Steelhead Management Console User's Guide

Version 6.5.2
August 2011

riverbed®

The content of this manual is furnished on a RESTRICTED basis and is subject to change without notice and should not be construed as a commitment by Riverbed Technology, Incorporated. Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in Subparagraphs (c) (1) and (2) of the Commercial Computer Software Restricted Rights at 48 CFR 52.227-19, as applicable. Riverbed Technology, Incorporated assumes no responsibility or liability for any errors or inaccuracies that may appear in this book.

riverbed

Riverbed Technology
199 Fremont Street
San Francisco, CA 94105

Phone: 415.247.8800
Fax: 415.247.8801
Web: http://www.riverbed.com

Part Number
712-00007-13

Contents

Preface

Welcome to the *Management Console User's Guide.* The Management Console makes managing your Steelhead appliance simpler through a Web browser interface. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, hardware and software dependencies, and contact information. It includes the following sections:

- "About This Guide" on page 13
- "Product Dependencies and Compatibility" on page 14
- "Riverbed Services Platform 32-Bit and 64-Bit Support" on page 15
- "SNMP-Based Management Compatibility" on page 16
- "CMC Compatibility" on page 16
- "Antivirus Compatibility" on page 17
- "Additional Resources" on page 17
- "Contacting Riverbed" on page 18

About This Guide

The *Steelhead Management Console User's Guide* describes how to configure and monitor the Steelhead appliance using the Management Console.

Audience

This guide is written for storage and network administrators familiar with administering and managing WANs using common network protocols such as TCP, CIFS, HTTP, FTP, and NFS.

This guide assumes you are familiar with using the Riverbed Command-Line Interface as described in the *Riverbed Command-Line Interface Reference Manual*.

Document Conventions

This manual uses the following standard set of typographical conventions.

Convention	Meaning
italics	Within text, new terms and emphasized words appear in italic typeface.
boldface	Within text, CLI commands and GUI controls appear in bold typeface.

Courier	Code examples appear in Courier font. For example:
	`login as: admin`
	`Riverbed Steelhead`
	`Last login: Wed Jan 20 13:02:09 2010 from 10.0.1.1`
	`amnesiac > enable`
	`amnesiac # configure terminal`

< >	Values that you specify appear in angle brackets. For example: interface <ipaddress>
[ ]	Optional keywords or variables appear in brackets. For example: ntp peer <addr> [version <number>]
{ }	Required keywords or variables appear in braces. For example: {delete <filename> | upload <filename>}
|	The pipe symbol represents a choice to select one keyword or variable to the left or right of the symbol. (The keyword or variable can be either optional or required.) For example: {delete <filename> | upload <filename>}

Product Dependencies and Compatibility

This section provides information about product dependencies and compatibility. It includes the following sections:

- "Hardware and Software Dependencies" on page 14
- "Riverbed Services Platform 32-Bit and 64-Bit Support" on page 15
- "SNMP-Based Management Compatibility" on page 16
- "CMC Compatibility" on page 16
- "Antivirus Compatibility" on page 17
- "Additional Resources" on page 17

Hardware and Software Dependencies

The following table summarizes the hardware and software requirements for the Steelhead appliance.

Riverbed Component	Hardware and Software Requirements
Steelhead appliance	19-inch (483 mm) two or four-post rack.
Steelhead Management Console, Steelhead Central Management Console	Any computer that supports a Web browser with a color image display. The Management Console has been tested with Mozilla Firefox version v2.x, v3.0.x and Microsoft Internet Explorer version v6.x and v7.x. **Note:** JavaScript and cookies must be enabled in your Web browser.

Riverbed Services Platform 32-Bit and 64-Bit Support

The following table describes the models that support 32-bit and 64-bit Virtual Machines (VMs). The 64-bit guest VMs (such as Windows Server 2008 R2) are not supported on the Models 250, 550, and the 1U xx20s because these models do not incorporate Virtual Technology (VT) support.

Steelhead Model	RiOS	RSP Image	32-bit Guest VMs	64-bit Guest VMs
250/550	32-bit	32-bit	Yes	No
520/1020/1520/2020	64-bit	64-bit	Yes	No
1050/2050	64-bit	64-bit	Yes	Yes (starting v5.5.3a)
3020/3520	64-bit	64-bit	Yes	Yes (starting v5.5.1)
5050/6050/7050	64-bit	64-bit	Yes	Yes

Ethernet Network Compatibility

The Steelhead appliance supports the following types of Ethernet networks:

- Ethernet Logical Link Control (LLC) (IEEE 802.2 - 2005)
- Fast Ethernet 100 Base-TX (IEEE 802.3 - 2005)
- Gigabit Ethernet over Copper 1000 Base-T and Fiber 1000 Base-SX (LC connector) and Fiber 1000 Base LX (IEEE 802.3 - 2005)

The Steelhead appliance ports support the following connection types and speeds:

- **Primary** - 10/100/1000 Base-T, auto-negotiating
- **Auxiliary** - 10/100/1000 Base-T, auto-negotiating
- **LAN** - 10/100/1000 Base-TX or 1000 Base-SX or 1000 Base-LX, depending on configuration
- **WAN** - 10/100/1000 Base-TX or 1000 Base-SX or 1000 Base-LX, depending on configuration

Note: 1000 Base-SX and 1000 Base-LX interface options are not available for the Steelhead appliance 250 and 550 models.

The Steelhead appliance supports VLAN Tagging (IEEE 802.1Q - 2005). It does not support the ISL protocol.

All copper interfaces are auto-sensing for speed and duplex (IEEE 802.3 - 2005).

The Steelhead appliance auto-negotiates speed and duplex mode for all data rates and supports full duplex mode and flow control (IEEE 802.3 – 2005).

The Steelhead appliance with a Gigabit Ethernet card supports Jumbo Frames on in-path and primary ports.

SNMP-Based Management Compatibility

The Steelhead appliance supports a proprietary Riverbed MIB accessible through SNMP. SNMPv1 (RFCs 1155, 1157, 1212, and 1215), SNMPv2c (RFCs 1901, 2578, 2579, 2580, 3416, 3417, and 3418), and SNMPv3 are supported, although some MIB items might only be accessible through SNMPv2 and SNMPv3.

SNMP support allows the Steelhead appliance to be integrated into network management systems such as Hewlett Packard OpenView Network Node Manager, BMC Patrol, and other SNMP-based network management tools.

CMC Compatibility

The Steelhead appliance has been tested with the following Central Management Console (CMC) versions.

