that nodes matching a pattern be hidden from view on the routing topology map (see [Trimming the Displayed Nodes](#) on page 186).
- **Unhide All Nodes** – Displays any hidden (trimmed) nodes. You can also click the **Unhide All Nodes** icon  in the lower right corner of the *Routing Topology Map* window (see [Figure 47](#)).
- **Trim Leaf Nodes** – Removes nodes that are on the edges of the map with a single link to simplify the map display (see [Trimming Leaves](#) on page 187).
- **Hide Failed Nodes** – Removes failed nodes from the display.
- **Zoom In** – Enlarges the center portion of the topology map to fill the window, reducing the amount of the map that is visible.
- **Zoom Out** – After **Zoom In** has been done, **Zoom Out** shrinks the routing topology map to increase the amount of the map that fits within the window.
- **Node Size Up** – Increases the size of node symbols and their accompanying text on the routing topology map. This function does not affect zoom level. Node size is saved when you save a layout.
- **Node Size Down** – Decreases the size of node symbols and their accompanying text on the routing topology map. This function does not affect zoom level.
- **Node Size Reset** – Resets the size of nodes and their accompanying text to the default size. This function does not affect zoom level.
- **Reset View to 1:1** – Resets the *Routing Topology Map* window to the original zoom level.

Tools Menu

Use the *Tools* menu to view and analyze the routing layout. The following options are included on the *Tools* menu:

- **History Navigator** – Opens the *History Navigator* window to allow “navigating in time” to analyze past routing data (see [Chapter 6, “The History Navigator”](#)).
- **RIB Browser** – Opens the *RIB Browser* window (see [RIB Browser](#) on page 247).
- **Path Reports** – Opens the *Path Reports* window (see [Chapter 11, “Path Reports”](#)).
- **Online Events** – Opens the All Events window displaying a detailed list of routing events (see [Understanding the Events List](#) on page 261).
- **VPN Reports** – Opens the *VPN Reports* window to allow configuring and monitoring a VPN (see [Chapter 11, “Path Reports”](#)). This option is only present if the RAMS unit has a license for the VPN protocol module.
- **Correlate Time Series** – Opens the *Time Series File Selection* dialog box to display a graph of external time series data in correlation with the routing history (see [Correlating Time Series Data](#) on page 278).
- **List Routers** – Opens the list of all routers on the routing topology map (see [Router List](#) on page 191).
- **List Links** – Opens the list of all links on the routing topology map (see [Links List](#) on page 192).
- **List Prefixes** – Opens the list of all prefixes announced by routers on the routing topology map (see [Prefix List](#) on page 193).
- **List OSI Prefixes** – Opens the list of all Prefix Neighbors and End System (ES) Neighbors announced by routers on the routing topology map (See [OSI Prefixes List](#) on page 195).

Note If OSI IS-IS is not detected by the appliance, the List OSI Prefixes option will not display.
- **List VPN Prefixes** – Opens the list of all prefixes that are part of a VPN (see [Chapter 10, “VPN Configuration and Reports”](#)). This option is only present if the RAMS unit has a license for the VPN protocol module.

- **Find Router** – Opens the *Find Node* dialog box that enables you to search for routers by IP address or router name (described in [Router List](#) on page 191).
- **List/Find Paths** – Opens the *List/Find Paths* dialog box that lets you highlight a path between a router and an Internet prefix or domain name (see [Finding a Route By Prefix](#) on page 209).
- **Highlight by Exit Router** – Finds the set of exit routers toward a specified prefix, NSAP Address, IP address or domain name, then color-codes 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. For more details, see [Viewing the Exit Routers from the Network for any Internet Destination](#) on page 196.

Note To highlight exit routers, the displayed network must include a route to the desired destination.
- **Router Name Repository** – Allows you to customize router names and to map this name to the routers' IP address. See [Assigning Router Names using Router Name Repository](#) on page 198.
- **Autonomous System (AS) Name Repository**—Allows you to customize the AS names as it appears in the appliance. See [Assigning AS Names using AS Name Repository](#) on page 202.
- **Assign BGP Autonomous Systems to Routers** – Allows you to assign routers manually to a BGP AS, primarily for a BGP confederation. See [Verify and Manually Assign BGP AS Assignments to Routers](#) on page 205.
- **Topology Diagnostics** – Presents a submenu to select display of several diagnostic tables (EIGRP topologies only). See [Topology Diagnostics Submenu](#) on page 172.
- **Options** – Opens the *RAMS Configuration Options* dialog box. Use this dialog box to configure the settings for the *Online Update Monitor* and *History Navigator* windows and set other preferences (see [Using the Configuration Options Dialog Box](#) on page 212).

Topology Diagnostics Submenu

The *Topology Diagnostics* submenu is present only for EIGRP topologies. It provides several options to display any problems found in the network configuration or in the modeling of the topology. These options are described in detail in [Diagnosing EIGRP Topology Errors](#) on page 219, and include the following:

- **List Topology Errors** – Opens a list of messages describing anomalies that were discovered during exploration of the EIGRP topology. Clicking on an entry in the list highlights the affected routers and links.
- **List Inaccessible Routers** – Opens a list of routers that were not accessible through Telnet during exploration of the EIGRP topology, including a reason such as authentication failure. Clicking on an entry in the list highlights the last accessible router on the path toward the inaccessible router.

Note Inaccessible routers are not shown on the topology map. They can cause incorrect routes to be displayed and can reduce the ability of RAMS to track changes in the network topology.

- **List Mismatched Distances** – Opens a list of prefixes for which the distance to the prefix reported by a router that peers with RAMS does not match the distance that RAMS calculates across its model of the topology. This list also shows when the periodic topology explorations start and end, along with summary statistics.
- **Find Invisible Links** – Runs a simulation on the RAMS topology model to determine if there are any links where a failure will not be immediately detected because the routers that peer with RAMS will not report an EIGRP distance change.

Note The simulation can take several hours to run on a large network topology. You can cancel it at any time.

Traffic Menu

The *Traffic* menu allows access to traffic-related functions like reports and the ability to import interface capacity data. This menu is described in detail in [Chapter 7, “Network Planning”](#) and [Chapter 12, “Traffic Reports.”](#)

Help Menu

The *Help* menu contains the following three options:

- The **C**ontents option, which displays the PDF version of this guide in a new window.
- The **A**bout option, which displays RAMS version information.
- The **L**egend option, which opens the *Legend* panel. See [Legend Panel](#) on page 173 for more information.

Legend Panel

The *Legend* panel, shown in [Figure 41](#), appears by default in the upper right corner of the map when you open a topology. You can move the *Legend* panel anywhere inside the topology map window. Click a symbol in the *Legend* panel to highlight like symbols on the topology map. For example, click the **R**outer symbol to highlight all routers on the topology map.

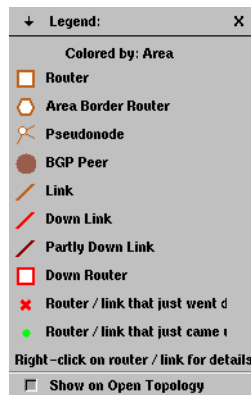





Figure 41 The Legend Panel



The content displayed on the *Legend* panel is dynamic. When you change the **C**olor **B**y option as described in [Colors Option Panel](#) on page 214, or from the *Colors Modes* submenu of the *View* menu, the *Legend* panel changes to reflect the colored elements on the map. For example, if you've chosen to color the map by **B**itrate, the *Legend* panel displays symbols corresponding to high bitrate links, medium bitrate links, and so on.

The following buttons are found on the *Legend* panel:

- The  **Arrow** button: Toggles between less detailed and more detailed symbol descriptions.
- The **Show on Open Topology** check box: Determines whether or not RAMS displays the *Legend* panel upon opening a topology. The box is selected by default.
- The  **Close** button: Closes the *Legend* panel. The next time you open a topology, the *Legend* panel will reappear unless you deselect the **Show on Open Topology** check box at the bottom of the panel. You can open the *Legend* panel at any time from the *Help* menu or by clicking  on the left toolbar.

Understanding the Topology Hierarchy

The *Topology Hierarchy* presents a hierarchical view (a tree structure) of the set of routing databases shown in the routing topology map. It lets you work with a subset of the topologies displayed in the RAMS window. This is particularly useful for large networks that contain a number of different IGP areas and BGP autonomous systems.


When requested, the Topology Hierarchy appears on the left side of the map window. You can display the Topology Hierarchy from the *View* menu or by clicking the  button on the toolbar. To close the Topology Hierarchy, click **Close (X)** in the upper right-hand corner of the Topology Hierarchy pane or click  again.

Next to each branch in the tree is a status indicator light. The color of each light indicates the state of the individual recorder. See [Table 4](#) for a list of colors and their corresponding meanings. You can also hover the mouse over each light to read a status message pertaining to the state of the recorder state.

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 menu bar, but they operate only on the areas selected from the tree structure. So, for example, by using the Topology Hierarchy *Tools* menu, it is possible to list the routers in just one area.


Alternatively, RAMS lets you display each individual area or AS in a separate window. Click **New View** to open a new topology map window containing only the selected areas.

To display an individual IGP area or BGP AS, perform the following steps:

- 1 Click  on the toolbar or select **Show Topology Hierarchy** from the *View* menu.

A tree structure appears displaying the names of all the administrative domains in the network.

- 2 For the network shown in [Figure 42](#), click Backbone in the tree structure.

- 3 Click **Select** in the Topology Hierarchy pane, and then click **Zoom In**  on the toolbar to expand the selected area of the topology map to fill the window. Alternatively, click **New View** in the Topology Hierarchy pane to open a new window containing just the selected area, as shown in [Figure 42](#).

Viewing Node and Link Details

RAMS stores many details about the nodes and links on the routing topology map besides the IP address and name labels that can be displayed next to the nodes. Since there is not enough space to display those details all at once, additional details about a particular node or link can be viewed by right-clicking on the object to open an information panel overlaid on the map.

Node Information Panel: Protocol Tab

Right-click on a node in the routing topology map to access the node information panel, as shown in [Figure 42](#). If the node is a pseudonode, the corresponding Designated Router (DR) is highlighted in orange on the topology map at the same time.

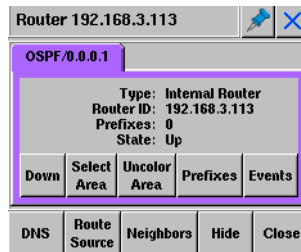


Figure 42 Node Information Panel

The title bar of the node information panel displays the name of the node and contains a tack button and a close button. Click the tack button to keep the panel open while opening the information panel for another object on the topology map. If you do not use the tack button, the current information panel disappears when you click anywhere outside the panel or on any of the panel buttons.

The body of the node information panel contains one or more tabs. Each tab displays the 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.

The details available for a router depend upon the protocol, but typically include the type of the node and one or more identifiers. In addition, the number of 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 state Up (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 following buttons are included on each tab:

- **Uncolor/Color Area** – Colors or removes colors from the area of the topology map that contains this node.
- **Select Area** – Selects all of the nodes within the routing area that contains this node. A bounding box is drawn around those nodes on the topology map, and then other operations can be performed on those nodes as described in the later sections of this chapter.
- **Prefixes** – Displays a list of the network prefixes announced by this router. The button is disabled when number of prefixes (shown on the information panel tab) is zero.
- **OSI Prefixes** – Displays a list of the OSI prefixes announced by this router.

Note If OSI IS-IS is not detected by the appliance, this option will not display.

- **Events** – Displays a list of the routing events originated by this router. If a time range has been 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.
- **Down** – In Design mode, brings down the node for planning purposes.

A second row of buttons at the bottom of the node information panel selects functions that apply to all instances of the node. The second row includes the following buttons:


- **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.
- **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 is displayed for another node, and the path between the two nodes is highlighted if you

click the button. For an illustration, see [Highlighting the IP Route Between Two Points in the Network](#) on page 207. This button is not present for pseudonodes.

Note Paths are generally highlighted in yellow; an orange highlight indicates that the route is not complete.

- **Neighbors** – Displays a list of neighboring routers and the interface addresses and metrics of the links connecting them.

Note For IS-IS routers that do not enable Traffic Engineering (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 be displayed in place of the source and destination interface addresses.

- **Hide** – Hides the selected node. To redisplay the node, select **Unhide All Nodes** from the *View* menu, or click the **Unhide All Nodes** icon  in the lower right corner of the *Routing Topology Map* window (see [Figure 47](#)). Click the button again to re-hide the nodes.
- **Close** – Closes the node information panel. The panel will also close if you click on the map.

Node Information Panel: Traffic Tab

Right-click on a node in the routing topology map to open a node information panel as described on [Viewing Node and Link Details](#) on page 177. If there is traffic information for the node, a **Traffic** tab appears beside the protocol tab(s) as shown in [Figure 43](#).

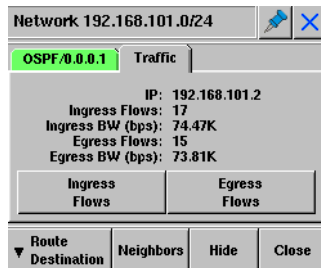


Figure 43 Node Traffic Tab

The body of the tab displays the IP address of the router, and information about ingress and egress flows. For example, *Egress Flows: 4* indicates that the router is exporting 4 traffic flows. *Egress BW (bps): 10* indicates that these 4 flows have a total bitrate of 10 bps.

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

Link Information Panel: Protocol Tab

Right-click on a link in the routing topology map to open a link information panel similar to the example shown in [Figure 44](#).

The title bar of the link information panel displays the name of the link and also contains a tack button and a close button. Click the tack button to keep the panel open while opening the information panel for another object on the topology map. If you do not use the tack button, the current information panel disappears when you click anywhere outside the panel or on any of the panel buttons.

Each link has two halves representing the two directions of communication. The router interface corresponding to the half that was clicked will appear on the left side of the link information panel, except for links to pseudonodes, in which case 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.




Figure 44 Link Information Panel

The link information panel can have one level of tabs, as shown in [Figure 44](#), or two levels of tabs, as shown in [Figure 45](#).

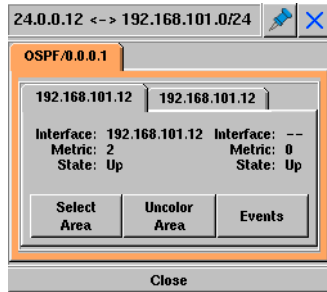


Figure 45 Link Information Panel, Multiple Parallel Links

In [Figure 45](#), the top-level tab indicates the instance of the link in a particular topology area (for example, OSPF area). If there is more than one top-level tab, 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 display 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. For example, in [Figure 45](#), the protocol is OSPF and the area identifier is 0.0.0.1. If the link is down, the tab label text will be red (or dark red if it is partially down).

Inside the top-level tab(s), there is an inner level of tabs if the link on the map represents multiple, parallel, physical links between the two routers. The inner tabs' labels indicate the interface of the router on the left side of the arrow in the title bar (for example, 192.168.101.12). The body of the tab displays details about the interfaces on the routers at each end of the link, including the following information:

- **Interface** – The address of the interface for the routers at either end of the link.
 - Note** For IS-IS routers that do not enable Traffic Engineering (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 be displayed in place of the source and destination interface addresses.
- **Metric** – The metric value for each interface. The link metric value helps traffic determine the best path to take through the network, and is based on bandwidth and delay, among other factors. For EIGRP protocol links, metric is shown as two components calculated from inverse bandwidth and delay.

- **State** – An 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 RAMS fails to receive information about BGP peering within a certain time frame, the link is marked Inactive, as RAMS cannot determine conclusively if the peering is down.
- **BW** (for EIGRP) – The bandwidth of the data traveling across the link.
- **Delay** (for EIGRP) – The 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 Area** – Colors or removes color from the area of the routing topology map that contains this link.
- **Select Area** – Selects all of the nodes within the routing area that contains this link. A bounding box is drawn around those nodes, and then other operations can be performed on those nodes as described in the later sections of this chapter.
- **Events** – Displays a list of the routing events related to adjacency changes on this link. If a time range has been selected in the *History Navigator* window, the same time range will be used for this list; otherwise, events occurring in the last 10 minutes will be listed.

The **Close** button at the bottom of the link information panel closes the panel. The panel will also close if you click on the topology map.

Link Information Panel: Traffic Tab

Right-click on a link in the routing topology map to open a link information panel as described on [Viewing Node and Link Details](#) on page 177. If there is traffic information for the link, a Traffic tab appears beside the protocol tab(s) as shown in [Figure 46](#).




Figure 46 Link Traffic Tab

The **Traffic** tab displays 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 bits per second (*bps*).

You can click **Flows** to display the *Details by Flow* table.

Viewing Complex Routing Hierarchies

In a complex network, it can be very difficult to see exactly what is happening. RAMS provides tools that let you tailor your view of the network so that the display contains just what you need to see.

Selecting Nodes in a Specific Area

In some situations, your investigations might focus on a specific area of routers in the network. RAMS lets you select specific nodes on the routing topology map and analyze the details of the selected nodes.

Note Be sure to save the routing layout (see [page 163](#)) before making any changes to the layout.

To select multiple nodes individually, perform the following steps:


- 1 Left-click on any node on the map. A selection box is drawn around the node.
- 2 Add nodes to the selection by holding down the **Ctrl** key and left-clicking on the desired nodes. The selection box expands to contain the additional node.

To select multiple nodes by dragging a selection box, perform the following steps:

- 1 Position the mouse cursor at one corner of the rectangular area bounding the nodes to be selected. The cursor must not be on a node.
- 2 Hold the left button and drag the cursor to the diagonally opposite corner of the rectangular area so that all desired nodes are included within the bounding box you have drawn.
- 3 Release the left button.

After you select the desired nodes by either method, you can perform various operations on those nodes, including the following:

- You can move the selected nodes to a different position on the map. With the mouse cursor inside the bounding box, hold the left button and drag the bounding box to the new position.

- You can zoom in on the selected nodes by clicking **Zoom In**  on the toolbar or selecting **Zoom In** from the *View* menu. **Figure 47** shows the result of zooming in on the lower portion of the network.
- You can access additional operations by right-clicking in an open space inside the bounding box area to display the *Selection* menu, as shown in **Figure 47**.

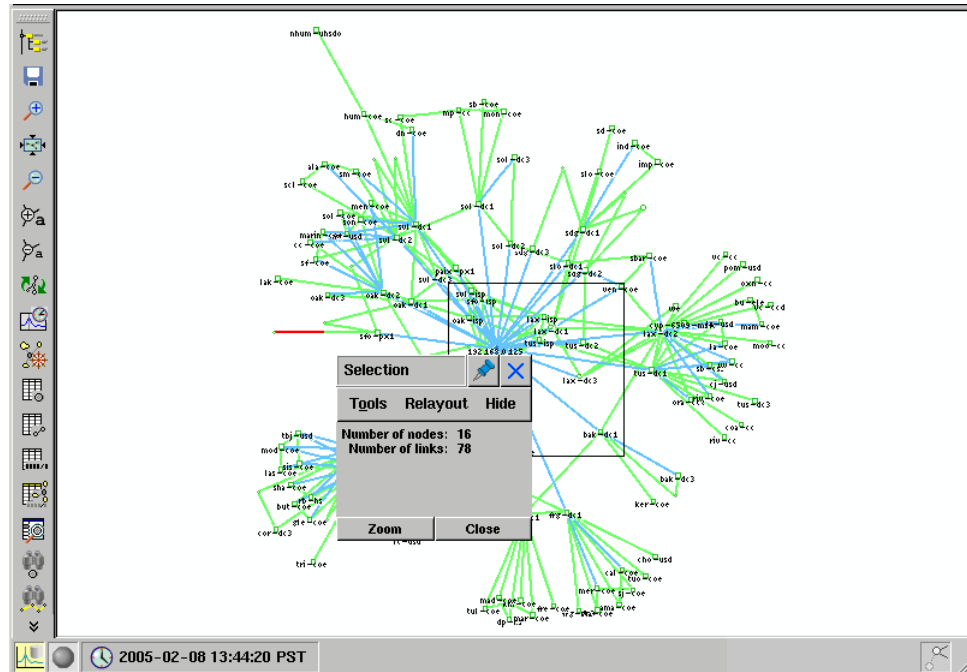



Figure 47 Selection Menu and Bounding Box

Note Only the elements completely within the bounding box are included in the counts shown in the *Selection* menu. Each direction of a link is counted separately, since one or both halves may be inside the bounding box.

The following options are available from the *Selection* menu:

- **Tools** – Use this option to open a drop-down menu of commands used to list routers, links, or prefixes that are within the bounding box.
- **Relayout** – Use this option to relayout the selected nodes, or relayout the unselected nodes.

- **Hide** – Use this option to hide the nodes in the selected area, hide the nodes outside the selected area, or unhide all nodes. To reveal hidden nodes, click the **Unhide All Nodes** icon  in the lower right corner of the *Routing Topology Map* window (see [Figure 47](#)). Click the button again to re-hide the nodes. Hidden nodes are saved with the layout.
- **Zoom** – Use this option to zoom in on the selected nodes.
- **Close** – Use this option to close the *Selection* menu.

Note If nodes were selected explicitly, only those nodes are affected by the specified action even if other nodes appear within the area of the bounding box.

In addition, the *Selection* menu contains a tack button. Click the tack button to keep the menu open while you perform other operations in the topology map. If you do not use the tack button, the *Selection* menu closes as soon as you click anywhere outside the menu.

Trimming the Displayed Nodes

You can instruct RAMS to display a particular class of router on the routing 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 **Trim Nodes** menu option to display only these routers.

To trim the displayed nodes, perform the following steps:

- 1 Select **Trim Nodes** from the *View* menu to open the *Trim Nodes* dialog box, as shown in [Figure 48](#).

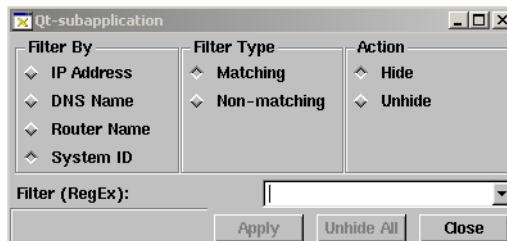


Figure 48 Trim Nodes Dialog Box

- 2 Choose among filtering by IP address, DNS name, router name, or System ID.


- 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 from the display, or click **Unhide** to restore them.
- 5 In the **Filter (RegEx)** box, type in a regular expression to select the class of routers to be matched. In the example shown in [Figure 48](#), `-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.

Note that extended regular expression syntax are used. Therefore, the following are metacharacters: `?`, `+`, `{`, `|`, `(`, and `)`. The syntax is not “glob” pattern matching, so use of `*-core-gw` is not correct.
- 6 Click **Trim**.

Trimming Leaves

Another useful trimming operation is leaf trimming. 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. Each time **Trim Leaf Nodes** is selected, edge nodes with only one link are removed.

Select **Trim Leaf Nodes** from the *View* menu to trim the nodes on the edge of the network. [Figure 49](#) illustrates the result of the Trim Leaf Nodes operation performed twice in succession. Trimmed (or hidden) nodes are saved with the topology layout.

To restore hidden nodes, select **Unhide All Nodes** in the *View* menu, or click the **Unhide All Nodes** icon  in the lower right corner of the *Routing Topology Map* window (see [Figure 49](#)).

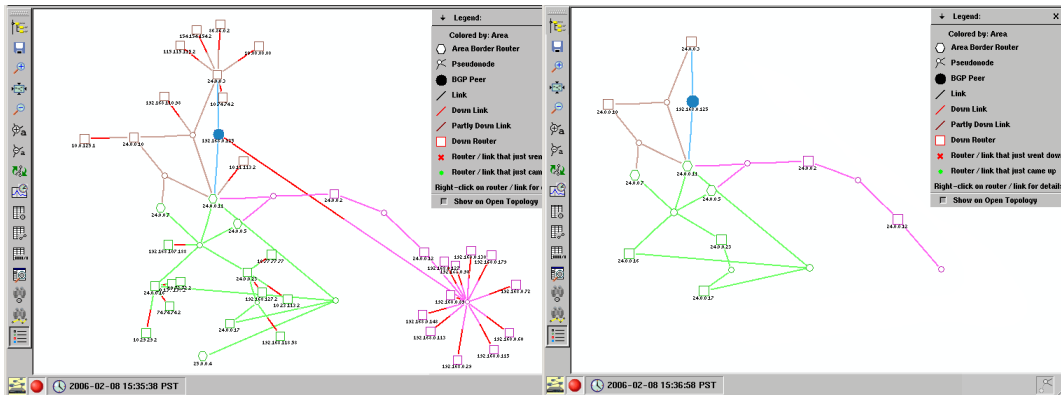


Figure 49 Trimming Leaves - Before and After

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

Finding a Router

When you know the name or address of a router and want to find where it is on the map, use the method described here. If you would prefer to scan the list of router names and addresses to find a router, see [Router List](#) on page 191.

To find a router using the Find Router method, perform the following steps:

- 1 Select **Find Router** from the *Tools* menu to open the *Find Router or LAN Node* dialog box, as shown in [Figure 50](#).




Figure 50 Find Router Dialog Box

- 2 In the *Search For* text box, type the IP address, name or System ID of a router, or the prefix of a LAN pseudonode. If the name entered is the initial portion of multiple router names, then all those routers will be matched.
- 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 steps 2 and 3.

Viewing Network Anomalies At A Glance

The RAMS topology map enables you to quickly see possible points of failure, failed links and routers, and other anomalies. For example, when a link goes down, that link turns red on the routing topology map.

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

- 1 Open the desired routing topology in RAMS.
- 2 Look for links that are shown in red. If a link is bright red, it means that the link has gone down. If it is dark red, then the link represents multiple parallel physical links only some of which are down.
- 3 Look for parts of the network that route through a single router (square or hexagon) or LAN (circle).

Viewing a Complete Up-to-the-Minute Network Inventory


RAMS can display a complete, up-to-date 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, and so on, and each entry in any list can be traced back to the associated router in the topology map with a single click.

Router List

You can use the *List of All Routers* dialog box to perform the following tasks:

- 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).
- Display 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* dialog box, perform the following steps:


- 1 Click the **List Routers** button  on the toolbar or select **List Routers** on the *Tools* menu to open the *List of All Routers* window.
- 2 Use the **Filter by** drop-down list at the top of the window to filter the routers displayed in the window (see [Using Filters](#) on page 280).
- 3 Scroll down the list to find the desired router.
- 4 To identify where a particular router is 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



The *List of All Links* dialog box displays 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.

Note For IS-IS routers that do not enable Traffic Engineering (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 be displayed in place of the source and destination interface addresses.

To view the List of All Links dialog box, perform the following steps:

- 1 Click the **List Links** button  on the toolbar or select **List Links** from the *Tools* menu.

The *List of All Links* window opens, as shown in [Figure 51](#).

- 2 Use the **Filter by** drop-down list at the top of the window to filter the links displayed in the window (see [Using Filters](#) on page 280).
- 3 Perform any of the following operations using the *List of All Links* dialog box:
 - Select a link in the list, which causes the link to flash yellow on the map.
 - Copy a single row of the table using **Ctrl-C** or the entire table using **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 the  button. This operation copies to the clipboard a subset of the data in each row, and does so in the format required for import in the *Router/Link Edits* window.
 - If the routing topology has changed since the dialog box was opened, refresh the contents of the dialog box by clicking the  button.
- 4 Close the dialog box by clicking **Close**.

Link	Interface Name	Source Interface	Destination Interface	Bandwidth	Delay	Metric (EIGRP: bw+dly)	State	Area or AS
IGP Link Router1								
Router16 → Router16	Fa1/0	192.168.116.16/24	192.168.116.26/24	100000 Kbps	1000 us	25600+25600	Up	TestLab.EIGRP/AS
Router16 ← Router16	Fa0/0	192.168.116.26/24	192.168.116.0/24	100000 Kbps	100 us	25600+2560	Up	TestLab.EIGRP/AS
IGP Link Router1								
Router16 → Router16	Fa1/1	192.168.107.16/24	192.168.107.23/24	20000 Kbps	200 us	128000+5120	Up	TestLab.EIGRP/AS
Router16 ← Router16	Fa1/0	192.168.107.23/24	192.168.107.16/24	100000 Kbps	100 us	25600+2560	Up	TestLab.EIGRP/AS
IGP Link Router1								
Router16 → Router16	Se2/0	192.168.130.2/24	192.168.130.1/24	44210 Kbps	200 us	57856+5120	Up	TestLab.EIGRP/AS
Router16 → Router16	Se2/1	--	--	44210 Kbps	200 us	57856+5120	Up	TestLab.EIGRP/AS
Router16 ← Router16	Se2/1	--	--	44210 Kbps	200 us	57856+5120	Up	TestLab.EIGRP/AS
Router16 ← Router16	Se2/0	192.168.130.1/24	192.168.130.2/24	44210 Kbps	200 us	57856+5120	Up	TestLab.EIGRP/AS
IGP Link Router2								
Router23 → Router23	Tu4	--	--	9 Kbps	500000 us	444416+12800000	Up	TestLab.EIGRP/AS
Router23 ← Router23	Tu3	--	--	9 Kbps	500000 us	444416+12800000	Up	TestLab.EIGRP/AS
IGP Link Router1								
Router16 → Router16	Fa0/0	192.168.109.16/24	192.168.109.20/24	10000 Kbps	2500 us	256000+64000	Up	TestLab.EIGRP/AS
Router16 ← Router16	Et0	172.33.2.0/24	192.168.109.16/24	10000 Kbps	10 us	256000+256	Up	TestLab.EIGRP/AS
IGP Link Router2								
Router26 → Router26	Fa1/0	192.168.112.26/24	192.168.112.20/24	100000 Kbps	100 us	25600+2560	Up	TestLab.EIGRP/AS
Router26 ← Router26	Et1	192.168.112.20/24	192.168.112.26/24	10000 Kbps	100 us	256000+2560	Up	TestLab.EIGRP/AS
IGP Link Router1								
Router16 → Router16	Fa1/0	192.168.116.16/24	192.168.116.17/24	100000 Kbps	1000 us	25600+25600	Up	TestLab.EIGRP/AS
Router16 ← Router16	Et1	192.168.116.0/24	192.168.116.17/24	10000 Kbps	1000 us	256000+25600	Up	TestLab.EIGRP/AS
IGP Link Router2								
88 top level entries, 280 total entries								


Figure 51 List of All Links Window


Prefix List

Use the *List of Prefixes* dialog box to obtain the following information:

- View the routers in a network that are advertising default routes.
- 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 are properly advertising prefixes.
- Display the advertised BGP prefixes and the list of attributes associated with each BGP prefix.

To view the List of Prefixes dialog box, perform the following steps:

- 1 Click the **List Prefixes** button  on the toolbar or select **List Prefixes** from the *Tools* menu to open the *List of Prefixes* window as shown in [Figure 52](#).
- 2 Use the **Filter by** drop-down list at the top of the window to filter the prefixes displayed in the window (see [Using Filters](#) on page 280).
- 3 You can perform the following operations using this dialog box:
 - Select a router in the list, and the node will flash yellow on the map.
 - If the routing topology has changed since the dialog box was opened, refresh the contents of the dialog box by clicking **Reload**.
- 4 Close the dialog box by clicking **Close (X)**.



Prefix	Router/Net	Attributes	State	Area or AS
0.0.0.0/0	Router23	Metric: bw=256000 dly=2560 (External)	Up	trunk4062.RA148.EIGRP/AS1
0.0.0.0/0	24.0.0.12	Metric: 1 (AS External) Forward Addr: 192.168.0.254	Up	trunk4062.RA148.OSPF/0.0.0.1
0.0.0.0/0	24.0.0.12	Metric: 1 (AS External) Forward Addr: 192.168.0.254	Up	trunk4062.RA148.OSPF/0.0.0.3
0.0.0.0/0	24.0.0.12	Metric: 1 (AS External) Forward Addr: 192.168.0.254	Up	trunk4062.RA148.OSPF/Backbone
10.23.89.0/30				
10.23.89.0/30	24.0.0.5	Metric: 11112 (Area External)	Up	trunk4062.RA148.OSPF/0.0.0.1
10.23.89.0/30	24.0.0.11	Metric: 11113 (Area External)	Up	trunk4062.RA148.OSPF/0.0.0.1
10.23.89.0/30	24.0.0.7	Metric: 15207 (Area External)	Up	trunk4062.RA148.OSPF/0.0.0.3
10.23.89.0/30	24.0.0.11	Metric: 11113 (Area External)	Up	trunk4062.RA148.OSPF/0.0.0.3
10.23.89.0/30	24.0.0.23	Metric: 11111	Up	trunk4062.RA148.OSPF/Backbone
10.23.89.0/30	Router23	Metric: bw=256000 dly=2560 (External)	Up	trunk4062.RA148.EIGRP/AS1

147 top level entries, 865 total entries

Figure 52 List of Prefixes Window

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

- 1 Locate the node on the map.
- 2 Right-click on the node to open the node information panel.
- 3 Click **Prefixes**.

The *List of Prefixes* dialog box opens, showing all prefixes advertised by the node you selected. For each prefix, all nodes that advertise the prefix are listed.


Note The names and addresses in the **Router/Net** column are the routers that are advertising the prefix. The way that the routers are displayed 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

Use the *List of OSI Prefixes* dialog box to obtain the following information:

- View the routers in a network that are advertising default routes.
- Determine if a prefix is currently advertised or not, and by which routers.
- See what metric is advertised with each prefix.

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

- 1 Click the **List Prefixes** button  on the toolbar or select **List of OSI Prefixes** from the *Tools* menu. to open the **List of OSI Prefixes**.
- 2 Use the **Filter by** drop-down menu list at the top of the window to filter the OSI Prefixes and ES Neighbors that are displayed in the window.
- 3 Close the window by clicking the **X** icon found at the bottom of the window.

Viewing the Exit Routers from the Network for any Internet Destination

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 helps you infer the flow of traffic to service providers or peers, helping balance flows to achieve optimum performance or minimize monthly cost in transit fees and bandwidth costs.

RAMS 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 209.

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

- 1 Select **Highlight by Exit Router** from the *Tools* menu.
- 2 Type the desired Internet prefix, NSAP address, or domain name in the *Highlight By Exit Router* dialog box (shown in [Figure 53](#)).
- 3 Click **OK**.




Figure 53 Highlight by Exit Router Dialog Box

Using the Network Summary Option

The network summary option in the *View* menu summarizes network information for all of the topologies found in your network.

To access this option, select **Network Summary** from the *View* menu. The Network Summary window will appear at the bottom-left portion of your screen.

The following describes the information displayed for each protocol:

- **IGP** — Displays the number of nodes, links, and prefixes on the network.
- **BGP** — Displays the number of active prefixes, active routes, baselines, next hop, and Neighbor AS's on the network.
- **EIGRP** — Displays the number of nodes, links, prefixes, and the number of unique prefixes on the network.
- **IS-IS** — Displays the number of nodes, pseudonodes, overloaded nodes, links, prefixes, unique prefixes, and the number of OSI prefixes on the network.
- **OSPF** — Displays the number nodes, pseudonodes, links, prefixes, and the number of unique prefixes on the network.
- **VPN** — Displays the number of active prefixes, active routes, baselines, MP next hops, neighbor AS's on the network.


This information will enable you to view up-to-date counts of your network, and monitor the effect of network changes for various network counts.

Assigning Router Names using Router Name Repository

The Router Name Repository tool allows you to have control over how the routers in your system are identified by the appliance. By default, the IP Address of the router is used to identify the router.

To access Router Name Repository, select **Router Name Repository** from the *Tools* menu.

The Router Name Repository Page opens, as shown in [Figure 54](#).



The screenshot shows a window titled "Router Name Repository" with a filter set to "Any". The window contains a table with the following columns: IP Address, System ID, Router Name, DNS Name, and User Specified Name. The table lists 37 entries, each with an IP address in the first column and empty cells in the others. At the bottom of the window, there are buttons for "Save", "Resolve DNS", "Import", and "Close", along with a status bar indicating "37 entries".

IP Address	System ID	Router Name	DNS Name	User Specified Name
24.0.0.2				
24.0.0.3				
25.0.0.4				
24.0.0.5				
24.0.0.7				
24.0.0.10				
24.0.0.11				
24.0.0.12				
24.0.0.16				
24.0.0.17				
192.168.0.22				
24.0.0.23				
192.168.0.29				
192.168.0.228				
192.168.0.36				
192.168.0.53				
10.22.22.22				
192.168.0.60				
192.168.0.68				
192.168.0.70				
192.168.0.72				
10.87.87.2				
10.228.228.2				
192.168.0.87				
10.29.29.2				
192.168.0.98				
192.168.0.113				
192.168.0.115				
10.20.20.20				
192.168.0.125				
192.168.0.138				
10.88.88.2				
192.168.0.161				
192.168.0.171				
10.23.113.2				

Figure 54 Router Name Repository

The following list describes the columns in this page:

- **IP Address** — Displays the IP Address for a particular router.

- **System ID** — Displays the System ID for the router received from the router (otherwise, this column will be empty).
 - Note** The System ID column will appear only if IS-IS is detected by the appliance.
- **Router Name** — Displays the router name derived from the routing protocol.
- **DNS Name** — Displays the name resolved with the DNS server using the routers' IP address.
- **User Specified Name** — Column where you can enter a unique name for the router.

The following lists the order in which the appliance prioritizes router names:

- 1 **User Specified Name** (highest priority)
- 2 DNS Name
- 3 Router Name
- 4 **Router IP Address** (lowest priority)

The following list describes the tabs that appear in this page:

- **Filter By** — Choose which router entries you want to show or hide in the table with the following choices:
 - **Any** — Selecting this will show all rows.
 - **Router IP Address** — Selecting this choice displays a drop-down menu as shown in [Figure 55](#). This menu allows you to enter the IP address of the router, and also displays previously entered IP addresses used for the router. If you select the routers' IP address and select **Show**, the system will display only the row with the IP address you entered. If you select **Hide**, the system will show all the router ip addresses except the one you entered.

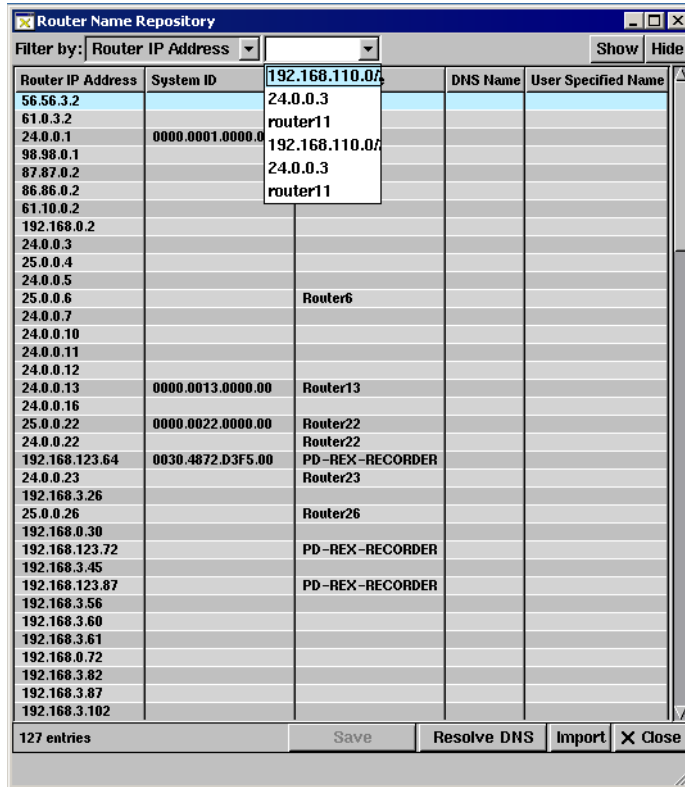


Figure 55 Router Name Drop-Down Menu

- **Router Names** — Selecting this choice displays a drop-down menu which allows you to enter a new name for the router, and also displays previously used names for the router. Simultaneously, an *Options* drop-down menu opens. This menu allows you to filter by following choices:
 - Substring: Filters the routers with the given string as a substring for either Router Name, DNS Name, or User Specified Name.
 - Exact Match: Filters the routers with the given string as an exact match either of its Router Name, DNS Name, or User Specified Name.
 - Begins With: Filters the routers with the beginning string used for either Router Name, DNS Name, or User Specified Name.

Show — Select this to show all router entries you want to view.

Hide — Select this to hide router entries.

Save — Select this to save information you edit. This button will be active only when you have edited information on the screen.

Import — Select this when you want to edit several router names.

Close — Select this to exit the Router Name Repository page.

To change a router's name

1 From the *Tools* Menu, select *Router Name Repository*.

The *Router Name Repository* page opens.

2 Assign a new name one of two ways:

a If you have only one router you wish to rename, find the row the router is located on and enter the new name in the *User Specified Name* column.

b Click **Save**. The new router name now appears in the *User Specified Name* column, or

c Click **Import** if you have more than one router to re-name.

The *Import Router Names* dialog box opens, as shown in [Figure 56](#).

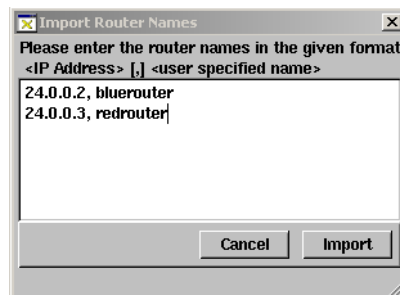


Figure 56 Import Router Names Dialog Box

3 Enter the router names in the specified format,

4 Click **Import**.


The new router names now appears in the *User Specified Name* Column.

Assigning AS Names using AS Name Repository

The Autonomous Systems (AS) repository allows you to assign a name to the AS. This name will take 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 access AS Name Repository, select **AS Name Repository** from the *Tools* drop-down menu.

The *AS Name Repository* Page opens, as shown in [Figure 57](#).



The screenshot shows a window titled "AS Name Repository" with a filter set to "Any". The table contains 35 rows of AS data. The columns are "AS Number", "AS Name (whois)", and "User Specified Name". The "User Specified Name" column is currently empty for all entries. At the bottom of the window, there are buttons for "Save", "Import", and "Close", and a status bar indicating "28050 entries".

AS Number	AS Name (whois)	User Specified Name
1	gntly_1	
2	dcn	
3	mit_gateways	
4	isi	
5	symbolics	
6	bull_hn	
7	uk_defence_research_agency	
8	rice	
9	cmu	
10	csnet_ext	
11	harvard	
12	nyu	
13	brl	
14	columbia_gw	
15	net_dynamics_exp	
16	lbl	
17	purdue	
18	utexas	
19	css	
20	ur	
21	rand	
22	nosc	
23	nisa_sip	
24	ames_nas_gw	
25	ucb	
26	cornell	
27	umdnet	
28	dfvtr_sys	
29	yale	
30	sri_aicnet	
31	cit	
32	stanford	
33	dec_vr1	
34	udelnet	
35	micaton	

Figure 57 AS Names Repository Page

The following list shows describes the columns of this page, and also shows the order in which the appliance prioritizes AS names:

- 1 **User Specified Name** — Enter a unique AS name in this column. (highest priority)
- 2 **AS Number** — Displays the number identifying the AS number.
- 3 **AS Name** — Displays the name of the AS derived from the Whois server. (lowest priority)

The following list describes the tabs that appear in this page:

- **Filter By** — Choose which AS Name and number entries you want to show and/or hide in the table with the following choices:
 - **Any** — Selecting this will show all rows.
 - **AS Number** — Selecting this choice displays a drop-down menu, which allows you to filter by entering the AS Number, while also displaying previously used AS Numbers. Simultaneously, an *Options* drop-down menu also opens, which allows you to filter with the following choices:
 - **Greater Than**: Displays AS numbers with a greater number than the one entered in the text box field.
 - **Exact Match**: Displays only exact matches entered in the text box field.
 - **Begins With**: Displays matches that begin with numbers entered in the text box field.
 - **AS Names** — Select this choice and a text box field opens where you can enter AS Names. Simultaneously, an *Options* drop-down menu 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 displayed.
- **Hide** — Select this to hide router entries.

- **Save** — Select this to save information you edit. This button will be 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 Name Repository*.

To change the AS Name:

- 1 From the Tools Menu, select *AS Name Repository*
The *AS Name Repository* page opens.
- 2 Assign a new name one of two ways:
 - a If you have only one AS name you wish to rename, find the row the AS name is located on and enter the new name in the *User Specified Name* column.
 - b Click **Save**. The new AS name now appears in the *User Specified Name* column, or
 - c Press **Import** if you have more than one AS name to be re-named.
The *Import AS Names* dialog box appears, as shown in [Figure 58](#).

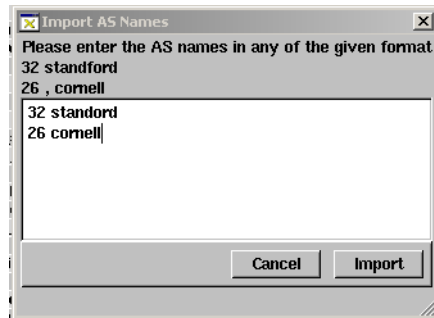


Figure 58 Import AS Names Dialog Box

Enter the AS names in the specified format.

- d Click **Import**.
The new AS names now appear in the *User Specified Name* Column.

Verify and Manually Assign BGP AS Assignments to Routers

RAMS can display the path resolved between two points in a network, as described in the next section. For network configurations that include BGP, RAMS requires correct BGP AS assignments to routers 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, RAMS can create BGP AS assignments for all routers automatically, but some routers may not be running BGP. For these cases, you can change the BGP AS assignment manually in the *BGP AS for Routers* dialog box.

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, RAMS 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:

- 1 On the *Tools* menu, click **Assign BGP AS to Routers** to open the *BGP AS for Routers* dialog box as shown in [Figure 59](#). Some routers may have AS numbers already assigned to them as detected by BGP peering or computed from network topology.
- 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 **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**.

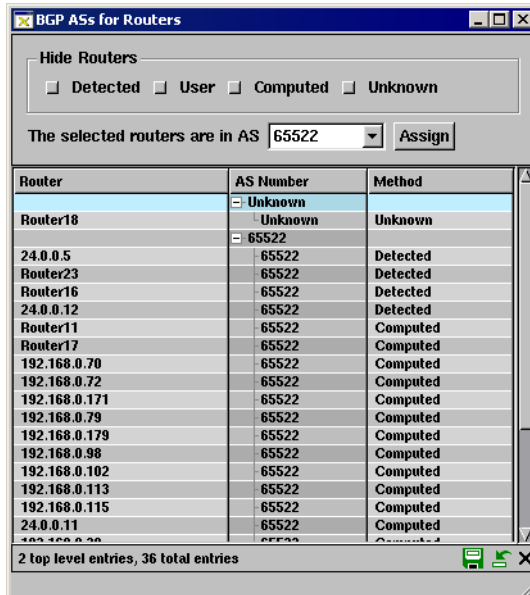


Figure 59 BGP AS to Routers Dialog Box

Highlighting the IP Route Between Two Points in the Network

Seeing the IP paths taken by traffic from a source router to a destination router in a multidomain network is very useful for network planning. There can be paths of high importance in a network such as a VoIP service path between an IP PBX and a PSTN media gateway that would not tolerate significantly increased delay resulting from a possible rerouting due to link or router failures.

RAMS can quickly display the path resolved between two nodes in a network at the current time or at any point in recorded topology history. The path from a router towards another highlights in yellow (or orange if the path is not complete), and the return route will highlight in green. Each segment of the path can be listed, along with the link metric and the prefix by which each next hop was resolved, using the *List/Find Paths* dialog box described in [Finding a Route By Prefix](#) on page 209.

The path is selected using the information panels for the source and destination nodes on the routing topology map. [Figure 60](#) shows the node information panels for the source and destination of a route.

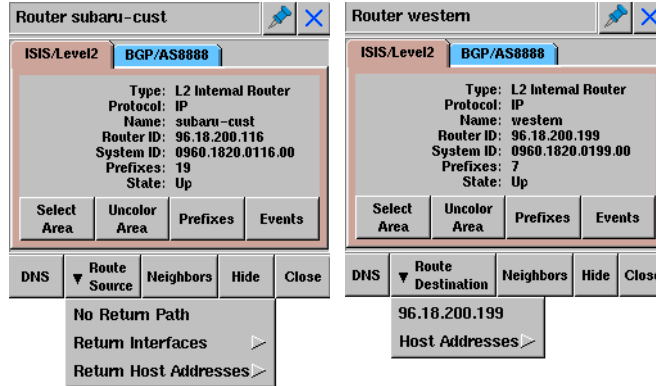


Figure 60 Source and Destination Node Information Panels

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

- 1 Right-click on the source node on the routing topology map.
The node information panel appears.

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

No Return Path — Displays route source only.


Return Interfaces — Specifies return interface for route source.

Return Host Addresses — Displays the Host address.

The route from one router toward another router is highlighted in yellow on the routing topology map. The return route highlights in green

- 3 Right-click on the destination node on the routing topology map.
- 4 Click **Route Destination** in the node information panel that appears, and then select an interface from the drop-down list or select **Host Addresses**.

The route from one router toward another router highlights in yellow on the routing topology map, with return path highlighting in green.

- 5 To see the path details, select  from the toolbar or **List/Find Paths** from the *Tools* menu.

Note 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 RAMS, 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, RAMS 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 Open the *Tools* menu and select **List/Find Paths**.

The *Find or List Paths* dialog box opens.

- 2 Type the IP address or System ID of the source router or name in the *Source Router* text box.

Alternatively, selecting a node on the topology map, and then selecting the *Source Router* button at the top left of the *Find or List Paths* dialog box specifies the router (or one of its interfaces) as the source.

- 3 Type the destination IP address, Internet prefix or domain name in the *Destination Prefix* text box.
- 4 Click **OK**.

The route is calculated to the destination prefix if internal or to the nearest exit router if external. The segments of the path are displayed in the lower section of the *Find or List Paths* dialog box, including the link metric and the prefix by which each next hop was resolved. The forward path highlights in yellow on the routing topology map and the return path displays in green.

Note The paths will flash momentarily on the topology map when its corresponding entry in the Find or List Paths dialog box is selected.

Multiple paths are shown for equal-cost multi-path routes.

Note 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 that case, the path might end at a LAN pseudonode.

Viewing the Highlighted Path Cost for EIGRP

To view the details of a highlighted path, select **List/Find Paths** from the *Tools* menu. An example for the EIGRP protocol is shown in [Figure 61](#).

The resulting *List/Find Paths* dialog box displays, for each link along the path, the source and destination nodes of the link, the metric cost of the link, the protocol used by multiprotocol routing to select that link, and the prefix used to resolve that hop.

The *Metric (bw)* and *Metric (delay)* columns for EIGRP show the cost in EIGRP metric units associated with the bandwidth and delay values that are configured for the link. The **Metric** column is not as obvious; it lists the amount that each link contributes to the overall path distance, or cost. The path cost for EIGRP is the sum of the *delay* values for each hop plus the maximum of the *bw* values (which would be the lowest bandwidth link since the *bw* values correspond to the inverse of the link bandwidth).

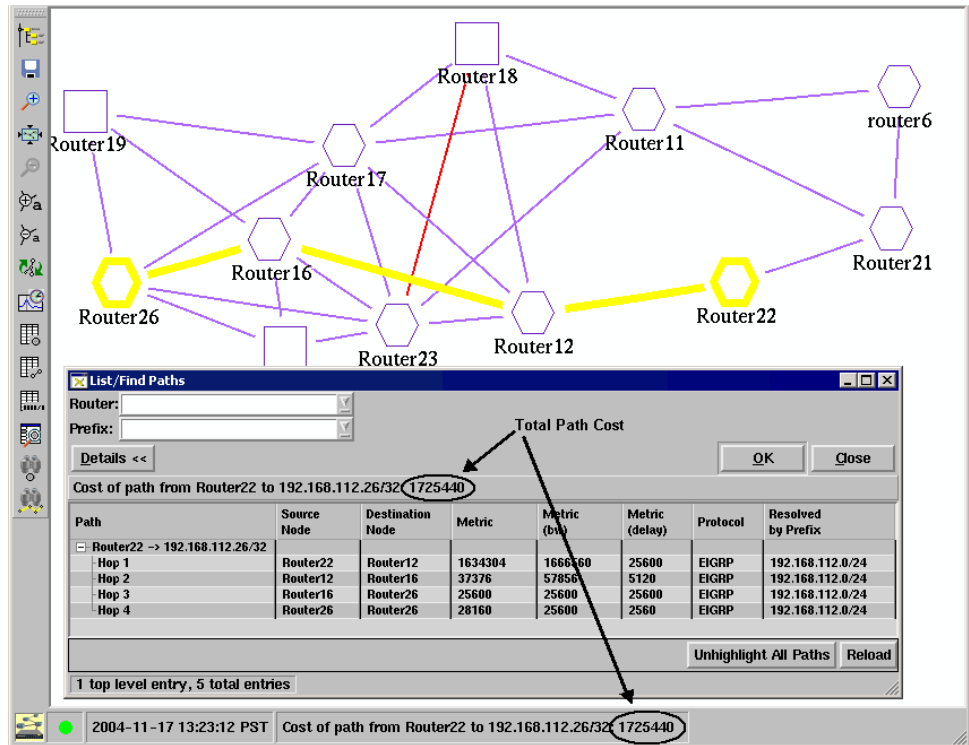


Figure 61 Highlighted Path Details for EIGRP

Since the EIGRP protocol calculates routes from the destination back towards the source, the **Metric** column needs to be read from bottom to top. In [Figure 61](#), hop 4 contributes both the **bw** and **delay** values to the total cost of 28160. At hop 3 the maximum **bw** value is unchanged, so hop 3 increases the total cost only by its **delay** value 25600. Hop 2 adds the amount by which its **bw** value is larger than the current maximum of the **bw** values (57856-25600) plus its **delay** value 5120, for a total of 37376. At hop 1, the maximum **bw** value increases again (1666560-57856) and a **delay** of 25600 is added. If all the values in the **Metric** column are added, the total path cost is the same as is reported in the *List/Find Paths* dialog box and in the status bar of the *Routing Topology Map* window.

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 and are therefore meaningless to add. In such cases, the status line indicates that the path cost cannot be calculated.

Using the Configuration Options Dialog Box

Use the Tools menu to select Options. The RAMS Configuration Options dialog box opens, where you can set preferences for various parameters of RAMS operation. The dialog box contains a hierarchical list of option categories. Select a category to display the options in that category in the right pane of the dialog box.

To close the dialog box and save any changes you have made, click **Close** at the bottom of the dialog box. To restore all options to their factory default values, click **Factory Settings** at the bottom of the dialog box.

Analysis Options Panel

There are two subcategories of analysis options. The first, **Visualization**, lets you customize the level of detail that is displayed in visualizations and animations of the BGP RIB as described in [RIB Visualization](#) on page 242. The second, **Algorithmic**, lets you set the thresholds used in Root Cause Analysis, which is described in [Root Cause Analysis](#) on page 235. In each case, a higher value decreases the level of detail, and a lower value increases it.

- 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 entities are as follows:
 - 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.
- Algorithmic options control thresholds used by the Root Cause Analysis function:

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

Node Labels Option Panel

The *Node Labels* option panel determines which node details to display on the routing topology map when it is first opened. By default, the **Automatic** option is enabled. In Automatic mode, RAMS automatically chooses a label for the node. For example, if no router name is available for the node, RAMS will use the DNS Name node label. Label options are prioritized in the following order:

- *Router Names* – Displays the name of the router obtained from the protocol (if available).
- *DNS Names* – Displays the router DNS name. If no DNS name resolution has been performed yet, selecting this mode will initiate DNS name resolution for all routers. In a large network, this can take some time.
- *IP Address* – Displays the router IP address.
- *System IDs (IS-IS only) or NSAP Address* – Displays IS-IS System IDs when IS-IS is present on the map. If **Show IS-IS NSAP Addresses** is selected on the *Miscellaneous* options panel as described in [Miscellaneous Options](#) on page 214, the label *NSAP Address* is used.

Selecting one or more of the options above disables Automatic mode.

You can select the following two options independently of Automatic mode:

- *Label Routers Only* – Does not label the pseudonodes representing LANs.
- *ID Numbers* – Displays internal ID number for the router, which may be useful as a shorthand reference.

In other words, when you enable the options above, you do not disable Automatic mode.

The menu options displayed depend on the network protocol. For example, all of the above options appear for an IS-IS network, while Router Names and System IDs do not appear for an OSPF network.

Colors Option Panel

The *Colors* option panel determines how RAMS uses color on the topology map.

- *Color Allocation Order* – Changes the default colors used to signify routers, links, and so on, when you click and drag color samples in the chart. For example, to make orange the first color RAMS uses to color elements 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 moves to position 5. Click **Set** to save the change or click **Default** to return to factory settings.
- *Color links with metrics greater than* - Use this option to set the links to a larger value to make them visually more noticeable.
- *Traffic Coloring Options* – Changes the default colors used to signify traffic flows when you click the **Color** box corresponding to a category of traffic flow (*High, Medium, or Low*) and select a new color from the palette. If desired, you can change the definition of traffic flow categories by editing the *Mbps or % Utilization* text box for *High, Medium, or Low*. Click **Set** to save the change or click **Default** to return to factory settings.
- *On Open Topology: Color Areas* – The nodes and links of each area are displayed in distinct colors, which are determined by the Color Allocation Order chart, when a topology is loaded.
- *On Open Topology: 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.
- *On Open Topology: 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.
- *On Open Topology: Don't Color* – All areas and traffic flows are displayed with black nodes and gray links when a topology is loaded.

Miscellaneous Options

The *Miscellaneous* panel allows you to set the following preferences:

- *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 manually, or set one automatically by saving a topology layout in the *Save Layout* dialog box invoked from the *Topology* menu and clicking the **Use as default layout** check box on that dialog box.
- *Show Legend on Open Topology* — Automatically displays the *Legend* panel when you load a topology. See [Legend Panel](#) on page 173 for information about the *Legend* panel.
- *Show IS-IS NSAP Addresses* — This option is displayed only when the loaded topology contains IS-IS protocol routers. When this option is enabled, NSAP addresses are displayed on the topology map rather than the system ID for the IS-IS router. See [Node Labels Option Panel](#) on page 213 for more information.
- *Event panel default time interval* — Sets the time period (default 600 seconds) that is included when you open a list of 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.
- *Hide DNS Suffix* — Determines how DNS names are displayed on the routing topology map. When DNS names of nodes are displayed, 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 ECMP (Equal Cost Multi-Path) routes to compute and display when you highlight a path between two routers or use the **List/Find Path** option from the *Tools* menu. The default is 4096; set to 1 to disable ECMP routes.

History Navigator Options

The *History Navigator* panel allows you to set default parameters for the *History Navigator* window.

- The initial collection of graphs to be displayed in the *History Navigator* window may be selected by checking the desired options. Only the *Events* graph is enabled by default.

- *Events* graph — Displays the number of routing protocol changes that occurred in the network between recorded snapshots. Example routing events are a neighbor adjacency going down or a new prefix being announced.
- *Routers* graph — Displays 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* graph — Displays the sum of router-to-router links plus the number of router-to-prefix links. Does not apply to BGP protocols.
- *Prefixes* — Displays the cumulative number of prefixes available in the network.
- *Routes* graph — Displays the number of routes advertised in the network. Does not apply to IGP protocols.
- *Value of playback step in seconds* — Sets the default value for a single step forward or backward in time during *History Navigator* playback. The factory setting is 600 seconds.
- *Max number of data points in event graph* — Limits the event graph to the specified number of data points. The factory default is 25,000 data points.
- *Online Mode Update interval* — Sets the number of seconds between updates. The factory default is 10 seconds.

Auto-Hide Options

Since RAMS cannot determine whether a node or link that goes down has temporarily failed or been decommissioned, an adjustable timeout interval determines when nodes and links that are down should be removed from the map.

The following options are available on the *Auto-Hide* panel of the *Configuration Options* dialog box:

- *Seconds to auto-hide detached nodes* — A node can become detached from the rest of the network when all of its links are auto-hidden due to one of the conditions above (for example, the node has failed temporarily or is decommissioned). Using this option, specify the number of seconds after which the detached node will be hidden.

The default setting is -1, which indicates that the option is disabled.

- *Seconds to auto-hide pseudonodes with one attachment* — This option is similar to *Seconds to auto-hide detached nodes*, except that it applies only to pseudonodes whose links are still visible. Often, such a condition is caused by a bug 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. By default, the option is set to -1 (disabled).

- *Seconds to auto-hide failed links* — When a link fails or is decommissioned from service permanently, it is colored red on the routing topology map. Specify how many seconds should pass before such links are no longer displayed on the map.

By default, this option is set to auto-hide failed links in 43,200 seconds (12 hours). To disable this option, change the default value to -1. Note that the link will be displayed once again if you use the *History Navigator* window to view a time when the link was up.

Note Since links between BGP routers and their clients are inferred, as described in [Understanding Links and Peerings on the Map](#) on page 162, RAMS cannot conclusively determine if such a peering has failed or is simply inactive. Therefore, if RAMS does not receive routing information about such a peering within a certain amount of time, the peering is marked Inactive rather than Down, and is not removed from the map.

- *Seconds to auto-hide failed nodes* — Analogous to the *Seconds to auto-hide failed links* option. Specify how many seconds should pass before such nodes should be considered decommissioned and are no longer displayed on the map.

By default, this option is set to auto-hide failed nodes after 43,200 seconds (12 hours). To disable this option, change the default value to -1.

- *Seconds to auto-hide failed links to pseudonodes* — Analogous to the *Seconds to auto-hide failed links* option, but applies where one end of the failed link is a pseudonode. By default, this option is set to zero (hide immediately); the pseudonode for a broadcast network will change whenever a new Designated Router is elected.

- *Seconds to auto-hide failed pseudonodes* — Analogous to *Seconds to auto-hide failed node* where the failed node is a pseudonode. This option defaults to zero (hide immediately); the pseudonode for a broadcast network will change whenever a new Designated Router is elected.

Diagnosing EIGRP Topology Errors

For EIGRP topologies, RAMS includes a *Topology Diagnostics* submenu in the *Tools* menu. You can select four different dialog boxes from this submenu to diagnose both anomalies in the network topology and problems with the modeling of the topology. The following four dialog boxes are accessible from the *Topology Diagnostics* submenu:

- **List Topology Errors** — Helps detect configuration anomalies.
- **List Inaccessible Routers** — Helps detect router accessibility problems that interfere with accurate discovery of the network topology by RAMS.
- **List Mismatched Distances** — Helps detect inaccuracies in the topology map by listing prefixes whose distance (metric) to the prefix reported by a router that peers with RAMS does not match the distance that RAMS calculates across its model of the topology.
- **Find Invisible Links** — Helps detect a failure of a real interface and/or route summarization.

These tables are each described in the following sections.

List Topology Errors

The Route Recorder can detect 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 for presentation in the *List of Topology Errors* dialog box. Only those anomalies detected since the start of the last full exploration are shown. These anomalies can indicate router configuration errors that should be corrected. Each row of the table includes the time when the anomaly was detected, a description of the problem, and for topologies with multiple autonomous systems, the AS where the anomaly was detected. Clicking on the table row for an error message will highlight the associated router(s) and/or link(s), assuming that those objects are present in the topology at the time currently being displayed. The *History Navigator* window can be used to change the current time back to the time of the error message so that other tables can be used to diagnose the problem.

[Figure 62](#) shows an example of the *List of Topology Errors* dialog box.

Time	Message	AS
2006-06-12 12:00:21	Duplicate interface address 10.148.148.1 on 192.168.109.11 and 24.0.0.23	Lab58.EIGRP/AS1
2006-06-12 12:00:21	2 duplicate interface addresses 192.168.118.23 ... on 24.0.0.23 and 24.0.0.23	Lab58.EIGRP/AS1
2006-06-12 12:00:46	Duplicate interface address 10.115.115.1 on 198.99.99.1 and 24.0.0.23	Lab58.EIGRP/AS1

3 entries

Figure 62 List of Topology Errors Dialog Box

The following anomalies are included in the table:

- Interface mask length mismatch – The address mask length is not the same for the interfaces on the two ends of a link.
- Duplicate router ID – 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 – Since the router ID is normally derived from an interface address on the router, and interface addresses are normally unique to one router, it is an anomaly for the router ID on one router to be the same as an interface address on another router.
- Duplicate interface address – 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 – If an external prefix is advertised by a router but is unreachable from that router, this can indicate that the prefix is configured for redistribution but the metric has not been configured.
- Variance not supported by RAMS – Indicates that the router is configured for equal-cost multi-path routing with a variance value other than one. RAMS will only include the paths with the lowest metric.
- Router ID unroutable in this AS – 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 – 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 – The Route Recorder can discover a routing loop either 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 display 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 RAMS peer routers to that prefix.

List Inaccessible Routers

During the EIGRP topology exploration, RAMS attempts to establish a Telnet/CLI connection to each router to collect information about neighbors, interface metrics, and external prefix attachments. If the connection to a router fails, RAMS cannot include that router in the topology, nor can it learn about other routers connected beyond that router. This might not matter if the routers are, for example, lab routers connected as a stub network. However, if there is a path between two accessible routers that passes through an inaccessible router, RAMS will not be able to find that path. Therefore, it is important to fix router accessibility problems so that the RAMS topology is correct.

An example of the List of Inaccessible Routers dialog box is shown in [Figure 63](#). The table lists the routers that were inaccessible since the start of the most recent full topology exploration. For each inaccessible router, there is a column indicating the address of the last router on the path to the inaccessible one, and clicking on the entry in this column highlights that router on the topology map. Another column indicates the gateway (first-hop router) that RAMS used in attempting the connection, in case the solution would be to specify a different default gateway. The next column indicates a possible reason for failure to access the router:


- Authentication failure – This occurs if the router does not accept the login password or user name/password configured in RAMS for the AS.
- Invalid input (unauthorized command?) – This error occurs if the TACACS account used by RAMS 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?) – If too many other Telnet sessions are open at the time Route Explore attempts its connection such that no free virtual terminal is available on the router, the connection is refused. RAMS

attempts several connections with exponentially increasing delay. This error can also occur if the connection is blocked by a firewall or other device between RAMS and the intended destination router.

- Connection timeout – If the router is unreachable, for example, if the path to the router hits a black hole, then the connection will time out.
- Telnet open failed – Additional details are listed if provided by Telnet.
- CLI parsing error – The output of the commands issued to the router in a query was not formatted as expected by RAMS. Please report this problem to technical support.
- Problem in recorder – RAMS was unable to issue the query for some reason. Please report this problem to technical support.

The final column 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.



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 Inaccessible Routers Dialog Box

List Mismatched Distances

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

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

An example of the *List of Mismatched Distances* dialog box is shown in [Figure 64](#).

Time	Source	Destination	Router's Metric (bv+dly)	Model's Metric (bv+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 actual path	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+4294967295	1666560+4262462465	Prefix 192.168.133.0/30 not converged at 192.168.118.18	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 prefix	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 prefix	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 prefix	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 prefix	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 prefix	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 prefix	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 prefix	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 prefix	Lab58.EIGRP
2006-06-12 04:11:08	192.168.118.23	24.1.7.0/24	256000+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	256000+130810	0+4261412865	Model and router behavior don't match	Lab58.EIGRP

210 entries

Figure 64 List Mismatched Distances Dialog Box

There are several reasons why a mismatch occurs:

- Unreachable router hides actual path — Some mismatches are caused when the actual path goes through a router that RAMS cannot access. The actual links traversed and their metrics are therefore unknown.
- Different equal-cost path chosen by model — If there are multiple paths whose total costs are equal, but which have different bandwidth and delay components of the metric, then RAMS might choose a different path than the routers actually use. The algorithm of the router is not always deterministic.
- Prefix not converged — When RAMS traces the actual path taken by the routers, but finds that one of the routers has no route to the prefix in question, this 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 — This message indicates that RAMS 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 — A mismatch can occur if RAMS’s modeling of the recorders’ behavior is not exact. However, a mismatch can also occur when a router becomes “confused” and reports inconsistent metric information (perhaps due to a bug in the router software). Some router configuration changes, such as changing an access list (ACL) used in a route filter, do not take effect until the routing process is restarted. RAMS will see the new value, but the router will not be using it, causing a distance mismatch. Please report this message to technical support if it persists across multiple full explorations.
- RAMS failed to query — This message indicates a possible bug in RAMS. Please report this problem to technical support.

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


RAMS periodically re-explores the topology to make sure no changes have been missed due to transitions that do not result in an EIGRP update being sent to RAMS or due to limitations in tracking network dynamics. The period is set as part of recorder configuration and defaults to 8 hours. At the end of each periodic topology exploration, a message is added to the List of Mismatched Distances dialog box telling the number of links and prefix attachments that were corrected during the last periodic topology exploration. Ideally, these numbers should also 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, RAMS 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 RAMS will not report an EIGRP distance change. During the simulation, RAMS 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, then RAMS will be able to detect a failure of the real interface. If not, then RAMS can only detect the interface failure during the next periodic topology exploration (by default every 8 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 by RAMS.

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 has changed or additional peer routers have been added, click  **Reload** on the *Invisible EIGRP Interfaces* dialog box to re-run the simulation.

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

6 The History Navigator

One of the most powerful features of RAMS is its ability to display the detailed routing history of a network. The routing history is recorded by the recording component of the appliance, Route Recorder, which listens to the routing protocols and records all protocol events in a database.

Every ten minutes the Route Recorder saves a complete, time-stamped snapshot of the routing topology in the database. In between the snapshots, all routing announcements are recorded with timestamps. With this data, RAMS is able to display a precise routing map of the network at any point in time.


The *History Navigator* window displays statistical summaries of the recorded data in graphical format. The graphs show many vital network statistics versus time, including the number of events between snapshots and, at each snapshot, the number of routers, routing adjacencies, and prefixes.

The *History Navigator* window provides powerful data analysis functions. You can perform root cause analysis on event data, display the contents of the Routing Information Base (RIB) or visual representations of the RIB at any point in time, and perform a before-and-after comparison of the state of the network at two different points in time.

The *History Navigator* window can also display a list of the routing events that occurred during a specified time period. This is useful when diagnosing a network outage or performing forensic analysis after an outage.

About the History Navigator Window

The *History Navigator* window, in combination with the topology map, allows you to display information about your network in a wide variety of ways. You can choose to display detailed lists of events; 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, or view a real-time graph of events as they occur.

You can select **History Navigator** from the *Tools* menu, or click  on the toolbar, to display the *History Navigator* window.

Note If a database contains multiple protocols, the *History Navigator* window displays multiple tabs, one for each protocol. As in the example shown in [Figure 65](#), a database containing BGP, EIGRP, and OSPF protocols has tabs for BGP, EIGRP, and OSPF.

Initially, the *History Navigator* window contains a graph of recorded network routing events. An example of the window is shown in [Figure 65](#).

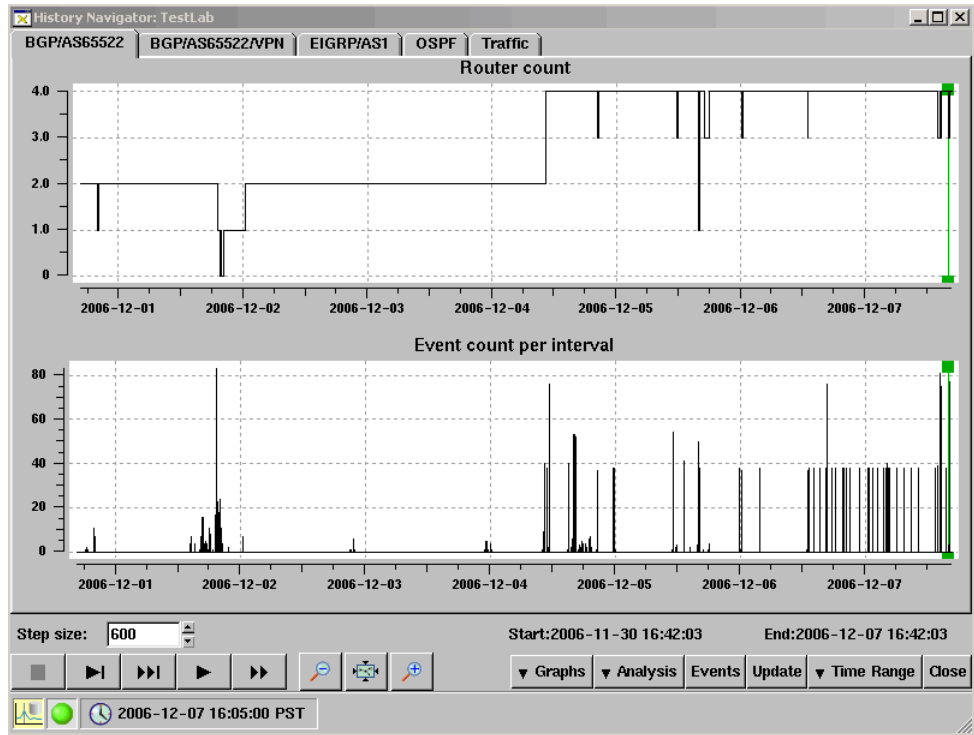


Figure 65 History Navigator Window in History Mode

If the topology map currently displayed by RAMS is an actively recording database, you can switch between History mode, Design mode and Online mode by clicking the mode icon located in the lower left-hand corner of the window. If the topology is not currently recording, the only modes available are Design mode and History mode.

When in Online mode, the History Navigator window displays routing events as they occur. Traffic events are delayed by 30 minutes in real time. Online mode is available when an active database is selected from the *Open Topology* dialog. While RAMS is displaying an active database, it continuously updates the routing topology map and the History Navigator window from the network routing announcements it receives.

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

History Navigator Controls

The *History Navigator* window controls, also shown in [Figure 65](#), allow you to navigate through the routing database and customize the presentation of data being displayed.

The elements on the window vary depending upon the current topology map mode. The following modes are available in the *History Navigator* window:



Online 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 in this mode.



History 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.



Design mode – A network icon indicates that planning features are enabled for the topology map.

You can switch between modes by clicking the mode icon. This has the effect of changing the Time Range shown on the graph from *Online* to *One Week* (see [Buttons](#) on page 231 for information about the values that can be set with the **Time Range** button).

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

Status Bar

The status bar at the bottom of the *History Navigator* window is the same as the status bar on the *Topology Map* window. See [Using the Status Bar](#) on page 165 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. The cursor indicates the currently displayed point in time within the routing topology history. There are several ways to move the cursor:

- You can use the mouse to drag the cursor to a different point on the time line. The topology map immediately displays the routing topology as it existed at that time. Note that traffic data is delayed by 30 minutes.
- You can right-click on a point in the time line. A pop-up appears asking if you want to move time to that point. Select **Yes**. The topology map immediately displays the routing topology as it existed at that time. If the time is within the last 30 minutes, traffic data is not displayed.
- You can move the cursor through time by stepping or animating (automated stepping) using the playback controls as described in [Playback Controls](#) on page 232. Any paths highlighted on the routing topology map will be recomputed and redisplayed at each step of the replay.

Note Because the routing topology database does not store nodes and links that are down, any objects that are down when the topology is first opened will not be 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, but colored 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 Options](#) on page 216 for information about node/link timeout intervals).

Buttons






The panel of buttons in the lower right-hand corner of the window access graphs, data analysis tools, and events lists, and allow you to specify the time range that the *History Navigator* window displays. From left to right, the buttons are:

- The **Graphs** button lets you select which graphs are displayed. See [Selecting History Graphs](#) on page 234 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*.
- The **Analysis** button offers a list of data analysis tools. See [Analyzing Historical Data](#) on page 235 for information about these tools.

- The **Events** button displays a detailed list of routing events. See [Understanding the Events List](#) on page 261 for information about the Events list.
- The **Update** button is only enabled 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.
- The **Time Range** button determines how much data is included in the *History Navigator* window. By default, the time range is set to *Online* when in Online mode and to *One Week* when in History mode. You can set the time range to one hour, one day, one week, or one month. Note that traffic data does not appear in Online mode and is delayed by 30 minutes in all other modes.
- The **Close** button closes the *History Navigator* window.

Playback Controls

The panel of buttons in the lower left-hand corner of the window control playback. From left to right, the buttons are:

-  The **Stop** button stops animated playback. The **Stop** button is enabled only in animated playback mode.
-  The **Step** button 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.
-  The **Fast Step** button 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.
-  The **Animate** button 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.
-  The **Fast Animate** button starts animated playback in fast mode, that is, 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.

Note 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 261, includes all routing changes.

Zoom Buttons

The panel of buttons toward the bottom center of the screen control zooming functions. From left to right, the buttons are:



The **Zoom In** button allows you to zoom into a subsection of the *History Navigator* graph.



The **Zoom Out** button allows you to zoom out of a subsection of the *History Navigator* graph.



The **Reset Zoom** button resets the view back to the standard viewer setting.

Zooming the Time Line

RAMS lets you 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. Zooming in the Y dimension may help you to see small changes in the number of objects in the Routers or Links graphs when the total number is large.

To zoom the time line, perform the following steps:

- 1 While holding the **Ctrl** key down, 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.

Selecting History Graphs

The primary feature of the History Navigator window is the Events graph that is shown in the default window. RAMS provides four additional graphs that may be used to display data. Together, these graphs offer a wide range of statistics about the state of the network.

- **Routers** – The Routers graph displays the number of physical entities in the network. For OSPF, this includes autonomous system (AS) Border Routers in other areas that are visible from the viewed area.
- **Routes** – The Routes graph 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.
- **Prefixes** – The Prefixes graph displays the number of prefixes available in the entire network.
- **Events** – The Events graph 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.
- **Links** – The Links graph 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 in the *History Navigator* window is a BGP AS.

To display any of these graphs, click **Graphs**, and then select the boxes beside the desired graphs.

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

You can configure RAMS to include any or all of the graphs in the default window. See [Using the Configuration Options Dialog Box](#) on page 212 for more information.

Analyzing Historical Data

RAMS provides powerful analysis tools through the Analysis button in the History Navigator window. These tools include the following:

- [Root Cause Analysis](#) on page 235
- [RIB Visualization](#) on page 242
- [RIB Browser](#) on page 247
- [RIB Browser Comparison](#) on page 252
- [Event Analysis](#) on page 256
- [Traffic Reports](#) on page 259
- [VPN Reports](#) on page 260

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.

Note Keep in mind that traffic data 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 5](#), “The Routing Topology Map.”

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

- 1 Click **Analysis**.
- 2 Select **Root Cause Analysis**.
- 3 Left-click in the graph just before the events occurred.
- 4 Left-click in the graph just after the events occurred.

If no significant incidents occurred during the time you specify, a dialog box appears indicating that the Root Cause Analysis algorithm did not find any major BGP problems. Adjusting the analysis options as described in [Chapter 5, “The Routing Topology Map,”](#) may increase the number of incidents the algorithm will find.

If incidents are found, the *Root Cause Analysis Results* dialog box opens. [Figure 66](#) shows an example of the dialog box.

Root Cause Analysis Results: 2006-04-05 11:21:38 to 2006-04-14 03:58:55	
Description	View Details
Recording restarts First event: 2006-04-05 11:21:38 Last event: 2006-04-13 21:50:43 Time elapsed: 202 h 29 m 5 s	Animation 296 Events 26 Prefixes
Router 24.0.0.3 has re-established peerings First event: 2006-04-05 11:25:12 Last event: 2006-04-14 03:58:38 Time elapsed: 208 h 33 m 26 s	Animation 73481 Events 18 Prefixes
Peering to 24.0.0.5 has flapped First event: 2006-04-05 11:25:20 Last event: 2006-04-14 03:58:55 Time elapsed: 208 h 33 m 35 s	Animation 14664 Events 1 Prefixes
Peering to 24.0.0.12 has flapped First event: 2006-04-06 15:43:15 Last event: 2006-04-14 03:58:36 Time elapsed: 180 h 15 m 21 s	Animation 40485 Events 8 Prefixes
Peering to 24.0.0.16 has flapped First event: 2006-04-12 13:32:16 Last event: 2006-04-14 03:57:44 Time elapsed: 38 h 25 m 28 s	Animation 4640 Events 1 Prefixes
5 entries	

Figure 66 Root Cause Analysis Results Dialog Box

The events that occurred during the time period you specified in Steps 3 and 4 are analyzed and correlated into groups. All of the BGP routing messages that apply to related events are summarized in the dialog box. 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, RAMS 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 RAMS establishes a BGP peering, the peer router sends RAMS all of its BGP routes. At the end of this sequence, RAMS writes a synchronized event.
- **Peering lost** — When a peering is lost, the connection between RAMS and the peer router has closed. In previous product releases, RAMS wrote peering withdrawal events before a peer was closed, and included any associated Announce and Withdraw messages. In the 5.20 version of the product, 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 RAMS loses its connection to a peer router, it may detect that other routers have also lost contact with the peer. As a result, RAMS 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** — RAMS 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. RAMS appears at the left side of the tree. The time range of the animation is indicated in the RAMS rectangle. To the right of the RAMS rectangle are its BGP peers; to their right are its BGP next hops; to their right are the autonomous systems (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 67](#) shows an example of the 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.

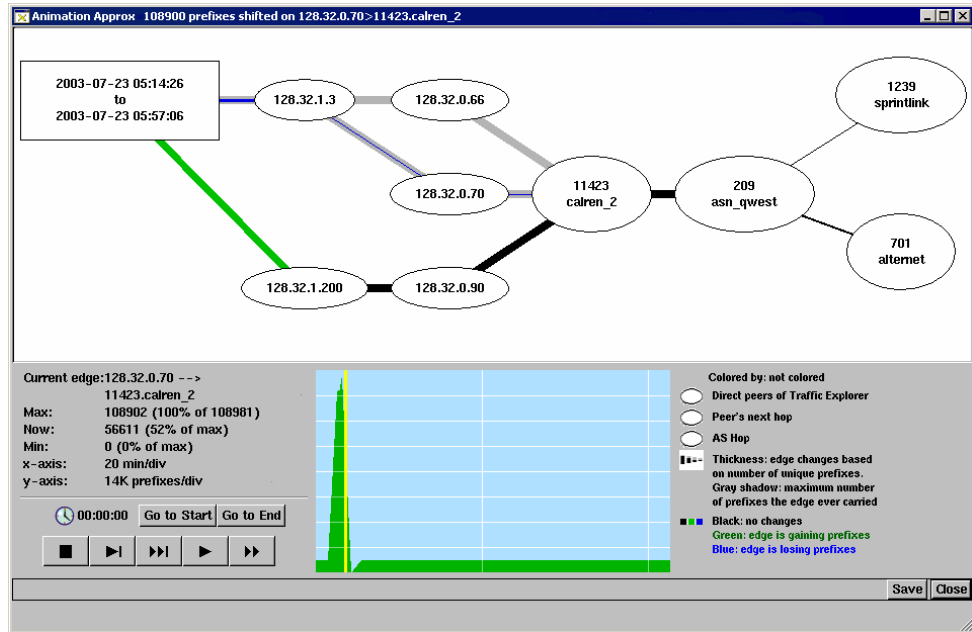


Figure 67 Root Cause Analysis Animation Window

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 232).
- **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 a playback mode:

- 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 twice the interval of step mode.
- 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.

- 2 Click the corresponding playback button to begin the animation.

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

To replay an animation, perform the following steps:

- 1 Click the **Go to Start** button to move the yellow cursor back to the left side of the graph.
- 2 Click a playback button to begin the animation.

To animate a different edge, perform the following steps:

- 1 In the visualization pane of the *Animation* window, click another edge to select it.

The graph pane of the *Animation* window now displays information about the selected edge, and the graph displays the change in prefixes on the selected edge.

- 2 Click a playback button to begin the animation.

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 (Scalable Vector Graphic, file extension `.svg`) on the RAMS 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 has a free SVG plug-in which can be downloaded from <http://www.adobe.com/svg/viewer/install/main.html>). HP has used the Adobe plug-in with a variety of browsers on Linux, Mac OS X and Microsoft Windows platforms.

Alternatively, the Apache Batik project has a standalone SVG viewer called *squiggle* which can be downloaded 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 a browser and navigate to the RAMS *Home* page on a Route Recorder. The *Login* page appears.
- 2 Type your user name and password. You do not need Administrator privilege to access saved animations.
- 3 Click the **Reports Portal** link on the top navigation bar.
- 4 On the left navigation bar, click the **BGP Animations** link. 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.
- 5 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 Information Base (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 Move the History Navigator cursor to the time desired.
- 2 Click **Analysis**.
- 3 Select **RIB Visualization**.

The *RIB Visualization* window opens, as shown in [Figure 68](#).

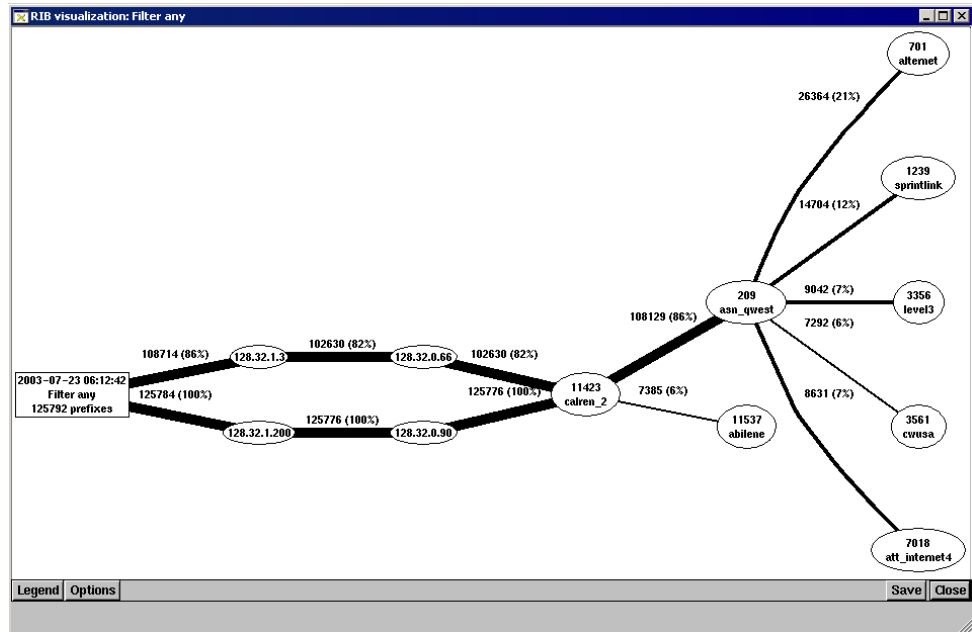


Figure 68 RIB Visualization Window

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 these trees into a single tree. RAMS appears at the left side of the tree. The RAMS 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. In the example shown in [Figure 68](#), “Filter any” indicates that all routes were included. You can create visualizations filtered to include only a subset of the routes from the *RIB Browser* as described on [page 250](#).

To the right of the RAMS rectangle are its BGP Peers, followed by its BGP next hops; to their right are the autonomous systems that the next hops serve; to the right of the autonomous systems are any downstream autonomous systems; to the right of the downstream autonomous systems 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 69](#) shows a visualization without the **Never Show an Edge to a Prefix** option selected. [Figure 70](#) shows the same visualization with **Always** option selected.




Figure 69 RIB Visualization without Show an Edge to a Prefix Option




Figure 70 RIB Visualization with Show an Edge to a Prefix Option

To change RIB visualization thresholds, perform the following steps:

- 1 In the *RIB Visualization* window, click **Options**.
The Visualization options are displayed in the left pane of the window.
- 2 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.

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

- 3 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 upper right corner of the *Visualization* window. The visualization is saved in SVG format (Scalable Vector Graphic, file extension `.svg`) on the RAMS 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 241 for available plug-ins.

To view a saved visualization, perform the following steps:

- 1 Open a browser and navigate to the RAMS *Home* page on a Route Recorder.
- 2 The *Login* page appears.
- 3 Type your user name and password. You do not need Administrator privilege to access saved visualizations.
- 4 Click the **Reports Portal** link on the top navigation bar.
- 5 On the left navigation bar, click the **BGP Animations** link.
- 6 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.
- 7 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 Routing Information Base (RIB) Browser to display the following types of information:

- For IGP, links and prefixes that are currently down.
- 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, additional fields are available if the OSI network is present.

To open the *RIB Browser* dialog box from the *History Navigator* window, click **Analysis** and select **RIB Browser**.


IGP Protocols

For IGP protocols, the *RIB Browser* dialog box displays a list of links and prefixes that are currently down at the bottom of the *History Navigator* window. [Figure 71](#) shows an example of the *RIB Browser* dialog box with IGP link data.

On the left side of the *RIB Browser* dialog box, click **Down Links** or **Down Prefixes** to view the list of links or prefixes, respectively, that are currently down. **Overloaded Routers** displays all the routers that currently have their overload bit set. The **Router** column identifies the router that corresponds to each link or prefix, while the **Number of Down Links** column displays the number of times the link or prefix that are down. To view the list of links or prefixes as a bar chart, click  **View as Bar Chart** in the upper right corner of the window.

Router/Net	IP Address	Number of Down Links
--	192.168.3.26	2
--	192.168.0.30	2
--	192.168.3.56	2
--	192.168.3.61	2
--	192.168.0.72	2
--	192.168.3.102	2
--	192.168.3.116	2
--	192.168.3.120	2
--	192.168.3.131	2
PD-REX-RECORDER	192.168.3.144	2
--	192.168.3.172	2
--	192.168.0.179	2
--	192.168.1.182	2
--	192.168.3.189	2
--	192.168.3.192	2
--	192.168.1.244	2
--	10.8.254.1	1
--	10.8.254.2	1
--	192.168.3.60	1
--	192.168.0.115	1
--	192.168.3.133	1
--	192.168.0.138	1
--	192.168.0.171	1

Figure 71 RIB Browser Dialog Box For IGP

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 times the link or prefix 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 detailed information about all the links or prefixes associated with that router.
- **Filter Analysis** – Displays a new *RIB Browser* window with data for the selected router only.

BGP Protocol

Figure 72 shows an example of the *RIB Browser* dialog box with BGP peer data.



Router	Peer BGP ID	Route Count
saratoga-border	47.150.22.242	131882
inglewood-border	96.18.255.3	121664
tucson-border	96.18.255.35	105427
oldsmobile-border	96.18.255.33	49328
--	96.18.255.5	37772
prefontaine-pf1	96.18.255.49	21195
--	96.18.255.24	9138
sebastopol-pf1	96.18.255.50	8152
--	96.18.255.12	4243
--	96.18.255.14	3113
sebastopol-border	96.18.255.36	2857
--	96.18.255.6	425
--	96.18.255.25	308
--	96.18.255.15	139
--	96.18.255.37	105
--	96.18.255.38	84
--	96.18.255.2	78
--	96.18.255.8	76
--	96.18.255.16	68

Figure 72 RIB Browser Dialog Box for BGP

For BGP protocol, the RIB Browser dialog box displays distributions of the advertised prefixes their attributes. A tree structure on the left side of the dialog box presents the following attribute options:

- **Peer** – For each peer, displays the number of routes advertised by that peer.
- **Nexthop** – For each Nexthop, displays the number of routes that list that Nexthop router among their attributes.
- **Originator** – For each Originator, displays the number of routes that list that Originator 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.
- **Communities** – For each community, displays the number of routes that list that community among their attributes.
- **Neighbor AS** – For each Neighbor AS, displays the number of routes that list that neighbor AS among their attributes.
- **2nd Hop AS** – For each 2nd Hop AS, displays the number of routes that list that 2ndHopAS among their attributes.
- **Origin AS** – For each Origin AS, displays the number of routes that list that origin AS among their attributes.

- **Any AS** – For each AS, displays the number of routes that list that AS among their attributes.
- **AS Peers** – For each pair of AS peers, displays the number of routes that list that peer-pairing among their attributes.
- **Prefixes** – For each prefix, 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 242).
- **Filter Analysis** – Displays a new *RIB Browser* window with data for the selected entity only.

VPN Protocol

For VPN protocol, the *RIB Browser* dialog box displays distributions of the advertised prefixes their attributes. A tree structure on the left side of the dialog box 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 Route Targets 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

As in IGP protocols, similar columns appear in the *RIB Before-N-After Comparison* dialog box for OSI IS-IS.

Note If OSI IS-IS is not detected by the appliance, this window will not display.

For OSI IS-IS protocol, the RIB browser dialog box displays a tree structure on the left side of the dialog box and presents the following attribute options:

- **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.

Use the **Filter By** drop-down list to select the filter parameters and the **Analyze Matching** or **Analyze Excluding** buttons to list only those links that match the filter criteria, or exclude those events that match the filter criteria, respectively.

For more information on using filters, see [Using Filters](#) on page 280.

RIB Browser 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 73 provides an example of one such instance.

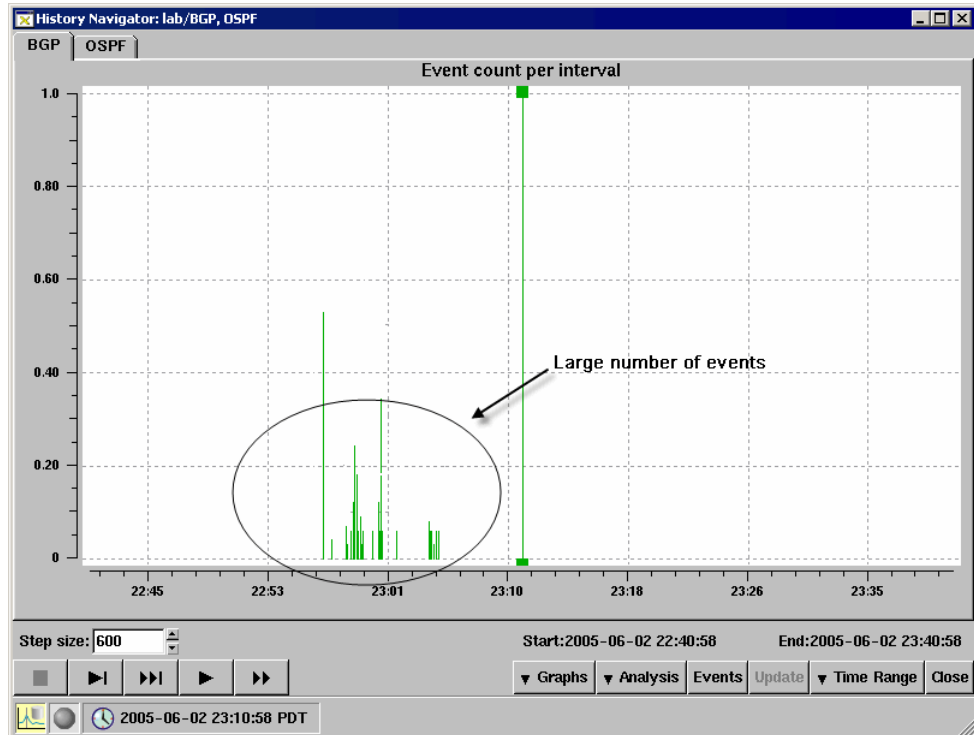


Figure 73 Large Number of Events in Short Space of Time

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 Browser Comparison, perform the following steps:

- 1 Click **Analysis**.
- 2 Select **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* dialog box opens. This dialog box is very similar to the *RIB Browser* dialog box.

IGP Protocols

For IGP protocols, the *RIB Before-N-After Comparison* dialog box includes columns for the link and prefix counts before and after the events, and also a column for the difference between the two. Figure 74 show an example of the *RIB Comparison* dialog box with IGP link data.

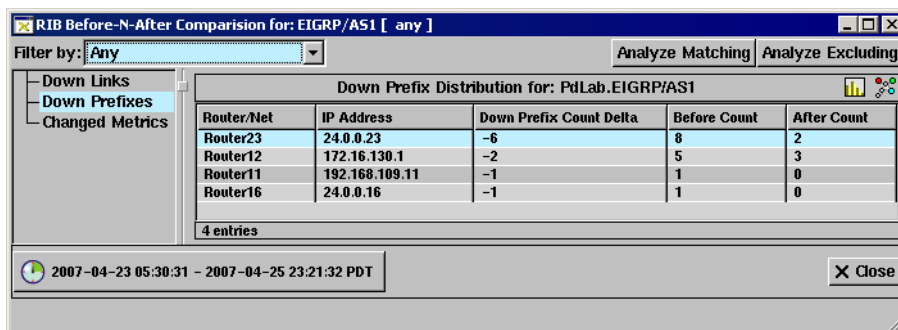




Figure 74 RIB Comparison Dialog Box for IGP

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 what states the bits were at the beginning and end time frames. **Changed Metrics** shows all the links whose metric values have changed between the time frames, along with what the values were at in the beginning and end time frames. 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 261.
- **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* dialog box as those that appear in the *RIB Browser* dialog box. In addition, *Before*, *After*, and *Delta* columns display the route count before and after the specified events, and the difference between the two counts.

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**.



Router	Peer BGP ID	Route Count
saratoga-border	47.150.22.242	131882
inglewood-border	96.18.255.3	121664
tuscon-border	96.18.255.35	105427
oldsmobile-border	96.18.255.33	49328
--	96.18.255.5	37772
prefontaine-pf1	96.18.255.49	21195
--	96.18.255.24	9138
sebastopol-pf1	96.18.255.50	8152
--	96.18.255.12	4243
--	96.18.255.14	3113
sebastopol-border	96.18.255.36	2857
--	96.18.255.6	425
--	96.18.255.25	308
--	96.18.255.15	139
--	96.18.255.37	105
--	96.18.255.38	84

Figure 75 RIB Comparison Dialog Box for BGP

In [Figure 75](#), the *RIB Comparison* dialog box is shown with the pop-up menu that opens when you right-click an entry in the list.

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* dialog box. See [Figure 76](#) for an example of the display.
- **Filter Analysis** – Displays a new *RIB Browser* window with data for the selected entity only.

Router Name	Peer BGP ID	Prefix	Before Attributes	After Attributes
--	24.0.0.3	192.168.120.0/24	--	Next Hop: 192.168.129.1 AS Path: 65533 11111 (IGP) Local-Pref: 100 MED: 222 Next Hop: 192.168.129.1
--	24.0.0.3	25.0.0.1/32	--	AS Path: (INCOMPLETE) Local-Pref: 100 MED: 0 Next Hop: 192.168.129.1
--	24.0.0.3	192.168.122.90/32	--	AS Path: (INCOMPLETE) Local-Pref: 100 MED: 0 Next Hop: 192.168.103.11
--	24.0.0.3	27.27.27.0/29	--	AS Path: 65533 11111 (INCOMPLETE) Local-Pref: 100 MED: 222 Next Hop: 192.168.129.1
--	24.0.0.3	192.168.103.0/24	--	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11

17 entries

Figure 76 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 in [VPN Protocol](#) on page 250.

The functions for *RIB Browser Before-N-After Comparison* for VPN protocol are the same as for BGP protocol (described in [BGP Protocols](#) on page 254) 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* dialog box for OSI IS-IS.

Note If OSI IS-IS is not detected by the appliance, this window will not display.

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. RAMS 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.


To analyze a series of events, perform the following steps:

- 1 Click **Analysis** to open a menu of analysis functions.
- 2 Select **Event Analysis** from the menu.
- 3 Click in the Events graph just before the events occurred.
- 4 Click in the Events graph just after the events occurred.