Steelhead RiOS Version	Recommended CMC Version	CMC v6.1.x	CMC v6.0.x	CMC v5.5.x
v6.5	v6.1.x partial support Parity in subsequent v6.5 release.	Manages only v6.1.x features, does not support QoS.	-	Not supported
v6.1.x	v6.1.0	Parity; includes Virtual Steelhead, does not support ProCurve.	Manages some RiOS v6.1.x features may be supported in subsequent point releases of CMC v6.0.	Not supported
v6.0.x	v6.1.0, 6.0.1	Parity; manages all Steelhead appliance v6.0.2 and later features, including RSP.	Parity; manages all Steelhead appliance v6.0.x features.	CMC v5.5.3 and later; manages only v5.5 Steelhead appliance features.
v5.5.x	v6.1.0, 6.0.1	Parity; manages all Steelhead appliance v5.5.4 only, does not support RSP.	Parity; manages all Steelhead appliance v5.5.4 features only, does not support RSP.	Parity; does not support RSP.

CSH The ESX Cloud Steelhead supports CMC.

Antivirus Compatibility

The Steelhead appliance has been tested with the following antivirus software with no impact on performance:

- Network Associates (McAfee) VirusScan v7.0.0 Enterprise on the server
- Network Associates (McAfee) VirusScan v7.1.0 Enterprise on the server
- Network Associates (McAfee) VirusScan v7.1.0 Enterprise on the client
- Symantec (Norton) AntiVirus Corporate Edition v8.1 on the server

The Steelhead appliance has been tested with the following antivirus software with moderate impact on performance:

- F-Secure Anti-Virus v5.43 on the client
- F-Secure Anti-Virus v5.5 on the server
- Network Associates (McAfee) NetShield v4.5 on the server
- Network Associates VirusScan v4.5 for multi-platforms on the client
- Symantec (Norton) AntiVirus Corporate Edition v8.1 on the client

Additional Resources

This section describes resources that supplement the information in this guide. It includes the following sections:

- "Release Notes" on page 17
- "Riverbed Documentation and Support Knowledge Base" on page 17

Release Notes

The following online file supplements the information in this manual. It is available on the Riverbed Support site at https://support.riverbed.com.

Online File	Purpose
<product>_<version_number> <build_number>.pdf	Describes the product release and identifies fixed problems, known problems, and workarounds. This file also provides documentation information not covered in the manuals or that has been modified since publication.

Please examine this file before you begin the installation and configuration process. It contains important information about this release of the Steelhead appliance.

Riverbed Documentation and Support Knowledge Base

For a complete list and the most current version of Riverbed documentation log in to the Riverbed Support Web site located at https://support.riverbed.com.

The Riverbed Knowledge Base is a database of known issues, how-to documents, system requirements, and common error messages. You can browse titles or search for key words and strings.

To access the Riverbed Knowledge Base, log in to the Riverbed Support site located at https://support.riverbed.com.

Contacting Riverbed

This section describes how to contact departments within Riverbed.

Internet

You can find out about Riverbed products through our Web site at http://www.riverbed.com.

Support

If you have problems installing, using, or replacing Riverbed products contact Riverbed Support or your channel partner who provides support. To contact Riverbed Support, please open a trouble ticket at https://support.riverbed.com or call 1-888-RVBD-TAC (1-888-782-3822) in the United States and Canada or +1 415 247 7381 outside the United States.

Professional Services

Riverbed has a staff of professionals who can help you with installation assistance, provisioning, network redesign, project management, custom designs, consolidation project design, and custom coded solutions. To contact Riverbed Professional Services go to http://www.riverbed.com or email proserve@riverbed.com.

Documentation

We continually strive to improve the quality and usability of our documentation. We appreciate any suggestions you may have about our online documentation or printed materials. Send documentation comments to techpubs@riverbed.com.

CHAPTER 1

Overview of the Management Console

This chapter provides an overview of the Management Console. It includes the following sections:

- "Using the Management Console" on page 19
- "Next Steps" on page 25

This chapter assumes you have installed and configured the Steelhead appliance. For details, see the *Steelhead Appliance Installation and Configuration Guide*.

This chapter also assumes you are familiar with the various deployment options available to you. For details, see the *Riverbed Deployment Guide*.

Using the Management Console

The following section describes how to connect to and navigate in the Management Console. It includes the following sections:

- "Connecting to the Management Console" on page 19
- "The Home Page" on page 20
- "Navigating in the Management Console" on page 21
- "Getting Help" on page 24

Note: If you prefer, you can use the CLI to perform configuring and monitoring tasks. For details, see the *Riverbed Command-Line Interface Reference Manual*.

Connecting to the Management Console

To connect to the Management Console you must know the URL and administrator password that you assigned in the configuration wizard of the Steelhead appliance. For details, see the *Steelhead Appliance Installation and Configuration Guide*.

To connect to the Management Console

1. Specify the URL for the Management Console in the location box of your Web browser:

```
protocol://host.domain
```

protocol is http or https. HTTPS uses the SSL protocol to ensure a secure environment. When you connect using HTTPS, you are prompted to inspect and verify the SSL certificate. This is a self-signed certificate used to provide encrypted Web connections to the Management Console. It is re-created when the appliance hostname is changed and when the certificate has expired.

Note: The secure vault does not protect the self-signed certificate used with HTTPS connections.

host is the hostname you assigned to the Steelhead appliance primary interface in the Configuration Wizard. If your DNS server maps that IP address to a name, you can specify the DNS name.

domain is the full domain name for the appliance.

Tip: Alternatively, you can specify the IP address instead of the host and domain name.

The Management Console appears, displaying the Login page.

Figure 1-1. Login Page

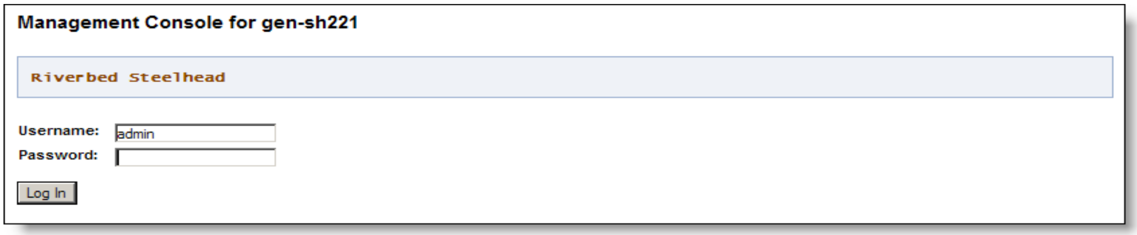

Management Console for gen-sh221

Riverbed Steelhead

Username: admin
Password:

2. In the Username text box, specify the user login: admin, monitor, a login from a RADIUS or TACACS+ database, or any local accounts created using the Role-Based Accounts feature. The default login is admin. For details on role-based accounts, see "Role-Based Accounts" on page 363.