The *Event Analysis* dialog box appears.

IGP Protocols

For IGP protocols, this dialog box displays the number of routing events that occurred in the specified time interval for each involved initiator, router, link, or prefix. [Figure 77](#) shows an example of this dialog box.

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 262 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.

Router/Net	IP Address	Event Count
--	24.0.0.11	407
192.168.0.0/24	192.168.0.2 DR	97
--	24.0.0.2	34
--	24.0.0.5	28
--	24.0.0.12	27
--	24.0.0.16	24
--	192.168.0.72	21
192.168.103.0/24	192.168.103.11 DR	12
--	24.0.0.3	10
--	192.168.0.38	10
192.168.101.0/24	192.168.101.2 DR	10
--	192.168.0.98	9
--	192.168.0.29	2
--	192.168.0.113	2

Figure 77 Event Analysis Dialog Box for IGP

BGP Protocol

For BGP, the *Event Analysis* dialog box 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 261 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 235 for information about the controls that appear in this window.
- **Filter Analysis** – Displays a window with event data for the selected entity only.

Figure 78 shows an example of the *Event Analysis* dialog box for BGP, and includes the pop-up menu that appears when you right-click an entry in the list.

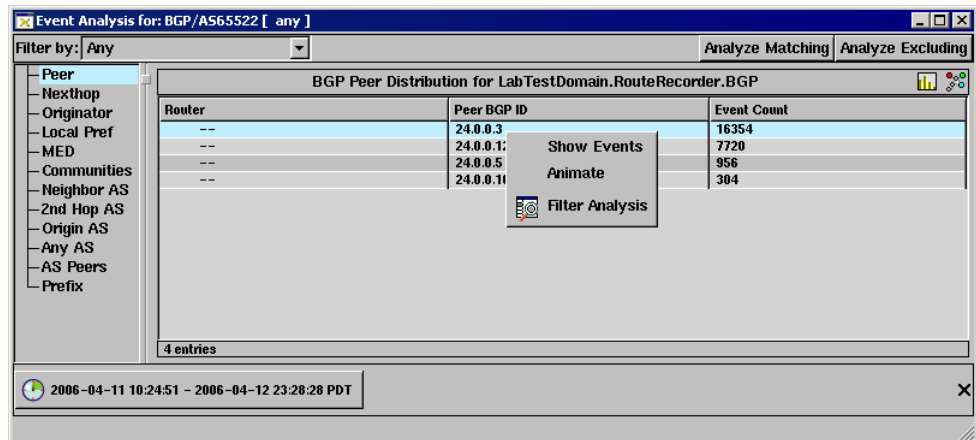


Figure 78 Event Analysis Dialog Box for BGP

VPN Protocol

In addition to the options found for BGP, *Event Analysis* dialog box 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 in [VPN Protocol](#) on page 250.

The functionality for Event Analysis for VPN is the same as the Event Analysis for BGP, described in [BGP Protocol](#) on page 257.

OSI IS-IS Protocol

For OSI IS-IS protocol, the *Events Analysis* dialog box 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. [Figure 79](#) shows an example of this dialog box.

Note If OSI IS-IS is not detected by the appliance, 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 261 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 235 for information about the controls that appear in this window.
- **Filter Analysis** – Displays a window with event data for the selected entity only.

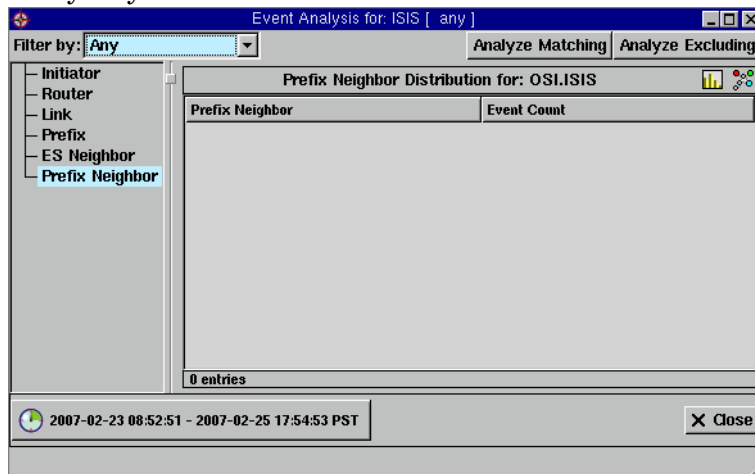


Figure 79 Event Analysis Dialog Box for OSI IS-IS

Traffic Reports

See [Chapter 12](#), “Traffic Reports.”

VPN Reports

See Chapter 10, “VPN Configuration and Reports.”

Understanding the Events List

After the general nature of a routing problem has been 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 selected 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 From the *History Navigator* window, click on **Events**
- 2 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.
- 3 Move the mouse cursor to the ending time and left-click again to mark that time.

The *All Events* window opens displaying details of the events that occurred within the selected time period, as shown in [Figure 80](#).

Note Another way to access the *All Events* window is to select **Online Events** from the *Tools* menu.



Caution 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.

Filtered events: LabTestDomain/RouteRecorder/OSPF [(--or Event change prefix Event add router)]					
Filter by: Event Operation		Options		Show	Hide
Time	Router	Operation	Neighbor/ Prefix	Attributes	Area or AS
2006-04-04 17:36:25.1488	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 2 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:36:25.1830	24.0.0.11	Change Prefix	10.88.88.0/24	Metric: 11113 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:36:25.1830	24.0.0.11	Change Prefix	10.79.79.0/24	Metric: 11113 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:36:30.0798	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 1 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:36:30.1150	24.0.0.11	Change Prefix	10.88.88.0/24	Metric: 11112 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:36:30.1150	24.0.0.11	Change Prefix	10.79.79.0/24	Metric: 11112 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:46:04.6575	192.168.0.171	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-04 17:46:04.6888	10.88.88.2	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-04 17:46:13.4894	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 2 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:46:13.5212	24.0.0.11	Change Prefix	10.88.88.0/24	Metric: 11113 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:46:13.5212	24.0.0.11	Change Prefix	10.79.79.0/24	Metric: 11113 (Area External)	LabTestDomain.RouteRecorder
2006-04-04 17:46:18.8811	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 1 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:43:25.7056	192.168.0.72	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-05 09:55:47.5291	192.168.0.2 DR	Add Neighbor	192.168.0.127	Metric: 0	LabTestDomain.RouteRecorder
2006-04-05 09:44:51.3011	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 2 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:44:56.5149	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 1 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:52:34.5317	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 2 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:52:40.3927	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 1 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:55:42.4624	192.168.127.2	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-05 09:55:47.4943	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 2 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 09:55:47.5291	192.168.0.127	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-05 09:55:52.8060	24.0.0.11	Change Prefix	192.168.116.0/24	Metric: 1 (Area External)	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	24.0.0.2	Add Router		Type: ASBR	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.101.2 DR	Add Router		Type: LAN Pseudo-Node	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.103.11 DR	Add Router		Type: LAN Pseudo-Node	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	24.0.0.5	Add Router		Type: Area BR, ASBR	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	24.0.0.11	Add Router		Type: Area BR, ASBR	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	24.0.0.12	Add Router		Type: ASBR	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.0.2 DR	Add Router		Type: LAN Pseudo-Node	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.0.179	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.0.176	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder
2006-04-05 10:01:27.2618	192.168.0.171	Add Router		Type: Internal Router	LabTestDomain.RouteRecorder


574 entries 2006-04-04 11:56:17 - 2006-04-07 02:18:57 PDT








Figure 80 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 267).

The following controls are arranged from left to right across the bottom bar of the *Events* window:

-  **Time Range** button – Opens the *Select Time Range* dialog box, 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** button – Refreshes the Events list with events that occurred within the past 10 minutes, with the exception of traffic data, which is delayed by 30 minutes. This button is disabled when the *History Navigator* window is in History mode.
-  **Show Current Event**,  **Stop Execution**,  **Execute One Event**, and  **Start Execution** buttons – See [Moving Time and Executing Events](#) on page 269 for information about using these buttons.
-  **Clear** button – Clears all events from the *Events* window. This button is only functional in Online mode.
-  **Close** button – 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 Link State Advertisement (LSA) packets or OSI IS-IS Link State Packets (LSP). EIGRP is a distance-vector protocol, so it does not communicate link-state changes directly. However, RAMS 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 and withdrawals. In addition to the events indicating network state changes, entries are inserted in to the events list when the peering between RAMS 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.

Name	Description
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 or Withdraw of a prefix.
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.
Area or AS	The OSPF area, OSI IS-IS level, or EIGRP or BGP AS where the event took place.

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 80](#), 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 Area Border Router. 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 Designated Router (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.

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” (Designated Router) following a router address or DNS name indicate events pertaining to the pseudonode representing 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, Area BR (Area Border Router), ASBR (Autonomous System Border Router), 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 been timed out without a refresh, or “Cause: Premature” when the router advertises a graceful withdrawal (for example, on shutdown). For protocol details, see [Appendix B, “Protocol Compliance.”](#)

- **OSI IS-IS:** Since OSI IS-IS is a link-state protocol like 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”. RAMS 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 (Traffic Engineering). For protocol details, see [Appendix B, “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 RAMS 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

RAMS 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. RAMS cannot infer the link change in this case.
- Not on old path; new path broken – If the internal routing topology model provided by RAMS 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, RAMS will not be able to infer a link failure that partitions some nodes from the viewpoint of RAMS.
- Query failed – A query by RAMS 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 RAMS 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 waiting for the nonresponding router to communicate, and the **Attribute** column reports the statistics from the *show ip eigrp neighbor* command on the last responding router on the route about the nonresponding router. 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 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 B](#), “Protocol Compliance.”

In addition to the protocol-specific events outlined above, there are Add Peering and Drop Peering events that indicate when RAMS 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 80](#), 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.

Some parameters require that you type a value in a text box (for example, if you filter by router, you must type 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 280 explains how to combine filter parameters using the **Expressions** option on the filter drop-down list.


See [Using Filters](#) on page 280 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 on the object of interest to display the node information panel or link information panel ([Node Information Panel: Protocol Tab](#) on page 177 and [Link Information Panel: Protocol Tab](#) on page 180, 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 2 and 3 in [Understanding the Events List](#) on page 261. You can adjust the time range as needed.

To adjust the start and end of the time range, perform the following steps:

- 1 Click  located in the lower left-hand corner of the *Events* window to display the *Set Time Range* dialog box.
- 2 You can then adjust the time range in any of the following ways:
 - Type new values into the *From* and *To* fields, or adjust the values with the up and down triangle buttons.

- Select one of the predefined time ranges from the *Select Time Range* menu: one hour, day, week or month. The time range will be centered around the currently displayed point in time.
- Select **Recent** from the *Select Time Range* menu. This displays a drop-down list of recently used time ranges from which you can select.
- Select **All** from the *Select Time Range* menu to include all events recorded in the database.

Note Traffic data is delayed by 30 minutes.

- 3 Click **OK** to accept the adjusted time range.






Caution 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.

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 80](#), the next event is the *Add Router* event at 11:02:48.

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 (X), 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 Browser Comparison](#) on page 252 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 use RAMS to help to narrow down the event churn to its root cause.

In the example shown in [Figure 81](#), 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.

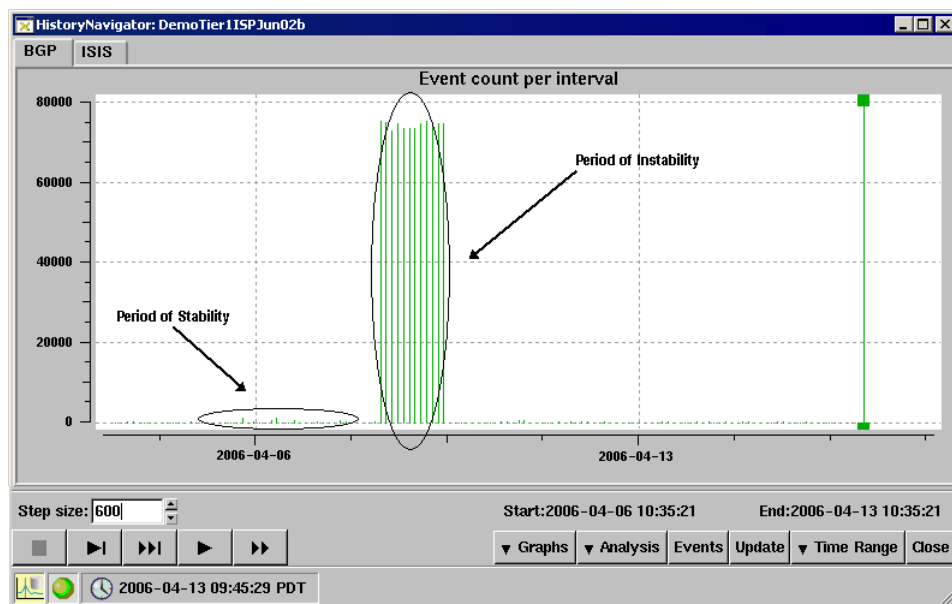


Figure 81 Stability and Instability

To perform an events analysis, perform the following steps:

- 1 Click **Analysis** in the *History Navigator* window.
- 2 Select **Event Analysis** from the pop-up menu that appears.
- 3 Mark the start and end time for the analysis using the blue cross-hairs.

The *Event Analysis* window appears, as shown in [Figure 82](#).

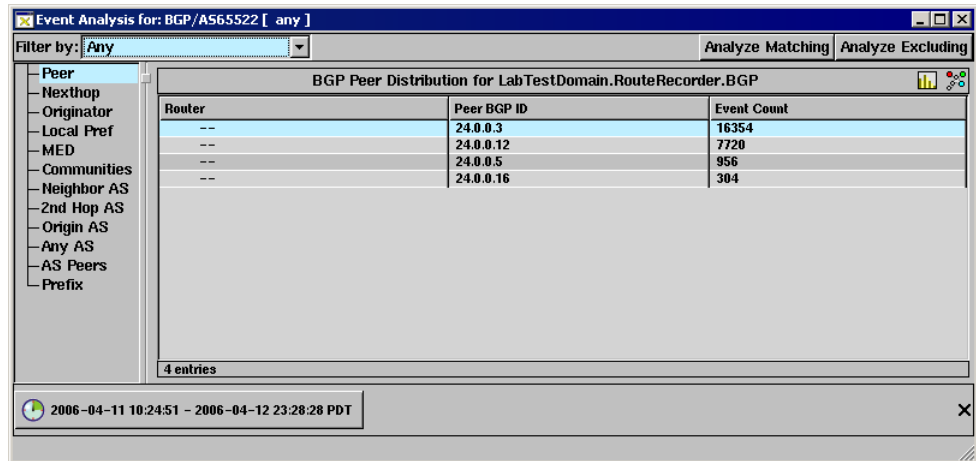


Figure 82 Event Analysis Window

The *Event Analysis* table can be filtered, sorted by column heading, or viewed as a bar chart.

When the **MED** option is selected, a small number of MEDs (three in this example) have a large number of events associated with them, as show in [Figure 83](#). This could represent a MED oscillation.

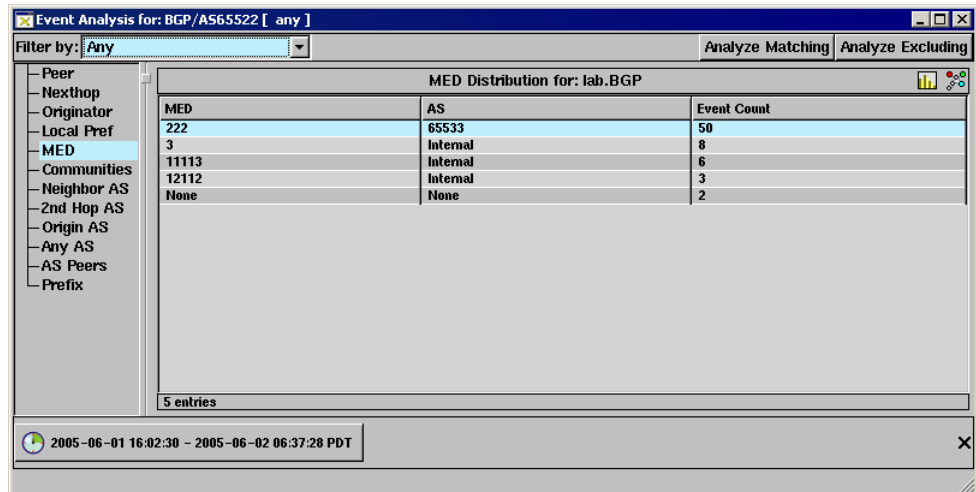



Figure 83 MEDs

To identify the prefixes affected by this possible MED oscillation, select the **Prefix** tab as shown in [Figure 84](#).



Prefix	Event Count
10.0.2.0/24	3
10.0.30.0/30	3
111.111.111.0/30	3
10.20.20.0/24	2
10.166.166.0/24	2
16.16.16.0/24	2
20.20.20.0/24	2
24.0.0.8/32	2
24.0.0.22/32	2
24.0.0.23/32	2
24.0.0.26/32	2
24.0.0.28/32	2
25.0.0.8/32	2
192.168.114.0/30	2
192.168.115.0/30	2
192.168.122.0/24	2
192.168.124.0/24	2
192.168.126.0/24	2
192.168.127.0/24	2
35 entries	

Figure 84 Prefix Details

A single prefix has a huge number of 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 Right-click the prefix in question.
- 2 Click **Filter Analysis** in the pop-up window that appears.

[Figure 85](#) displays the results of the drill-down filter analysis.

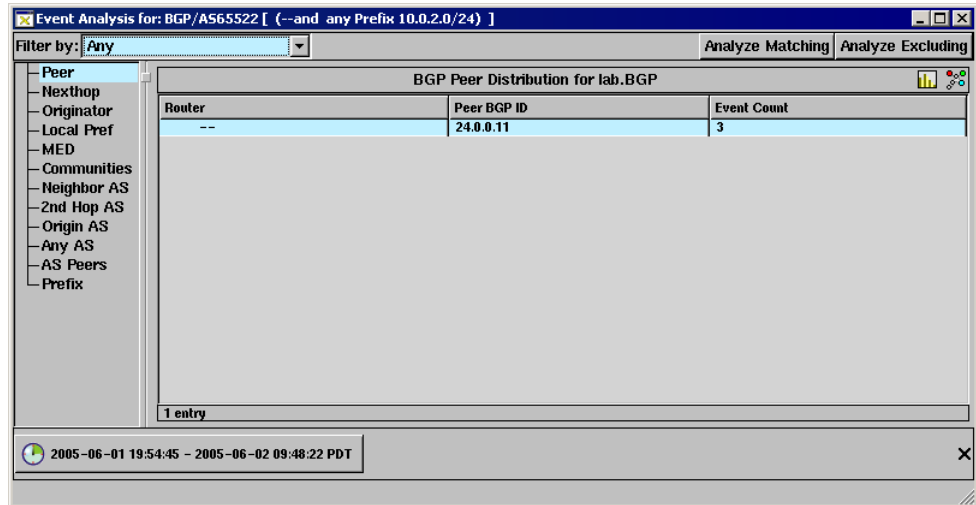


Figure 85 Filtered Event Distribution

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 RAMS RIB analysis identified the possibility of a MED oscillation, and the RAMS 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 In the *History Navigator* window, click **Analyze**, and then select **Event Analysis** from the pop-up menu that appears.
- 2 Select the desired start and end times with the blue cross-hairs.
- 3 To view the events associated with an individual table entry, right-click the entry, and then select **Show Events** in the pop-up menu that appears, as shown in [Figure 86](#).

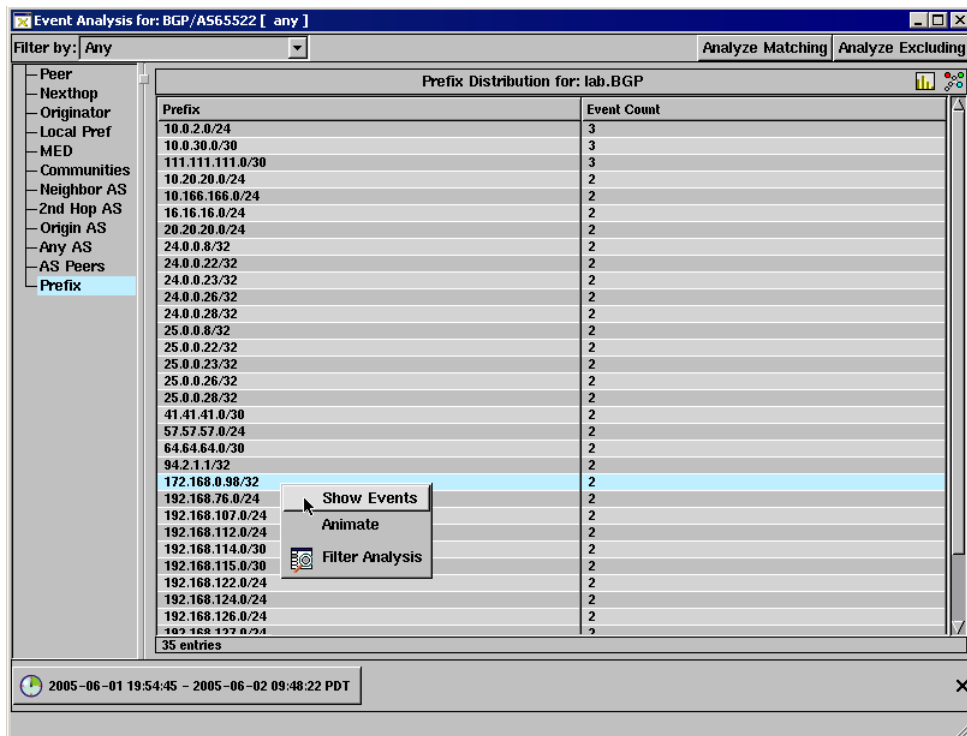


Figure 86 Show Details Button

The *Filtered Events* window appears, as shown in [Figure 87](#). This window lists all of the events associated with the selected table entry.

Filtered events: LabTestDomain/RouteRecorder/[BGP AS65522] [(--and Router 24.0.0.3 (--and any Prefix 192.168.103.0/24))]

Filter by: Any Show Hide

Time	Router	Operation	Neighbor/Prefix	Attributes	Area or AS
2006-04-05 11:21:49.503844	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-05 11:25:40.258417	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-05 11:34:54.091977	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-05 19:12:23.480202	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-05 19:13:39.313873	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-05 19:15:38.430817	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-06 15:41:44.330740	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP
2006-04-06 15:42:40.102240	24.0.0.3	Announce	192.168.103.0/24	AS Path: (IGP) Local-Pref: 100 MED: 2 Next Hop: 192.168.110.11	LabTestDomain.RouteRecorder.BGP

8 entries 2006-04-04 11:56:17 - 2006-04-06 18:37:14 PDT

Figure 87 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. This procedure is based upon the History Navigator example shown in [Figure 88](#).

To perform a Root Cause Analysis, perform the following steps:

- 1 In the *History Navigator* window, select **Root Cause Analysis** from the *Analysis* menu.
- 2 Right-click just before the period of instability and again just after the period of instability to specify the time-frame of the analysis.

The *Root Cause Analysis* window appears, as shown in [Figure 88](#). The window indicates that several prefixes are flapping, at least one of which is generating a very large number of events.

- 3 Click **Animation** for that prefix.

An *Animation* window appears. Playing the animation shows the network disturbance as it occurs. See [Root Cause Analysis](#) on page 235 for information about playback controls, the graph, and the procedures for manipulating the animation display.

Root Cause Analysis Results: 2006-12-01 20:37:32 to 2006-12-01 21:38:15	
Description	View Details
Peering to 24.0.0.5 has flapped	Animation
First event: 2006-12-01 20:37:32	6 Events
Last event: 2006-12-01 21:38:15	0 Prefixes
Time elapsed: 1 h 0 m 43 s	
1 entry	

Figure 88 Root Cause Analysis Window

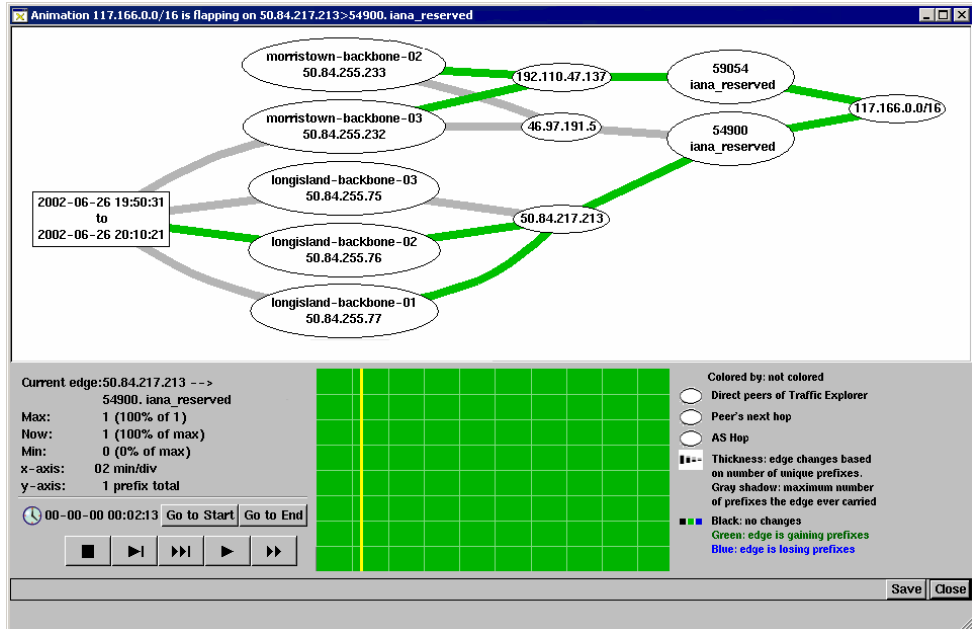


Figure 89 Root Cause Analysis Animation Window

The thickness of the gray lines indicate the maximum number of prefixes that were ever carried on each edge. Green lines indicate edges that are gaining prefixes. It is apparent that several of the backbone nodes experienced a high degree of instability, causing the routes to prefix 117.166.0.0/16 to flap. This information helps you to narrow down the possible sources of the instability to just a few nodes.

Correlating Time Series Data

RAMS extends the power of routing navigation 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 RAMS 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 into RAMS and performing a time correlation with routing events can help diagnose the root cause of the anomaly.

RAMS can import data in MRTG ASCII .log file format, Round Robin Database (RRDtool) .rrd binary 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 RAMS appliance. The administrator must first enable the FTP server on the RAMS appliance. See [Chapter 3, “Administration.”](#)

To upload files to the RAMS appliance, perform the following steps:

- 1 FTP to the RAMS appliance using the IP address of the appliance.
- 2 Log in as user `rexftp`. The administrator sets the password.

- 3 Change directory to `pub`.
- 4 Transfer one or more files to the RAMS appliance.

To display a time series file (all formats), perform the following steps:

- 1 Select **Correlate Time Series** from the *Tools* menu.
- 2 Select the file to be displayed, and then click **Select**.

The time series graph is displayed in a separate window.

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.

Note 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 feature is provided on several tables, including the RIB Browser and Events List, to allow you to focus in on items of interest. [Figure 90](#) 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 filter feature offers three levels of filtering:

- 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 283 for examples of the parameter syntax as illustrated using filter expressions.

Note With a simple filter, you type 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 91](#) 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.




Figure 90 Filter By Drop Down List

In many cases, the built-in **Filter by** selections, such as the ones shown in [Figure 90](#), provide sufficient flexibility in filtering. For example, the Router filter accepts a list of router addresses or names when several specific routers are of interest.

To set up a simple filter, perform the following steps:

- 1 Select an operator from the **Filter by** drop-down list. For example, assume you want to see only events reported by a node whose IP address is 192.168.167.166. Select **Router** in the **Filter by** drop-down list.

A text box appears to the right of the **Filter by** parameter.

- 2 Type the address of the node in the text box as shown in [Figure 92](#).



Figure 91 Parameter Text Box

- 3 Click **Show** to list only items that match the address, or click **Hide** to list only items that do not match the address.

To set up an advanced filter, perform the following steps:

- 1 Select **Advanced** from the **Filter by** drop-down list.

The *Advanced Filter* window appears, as shown in [Figure 92](#).




Figure 92 Advanced Filter Window

- 2 In the *Advanced Filter* window, select a filter operator from the **Filter by** drop-down list.

Some operators require that you type the parameter in a text box, others let you choose among items in a list.

- 3 Specify the appropriate parameter, either by typing it into a text box or choosing it from a list, depending upon the operator.
- 4 To exclude matching items, click **Not** for that operator.

- 5 To add another operator to the filter, click **Add** in the upper left corner of the window. Then repeat steps 2 through 4 to define the parameters for the new operator.
- 6 After you have added 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.
- 7 Click **Show** to list all events that match the filter, or click **Hide** to list only events that do not match the filter.

The filter is translated internally into a filter expression combining the filter operators with the parameters that you specified.

Multiple levels of advanced filters can be combined to construct any logical AND-OR expression you desire.

To enter a filter expression manually, perform the following steps:

- 1 Select **Expression** from the **Filter by** drop-down list.
- 2 Type the desired expression in the adjacent text box as shown in [Figure 93](#).

See [Expression Syntax](#) on page 282 for information about the syntax used to enter expressions, and [Expression Definitions](#) on page 283 for a complete list of operators and examples showing their use.




Figure 93 Expression Text Box

- 3 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 be composed of multiple terms (operators and parameters) to form a complicated filter.

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.

Instead, if the entries with LocalPref 100 were not interesting but other values were, then the expression could be modified as follows:

```
(peer 192.0.2.1 or peer 192.0.2.2) and not localPref 100
```

Expression Definitions

This section defines the filter operators used in RAMS, 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
```

bgpState <state>

Matches BGP routes in the specified state with respect to the baseline. The states are as follows:

<code>bgpState Dead</code>	(not in baseline and dead)
<code>bgpState Down</code>	(not in baseline and down)
<code>bgpState Up</code>	(not in baseline but up)

bgpState Down/B	(in baseline but down)
bgpState Up/B	(in baseline and up)

`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.

`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 rixpeering
eventType drop rixpeering
eventType change rixpeering
eventType drop *
eventType add *
```

```
eventType change *
eventType * router
eventType * neighbor
eventType * prefix
eventType * rexperring
```

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)
eventType withdraw	(route withdrawal)

`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 above for `router`. For example,

```
externaloriginator 192.168.0.36
```

`igpPrefixType <type>`

Matches the specified IGP prefix type. The available prefix types include the following:

- OSPF: Internal, AreaExt, ASEExt_Type2, or ASEExt_Type1
- IS-IS: Internal, ASEExt_Type2, ASEExt_Type1, or TE
- EIGRP: Internal, ASEExt_Type2, ASEExt_Type1, Loopback, Dialup, AutoSum, ManualSum, StaticInt, StaticExt, NotFoundInt, NotFoundExt

For example, in an OSPF network:

```
igpPrefixType AreaExt
```

`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:

```
and med 987 neighAS 208
```

`neighbor <router>`

Matches a neighbor router in an events list using any of the forms of router identification described above for `router`. For example,

```
neighbor labnet-gw
```

`neighborAS <asn>`

`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
```

`peer <router>`

Matches a specific BGP peer address using any of the applicable forms of router identification described above for `router`. For example,

```
peer 192.0.2.3
```

`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
```

`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. 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
```

`secondHopAS <asn>`

Matches the second hop AS, meaning, the one after the neighbor AS in an AS path. For example,

```
secondHopAS 288
```

`staticNextHopType <type>`

Matches the specified nexthop type for a static route. The available types are: Network, Interface, Gateway, Default. For example,

```
staticNextHopType Network
```

`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 10, “VPN Configuration and Reports”](#) 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
```

7 Network Planning

Network planning is an essential part of optimizing performance. Analyzing the needs of the network to identify and eliminate inevitable hot spots is key to staying ahead of potential failures in service.

To this end, RAMS provides a view of network demands derived from actual measurements, rather than synthetic models of network activity or limited link utilization measurements. Planning capability is expanded with features like failure analysis, the ability to change link metrics and the ability to move prefixes among border routers. By combining routing and traffic data, traffic trends and their impact on the available capacity and reliability of the network are easy to see across the entire network.

On top of this network-wide view of traffic, the planning tools provided by RAMS allow you to add, delete and change routers, links, flows and prefixes using Design mode. Then, you can view the edits, import and export data, or undo changes, as well as view generated reports in the *Planning Reports* window to provide a comparison of network activity before and after the applied changes took effect to analyze the differences in traffic measurements that resulted from simulated network modifications. Using *Planning Reports*, you can analyze how traffic changes not only across the entire network, but on any given node, interface, exit router, next-hop autonomous system (AS) or destination AS.

This chapter introduces the network planning process and describes the steps involved in making changes to the topology map and analyzing those changes using *Planning Reports*.

Network Planning and Analysis Tools

Display the planning toolbar by selecting **Design mode** from the *Topology*

menu, or clicking  on the toolbar in RAMS, as shown in [Figure 94](#).

Starting Design mode also enables the planning options listed under the *Edit* menu. Use the planning toolbar and the *Edit* menu to perform network planning tasks described in this section.

The Planning Toolbar

The buttons on the planning toolbar, which is docked on the right side of the Topology Map window by default, provide access to the planning functions provided by RAMS. 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 these buttons. The functions on the toolbar can also be found under the *Edit* menu.

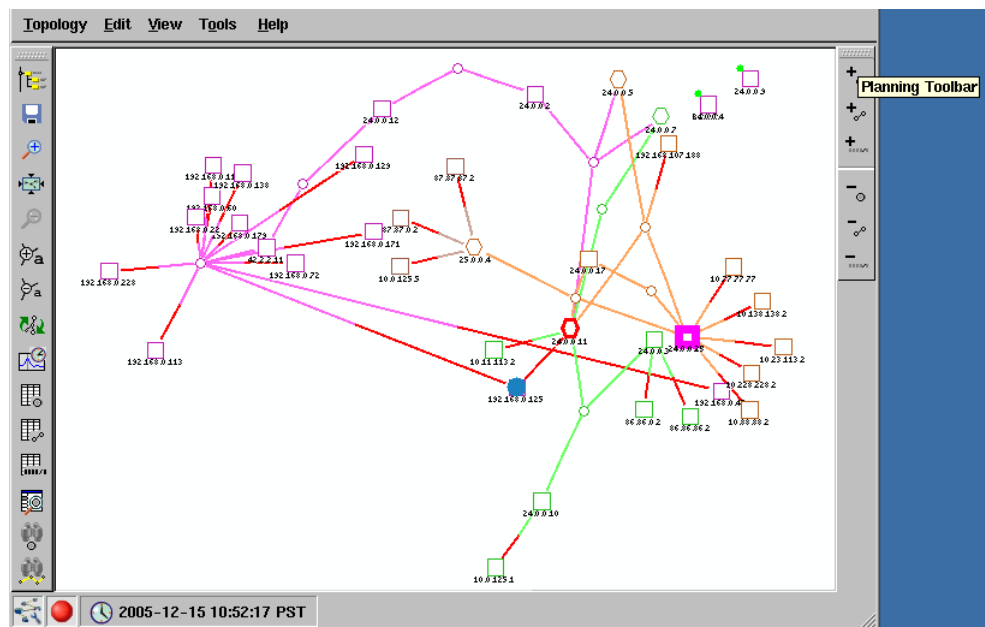


Figure 94 Planning Toolbar

From top to bottom, the following buttons are located on the toolbar:



Add Node — Places a new node on the topology map. See [Add a Node](#) on page 297 for a description of this task.



Add Peering — Creates a peering relationship between two nodes on the topology map. See [Create a Peering](#) on page 298 for a description of this task.



Add Prefix — Applies prefixes to a router on the topology map. For BGP routers, add prefixes manually or select a filtering method for the prefix. See [Add a Prefix](#) on page 301 for more information.



Add OSI Prefix — Applies OSI prefixes to an OSI IS-IS router on the topology map. See [OSI IS-IS Router](#) on page 306 for more information.



Add Traffic Flow — Creates one or more flows from one node to another on the topology map. See [Add Traffic Flow](#) on page 308 for a description of this task.



Change Prefix — Removes or changes the attributes of a prefix on a particular node on the topology map. Change multiple prefixes at once by selecting more than one prefix from the table. See [Change a Prefix](#) on page 308 for a description of this task.



Down Node — Changes the state of a node from “Up” to “Down,” simulating what would happen if the selected router should fail. See [Bring Down a Node](#) on page 314 for a description of this task.



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. See [Bring Down Peerings](#) on page 315 for a description of this task.



Down Prefix — Changes the state of one or more prefixes from “Up” to “Down” on a particular router, simulating what would happen if the selected prefix(es) should fail. See [Bring Down a Prefix](#) on page 316 for a description of this task.



Down OSI Prefix — Changes the state of one or more OSI prefixes from “Up” to “Down” on an OSI IS-IS router, simulating what would happen if the selected OSI prefixes should fail.

The Edit Menu

In addition to the functions described in [The Planning Toolbar](#) on page 292, the *Edit* menu provides access to the following traffic-related windows:

- The **Reports** menu item launches the *Planning Reports* window, where the edits made to the topology map are listed. Analyze, undo, import, and export these edits in the *Planning Reports* window. See [Analyzing Topology Edits](#) on page 328 for more information.
- The **Edit Traffic Flows** menu item launches the *Edit Flows* window, where you can make changes to flows, including adding, moving, and deleting a flow. See [Edit Traffic Flows](#) on page 310 for more information.

The Planning Reports Window

To access the *Planning Reports* window, click the *Edit* menu in RAMS, then choose **Reports** as shown in [Figure 95](#).



Figure 95 Reports Menu Option

Navigate through the *Planning Reports* window using the buttons in the upper right corner and by selecting a link from the list in the left pane, as shown in [Figure 97](#).








Figure 96 Planning Reports Window

Use the buttons in the upper right corner of the *Planning Reports* window to perform the following functions:

-  **Analyze Edits** — Allows you to see the effects the edits had on link traffic and utilization.
-  **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 Edits**— Allows you to import edits from either the clipboard or from a database. See [Importing and Exporting Planning Data](#) on page 340 for more information.
-  **Export Edits**— 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. See [Importing and Exporting Planning Data](#) on page 340 for more information.

Analyze the effects of edits to the topology map using the links in the left pane of the *Planning Reports* window. Click the **Analyze Edits** button, then click a link to display a corresponding table of data in the window. The following options are available for each table:

- View the data as a table, pie chart or bar chart. To toggle among these three views, click the **View As** button in the upper right corner of the window.
- Limit the list of results displayed on the table or chart. Use the **Filter by** drop-down list and click **Show** or **Hide**.
- Sort the table by clicking the label at the top of each column. For example, to sort by how much available capacity changed, click **Available Change** at the top of that column.

The Select **Traffic Group** tab, located at the bottom right portion of the Planning Reports window, allows you to choose which traffic group information you want to display. For more information about this feature, see [Select Traffic Groups Tab](#) on page 429.

After you click **Analyze**, the following links are available on the left pane of the *Planning Reports* window:

- **List of Edits** — Displays all the changes made to the topology map.
- **Link Utilization** — Displays the levels of utilization for each link in the network before and after edits were made. The *Link Utilization* table also shows the amount of change resulting from the edits. For more information, see [Changes in Link Utilization](#) on page 329.
- **Link Traffic** — Displays the amount traffic on each link in the network before and after edits were made. The *Link Traffic* table also shows the amount of change resulting from the edits. For more information, see [Changes in Link Traffic](#) on page 330.
- **Available Capacity** — Displays the amount of capacity available on each link in the network before and after edits were made. The *Available Capacity* table also shows the amount of change resulting from the edits. For more information, see [Changes in Available Capacity](#) on page 331.
- **Neighbor** — Displays the amount of traffic on routes between neighboring autonomous systems. For more information, see [Changes in Neighbor AS Traffic](#) on page 332.
- **Destination** — Displays the amount of traffic on routes to a destination AS. For more information, see [Changes in Destination AS Traffic](#) on page 333.
- **Exit Router Traffic** — Displays the amount of traffic on routes that exit from various points in the network. For more information, see [Changes in Exit Router Traffic](#) on page 333.
- **Community Traffic** — Displays the amount of traffic on routes between members of a BGP community. For more information, see [Changes in Community Traffic](#) on page 334.


Editing Network Topology

Working in Design mode, you can simulate changes to the network by editing the topology map in RAMS. Simulating the addition of a node, for example, allows you to see realistic effects of the new router on network traffic activity. These effects are analyzed in *Planning Reports*, as described in [Analyzing Topology Edits](#) on page 328.

To start planning changes to the network, first open a topology edit by performing the following steps:

- 1 From the *Topology* menu in RAMS, choose **Open Topology ...**.
- 2 Select a topology from the list that appears in the *Open Topology* dialog box.
- 3 Click **OK**.

Now start Design mode to access network planning features. There are two ways to enable this mode:

- From the *Topology* menu, choose **Design mode**, or
- Click  in the bottom left corner of RAMS.

In Design mode, the planning toolbar appears, which is described in [The Planning Toolbar](#) on page 292. In addition, the *Edit* menu is enabled, as described in [The Edit Menu](#) on page 294. Use the planning toolbar or the *Edit* menu to access the functions described in this section. After making an edit, you can view and undo the edit in the *Planning Reports* window (see [View and Remove Edits](#) on page 320).

Add a Node

Simulate the effect of adding one or more nodes to the network by using the Add Node function in Design mode. When you add a node, you must specify the protocol and properties of that node. Then, you can create peering relationships with other nodes on the network as described in [Create a Peering](#) on page 298. In the case of BGP routers, the node is automatically peered with RAMS.

To add a node, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Add Node**.

- 2 Click on the topology map to specify an insertion point for the node. You can move the node later if necessary.

The *Add Node* dialog box appears with a tab for each protocol in the network, as shown in [Figure 97](#).




Figure 97 Add Node Dialog Box

- 3 Choose a protocol for the node by clicking the appropriate tab.
- 4 The network you are currently editing appears in the first text box. 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 network specified in the text box.
- 5 Specify the router ID or System ID (for OSI IS-IS).
- 6 Optional: Type a name for the node in the *Router Name* text box.
- 7 Click **Apply** to save the node or **Cancel** to dismiss the dialog box without adding a node.

Create a Peering

Create a peering between two nodes to simulate the effect this relationship would have on the network.

To add a peering, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Add Peering**.
- 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 peering.

The *Add Peering* dialog box appears with a tab for each protocol in the network.

- 4 Choose a protocol for the peering by clicking the appropriate tab:
 - BGP (see [BGP Routers](#) on page 299)
 - IGP (see [IGP Routers](#) on page 300)

BGP Routers

If you chose the **BGP** tab in the *Add Peering* dialog box as described in [Create a Peering](#) on page 298, the BGP panel appears in the dialog box as shown in [Figure 98](#).

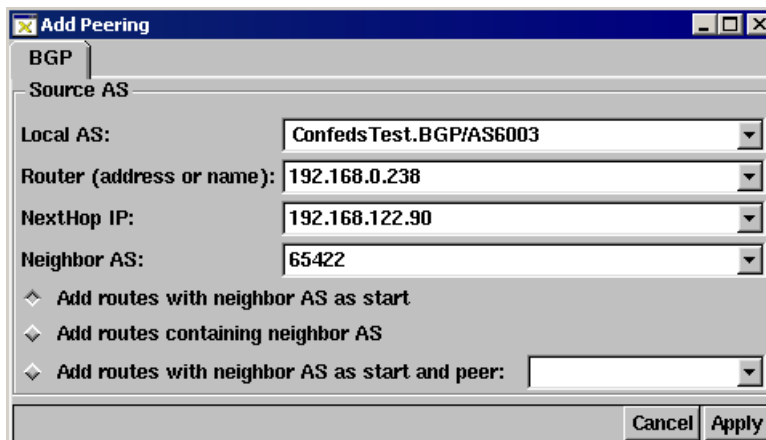


Figure 98 Add BGP Peering Dialog Box

To create a peering between two BGP routers, perform the following steps:

- 1 The autonomous system you are currently editing appears 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.
- 2 In the *Router* text box, specify the IP address or name of the source node in the peering.
- 3 In the *NextHop IP* text box, type or choose the next hop IP address of the source node.
- 4 The *Neighbor AS* text box is populated with the Neighbor AS number you are creating a peering for. You may edit this box by either entering other neighboring AS numbers, or selecting an AS number from the drop-down menu.

- 5 Select one of the following options:
 - Add routes with neighbor AS as start.
 - Add routes containing neighbor AS.
 - Add routes with neighbor AS as start and peer.
- 6 Specify information for the destination AS in the peering relationship following Steps 1 through 5.
- 7 Click **Apply** to save the peering or **Cancel** to dismiss the dialog box without adding a peering.
- 8 Click **Close (X)** to dismiss the dialog box.

IGP Routers

If you chose the **OSPF** tab in the *Add Peering* dialog box as described in [Create a Peering](#) on page 298, the OSPF panel appears in the dialog box as shown in [Figure 99](#).

OSPF	
Network:	RouteRecorder.OSPF/0.0.0.1
Src Router: (name or address)	24.0.0.12
Dst Router: (name or address)	192.168.0.138
Src Interface Address:	255.255.255.255/32
Dst Interface Address:	255.255.255.255/32
Bandwidth/Capacity (Kbps):	1
Metric:	1
<input type="button" value="Cancel"/> <input type="button" value="Apply"/>	

Figure 99 Add IGP Peering Dialog Box

To create a peering between IGP routers, perform the following steps:

- 1 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.
- 2 The *Src Router* text box is populated with the IP address or name of the source node in the peering. Edit this text box if necessary.

- 3 The *Dst Router* text box is populated with the IP address or name of the destination node in the peering. Edit this text box if necessary.
- 4 In the *Src Interface Address* text box, a default IP address appears. Replace the default value with the address of the source interface. This address can be followed by a mask length, such as 192.168.1.101/24.
- 5 In the *Dst Interface Address* text box, a default IP address appears. Replace the default value with the address of the destination interface. This address can be followed by a mask length, such as 192.168.1.101/24.
- 6 In the *Bandwidth / Capacity* text box, type the amount of available bandwidth allocated for this peering.
- 7 For EIGRP, in the *Delay* text box, enter the time required to move traffic from the source to destination (in microseconds).
- 8 For OSPF and IS-IS, in the *Metric* text box, type 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.
- 9 Click **Apply** to save the peering or **Cancel** to dismiss the dialog box without adding a peering.
- 10 Click **Close (X)** to dismiss the dialog box.

Add a Prefix

Simulate the effect of adding one or more prefixes to the topology by using the Add Prefix function in Design mode. When you add a prefix, its attributes can be specified manually or using a filter. Each of these methods are described in this section.

To begin adding a prefix, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Add Prefix**.
- 2 Click a node on the topology map to specify the router that will advertise the prefix. The prefix can be moved later if necessary.

The **Add Prefix** dialog box appears with a tab for each protocol in the network.
- 3 Choose a protocol for the prefix by clicking the appropriate tab:

- BGP (see [BGP Routers](#) on page 302)
- IGP (see [IGP Routers](#) on page 304)

BGP Routers

If you chose the **BGP** tab in the *Add Prefix* dialog box as described above, specify either of two methods:

- Type the attributes of the prefix manually as described on [page 302](#).
- Use one or more filtering methods as described on [page 303](#).

To manually add a prefix, perform the following steps:

- 1 Click the **Manually** option at the top of the **Add Prefix** dialog box as shown in [Figure 100](#).

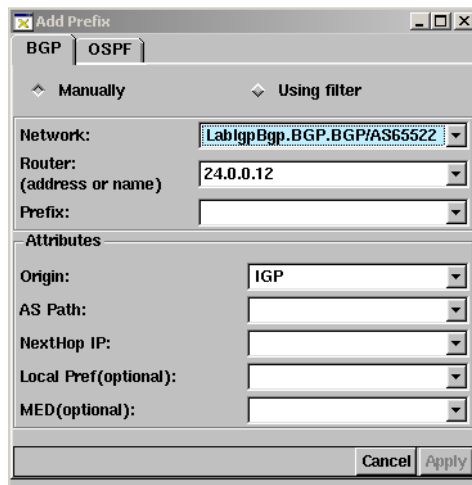


Figure 100 Add BGP Prefix Manually Dialog Box

- 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, type the IP address or name of the router that advertises the prefix.
- 4 In the *Prefix* text box, type the address of the new prefix.

- 5 Set the attributes of the prefix by specifying the Origin (which needs to be populated with one of three strings: IGP, BGP, or Incomplete, AS path name, the next hop IP address, the local preference (optional), MED value (optional), and Originator ID (this attribute displays only if the originating node is BGP).
- 6 Click **Apply** to save the prefix or **Cancel** to dismiss the dialog box without adding a prefix.

To add a prefix using a filter, perform the following steps:

- 1 Click the **Using a Filter** option at the top of the *Add Prefix* dialog box as shown in [Figure 101](#).

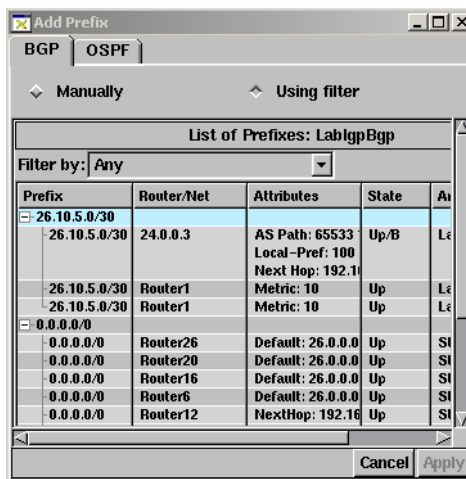


Figure 101 Add BGP Prefix Using Filtering Dialog Box

- 2 Use the **Filter by** drop-down list to choose which prefixes to show or hide in the table.

The **Filter by** feature offers three levels of filtering:

- 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 280 in [Chapter 6, “The History Navigator”](#) for examples of the parameter syntax using filter expressions described in [Expression Syntax](#) on page 282.

Note With a simple filter, you type 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.

- 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.

See [Using Filters](#) on page 280 in [Chapter 6, “The History Navigator”](#) for more detailed information.

- 3 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.
- 4 Click on a prefix in the table to highlight it.
- 5 Choose a filtering method for the highlighted prefix from the **Attributes** drop-down list and type a corresponding value in the text box to the right.
- 6 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, type 99 in the text box and choose **Set** from the *Operation* list.

- 7 Click **More** to apply additional filtering methods to the prefix or **Remove** to remove the last filtering method applied to the prefix.
- 8 Click **Apply** to save the changes or **Cancel** to dismiss the dialog box without adding a prefix.
- 9 Click **Close (X)** to close the *Add Prefix* dialog box.

IGP Routers

If you chose the **OSPF** tab in the *Add Prefix* dialog box as described in [Add a Prefix](#) on page 301, the following panel appears in the dialog box:

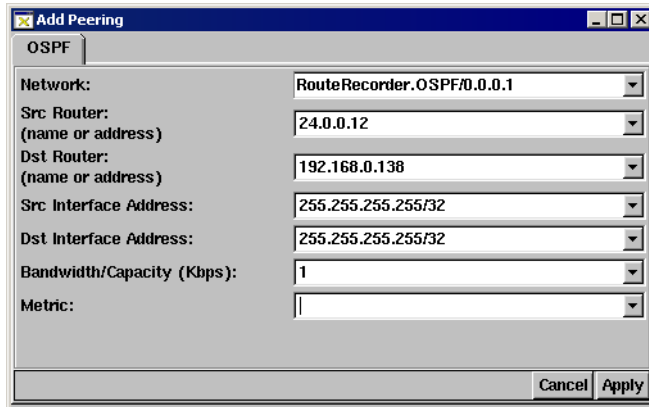


Figure 102 Add IGP Prefix Tab

To add a prefix to an IGP router, perform the following steps:

- 1 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 prefix will be added to the network specified in the text box.
- 2 In the *Prefix* text box, type the address of the new prefix.
- 3 The *Router* text box is populated with the IP address or name of the node that advertises the prefix. Edit this text box if necessary.
- 4 For EIGRP, in the *Bandwidth* text box, type the amount of bandwidth allocated for this peering.
- 5 For EIGRP, in the *Delay* text box, enter the time required to move traffic from the source to destination (in microseconds).
- 6 For OSPF and IS-IS, in the *Metric* text box, type the metric value for the peering.

Metric values help traffic determine the best path to take through the network.

- 7 From the **Metric Type** drop-down list, choose the type of prefix to add:
 - **Internal:** This metric type applies to area internal prefixes. Use this metric type if you are adding a prefix corresponding to a LAN or a subnet.

- **AS External Comparable:** This metric type is present for OSPF and IS-IS protocols, and applies to external routes that are redistributed into either OSPF or IS-IS areas. The metric of an AS External Comparable prefix can be directly compared with the metric of an Internal prefix.
- **AS External:** This metric type applies to external routes that are redistributed into EIGRP, OSPF, or IS-IS areas. For OSPF and IS-IS, the metric of an Internal prefix cannot be directly compared with the metric of an AS External prefix. For EIGRP, the metric of an Internal prefix is comparable to the metric of an AS External prefix.

Metrics and metric types are used to determine the order of route preference. The rules of preference for IS-IS route resolution are listed in RFC 1195 and RFC 2966. For OSPF route resolution rules, refer to RFC 2328.

- 8 For IS-IS, indicate if the prefix should be redistributed from Level 2 (L2) to all Level 1 (L1) areas. Toggle between **Yes** and **No** by selecting the check box.

Yes adds the prefix to all L2L1 routers. Distributing the prefix across the domain increases the granularity of routing information, which can improve the quality of the resulting routes.

- 9 Click **Apply** to save the prefix or **Cancel** to dismiss the dialog box without adding a prefix to the router.
- 10 Click **Close (X)** to dismiss the dialog box.

OSI IS-IS Router

To add a prefix to an OSI IS-IS router, perform the following steps:

- 1 Follow the steps for [Add a Prefix](#) on page 301.
- 2 Click on an OSI IS-IS router on the topology map.

The *Add Osi Prefix* dialog box appears as shown in [Figure 103](#).

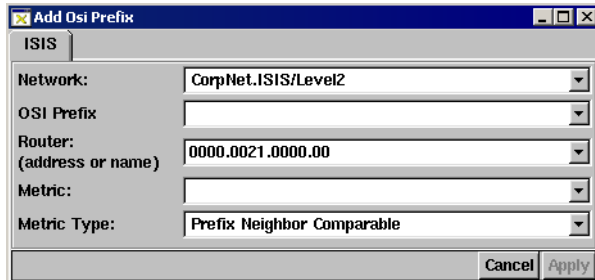


Figure 103 Add OSI Prefix

- 3 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.
- 4 Type the address of the new prefix in the *OSI Prefix* text box. The OSI prefix will be added to the network specified in the text box.
- 5 The *Router* text box is populated with the system ID or name of the node that advertises the OSI prefix. Edit this text box if necessary.
- 6 Type the metric value for the peering in the *Metric* text box.
Metric values help traffic determine the best path to take through the network.
- 7 From the **Metric Type** drop-down list, choose the type of OSI prefix to add:
 - **ES Neighbor:** Use this metric type if you are adding an ES Neighbor for Layer 1.
 - **Prefix Neighbor Comparable:** This metric type is present for Layer 2/backbone area and applies to external routes that are redistributed into an IS-IS areas. The metric of a Prefix Neighbor Comparable prefix can be directly compared with the metric of an Internal prefix.
 - **Prefix Neighbor:** This metric type applies to external routes that are redistributed into IS-IS areas. For OSI IS-IS, the metric of an Internal prefix cannot be directly compared with the metric of a Prefix Neighbor prefix.
- 8 Click **Apply** to save the OSI prefix or **Cancel** to dismiss the dialog box.

Change a Prefix

Change an existing BGP prefix using many of the same steps described in [Add a Prefix](#) on page 301. To change an IGP or EIGRP prefix, use the *Set Metrics* dialog box, as described in [Set or Change Link Metrics and Bandwidth](#) on page 319.

To change a BGP prefix, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Change BGP Prefix**.
- 2 Click the router on the topology map that advertises the prefix to change.
The *Change Prefix* dialog box appears.
- 3 Use the **Filter By** drop-down list to show or hide of prefixes in the table as described in Step 2 of [Add a Prefix](#) on page 301.
- 4 Click on a prefix in the table to highlight it.
- 5 Modify the attributes of the highlighted prefix as described in Step 5 of [Add a Prefix](#) on page 301.
- 6 Use the **Operation** drop-down list to set, append or prepend the modification to an existing prefix as described in Step 6 of [Add a Prefix](#) on page 301.
- 7 Click **More** to create further modifications to prefix attributes and the **Remove** button to create fewer modifications to prefix attributes.
- 8 Click **Apply** to save the changes or **Cancel** to dismiss the dialog box without adding a prefix.
- 9 Click **Close (X)** to close the *Change Prefix* dialog box.

Add Traffic Flow

Add single or multiple traffic flows to the topology map using the Add Traffic Flow function. When adding a traffic flow, specify the source and destination prefixes as well as 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* dialog box as described in [Edit Traffic Flows](#) on page 310.

To add a traffic flow, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Add Traffic Flow**.

You can also add traffic flow in the *Edit Flows* dialog box using the **Add** button described in [Edit Traffic Flows](#) on page 310.

- 2 Click a node on the topology map to specify where you'll add traffic flow.

The *Add Traffic Flow* dialog box appears with tabs for single and multiple flows, as shown in [Figure 104](#).

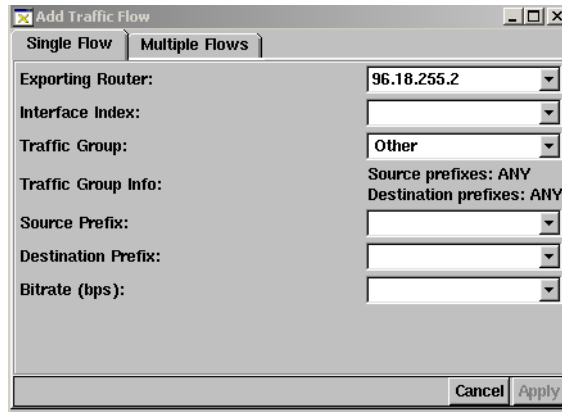


Figure 104 Add Traffic Flow Dialog Box

- 3 In the *Exporting Router* text box, the address of the router chosen on the topology map appears. 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.
- 4 In the *Interface Index* text box, type the value of the interface index to associate with the traffic flow. You can type the interface index of an existing exporting router, or you can enter an arbitrary value.

The interface index determines how the new traffic flow is grouped in the *Flow Exporters* table, which is described in [Chapter 12, “Traffic Reports.”](#)

- 5 In the *Traffic Group* text box, select the traffic group the flow belongs to.
- 6 In the *Source Prefix* text box, type the address of the prefix where the new traffic flow originates.
- 7 In the *Destination Prefix* text box, type the address of the destination prefix for the new traffic flow.
- 8 In the *Bandwidth* text box, type the bitrate for the new traffic flow.
- 9 Click the **Multiple Flows** tab to add more than one traffic flow at a time, and type the values described in steps 3-8.

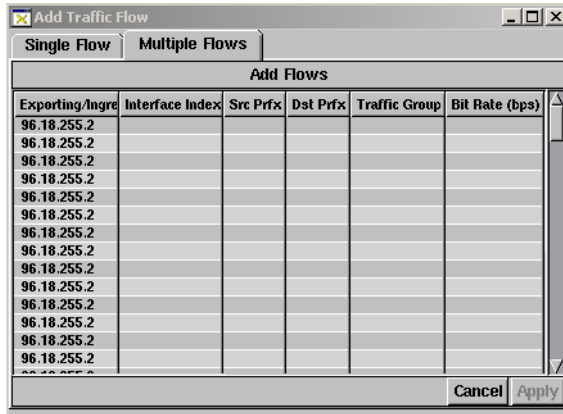


Figure 105 Add Multiple Traffic Flow Dialog Box

Each line in the table represents one new traffic flow.

- 10 Select a row and double-click the left mouse.
A drop down menu will open.
- 11 Select the traffic group that you want the traffic flow to apply to.
- 12 Click **Apply** to save the changes or **Cancel** to dismiss the dialog box without adding a traffic flow. to
- 13 Click **Close (X)** to close the *Add Traffic Flow* dialog box.

Edit Traffic Flows

The ability to edit traffic flows in Design mode allows you to plan the most effective routing scenario for your network. You can manipulate traffic flows using the functions available in the *Edit Flows* window. For example, you can move a traffic flow from one exporting router to another or remove a flow from the topology.

To edit traffic flows, perform the following steps:

- 1 From the *Edit* menu, choose the **Edit Traffic Flows** menu option.
The *Edit Flows* window appears as shown in [Figure 106](#).




Figure 106 Edit Flows Dialog Box

- 2 If desired, limit the list of traffic flows shown using the **Filter** by drop-down list.

You can perform the following functions using the *Edit Flows* window:

- [Change the Bitrate of a Flow](#) on page 311
- [Move a Traffic Flow from One Router to Another](#) on page 312
- [Delete a Flow from a Router](#) on page 313
- [Add a Flow Server](#) on page 313
- [Add a Traffic Flow](#) on page 314

Change the Bitrate of a Flow

Edit the bitrate of a flow to increase or decrease the speed of traffic flowing between the source and destination prefixes. Bitrate is in bits per second (bps).

- 1 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.
- 2 Right-click the highlighted flow(s) and choose **Change flow** from the pop-up menu.

You can also perform this function using the **Change** button at the bottom of the *Edit Flows* dialog box. From the menu of the **Change** button, choose one of the following options:

- **Selected** to change only highlighted flows.
- **All** to change every flow in the table.

The *Change flow bitrate* dialog box appears.

- 3 In the *Change flow bitrate* dialog box, select **Set**, **Scale**, or **Add** and type a value in the text box.

- **Set:** 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:** 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:** 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 type 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:** 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 type 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.
- 4 Click OK to save your changes.

Move a Traffic Flow from One Router to Another

Change the exporting/ingress router of one or more traffic flows.

- 1 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.
- 2 Right-click the highlighted flow(s) 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* dialog box. From the menu of the **Move** button, choose one of the following options:

- **Selected** to move only highlighted flows.
- **All** to move every flow in the table.

The *Change ingress router to* dialog box appears.

- 3 In the *Change ingress router to* dialog box, type the address of the router where you'll move the traffic flow. For example, to move the flow from Router A (192.168.107.11) to Router B (24.0.0.10), type the address of Router B (24.0.0.10) in the *IP Address* text box as shown in [Figure 107](#).



Figure 107 Move Flow Dialog Box

- 4 Click OK to save your changes.

Delete a Flow from a Router

Remove one or more traffic flows from the *Edit Flows* table. The flow will be removed from the topology.

- 1 Click a line in the table to highlight the flow.
Hold down the **CTRL** key and click to highlight more than one flow at a time.
- 2 Right-click the highlighted flow(s) 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* dialog box. From the menu of the **Delete** button, choose one of the following options:

- **Selected** to delete only highlighted flows.
- **All** to delete every flow in the table.

The *Confirm deleting selected table entries* dialog box appears.

- 3 Click **Yes** to complete deletion of the flow(s) or click **No** to cancel the deletion.

Undo the deletion using the *Planning Reports* window **List of Edits** as described in [View and Remove Edits](#) on page 320.

Add a Flow Server

- 1 Click the **Flow Server** button at the bottom of the *Edit Flows* dialog box to add a flow server to the simulated network.

The *Add Flow Distribution* dialog box appears as shown in [Figure 108](#).

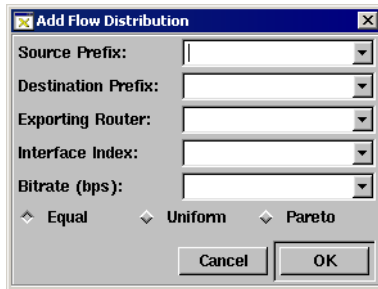


Figure 108 Add Flow Distribution Dialog Box

- 2 Type the source prefix, destination prefix, exporting router, interface index, and bandwidth of the new server.

See Step 3 in [Add Traffic Flow](#) on page 308 for more information about each of these text boxes.

- 3 Choose how traffic will flow from the new server to its clients by clicking **Equal**, **Uniform** or **Pareto**.
- 4 Click **OK** to save the server settings or **Cancel** to dismiss the dialog box without adding a server.

Add a Traffic Flow

- 1 Click **Add** at the bottom of the *Edit Flows* dialog box.

The *Add Traffic Flow* dialog box appears as shown in [Figure 104](#) on [Add Traffic Flow Dialog Box](#) on page 309.

- 2 Type the new flow information as described in [Add Traffic Flow](#) on page 308.

You can view and undo this edit using the *Planning Reports* window **List of Edits** as described in [View and Remove Edits](#) on page 320.

Bring Down a Node

The Down Node function changes the state of a non-BGP node from “up” to “down,” simulating the impact on network activity should the selected router fail.

To bring down a node, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, click **Down Node**.
- 2 On the topology map, click a node to bring down.

As an alternative to steps 1 and 2, you can right-click an IGP or EIGRP node to bring up its information panel, then click **Down** on the panel.

The node turns red, indicating that its state is Down.

- 3 To bring an IGP or EIGRP node back up, right-click it.

The node information panel appears as shown in [Figure 109](#).




Figure 109 Node Information Panel

- 4 Click **Up**.

The node is no longer red, indicating that its state is Up.

Note that BGP protocol nodes do not have **Up** or **Down** buttons on the node information panel and cannot be brought down or back up using this procedure.

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 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 From the planning toolbar or the *Edit* menu, click **Down Peering**.
- 2 On the topology map, click the node where you’ll bring down a peering.

The *Bringing Down Peerings* dialog box appears as shown in [Figure 110](#).

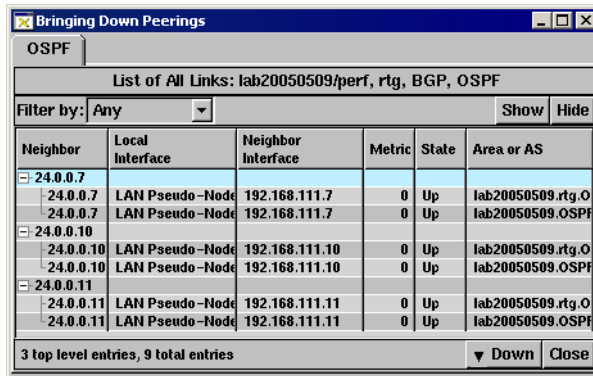


Figure 110 Bringing Down Peerings Dialog Box

- 3 Click the tab for the appropriate protocol.
- 4 Use the **Filter By** drop-down list to choose the attribute you'll use to narrow the list of peerings shown in the table. To use an advanced filter or to write an expression, see [Using Filters](#) on page 280 in [Chapter 6](#), “[The History Navigator](#).”
- 5 Click on a peering to highlight it in the table.
- 6 Click **Down** or right-click the peering(s).
- 7 Choose **All** to bring down all peerings in the table or **Selected** to bring down only highlighted peering(s).
- 8 In the *Confirm* dialog box that appears, click **Yes** to complete the process of bringing down the peering(s). Click **No** to return to the *Bringing Down Peerings* dialog box without bringing down the peering(s).
- 9 Click **Close** to dismiss the *Bringing Down Peerings* dialog box.

Alternatively, you can bring down a peering by right-clicking an IGP or EIGRP 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 fail.

To bring down a prefix, perform the following steps:

- 1 From the planning toolbar or *Edit* menu, choose **Down Prefix**.
- 2 On the topology map, click the node that advertises the prefix you'll bring down.

The *Bringing Down Prefixes* dialog box appears as shown in [Figure 111](#).

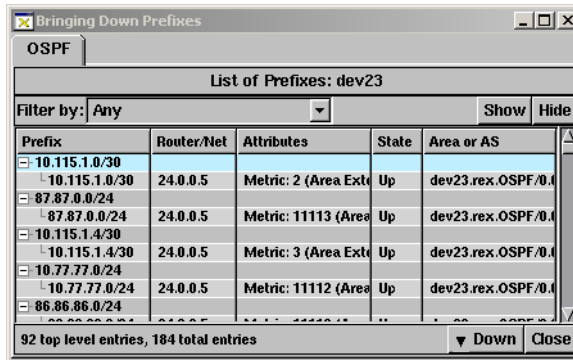


Figure 111 Bringing Down Prefixes Dialog Box

- 3 Click the tab for the appropriate protocol.
- 4 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 301.
- 5 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.
- 6 Click a line in the table to highlight the prefix to bring down.

Hold down **Ctrl** and click the mouse button to highlight more than one prefix at a time.

- 7 Right-click the highlighted prefix(es) and choose **All** or **Selected** from the pop-up menu.

You can also click **Down** at the bottom of the *Bringing Down Prefixes* dialog box. **All** deletes every prefix listed in the table. **Selected** deletes only highlighted prefixes.

- 8 Click **Yes** in the *Confirm* dialog box to complete the process of bringing down the prefix(es) or click **No** to cancel the process.

Bring Down an OSI Prefix

The Down OSI Prefix function changes the state of an OSI prefix from “up” to “down,” simulating the impact on network activity should the selected OSI prefix fail.

To bring down an OSI Prefix, perform the following steps:

- 1 From the planning toolbar or the *Edit* menu, choose **Down OSI Prefix**.
- 2 Click the node that advertises the OSI prefix you’ll bring down.

The *Bringing Down OSI Prefixes* dialog box opens as shown in [Figure 112](#).




Figure 112 Bringing Down OSI Prefixes

- 3 Use the Filter by drop-down list to choose the attribute you’ll use to narrow the list of OSI prefixes shown in the table.
- 4 Click **Show** to list only the items that match the parameters of the filter, or click **Hide** to list only items that do not match the filter parameters.
- 5 Click a line in the table to highlight the OSI prefix you want to bring down.

Note Hold **Ctrl** and click the mouse button to highlight more than one prefix at a time.

- 6 Right-click the highlighted prefix(es) and choose **All** or **Selected** from the pop-up menu.

You can also click **Down** at the bottom of the Bringing Down OSI Prefixes dialog box. **All** deletes every prefix listed in the table. **Selected** deletes only highlighted prefixes.

- 7 Click Yes in the Confirm dialog box to complete the process of bringing down the prefix(es).

Set or Change Link Metrics and Bandwidth

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. The *Set Metric* dialog box allows you to simulate a change to the metric for a physical link in one or both directions. The routing table for the topology is then recalculated using the simulated metric. Any highlighted paths affected by the metric change are redrawn.

To set or change a metric and bandwidth values for a link, perform the following steps:

- 1 Right-click the link on the topology map.

The link information panel appears as shown in [Figure 113](#).




Figure 113 Link Information Panel

- 2 Click **Set Metrics**.

The *Set Metric* dialog box appears as shown in [Figure 114](#).

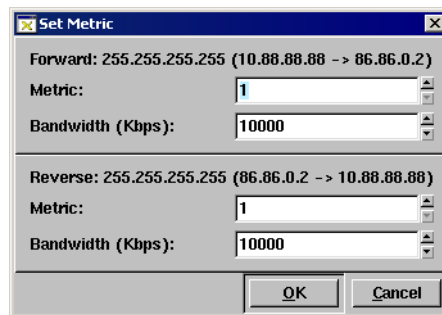


Figure 114 Set Metric Dialog Box - OSPF

[Figure 114](#) shows an example of an OSPF link. For EIGRP links, shown in [Figure 115](#), you can set the bandwidth and delay metrics.

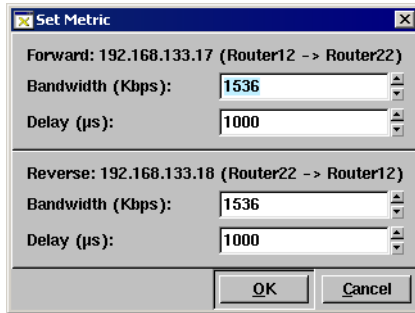



Figure 115 Set Metric Dialog Box - EIGRP

View and Remove Edits

Each edit made to the topology map is displayed in the *Planning Reports* window. Edits can be viewed and removed as described in this section. You can also import and export the list of edits to a database or text file as described in [Importing and Exporting Planning Data](#) on page 340.

To view an edit, choose **Reports** from the *Edit* menu. The *Planning Reports* window appears.

To undo one or more edits from the topology map, perform the following steps:

- 1 Open the *Planning Reports* window as described in the previous paragraph.
- 2 You can undo all edits at once or incrementally:
 - a To undo all edits at once, click the  **Undo All** button in the top right corner of the window. A confirmation dialog box appears. Click **Yes**.
 - b To undo edits incrementally, right-click an edit in the table. A pop-up message appears. Click the pop-up message to confirm that the highlighted edit and all the edits below it will be undone. Any edits above the selected edit are preserved.
- 3 Choose **Yes** from the dialog box to undo the edits or **No** to cancel.

The edits are removed from the list of edits as well as the topology map.

Trending and Estimation

RAMS allows you to predict the trend of your traffic into the future. The system uses the Robust Linear Regression on Vector Time Series trending algorithm. You can predict the trend of IGP, BGP, and total traffic in your network in Design mode.

The flows that RAMS collects from the flow collectors are maintained in the form of a Vector Time Series per flow. The Linear Regression algorithm aggregates the segments of the Vector Time Series using Least Squares, generates a single trend from all the flow data, and then extrapolates this trend to a future date and time you specify.

The algorithm preserves the area of your original flows by extrapolating the trend as segments instead of a single line. The maximum length of each segment is equal to the length of your original flows. In [Figure 116](#), the example trend is not a single straight line but a set of lines, each of which has the same slope. Each line is the trend over a time period that is equal to the length of your original flows.

This trending algorithm provides the most accurate trend when the length of time between the captured flow data end time and your future date is less than the recorded flow data time. For example, if your recorded link traffic data is over a two-month period, and you want to predict the value of the trend less than two months from today, your trend will be highly accurate. If you predict the value of the trend for six months, then you will see the same two-month trend line repeated three times.

Link Traffic Estimation

You can estimate and predict network traffic for IGP links and BGP peerings. The example below shows how to estimate traffic for IGP links. The functionality in the example is the same for estimating BGP peerings.

To estimate link traffic, perform the following steps:


- 1 Open a topology.
- 2 Click the Design mode button.
- 3 From the *Traffic* menu, click **Reports**.

The *Traffic Reports* window appears.

- 4 In the left pane tree, click **IGP Link**.

Note If you want to estimate BGP neighbor, destination, or exit router links, click **Neighbor**, **Destination**, or **Exit Router**, respectively.

The *Link Utilization* table appears at the top of the right pane, and the History tab appears underneath the table. The History tab displays the Traffic report chart.

- 5 In the Traffic report chart title bar, click the  **Click to estimate** button.
- 6 In the menu that appears, click **Predict**.

In the lower left corner of the History tab, the date and time boxes appear. **OK** and **Cancel** buttons appear to the right of the time box.

- 7 Type or select the future date and time at which you want to estimate the link bandwidth in the Date and Time boxes, respectively.
- 8 Click **OK**.

The estimated link traffic curve appears in the Traffic chart. This curve shows the estimated link bandwidth until the future date and time.

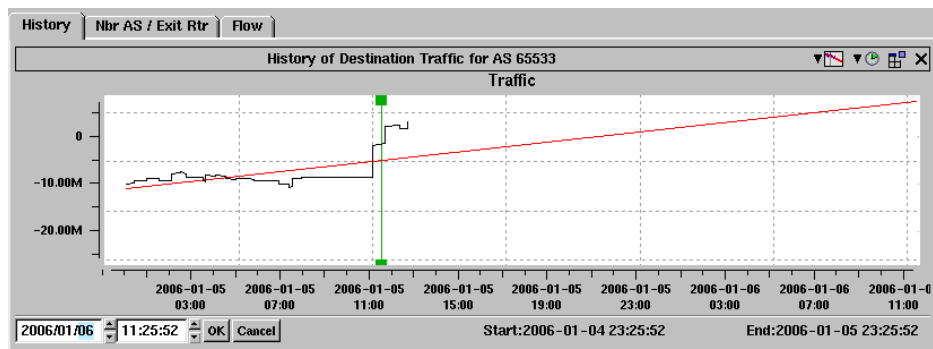




Figure 116 Estimated Link Traffic Curve

You can delete the curve by clicking the  **Click to estimate** button and then clicking **Clear** in the menu.

Total Traffic Estimation

You can estimate total traffic throughout the entire network using a method similar to estimating traffic for IGP links and BGP peering.

To estimate total traffic, perform the following steps:

- 1 Open a topology.
- 2 Click the Design mode button.
- 3 From the *Traffic* menu, click **Reports**.
The *Traffic Reports* window appears.
- 4 In the left pane tree, click Flow Collectors.
The Flow Collector reports appear in the right pane. The Traffic report area appears in the upper left corner of the pane.
- 5 In the Traffic report area title bar, click the  **Click to estimate** button.
- 6 In the menu that appears, click **Predict**.
In the lower left corner of the Traffic report area, the date and time boxes appear. **OK** and **Cancel** buttons appear to the right of the time box.
- 7 Type or select the future date and time at which you want to estimate the link bandwidth in the date and time boxes, respectively.
- 8 Click **OK**.

The total link traffic curve appears in the Traffic chart as shown in [Figure 117](#). This curve shows the estimated bitrate at the future date and time you specified in Step 7.

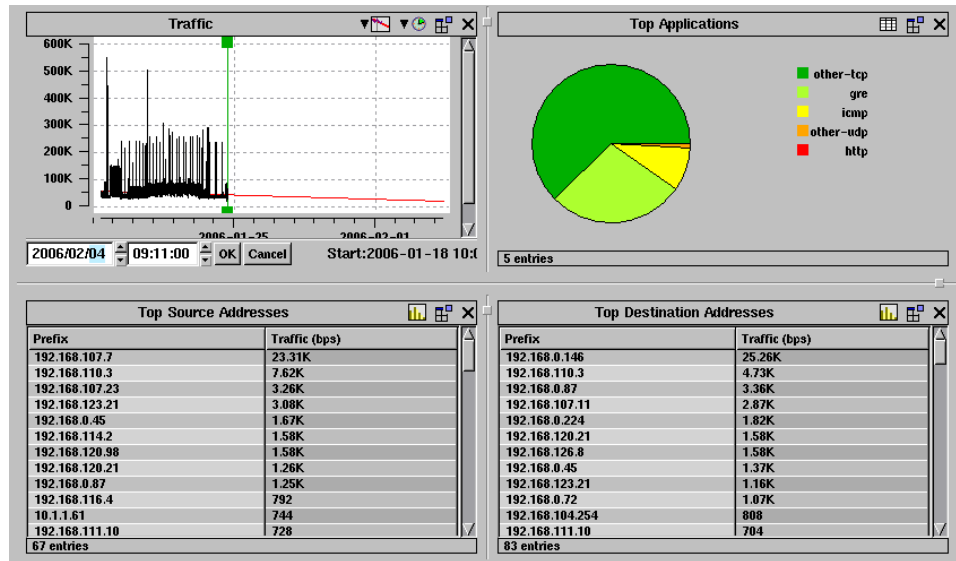



Figure 117 Total Link Traffic Curve

- A dialog box also appears that displays the percentage change in total traffic as the result of the trending analysis, and asks if you want to scale your entire traffic matrix to reflect this percentage change.

If you click **Yes**, RAMS creates a traffic estimation edit that scales all your traffic flows to reflect the percentage change.

You can delete the curve by clicking the  **Click to estimate** button and then clicking **Clear** in the menu.

Link Utilization Prediction

The *Link Utilization Prediction* window, shown in [Figure 118](#), provides a way for you to plan your networks' needs by creating a way to forecast your networks traffic. By selecting a link (identified by the source router and destination router) and time, you can predict what the future traffic for a particular link will be.

Source	Destination	Capacity	Predicted Traffic (bps)	Utilization
Router11	192.168.103.0/24	10.00M	10.00M	100.0%
Router11	Router23	10.00M	297.16K	3.0%
Router11	Router17	10.00M	6.89K	0.0%
Router23	Router16	100.00M	1.77K	0.0%
192.168.103.0/24	24.0.0.12	NA	0	NA
24.0.0.12	192.168.0.0/22	NA	0	NA
Router11	192.168.111.0/24	NA	0	NA
Router11	192.168.110.0/24	NA	0	NA

Figure 118 Link Utilization Prediction Window

Narrow the list of link utilization data using the **Filter By** drop-down list. Filtering allows you to create a more targeted analysis, showing or hiding links based on criteria you specify. The filters choices shown are based on predicted traffic and utilization.

Filter expressions lets you manually enter an expression that is too complex to be set up with either simple or advanced filter menus. For in-depth information about filters and advanced filtering, see [Using Filters](#) on page 280).

The time range bar, located at the bottom left of the *Link Utilization Prediction* window, shown in [Figure 119](#), allows you to set the time for the traffic you want to forecast.



Figure 119 Time Range Bar

Hide — Hides the rows chosen by the filter.

Show — Shows the rows chosen by the filter.

Right-click on a row, and a pop-up menu appears with the options, **Show history**, and **Details by Flow**.

When you select **Show history**, the *History of Utilization* window launches.

This window allows you to go into the detailed graph of how the traffic changed over time for a particular link. You can also individually select a time to predict traffic for. The red line in the graph then shows the trend line that displays on that graph.

The traffic estimation icon (shown in [Figure 120](#)) launches the graphs for the forecasted traffic window (shown in [Figure 121](#)).



Figure 120 Traffic Estimation Icon

The red line corresponds to the predicted traffic for the time frame you selected. The green vertical line represents current traffic conditions.

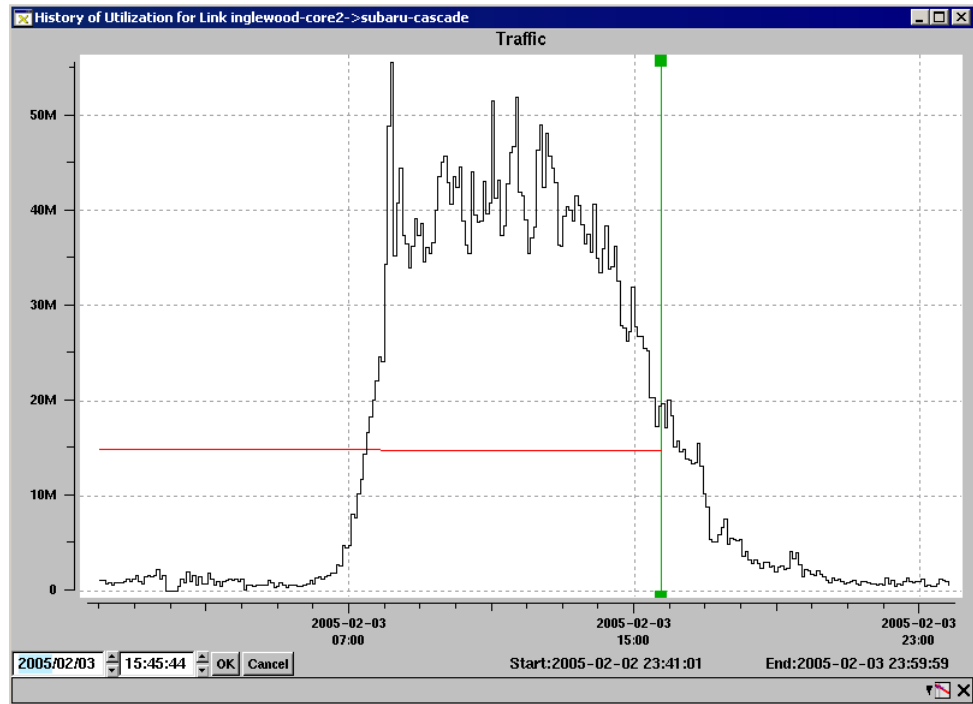


Figure 121 Forecasted Traffic for Link Utilization Prediction

If you select **Details by Flow** after right-clicking in a row, the **Details by Flow For Link** window opens, as shown in [Figure 122](#). This window provides a convenient way for you to see the current flows for the link.

Source Prefix	Destination Prefix	Exporting/Ingress Router	Traffic Group	Traffic (bps) (Current)
37.211.0.0/16	74.32.0.0/16	96.18.255.3:93	Other	2.45M
207.0.0.0/8	74.32.0.0/16	96.18.255.3:23	Other	1.21M
187.101.176.0/20	74.32.0.0/16	96.18.255.3:23	Other	1.19M
169.90.0.0/16	74.32.0.0/16	96.18.255.3:93	Other	859.85K
12.0.0.0/8	74.32.0.0/16	96.18.255.3:93	Other	625.07K
105.189.88.0/21	74.32.0.0/16	96.18.255.3:23	Other	605.87K
49.130.80.0/20	74.32.0.0/16	96.18.255.3:23	Other	530.13K
49.93.176.0/20	74.32.0.0/16	96.18.255.3:92	Other	418.15K
56.97.64.0/23	74.32.0.0/16	96.18.255.3:93	Other	320.00K
171.199.96.0/19	74.32.0.0/16	96.18.255.3:23	Other	312.16K
46.105.110.0/24	74.32.0.0/16	96.18.255.3:93	Other	304.11K
36.10.0.0/16	74.32.0.0/16	96.18.255.3:93	Other	301.21K
67.86.192.0/19	74.32.0.0/16	96.18.255.3:23	Other	278.17K
214.96.0.0/13	74.32.0.0/16	96.18.255.3:23	Other	273.01K
214.80.0.0/12	74.32.0.0/16	96.18.255.3:93	Other	272.12K
168.22.1.0/24	74.32.0.0/16	96.18.255.3:23	Other	266.56K
47.232.0.0/14	74.32.0.0/16	96.18.255.3:23	Other	265.33K
171.251.0.0/16	74.32.0.0/16	96.18.255.3:93	Other	252.10K
169.232.96.0/20	74.32.0.0/16	96.18.255.3:23	Other	239.95K
214.32.0.0/12	74.32.0.0/16	96.18.255.3:93	Other	230.94K
57.91.17.0/24	74.32.0.0/16	96.18.255.3:23	Other	227.20K
67.162.239.0/24	74.32.0.0/16	96.18.255.3:23	Other	220.54K
168.28.0.0/14	74.32.0.0/16	96.18.255.3:23	Other	218.04K
171.116.0.0/23	74.32.0.0/16	96.18.255.3:23	Other	192.53K
169.186.0.0/16	74.32.0.0/16	96.18.255.3:93	Other	180.91K
169.90.240.0/20	74.32.0.0/16	96.18.255.3:93	Other	163.47K
214.205.160.0/19	74.32.0.0/16	96.18.255.3:23	Other	160.00K
168.166.184.0/21	74.32.0.0/16	96.18.255.3:93	Other	160.00K
248.0.0.0/9	74.32.0.0/16	96.18.255.3:93	Other	158.46K
169.44.184.0/21	74.32.0.0/16	96.18.255.3:23	Other	155.73K
105.189.72.0/21	74.32.0.0/16	96.18.255.3:23	Other	151.47K
46.3.107.0/24	74.32.0.0/16	96.18.255.3:92	Other	130.09K
47.231.240.0/20	74.32.0.0/16	96.18.255.3:93	Other	126.72K
39.174.64.0/19	74.32.0.0/16	96.18.255.3:23	Other	121.05K
168.110.128.0/19	74.32.0.0/16	96.18.255.3:23	Other	119.47K
77.219.64.0/19	74.32.0.0/16	96.18.255.3:23	Other	112.91K
224 entries				
224 flows, total bitrate 17.86M				

Figure 122 Details by Flow for Link

Analyzing Topology Edits

RAMS uses present-time statistics derived from traffic data to help plan realistic changes to the network. Use the resulting information to create an optimal configuration for the network, one that will maximize available resources and provide consistent quality service to users.

After using Design mode to make changes to the topology as described in the previous section, analyze the effects of those simulated changes on such variables as link utilization, bandwidth and available capacity using the *Planning Reports* window. This data is displayed per node, interface, exit router and more. 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, after it was added and the amount traffic changed as a result of the edit.

To begin analyzing topology edits in RAMS, use the *Edit* menu to choose **Reports**. By default, *Planning Reports* opens with a table of all the changes you've made to the topology map displayed in the *List of Edits* table as shown in the [Figure 123](#).




Figure 123 Planning Reports Window

Click **Analyze** to see the effects the edits had on the following areas of interest:

- [Changes in Link Utilization](#) on page 329
- [Changes in Link Traffic](#) on page 330
- [Changes in Available Capacity](#) on page 331
- [Changes in Neighbor AS Traffic](#) on page 332
- [Changes in Destination AS Traffic](#) on page 333
- [Changes in Community Traffic](#) on page 334
- [Changes in Exit Router Traffic](#) on page 333

The functions available in this window are described in [The Planning Reports Window](#) on page 294.

Changes in Link Utilization


Identify under- and over-utilized links in the network through the *Planning Reports* window. This information allows more effective traffic routing among links. The *Link Utilization* table displays utilization levels of each link and shows the amount of change resulting from the edits.

RAMS calculates link utilization values by dividing the amount of traffic on a link by the capacity of that link. For more information about link capacity, see [Using the Set Interface Capacities Table](#) on page 336.

To analyze changes in link utilization, perform the following steps:

- 1 Click **Analyze**.
- 2 Click **Link Utilization** in the left pane.

The *Link Utilization* table appears as shown in [Figure 124](#).



Link Utilization					
Filter by: Any					
Source	Destination	Capacity (bps)	Utilization Before	Utilization After	Utilization Change
Router11	192.168.107.0/24	10.00M	0.0%	109.3%	109.3%
192.168.107.0/24	Router16	10.00M	0.0%	99.0%	99.0%
Router11	192.168.103.0/24	10.00M	0.0%	51.4%	51.4%
192.168.103.0/24	24.0.0.5	10.00M	0.0%	50.2%	50.2%
Router11	192.168.110.0/24	10.00M	0.0%	22.4%	22.4%
192.168.110.0/24	24.0.0.3	10.00M	0.0%	22.0%	22.0%
192.168.103.0/24	24.0.0.12	10.00M	0.0%	1.1%	1.1%
24.0.0.12	192.168.0.0/22	10.00M	0.0%	1.1%	1.1%
Router11	192.168.111.0/24	10.00M	0.0%	1.0%	1.0%
192.168.111.0/24	24.0.0.10	10.00M	0.0%	0.8%	0.8%
Router11	192.168.116.0/24	10.00M	0.0%	0.1%	0.1%
192.168.107.0/24	Router23	10.00M	0.0%	0.0%	0.0%
Router23	192.168.118.0/24	10.00M	0.0%	0.0%	0.0%
192.168.116.0/24	25.0.0.4	10.00M	0.0%	0.0%	0.0%
192.168.107.0/24	24.0.0.7	10.00M	0.0%	0.0%	0.0%

Figure 124 Link Utilization Table

The *Link Utilization* table displays the address of the source and destination routers, as well as the capacity of the link. The *Utilization Change* field contains the difference between the amount of utilization before and after the edit. In other words, *Utilization Change* shows the effect the topology edit had on the level of link utilization.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of link utilization data using the **Filter By** drop-down list. Filtering allows you to create a more targeted analysis, showing or hiding links based on criteria you specify.

Changes in Link Traffic

Identify how much traffic a link is carrying by viewing the *Link Traffic* table in *Planning Reports*. If particular links are over- or under-utilized, edit the topology map to redistribute traffic, making better use of resources. The *Link Traffic* table displays how much link traffic changes as a result of the edits.

RAMS calculates link utilization values by dividing the amount of traffic on a link by the capacity of that link. For more information about link capacity, see [Using the Set Interface Capacities Table](#) on page 336.

To analyze changes in link traffic, perform the following steps:

- 1 In the *Planning Reports* window *List of Edits* table, click **Analyze**.
- 2 Click **Link Traffic** in the left pane.

The *Link Traffic* table appears as shown in [Figure 125](#).

Source	Destination	Capacity (bps)	Traffic Before (bps)	Traffic After (bps)	Traffic Change (bps)
Router11	192.168.107.0/24	10.00M	0	10.93M	10.93M
192.168.107.0/24	Router16	10.00M	0	9.90M	9.90M
Router11	192.168.103.0/24	10.00M	0	5.14M	5.14M
192.168.103.0/24	24.0.0.5	10.00M	0	5.02M	5.02M
Router11	192.168.110.0/24	10.00M	0	2.24M	2.24M
192.168.110.0/24	24.0.0.3	10.00M	0	2.20M	2.20M
192.168.103.0/24	24.0.0.12	10.00M	0	109.77K	109.77K
24.0.0.12	192.168.0.0/22	10.00M	0	108.92K	108.92K
Router11	192.168.111.0/24	10.00M	0	97.99K	97.99K
192.168.111.0/24	24.0.0.10	10.00M	0	83.64K	83.64K
Router11	192.168.116.0/24	10.00M	0	10.13K	10.13K
192.168.107.0/24	Router23	10.00M	0	5.91K	5.91K
Router23	192.168.118.0/24	10.00M	0	5.91K	5.91K
192.168.116.0/24	25.0.0.4	10.00M	0	5.01K	5.01K
192.168.107.0/24	24.0.0.7	10.00M	0	0	0

Figure 125 Link Traffic Table

The *Link Traffic* table displays the address of the source and destination routers, as well as the capacity of each link. In addition, the *Traffic Change* field contains the difference between traffic before and after the edit. In other words, *Traffic Change* shows the effect the topology edit had on the level of link traffic.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of link traffic data using the **Filter By** drop-down list. Filtering allows you to create a more targeted analysis, showing or hiding links based on criteria you specify.

Changes in Available Capacity

RAMS calculates available capacity values by subtracting the amount of traffic on a link from the capacity of that link. RAMS calculates link utilization values by dividing the amount of traffic on a link by the capacity of that link. For more information about link capacity, see [Using the Set Interface Capacities Table](#) on page 336.

To analyze changes in available link capacity, perform the following steps:

- 1 In the *Planning Reports* window *List of Edits* table, click **Analyze**.
- 2 Click **Available Capacity** in the left pane.

The *Available Capacity* table appears as shown in [Figure 126](#).

Source	Destination	Capacity (bps)	Available Before (bps)	Available After (bps)	Available Change (bps)
Router11	192.168.107.0/24	10.00M	10.00M	-936.22K	-10.93M
192.168.107.0/24	Router16	10.00M	10.00M	101.74K	-9.90M
Router11	192.168.103.0/24	10.00M	10.00M	4.86M	-5.14M
192.168.103.0/24	24.0.0.5	10.00M	10.00M	4.98M	-5.02M
Router11	192.168.110.0/24	10.00M	10.00M	7.76M	-2.24M
192.168.110.0/24	24.0.0.3	10.00M	10.00M	7.80M	-2.20M
192.168.103.0/24	24.0.0.12	10.00M	10.00M	9.89M	-109.77K
24.0.0.12	192.168.0.0/22	10.00M	10.00M	9.89M	-108.92K
Router11	192.168.111.0/24	10.00M	10.00M	9.90M	-97.99K
192.168.111.0/24	24.0.0.10	10.00M	10.00M	9.92M	-83.64K
Router11	192.168.116.0/24	10.00M	10.00M	9.99M	-10.13K
192.168.107.0/24	Router23	10.00M	10.00M	9.99M	-5.91K
Router23	192.168.118.0/24	10.00M	10.00M	9.99M	-5.91K
192.168.116.0/24	25.0.0.4	10.00M	10.00M	9.99M	-5.01K
192.168.107.0/24	24.0.0.7	10.00M	10.00M	10.00M	0

Figure 126 Available Capacity Table

The *Available Capacity* table displays the address of the router and the interface, as well as the capacity of the link. In addition, the *Available Change* field contains the difference between Available Before the edit (which is calculated by the capacity minus the traffic before) and Available After (calculated by the capacity the edit). In other words, *Available Change* shows the effect the topology edit had on the level of link capacity.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of available capacity data using the **Filter By** drop-down list. Filtering allows you to create a more targeted analysis, showing or hiding links based on criteria you specify.

Changes in Neighbor AS Traffic

Analyze any changes in outbound traffic distribution to neighbor autonomous systems using the *Neighbor AS Traffic* table in *Planning Reports*.

To analyze changes in neighbor AS traffic, perform the following steps:

- 1 Click **Analyze**.
- 2 Click **Neighbor AS** in the left pane.

The *Neighbor AS Traffic* table appears as shown in [Figure 127](#).

AS	Transit Before (bps)	Transit After (bps)	Destination Before (bps)	Destination After (bps)	Traffic Before (bps)	Traffic After (bps)	Traffic Change (bps)
private (65533)	0	17.92K	0	2.80K	0	20.73K	20.73K

Figure 127 Neighbor AS Traffic Table

The *Neighbor AS Traffic* table displays the name of the AS, as well as before and after values for transit, destination, the total of transit plus destination and the total traffic change.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of neighboring AS data displayed using the **Filter By** drop-down list. Filtering allows you to create a more targeted analysis, showing or hiding links based on criteria you specify.


Changes in Destination AS Traffic

Analyze changes in the amount of traffic sent to every unique destination AS using the *Destination AS Traffic* table.

To analyze changes in destination AS traffic, perform the following steps:

- 1 Click **Analyze**.
- 2 Click **Destination AS** in the left pane.

The *Destination AS Traffic* table appears as shown in [Figure 128](#).



The screenshot shows a window titled "Planning Reports" with a sidebar on the left containing a tree view of navigation options: "List of Edits", "IGP Link" (with sub-items "Link Utilization", "Link Traffic", "Available Capacity"), "BGP Peering" (with sub-items "Neighbor", "Destination", "Exit Router"). The main area displays a table titled "Destination AS Traffic". Above the table is a "Filter by:" dropdown menu set to "Any" and "Show" and "Hide" buttons. The table has four columns: "AS", "Traffic Before (bps)", "Traffic After (bps)", and "Traffic Change (bps)". The data rows are:

AS	Traffic Before (bps)	Traffic After (bps)	Traffic Change (bps)
thomson_healthcare (11111)	0	17.92K	17.92K
private (65533)	0	2.80K	2.80K

Figure 128 Destination AS Traffic Table

The *Destination AS Traffic* table displays the name of the autonomous system and the traffic activity before and after edits made to the topology map. The *Total Change* field contains the difference between total before and total after.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of destination AS data displayed using the **Filter By** drop-down list.


Changes in Exit Router Traffic

Analyze changes in traffic on routes that exit from various points in the network using the *Exit Router Traffic* table.

To analyze changes in exit router traffic, perform the following steps:

- 1 Click **Analyze**.
- 2 Click **Exit Router Traffic** in the left pane.

The *Exit Router Traffic* table appears as shown in [Figure 129](#):



The screenshot shows a window titled "Planning Reports" with a tree view on the left containing "List of Edits", "IGP Link", "BGP Peering", and "Community". The "Exit Router Traffic" table is displayed in the main area, filtered by "Any". The table has five columns: "Exit Router", "NextHop Router", "Traffic Before (bps)", "Traffic After (bps)", and "Traffic Change (bps)".