   Users with administrator (admin) privileges can configure and administer the Steelhead appliance. Users with monitor (monitor) privileges can view the Steelhead appliance reports, user logs, and change their own password. A monitor user cannot make configuration changes.

3. In the Password text box, specify the password you assigned in the configuration wizard of the Steelhead appliance. (The Steelhead appliance is shipped with the default password: password.)

4. Click **Log In** to display the Home page.

The Home Page

The Home page lists the system hostname, system up time, service up time, temperature, and the CMC hostname (if you have one in your network). It also displays the following reports:

- **Optimized LAN Throughput Over Last Week** - Summarizes the throughput or total data transmitted for all applications in the last week.
- **Bandwidth Summary Over Last Week** - Provides a 3-dimensional view of traffic patterns (byte counts) over the last week. Each column represents the number of bytes, the time of day, and the day of the week. For example, the report might display that there were 4 GBs of WAN traffic from 12 PM to 3 PM on Wednesday of the prior week.

The top of every page displays the menu bar. The current state of the system appears to the right of the menus: Healthy, Admission Control, Degraded, or Critical and is always visible. For details, select the current system status to display the Alarm Status page, or see "SNMP Traps" on page 498.

Figure 1-2. The Home Page

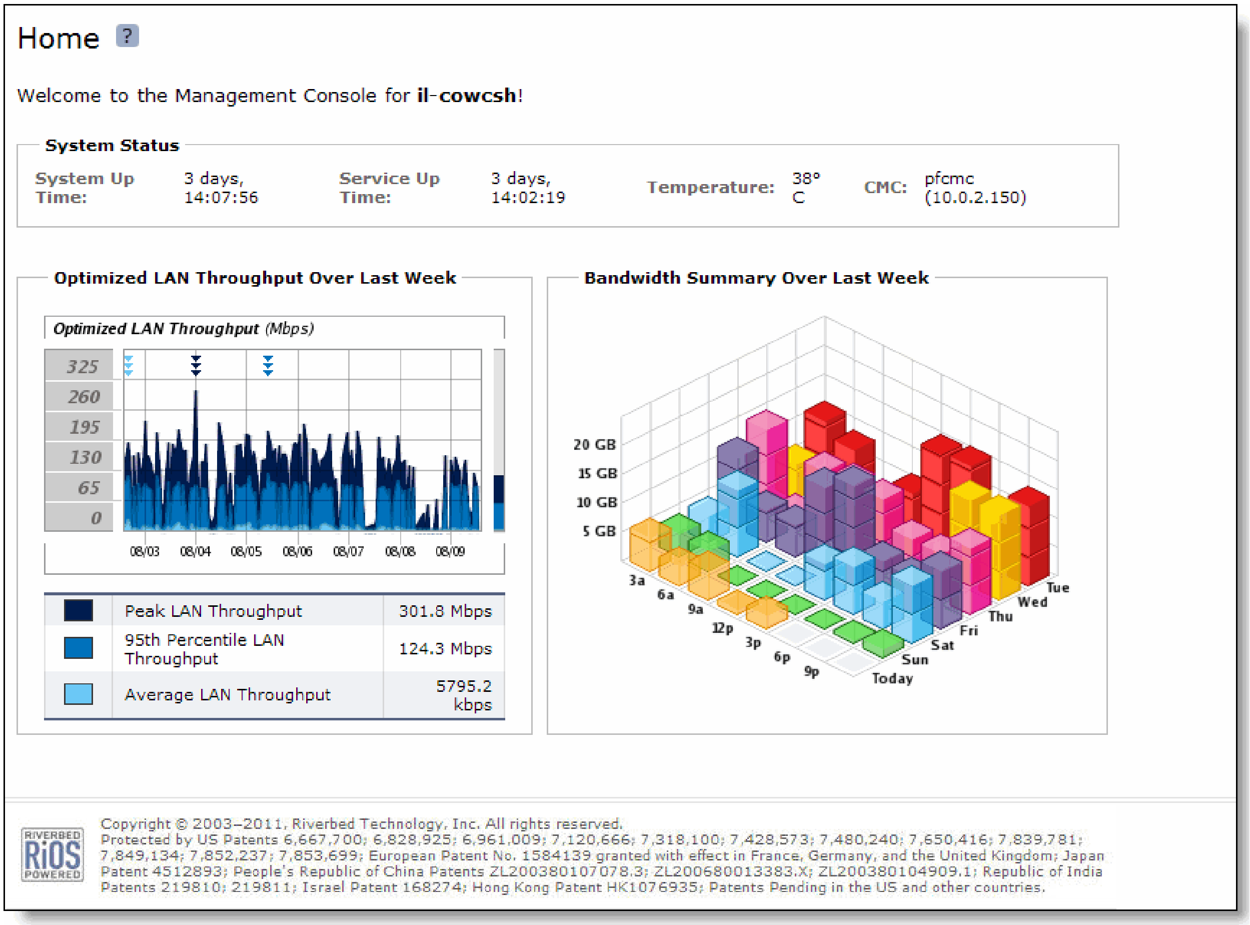

Home ?

Welcome to the Management Console for **il-cowcsh**!

System Status

System Up Time:	3 days, 14:07:56	Service Up Time:	3 days, 14:02:19	Temperature:	38° C	CMC:	pfcmc (10.0.2.150)

Optimized LAN Throughput Over Last Week

Optimized LAN Throughput (Mbps)

■	Peak LAN Throughput	301.8 Mbps
■	95th Percentile LAN Throughput	124.3 Mbps
■	Average LAN Throughput	5795.2 kbps

Bandwidth Summary Over Last Week

Navigating in the Management Console

You navigate to the tools and reports available to you in the Management Console using cascading menus.

To display cascading menus

1. Select the Configure and Reports menus to display the submenus. For example, select Reports to display the submenus Optimization, Branch Services, Diagnostics, and Export. The menu item that is currently active is highlighted.

2. To go to a page, slide your cursor down to the submenu item you want to display and select the menu name. For example, under Reports > Optimization go to Bandwidth Optimization and select Bandwidth Optimization to display the page.

Figure 1-3 shows the cascading menus.