Exit Router	NextHop Router	Traffic Before (bps)	Traffic After (bps)	Traffic Change (bps)
24.0.0.3	192.168.129.1	0	17.92K	17.92K
Router11	192.168.127.27	0	2.80K	2.80K

Figure 129 Exit Router Traffic Table

Because traffic exiting through a particular router may go to any number of BGP next hops, and one next hop address usually corresponds to a peering link, it is valuable to keep track of traffic per next hop address. The *Exit Router Traffic* table displays both the address of the exit router and the next hop router. It also displays traffic activity before and after edits were made to the topology map. The *Traffic Change* field contains the difference between *Traffic Before* and *Traffic After*.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of exit router traffic data displayed using the **Filter By** drop-down list.

Changes in Community Traffic

Analyze the changes in traffic among members of a BGP community using the *Community Traffic* table.

To analyze changes in community traffic, perform the following steps:

- 1 Click **Analyze**.
- 2 Click **Community Traffic** in the left pane.

The *Community Traffic* table appears as shown in [Figure 130](#).

Community Traffic			
Filter by: Any			
Community	Traffic Before (bps)	Traffic After (bps)	Traffic Change (bps)

Figure 130 Community Traffic Table

The *Community Traffic* table displays the community, the address of the exit router, and the traffic activity before and after an edit was made to the topology map. The *Traffic Change* field contains the difference between traffic before the edit was made and traffic after the edit was made.

The view segregation per traffic group, click on a row, and the *Traffic Group* portion of the window will open. This section of the window displays segregation per traffic group.

Narrow the list of community traffic data displayed using the **Filter By** drop-down list.

Using the Set Interface Capacities Table

Link capacity is the amount of traffic a link can transport. Some links have multiple interfaces. You set the capacity of these interfaces using the *Set Interface Capacities* table. As RAMS initially creates a topology map, the Modeling Engine discovers the capacity of each interface in the network, and then displays this information in the *Set Interface Capacities* table.

Customizing Interface Capacity Values

When you open a topology for the first time, the *Set Interface Capacities* table contains the capacities discovered by the **Modeling Engine in the Discovered Capacity** column. The RAMS graphical interface allows you to manipulate this data. You can then store the manually configured capacity values in the database along with the topology map. When you open the saved topology map, the Modeling Engine uses your customized interface capacity values to populate the **Configured Capacity** column in the *Set Interface Capacities* table.

Note In Design mode, the ability to customize capacities using the *Set Interface Capacities* table is disabled. See [To Customize Interface Capacity Values in Design Mode on page 337](#) for more information.

To customize interface capacity values in History or Online mode, perform the following steps:

- 1 After opening a topology map, go to the *Traffic* menu and select **Set Interface Capacities**.

The *Set Interface Capacities* table appears.

- 2 Edit the interface capacity values using one of the following methods:

Edit individual capacity values.

- a Double-click the *Configured Capacity (bps)* table cell of the desired interface.
- b Delete the existing value.
- c Type the new capacity value.

Edit multiple interfaces to use the same capacity value.

- a Hold down the **CTRL** key and click the desired rows to highlight them.

- b Click the **Change** button and choose **Selected**.
- c In the *Change capacity to* dialog box, type the capacity value to apply to all selected rows.
- d Click **OK**.

Edit all interfaces in the table to use the same capacity value.

- a Click the **Change** button and choose **All**.
 - b In the *Change capacity to* dialog box, type the capacity value to apply to all rows.
 - c Click **OK**.
- 3** Click **Apply**.
- 4** To save the new values to the database, click **Save**.

The new values are reflected in the **Configured Capacity** column whenever you load the topology.

To customize interface capacity values in Design mode, perform the following steps:

- 1** After opening a topology map, right-click the interface whose capacity you want to edit.

An information panel appears as shown in [Figure 131](#).




Figure 131 Link Information Panel

- 2** Click **Set Metrics**.
- The *Set Metric* dialog box appears as described in [Set or Change Link Metrics and Bandwidth](#) on page 319.
- 3** Type the customized values and click **OK** to save the new values to the database.

The new values are reflected in the *Set Interface Capacities* table until you exit Design mode. Note that the way the values are displayed differs from protocol to protocol and by the nature of the edit:

- For link-state protocols (OSPF, IS-IS), when you create an interface and set metrics using the information panel, the new metric value appears in the **Discovered Capacity** column of the *Set Interface Capacities* table.
- For link-state protocols (OSPF, IS-IS), when you set metrics on an existing interface using the information panel, the new value is displayed in the **Configured Capacity** column of the *Set Interface Capacities* table.
- For distance vector protocols (EIGRP), when you create an interface or set metrics on an existing interface using the information panel, the new value is displayed in the **Discovered Capacity** column of the *Set Interface Capacities* table.

Clearing Manual Interface Capacity Changes

Manually configured interface capacities are used in preference to discovered capacities. If you want to clear all manual changes, you can enter **0** for all interface capacities, which erases the customized capacity values from the database.

To clear all manual changes from the *Set Interface Capacities* table, perform the following steps:

- 1 In the *Set Interfaces Capacities* table, click the **Change** button and choose **All**.
- 2 In the *Change capacity to* dialog box, type **0**.
- 3 Click **OK**.
- 4 Click **Save**.

RAMS replaces all configured interface capacity values in the table with 0.

Reverting to Saved Interface Capacity Values


After making a series of unsaved edits to the *Set Interface Capacities* table, you can revert to a saved set of capacity values using the **Import** button. Clicking **Import** replaces manually configured, unsaved data in the table with data stored in the database.

Note Importing saved interface capacities does not affect capacities that were initially 0 and were edited to become non-zero values. For example, if you change an interface capacity from 0 to 256, the capacity is still 256 after saved data is imported from the database.

To set interface capacity values and revert manual changes, perform the following steps:

- 1 In RAMS, click the *Traffic* menu.
- 2 Choose **Set Interface Capacities**.

The *Set Interface Capacities* table appears, as shown in [Figure 132](#).



Source	Destination	Interface	Area	Discovered Capacity (bps)	Configured Capacity (bps)
IGP Link Router26 -> Router26	Router16	192.168.116.26	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router26 -> Router26	Router20	192.168.112.26	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router26 -> Router26	Router19	192.168.116.26	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router26 -> Router26	Router16	192.168.116.26	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router26 -> Router26	Router17	192.168.116.26	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router23 -> Router23	Router17	192.168.118.23	LabTestDomain.RouteRec	100.00M	NA
IGP Link Router20 -> Router20	Router16	192.168.109.20	LabTestDomain.RouteRec	10.00M	NA
IGP Link Router19 -> Router19	Router16	192.168.116.19	LabTestDomain.RouteRec	10.00M	NA
IGP Link Router19 -> Router19	Router16	192.168.116.19	LabTestDomain.RouteRec	10.00M	NA

Figure 132 Set Interface Capacities Table

- 3 Click the **Import** button in the bottom right corner of the window.
- 4 Click **Apply**.
- 5 Verify the information in the *Set Interface Capacities* table and click **Save**.

The capacity data you've imported is now reflected in the table and will be stored with the topology map.

Importing and Exporting Planning Data


The ability to import and export edits allows you to plan traffic scenarios corresponding to varying network conditions. For example, if the network experiences heavy traffic activity at a particular time of day, you can develop a plan to effectively handle this activity and maintain another set of edits to correspond to normal network activity.

Store planning information in a database or text file, either by exporting it from the *Planning Reports* window or by creating the file using the appropriate language as described in [Export Topology Edits](#) on page 340. Note that RAMS stores only one set of edits in its database at a time. To store multiple sets of data, export edits to the clipboard, paste them into an external editor, then save the file.

Export Topology Edits

After making changes to network topology, you can export the list of edits to a clipboard or database. Exporting the edits allows planning and storage of different scenarios for the topology, perhaps based on time of day or other variables affecting network activity. You can either export edits to the clipboard, then paste them from the clipboard into a text editor or external database program, or you can export edits to the internal database on the appliance. These scenarios can be loaded later by importing the edits back into RAMS as described in [Import Topology Edits](#) on page 343.

To export edits, perform the following steps:

- 1 In the *Planning Reports* window, click the  **Export** button in the top right corner of the window.
- 2 Choose the destination of the exported edits — clipboard or database — from the pop-up menu.

A dialog box appears confirming that edits are either saved to a database or copied to the clipboard. Note that when you export edits to the database, RAMS stores only one table at a time. In other words, the next time you export a set of edits to the database, your previous set is overwritten.

- 3 Click **OK**.

Exported edits follow a prescribed format. Information in the table is not exported literally as it appears in the list of edits. Instead, it is formatted into an external table format suitable for a person to interpret and manipulate in a spreadsheet. Each row in the table is a separate edit, where the columns specify the parameters of the edit.

For example, the line of text below represents the addition of a traffic flow:

```
add flow -area -srcPrefix -destPrefix -exporterIP -bw
```

In keeping with this format, the following values would be exported for the highlighted flow shown in [Figure 133](#):

```
add flow -area lab.plan/box -srcPrefix 192.168.101.0/24
-destPrefix 192.168.111.0/24 -exporterIP 192.168.107.11 -bw
74
```

Src Prfx	Dst Prfx	Exporting/Ingress Router	Bit Rate (Bps)
192.168.101.0/24	192.168.111.0/24	192.168.107.11	74
192.168.107.0/24	192.168.111.0/24	192.168.107.11	66
192.168.116.0/24	192.168.116.0/24	192.168.107.11	428

Figure 133 List of Flows Window

The syntax for each type of edit is included in [Table 6](#).

Table 6 Syntax for Topology Edits


Operation	Type	Options
add	router (IGP)	-area -proto -rtrID -rtrName
add	router (BGP)	-area -proto -rtrID
add	peering (IGP)	-area -proto -srcRtrID -destRtrID (for pseudonodes, destRtrID will be a prefix) -sif -dif (for pseudonodes, dif will be ignored) For EIGRP: -bw -delay For All Others: -metric
add	peering (BGP)	-area -proto -rtrID -destAS -nbrASStart -peerIP (optional) -nbrASContain -nextHopIP

Operation	Type	Options
add	prefix (IGP)	-area -proto -rtrID -net -bw -delay -metric
add	prefix (BGP)	-area -proto -rtrID -net -ASPATH -origin -localPref -MED -nextHopIP
add	flow	-area -srcPrefix -destPrefix -exporterIP -bw -area -srcPrefix -destPrefix -exporterIP -bw -index -distrib
change	prefix (BGP)	-area -proto -rtrID -net -ASPATH -origin -localPref -MED -nextHopIP (not all MED values needed)
change	peering (IGP)	-area -proto -srcRtrID -destRtrID -sif -dif -metric -bw -delay
change	flow	-area -srcPrefix -destPrefix -exporterIP -index -changeBwOp -value -area -srcPrefixFilter -changeBwOp -value -area -destPrefixFilter -changeBwOp -value -area -srcPrefix -destPrefix -exporterIP -index -newExporterIP -area -srcPrefixFilter -exporterIP
up/down	router	-area -proto -rtrID
up/down	peering (IGP)	-area -proto -srcRtrID -destRtrID -sif -dif
up/down	peering (BGP)	-area -proto -rtrID -destAS -nextHopIP
down	prefix (BGP)	-area -proto -rtrID -net
down (drop)	flow	-area -srcPrefix -destPrefix -exporterIP -index -area -srcPrefixFilter -area -destPrefixFilter
load		-time ("year-month-day hour:minute:second") <area1> <area2> <area3>

Import Topology Edits

Planning information can be imported to RAMS using the **Import Edits** button. The imported information must follow the syntax described in the previous section.

To import topology edits, perform the following steps:

- 1 In the *Planning Reports* window, click the  **Import** button in the top right corner of the window.
- 2 Choose the source of the edits — clipboard or database — from the pop-up menu.
 - If you choose database, a window appears containing the edits in text format. You can make any modifications to the text in the window before applying the data to the table.
 - If you choose the clipboard, a blank pop-up window appears. Paste the edits into the window from an external program.
- 3 Click **Apply** to add the edits to the table and dismiss the pop-up window.

The *Edits* table is populated with the imported edits.

8 BGP Reports

RAMS features a comprehensive set of predefined reports that provide a quick, easy-to-understand view of routing events and problems in the network. When you require a report, select the desired time period from the RAMS database and generate a report for this time period. Reports are available in HTML format on Route Recorders and, in a deployment with multiple Route Recorders, from the centralized Modeling Engine.

This chapter details how to run BGP reports, describes the data each report provides, and illustrates the situations in which each report is useful.

BGP reports fall into two basic categories:

- Activity Reports

These are high level reports that include data about recent BGP event activity. You will often use these reports to quickly check the general status on a day-to-day or shift-to-shift basis and to quickly identify potential problems and areas of particular interest.

- BGP Activity Summary
- BGP Activity by AS
- BGP Activity by Peer
- Route Flap Report
- Prefix Event Detail

- Logical Topology

These reports provide an in-depth insight into the state of the routing tables for a particular point in time. These reports are typically used during problem identification and resolution. These reports also provide an easy-to-read summary that is accessible in a few seconds, so you do not need to go to each individual router to collect the data.

- 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

When RAMS is installed for the first time, it is a good idea to run and print all of the reports. This provides an overview of the network status and should prove useful for network baseline comparison in the future.

Accessing the BGP Report Pages

IGP and BGP reports pages are available on Route Recorders and, in a deployment with multiple Route Recorders, from the centralized Modeling Engine. It is most effective to access report data from the centralized Modeling Engine. The Modeling Engine is physically closer to the user, which reduces the amount of time required for data to travel over the network. In addition, requesting report data from the Modeling Engine reduces the amount of work required of the Route Recorder, whose primary duty is to record data.

When you obtain reports directly from a Route Recorder, RAMS returns information local to the area and/or protocol being monitored. When you obtain reports from the centralized Modeling Engine, RAMS returns information collected from recorders across entire network.

To access the report pages, use a browser to connect to the RAMS Home page on a Route Recorder or Modeling Engine. Administrator privileges are not required. From the *Home* page, click **Reports Portal** on the top navigation bar. On the *Daily Reports* page that opens by default, click **BGP Reports** on the left navigation bar. The *BGP Reports Activity Summary* page appears, as illustrated in [Figure 134](#). Report data is not available until 15 minutes after recording has begun.

Note The browser must accept cookies to login.




Figure 134 BGP Reports Page

Creating BGP Activity Reports

This section describes the following BGP activity reports:

- [BGP Activity Summary Report](#) on page 348
- [BGP Activity by AS Report](#) on page 349
- [BGP Activity by Peer Report](#) on page 350
- [Route Flap Report](#) on page 350
- [Prefix Event Detail](#) on page 351

BGP Activity Summary Report

The *BGP Activity Summary* report provides a high-level overview of BGP network activity over a specified period of time. Any deviations, positive or negative, indicate changes or problems with the network. These deviations can include new peering sessions or routers appearing on the network.

This report is often run on a daily or per-shift basis as it provides a way 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. If BGP activity is high and stays high, this might indicate that a configuration error occurred 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.

This report displays a list of the prefixes that have oscillated between announced and withdrawn from the BGP protocol.

To configure the BGP Activity Summary report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar.

The *BGP Activity Summary* page appears by default with **Time** and **Database** drop-down lists.

- 2 Select the desired time period from the **Time** drop-down list.

Note 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.

- 3 Select the desired database from the **Database** drop-down list.
- 4 Click **Create Report**.

If there was any BGP network activity for the time period selected, a graph appears with the recorded data displayed as lines on the graph. The lines represent the following data:

- 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.

To determine which routers are responsible for activity spikes, check the [Route Flap Report](#) on page 350.

BGP Activity by AS Report

The *BGP Activity by AS* report provides a filtered view of BGP activity for individual autonomous systems (AS) that comprise the entire network. This report lets you drill down into each AS and identify specific sources of instability or excessive activity. This report is especially useful for private enterprise internets where more than one AS is in use and connected by BGP.

To configure the BGP Activity by AS report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar, then click **Activity by AS**.

The **Database** drop-down list appears on the *Activity by AS* page.

- 2 Select the desired database from the drop-down list.
- 3 Click **List AS**.

If there was any BGP network activity between the selected databases for the time period, a graph opens displaying incoming updates and withdrawals and outgoing updates and withdrawals.

After viewing this report, you may want to refer to the [BGP Activity Summary Report](#) on page 348 to view specific sources of instability. Alternatively, you can open the *History Navigator* window and perform an event analysis for the time period in question.

BGP Activity by Peer Report

The *BGP Activity by Peer* report is useful for quickly identifying which BGP peers are most active and for diagnosing internal churn. Since this report provides a greater level of detail, you will normally start with the *BGP Activity Summary Report* and then run this report if you notice an unusual amount of churn activity. This report provides a quick way of pinpointing activity that exceeded the normal expected range.

To configure the BGP Activity by Peer report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar, then click **Activity by Peer**.
The **Database** drop-down list appears on the *Activity by Peer* page.
- 2 Select the desired database from the drop-down list.
- 3 Click **List Peer**.
The **Peer** and **Time** drop-down lists appear.
- 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 there was any BGP network activity between the database and the selected peer, the report is displayed as a line chart over time, and displays the update churn (updated prefixes) and withdrawn churn (prefixes withdrawn).

Route Flap Report

The *Route Flap* report provides a list of prefixes that have oscillated 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. In general, you would run this report as a periodic status check to determine if a problem requires further investigation.

To configure the Route Flap report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar, then click **Route Flap Report**.
The **Time** and **Database** drop-down lists appear on the *Route Flap Report* page.
- 2 Select the desired time period from the **Time** drop-down list.
- 3 Select the desired database from the **Database** drop-down list.
- 4 Click **Create Report**.

If there was any BGP network activity, the report displays a table that provides the prefix ID, the peer that announced or withdrew the prefix, the date/time of the last update, and the current prefix state (announced or withdrawn).

The table can be sorted by prefix, the peer which announced/withdrew the prefix, the number of events, or the current state. To change the sort order, click the heading of the appropriate column.

Prefix Event Detail

The *Prefix Event Detail* report is a detailed report that provides insight into how long a problem has been happening and identifies which prefixes are affected. For example, you might run this report if the Route Flap Report identifies a prefix that is flapping. The *Prefix Event Detail* report provides insight into how long a prefix has been experiencing intermittent service and identifies the customers that have been affected by the associated service degradation. This report can save substantial time because you do not have to log into each individual router and view the routing tables to try to identify the problem.

To configure the Prefix Event Detail report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar, then click **Prefix Event Detail**.
- 2 On the *Prefix Event Detail* page, select the desired database from the **Database** drop-down list.
- 3 Select the desired start (**From**) and end (**To**) times from the respective **Time** drop-down lists.

4 Type the prefix number in the *Prefix* text box.

5 Click **Create Report**.

The report is displayed in a tabular format, which can be sorted based on the headers of each column. The report provides the time of event, direct peer address, Op (type of event), and attributes (BGP announcement attributes of routes control mechanisms).

Creating BGP Logical Topology Reports

This section describes the following BGP logical topology reports:

- [Route Distribution Detail Report](#) on page 353
- [Redundancy by Prefix Report](#) on page 354
- [Baseline Redundancy by Prefix Report](#) on page 355
- [AS Reachability Report](#) on page 355
- [Baseline AS Reachability Report](#) on page 355
- [Prefix Reachability Report](#) on page 356

Each of the reports in this section are accessible from the *Reports Portal* page on a Route Recorder or Modeling Engine, and are configured by following the same set of steps.

To configure a BGP Logical Topology report, perform the following steps:

- 1 On the *Reports Portal* page of a Route Recorder or Modeling Engine, locate **BGP Reports** on the left navigation bar, then locate the subheading, **Logical Topology**.
- 2 Click the name of the report to configure.
- 3 On each report page, select the desired database from the **Database** drop-down list.
- 4 Select the desired date and time information from the **Time** drop-down lists.
- 5 Click **Create Report**.

Route Distribution Detail Report

The *Route Distribution Detail* report provides an insight into the distribution of BGP routes as determined by the BGP path selection algorithm. These distributions are key in traffic engineering, capacity planning, and configuring maintenance activities. 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.

The report is displayed in a tabular format and can be sorted based upon Next Hop, Origin Router, or Next Hop AS. These three variables are critical in determining the BGP routing policies that will provide optimal routing across internal infrastructure, as well as network exits. These policies influence traffic distributions and can have dramatic effects on the costs and performance of the network as a whole.

The report specifies the prefix ID, next hop address(es) on the BGP path, address of the router that announced the prefix, AS number associated with the next hop, local preference (BGP mechanism that selects best path), AS path, MED, and BGP community.

The *Route Distribution Detail* report can be customized by sorting on the key variables of Next Hop (the address the router traffic will be sent to based on the BGP path selection algorithm), Origin Router (the router that announced the prefix), and Next Hop AS (the autonomous system number associated with the next hop determined by the BGP path selection algorithm).

The *Next Hop Detail* report is a drill-down report that is accessed from the *Route Distribution Detail* report and provides a view of all prefixes over each BGP route.

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. This report is normally run 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.

The report can be sorted by the number of available next hops, so you can quickly and easily identify prefixes that are only available through a single point path.

The *Redundancy by Prefix* report is displayed in a tabular format and contains the following data; Prefix ID, baseline number of next hops, number of next hops (based upon period for which report is run), next hop prefix, next hop AS ID.

Note Baseline number of next hops is calculated based on an average sampled over a 1-week period (last 7 days) and any that have been in the routing table 80% of the time are included.

Baseline Redundancy by Prefix Report

The *Baseline Redundancy by Prefix* report identifies whether or not each routed prefix on the network is available from the RAMS appliance. Run this report to ensure that all network prefixes are available from the RAMS appliance.

The *Baseline Redundancy by Prefix* report is displayed in a tabular format.

AS Reachability Report

The *AS Reachability* report indicates the degree of connectivity, next hops, and AS paths toward all reachable autonomous systems. The report enables you to quickly sort through the list of reachable autonomous systems and the paths taken to reach them. This report can be used to validate security and routing policies and to ensure that there are no single points of failure between the network and key external autonomous systems. You can refer to this report during planning to see whether there is adequate network redundancy.

The report is displayed in a tabular format with the following data: Baseline number of next hops, number of next hops (based upon time query), ID of next hops, and AS path.

Note Baseline number of next hops is calculated based on an average sampled over a 1-week period (last 7 days) and any that have been in the routing table 80% of the time are included.

Baseline AS Reachability Report

The Baseline AS Reachability report provides an insight into whether an entire AS is reachable from the RAMS appliance and through what AS paths. This report is normally run to ensure that RAMS is monitoring all autonomous systems as planned (baseline) and to assist in network planning.

The report is displayed in a tabular format with the following data: AS number for next hop, next hops in baseline, AS path in baseline.

Note Baseline number of next hops is calculated based on an average sampled over a 1-week period.

Prefix Reachability Report

The *Prefix Reachability* report indicates the degree of connectivity and BGP attributes for prefixes that are routable by BGP across the entire network. The report can be sorted by Prefix, Destination, AS, or Next Hop, so you can quickly 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.

The report is displayed in a tabular format with the following data: prefix address, destination AS, next hop, AS path.

9 IGP Reports

RAMS features a comprehensive set of predefined reports that provide a quick, easy-to-understand view of routing events and problems in the network. When you require a report, select the desired time period from the RAMS database and generate a report for this time period.

This chapter describes how to generate IGP reports, lists the data each report provides, and illustrates the situations in which each report is useful.

IGP reports fall into three general categories:

- Summary Reports

These reports provide a high-level view of network activity over a specified time period. Typically, these reports are run daily to display day-to-day changes and quickly flag potential problems. There are two types of summary reports:

- Network Events Summary
- Network Churn

- Drill-Down Reports

Drill-down reports provide in-depth information and are normally run after scheduled maintenance to verify that configuration changes were made correctly. They are also used to pinpoint and troubleshoot network problems once the cause of a problem has been identified. The various types of drill-down reports are:

- Changed Metrics
- Flapping Links
- Prefix Origination Changes
- New Prefixes
- New Routers and Links
- Prefixes Withdrawn

- Inventory Reports

This group of reports is used to assess the resources available on the network for a certain time period. There are two types of inventory reports:

- Prefix List
- Prefix Origination from Multiple Sources

After you install RAMS for the first time, it is a good idea to generate and print all of the reports. This provides an overview of the network status, which is useful for baseline comparison in the future.

Preparing to Create IGP Reports

Before generating any reports, ensure that the appropriate database is active.

To find out which databases are active, perform the following steps:

- 1 Open the RAMS application using either the X Window System or a VNC viewer.
- 2 Select **Open Topology** from the *Topology* menu.
The database list opens.
- 3 Scroll through the list. Databases with names in green letters are actively recording data.
- 4 If the desired database is not recording data, ask the administrator to start the recording process for the desired database.

Note For each report, you select a database from the **Administrative Domain** drop-down list. For database configurations with more than one level of administrative domains, the report will cover all databases sharing the same first-level administrative domain name.

Accessing the IGP Report Pages

IGP and BGP reports pages are available on Route Recorders and, in a deployment with multiple Route Recorders, from the centralized Modeling Engine. It is most effective to access report data from the centralized Modeling Engine. The Modeling Engine is physically closer to the user, which reduces the amount of time required for data to travel over the network. In addition, requesting report data from the Modeling Engine reduces the amount of work required of the Route Recorder, whose primary duty is to record data.

When you obtain reports directly from a Route Recorder, RAMS returns only information local to the area and/or protocol being monitored. When you obtain reports from the centralized Modeling Engine, RAMS returns information collected from recorders across the entire network.

To access the report pages, use a browser to connect to the RAMS Home page on a Route Recorder or centralized Modeling Engine. Administrator privileges are not required. From the RAMS Home page, click Reports Portal, then click IGP Reports on the left navigation bar. The IGP Reports page appears, as shown in [Figure 135](#).




Figure 135 IGP Reports Page

Note Your browser must accept cookies to login.

The *IGP Reports* page contains the **Select Report** drop-down list and the **Configure Report** button.

It is recommended that you run the Networks Events Summary report as the first step when trying to identify network problems, because it provides an overall view of network activity. After you identify an area of difficulty, you can run a report that focuses on that area to further assess the problem. For example, if the Networks Events Summary report displays a large number of flapping links, run a Flapping Links report to further analyze this problem.

Creating IGP Reports

This section describes the reports available from the **Select Report** drop-down list on the *IGP Reports* page.

Network Events Summary

The *Network Events Summary* report summarizes network changes. This report presents an overview of network status for a specific period of time and is a good place to start when diagnosing network problems or checking the general network status.

The report displays the following information:

- Number of Link Flaps
- Number of Links with Changed Metrics
- Number of Events
- Number of Prefix Origination Changes
- Number of Prefixes Withdrawn
- Number of New Prefixes Advertised

Configuring the Report

To configure the *Network Events Summary* report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, choose **Network Events Summary**, then click **Configure Report**.

The *Network Events Summary* section appears on the *IGP Reports* page, as shown in [Figure 136](#).

- 3 Select the desired database from the **Administrative Domain** drop-down list.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.

- Select the bottom option button to specify an exact start and end time.

5 Click the **Create Report button.**

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.

A 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.

6 Click **Reconfigure Report to change the time period for the report and run it again.**

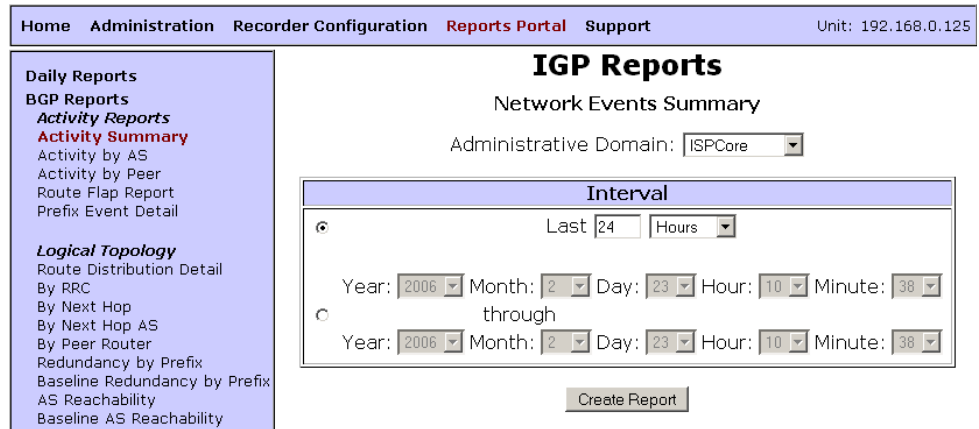


Figure 136 Network Events Summary Section on Reports Page

Understanding the Report

The *Network Events Summary* report contains the fields described in [Table 7](#).

Table 7 Network Events Summary Fields

Field	Description
Number of Link Flaps	The number of times the interface goes up and down.

Field	Description
Number of Links with Changed Metrics	The number of links that have had a metric change between the two time frames, along with what the values were at the beginning and end time time frames.
Number of Events	The number of events recorded in the database.
Number of Prefix Origination Changes	The number of prefixes advertised on a different router.
Number of Prefixes Withdrawn	The number of prefixes that have been withdrawn from the network.
Number of New Prefixes Advertised	The number of prefix advertisements recorded.

Changed Metrics

The *Changed Metrics* report provides a quick summary of all link metrics that have changed in the network. It identifies the router that is advertising the changed metric, along with the original metric, the new metric, and the time when it changed.

This report is normally run after a scheduled maintenance to provide an easy way to verify that planned metric changes have occurred.

Configuring the Report

To configure the *Changed Metrics* report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **Changed Metrics**, and then click **Configure Report**.
The *Changed Metrics* Configuration section appears.
- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:

- Select the top option button to specify a time period up till the current time.
- Select the bottom option button to specify an exact start and end time.

Note If the time specified is earlier than the first entry in the database, the report will begin with the first entry in the database.

5 Click Create Report.

A 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.

6 Click Reconfigure Report to change the time period for the report and run it again.

Understanding the Report

Table 8 details the fields in the **Changed Metrics report**.

Table 8 Changed Metric Report Fields

Field	Description
Link	Link from source router to the destination router.
Source Interface	The source of the link.
Destination Interface	Destination of the link.
Original Metric	The metric originally assigned to the interface.
New Metric	The new metric assigned to the interface.
Time	The date and time the change occurred.

Flapping Links

The *Flapping Links* report lists all routing links that have gone down and come up recently, facilitating the rapid isolation of problem links in the network. This report indicates 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 so that the links with the highest flap count are listed at the top of the report for easy identification.

You would normally run this report after there is an indication of a problem with the links in the network. For example, the *Network Events Summary* report may display a high incidence of flapping link events. Use this report to find out 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. Run the *Flapping Links* report to quickly display the information required to identify exactly where the problem links are.

In cases where a route is flapping an abnormally high number of times, it is possible to set up a RAMS alert on the link to monitor it more closely for a period of time to ensure that an outage is avoided.

Configuring the Report

To configure the Flapping Links report, perform the following steps:

- 1 On the left navigation bar, click **IGP Reports**.

The *IGP Reports* page appears.

- 2 From the **Select Report** drop-down list, choose **Flapping Links**, then click **Configure Report**.

The Flapping Links Configuration section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.

- 4 Specify the time period for the report in the Interval section by choosing one of the following options:

- Select the top option button to specify a time period up till the current time.
- Select the bottom option button to specify an exact start and end time.

- 5 Set the minimum number of flaps required for a link to be included in the report.

- 6 Click **Create Report**.

A 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.

- 7 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

Table 9 details the fields in the *Flapping Links* report.

Table 9 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 the number of routing events that took place for a selected time period. This report identifies all of the sources (routers) that generated events and the number of events attributed to each source. The events are broken down into those describing the router itself, those related to prefixes, and those related to neighbor adjacencies. Routing events tabulation excludes “hello” packets, which are exchanged periodically.

You would normally run this report after the *Network Events Summary* report displays an unusually high level of network events. This report identifies the sources of churn to help isolate the location of a problem so that it can be further investigated. This high level summary report provides a unique insight into the overall level of routing control plane activity throughout the network.

Configuring the Report

To configure the Network Churn report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **Network Churn**, then click **Configure Report**.

The Network Churn Configuration section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**.
A 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.
- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

The report has two columns: *Source Router* and *Number of Events*. [Table 10](#) describes the fields in the **Router** column. [Table 11](#) describes the fields in the **Number of Events** column.

Table 10 Router Column Fields

Field	Description
Name	The router name provided by the routing protocol, if available.
IP Address	Address of the router originating the events.

Table 11 Number of Events Column Fields

Field	Description
Total	Cumulative total of all events.
Router	Number of router events.
Prefix	Number of prefix events for this router.
Link	Number of link events for this router.

Router events captured for this report include the following:

- Router dynamic hostname change (IS-IS only).
- IS-IS overload bit change.
- Router type change, for example between Internal and Area Border Router (ABR) for OSPF.

Prefix events captured for this report include the following:

- Addition and dropping of prefix adjacencies.
- Changes in the metric to a prefix.

Link events captured for this report include the following:

- Addition and dropping of neighbor router adjacencies, including RAMS peering.
- Changes in the metric on the link to a neighbor.

New Prefixes

The New Prefixes report lists all of the newly advertised prefixes during the report period, and how many times. Sources of new prefixes include new links, networks, tunnels, and routers. This report is useful after scheduled maintenance because it lets you quickly verify that planned routing topology changes were made appropriately. This report is valuable as it helps you to ensure that customers continue to receive the appropriate level of service after scheduled maintenances are performed.

Configuring the Report

To configure the New Prefixes report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **New Prefixes**, then click **Configure Report**.

The New Prefixes Configuration section opens.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:

- Select the top option button to specify a time period up till the current time.
- Select the bottom option button to specify an exact start and end time.

5 Click Create Report.

A 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.

To sort the report by the totals in the **Count** column rather than by router prefix, click **Count**. This sorts the report in ascending order by count. Click **Count** again to resort in descending order by count.

6 Click Reconfigure Report to change the time period for the report and run it again.

Understanding the Report

Table 12 describes the fields in the *New Prefixes* Report

Table 12 New Prefixes Fields

Field	Description
Prefix	The IP address of the new prefix.
System ID	The System Identifier for the router.
Name	The name of the router.
IP Address	The IP address of the router. This column may not display if the routers do not have an IP address attributed to them.
Time	The time the new prefix first appeared on the network.
Count	The 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. You would normally run this report after new routers are inserted into the network or new links are set up. This report provides you a quick way to verify that changes to the network have taken place as planned.

Configuring the Report

To configure the **New Routers and Links** report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **New Routers and Links**, then click **Configure Report**.

The **New Routers and Links Configuration** section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the **Interval** section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**.

A 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.

- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

[Table 13](#) describes the **New Router** fields in the *New Routers* section of the report.

Table 13 New Routers Table Fields

Field	Description
IP Address	IP Address of the newly advertised router.

Type	Could be: <ul style="list-style-type: none"> • unknown • internal • area-external • AS-external • both internal and area-external • both internal and AS-external • internal, area-external and AS-external
Name	The router name provided by the routing protocol, if available.

Table 14 describes the **New Links** fields in the *New Links* section of the report.

Table 14 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 prefixes currently advertised in the network, or advertised at any point in the recorded history. The report shows reachability into and out of the network.

This report is often run 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.

Configuring the Report

To configure the Prefix List report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **Prefix List**, then click **Configure Report**.

The Prefix List Configuration section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**. A 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.
- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

Table 15 describes the fields in the *Prefix List* report.

Table 15 Prefix List Fields for an IGP Domain

Field	Description
Prefix	The address of the prefix.

Type	Could be: <ul style="list-style-type: none"> • 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	The area or autonomous system where the prefix is located.
Name	The router name provided by the routing protocol, if available.
IP Address	IP address of the router advertising the prefix.

If you chose OSI as the Administrative Domain, [Table 16](#) describes the fields in the *Prefix List* report.

Table 16 Prefix List Fields for an OSI Domain

Field	Description
Prefix Neighbor/ES Neighbor	The address of the prefix neighbor or ES neighbor.
Type	Could be <ul style="list-style-type: none"> • 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	The area or autonomous system where the prefix is located.
System ID	The System Identifier for the router.
Name	The 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.

You would normally run this report every few days, as it provides a quick insight into potential problems. Problems like a router losing an interface or a flapping link will become apparent.

Configuring the Report

To configure the Prefix Origination Changes report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **Prefix Origination Changes** and click **Configure Report**.

The Prefix Origination Changes Configuration section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**.

A 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.

- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

[Table 17](#) describes the fields in the *Prefix Origination Changes for Current Routers* and [Table 18](#) describes the fields in the *Prefix Origination Router Changes* report for IGP.

Table 17 Prefix Origination Changes Current Routers Report Fields for IGP Domain

Field	Description
Prefix	The prefix address of the network.
Name	The name of the router (if available in the routing protocol).
IP Address	The IP address (router ID) in dotted decimal notation.

Table 18 Prefix Origination Router Changes Report Fields for IGP Domain

Field	Description
Name	The name of the router (if available in the routing protocol).
IP Address	The IP address (router ID) in dotted decimal notation.
Advertised	The number of times the router advertised the prefix. *indicates whether the route was recently advertised.
Withdrawn	The number of times the router withdrew the prefix. *indicates whether the route was recently withdrawn.
Time	The time when the route was most recently announced or withdrawn.

If you chose OSI as the Administrative Domain, [Table 19](#) describes the fields in the *Prefix Origination Changes Current Router Report* Fields for OSI and [Table 20](#) describes the fields for the *Prefix Origination Router Changes* for OSI.

Table 19 Prefix Origination Changes Current Routers Report Fields for OSI Domain

Field	Description
Prefix	The prefix neighbor and ES neighbor.
System ID	System Identifier for the current router.

Field	Description
Name	The router name provided by the routing protocol, if available.

Table 20 Prefix Origination Router Changes Report Fields for OSI

Field	Description
Prefix	The prefix neighbor and ES neighbor.
System ID	System Identifier for the changed router
Name	The router name provided by the routing protocol, if available.
Advertised	The number of times the router advertised the prefix. *indicates whether the route was recently advertised.
Withdrawn	The number of times the router withdrew the prefix. *indicates whether the route was recently withdrawn.
Time	The 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. Reviewing this report provides a quick way to determine if redundant links or hosts (for example, redundant DNS servers, “Anycast” IP Multicast Rendezvous Points, and so on) are operating normally. The absence of a redundant link or host from the list indicates that a redundant link or host is down, possibly causing reduced service levels or other problems within the network. This report can also detect configuration errors.

This drill-down report is a logical next step after the *Prefix Origination Changes* report identifies a problem. This report may also be the first place to look when an alert is received that a redundant link has failed.

Configuring the Report

To configure the Prefix Origination from Multiple Sources report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.

- 2 From the **Select Report** drop-down list, click **Prefix Origination from Multiple Sources**, then click **Configure Report**.

The Prefix Origination from Multiple Sources Configuration section appears.

- 3 From the **Administrative Domain** drop-down list, select the database for the report.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**.

A 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.

- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

Table 21 describes the fields in the *Prefix Origination from Multiple Sources* report.

Table 21 Prefix Origination from Multiple Sources Report Fields

Field	Description
Prefix	The IP Address of the network prefix.
Type	Could be: <ul style="list-style-type: none">• 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	The area or autonomous system where the prefix is located.
Name	Name of the router (if available in the routing protocol).
IP Address	The IP address of the router.

If you selected OSI for the Administrative Domain, [Table 22](#) describes the fields in the *Prefix Origination from Multiple Sources* report.

Table 22 Prefix Origination from Multiple Sources Router Report Fields for OSI

Field	Description
Prefix	The IP Address of the network prefix.
Type	Could be: <ul style="list-style-type: none"> • 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	The area or autonomous system where the prefix is located.
System ID	The System identifier of the router.
Name	The 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. Use this report to quickly identify customers that can no longer access the network.

You would normally run this report after a scheduled maintenance to verify that no prefixes have been dropped unintentionally.

Configuring the Report

To configure the *Prefixes Withdrawn* report, perform the following steps:

- 1 On the navigation bar, click **IGP Reports**.
- 2 From the **Select Report** drop-down list, click **Prefixes Withdrawn**, then click **Configure Report**.

The *Prefixes Withdrawn* Configuration section appears.

- 3 Select the database for the report from the **Administrative Domain** drop-down list.
- 4 Specify the time period for the report in the Interval section by choosing one of the following options:
 - Select the top option button to specify a time period up till the current time.
 - Select the bottom option button to specify an exact start and end time.
- 5 Click **Create Report**.

A 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.

- 6 Click **Reconfigure Report** to change the time period for the report and run it again.

Understanding the Report

Table 23 describes the fields in the *Prefixes Withdrawn* report.

Table 23 Prefixes Withdrawn Report Fields

Field	Description
Prefix	The IP Address of the network prefix.
Name	The name of the router (if available in the routing protocol).
IP Address	The IP Address of the router, if available. This field may not display if you chose OSI as the Administrative Domain.
Time	The time the prefix was withdrawn.
Count	The number of times the prefix was withdrawn.

10 VPN Configuration and Reports

This chapter describes the procedures for configuring the RAMS MPLS VPN (Multi-Protocol Label Switching Virtual Private Network) protocol module and displaying VPN reports.

The information in this chapter only applies to RAMS units that have licenses for both the BGP protocol and the VPN protocol.

The BGP/MPLS VPN, as described in RFC 4364, is the most common form of service-provider-managed 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 target 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, resulting in a scalable solution both in terms of the effort needed to set up the tunnel mesh, and also the number of VPNs that can be supported on the service provider infrastructure.

However, a major source of complexity is the fact that 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 any problems with BGP routes can potentially affect normal Internet connectivity.

Therefore, to manage a large-scale deployment of VPNs in a robust manner, it is critical to tie the protocol diagnostics capability with the VPN service metrics such as reachability and participation. RAMS has extended many of the BGP diagnostic tools to include the VPN protocol.

The RAMS VPN protocol module provides the following benefits:

- A VPN topology overlay that lets you visualize the VPN topology on a per-customer basis.
- Summary and detailed views of the reachability (that is, the existence of routes) for the prefixes advertised by the customer enterprise and the status of the PE routers that participate in a VPN help you to monitor and manage the VPN.
- Reports and customizable alarms alert you to potential or ongoing problems in the VPN.
- Integrated VPN and BGP routing diagnostics enable you to isolate reachability issues down to a single prefix, to determine the routers participating in any VPN, and to isolate and debug complex routing problems.

Understanding the Reachability and Participation Index

RAMS dynamically tracks every prefix that is advertised from each customer on every VPN. RAMS also tracks the PE routers that participate in each VPN.

Although from a customer perspective, the reachability and participation metrics are expected to remain stable, from a service provider's perspective, these numbers can change constantly. The changes might be due to the periodic addition of new customer sites or VPN prefixes being added to existing sites, the addition of new PE routers to the network or the 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 the stability of a VPN with respect to reachability and participation, RAMS establishes a baseline of 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, RAMS looks at routes that have been “up” for more than 80% of the time over the course of 7 days. (Before a route reaches the 7-day mark, its baseline is determined on a day-to-day basis — 80% of 24 hours, 80% of 48 hours, and so on.) This number is assigned a stability index of 100. As the number of prefixes or routers change, the corresponding index number changes accordingly.

Note If you begin recording new routes to an existing database, RAMS 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.

RAMS displays the change in the reachability and participation index in both graphical and text-based reports. You can use these reports to prioritize the allocation of technical resources to resolve customer VPN problems. See [Configuring VPN Reports](#) on page 392 for information about the available reports.

Using VPN Reports

The *VPN Reports* window provides all of the functions that you use to configure and monitor VPNs.

From the *Routing Topology Map* window, you can use either of the following methods to open the *VPN Reports* window:

- Select **VPN Reports** on the *Tools* menu.
- Click the VPN Reports icon on the toolbar.

In a deployment with multiple Route Recorders and a centralized Modeling Engine, you can open the *VPN Reports* window on a Route Recorder or Modeling Engine. Reports should be obtained from the Modeling Engine, which provides a network-wide view of recorded data. Reports are also available directly from Route Recorders, though querying the Route Recorder returns only information local to the area and/or protocol being monitored by the individual recorder.

By default, the *VPN Reports* window displays a summary of reachability and participation in pie-chart format, lists of deviation from a baseline index, and a list of the most recent VPN alarms. The other VPN reports and utilities are listed in the tree structure in the left-hand pane of the window. See [VPN Summary Report](#) on page 392 for more information about the VPN Summary report.

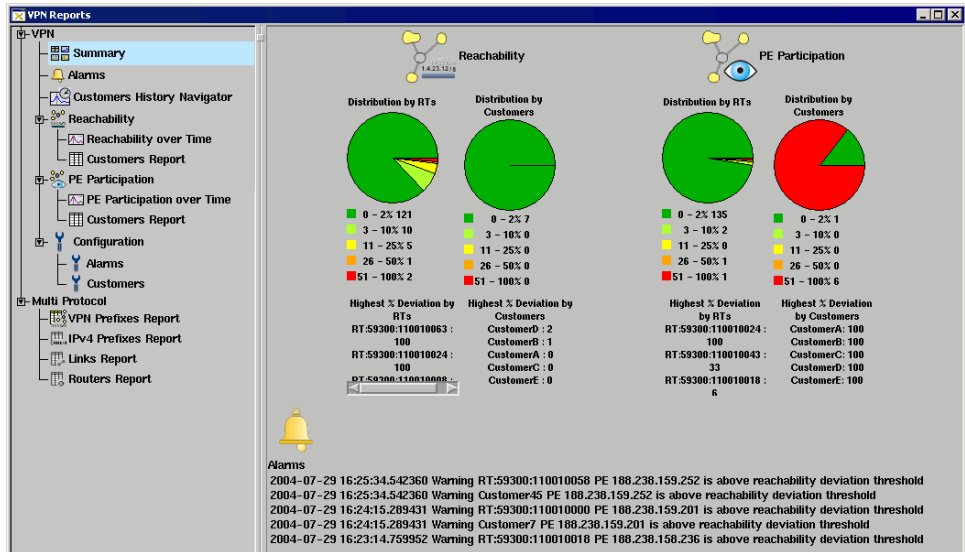


Figure 137 VPN Reports Window

Creating Customer and Route Target 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 RAMS appliance in the network. See [Configure a BGP Instance](#) on page 60 for information about how to enable VPN when establishing peering.

However, to display summary reports on a per-customer or per-PE-router basis, you can associate a customer identifier with one or more Route Targets (RTs) by entering the association information manually in the *VPN Reports* window. To configure many customer/RT associations, it is easier to copy and paste a comma-separated value (CSV) file.

You can specify one of two 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 device 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 route target 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 alerts and watch lists, except alerts requiring information from more than one recorder (for example, Route Change), you should enable queries on the destination Route Recorder.
- 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 RAMS or centralized Modeling Engine *Home* page, click **Administration**, and then click **Queries** on the left navigation bar.
- 2 Select **Enable queries**.
- 3 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.
- 4 Click **Update**.

To set up customer/RT associations manually, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under **Configuration** in the left pane of the window, select the **Customers** subcategory.

The *VPN Customer Configuration* table is displayed in the right pane of the window, with a text box above it into which you can type new associations.

- 3 In the text box, type 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, type each association on a separate line.

- 4 After entering all of the required customer/RT associations, click **Submit**. You may have to scroll down to see the **Submit** button.

The new associations appear in the table in the lower portion of the pane. 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 dialog box containing the complete list.

To set up associations by copy and pasting a file, perform the following steps:

- 1 Create a text file.