Figure 1-3. Cascading Menus

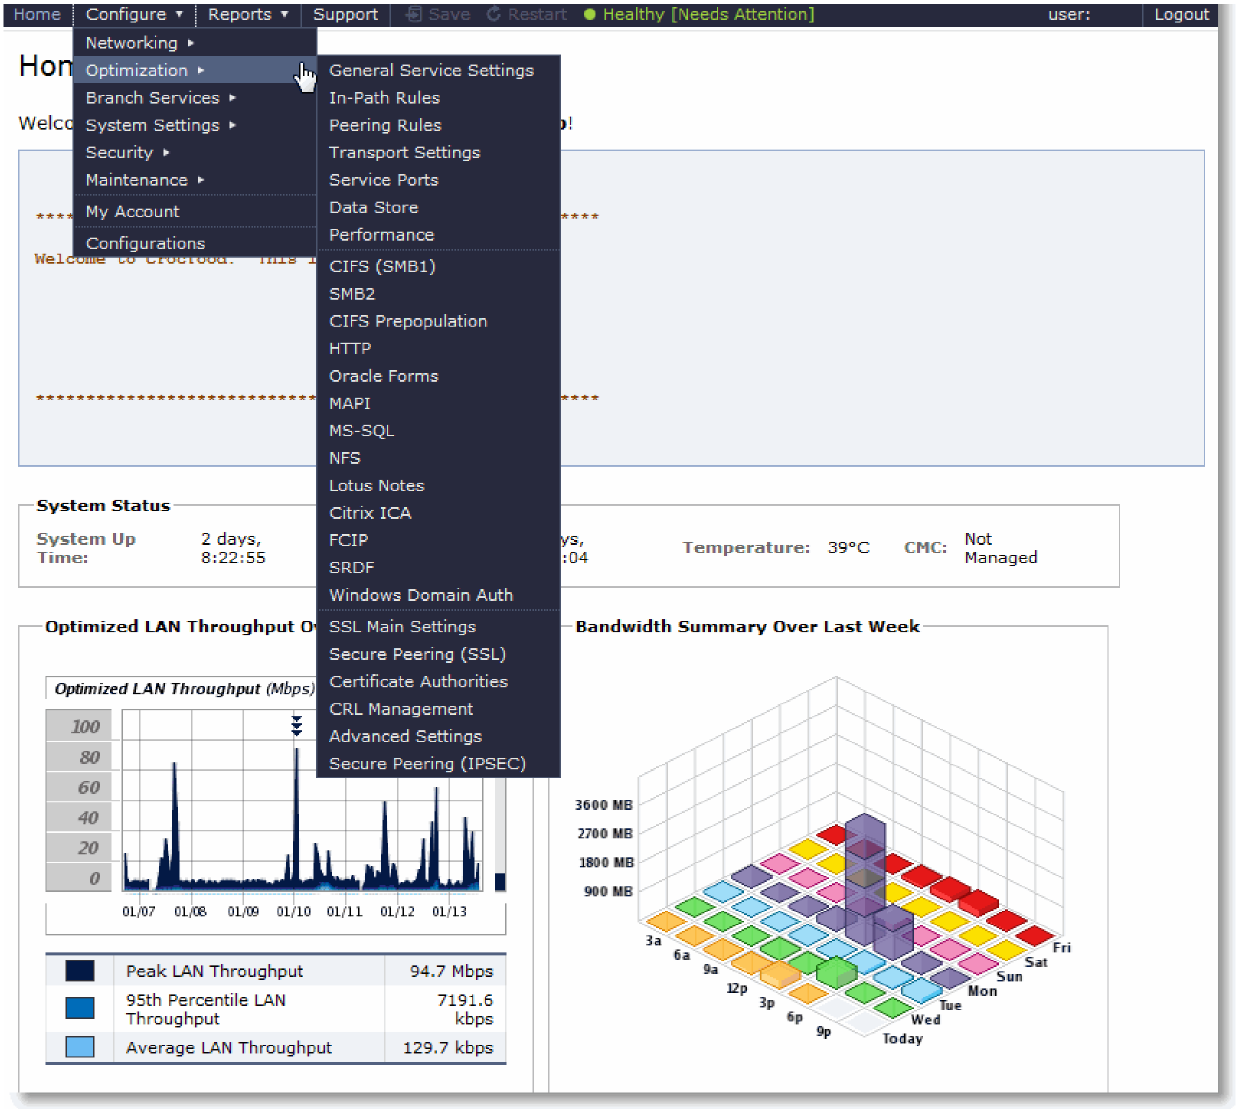

System Status

System Up Time: 2 days, 8:22:55 ys, :04 Temperature: 39°C CMC: Not Managed

■	Peak LAN Throughput	94.7 Mbps
■	95th Percentile LAN Throughput	7191.6 kbps
■	Average LAN Throughput	129.7 kbps

The following table summarizes the cascading menus.

Menu	Submenus
Home	Displays the Home page.
Configure	**Networking** - Configure host settings, base interfaces, asymmetric routing, connection forwarding, encryption, flow export, QoS, simplified routing, port labels, Windows Domain, WCCP, and subnet side rules from this menu.
	Optimization - Configure optimization features such as in-path rules, protocols, datastore, high-speed TCP, peering rules, CIFS prepopulation, service ports, Oracle Forms, MS-SQL, NFS, MAPI, Lotus Notes, Citrix ICA, FCIP, SRDF, Windows Domain authentication, and SSL from this menu.
	Branch Services - Configure branch service features such as PFS, DNS caching, and RSP from this menu.
	System Settings - Configure alarm settings, announcements, email settings, log settings, monitored ports, and SNMP settings from this menu.

Menu	Submenus
	Security - Configure general security parameters, RADIUS, TACACS+, the secure vault, Management ACL, Web settings, and user permissions from this menu.
	Maintenance - Start and stop Steelhead services, schedule jobs, manage and update licenses, upgrade software, and reboot or shutdown the appliance from this menu.
	My Account - Change your login password and view user permissions from this menu.
	Configurations - Manage, import, delete, and change your configuration files for the appliance from this menu.
Reports	**Networking** - Create and display reports such as current connections, connection history, connection forwarding, QoS statistics for dropped and sent data packets, top talkers, traffic summary, interface counters, and TCP statistics from this menu.
	Optimization - View optimization reports such as optimized throughput, bandwidth optimization, data reduction, connected appliances, connection pooling, HTTP statistics, NFS statistics, SSL servers, and datastore reports from this menu.
	Branch Services - Create and display PFS, DNS caching, and RSP VNI IO reports from this menu.
	Diagnostics - Display and download diagnostic reports such as CPU utilization, memory paging, user and system logs, alarms status, system snapshots, system dumps, TCP dumps, and the Steelhead appliance health status from this menu.
	Export - Export reports from this menu.
Support	Display online help and appliance documentation, contact information for Riverbed Support, appliance details such as the model number, revision type, serial number, software version, and appliance MIB files from this menu.

Saving Your Configuration

As you **Apply** page settings, the values are applied to the running configuration. Most Management Console configuration pages include an **Apply** button for you to commit your changes. When you click **Apply**, the Management Console updates the running configuration. Your changes are only written to disk when you save your configuration.

The Save icon on the menu bar alerts you if the changes you have made require you to save them to disk.

To save your changes, click **Save** to save the changes to disk.

A red dot in a control indicates that the field is required. You must specify a valid entry for all of the required controls on a page before submitting the changes to the system.

Restarting the Steelhead Service

The Steelhead service is a daemon that executes in the background, performing required operations.

Some configuration settings apply to the Steelhead service. When you change settings for features that depend on the Steelhead service, you must restart the service for the changes to take effect.

To restart the service, click **Restart** to display the Services page or choose Configure > Maintenance > Services and then restart the service from the Services page. For details, see “Starting and Stopping the Optimization Service” on page 345.

Logging Out

Click **Logout** in the upper-right corner of the screen to log out of the current session.

Printing Pages and Reports

You can print Management Console pages and reports using the print option on your Web browser.

To print pages and reports

- Choose File > Print in your Web browser to open the Print dialog box.

Getting Help

The Support page provides the following options:

- **Online Help** - View browser-based online help.
- **Support** - View links and contact information for Riverbed Support.
- **Appliance Details** - View appliance information such as the model number, hardware revision type, serial number, and the software version number currently installed on the appliance.
- **MIB Files** - View Riverbed and appliance MIB files in text format.

Displaying Online Help

The Management Console provides page-level help for the appliance.

To display online help in the Management Console

- Click the Question Mark icon next to the page title. The help for the page appears in a new browser window.

Downloading Documentation

The Riverbed Support Site contains PDF versions of the *Steelhead Management Console User's Guide* and the *Riverbed Command-Line Interface Reference Manual*.

To download the PDF versions of the User's Guide or Command-Line Interface Reference Manual

1. Select Support in the menu bar to display the Support page.
2. You must be registered on the Riverbed Support site to download the documentation. Go to one of the following links:
   - To register on the Riverbed Support site:

     https://support.riverbed.com/account/registration/register.htm
   - If you are registered on the Riverbed Support site:

     https://support.riverbed.com/docs/index.htm
3. Go to the PDF document.
4. Select the document name to download the document.

Logging Out

In the menu bar, click **Logout** to end your session.

Next Steps

The following table describes a basic approach to configuring the Steelhead appliance.

Task	Reference
1. Become familiar with basic and advanced deployment types.	*Riverbed Deployment Guide*
2. Make decisions about where to deploy Steelhead appliances, and what features to use.	Riverbed Professional Services
3. Install appliances and optional interface cards.	*Steelhead Appliance Installation and Configuration Guide* *Network Interface Card Installation Guide*
4. Configure optimization traffic with in-path rules.	"Configuring In-Path Rules" on page 28
5. Enable optimization features related to your deployment.	"Configuring SSL Server Certificates and Certificate Authorities" on page 213 (if applicable) "Configuring Optimization Features" on page 59 "Configuring Network Integration Features" on page 251 (if applicable) "Configuring Branch Services" on page 161 (if applicable) "Managing Licenses and Model Upgrades" on page 349 (if necessary)
6. Distribute administrative responsibility by configuring secure access for other administrators, monitor users, or other types of users you choose to create.	"Configuring General Security Settings" on page 361 (if applicable)
7. Modify default system administration settings.	"Configuring Alarm Settings" on page 323 (if desired)
8. Modify host and network settings you initially set with the installation wizard.	"Modifying Host and Network Interface Settings" on page 39 (if desired)
9. Save your configuration changes and restart services as necessary.	"Starting and Stopping the Optimization Service" on page 345 (as necessary) "Managing Configuration Files" on page 358 (as necessary)
10. View logs and reports to verify your deployment.	"Viewing Current Connections" on page 385
11. Troubleshoot (if necessary).	*Riverbed Deployment Guide* Riverbed Support

CHAPTER 2 Configuring In-Path Rules

This chapter describes how to configure in-path rules. It includes the following sections:

- "About In-Path Rules" on page 27
- "About Default In-Path Rules" on page 28
- "Configuring In-Path Rules" on page 28

About In-Path Rules

In-path rules are used only when a connection is *initiated*. Because connections are usually initiated by clients, in-path rules are configured for the initiating, or client-side Steelhead appliance. In-path rules determine Steelhead appliance behavior with SYN packets.

In-path rules are an ordered list of fields a Steelhead appliance uses to match with SYN packet fields (for example, source or destination subnet, IP address, VLAN, or TCP port). Each in-path rule has an *action* field. When a Steelhead appliance finds a matching in-path rule for a SYN packet, the Steelhead appliance treats the packet according to the action specified in the in-path rule.

In-path rules are used only in the following scenarios:

- TCP SYN packet arrives on the LAN interface of physical in-path deployments.
- TCP SYN packet arrives on the WAN0_0 interface of virtual in-path deployments.

Both of these scenarios are associated with the first, or *initiating*, SYN packet of the connection. Because most connections are initiated by the client, you configure your in-path rules on the client-side Steelhead appliance. In-path rules have no effect on connections that are already established, regardless of whether the connections are being optimized.

In-path rule configurations differ depending on the action. For example, both the fixed-target and the auto-discovery actions allow you to choose what type of optimization is applied, what type of data reduction is used, what type of latency optimization is applied, and so on.

You can configure optional settings to support a variety of deployment needs, including:

- **Optimization Policies** - Optimize connections using scalable data reduction, compression, both, or none.
- **VLAN Tags** - Apply a rule to a specific VLAN or all VLANs.
- **Preoptimization Policies** - Special handling required for Oracle Forms over SSL support.

- **Latency Policies** - Set to normal, none, or HTTP to support HTTP traffic. Special handling required for Oracle Forms over SSL support.
- **Neural Framing Requirements** - Specify never, always, TCP Hints, or Dynamic.
- **WAN Visibility** - Preserve TCP/IP address or port information.

For details, see the configuration descriptions included in "Configuring In-Path Rules" on page 28.

About Default In-Path Rules

Three types of default in-path rules ship with Steelhead appliances. These default rules pass through certain types of traffic unoptimized. The primary reason that these types of traffic are passed through is because you are likely to use these types of protocols (telnet, ssh, https) when you deploy and configure your Steelhead appliances. The default rules allow the following traffic to pass through the Steelhead appliance without attempting optimization:

Port Type	Description and Ports
Interactive traffic	Ports 7, 23, 37, 107, 513, 514, 3389, 5631, 5900-5903, 6000. This default rule automatically passes traffic through on interactive ports (for example, Telnet, TCP ECHO, remote logging, and shell).
Riverbed Protocols	Ports 7744 (datastore synchronization), 7800-7801 (in-path), 7810 (out-of-path), 7820 (failover), 7850 (connection forwarding), 7860 (Interceptor appliance), 7870 (Steelhead Mobile Controller). This default rule automatically passes traffic through on ports used by the system.
Secure, encrypted traffic	Ports 22, 443, 465, 563, 585, 614, 636, 989, 990, 992, 993, 995, 1701, 1723, 3713. This default rule automatically passes traffic through on commonly secure ports (for example, ssh, https, and smtps).