- 2 Type customer data in the text file 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 type 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
```

- 3 Copy the contents of the file and paste it into the text box in the *VPN Customer Configuration* window.
- 4 Click **Submit**. You may have to scroll down to see the **Submit** button. The new associations are added to the table in the lower portion of the pane.

Configuring VPN Alarms

When RAMS detects an exception from the baseline index, an alarm (also known as an alert) is generated. You can configure the threshold and severity of the alarms. The following types of alarms are available for configuration:

- **Customer Reachability** – This alarm category notifies you if the network-wide number of routes, or route count, for each customer deviates from the baseline route count. For example, if the baseline number of routes is 100, and the threshold is set to 10%, an alert is triggered when the number of active routes decreases to 90 or increases to 110. Changes in route count occur when a route is announced or withdrawn. Note that this alarm is not configured per PE. In other words, if 10 routes move from one PE to another, an alarm is not triggered.
- **Customer Reachability by PE** – This alarm category notifies you if the number of routes to any PE router changes. Unlike the Customer Reachability alarm, this alarm is not based on the network-wide route count. Rather, it is configured per PE. If 10 routes move from one PE to another, two alarms are triggered, one for each PE.
- **Customer PE Participation** – This alarm category notifies you if the number of active PEs for a customer deviates from the baseline number of PEs for that customer. That is, if the route count for a PE drops to 0, an alert is triggered. A route count of 0 can be caused by a PE crashing, a link between a CE and a PE going down, a CE withdrawing its routes, or the connection between a PE connection and RAMS being severed.
- **New Customer PE** – This alarm category notifies you if a router that is not included in the baseline begins to participate in a VPN. Initially, no routers are in the baseline, so all routers participating in a VPN generate alarms; once the baseline is established, however, only a new router that begins to participate in a VPN generates an alarm. This alarm detects newly added PEs and PE misconfigurations, both of which can result in privacy violations.

To configure an alarm, perform the following steps:

- 1 Open the *VPN Reports* window on a Route Recorder. It is not recommended that alarms be configured on the centralized Modeling Engine.
- 2 Under **Configuration** in the left pane of the window, select the **Alarms** subcategory.

Note You must enable queries on a VPN client before you can edit alarm settings. See [page 387](#) for instructions for enabling queries.

The *Configure VPN Alarms* table is displayed in the right-hand pane of the window. Initially, this table is empty.

3 Click **Add** to begin configuring alarms.

Two drop-down menus, a text box, a **Set** button, and a **Remove** button appear in the *Alarms* table (see [Figure 138](#) for an example).

- The *Category* menu lists the alarm categories from which you can choose.
- The *Severity* menu lists the severity levels that you can assign to an alarm category.
- The *Threshold* text box specifies the threshold for the alarm.

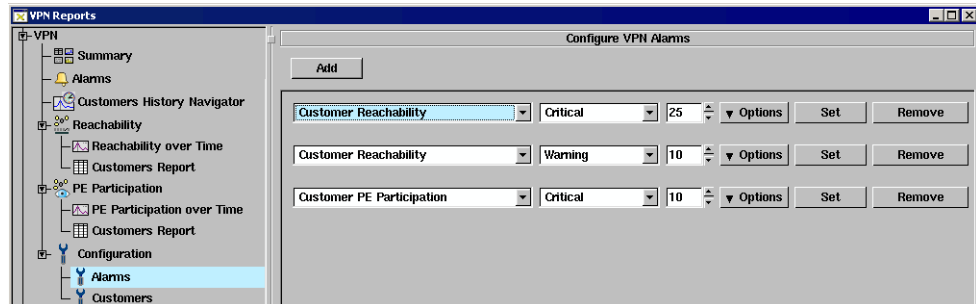


Figure 138 Configure VPN Alarms Window

- 4 Select an alarm category from the *Category* menu.
- 5 Select a severity level for that category from the *Severity* menu. The levels are as follows:
- Critical – Alarms at this level appear in red type in the Alarms report.
 - Warning – Alarms at this level appear in orange type in the Alarms report.
 - Notice – Alarms at this level appear in black type in the Alarms report.
 - Normal – Not an alarm. Does not appear in the Alarms report. Use this level to render a configured alarm inactive.
- 6 Specify an alarm threshold in the *Threshold* text box, either by entering a number or by using the up and down arrows to select a number. If an index varies from the baseline by this percentage, RAMS generates an alarm.

Note You can specify a severity level for the New Customer PE alarm category, but not a threshold. Any router that begins to participate in a VPN generates an alarm at the severity level specified (by default, severity is set to Critical).

- 7 From the *Options* menu, select one or both of the following delivery options:
 - **SNMP Trap**
 - **Remote Syslog**
- 8 Click **Set**.
- 9 Repeat Steps 4 through 8 for each alarm you want to configure.

You can configure different severity levels and corresponding thresholds for a given category. For example, you might want to receive Critical alarms when Customer Reachability reaches a threshold of 25%, Warning alarms when Customer Reachability reaches a threshold of 10%, and Notice alarms when Customer Reachability reaches a threshold of 5%.

To permanently remove an alarm, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under **Configuration** in the left pane of the window, select the **Alarms** subcategory.
- 3 Click **Remove** next to the alarm definition you want to remove.

Configuring VPN Reports

Two groups of VPN reports are available. The first group of reports display information specific to the VPN protocol, and the second group of reports display multiprotocol data. In a deployment with multiple Route Recorders and a centralized Modeling Engine, reports should be obtained 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 display a report, select the appropriate item in the hierarchical tree in the left pane of the *VPN Reports* window.

VPN Summary Report

The VPN Summary report is the default report present when you open the *VPN Reports* window. The report displays three types of information.

- At the top of the pane, the report contains four 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).
- At the bottom of the window, the report lists the most recent five alarms.

The summary report presents a snapshot of VPN data at the time the report is launched. To refresh the data, select a different report, and then select the summary report again.

VPN Alarms Report

The Alarms report lists all VPN alarms.

You can clear an alarm, for example, when changes in the index are caused by planned maintenance.

To display a list of recent alarms, perform the following steps:

- 1 Open the *VPN Reports* window.

- 2 Select **Alarms** in the left pane of the window.

The *Customer Alarms* report is displayed in the right-hand pane of the window.

By default, the report is sorted by time. You can sort the list by severity level, customer name, or message text by clicking the corresponding column heading. To change the sort order (that is, ascending versus descending order), click the column heading a second time.

In addition, you can filter the report by severity level, customer or RT name, or using a simple or advanced filter expression (see [Using Filters](#) on page 280).

Note The **Route Target** column in the Alarms report 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.

To view routing events around the time of an alarm, perform the following steps:

- 1 Right-click the alarm in the *Customer Alarms* report.
- 2 In the pop-up menu that appears, select **Move Time to Here**.
- 3 In the left pane, under **History Navigator**, select the **Customers Report** subcategory. The *History Navigator* pane appears, with its cursor set to the time of the alarm.

To list the routing events that occurred around the time of the alarm, click **Events** and set the beginning marker just to the left of the cursor and the ending marker just to the right of the cursor. The *Events* window lists all events occurring in that time period.

To replay the events that occurred around the time of the alarm, use the playback controls in the History Navigator to step through the events. As playback proceeds, the changes to the routing topology are reflected in the *Topology Map* window.

To clear an alarm, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Click **Alarms**.

The *Customer Alarms* report is displayed in the right pane of the window.

- 3 Click the check box to the right of the message you want to clear, or click **Check All** to clear multiple alarms at once.

To clear all alarms at once, click **Clear All**. A confirmation dialog box displays the number of alarms that will be cleared.

- 4 Click **Apply** to redisplay the report without the cleared items.

To refresh the Customer Alarms list with any new alarms that may have been generated, click **Reload**.

VPN Reachability Reports

You can display reachability data as a graph or as a summary.

To display the Reachability graphs, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under **Reachability** in the left pane, click **Reachability over Time**.

The reachability graphs describing deviation from the baseline by RT and by customer are displayed in the right-hand pane of the window. In these graphs, the x axis is time and the y axis is percentage of deviation.

To display the Reachability summary, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under the category **Reachability** in the left pane, click **Customers Report**.

The reachability summary report is displayed in the right-hand pane of the window. The report includes the customer identifier or RT identifier, the number of PE routers that advertise that RT, and the numbers of active routes, baseline routes, withdrawn routes, new routes, and the deviation from the baseline.

Radio buttons at the top of the pane let you specify whether you want to view data by customer or by RT.

The *Filter By* drop-down menu lets you filter the report to include only the data you want to see. In addition, you can re-sort the data by clicking any column heading in the report.

Note The **Route Target** column in the Reachability summary report 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.

VPN Participation Reports

You can display participation data as a graph or as a summary.

To display the Participation graphs, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under the category **PE Participation** in the left pane, click **PE Participation over Time**.

The participation graphs describing deviation from the baseline by RT and by customer are displayed on the right side of the window. In these graphs, the *x* axis is time and the *y* axis is number of customers experiencing a given percentage of deviation.

To display the Participation summary, perform the following steps:

- 1 Open the *VPN Reports* window.
- 2 Under the category **PE Participation** in the left pane, click **Customers Report**.

The participation summary report is displayed in the right-hand pane of the window. 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.

Radio buttons at the top of the pane let you specify whether you want to view data by customer or by RT.

The *Filter By* drop-down menu lets you filter the report to include only the data you want to see. In addition, you can re-sort the data by clicking any column heading in the report.

Note The **Route Target** column of the Participation summary report 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.

Summary Reports for More Detailed Information

Several tools are available in the Reachability and Participation Summary reports that allow you to display more detailed information and to focus your attention on specific RTs or customers.

To drill down in a summary report, perform the following steps:

- 1 In the left pane, under **Reachability** or **PE Participation**, click **Customers Report** to open the *Reachability Summary* report or the *Participation Summary* report.
- 2 Right-click an entry in the report. A pop-up menu appears containing four options:
 - [Details by PE](#) on page 396
 - [List Routes](#) on page 397
 - [Highlight PEs or Highlight VPN](#) on page 397
 - [PE Participation over Time or Reachability Over Time](#) on page 397
 - [History Navigator](#) on page 399

Details by PE

This option opens a new table at the bottom of the *VPN Reports* window that contains a list of all PE routers associated with the selected RT or customer. For each PE router, the report includes the identifier of the PE router, IP address, type, state, and area or autonomous system (AS).

The *Details by PE* table contains the following buttons:

- **Tear Off** – Creates a new window for the *Details by PE* report, so you can have more than one *Details by PE* report open. If you do not use the **Tear Off** button, the next report you request replaces the current *Details by PE* report.

- **Highlight PE** – Highlights the PE routers listed in the summary report on the topology map in orange.
- **Reload** – Updates the report with new data.
- **Close** – Closes the *Details by PE* report.

List Routes

This option opens a new table at the bottom of the *VPN Reports* window that contains the routes for the VPN customer.

Highlight PEs or Highlight VPN

This option highlights the PE routers that advertise the selected RT or are associated with the selected customer VPN on the routing topology map.

PE Participation over Time or Reachability Over Time

This option opens a new pane at the bottom of the *VPN Reports* window that contains a graph of participation or reachability data for the selected RT or customer. See [Figure 139](#) for an example.

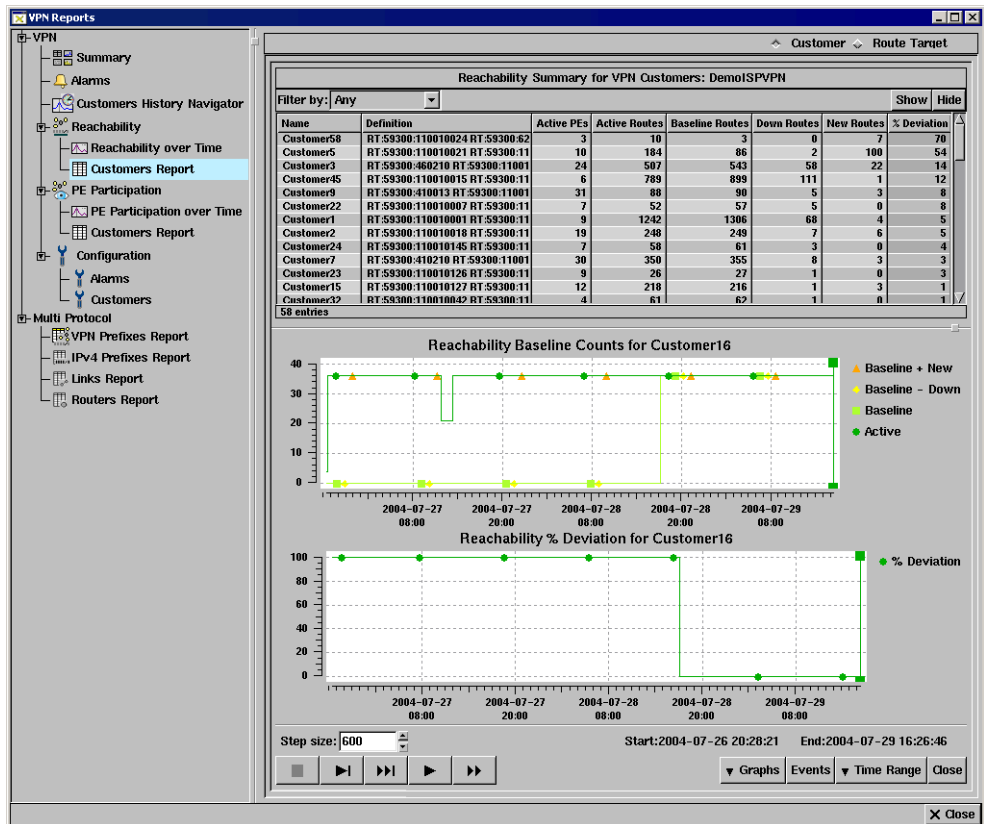


Figure 139 RT Reachability Summary with Graph

The graph pane contains playback controls and buttons similar to the *History Navigator* window, as well as a **Close** button. See [About the History Navigator Window](#) on page 228 for more information. The current time is indicated in the graph by a green vertical line with green grab handles at top and bottom. You can move this time line by dragging it with the mouse.

Note If you move the time line, the data in the summary report above the graph does not change to reflect the new time until you click **Reload** in the summary report pane.

In the graph, the light green line represents the baseline route count for the selected RT, the dark green line represents the count of active routes for that RT, the yellow line represents the baseline route count minus the count of

down routes for that RT, and the orange line represents the baseline route count plus the count of new routes for that RT. All counts are based on the indicated current time.

History Navigator

This option opens a new pane at the bottom of the *VPN Reports* window for the *VPN History Navigator* showing events related to the selected RT or customer (see [Using the VPN History Navigator](#) on page 401 for more information).

Note Events related to a given RT that occurred before a customer/RT association is configured for that RT will show up in a per-RT report displayed later but not in a per-customer report. This is because the data for the report is recorded at the time of the event, and is not updated at the time the association is configured or the report is displayed. Customer/RT association information is recorded only for events that occur after the association is configured.

VPN Prefixes Report

The *VPN Prefixes* report lists all prefixes advertised by VPNs in the network. For each prefix listed, the report lists the router, attributes, state, and area or AS.

The *Filter By* menu lets you filter the report to include only the prefixes you want to see. In addition, you can re-sort the data by clicking any column heading in the report. See [Using Filters](#) on page 280.

IPv4 Prefixes Report

The *IPv4 Prefixes* report lists all prefixes in the network. This report is equivalent to the list generated by the **List Prefixes** option on the *Tools* menu of the *Topology Map* window or the **List Prefixes** button on the topology map toolbar.

Links Report

The *Links* report lists all links in the network. This report is equivalent to the list generated by the **List Links** option on the *Tools* menu of the *Topology Map* window or the **List Links** button on the topology map toolbar.

Routers Report

The *Routers* report lists all routers in the network. This report is equivalent to the list generated by the **List Routers** option on the *Tools* menu of the *Topology Map* window or the **List Routers** button on the *Topology Map* toolbar.

Using the VPN History Navigator

Open the *VPN History Navigator* window by clicking the **Customers History Navigator** link in the left pane of the *VPN Reports* window.

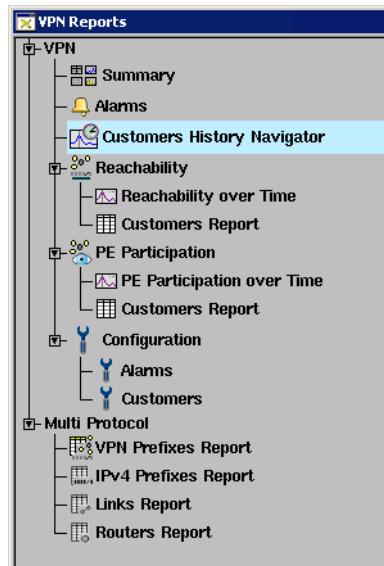


Figure 140 Left Panel of Customers History Navigator

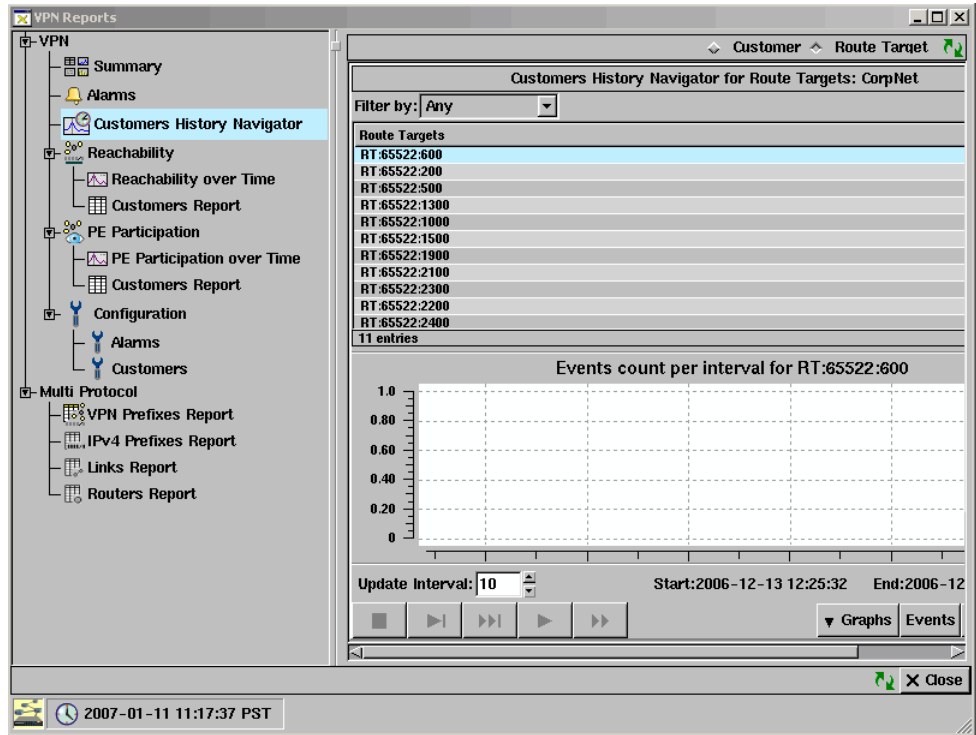


Figure 141 Customers History Navigator Window

To view customer information, perform the following steps:

- 1 Select the **Customer** or the **Route Target** radio button at the top of the *Customers History Navigator* window.
If you have not yet configured any customer identifiers, the display defaults to per-route target information.
- 2 Sort the resulting list of customers by name, definition, or route target by clicking the corresponding column heading.
- 3 Filter the results by VPN customer, route target, expression, or advanced. For more information about filters, see [Using Filters](#) on page 280.

The *VPN Customers History Navigator* window is similar to the *History Navigator* window described in [Chapter 6, “The History Navigator.”](#) See that chapter for descriptions of the controls and buttons in the lower portion of the *VPN Customers History Navigator* window.

11 Path Reports

Use RAMS to analyze network connectivity and optimize routing performance with the Path Reports tool. 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 provides a summary of the paths between routers in a selected topology, and allows you to break down the summary by area of interest: asymmetric links, unused links, source routers, and so on.

Selecting a Topology for Path Reports Analysis

The Path Reports tool computes paths between pairs of routers and generates reports for analysis. The reports are organized by type of analysis. For example, you can view an analysis of all paths in the selected topology, or you can choose to view only asymmetric paths.

To access the Path Reports window, perform the following steps:

- 1 Open a topology map in RAMS as described in [Chapter 5, “The Routing Topology Map.”](#)

- 2 From the *Tools* menu, choose **Path Reports**.

Note Path Reports is disabled in Online mode.

The *Select Topologies and Routers for Path Analysis* dialog box opens.

- 3 Choose one or more databases from the list provided. Selected items are highlighted.

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.

Note 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.

- 4 If desired, select one or both of the following check boxes:
 - **ECMP Paths:** Selecting this check box enables the Equal Cost Multi-Path analysis option, which finds and lists multiple paths of the same cost. See [ECMP Paths Analysis](#) on page 413 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:** When this check box is selected, RAMS systematically fails each link in the selected database, one link at a time, to help determine which link failures are costliest. See [Failure Analysis](#) on page 422 for more information.

Note When you enable failure analysis, you increase computation time significantly compared to reports generated without failure analysis. For example, when you enable failure analysis for a 300-node, non-BGP topology, you can add up to 8 minutes to the computation time.

5 Select which source routers to include in the path analysis:

- **All:** Every source router in the selected database(s) will be included in the analysis.
- **Specify:** Choose which routers to include in the path analysis.

When you click **Specify**, a dialog box appears with 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 dialog box. Click to select one or more routers from the **Available Routers** column, then click → to move the specified routers to the **Selected Routers** column.

6 Select which destination routers to include in the path analysis:

- **All:** Every destination router in the selected database(s) will be included in the analysis.
- **Specify:** Choose which routers to include in the path analysis, as described in Step 5.
- **Same as Source Routers:** Destination routers included in the analysis will be the same as the source routers you chose in Step 5.

7 Select which routed protocols to include in the path analysis:

- **IP:** Route Resolution for IP Prefixes.
- **ISO OSI:** Route Resolution for OSI Prefixes.

Note If OSI IS-IS is not detected by the appliance, the options in Step 7 will not display.

8 Click **OK** to begin generating reports.

You can cancel the report generation at any time. Partial results will be displayed in the *Path Reports* window.

Using the Path Reports Window

The *Path Reports* window is displayed as shown in [Figure 142](#). Path Reports are divided into two main categories of analysis: Path Analysis, described on [page 407](#), and Network Element Analysis, described on [page 418](#). Additionally, if you selected **Failure Analysis** on the *Select Topology* dialog box, a corresponding third category appears, described in [Failure Analysis](#) on [page 422](#).

Source Node	Destination Node	Target NSAP	Paths	Hops	Metric	
Router13	Router8	27.0001.0000.0008.0000.00		1	6	41
Router8	Router13	47.0025.0000.0013.0000.00		1	6	32
Router6	Router8	27.0001.0000.0008.0000.00		1	4	31
Router1	Router13	47.0025.0000.0013.0000.00		1	6	30
Router8	Router6	47.0025.0000.0006.0000.00		1	4	22
Router21	Router8	27.0001.0000.0008.0000.00		1	3	21
Router1	Router6	47.0025.0000.0006.0000.00		1	4	20
Router21	Router13	47.0025.0000.0013.0000.00		1	4	20
Router13	Router21	47.0025.0000.0021.0000.00		1	4	20
Router8	Router21	47.0025.0000.0021.0000.00		1	3	12
Router8	Router8	27.0001.0000.0008.0000.00		1	3	10
Router1	Router21	47.0025.0000.0021.0000.00		1	3	10
Router21	Router6	47.0025.0000.0006.0000.00		1	2	10
Router13	Router6	47.0025.0000.0006.0000.00		1	3	10
Router6	Router21	47.0025.0000.0021.0000.00		1	2	10
Router6	Router13	47.0025.0000.0013.0000.00		1	3	10
Router1	COA8.7857.0000.00	47.0023.COA8.7857.0000.00	Path not found		NA	NA
Router8	Router1	47.0024.0000.0001.0000.00	Incomplete Path		NA	NA
Router8	COA8.7857.0000.00	47.0023.COA8.7857.0000.00	Path not found		NA	NA
Router21	Router1	47.0024.0000.0001.0000.00	Path not found		NA	NA
Router21	COA8.7857.0000.00	47.0023.COA8.7857.0000.00	Path not found		NA	NA

Figure 142 Path Reports Window

The following icons appear in the *Path Reports* window. Not every icon appears in every table.



Color paths — Highlight paths on the routing topology map. Click a particular row in the *Path Reports* table and the corresponding path(s) between the selected source and destination routers is highlighted on the map, as shown in [Figure 144](#).



Show paths — View a table listing all paths between the selected source and destination router pair. The paths are broken down by hop or link as shown in [Figure 143](#).



Reload — Refresh the table.



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 — Put the table you opened in a new window back into the original window.



Show detailed path statistics — View more details about the paths originating or ending with a selected router. For example, in [Figure 145](#), clicking this icon opens an *Analysis of Paths* table in the lower half of the *Path Reports* window. 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.

Path Analysis

When you open the *Path Reports* window, the *Path Statistics* table is displayed by default. *Path Statistics* is one of several tables in the Path Analysis category of reports. Other tables include the *Asymmetric Path Analysis* set of reports, described on [page 415](#). Optionally, *ECMP Analysis*, described on [page 413](#), and *Single Path Analysis*, described on [page 414](#), are displayed if you selected the **ECMP Paths** check box on the *Select Topologies for Path Analysis* dialog box.

Path Statistics Report

The *Path Statistics* table lists every path between router pairs in the selected database. If you chose specific source and destination routers to analyze, as described in [Selecting a Topology for Path Reports Analysis](#) on page 404, 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 has the following columns:


- **Source Node** — Router where the path originates.

- **Destination Node** — Name or ID of the router where the path ends. Since a router may advertise multiple prefixes, RAMS uses a target prefix, described below, as the destination IP address for the path.
- **Target Prefix** — RAMS chooses a prefix unique to each router and uses that prefix as the target for the destination node. In other words, if a node advertises 20 prefixes, RAMS 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, RAMS implements the following rules to determine the target:
 - 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 RAMS will consider non-unique prefixes. However, if RAMS is unable to determine a target prefix, 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 420.

- **Target NSAP** — Displays NSAP address of the router.
 - Note** If OSI IS-IS is not detected by the appliance, the Target NSAP option will not display.
- **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 a path between the source and destination router cannot be computed, 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 target 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. RAMS detects a loop that prevents the path from being computed 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 3 paths between Router A and Router B, with a range of 4-9 hops, Path 1 is made up of 4 hops, Path 2 is made up 5 hops, and Path 3 is made up of 9 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, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Paths* pop-up menu item. The *List of Paths* table appears in the lower half of the *Path Reports* window as shown in [Figure 143](#).

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 RAMS finds four paths between Router A and Router B, each of the four paths and the associated hops are listed along with the following information:

- **Path** — Path(s) 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 appears in this column as well.
- **Source Node** — Router where the link originates.
- **Destination Node** — Router where the link ends.
- **Metric** — Metric value for the link.
- **Protocol** — Routing protocol associated with the link.

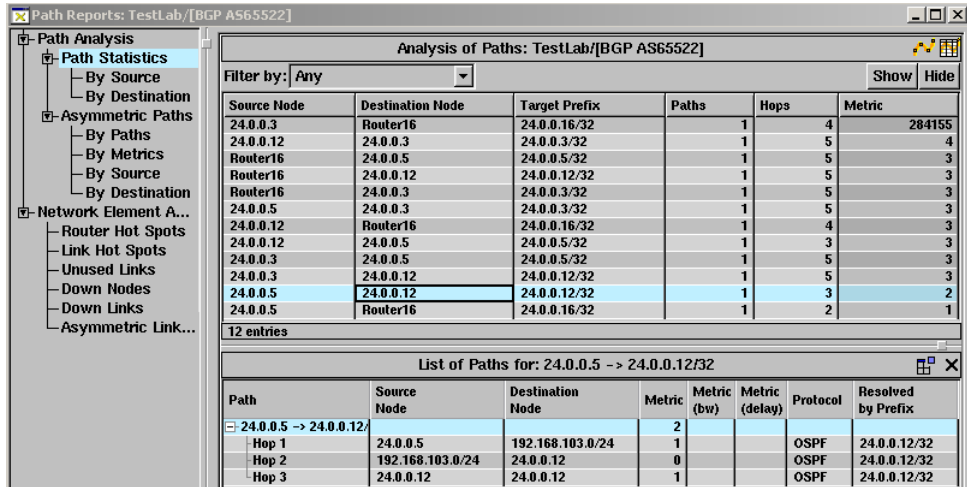


Figure 143 Analysis of Paths with List of Paths Table

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 map as shown in [Figure 144](#).

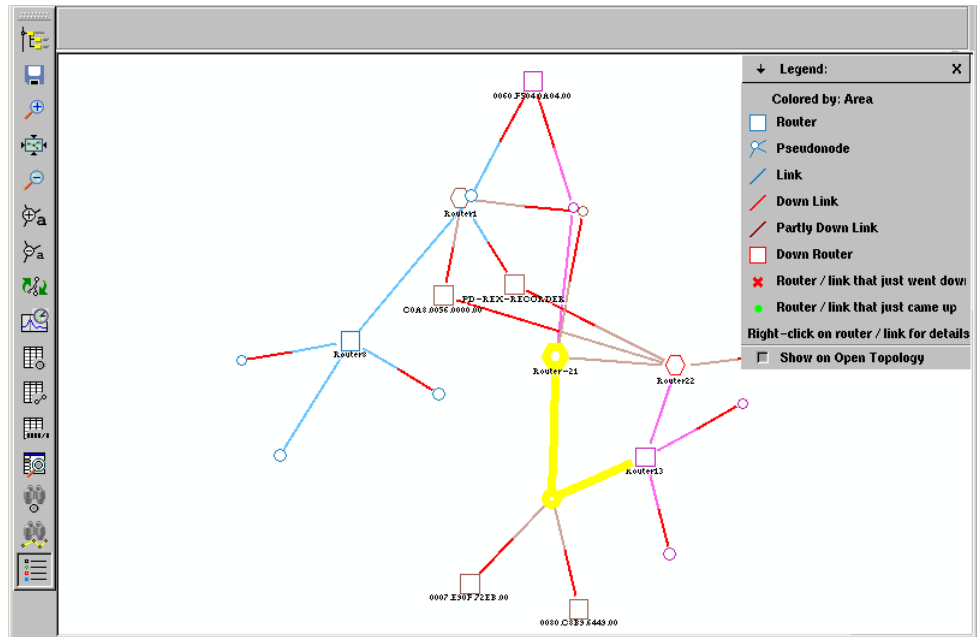



Figure 144 Path Highlighted on the Routing Topology Map

Path Statistics by Source

To view the number of paths reachable by each source node, click **By Source** in the left pane to organize data accordingly. The following information is displayed in the table:

- **Source Node** — Router where a path originates.
- **Reachable Destinations** — Number of routers that the source node 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 node. 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 node.

- **Metric** — The metric value for the path(s) 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 node.

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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Paths* table appears in the lower half of the *Path Reports* window as shown in [Figure 145](#).

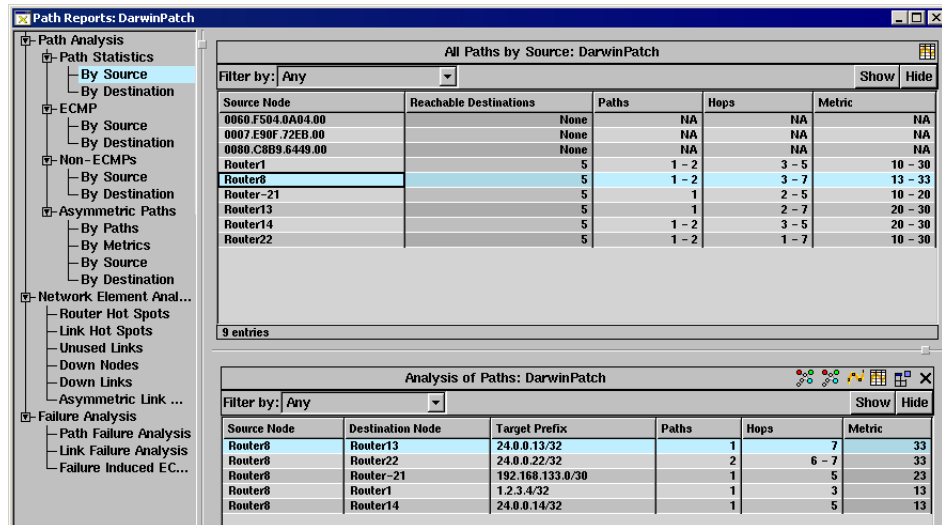



Figure 145 All Paths by Source with Analysis of Paths Table


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 407.


Path Statistics by Destination

To view the number of paths reachable by each destination node, click **By Destination** in the left pane to organize data accordingly. The following information is displayed in the table:

- **Destination Node** — Node where a path ends.

- **Reachable by** — Number of routers that can reach the specified destination node.
- **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 node. 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 node.
- **Metric** — The metric value for the path(s) 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 node.

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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Paths* table appears in the lower half of the *Path Reports* window as shown in [Figure 145](#).

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 407.


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 RAMS finds two equal-cost paths between Router A and Router B, you might determine that two paths do not provide enough redundancy. As a result, you might 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 Node** — Router where the path originates.
- **Destination Node** — Name or ID of the router where the path ends. Since a router may advertise multiple prefixes, RAMS uses a target prefix as the destination IP address for the path.

- **Target Prefix** — RAMS chooses a prefix unique to each router and uses that prefix as the target for the destination node. For more information about target prefixes, see [Path Statistics Report](#) on page 407.
- **Paths** — Number of equal-cost paths between the source and destination nodes.
- **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 node, click *By Source* or *By Destination* in the left pane to organize data accordingly, then see [Path Statistics by Source](#) on page 411 and [Path Statistics by Destination](#) on page 412 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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of ECMP Paths* table appears in the lower half of the *Path Reports* window.

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 [Selecting a Topology for Path Reports Analysis](#) on page 404.

The *Analysis of Single Paths* table has the following columns:

- **Source Node** — Router where the path originates.
- **Destination Node** — Name or ID of the router where the path ends. Since a router may advertise multiple prefixes, RAMS uses a target prefix as the destination IP address for the path.
- **Target Prefix** — RAMS chooses a prefix unique to each router and uses that prefix as the target for the destination node. For more information about target prefixes, see [Path Statistics Report](#) on page 407.

- **Paths** — Number of paths between the source and destination nodes 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 node, click **By Source** or **By Destination** in the left pane to organize data accordingly, then see [Path Statistics by Source](#) on page 411 and [Path Statistics by Destination](#) on page 412 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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Single Paths* table appears in the lower half of the *Path Reports* window.

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.

The *Analysis of Asymmetric Paths* table has the following columns:

- **Source Node** — Node where the path originates.
- **Destination Node** — Node where the path ends.
- **Forward Paths** — Number of paths from the source node to the destination node.
- **Reverse Paths** — Number of paths from the destination node to the source node.
- **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 413.

The screenshot shows a window titled "Path Reports: DarwinPatch" with a tree view on the left and a table on the right. The table is titled "Analysis of Asymmetric Paths: DarwinPatch" and has a filter set to "Any". The table columns are: Source Node, Destination Node, Forward Paths, Reverse Paths, Forward Hops, Reverse Hops, Forward Metric, Reverse Metric, and Metric Difference. The data rows are as follows:

Source Node	Destination Node	Forward Paths	Reverse Paths	Forward Hops	Reverse Hops	Forward Metric	Reverse Metric	Metric Difference
Router1	Router14	1	1	3	3	10	20	10
Router14	Router-21	1	1	3	3	20	10	10
Router14	Router8	1	1	5	5	20	13	7
Router1	Router8	1	1	3	3	10	13	3
Router13	Router8	1	1	7	7	30	33	3
Router-21	Router8	1	1	5	5	20	23	3

Figure 146 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 organize data accordingly.


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 organize data accordingly.

Asymmetric Paths by Source


To view the *Analysis of Asymmetric Paths* table by source node, click **By Source** in the left pane to organize data accordingly. See [Path Statistics by Source](#) on page 411 for a description of this table.


To drill down further and view detailed path statistics for a router in the *Asymmetric Paths by Source* table, highlight the router, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Asymmetric Paths* table appears in the lower half of the *Path Reports* window.

To view details for each of the paths listed in the *Analysis of Asymmetric Paths* table, click the **Show Paths** icon  as described in [Path Statistics Report](#) on page 407.

Asymmetric Paths by Destination

To view the *Analysis of Asymmetric Paths* table by destination node, click **By Destination** in the left pane to organize data accordingly. See [Path Statistics by Destination](#) on page 412 for a description of this table.

To drill down further and view detailed path statistics for a router in the *Asymmetric Paths by Destination* table, highlight the router, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Asymmetric Paths* table appears in the lower half of the *Path Reports* window.

To view details for each of the paths listed in the *Analysis of Asymmetric Paths* table, click the **Show Paths** icon  as described in [Path Statistics Report](#) on page 407.

Network Element Analysis


Determining which network elements play too large a part in network routing and which are playing no part at all is critical in the optimization of network performance. “Hot” spots 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, “cold” spots 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.

Router Hot Spots

This table is sorted by the routers that are used most frequently in paths on the network.

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.




The screenshot shows a software window titled "Path Reports: DarwinPatch". On the left is a tree view with categories like "Path Analysis", "ECMP", "Non-ECMPs", "Asymmetric Paths", and "Network Element Anal...". The "Router Hot Spots" item is selected. The main area displays a table titled "Hot Nodes: DarwinPatch" with a "Filter by: Any" dropdown and "Show" and "Hide" buttons. The table has two columns: "Node" and "Paths".

Node	Paths
Router-21	10
Router1	8
Router14	4
Router22	1
0060.F504.0A04.00	0
Router13	0
0080.C8B9.6449.00	0
0007.E90F.72EB.00	0
Router8	0

Figure 147 Router Hot Spots Table

To drill down further and view detailed path statistics for a router in the *Hot Nodes* table, highlight the router, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Paths* table appears in the lower half of the window.


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 407.

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. 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.




The screenshot shows a software window titled "Path Reports: DarwinPatch". On the left is a tree view with categories like "Path Analysis", "Path Statistics", "ECMP", "Non-ECMPs", "Asymmetric Paths", and "Network Element Anal...". The "Link Hot Spots" item is selected. The main area displays a table titled "Hot Links: DarwinPatch" with a "Filter by: Any" dropdown and "Show" and "Hide" buttons. The table has two columns: "Link" and "Paths".

Link	Paths
192.168.120.0/24 -> Router-21	5
Router-21 -> 192.168.120.0/24	5
Router1 -> 192.168.120.0/24	4
192.168.120.0/24 -> Router1	4
192.168.120.0/24 -> Router1	4
Router1 -> 192.168.120.0/24	4
0000.0021.0000.02 -> Router-21	3
Router-21 -> 0000.0021.0000.02	3
Router14 -> 192.168.120.0/24	2
0000.0014.0000.03 -> Router14	2
Router14 -> 0000.0014.0000.03	2
192.168.120.0/24 -> Router14	2
0000.0014.0000.03 -> Router22	1

Figure 148 Link Hot Spots Table

To drill down further and view detailed path statistics for a link in the *Hot Links* table, highlight the link, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Paths* table appears in the lower half of the window.

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 407.

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 target 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 target 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 on [page 415](#).

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 [Selecting a Topology for Path Reports Analysis](#) on page 404, RAMS will systematically fail links, one by one, along every path in the selected database. The results of the simulated link failure appear in the *Failure Analysis* category of tables, which helps isolate the links that would be most damaging in the event of a failure.

Path Failure Analysis

This table lists each path that failed as part of the failure analysis.

The *Path Failure Analysis* table has the following columns:


- **Source Node** — Node where the path originates.
- **Destination Node** — Name or ID of the node where the path ends. Since a router may advertise multiple prefixes, RAMS uses a target prefix as the destination IP address for the path.
- **Target Prefix** — RAMS chooses a prefix unique to each router and uses that prefix as the target for the destination node. For more information about target prefixes, see [Path Statistics Report](#) on page 407.
- **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.
- **Damaging Link** — Address of the link whose failure caused path cost to increase most of all failed links.
- **Damage Metric** — The difference between the original metric of the path and largest metric in the *List of Paths* table resulting from the link failure.

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  in the upper right corner of the *Path Reports*

window. Alternatively, right-click the row and choose the *Show Effect of Link Failures* pop-up menu item. The *Link Failure Analysis for Path* table appears in the lower half of the *Path Reports* window.

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 the path and the largest metric resulting from link failure.


Additionally, you can show path details for each row in the *Link Failure Analysis* table by clicking the **Show Paths** icon . The first entry in the *List of Paths* table reflects the original path, and the second entry reflects the state of the path after the link failure.


Link Failure Analysis

This table lists each failed link and the number of paths that were damaged as a result of the failure of that link. “Damaged” means the path metric, or cost, increased after the link failed.

The *Link Failure Analysis* table has the following columns:

- **Link** — Address of the failed link.
- **Damaged Paths** — Number of paths whose cost increased after the link failed.


To drill down further and view detailed path statistics for a link in the *Link Failure Analysis* table, highlight the link, then click the **Show Paths** icon  in the upper right corner of the *Path Reports* window. Alternatively, right-click the row and choose the *Show Detailed Path Statistics* pop-up menu item. The *Analysis of Paths* table appears in the lower half of the window.

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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the path row and choose the *Show Effect of Link Failures* pop-up menu item. The *Link Failure Analysis for Path* table appears in the lower half of the *Path Reports* window. This table is described in [Path Failure Analysis](#) on page 422.

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 413 for more information about equal-cost paths.

The *Failure Induced ECMP* table contains the same information as that in the *Path Statistics* table, described on [page 407](#).

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  in the upper right corner of the *Path Reports* window. Alternatively, right-click the path row and choose the *Show Effect of Link Failures* pop-up menu item. The *Link Failure Analysis for Path* table appears in the lower half of the *Path Reports* window as described in [Path Failure Analysis](#) on page 422.

12 Traffic Reports

RAMS calculates reports that provide details about traffic flow using data from the Flow Collectors. In most cases, this data is processed by the Flow Analyzer, a dedicated HP appliance that correlates data collected by the Route Recorder and the Flow Collectors. Use these reports to view and analyze traffic distribution throughout the entire network. You can customize traffic reports by specifying a time range from the last 5 minutes to the entire time traffic has been recorded. RAMS provides 10 different time range options, each of which is described later in this chapter.

Viewing Traffic Reports

After you open a topology that contains traffic routing information, the *Traffic* menu appears. To access the *Traffic Reports* window, click the *Traffic* menu, then click **Reports** as shown in [Figure 149](#).



Figure 149 Traffic Menu

Note Traffic reports are available only in History and Design mode. In Online mode, some *Traffic* menu options are not enabled.

The *Traffic Reports* window appears as shown in [Figure 150](#).

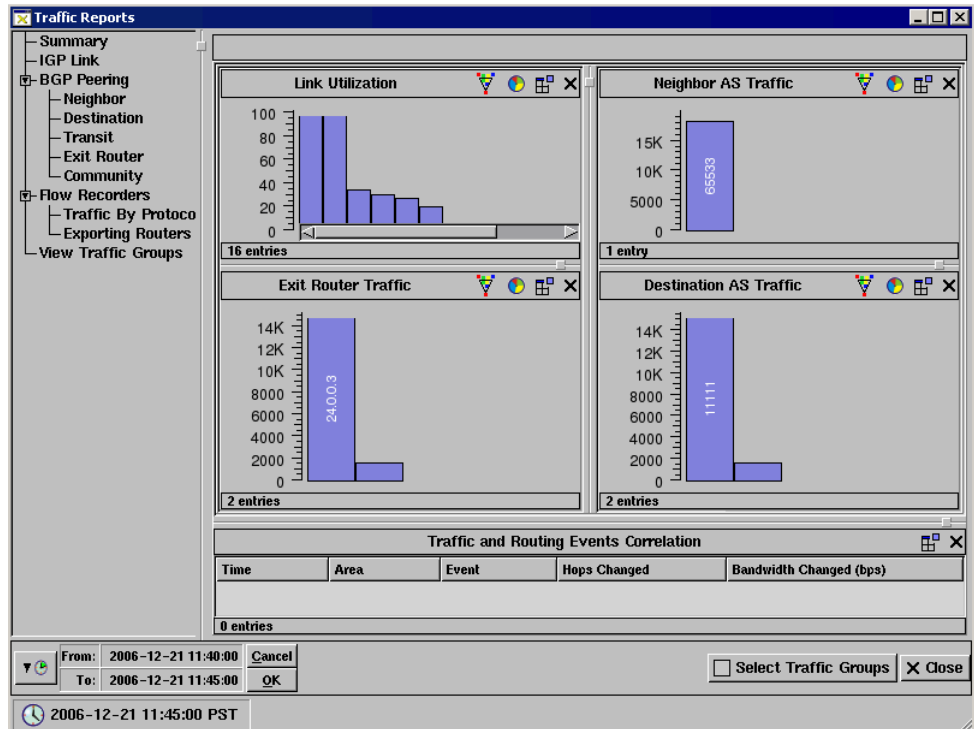


Figure 150 Traffic Reports Summary Window

Using the Traffic Reports Summary Window

The *Traffic Reports* window displays the *Summary* page by default. This screen displays the five most significant traffic reports that you can utilize. These reports are discussed in detail later in this chapter.

From the tree in the left pane, you can choose the type of network traffic information to view by clicking one of the following options:

- **IGP Link** — Shows the amount of traffic for each link in the network. See [Using the IGP Link Report Window](#) on page 438 for more information.
- **BGP Peering** — Shows traffic distribution between one autonomous system (AS) and other autonomous systems. See [Using the BGP Peering Report Window](#) on page 443 for more information.

- **Neighbor** — Displays outbound traffic distribution to each neighbor AS. See [Neighbor AS Traffic Report](#) on page 443 for more information.
- **Destination** — Displays the amount of traffic to every unique destination AS. See [Destination AS Traffic Report](#) on page 444 for more information.
- **Transit** — Shows transitory autonomous systems the flow used as an intermediary hop before reaching the destination. See [Transit AS Traffic Report](#) on page 445 for more information.
- **Exit Router** — Shows the amount of traffic exiting the customer's network through a given router at one point in time. See [Exit Router Report](#) on page 446 for more information.
- **Community** — Displays peering traffic by community attributes. See [Distribution by BGP Community Report](#) on page 446 for more information.
- **Flow Collectors** — Displays top source and destination summaries in four different reports. See [Using the Flow Collectors Report Window](#) on page 448 for more information.
- **Traffic by Protocol** — Provides traffic information according to the protocol of the traffic broken down for the source and destination ports for TCP/UDP. See [Traffic By Protocol Report Window](#) on page 450 for more information.
- **Exporting Routers** — Provides the breakdown of traffic by NetFlow exporting router. See [Using the View Traffic Groups Window](#) on page 452 for more information.
- **View Traffic Groups** — Displays traffic group configuration details. See [Using the View Traffic Groups Window](#) on page 452 for more information.

Using Filters

Most of the tables listed above have filtering (**Filter by ...**), grouping, and/or sorting capabilities. To access these options, right-click the title bar of a column. For example, the *Exit Router Traffic* table filter provides the following options:

- **Sort** — Sort the data by the column you clicked (for example, **Exit Router**).
- **Group** — Organize sub-routers into groups according to parent router.
- **Ungroup** — “Undo” groups to list sub-routers individually.

- **Collapse All** — When routers are grouped, hide sub-routers under the parent router.
- **Expand All** — When routers are grouped, show all sub-routers under the parent router.
- **Hide** — Hide (or remove) the column you clicked.
- **Show** — Restore hidden columns to the table.

For further information about using filters, see [Using Filters](#) on page 280 in [Chapter 6, “The History Navigator.”](#)

Select Traffic Groups Tab

Most of the windows listed above also feature the *Select Traffic Group* tab, which is located at the bottom right page of the window. This tab allows you to choose which traffic group information you want to display and shows the sum traffic belonging to these selected traffic groups. For example, If you have a large number of traffic groups configured but only want to view a few of these groups select *Clear All*, and then select the groups with the information you want to view. If you want to view a large number of traffic groups, select *Check All* and then deselect the traffic groups whose details you do not want to view.

The default view displays all traffic groups, as shown in [Figure 151](#).

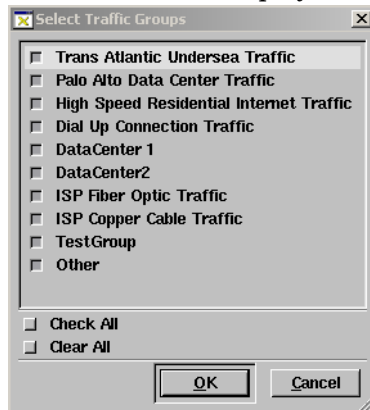


Figure 151 Select Traffic Dialog Box

Display Modes

At the bottom of the *Traffic Reports* window is a status bar. The icon at the left edge indicates the mode:



Online mode – A network icon indicates that the topology is currently being recorded and updates to the routing database are shown on the topology map as they occur. Note that traffic data is delayed by 30 minutes.



History mode – A graph icon indicates that only previously recorded information in the routing database is shown on the topology map. Note that traffic data is delayed by 30 minutes.



Design mode – A network icon indicates that planning features are enabled for the topology map. Note that traffic data is delayed by 30 minutes.

You can switch between modes by clicking the mode icon.

Note Some saved reports require that you open two entries from the topology tree to properly view the report. One topology entry must contain routing information and the other must contain traffic information. Although routing information can be opened on its own, to obtain traffic information you must also open routing.

Window Icons

All *Traffic Reports* pages are divided into one or more report areas. Each area has its own title bar, and you can click one or more of the following icons shown to change the representation of the area.

The following icons are available in the *Traffic Reports* window:



View as Pie Chart: Displays the report as a pie chart.



View as Table: Displays the report in table format.



View as Bar Chart: Displays the report as a bar chart, which is the default view type.



Select Time Range: Select one of the time ranges for the report. See [Creating Time Range Reports](#) on page 432 for more information.



Export Traffic: Exports all visible data displayed on the History graph to the clipboard as a sequence of (x,y) points in ASCII. See the [History Tab](#) on page 439 for more information.



Tear Off: Makes the report area a separate window.



Put Back: Closes the separate report window and restores the report area in the page.



Restore Original View: Restores the default view for the report page.



Close: Closes the separate report window or the report area. When you close the report area, other report areas expand to fill the page.

Creating Time Range Reports

You can generate a report for a specific interval of time. RAMS creates reports based on the following criteria:

- **Last 5 minutes:** Includes the most recent 5 minutes of recorded data.
- **Last hour:** Includes the most recent 60 minutes of recorded data.
- **Last 24 hours:** Includes the most recent 24 hours of recorded data.
- **Yesterday:** Includes data recorded between 12:00 AM and 11:59 PM the previous day.
- **Last 7 days:** Includes data going back 7 days in history, starting with the most recently recorded data.
- **Last week:** Includes data recorded starting on the most recent Sunday at 12 AM, going back to the previous Saturday at 11:59 PM.
- **Last 30 days:** Includes data going back 30 days in history, starting with the most recently recorded data.
- **Last month:** Includes data recorded in the last calendar month, starting on the first day of that month at 12 AM, and ending on the last day of the month at 11:59 PM.
- **All:** Includes all recorded data for the topology.
- **Recent:** When you select this option, a drop down list is displayed with all recent reports available for selection.
- **Custom:** When you select this option, a time range dialog box is displayed. Use the cursor to choose a time range on the graph, and then click **OK**. A report is generated for the specified time range.

Note For some reports (For IGP, the flow report; for BGP, the Flow, Destination AS, and the Exit router report) will be disabled if you change time reports to a time other than the last five minutes.

To generate reports based on a time range, perform the following steps:

- 1 In the bottom left corner of the *Traffic Reports* window, click the **Select Time Range** bar (see [Figure 152](#)) to display the list of available time ranges.




Figure 152Time Range Bar

- 2 Select a time range and confirm the start and end times in the From and To boxes.

Another option to set the time is to use the drop-down menu (see [Figure 153](#)) that appears after you click in the time range bar in Step 1, and selecting the time interval of your choice. The descriptions for each interval listed in this menu is described in [Creating Time Range Reports](#) on page 432.




Figure 153Time Range Drop-Down Menu

- 3 Click **OK** to produce traffic reports for the specified period or **Cancel** to start over.

Using the Summary Report Window

All summary reports include data from all Flow Collectors known to the Flow Analyzer. You cannot dynamically change the set of Flow Collectors. In other words, the summaries always show a network view that is as complete as possible. By default, these reports cover the most recently recorded 5 minutes of traffic activity

The *Summary Report* window in [Figure 149](#) displays the five most crucial traffic reports:

- Link Utilization Report
- Neighbor AS Traffic
- Exit Router Traffic
- Destination AS Traffic
- Traffic and Routing Events Correlation

The *Traffic and Routing Events Correlation* report displays in table format only. You can view the other four reports on this page as the default bar chart, pie chart, or table.

Traffic and Routing Events Correlation Report

The *Traffic and Routing Events Correlation* report table appears at the top of the *Summary Report* window as shown in [Figure 154](#).

Traffic and Routing Events Correlation				
Time	Area	Event	Hops Changed	Bandwidth Changed(bps)
2005-12-19 15:56:47	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Withdraw Neighbor (24.0.0.12,192.168.0.2 DR)	8	528
2005-12-19 15:57:28	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (24.0.0.12,192.168.104.0/24)	9	132
2005-12-19 15:57:28	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (192.168.0.2 DR,192.168.0.0/2)	36	11.10K
2005-12-19 16:00:52	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (24.0.0.12,192.168.0.0/24)	42	69.38K
2005-12-19 16:00:52	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (24.0.0.12,192.168.104.0/24)	6	118
2005-12-19 16:00:53	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Withdraw Prefix (24.0.0.12,192.168.0.0/24)	42	69.40K
2005-12-19 16:00:53	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (192.168.0.2 DR,192.168.0.0/2)	56	69.40K
2005-12-19 16:13:10	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Withdraw Neighbor (24.0.0.12,192.168.0.2 DR)	32	1.04K
2005-12-19 16:13:50	Beta4037.PDRouteRecorder.OSPF/0.0.0.1	Announce Prefix (24.0.0.12,192.168.0.0/24)	27	11.54K

Figure 154 Traffic and Routing Events Correlation Report Table

Each row in the table represents the correlation between network traffic and routing events. The report data is organized in five columns, displaying the time each event occurred, in which area, and a description of the event. The **Hops Changed** and **Bandwidth Changed (bps)** columns show how the event affected the number of hops and the bandwidth for the route.

Link Utilization Report

The *Link Utilization* report, shown in [Figure 155](#), displays the amount of traffic for each link in the network.

Link Utilization					
Source	Destination	Capacity	Current Traffic (bps)	Average Traffic (bps)	Utilization
24.0.0.11	192.168.103.0/24	10.00M	4.16M	0	0.0%
192.168.103.0/24	c-24-0-0-5.hsd1.nj	NA	3.67M	0	NA
24.0.0.11	192.168.110.0/24	NA	2.84M	0	NA
192.168.110.0/24	c-24-0-0-10.hsd1.nj	NA	74.19K	0	NA
192.168.103.0/24	24.0.0.12	NA	1.49K	0	NA
24.0.0.11	192.168.107.0/24	100.00M	66.03K	0	0.0%
24.0.0.11	192.168.116.0/24	10.00M	6.07K	0	0.0%
192.168.110.0/24	c-24-0-0-3.hsd1.nj	NA	2.77M	0	NA
24.0.0.11	192.168.111.0/24	NA	817	0	NA
192.168.107.0/24	c-24-0-0-7.hsd1.nj	NA	143	0	NA

13 entries

Figure 155 Link Utilization Report

The **Capacity** column lists the amount of egress traffic the link is capable of handling, in bits per second (bps). The **Current Traffic (bps)** column lists the current traffic flow for the selected link. RAMS divides the **Average Traffic** value of a link by the **Capacity** value of that link and displays the result in the **Utilization** column. For example, if the capacity of the link is 100 bps and RAMS finds that egress traffic for the link is 79 bps, then 79% of the egress capacity of the link is being utilized.

You can compare utilization values of the links listed in the table to find over- and under-utilized links and make changes to help route traffic more evenly among the links.

To configure a visual representation of link utilization on the routing topology map, use the *Tools* menu to open the *Configuration Options* dialog box. Select **Colors**, then set Traffic Coloring Options.

Neighbor AS Traffic Report

The *Neighbor AS Traffic* report, shown in [Figure 156](#), displays outbound traffic distribution to each neighbor AS.

Neighbor AS Traffic				
AS	Current Traffic (bps)	Average Traffic (bps)	Average Transit (bps)	Average Destination (bps)
PdLabRight (65533)	794.69K	0	0	0

Figure 156 Neighbor AS Traffic Report

In its table format, the *Neighbor AS Traffic* report allows you to view the egress traffic sent from the AS, based on hourly, daily, weekly, and monthly amounts of traffic (in bits per second). The average traffic is displayed in the **Average Traffic (bps)** column. The average transit traffic is displayed in the **Average Transit (bps)** column, and the average traffic received at the Neighbor AS is reported in the **Average Destination (bps)** column.

Exit Router Traffic Report

The *Exit Router Traffic* report shows the amount of traffic exiting the customer network through a given router at a single point in time. By default, this report appears in its table format as shown in [Figure 157](#).




Figure 157 Exit Router Traffic Report

The *Exit Router Traffic* report displays exit router traffic information, including the traffic speed in bits per second.

Traffic exiting at this router may go to many different BGP next hops. Because one next hop address usually corresponds to a peering link, it is important to keep track of traffic on the next hop address level. Therefore, this table shows exit traffic per the exit router and per the next hop address.

Destination AS Traffic Report

The *Destination AS Traffic* report, shown in [Figure 158](#), displays the average amount of traffic to every unique destination AS.



Destination AS Traffic  		
AS	Current Traffic (bps)	Average Traffic (bps)
thomson-healthcare (11111)	794.69K	0

Figure 158 Destination AS Traffic Report

Using the IGP Link Report Window

The *IGP Link* report window contains several reports that allow you to view traffic activity on IGP links. The *Link Utilization* report appears at the top of the window as shown in [Figure 159](#).

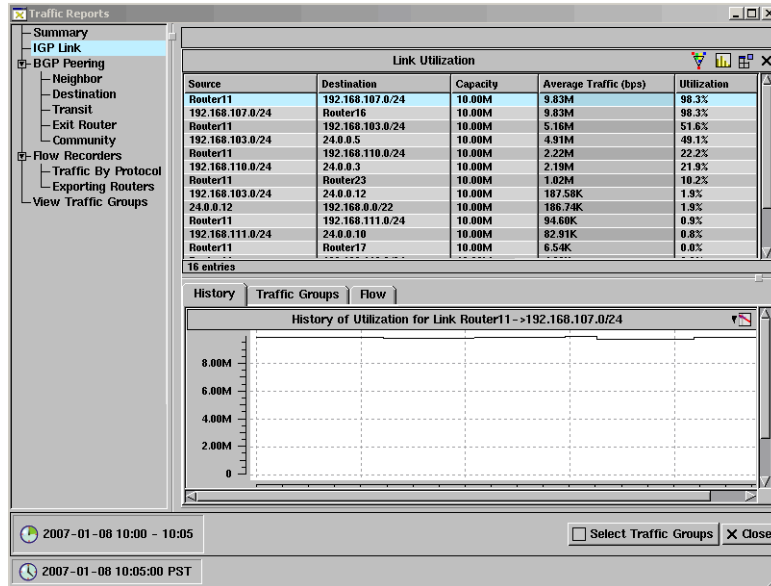


Figure 159 IGP Link Report Window

The following three tabs appear below the *Link Utilization* table:

- **History:** Contains a graph depicting link utilization for a particular time period.
- **Traffic Groups:** Displays the amount of traffic activity for a link selected in the *Link Utilization* window.
- **Flow:** Shows flow details for the link selected in the *Link Utilization* table.

When you click a link in the *Link Utilization* table, the information displayed in the **History**, **Traffic Group**, and **Flow** tabs changes to reflect the network activity of that link. For example, in [Figure 159](#), a link is highlighted in the *Link Utilization* table. Below, on the **Flow** tab, the activity of that link is

broken down into three reports: *Details by Flow*, *List of Prefixes*, and *Find or List Paths*. If you click another link in the *Link Utilization* table, the **Flow** tab changes to reflect the newly highlighted link.

You can also access History, Traffic Group Details, and Flow information when you right-click a link in the *Link Utilization* table. A pop-up menu appears with the options **Show history**, **Traffic Group Details** and **Details by Flow**.


- Choose **Show history** to launch the *History of Utilization* window, which is a standalone version of the **History** tab.
- Choose **Traffic Groups** to display configuration details for the traffic group selected in the Link Utilization window.
- Choose **Details by Flow** to launch the *Details by Flow* window, which is a standalone version of the **Flow** tab.

The **History**, **Traffic Groups**, and the **Flow** tab are described in more detail in the following sections. These descriptions also apply to the *History* and *Details by Flow* windows.

History Tab

The **History** tab displays a chart of traffic activity for the link highlighted in the *Link Utilization* table. You specify a time range to view using the **Select Time Range** bar as described in [Creating Time Range Reports](#) on page 432. The **History** tab then displays a graph representing traffic activity during that time range, as shown in [Figure 160](#). You can also select a time range by clicking on the green cursor and then clicking on the green cursor to a point on the graph.

When selecting a time range, keep in mind that traffic data is delayed by 30 minutes in real time, and is defined by the range over which the traffic is averaged over routing elements (for example, routers and links). Traffic activity that takes place within the last 30 minutes will not appear on the chart. The start and stop times for the report appear below the chart. Current time is the time at which the routing topology is loaded, and includes routing elements over which the traffic is running. The current time is reflected by green cursor shown on screen.

To export traffic history data for use in another application, click  **Export Traffic**. RAMS exports all visible data displayed on the graph to the clipboard as a sequence of (x,y) points in ASCII:

```
x1 y1
x2 y2
...
```




Figure 160 History Tab

Traffic Groups Tab

The **Traffic Groups** tab ([Figure 161](#)) displays the current and average amount of traffic for a particular traffic group flowing over a selected link. The sum traffic total for the traffic group appears in the *Link Utilization* window, but the breakdown appears in the **Current Traffic (bps)** and **Average Traffic (bps)** columns. The **Percent** column reflects the percentage of the total traffic flowing for the traffic group.

You specify a time range to view using the **Select Time Range** bar, as described in [Creating Time Range Reports](#) on page 432. When selecting a time range, keep in mind that traffic data is delayed by 30 minutes in real time. Traffic activity that takes place within the last 30 minutes will not appear on the chart. The start and stop times for the report appear below the chart.

Traffic Groups for Link Router16 -> 192.168.116.0/24			
Traffic Group	Current Traffic (bps)	Average Traffic (bps)	Percent
Other	2.62K	0	NA


Figure 161 Traffic Groups Tab

To drill down further and obtain *History*, *Traffic Group Details*, or *Details by Flow*, right-click in a selected row for one of the traffic groups, then refer to the reports you wish to view.

Flow Tab

The **Flow** tab allows you to drill down to obtain three types of traffic information for each link displayed in the *Link Utilization* table:

- [Details by Flow Report](#) on page 441
- [Show Prefixes for Selected Flow Report](#) on page 442
- [Show Paths for Selected Flow Report](#) on page 442



Source Prefix	Destination Prefix	Exporting/Ingress Router	Traffic Group	Traffic (bps) (Current)
194.10.2.1/32	192.168.135.0/24	192.168.107.16:5	Other	2.62K

Figure 162 Flow Tab

The Details by Flow window displays details for the Source Prefix, Destination Prefix, Exporting /Ingress Router, the Traffic Group, and Traffic for the last 5 minutes.

Details by Flow Report

Highlight a link in the *Link Utilization* table, then refer to the *Details by Flow* report to view the source, destination, router, and bitrate for each traffic flow on that link. To drill down further and obtain a list of prefixes or paths for individual flows, click a flow in the *Details by Flow* table, then refer to the *List of Prefixes* and *Find or List Paths* reports respectively.

Show Prefixes for Selected Flow Report

The *Show Prefixes for Selected Flow* report table displays a list of destination prefixes and sub-prefixes for the selected flow in the *Details by Flow* report area. The table shows sub-prefixes in a tree structure underneath the parent destination prefix in the **Prefix** column, and displays sub-prefix information.

Show Paths for Selected Flow Report

The *Show Paths for Selected Flow* report table displays a list of hops the link took along a specific path. The path and hops appear in a tree format, where the hops in each path appear underneath the path in the **Path** column.

Using the BGP Peering Report Window

The *BGP Peering* report window enables you to view several reports listing traffic activity among BGP peers.


These reports appear in table format by default and include the following five areas:

- Neighbor AS Traffic report
- Destination AS Traffic report
- Transit AS report
- Exit Router Traffic report
- Distribution by Community report

Each of these reports are described in this section.

Neighbor AS Traffic Report

In the left pane of the *Traffic Reports* window, click **Neighbor** to display the *Neighbor AS Traffic* report window shown in [Figure 163](#).



Neighbor AS Traffic				
AS	Current Traffic (bps)	Average Traffic (bps)	Average Transit (bps)	Average Destination (bps)
PdLabRight (65533)	794.69K	0	0	0

1 entry

History | Destination AS/ Exit Router | Traffic Groups | Flow

Traffic Groups for Neighbor AS 65533			
Traffic Group	Current Traffic (bps)	Average Traffic (bps)	Percent
Other	794.69K	0	NA

Figure 163 Neighbor AS Traffic Report Window

The *Neighbor AS Traffic* report table lists each neighbor AS and its corresponding traffic data, including the ID of the transmitting AS (AS), the current egress traffic (**Current Traffic (bps)**), the average egress traffic (**Average Traffic (bps)**) flowing from the AS (**Total**), the amount of traffic in transit across the AS (**Average Transit**), and the average amount of traffic whose final destination is the AS (**Average Destination**).

Four tabs appear below the *Neighbor AS Traffic* table:

- **History:** Displays a graph depicting traffic processed by the neighbor AS highlighted in the table. See [History Tab](#) on page 439 for more information.
- **Destination AS/Exit Router: Contains the *Neighbor AS Traffic* and *Exit Router* report areas. The *Neighbor AS Traffic* report displays destination traffic information for the neighbor AS. The *Exit Router* report displays exit router and sub-router details for the AS specified in the report area title bar.**
- **Traffic Groups:** Displays the breakdown of the average traffic and the percentage of the traffic flow for a selected traffic group.
- **Flow:** The Details by Flow window displays details for the Source Prefix, Destination Prefix, Exporting /Ingress Router, the Traffic Group, and Traffic for the last 5 minutes.

Destination AS Traffic Report

In the left pane of the *Traffic Reports* window, click **Destination** to display the *Destination AS Traffic* report window shown in [Figure 164](#).

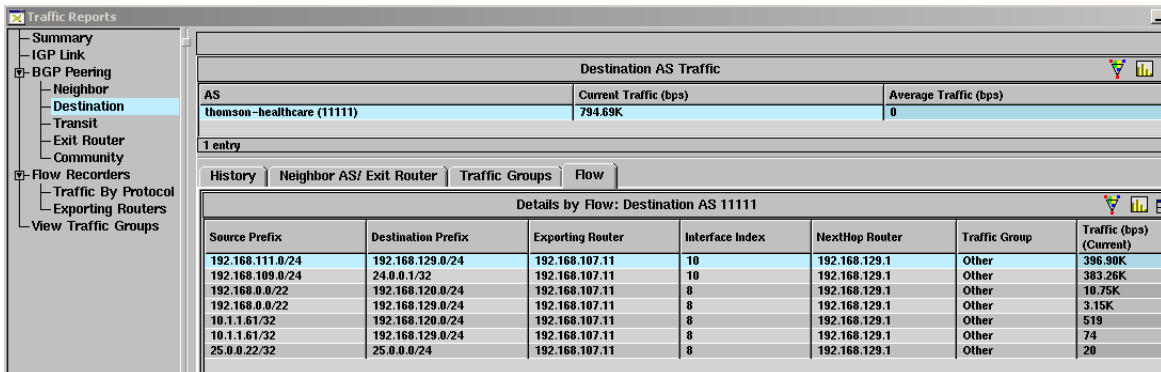


Figure 164 Destination AS Traffic Report Window

The *Destination AS Traffic* report table lists each destination AS and its corresponding traffic data, including the ID of the destination AS (**AS**), the **Current Traffic (bps)** and **Average Traffic (bps)** received at the destination AS (**Total**).

Four tabs appear below the *Destination AS Traffic* table:

- **History:** Displays a graph depicting traffic processed by the destination AS highlighted in the table. See [History Tab](#) on page 439 for more information.
- **Neighbor AS/Exit Router: Contains the *Neighbor AS Traffic* and *Exit Router* reports for the highlighted link. The *Neighbor AS Traffic* report displays destination traffic information for the neighbor AS. The *Exit Router* report displays exit router and sub-router details for the AS specified in the report area title bar.**
- **Traffic Groups:** Displays the breakdown of the average traffic and the percentage of the traffic flow for a selected traffic group.
- **Flow:** The Details by Flow window displays details for the Source Prefix, Destination Prefix, Exporting /Ingress Router, the Traffic Group, and Traffic for the last 5 minutes.

Transit AS Traffic Report

In the left pane of the *Traffic Reports* window, click **Transit** to display the *Transit AS Traffic* report window shown in [Figure 165](#).

Transit AS Traffic		
AS	Transit (bps) (Current)	Destination (bps) (Current)
PdLabRight (65533)	794.69K	794.69K
1 entry		

Traffic Groups		
Traffic Group	Traffic (bps) (Current)	Percent
Other	794.69K	794.69K

Figure 165 Transit AS Traffic Report


The *Transit AS Traffic* report table lists each transitory AS that was used by the flow as an intermediary hop before reaching the destination.

The **Traffic Groups** and **Flow** tabs appears below the *Transit AS Traffic* report table. See [Traffic Groups Tab](#) on page 440 and [Flow Tab](#) on page 441.

After you click a row in the report table, the associated list of prefixes and path information appears in the *List of Prefixes for Selected Flow* and *Find or List Paths for Selected Flow* report tables, respectively.

Exit Router Report

In the left pane of the *Traffic Reports* window, click **Exit Router** to display the *Exit Router Traffic* report window shown in [Figure 157](#).



Exit Router	NextHop Router	Current Traffic (bps)	Average Traffic (bps)
192.168.103.0/24			
192.168.103.0/24	192.168.103.5	394.43K	0
192.168.103.0/24	192.168.0.128	1.76K	0

2 top level entries, 5 total entries

History | Neighbor AS/ Destination AS | Traffic Groups | Flow

Traffic Group	Current Traffic (bps)	Average Traffic (bps)	Percent
Other	394.43K	0	NA

Figure 166 Exit Router Report

The *Exit Router* report window displays the IP address of the exit router, the nexthop router, and the current and average traffic for the exit router. You can group the table by exit router or nexthop router. When grouped, the table shows sub-routers in a tree structure underneath the parent router. In this case, when you click **Traffic (bps)** to sort traffic data, the data is sorted per parent router. Rather than sorting by individual rows, each group of sub-routers is sorted by the amount of traffic exported within the group.

Three tabs appear below the *Exit Router Traffic* table:

- **History:** Displays a graph depicting traffic processed by the exit router highlighted in the table. See [History Tab](#) on page 439 for more information.
- **Nbr AS/Dest AS: Contains the Neighbor AS Traffic and Destination AS Traffic reports.** The *Neighbor AS Traffic* report displays destination traffic information for the neighbor AS. The *Destination AS Traffic* report displays exit router and sub-router details for the AS specified in the report area title bar.
- **Flow:** Displays details for the Source Prefix, Destination Prefix, Exporting /Ingress Router, the Traffic Group, and Traffic for the last 5 minutes.

Distribution by BGP Community Report

In the left pane of the *Traffic Reports* window, click **Community** to display the *Distribution by BGP Community* report window.

Distribution by BGP Community Report		
First 2 Octets	Second 2 Octets	Out Traffic (bps) (Last 5 mins.)
57822	4315	334.42M
58538	60735	334.38M
58538	61187	334.38M
57822	4315	334.18M
57822	4311	333.73M
58538	61307	333.73M
58538	61284	333.69M
58538	61270	333.59M
58538	60826	217.48M
431 entries		

Details by Flow: BGP Community 57822:4315					
Source Prefix	Destination Prefix	Exporting Router	Interface Index	NextHop Router	Traffic (bps) (Last 5 mins.)
99.161.0.0/16	51.221.64.0/18	96.18.255.3	36	169.42.80.9	11.97M
99.161.0.0/16	50.228.64.0/18	96.18.255.3	36	169.42.80.9	11.74M
96.53.0.0/16	170.192.0.0/12	96.18.255.3	52	214.39.79.157	716.59K
96.216.0.0/16	171.173.128.0/18	96.18.255.3	49	169.42.80.9	580.27K
96.53.0.0/16	173.146.0.0/15	96.18.255.3	52	169.42.80.9	560.00K
1361 entries					

Prefixes for Flow Destination: Prefix 51.221.64.0/18				
Prefix	Router/Net	Attributes	State	Area or AS
51.221.64.0/18	96.18.255.3	AS Path: 57884 60	Up/B	ISPTraf1.BGP/AS8
1 top level entry, 4 total entries				

Paths for Flow Destination: Prefix 51.221.64.0/18					
Path	Source Node	Destination Node	Metric	Protocol	Resolved by Prefix
inglewood-					
1 top level entry, 3 total entries					

Figure 167 Community Report

The *Distribution by BGP Community* report table shows how much traffic each community receives. The report displays a list of outbound traffic to BGP communities.

The **Flow** tab appears below the report table and displays the three transit flow reports. See [Traffic Groups Tab](#) on page 440 for more information.

Using the Flow Collectors Report Window

Because some data is masked by prefix-level aggregation, only Flow Collectors — not the Flow Analyzer — can generate reports for this masked data. Examples of masked data include the following:

- Top source addresses
- Top destination addresses
- Traffic distribution per protocol
- Traffic distribution per Flow Recorder

The Flow Analyzer aggregates all reports generated by Flow Collectors. The flow recorder reports generated are averaged over the time range selected.

Note The reports generated by the flow recorder are independent from the traffic group configuration. Therefore, the Select Traffic Groups tab is disabled for these reports.

In the left pane of the *Traffic Reports* window, click **Flow Collectors** to display the Flow Collectors report window shown in [Figure 168](#).

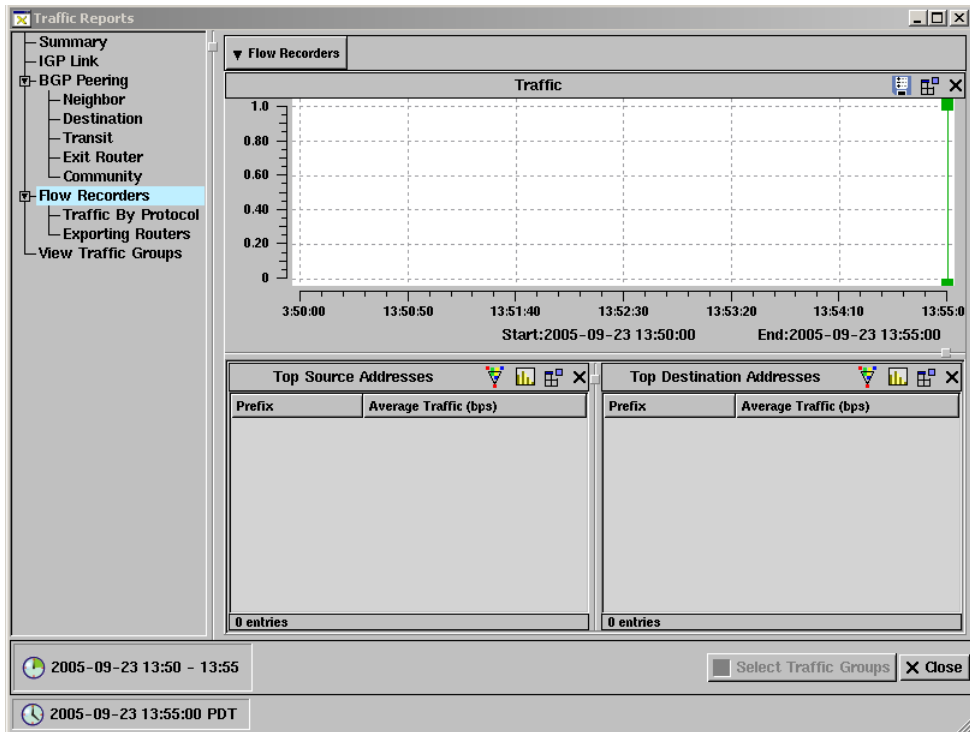


Figure 168 Flow Collectors Report

This page is divided into three report areas, displaying information about each Flow Collector in the configuration. To select which Flow Collector report to view, click the **Flow Collectors** drop-down box in the upper right corner of the window, and select a Flow Collector from the list. A selected check box appears next to the recorder you have chosen. The default setting for this drop-down menu will show all Flow Recorders in the configuration.

Time ranges for these reports is controlled by the time-range bar, where you can select the specific interval of time needed for your reports. See [Creating Time Range Reports](#) on page 432 for more information.

The three report areas are populated with data reflecting the activity of the selected Flow Collector, and include the following information:

- **Traffic:** Displays a chart of traffic activity for a specified time period. The start and end times for the report appear below the chart.

- **Top Source Addresses:** Displays the top 100 source addresses per Flow Recorder that generate the most traffic. This report appears only as a table.
- **Top Destination Addresses:** Displays the top 100 destination address prefixes that receive the most traffic. This report appears only as a table.

Traffic By Protocol Report Window

In the navigation bar of the Traffic Reports window, click *Traffic By Protocol* to display the Traffic By Protocol report window, shown in [Figure 169](#).

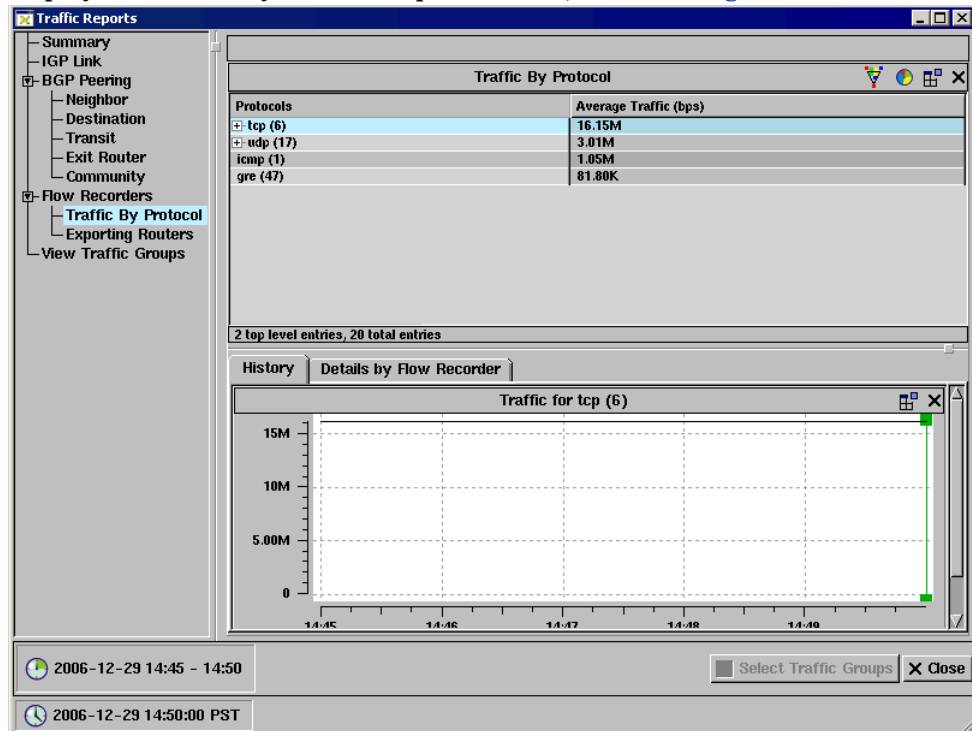


Figure 169 Traffic By Protocol

This window allows you to view flow recorder traffic breakdown by IP protocol. The report displays the average amount of traffic over the selected time range (see [page 433](#) for instructions for selecting a time range).

Note The traffic views are enabled for time frames less than or equal to one hour.

This report displays a table by default. The table displays the protocols used and the average traffic for these protocols. For TCP and UDP protocols, click on the “+” sign to view further breakdown of the traffic by port.

Note When you break the traffic down by port, and the number exceeds six thousand, the breakdown for the destination port will be displayed.

The History tab at the bottom portion of this window displays the traffic history for the entry you selected in the *Traffic By Protocol* window. The Details by Flow Recorder tab breaks down the traffic for the selected protocol across all the configured flow recorders in the network.

Exporting Routers Report Window

In the navigation bar of the *Traffic Reports* window, click **Exporting Routers** to display the report window, shown in [Figure 170](#).

The screenshot shows a window titled "Distribution by Flow Exporter" with a filter set to "Any". The table below lists the data shown in the window.


Flow Exporter Name	Flow Exporter Address : Interface Index	Average Traffic (bps)
[- Router11		
Router11	192.168.107.11:1	26.22K
Router11	192.168.107.11:7	89.25K
Router11	192.168.107.11:9	3.02M
Router11	192.168.107.11:8	17.47K
Router11	192.168.107.11:6	22.50K
Router11	192.168.107.11:10	3.97M
Router11	192.168.107.11:5	114.06K
Router11	192.168.107.11:2	512
Router11	192.168.107.11:13	0
Router11	192.168.107.11:15	0

Figure 170 Exporting Routers Report

This report appears as a table by default. The parent exporter appears at the top of the table with sub-exporter information underneath the parent exporter in tree format. Sub-exporter information is organized in three columns:


- Flow Exporter Name — Displays the name of the flow recorder.
- Flow Exporter Address: Interface Index— Displays the IP address and the interface index of the flow recorder.
- Average Traffic — Displays the average traffic flow per bps.

Note If an exporter is not present in the topology, the **Flow Exporter Name** column displays Unknown.

You can click the  **Highlight** button to highlight exporting routers in yellow on the topology map. Click the button again to unhighlight exporting routers.

Using the View Traffic Groups Window

In the navigation bar of the Traffic Reports window, click **View Traffic Groups** (Figure 171) to display configuration details for the traffic groups you configured on the web.



Traffic Groups Configuration						
Order	Traffic Group Name	Source Prefixes	Destination Prefixes	Source AND/OR Destm	Protocols	Traffic Classes
20	ipsec-1st	192.168.133.0/24	241.0.0/24	99 OR 103	17	ANY
21	Teladidicom	192.168.107.0/24	195.10.1.0/24	1024 OR 21	6	ANY
23	vand-been	192.168.119.0/24	194.10.102.0/24	17 OR 80	6	ANY
24	corpnetc	192.168.107.0/24	195.10.0.0/24	16 OR 25	6	ANY
25	portTraff	192.168.116.0/24	194.10.107.0/24	85 OR 99	17	ANY
27	ipTraffic	192.168.107.0/24	24.0.0.10/32	ANY OR ANY	ANY	ANY
28	calif	192.168.135.0/24	194.10.0.0/16	ANY OR ANY	ANY	ANY
29	classC_rtr3	192.168.135.0/24	194.10.0.0/16	ANY OR ANY	6	ANY
30	komp-traffic	192.168.135.0/24	192.168.110.0/24	ANY OR ANY	1	ANY
31	komp-ppp	192.168.103.0/24	195.10.1.0/24	ANY OR ANY	1	ANY
32	us-hyd_DeletedOn_200	192.168.116.0/24	194.10.2.1/32	ANY OR ANY	ANY	ANY
33	us-blr_DeletedOn_200	192.168.135.0/24	194.10.102.1/32	ANY OR ANY	ANY	ANY
34	FDS3_DeletedOn_2006	ANY	ANY	ANY OR ANY	48	ANY
35	New Group_DeletedOn	ANY	ANY	ANY OR ANY	ANY	ANY
36	BharT_DeletedOn_200	192.168.110.0/24	192.168.103.0/24	4112 OR 4113	6	ANY
37	local-hyd_DeletedOn_2	192.168.116.0/24	194.10.103.0/24	85 OR 99	17	ANY
38	Other	ANY	ANY	ANY OR ANY	ANY	ANY

Figure 171 View Traffic Groups Window

For information on creating Traffic Groups, see [Creating Traffic Groups](#) on page 94.

Using the Traffic Menu

In addition to traffic reports, the *Traffic* menu also includes four menu options for viewing and changing links:

- **List Flows:** Displays all traffic flows in the *List of All Flows* window.
- **Link Utilization:** Displays all links that utilize traffic in the *Link Utilization* window. See [Link Utilization Report](#) on page 435.
- **Link Utilization Prediction:** Displays the predicted value of traffic set for a specified future date and time. For more information, see [Chapter 7, “Network Planning.”](#)
- **List Exporters:** Displays all traffic flows distributed by each NetFlow exporter in the *Distribution by Flow Exporter* window. See [Exporting Routers Report Window](#) on page 451.
- **Set Interface Capacities:** Allows you to open, edit, and save interface capacity files and apply those capacities to the current topology from the *Set Interface Capacities* window. For more information, see [Chapter 7, “Network Planning.”](#)

13 Alerts

RAMS alerts are used to obtain immediate notification of any changes to routes, network traffic or routers based upon configurable thresholds. Alerts can be sent as SNMP traps to an SNMP-based network management system integrated with other network operations, or logged to a Syslog file.

By default, alerts are not enabled in RAMS. You must set and configure alerts that identify specific variables. These variables allow you to customize alerts and to turn alerts on and off for various types of routing and traffic events.

Alerts are set globally and affect all routing and traffic events. For example, if you are monitoring a single or multi-area OSPF network, the alerts you select are applied to all the OSPF routing events for all the OSPF areas being monitored.

There are three groups of alerts, one for IGP protocols (this group includes OSPF, IS-IS, and EIGRP), one for BGP protocols, and one for traffic. You can configure alerts on the Flow Analyzer (traffic alerts) and on Route Recorders (IGP and BGP alerts). In a deployment with multiple Route Recorder units and a centralized Modeling Engine, IGP and BGP alerts should be configured per Route Recorder rather than on the Modeling Engine. When alerts are configured on each recorder, traps are sent to the destination local to the area or protocol being monitored, which allows you to respond quickly to time-critical events. Alerts requiring information from more than one recorder, however, should be enabled on the Modeling Engine. For example, the Route Change alert should be configured on the Modeling Engine when the source and destination routers are being seen by different recorders.

This chapter provides an overview of Route Explorer alerts and explains how to configure these alerts using the web browser interface. It also describes the data that is contained in the various alerts and how alerts can be used to improve service availability and customer satisfaction.

Configuring an SNMP Server

Before setting any alerts, you must configure the global settings on the *Alerts* page of the Route Recorder or Flow Analyzer. Additionally, you must download the *Packet Design* Structure of Management Information (SMI) and the *pdRouteExplorer* Management Information Base (MIB) to the SNMP Agent software.

Note The Route Explorer SNMP daemon 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 Use a browser to connect to the Route Explorer *Home* page of the Route Recorder or Flow Analyzer.
- 2 Type the user name and password, and then click **Login**.
The *Home* page appears.
- 3 Click **Administration**, then click **Alerts** in the left navigation bar.
The *Alerts* page appears as shown in [Figure 172](#).

Alerts

Configure SNMP

SNMP Trap Destination

Community String

Send Test SNMP Alert

Configure Remote Syslog

Remote Syslog Address

Test Syslog Configuration

Figure 172 Alerts Page

The *Alerts* page on a Flow Analyzer shows only Traffic alerts, as shown in [Figure 173](#).

Traffic Route Correlation Alert Configuration

Alert Notification Options

SNMP Trap Remote Syslog Both None

Traffic Route Correlation threshold (bits/sec)

Figure 173 Alerts Page for Traffic Route Correlation Alert

- 4 At the bottom of the RAMS *Alerts* page, click the **Packet Design SMI** link.
The Packet Design SMI appears in the browser window.
- 5 From the *File* menu on your browser, choose **Save As** and save the SMI file in the directory containing the SNMP Agent software.

- 6 Click the **Packet Design Products MIB** link at the bottom of the Route Explorer *Alerts* page.

The Packet Design MIB appears in the browser window.

- 7 From the *File* menu on your browser, choose **Save As** and save the MIB file in the directory containing the SNMP Agent software.

- 8 Click the **PD RAMS MIB** link at the bottom of the RAMS *Alerts* page.

The Packet Design RAMS MIB appears in the browser window.

- 9 From the *File* menu on your browser, choose **Save As** and save the MIB file in the directory containing the SNMP Agent software.

To configure the SNMP Manager address, perform the following steps:

- 1 Log in as an administrator and navigate to the *Alerts* page if you are not already there.
- 2 In the *SNMP Trap Destination* text box, type the IP address or fully qualified domain name of the machine with the SNMP Agent software.
- 3 In the *Community String* text box, type the community string.
- 4 Click **Configure SNMP**.

This configures the SNMP manager address.

Configuring the Remote Syslog and RAMS

Before configuring RAMS to send alerts, you must set up the *syslogd* on the servers 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.
- The *syslogd* may default to starting with no remote reception capability. In order to receive remote syslog messages, you may need to restart *syslogd* with the *-r* option. If in doubt, look in `/var/log/messages` for a “*syslogd started < remote reception >*” log entry.

After you have configured the remote syslog, you can configure the syslog settings on the Route Explorer *Alerts* page.

To configure syslog settings on the Route Recorder and Flow Analyzer, perform the following steps:

- 1 Log in as an administrator on the Route Recorder or Flow Analyzer and click **Administration**.
- 2 Click **Alerts** in the left navigation bar.
- 3 In the *Remote Syslog Address* text box, type the IP address of the remote system log.
- 4 Click **Configure Remote Syslog**.

Setting an Alert

This section provides an example of how a specific alert is set and an agent is configured. In this instance, the agent is HP Network Node Manager (NNM) and a prefix flap alert is set.

Note The method used to configure all IGP alerts is virtually the same. The only difference is that some alerts contain an additional threshold configuration and some do not.

To set a Prefix Flap alert and configure an SNMP trap, perform the following steps:

- 1 On the Route Recorder *Home* page, log in as an administrator and click **Administration**.
- 2 Under the **Alerts** heading in the left navigation bar, click **IGP**, then **Prefix Flap**.

The *Prefix Flap Alerts* page appears, as shown in [Figure 174](#).

- 3 In the Prefix Flap Threshold Configuration section, type the number of events that need to occur in the specified time for an alert to occur.

For example, to send an alert if a minimum of ten events occur in thirty seconds, type 10 in the *Events* text box and 30 in the *Time Scale* text box.
- 4 Click **Configure Threshold**.
- 5 (Optional) Type network and mask details in the Watch List Configuration section. If configured, the watch list limits alerts to events occurring in the specified network.
- 6 Select the desired alert type (SNMP Trap in this example) in the Alert Notification Options section, and then click **Submit Notification Options**.
- 7 Click **Update**.

Go to Master Home **Administration** Recorder Configuration Reports Portal Support

Alerts

- Destinations
- SNMP Test
- IGP
- Adjacency Lost
- Adjacency Established
- Adjacency Flap
- Prefix Change
- Prefix Origination Change
- Prefix Flap**
- Route Change
- Routing Event
- Excess Churn
- Peer Change
- BGP
- Application**
- VNC Configuration
- Layout Backgrounds
- Queries
- Diagnostics**
- System Diagnostics

Prefix Flap Alert Configuration

Alert Notification Options

SNMP Trap
 Remote Syslog
 Both
 None

Prefix Flap Threshold Configuration

Number of Events

Time Scale (seconds)

Watch List Configuration [Optional]

Network Mask

Figure 174 Prefix Flap Alerts Page

The next step is configuring the SNMP trap that will display the alarm in the network management system. In this example, the network management system is HP Network Node Manager (NNM).

To configure the SNMP trap in NNM, perform the following steps:

- 1 In the *HP* main window, select **Event Configuration**.
- 2 Select the RAMS MIB.
- 3 Select the **Prefix Flap** event.
- 4 From the *Edit* menu, select **Modify Event**.
- 5 Enable the trap display as an Error Alarm.
- 6 Set the severity of the alarm
- 7 Set the event log pop-up notification.
- 8 Click **OK**.
- 9 From the *Edit* menu, select **Save**.

Now, when a prefix flap is detected, Route Explorer sends the trap to NNM. Then, NNM displays the corresponding alarm in the alarm log and displays the alarm pop-up notification.

You can configure a watch list for most alerts. A watch list allows you to specify routers or prefixes to monitor. In addition, you can configure a threshold for some alerts, like *Prefix Flap*. Threshold configuration is described later in this chapter.

Understanding Alert Formats

Since all alerts have a different format, all of the possible alert formats are not covered in this guide. This section presents one example of a SNMP Trap alert and one example of a system log alert.

SNMP Trap Alert Format

Figure 175 displays an example of a Prefix Origination Change alert sent from Route Explorer to HP Network Node Manager.

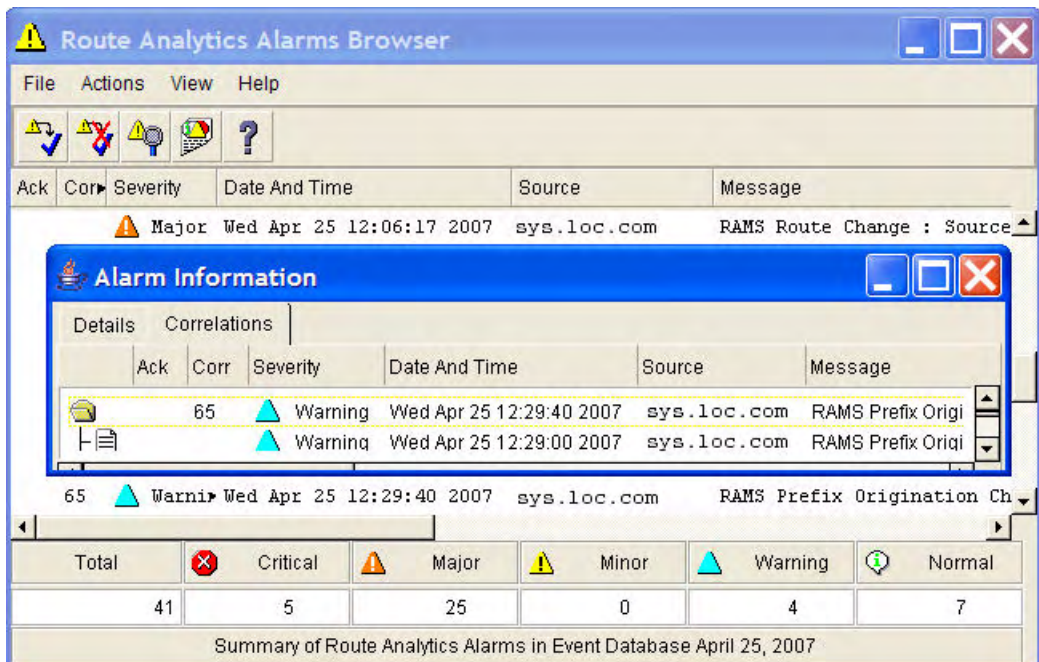


Figure 175 HP Network Node Manager Alarm

Syslog Alert Format

All syslog messages are sent with facility *Local0* and severity *Alert*. This is not currently configurable.

The following example illustrates an alert for a Prefix Origination Change alert sent from Route Explorer to a Syslog file, and assumes that the Syslog host has inserted timestamps.