Riverbed recommends you retain the default rules. However, you can remove or overwrite the default in-path rules by altering or adding other rules to the in-path rule list, or by changing the port groups that are used.

For details on changing port labels, see "Configuring Port Labels" on page 94.

Configuring In-Path Rules

You can review, add, edit, and remove in-path rules in the Configure > Optimization > In-Path Rules page. The In-Path Rules table lists the order and properties of the rules set for the running configuration.

For an overview of in-path rules, see "About In-Path Rules" on page 27.

To configure in-path rules

1. Choose Configure > Optimization > In-Path Rules to display the In-Path Rules page.

Figure 2-1. In-Path Rules Page

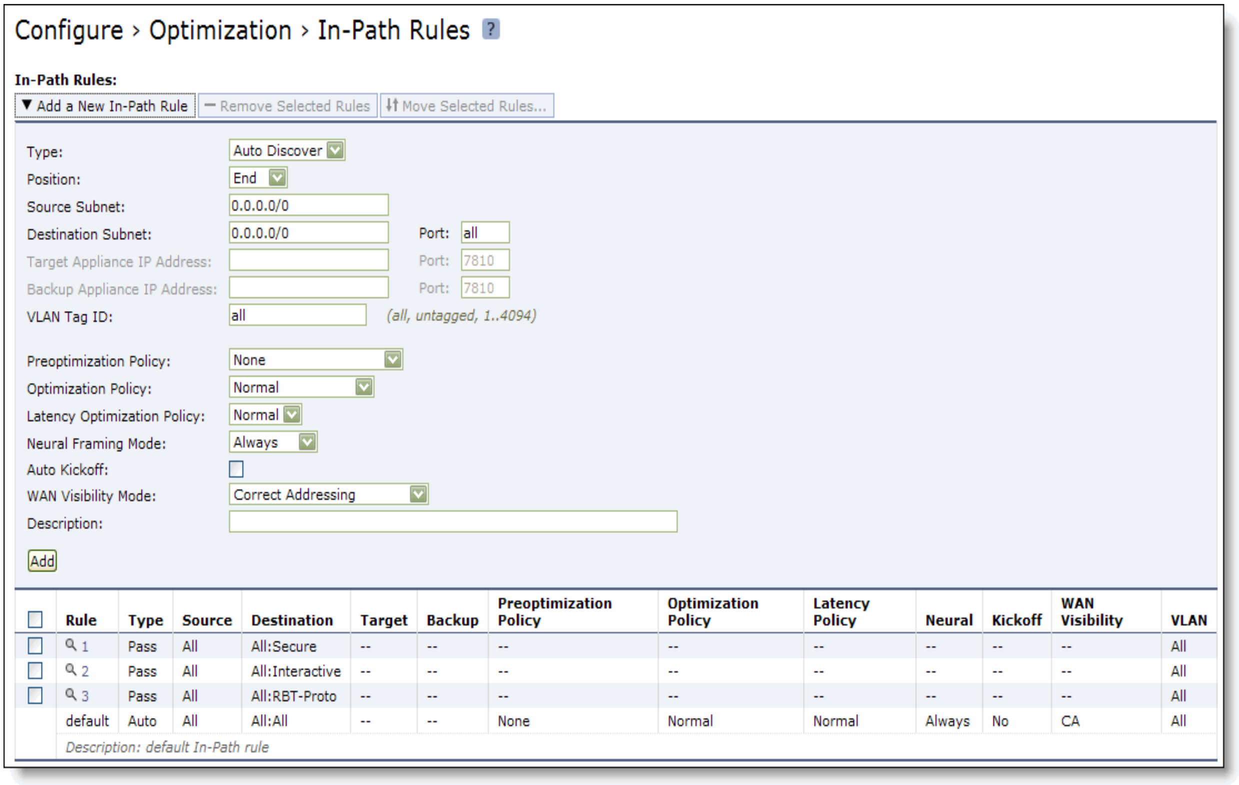

Configure > Optimization > In-Path Rules

In-Path Rules:

Type: Auto Discover

Position: End

Source Subnet: 0.0.0.0/0

Destination Subnet: 0.0.0.0/0 Port: all

Target Appliance IP Address: Port: 7810

Backup Appliance IP Address: Port: 7810

VLAN Tag ID: all *(all, untagged, 1..4094)*

Preoptimization Policy: None

Optimization Policy: Normal

Latency Optimization Policy: Normal

Neural Framing Mode: Always

Auto Kickoff: □

WAN Visibility Mode: Correct Addressing

Description:

□	Rule	Type	Source	Destination	Target	Backup	Preoptimization Policy	Optimization Policy	Latency Policy	Neural	Kickoff	WAN Visibility	VLAN
□	1	Pass	All	All:Secure	--	--	--	--	--	--	--	--	All
□	2	Pass	All	All:Interactive	--	--	--	--	--	--	--	--	All
□	3	Pass	All	All:RBT-Proto	--	--	--	--	--	--	--	--	All
	default	Auto	All	All:All	--	--	None	Normal	Normal	Always	No	CA	All
	Description: default In-Path rule												