```
Nov 13 13:23:05 dhcp-168-0-28 RouteAnalyzer[1264]: - Prefix  
Origination Change 25.0.0.2/32
```

```
Nov 13 13:23:05 dhcp-168-0-28 RouteAnalyzer[1264]: -  
rexPrefixFlap: rexPrefixNet = 25.0.0.2, rexPrefixMask =  
255.255.255.255
```

```
Nov 13 13:23:05 dhcp-168-0-28 RouteAnalyzer[1264]: -  
Adjacency Lost: Router(192.168.103.11) to  
PseudoNode(192.168.103.2)
```


Configuring IGP Alerts

This section describes the types of IGP Alerts that you can activate on a Route Recorder and, in some cases, the configurable threshold variables. The Route Recorder *Alerts* page, shown in [Figure 172](#), has two groups of alerts. The alerts listed under **IGP** are described in this section.

Adjacency Lost Alert

When this alert is enabled, an alert is sent when an adjacency is lost between two routers in the network. This alert is normally used to monitor key adjacencies, like adjacencies between areas, adjacencies between a router and a server farm, and point-to-point links to remote areas on the backbone. When the alert is generated, the following information is included in the alert:

- Time and date the adjacency was lost.
- The IP address of the source and destination routers where the adjacency was lost.

When an Adjacency Lost alert is received, you can open the routing topology map to quickly view what effect the lost adjacency had. The *History Navigator* window is also useful for viewing routing events that may have taken place immediately preceding the adjacency loss to further analyze the problem.

To configure an Adjacency Lost alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Adjacency Lost**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, a watch list limits alerts to events occurring in the specified network.
 - a In the *Source Router* and *Destination Router* fields, type the interface IP addresses for the two ends of the adjacency.

In the case of a pseudonode, use the IP address of the LAN interface of the Designated Router.

Note Source and Destination Router ID fields will display if a router ID is detected by the appliance. These fields will be populated with the Router ID detected by the appliance.

- b Choose among the operators None, And or Or.

For example, if you choose Or, an alert will be generated if either end of the adjacency is lost.

- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Established Peer Alert

When this alert is enabled, an alert is sent when a peering is established between the appliance and a router in the network. When the alert is generated, the following information is included in the alert:

- Time and date the peering was lost.
- The IP address routers where the peer was established.

Note Source and Destination Router ID fields will display if a router ID is detected by the appliance. These fields will be populated with the Router ID detected by the appliance.

When an Established Peer alert is received, you can open the routing topology map to quickly view what effect the established peer had.

To configure an Established Peer alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **IGP** heading, click **Established Peer**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.

Adjacency Flap Alert

This alert is triggered when the adjacency between two routers is flapping. RAMS lets you define the adjacency flap threshold by specifying the number of transitions per second that will initiate an alert. When the alert is generated, the following information is included in the notification:

- The time and date of the adjacency flap.
- The IP address of the flapping route.

When an Adjacency Flap alert is received, you can quickly view what effects the flapping route has on the network in the real-time routing topology map. The *History Navigator* window is also useful for viewing routing events that may have taken place immediately preceding the alert to determine how long the problem has been occurring.

Note This alert is not effective for EIGRP, because the derivation of link-state events, such as an adjacency state change, from EIGRP events takes multiple seconds.

To configure an Adjacency Flap alert, perform the following steps:

- 1 On the RAMS Alerts page, under the **IGP** heading, click **Adjacency Flap**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 Type the adjacency flap threshold in transitions per second in the *Adjacency Flap Threshold Configuration* text box.
- 5 Click **Configure Threshold**.
- 6 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified adjacencies.
 - a In the *Between* field, select either **Interface** or **Router** from the drop-down menu.

Select **Router** if you want to specify the endpoint(s) of an adjacency as any of the interface(s) of the router(s).

Select **Interface** if you want to specify the endpoint(s) of an adjacency in terms of interface(s) address(es).
 - b In the *Source* and *Destination* fields, type the interface or router IP addresses for the source and destination of the adjacency.

In the case of a pseudonode, use the IP address of the LAN interface of the Designated Router.

Note Source and Destination Router ID fields will display if a router ID is detected by the appliance. These fields will be populated with the Router ID detected by the appliance.
 - c In the Operation field, choose among the operators **None**, **And** or **Or**.

Choose **And** to receive an alert if an adjacency is lost at which has at its source and destination the designated source and destination endpoints (for example, routers or interfaces).

Choose **Or**, to receive an alert if an adjacency is lost which has as its source the specified source endpoint or as its destination the specified destination endpoint.

Choose **None** if you want to receive an alert if an adjacency is lost which has at its source the specified source endpoint or as its destination the specified destination point.

- d In the Destination field, enter the IP address of the destination adjacency.
- 7 Click **Update**.
- 8 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Change Alert

This alert is sent when a prefix attribute, such as a metric or prefix type, changes. Prefix attribute changes normally occur during scheduled maintenance. This alert is used to carefully monitor that the correct changes have been made and that changes affecting critical routes continue to receive the appropriate service levels. When the alert is generated, the following information is included in the notification:

- The time and date of the prefix change.
- The prefix where the change occurred.

After receiving this alert, you can go to the router in question to see what metric change took place. Review the service log to see if the prefix change was planned, and if so, verify the change.

To configure a Prefix Change alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Prefix Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.

- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Origination Change Alert

This alert is sent when a prefix has been withdrawn or is newly advertised. A prefix withdrawal indicates that services have been disrupted. This may present a number of problems, such as dropped packets, slowed service, and so on. Being notified of a prefix origination change allows significant time savings compared to the process of telnetting to each router to determine which prefix was withdrawn or added.

When the alert is generated, the following information is included in the notification:

- The time and date the prefix was withdrawn or advertised.
- The prefix ID.

After receiving a Prefix Origination Change alert, you can open the *Online Update Monitor* in RAMS and view the Events panel to locate the prefix in question. You can then drill down and find out the router, neighbor/prefix, attributes, and area for the prefix in question.

To configure a Prefix Origination Change alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **IGP** heading, click **Prefix Origination Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Flap Alert

This alert indicates that a prefix is flapping, providing early indication of a service becoming unavailable. By pinpointing the flapping prefix, RAMS saves significant time and effort in resolving the problem. During the setup of this alert, you define the prefix flap threshold by entering the number of transitions that must occur within a number of seconds to initiate an alert. When the alert is generated, the following information is included:

- The time and date of the prefix flap.
- The prefix that is flapping.

When a Prefix Flap alert is received, use the Tools menu in RAMS to select **List Prefixes**. Here, you can obtain a list of all network prefixes, then select the prefix in question to identify the routers that are sourcing the prefix. This provides quick router identification so you can troubleshoot the problem further.

To configure a Prefix Flap alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **IGP** heading, click **Prefix Flap**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Route Change Alert

This alert indicates that the route between an explicitly configured source and destination has changed, providing early indication of service degradation or even an outage. Configure this alert on the centralized Modeling Engine. You can define a set of source and destination router pairs to monitor their end-to-end routes. When the alert is generated, the following information is included in the notification:

- The time and date of the route change.

- The endpoints of the route that has changed.

When a Route Change alert is received, you can investigate the source of the problem by opening the *History Navigator* window and performing an event analysis for the time period in question. The *Event Analysis* window gives you greater detail about the route change, showing routers, links, and prefixes. Alternatively, you can run a *Prefix Origination Change* report or a *Prefix Withdrawn* report for the time period in question to further identify which routers changed sourcing.

To configure a Route Change alert, perform the following steps:

- 1 On the RAMS Alerts page, under the **IGP** heading, click **Route Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 Type the watch list configuration details, which are the source and destination routers for the routes being monitored.
 - The source is a router IP address, which can be the router ID.
 - The destination must be a routable IP address (not all router IDs are routable addresses).
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Routing Event Alert

This alert is triggered when a routing event is received after a *quiet period*. You can use this alert to catch anomalies or unexpected events occurring on an otherwise quiet network. If a particular router is experiencing problems, you can add that router to a watch list to ensure it is operating correctly. The Routing Event alert lets you identify a potential problem early, often before it causes a more serious problem.

All routing events are included in the total routing event count. You can set the threshold for the number of routing events per second that trigger the Routing Event alert. When no routing events have occurred for the threshold number of seconds, the next routing event to occur triggers the alert.

When the alert is generated, it includes the time and date when the first routing event occurred after a quiet period.

To configure a [Routing Event](#) alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Routing Event**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 In the *Hold Time* text box, type the event threshold.
- 5 Click **Configure Threshold**.
- 6 If an SNMP trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Excess Churn Alert

This alert is sent when the number of routing messages transmitted in the network reaches a threshold you define. Enable this alert to notify you when activity levels are higher than normal, which could be an early indicator of a problem.

Keep the following points in mind when configuring this alert:

- The Excess Churn alert is based upon number of *routing messages*, not *routing events*.
- A routing message may contain information about multiple events.
- Routing messages may be redundant. For example, multiple routers may report the same event.
- All routing messages (not Hellos) are included in the total message count. For EIGRP, this alert is based directly on the distance-vector EIGRP Update messages, not the derived link-state events.
- The number of routing messages per second that are generated during normal operation vary depending on factors such as the size of the network and the protocols being used. EIGRP, for example, typically generates more routing messages than other IGP protocols.

Determining an appropriate threshold for this alert may require some experimentation. Generally, you should set the threshold at the upper bound of the expected number of events that occur during normal network operation. If the number of alerts generated is too high (or low), adjust the threshold.

After receiving this alert, you should open the *History Navigator* and view the Events List to obtain more detail about the sources of the unusual activity, and to determine if further investigation is required.

To configure the *Excess Churn* alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Excess Churn**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 In the *Excess Churn Threshold Configuration* text box, type the event threshold.
- 5 Click **Configure Threshold**.
- 6 If an SNMP trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Peer Change Alert

This alert is sent if the adjacency between RAMS and a peer router has changed by going up or down. You might want to enable this alert at all times so you are notified when RAMS is unable to monitor certain parts of the network. When the alert is generated, the following information is included in the notification:

- The time and date of the peer change.
- The IP address of the new peer.
- The protocol of the peering.

When this alert is received, it is usually easiest to log into the peered router to check on the link status of the link in question.

To configure a *Peer Change* alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Peer Change**.
- 2 Select the desired alert notification option.

- 3 Click **Submit Notification Options**.
- 4 If an SNMP trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Configuring BGP Alerts

The Route Recorder *Alerts* page, shown in [Figure 172](#), has two groups of alerts. This section describes the alerts listed under **BGP** and the configurable threshold variables for these alerts.

BGP Route Flap Alert

This alert is triggered when a route is repeatedly updated or withdrawn from the BGP table. When you enable this alert, you can receive early indication of a critical service becoming unavailable or of a degradation in performance.

You must identify the following values to enable this alert:

- A flap period (for example, 10 minutes)
- A threshold (the number of up and down cycles in a flap period that will trigger the alert, for example, 5).
- The network and mask to add to the watch list.

When the alert is generated, the following information is included in the notification:

- The time and date when the flap reached the threshold.
- The prefix length.
- The peer router.
- The peer autonomous system (AS).
- The next hop.
- The next hop AS.
- The current state (active/withdrawn).
- The time and date.

When you receive a Route Flap alert, you can run the *BGP Prefix Events Detail* report for the prefix in question to determine the duration of the flapping problem.

BGP Lost Redundancy Alert

This alert is triggered when a redundant link is lost, which can be an early indicator of a reduction in service level or higher cost for service. The ability to immediately become aware of a lost redundant link can also be critical in maintaining acceptable levels of customer satisfaction and SLAs. You can set the threshold for this alert, which is the minimum number of next-hops that will trigger the alert. When the number of next hops equals the threshold (calculated by baseline), the alert is triggered.

Note The baseline is calculated based on an average redundancy for a prefix sampled over the last 7 days.

When the alert is generated, the following information is included:

- The prefix.
- The length of the redundancy.
- The source AS number.
- The baseline period.
- The number of baseline next-hops.
- The current number of next-hops.
- The time and date when the prefix redundancy was lost.

After receiving this alert, you can run the *Route Distribution Detail* report to view next hops from the prefix in question to see which peers withdrew routes and why.

Note In networks with two or more Internet connections to different ISPs, the BGP Lost Redundancy Alert may produce an excessive number of alerts. This alert may not be as effective in such a network environment.

BGP Lost Peer Alert

This alert is initiated when the connection between RAMS and a BGP peer is lost. The Lost Peer alert could indicate that the router is down, traffic going to this router is being blackholed. Alternatively, the router may still be up, but RAMS is no longer able to listen to it and consequently cannot manage it.

When the alert is generated, the following information is included:

- The lost peer IP address.
- Time and date the peer was lost.

When you receive a Lost Peer alert, you can run the *Route Distribution Detail* report and begin looking into the possible cause.

BGP Prefix Flood and BGP Prefix Drought Alerts

These alerts are initiated when a routing table size increases or decreases by a significant amount, relative to the calculated baseline. This type of alert could indicate that a router is misconfigured or possibly that you've experienced a security breach. You have the flexibility to define a threshold for this alert.

Baselines are calculated by the amount of time the prefixes are up. For example, baselines are stable as long as the prefixes are up 80% of the time for the first 24 hours they are running. If the prefixes fall below 80%, they fall out of the baseline and an alert is generated. This pattern continues as each day passes (for example, baselines need to stay up 80% for 48 hours, and 80% of 72 hours, etc)

Note The baseline is calculated based on an average sampled over a one week period (last 7 days).

When the alert is generated, the following information is included:

- The baseline period
- The baseline table size.
- The current table size.
- The percentage delta in the table size
- The time and date.
- The peer IP address.

When you receive a Prefix Flood or Drought alert, you can telnet to the router in question to review the routing table, verify the problem, and begin looking into the possible cause.

BGP Acquired Redundancy Alert

The BGP Acquired Redundancy Alert is triggered when the number of next hops for the prefix becomes greater than the threshold for the BGP Lost Redundancy Alert. In some cases, the BGP Acquired Redundancy Alert could indicate that a problem previously identified by a BGP Lost Redundancy Alert has been corrected.

The following information is included in this alert:

- The prefix.
- The peer IP address.
- The source AS (or the number associated with the AS where the link originated).
- The number of baseline next-hops.
- The current number of next-hops.
- The time and date when the prefix redundancy was acquired.

After you receive this alert, you would typically run the *Route Distribution Detail* report to view next hops from the prefix in question to see which peers added routes.

Note In networks with two or more Internet connections to different ISPs, the BGP Acquired Redundancy Alert may produce an excessive number of alerts. This alert may not be as effective in such a network environment.

BGP Established Peer Alert

When this alert is enabled, an alert is sent when a peering has been established between two routers in the network. This alert is normally used to monitor key peerings. When this alert is generated, the following information is included in the alert:

- Time and date peering was lost
- IP address of the down router

When an Established Peer alert is received, you can then investigate why the router went down, and the affect this had on the network.

BGP AS Path Longer Alert

Each BGP prefix has an AS path associated with it. When a prefix begins advertising an AS path that is longer than previously advertised, the AS Path Longer alert is initiated. This alert indicates that traffic for this prefix will now take a longer path to arrive at the destination, possibly causing delays in service.

Note The BGP AS Path Longer alert is of limited utility in this release. Be aware that when you enable this alert, you may cause RAMS to produce an excessive number of alert messages.

The following information is included in this alert:

- The prefix.
- The peer IP address.
- The source AS number.
- The time and date when the alert was triggered.

After you receive this alert, you would typically run the *Route Distribution Detail* report to view next hops by AS.

BGP Down to One Path and BGP Down to Zero Paths Alerts

These alerts perform similarly to the BGP Lost Redundancy Alert, with the following exceptions:

- The threshold for the BGP Down to One Path alert is pre-set to 1.
- The threshold for the BGP Down to Zero Paths alert is pre-set to 0.

Note The BGP Down to One Path and Down to Zero Paths alerts are of limited utility in this release. Be aware that when you enable these alerts, you may cause RAMS to produce an excessive number of alert messages.

Configuring OSI IS-IS Alerts

ES Neighbor Flap Alert

This alert indicates that an ES Neighbor is flapping, providing early indication of a service becoming unavailable. By pinpointing the flapping ES Neighbor, RAMS saves significant time and effort in resolving the problem. During the setup of this alert, you define the ES Neighbor flap threshold by entering the number of transitions that must occur within a number of seconds to initiate an alert. When the alert is generated, the following information is included:

- The time and date of the ES Neighbor flap.
- The ES Neighbor that is flapping.

When a ES Neighbor Flap alert is received, use the Tools menu in RAMS to select **List OSI Prefixes**. Here, you can obtain a list of all OSI prefixes, then select the ES Neighbor in question to identify the Intermediate Systems that are sourcing the ES Neighbor. This provides quick Intermediate System identification so you can troubleshoot the problem further.

To configure a ES Neighbor Flap alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **OSI** heading, click **ES Neighbor Flap**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

ES Neighbor Change Alert

This alert is sent when an ES Neighbor attribute, such as a metric changes. This alert is used to carefully monitor that the correct changes have been made and that changes affecting critical routes continue to receive the appropriate service levels. When the alert is generated, the following information is included in the notification:

- The time and date of the ES Neighbor change.
- The ES Neighbor where the change occurred.

After receiving this alert, you can go to the Intermediate System in question to see what metric change took place. Review the service log to see if the ES Neighbor change was planned, and if so, verify the change.

To configure a ES Neighbor Change alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **OSI** heading, click **ES Neighbor Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

ES Neighbor Origination Change Alert

This alert is sent when an ES Neighbor has been withdrawn or is newly advertised. An ES Neighbor withdrawal indicates that services have been disrupted. This may present a number of problems, such as dropped packets, slowed service, and so on. Being notified of an ES Neighbor origination change allows significant time savings compared to the process of telnetting to each Intermediate System to determine which ES Neighbor was withdrawn or added.

When the alert is generated, the following information is included in the notification:

- The time and date the ES Neighbor was withdrawn or advertised.
- The ES Neighbor ID.

After receiving an ES Neighbor Origination Change alert, you can open the *Online Update Monitor* in RAMS and view the Events panel to locate the ES Neighbor in question. You can then drill down and find out the Intermediate System, neighbor, attributes, and area for the ES Neighbor in question.

To configure an ES Neighbor Origination Change alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **OSI** heading, click **ES Neighbor Origination Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Neighbor Flap Alert

This alert indicates that a prefix neighbor is flapping, providing early indication of a service becoming unavailable. By pinpointing the flapping prefix neighbor, RAMS saves significant time and effort in resolving the problem. During the setup of this alert, you define the prefix neighbor flap threshold by entering the number of transitions that must occur within a number of seconds to initiate an alert. When the alert is generated, the following information is included:

- The time and date of the prefix neighbor flap.
- The prefix neighbor that is flapping.

When a Prefix Neighbor Flap alert is received, use the *Tools* menu in RAMS to select **List OSI Prefixes**. Here, you can obtain a list of all OSI prefixes, then select the Prefix Neighbor in question to identify the intermediate systems that are sourcing the prefix neighbor. This provides quick intermediate system identification so you can troubleshoot the problem further.

To configure a Prefix Neighbor Flap alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **OSI** heading, click **Prefix Neighbor Flap**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Neighbor Change Alert

This alert is sent when a prefix neighbor attribute, such as a metric or prefix neighbor type, changes. Prefix neighbor attribute changes normally occur during scheduled maintenance. This alert is used to carefully monitor that the correct changes have been made and that changes affecting critical routes continue to receive the appropriate service levels. When the alert is generated, the following information is included in the notification:

- The time and date of the prefix neighbor change.
- The prefix neighbor where the change occurred.

After receiving this alert, you can go to the intermediate system in question to see what change took place. Review the service log to see if the prefix neighbor change was planned, and if so, verify the change.

To configure a Prefix Neighbor Change alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **OSI** heading, click **Prefix Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.

- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Prefix Neighbor Origination Change Alert

This alert is sent when a prefix neighbor has been withdrawn or is newly advertised. A prefix neighbor withdrawal indicates that services have been disrupted. This may present a number of problems, such as dropped packets, slowed service, and so on. Being notified of a prefix neighbor origination change allows significant time savings compared to the process of telnetting to each intermediate system to determine which prefix neighbor was withdrawn or added.

When the alert is generated, the following information is included in the notification:

- The time and date the prefix neighbor was withdrawn or advertised.
- The prefix neighbor ID.

After receiving a Prefix Neighbor Origination Change alert, you can open the *Online Update Monitor* in RAMS and view the Events panel to locate the prefix neighbor in question. You can then drill down and find out the intermediate system, neighbor, attributes, and area for the prefix in question.

To configure a Prefix Neighbor Origination Change alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **OSI** heading, click **Prefix Neighbor Origination Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 (Optional) Type the watch list configuration details. If configured, the watch list limits alerts to events occurring in the specified network.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

OSI Route Change Alert

This alert indicates that the route between an explicitly configured source and destination has changed, providing early indication of service degradation or even an outage. Configure this alert on the centralized Modeling Engine. You can define a set of source and destination router pairs to monitor their end-to-end routes. When the alert is generated, the following information is included in the notification:

- The time and date of the route change.
- The endpoints of the route that has changed.

When an OSI Route Change alert is received, you can investigate the source of the problem by opening the *History Navigator* window and performing an event analysis for the time period in question. The *Event Analysis* window gives you greater detail about the route change, showing routers, links, and prefixes. Alternatively, you can run a *Prefix Origination Change* report or a *Prefix Withdrawn* report for the time period in question to further identify which routers changed sourcing.

To configure a Route Change alert, perform the following steps:

- 1 On the *RAMS Alerts* page, under the **IGP** heading, click **Route Change**.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 Type the watch list configuration details, which are the source and destination routers for the routes being monitored.
 - The source is a System ID.
 - The destination must be an NSAP address.
- 5 Click **Update**.
- 6 If an SNMP Trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

Configuring Traffic Alerts

This section describes the types of traffic alerts that you can set on the Flow Analyzer and the configurable threshold variables for these alerts. You can combine each of these alerts with routing analysis to show traffic routing changes before and after the alarm is generated.

The *Traffic Alerts* section lists two types of alerts on the navigation bar: **Link Utilization** and **Route Correlation**. These alerts are described in this section.

Traffic Link Utilization Alert

Link utilization is the amount of traffic on a link or traffic group divided by the capacity of the link. When the Link Utilization alert is enabled, the appliance monitors the amount of aggregate traffic or traffic group(s) on each link across the network, and generates an alert when the utilization of a link drops below or exceeds the user-defined thresholds. The threshold is specified in percent utilization on the *Alerts* page. This alert is normally used to discover unusual patterns in link utilization levels, which may indicate the presence of a virus, worm, DOS (Denial of Service) attack or other harmful attacks on the network.

The following information is included in the generated alert:

- The IP address of the source router on the link
- The IP address of the destination router on the link.
- The traffic in Mbps

The following describes the buttons shown on this page:

- **Reset** — Clears the screen of input values.
- **Submit** — Submits information.
- **Show All** — Shows all configured Traffic Groups.
- **Show Enabled** — Shows enabled Traffic Groups.
- **Show Disabled** — Shows disabled Traffic Groups.

To configure the Traffic Link Utilization alert, perform the following steps:

- 1 On the RAMS Alerts page, under the **Traffic** heading, click **Link Utilization**.

The Link Utilization Alert Configuration page appears, as shown in [Figure 176](#).

Link Utilization Alert Configuration

Alert Notification Options

SNMP Trap Remote Syslog Both None

Submit Notification Options

Aggregate Alert Threshold Parameters

Utilization		BitRate	
Low (%)	High (%)	Low (Mbps)	High (Mbps)
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

Submit

Traffic Groups Alert Threshold Parameters

Name	Utilization		BitRate	
	Low (%)	High (%)	Low (Mbps)	High (Mbps)
Group One	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

Reset Submit Show All ShowEnabled Show Disabled

Figure 176 Link Utilization Alert Configuration page

- 2 In the *Alert Notification Options* section, select the desired alert notification option.

- 3 Click **Submit Notification Options**.
- 4 In the *Aggregate Alert Threshold Parameters* text boxes, type the percentage of utilization and bit rate allowed on a link before an alert is initiated.
- 5 Click **Submit**.
- 6 In the *Traffic Groups Alert Threshold Parameters* section, you can view and edit all, enabled, or disabled traffic groups Utilization and BitRate thresholds. All configured traffic groups will appear in this box. See [Chapter 3, “Administration”](#) for more information.

Note If no traffic groups are configured, traffic will be listed as “other” in the Traffic Group Alert Threshold Parameter box.
- 7 If an SNMP trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

When you receive a Link Utilization alert, use the **Traffic** tab in the *History Navigator* window to view activity that may have taken place immediately preceding the increase in traffic link utilization to further analyze the problem.

Note The traffic value of the correlator is now reported as a DisplayString with text describing the event.

Traffic Route Correlation Alert

RAMS receives notification of standard routing events such as prefix announcements and peering establishments. When a routing event is associated with a change in traffic flow greater than 100 bps on a link, RAMS compares the change in flow with the user-defined threshold for the alert. If the change in traffic level exceeds the threshold, RAMS generates an alert. The threshold is defined on the *Alerts* page.

The following information is included in the Traffic Route Correlation alert:

- The time and date the alert was triggered.
- The correlator ID for the associated routing event.
- The number of hops changed.
- The amount traffic has changed.

To configure the Traffic Route Correlation alert, perform the following steps:

- 1 On the RAMS *Alerts* page, under the **Traffic** heading, click **Route Correlation**.
The Traffic Route Correlation Alert page appears.
- 2 Select the desired alert notification option.
- 3 Click **Submit Notification Options**.
- 4 In the *Traffic Route Correlation threshold* text box, type the maximum change in flow allowed on a link before an alert is generated.
- 5 Click **Submit Query**.
- 6 If an SNMP trap was selected, follow the steps in [Setting an Alert](#) on page 460 to configure the SNMP trap.

When you receive a Route Correlation alert, view the *Traffic and Routing Events Correlation* window in the *Traffic Reports* window to view the events occurring around the time of the alert.

Note The traffic value of the correlator is now reported as a DisplayString with text describing the event.

A License Key Details

This appendix describes the license keys associated with RAMS. For information on how to apply license keys to an appliance, see [Chapter 2, “Configuration and Management.”](#)

Each license key shows an expiration date for the applied license, or “--” for a permanent license. In a typical installation, all licenses will show the same expiration date. However, if you purchase an incremental license, that license will have a different expiration date. A license with a short-duration expiration date is often provided during product evaluations. When the license expires, data recording is disabled, protocol configuration is disabled, and RAMS displays a warning message.

Table 24 Available License Keys

License Name	Description
Protocol Licenses	Required for a RAMS unit to monitor particular protocols. If the license for a given protocol (OSPF, IS-IS, EIGRP, BGP, or MPLS VPN) is not enabled, an instance of that protocol cannot be created on the <i>Recorder Configuration</i> page. If you attempt to add or edit a protocol instance that does not have a license, you will receive an error message indicating the lack of license.
Route Recorder	Required for a RAMS unit to act as a Route Recorder, which stores information obtained from the licensed routing protocols (OSPF, IS-IS, EIGRP, BGP, or MPLS VPN).
Flow Collector	Required for a RAMS unit to collect traffic flow data, which is then analyzed by the Flow Analyzer.
Flow Analyzer	Required for a RAMS unit to generate reports using data collected from Flow Collectors and Route Recorders.

GUI	Required for a RAMS unit to accept VNC or X Window System connections to display the graphical user interface.
RAMS Traffic SPI	Applied to the Modeling Engine to enable viewing of traffic and routing data and traffic reports.
Route Analyzer Alerts and Reports	Required for a RAMS unit to act as a Route Analyzer. Route Analyzer generates reports using data collected from Route Recorders. The licenses for Alerts and Reports are typically enabled in a RAMS unit that is configured as a Route Recorder.
Database Server	Required for a RAMS unit that records routing or traffic data to act as a database server. Machines that are configured to act as database clients can connect to this machine.
Database Client	Required for a RAMS unit to act as a database client which can then connect to a database server to access routing and traffic databases.
Master Capability	Required for a RAMS unit to be designated as the master appliance in a distributed configuration environment. The master is then used to configure a set of clients.
Router Count	Sets the number of routers you are able to monitor with RAMS. When the number of routers you are monitoring exceeds the number of routers allowed by the license, RAMS displays a warning message. Sets the number of routers you are able to monitor with Traffic Explorer. When the number of routers you are monitoring exceeds the number of routers allowed by the license, Traffic Explorer displays a warning message. If you exceed the number by more than 10, then the Open Topology operation in the GUI will abort unless you select a subset of the databases containing a smaller number of routers.

MPLS VPN Prefix Count	Sets the number of VPN prefixes you are able to monitor. If you exceed the number supported by the license, RAMS displays a warning message. Count Sets the number of VPM prefixes you are able to monitor. If you exceed the number supported by the license, RAMS displays a warning message. If you exceed the number by more than 10%, then the Open Topology operation in the GUI will abort unless you deselect all VPN databases.
Software Update	Indicates whether the unit can obtain software updates from HP. Software Update is enabled with all licenses, including the temporary license that can be activated on a new unit.Updates the software to the most current version.

B Protocol Compliance

This appendix lists the protocol compliance details for RAMS.

RAMS supports the following implementations:

- OSPF:
 - RFC 2328, OSPF Version 2
 - RFC 3101, The OSPF Not-So-Stubby Area (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
- 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>).

In RAMS 5.20, 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

A

Acquired Redundancy alert, 478

Activity by AS report, 349

Activity by Peer report, 350

Activity Summary report, 348

add

- client, 35

- node, 293

- peering, 293

- prefix, 293

- traffic flow, 293

Adjacency Established alert, 466

adjacency flap

- alert, 467

- EIGRP, 467

Adjacency Flap alert, 466

Adjacency Lost alert, 465

administration interface, selecting, 42

Administration pages, 25, 85

administrative domain

- creating, 48

- deleting, 66

- top-level, 48

administrator default user name and password, 29

alarms

- clearing, 393

- VPN, 389

alerts

- BGP, 475

- configuration page, 456

- IGP, 465

- setting, 460

- syslog, 463

- traffic, 486

- VPN, 389

alias

- BGP AS, 60

- interface, 42

all events list, 261

analysis

- edits, 295

- event, 256

- history, 231

- options, 212

- root cause, 235, 236

animation, 238

- button, 232

- clock, 239

- fast, 232

- graph, 239

- mode, 232, 240

- playing, 240

- replaying, 241

- saving, 241

- window, 238

anomalies, network, 190

- appliance
 - reboot, 144
 - reset, 145
 - shutdown, 144
- application interface, 23
- Archival Configuration and Remote Storage page, 119
- archiving, 119
 - automatic, 119
 - configuration, 119
 - enable, 120
 - end time, 122
 - filename format, 121
 - frequency of, 119
 - manual, 121
 - restoring data, 122
 - start time, 122
- areas
 - contiguous, 55
 - multiple, 55
- AS
 - BGP, 60
 - configuring for EIGRP protocol instance, 58
- AS Path Longer alert, 479
- AS Reachability report, 355
- asymmetric paths, 415
- authentication
 - MD5, 81
 - OSPF, 57, 67, 81
 - user, 29
- auto-hide
 - nodes, 216
 - options, 216
- automatic labels, 213
- autonomous system
 - see AS

B

- Backup and Restore page, 107
- backup file
 - creating, 107, 109
 - deleting, 112
 - downloading, 111
 - FTP, 115
 - restoring, 111, 112
 - restoring from remote server, 114
 - saving, 110
 - transferring, 115
 - uploading, 111
- backup on unit's disk, 109
- bandwidth
 - and delay
 - EIGRP, 182
 - setting, 182
 - changing, 311, 319
 - setting, 319
- bar charts and tables, 256
- Baseline AS Reachability report, 355
- Baseline Redundancy by Prefix report, 355
- Before-N-After comparison, RIB, 253

BGP

- adding prefix, 302
- add peering, 299
- alerts, 475
 - Acquired Redundancy, 478
 - AS Path Longer, 479
 - Down to One Path, 479
 - Down to Zero Paths, 479
 - Lost Peer, 476
 - Lost Redundancy, 476
 - Prefix Drought/Prefix Flood, 477
 - Route Flap, 475
- aliases, AS, 60
- autonomous system (AS), 60
- configuring, 60
- display AS, 175
- distribution by community report, 446
- peering report, 443
- protocol
 - Before-N-After comparison, 254
 - instances, 60
 - RIB browser, 247
- reports, 345
 - accessing, 347
 - Activity by AS, 349
 - Activity by Peer, 350
 - Activity Summary, 348
 - AS Reachability, 355
 - Baseline AS Reachability, 355
 - Baseline Redundancy by Prefix, 355
 - Prefix Event Detail, 351
 - Prefix Reachability, 356
 - Redundancy by Prefix, 354
 - Route Distribution Detail, 353
 - Route Flap, 350

BGP alerts

- Established Peer, 478

BGP Peering Report window, 443

BGP Reports, 24

bitrate, changing, 311

bounding box, 178, 182, 185

browser

- supported, 27

button

- Add Node, 293
- Add Peering, 293
- Add Prefix, 293
- Add Traffic Flow, 293
- Analysis, 231, 235
- Analyze, 328
- Analyze Edits, 295
- Animation, 238
- Change Prefix, 293
- Clear, 263
- Clear All, 394
- Click to Estimate, 323
- Down Node, 293
- Down Peering, 293
- Down Prefix, 293
- Events, 232, 238
- Execute One Event, 263, 270
- Export, 295
- Export Traffic, 431
- Fast Animate, 232
- Fast Step, 232
- Go to End, 239
- Go to Start, 239
- Graphs, 231
- Hide, 251, 268
- History Navigator, 231
- Import, 295, 343
- List Links, 192
- List Prefixes, 194, 195
- List Routers, 191
- Make Master, 34
- Online Update, 263
- Planning Reports, 295
- Prefixes, 238
- Re-apply All, 32
- Relinquish Master Role, 36, 37
- Restore All Edits, 295
- Restore Now, 123
- Select Time Range, 430, 439, 440
- Set, 391
- Show, 251, 268

- Show Current Event, 263, 270
- Start Execution, 263, 269
- Step, 232
- Stop, 232
- Stop Execution, 263, 269
- Time Range, 232, 262, 268
- Topology Map toolbar, 163
- Update, 232
- Users, 136

C

- callback, enable, 83
- capacity, 296
 - available, 331
 - custom, 336
 - interface, 336
- CE router, 381
- Changed Metrics report, 363
- Cisco router loopback interface, 82
- class, changing user, 138
- clearing events, 263
- CLI Access, 136
- client, 21, 25
 - adding, 35
 - removing, 37
- client list, 34
- clock
 - animation, 239
 - icon, 262, 268
 - setting, 38
- cluster event analysis, 256
- cold spots, 418
- collector
 - traffic flow, 19

- color by
 - area, 169
 - bitrate, 169
 - traffic, 169
 - utilization, 169
- columns, hidden, 429
- community traffic, 296, 334
- comparison
 - RIB Before-N-After, 253
 - RIB Browser, 252
- components
 - data flow, 20
 - RAMS, 19
- configuration
 - additional tasks, 79
 - administration interface, 42
 - alerts, 460
 - BGP, 60
 - client, 25
 - database
 - deleting, 104
 - renaming, 105
 - DHCP, 41
 - distributed, 21
 - Flow Analyzer, 75
 - Flow Collector, 68
 - for recording, 47
 - FTP server, 134
 - GRE tunnels, 79
 - hierarchy, 47, 48
 - hostname, 41
 - loopback interface, 82
 - master, 25
 - option card, 44
 - options dialog box, 212
 - overview, 26
 - preferences, 212
 - Route Recorder, 47, 61
 - SMI and MIB, 456
 - SNMP
 - manager address, 458
 - server, 456
 - traps in NNM, 461
 - static route, 42
 - support access, 83
 - syslog, 459
 - traffic instance, 68
 - tree, 47
 - users, 30, 136
 - VNC, 99
 - VPN, 386
- controls
 - History Navigator, 230
 - playback, 239

- cookies, 29
 - reports and, 360
- cost, path, 210
- cursor
 - dragging to change time view, 231
 - History Navigator, 231

D

- database
 - administration page, 102
 - archiving, 120
 - backing up, 107
 - deleting, 104
 - in green letters, 158
 - online vs. offline, 103
 - recording, 158
 - renaming, 105, 122
 - restored filename, 123
 - restoring, 107
 - restoring archived, 122
 - selecting, 76
 - trimming, 106
 - unarchived, 120
- Database Client, 492
- Database Server, 492
- Data Configuration page, 116
- data source
 - choosing, 116
- default, factory, 144
- designated router, 177
- Design mode, 23, 165, 230, 292, 297, 430
- destination AS, 444
 - traffic, 436
- destination traffic, 333
- Details by flow window, 439
- DHCP, 41

- diagnostics
 - ping, 143
 - topology, 172
 - traceroute, 143
- distributed configuration, 21
 - access privileges for, 137
 - applying license keys, 31
 - archival settings for, 120
 - archiving data for, 119
 - choosing data source for, 116
 - clients in, 25
 - configuring route recorder in, 22
 - deleting a database within, 104
 - designating master and client roles for, 33
 - IGP and BGP reports for, 24
 - master unit in, 25
 - updating software for, 124
 - updating software with internet access for, 125
 - updating software without internet access, 126
- distributed configuration master and client units in, 21
- Distribution by BGP Community report, 446
- DNS
 - displaying node name, 178
 - server, 41
- download
 - MIB, 456
- downloading
 - software updates, 124
- Down to One Path alert, 479
- Down to Zero Paths alert, 479

E

- Edit Flows window, 294
- Edit menu, 168

edits

- analysis, 295
- analyzing, 328
- exporting, 295
- importing, 295
- incremental restore, 320
- list of, 296, 328
- removing, 320
- restoring, 295, 320
- topology, 297, 328
 - exporting, 340
 - importing, 343
- undoing, 295, 320
- viewing, 320

EIGRP

- adjacency flap ineffective for, 467
- and excess churn alert, 472
- bandwidth and delay, 182
- distances, mismatched, 222
- highlighted path cost, 210
- multiple areas, ASs, 55
- multiple interfaces, coverage and, 57
- path cost, defined, 210
- prefix types, 265
- protocol instance, configuring ASs, 55
- topology diagnostics, 219
- update event, 265, 269, 472

enable archival, 120

enable XML RPC queries, 387

Established Peer alert, 478

estimate

- total traffic, 322
- traffic, 321

event, 232, 238

adjusting time range, 268

analysis, 256

clearing, 263

details, 263

BGP, 267

EIGRP, 265

IS-IS, 263, 265

OSPF, 263, 265

execute, 270

executing, 263, 269

start, 263, 269

stop, 263, 269

filtering, 267

graph, 234

highlighting associated nodes, 267

list all, 261

listing all, 221, 261

matching time, 268

online, 221

showing current, 263, 270

Excess Churn alert, 472

exit router, 196

highlighting by, 196

report, 446

traffic, 296, 333, 436

export

edits, 295

log, 141

traffic, 440

exporting routers report, 451

expression

definitions, 283

filter, 280

regular, 187

syntax, 282

F

factory defaults, 144

- filename, restored database, 123
- file upload, 278
- filter, 303
 - definitions, 283
 - events, 267
 - expression examples, 283
 - expressions, 280
 - syntax, 282
 - using, 280
- Filter by, 295, 303
- find
 - paths, 210
- flapping, prefix, 237
- Flapping Links report, 364
- Flow Analyzer, 19, 33, 491
 - configuration, 75
- flow analyzer
 - settings, 77
- Flow Collector, 19, 33, 448, 491
 - configuration, 68
 - report window, 448
- flows
 - details, 439
- flows, list, 341
- flow server, adding, 313
- Flow tab, 438, 441
- FTP
 - file transfer, 134
 - file upload, enabling, 134
 - file uploads, 134
 - server configuration, 134

G

- graph
 - animation, 239
 - button, 231
 - displaying, 234
 - events, 234
 - History Navigator, 234
 - links, 234
 - network events, 178, 182
 - online update monitor, 171
 - routers, 234
 - routes, 234
 - VPN participation, 395
 - VPN reachability, 394
- GRE tunnels, 19, 79

H

- Help menu, 173
- hidden columns, 429
- hidden nodes, 169, 179, 186, 187, 216
- hierarchy
 - configuration, 47, 48
 - topology, 168
- History mode, 23, 165, 230, 430

- History Navigator, 23, 227
 - and EIGRP topology errors, 219
 - button, 231
 - controls, 230
 - cursor, 231
 - example of use, 271
 - fast step, 232
 - graphs, 234
 - mode, 163, 230
 - options, 215
 - overview, 227
 - playback controls, 232
 - replay, 241
 - status bar, 230
 - update, 232
 - zoom timeline, 233

History tab, 438, 439

Home page, 27

hot spots, 418

HP Software, 461

HTTP POST, 34

I

ICMP redirects, 43

IGP

- add peering, 300
- add prefix, 304
- alerts, 465
 - Adjacency Established, 466
 - Adjacency Flap, 466
 - Adjacency Lost, 465
 - Excess Churn, 472
 - Peer Change, 473
 - Prefix Change, 468
 - Prefix Flap, 460, 470
 - Prefix Origination Change, 469
 - Route Change, 470
 - Routing Event, 471

display area, 175

link report, 438

protocols

- Before-N-After comparison, 253

- RIB browser, 247

reports, 357

- accessing, 360

- Changed Metrics, 363

- Flapping Links, 364

- Network Churn, 366

- Network Events Summary, 361

- New Prefixes, 368

- New Routers and Links, 369

- Prefixes Withdrawn, 378

- Prefix List, 371

- Prefix Origination Changes, 374

- Prefix Origination from Multiple Sources, 376

IGP Link Report window, 438

IGP Reports, 24

import

- edits, 295

- time series data, 278

inaccessible routers

- EIGRP, 221

- list of, 172, 221

- inactive node, 217
- incremental restore, 320
- information panel, 177
 - link, 180, 182
 - link traffic, 182
 - node, 177, 189
 - traffic, 179
- interface
 - administration, 42
 - alias, 42
 - application, 23
 - capacity, 336
 - removing from protocol instance, 66
 - status, link details, 181
 - web, 25
- interface capacity, 336
- Internet Explorer, 27
- invisible links, finding, 172, 224
- IP address, 42

K

- keys, license, 491

L

- labels
 - node, 169
- layout
 - managing topology map, 159
 - saving, 168
 - topology
 - default, 215
- leaf trimming, 187
- Legend, 164, 173
- legend, 173

- license
 - applying, 31
 - install new key, 32
 - key, 31
- license key, 491
 - database client, 492
 - database server, 492
 - flow analyzer, 491
 - flow collector, 491
 - master capability, 492
 - modeling engine, 492
 - MPLS VPN prefix count, 493
 - protocol, 491
 - route analyzer, 492
 - router count, 492
 - Route Recorder and GUI, 491
 - software update, 493
- License Update page, 32

link

- details, 177
- finding invisible, 172, 224
- graph, 234
- hot and cold, 418
- IGP report, 438
- information panel, 180, 182
- invisible, in EIGRP, 224
- list all, 192
- report, 400
- set metric, 319
- traffic, 296, 330
- traffic estimation, 321
- utilization, 296, 322, 329
- utilization report, 435

- list
 - all events, 261
 - client, 34
 - edits, 328
 - flows, 341
 - inaccessible routers, 172
 - mismatched distances, 172
 - paths, 210
 - topology errors, 172
- log
 - export as plain text, 141
 - viewing, 140
- logging in, 29
- login page, 29
- loopback, 266
 - interface, 57, 82
 - interface, on Cisco router, 82
- Lost Peer alert, 476
- Lost Redundancy alert, 476

M

- Management Information Base
 - see MIB

- map
 - layout, 157
 - Routing Topology, 23, 157
 - symbols, 161
 - topology, 292

- master, 21, 25
 - assigning, 34
 - license key, 492
 - password, 34
 - relinquish, 36
 - removing client, 37

- master access password, 34

- Master Capability, 492

- master-client, 21

- MD5 authentication, 81

- MED, 249, 250, 272

- menu bar
 - RAMS window, 167

- metric
 - changing, 319
 - setting, 319
 - setting EIGRP, 319
 - setting OSPF, 319

- MIB

- MIB, downloading, 456

- mismatched distances
 - EIGRP, 222
 - listing, 172, 222

- mode
 - animation, 240
 - Design, 23, 165, 230, 292, 297, 430
 - fast stepping, 240
 - History, 23, 165, 230, 430
 - History Navigator, 163
 - Online, 23, 165, 230, 430
 - Step, 232, 240
 - switching, 230

- Modeling Engine, 19, 33, 492

- modes
 - switching, 230

- Mozilla, 27

- MPLS, 381

- MPLS VPN Prefix Count, 493

- MRTG, 278
 - file formats, 278

- multi-exit discriminator, 249, 250, 272

- Multiprotocol Label Switching
 - see MPLS

- Multi Router Traffic Grapher, see MRTG

N

NBMA networks, 161

neighbor AS, 435, 443
traffic, 332

NetFlow, 19

Netscape, 27

network

- anomalies, 190
- configuration, 40
- DHCP, 41
- events, 178
- events, graph, 182
- inventory, 191
- large, viewing, 196
- zooming in, 185

Network & Interface Configuration page, 40

Network Churn report, 366

Network Events Summary report, 361

Network Node Manager, 460, 461

New Prefixes report, 368

New Routers and Links report, 369

NNM, 460, 461

node information panel, 189

nodes

- adding, 293
- auto-hide, 216
- bringing down, 293, 314
- bringing up, 315
- details, 177
- displaying DNS name, 178
- hidden, 169, 179, 186, 187, 216
- hiding failed, 169
- highlighting, 267
- inactive, 217
- information panel, 177, 207
- labels, 169, 213
- re-hide, 179, 186
- restoring visibility, 187
- selecting, 184
- selecting by dragging, 184
- trim, 169, 186, 187
- trim leaf, 169
- unhide, 186, 187
- unhide all, 169
- working with, 184

NSAP address, 213, 215

NTP server, 38

O

online/offline databases, 103

online events, 221

Online mode, 23, 165, 230, 430

online update, 263

- button, 263
- monitor, 171

option cards, configuring, 44

options

- analysis, 212
- auto-hide, 216
- configuring, 212
- miscellaneous, 214
- visualization, 212

OSPF

- area ID format, 65
- authentication, 57, 67, 81
- contiguous areas, 53
- loopback interface, 57
- protocol instance, 57

overload router, 178, 368

overwrite layouts, 112, 115

P

participation, graph, 395

password

- changing user, 138
- Master Access, 34

Path Reports, 24

Path reports, 403

paths, 209

- asymmetric, 415
- cost, 210
- finding, 210
- hot and cold, 418
- listing, 209, 210
- reports, 415

Peer Change alert, 473

peering

- adding, 293, 298
- adding BGP, 299
- adding IGP, 300
- BGP report, 443
- bringing down, 293, 315

peers, status of, 165

PE router, 381

ping, 143

planning

- exporting data, 340
- importing data, 340
- network, 291
- reports, 294
- toolbar, 292

Planning Reports window, 23, 294

Planning toolbar, 292

playback

- animation mode, 240
- controls, 232, 239
- fast step, 240
- set step, 216
- Step mode, 240

prefix

- adding, 293, 301
 - BGP, 302
 - IGP, 304
 - manually, 302
 - with filter, 303
- bringing down, 293, 316
- button, 238
- changing, 293, 308
- finding route by, 209
- flapping, 237
- list of, 194, 195
- report, 399
- shifting, 236
- types, EIGRP, 265

Prefix Change alert, 468

Prefix Drought/Prefix Flood alert, 477

Prefixes Withdrawn report, 378

Prefix Event Detail report, 351

Prefix Flap alert, 460, 470

Prefix List report, 371

Prefix Origination Change alert, 469

Prefix Origination Changes report, 374

Prefix Origination from Multiple Sources
report, 376

Prefix Reachability report, 356

previous version, revert to, 128

protocol

BGP, 60

instance, 57

instance, deleting, 66

license keys, 491

link-state, 18

multiple, in history navigator, 228

supported, 19

Protocol Licenses, 491

pseudonode, 177, 180, 189, 217

R

RADIUS, 59

RAMS

components of, 19

operation of, 18

reverting to previous version, 128

using, 22

web server, 27

reachability

and participation index, 383

graph, 394

re-apply all, 32

reboot, 144

rebooting the system, 144

recorder

NetFlow, 19

routing, 19

Recorder Configuration page, 47

recording, configuration, 47

Redundancy by Prefix report, 354

RegEx, 187, 405

regular expression, 187, 405

re-hide nodes, 179, 186

remote server, adding, 113

replay an animation, 241

reports

BGP, 24, 345, 347

cookies and, 360

deleting instance, 78

IGP, 24, 357, 360, 361

Links, 400

Path, 24, 403

Planning, 23, 294

Prefixes, 399

Routers, 400

time range, 432

Traffic, 23, 425, 434

VPN, 392

reset to factory defaults, 145

restore

archived data, 122

databases, 122

edits, 295

from backup on unit's disk, 109

nodes, 187

Restore from Archive page, 122

revert to previous version, 128

RIB

Before-N-After comparison dialog, 253

browser, 247

visualization window, 242

root cause analysis, 235, 236

Round Robin Database (RRDtool), 278

- route
 - finding, 189
 - finding by prefix, 209
 - redistributing, 220
 - static, configuring, 42
 - summarization, and invisible links, 219, 225
- Route Analyzer, 492
- Route Change alert, 470
- Route Distribution Detail report, 353
- Route Flap
 - alert, 475
 - report, 350
- router
 - CE, 381
 - count, 492
 - exit, 196, 446
 - exit, traffic, 333, 436
 - exporting, 451
 - finding, 189, 191
 - graph, 234
 - inaccessible, EIGRP, 221
 - list all, 191
 - login for EIGRP, 59
 - naming conventions, hiding by, 186
 - PE, 381
- Router Count, 492
- Route Recorder, 19, 33
 - configuration page, 47
 - status indicator, 165
- Route Recorder and GUI license, 491
- route reflectors, 61
- Routers report, 400
- routes graph, 234
- routing
 - hierarchies, viewing, 184
 - history, 227
 - loop, 221

- Routing Event alert, 471
- Routing Information Base
 - see RIB
- Routing Topology Map, 23, 157
 - menu, 167
 - opening, 167
 - selecting, 158, 167
 - toolbar, 163

S

- saving
 - animation, 241
 - topology layout, 168
 - visualization, 246
- Selection menu, 185
- select time range, 439, 440
- server
 - database, 492
 - DNS, 41
 - flow, 313
 - FTP, 134
 - NTP, 38
 - remote, 113
 - SNMP, 456
 - VNC, 99
- service availability, alerts and, 455
- session timers, 139
- set interface capacity, 336
- set metrics, 319
- shifting, prefix, 236
- shutdown
 - options, 144
 - page, 144
- SMB, enabling, 113
- SMI
 - downloading, 456

- SNMP
 - configuring in NNM, 461
 - manager address, configuring, 458
 - server configuration, 456
 - trap, 460
- software
 - reinstalling previous version, 128
 - updating, 493
- Software Update
 - license, 493
 - page, 124
- SSH, 43, 59, 83, 136
- state changes, viewing, 263
- static route, configuration, 42
- status bar
 - History Navigator, 230
 - Routing Topology Map, 165
- status icon
 - peers, 165
 - recorders, 165
- status indicator
 - colors, 165
- status table, 71
- Structure of Management Information
 - see SMI
- summarization boundaries, and EIGRP, 219, 225
- support access, configuring, 83
- symbols, 161
- syntax
 - filter expression, 282
- syslog
 - alert format, 463
 - configuration, 459

- system
 - reboot, 144
 - shutdown, 144
- System Diagnostics page, 143
- system settings, viewing, 46

T

- TACACS, 59
- time
 - adjusting range, 268
 - moving, 268
 - selecting range, 439, 440
 - setting, 38
 - setting range, 232, 262, 268
- Time and Date page, 38
- timeline, zooming, 233
- time range reports, 432
- time series data, 278
 - correlating, 278
 - importing, 278
 - uploading, 278
- Tools menu, 170

- topology
 - diagnostics, 172, 219
 - EIGRP, 172
 - submenu, 172
 - editing, 297
 - edits, importing, 343
 - errors
 - EIGRP, 219
 - list, 172, 219
 - hierarchy, 168
 - history, 227
 - map, 23
 - map, colors, 214
 - menu, 167
 - opening, 297
 - saving a layout, 168
 - selecting, 158
 - status bar, 165
 - window, 163, 168, 175
- Topology Map, 292
- Topology menu, 292
- traceroute, 143
- traffic
 - alerts, 486
 - Link Utilization, 486
 - Route Correlation, 488
 - color by, 169
 - colors, setting, 214
 - community, 296, 334
 - configuration status, 71
 - curve, 324
 - data delay, 231, 232, 263
 - deleting instance, 74
 - destination, 333
 - destination AS, 436, 444
 - estimating, 321, 322
 - exit router, 296, 333, 436, 446
 - export, 440
 - information panel, 179
 - instance, 68
 - link, 321, 330
 - link information panel, 182
 - loading, 235
 - neighbor AS, 332, 435, 443
 - reports, 321, 425, 434
 - viewing, 426
 - transit AS, 445
- Traffic and Routing Events Correlation report, 434
- traffic flow
 - adding, 308, 314
 - changing bitrate, 311
 - deleting, 313
 - editing, 294, 310
 - list of, 341
 - moving, 312
- Traffic Groups
 - creating, 94
- Traffic menu, 453
- Traffic Reports, 23, 321
- transit AS, 445

- trap, SNMP, 460
- tree, configuration, 47
- trekking, traffic, 321
- trim
 - leaf, 187
 - leaf nodes, 169
 - nodes, 169, 186, 187
- tunnels
 - GRE, 19, 79
 - VPN, 381

U

- unhide nodes, 186, 187
- Units page, 34
- upload
 - FTP, 134
 - to appliance, 278
- User Administration page, 30, 136
- user interface
 - application, 23
 - web, 22
- user name and password, default, 29
- users
 - adding, 30, 138
 - changing class, 138
 - changing password, 138
 - CLI access class, 136
 - deleting, 138
 - removing, 138
- utilization
 - link, 329
 - report, 435

V

- View Configuration page, 46

- viewer, 22
 - VNC, 30
 - X Window System, 30
- View menu, 168
- visualization
 - options, 212
 - RIB, 242
 - saving, 246
- VNC, 22, 30
 - configuration, 99
 - server, 99
 - start, stop server, 99

VPN

- alarms, 389
- alerts, 389
- clearing alarms, 393
- configuration, 386
- explorer window, 384
- History Navigator window, 401
- participation graph, 395
- protocol, 381
- reachability
 - and participation index, 383
 - graph, 394
- removing alarms, 391
- reports, 392
 - Alarms, 392
 - Participation, 395
 - Prefixes, 399
 - Reachability, 394
 - Summary, 392
- tunnels, 381

W

- web browser
 - supported, 27
- web browser, cookies, 29
- web interface, 22, 25
- web server, 27

X

X Window System, 22, 30
 evaluation license, 30

Z

zoom

 History Navigator timeline, 233
 in, 185
 out, 185