2. Configure the rules as described in the following table.

Control	Description
Add a New In-Path Rule	Displays the controls for adding a new rule.
Type	Select one of the following rule types from the drop-down list: • **Auto-Discover** - Uses the auto-discovery process to determine if a remote Steelhead appliance is able to optimize the connection attempting to be created by this SYN packet. By default, auto-discover is applied to all IP addresses and ports that are not secure, interactive, or default Riverbed ports. Defining in-path rules modifies this default setting. • **Fixed-Target** - Skips the auto-discovery process and uses a specified remote Steelhead appliance as an optimization peer. You must specify at least one remote target Steelhead appliance to optimize (and, optionally, which ports and backup Steelhead appliances), and add rules to specify the network of servers, ports, port labels, and out-of-path Steelhead appliances to use. • **Pass-Through** - Allows the SYN packet to pass through the Steelhead appliance unoptimized. No optimization is performed on the TCP connection initiated by this SYN packet. You define pass-through rules to exclude subnets from optimization. Traffic is also passed through when the Steelhead appliance is in bypass mode. (Pass through of traffic might occur because of in-path rules or because the connection was established before the Steelhead appliance was put in place or before the Steelhead service was enabled.) • **Discard** - Drops the SYN packets silently. The Steelhead appliance filters out traffic that matches the discard rules. This process is similar to how routers and firewalls drop disallowed packets: the connection-initiating device has no knowledge of the fact that its packets were dropped until the connection times out. • **Deny** - Drops the SYN packets, sends a message back to its source, and resets the TCP connection being attempted. Using an active reset process rather than a silent discard allows the connection initiator to know that its connection is disallowed.
Position	Select Start, End, or a rule number from the drop-down list. Steelhead appliances evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted. In general, list rules in the following order: 1. Deny 2. Discard 3. Pass-through 4. Fixed-target 5. Auto-Discover **Note:** The default rule, Auto-Discover, which optimizes all remaining traffic that has not been selected by another rule, cannot be removed and is always listed last.
Source Subnet	Specify the subnet IP address and netmask for the source network. Use the following format: XXX.XXX.XXX.XXX/XX Or, you can specify all or 0.0.0.0/0 as the wildcard for all traffic.

(1 of 7)

Control	Description
Destination Subnet	Specify the subnet IP address and netmask for the destination network. Use the following format: XXX.XXX.XXX.XXX/XX Or, you can specify all or 0.0.0.0/0 as the wildcard for all traffic.
	Port - Specify the destination port number, port label, or all.
Target Appliance IP Address	Specify the target appliance address for a fixed-target rule.
	Port - Specify the target port number for a fixed-target rule.
Backup Appliance IP Address	Specify the backup appliance address for a fixed-target rule.
	Port - Specify the backup destination port number for a fixed-target rule.
VLAN Tag ID	Specify a VLAN identification number from 0 to 4094, enter all to apply the rule to all VLANs, or enter untagged to apply the rule to non-tagged connections. RiOS supports VLAN v802.1Q. To configure VLAN tagging, configure in-path rules to apply to all VLANs or to a specific VLAN. By default, rules apply to all VLAN values unless you specify a particular VLAN ID. Pass-through traffic maintains any pre-existing VLAN tagging between the LAN and WAN interfaces.
Preoptimization Policy	Select a traffic type from the drop-down list: • **None** - If the Oracle Forms, SSL, or Oracle Forms over SSL preoptimization policy is turned on and you want to turn it off for a port, select none. This is the default setting. In RiOS v6.0 and later, traffic to port 443 always uses a preoptimization policy of SSL even if an in-path rule on the client-side Steelhead appliance sets the preoptimization policy to None. To disable the SSL preoptimization for traffic to port 443, you can either: 1. Disable the SSL optimization on the client or server-side Steelhead appliance. —or— 2. Modify the peering rule on the server-side Steelhead appliance by setting the SSL Capability control to No Check. • **Oracle Forms** - Enables preoptimization processing for Oracle Forms. • **Oracle Forms over SSL** - Enables preoptimization processing for both the Oracle Forms and SSL encrypted traffic through SSL secure ports on the client-side Steelhead appliance. You must also set the Latency Optimization Policy to HTTP. **Note**: If the server is running over a standard secure port, for example, port 443, the Oracle Forms over SSL in-path rule needs to be *before* the default secure port pass-through rule in the in-path rule list. • **SSL** - Enables preoptimization processing for SSL encrypted traffic through SSL secure ports on the client-side Steelhead appliance.

(2 of 7)

Control	Description
Optimization Policy	Optionally, if you have selected Auto-Discover or Fixed Target, you can configure the following types of optimization policies: • **Normal** - Perform LZ compression and SDR. • **SDR-Only** - Perform SDR; do not perform LZ compression. • **SDR-M** - Performs data reduction entirely in memory, which prevents the Steelhead appliance from reading and writing to and from the disk. Enabling this option can yield high LAN-side throughput because it eliminates all disk latency. This optimization policy is useful for – a very small amount of data; for example, interactive traffic. – point-to-point replication during off-peak hours when both the server-side and client-side Steelheads are the same (or similar) size. Both Steelhead appliances must be running RiOS v6.0.x or later. • **Compression-Only** - Perform LZ compression; do not perform SDR. • **None** - Do not perform SDR or LZ compression. To configure optimization policies for the FTP data channel, define an in-path rule with the destination port 20 and set its optimization policy. Setting QoS for port 20 on the client-side Steelhead appliance effects passive FTP, while setting the QoS for port 20 on the server-side Steelhead appliance effects active FTP. To configure optimization policies for the MAPI data channel, define an in-path rule with the destination port 7830 and set its optimization policy.
Latency Optimization Policy	Select one of the following policies from the drop-down list: • **Normal** - Perform all latency optimizations (HTTP is activated for ports 80 and 8080). This is the default setting. • **HTTP** - Activate HTTP optimization on connections matching this rule. • **Outlook Anywhere** - Activate RPC over HTTP(S) optimization for Outlook Anywhere on connections matching this rule. To auto-detect Outlook Anywhere or HTTP on a connection, select the Normal latency optimization policy and enable the Auto-Detect Outlook Anywhere Connections option in the Configure > Optimization > MAPI page. The auto-detect option in the MAPI page is best for simple Steelhead configurations with only a single Steelhead at each site and when the IIS server is also handling Web sites. If the IIS server is only used as RPC Proxy, and for configurations with asymmetric routing, connection forwarding or Interceptor installations, add in-path rules that identify the RPC Proxy server IP addresses and select this latency optimization policy. After adding the in-path rule, disable the auto-detect option in the Configure > Optimization > MAPI page. • **None** - Do not activate latency optimization on connections matching this rule. For Oracle Forms over SSL encrypted traffic, you must set the Latency Optimization Policy to HTTP. **Tip:** Setting the Latency Optimization Policy to None excludes *all* latency optimizations such as HTTP, MAPI, and SMB.

(3 of 7)

Control	Description
Neural Framing Mode	Optionally, if you have selected Auto-Discover or Fixed Target, you can select a neural framing mode for the in-path rule. Neural framing enables the system to select the optimal packet framing boundaries for SDR. Neural framing creates a set of heuristics to intelligently determine the optimal moment to flush TCP buffers. The system continuously evaluates these heuristics and uses the optimal heuristic to maximize the amount of buffered data transmitted in each flush, while minimizing the amount of idle time that the data sits in the buffer. You can specify the following neural framing settings: • **Never** - Never use the Nagle algorithm. All the data is immediately encoded without waiting for timers to fire or application buffers to fill past a specified threshold. Neural heuristics are computed in this mode but are not used. In general, this setting works well with time-sensitive and chatty or real-time traffic. • **Always** - Always use the Nagle algorithm. This is the default setting. All data is passed to the codec which attempts to coalesce consume calls (if needed) to achieve better fingerprinting. A timer (6 ms) backs up the codec and causes leftover data to be consumed. Neural heuristics are computed in this mode but are not used. • **TCP Hints** - If data is received from a partial frame packet or a packet with the TCP PUSH flag set, the encoder encodes the data instead of immediately coalescing it. Neural heuristics are computed in this mode but are not used. • **Dynamic** - Dynamically adjust the Nagle parameters. In this option, the system discerns the optimum algorithm for a particular type of traffic and switches to the best algorithm based on traffic characteristic changes. For different types of traffic, one algorithm might be better than others. The considerations include: latency added to the connection, compression, and SDR performance. To configure neural framing for an FTP data channel, define an in-path rule with the destination port 20 and set its optimization policy. To configure neural framing for a MAPI data channel, define an in-path rule with the destination port 7830 and set its optimization policy.

(4 of 7)

Control	Description
Auto Kickoff	Enables kickoff, which resets pre-existing connections to force them to go through the connection creation process again. If you enable kickoff, connections that pre-exist when the optimization service is started are re-established and optimized. Generally, connections are short lived and kickoff is not necessary. It is suitable for certain long-lived connections, such as data replication, and very challenging remote environments. For example, in a remote branch-office with a T1 and a 35 ms round-trip time, you would want connections to migrate to optimization gracefully, rather than risk interruption with kickoff. RiOS v6.1 and later provides three ways to enable kickoff: • Globally for all existing connections in the Configure > Optimization > General Service Settings page. • For a single pass-through or optimized connection in the Current Connections report, one connection at a time. • For all existing connections that match an in-path rule and the rule has kickoff enabled. In most deployments, you do not want to set automatic kickoff globally because it disrupts *all* existing connections. When you enable kick off using an in-path rule, once the Steelhead detects packet flow that matches the IP and port specified in the rule, it sends an RST packet to the client and server maintaining the connection to try to close it. Next, it sets an internal flag to prevent any further kickoffs until the optimization service is once again restarted. **Note:** If no data is being transferred between the client and server, the connection is not reset immediately. It resets the next time the client or server tries to send a message. Therefore, when the application is idle, it may take a while for the connection to reset. By default, auto kickoff per in-path rule is disabled. The service applies the first matching in-path rule for an existing connection that matches the source and destination IP and port; it does not consider a VLAN tag ID when determining whether to kick off the connection. Consequently, the service automatically kicks off connections with matching source and destination addresses and ports on different VLANs. The source and destination of a pre-existing connection cannot be determined because the Steelhead appliance did not see the initial TCP handshake whereas an in-path rule specifies the source and destination IP address to which the rule should be applied. Hence this connection for this IP address pair is matched twice, once as source to destination and the other as destination to source to find an in-path rule. For example, the following in-path rule will kick off connections from 10.11.10.10/24 to 10.12.10.10/24 and 10.12.10.10/24 to 10.11.10.10/24. Src 10.11.10.10/24 Dst 10.12.10.10/24 Auto Kickoff enabled The first matching in-path rule will be considered during the kickoff check for a pre-existing connection. If the first matching in-path rule has kickoff enabled, then that pre-existing connection will be reset. **Important:** Specifying automatic kickoff per in-path rule enables kickoff even when you disable the global kickoff feature. When global kickoff is enabled, it overrides this setting. You set the global kickoff feature using the Reset Existing Client Connections on Start Up feature, which appears in the Configure > Optimization > General Service Settings page. **Note:** This feature pertains only to auto-discover and fixed-target rule types and is dimmed and unavailable for the other rule types.

(5 of 7)

Control	Description
WAN Visibility Mode	Enables WAN visibility, which pertains to how packets traversing the WAN are addressed. RiOS v5.0 or later offers three types of WAN visibility: correct addressing, port transparency, and full address transparency. You configure WAN visibility on the client-side Steelhead appliance (where the connection is initiated). The server-side Steelhead appliance must also support WAN visibility (RiOS v5.0 or later). Select one of the following modes from the drop-down list: • **Correct Addressing** - Turns WAN visibility off. Correct addressing uses Steelhead appliance IP addresses and port numbers in the TCP/IP packet header fields for optimized traffic in both directions across the WAN. This is the default setting. • **Port Transparency** - Port address transparency preserves your server port numbers in the TCP/IP header fields for optimized traffic in both directions across the WAN. Traffic is optimized while the server port number in the TCP/IP header field appears to be unchanged. Routers and network monitoring devices deployed in the WAN segment between the communicating Steelhead appliances can view these preserved fields. Use port transparency if you want to manage and enforce QoS policies that are based on destination ports. If your WAN router is following traffic classification rules written in terms of client and network addresses, port transparency enables your routers to use existing rules to classify the traffic without any changes. Port transparency enables network analyzers deployed within the WAN (between the Steelhead appliances) to monitor network activity and to capture statistics for reporting by inspecting traffic according to its original TCP port number. Port transparency does not require dedicated port configurations on your Steelhead appliances. **Note:** Port transparency only provides server port visibility. It does not provide client and server IP address visibility, nor does it provide client port visibility. • **Full Transparency** - Full address transparency preserves your client and server IP addresses and port numbers in the TCP/IP header fields for optimized traffic in both directions across the WAN. It also preserves VLAN tags. Traffic is optimized while these TCP/IP header fields appear to be unchanged. Routers and network monitoring devices deployed in the WAN segment between the communicating Steelhead appliances can view these preserved fields. If both port transparency and full address transparency are acceptable solutions, port transparency is preferable. Port transparency avoids potential networking risks that are inherent to enabling full address transparency. For details, see the *Riverbed Deployment Guide*. However, if you must see your client or server IP addresses across the WAN, full transparency is your only configuration option. **Important:** Enabling full address transparency requires symmetrical traffic flows between the client and server. If any asymmetry exists on the network, enabling full address transparency might yield unexpected results, up to and including loss of connectivity. For details, see the *Riverbed Deployment Guide*.

(6 of 7)

Control	Description
WAN Visibility Mode (*continued*)	RiOS v6.0 and later includes an option for using Full Transparency with a stateful firewall. A stateful firewall examines packet headers, stores information, and then validates subsequent packets against this information. If your system uses a stateful firewall, the following option is available: • **Full Transparency w/Reset** - Enables full address and port transparency and also sends a forward reset between receiving the probe response and sending the transparent inner channel SYN. This ensures the firewall does not block inner transparent connections because of information stored in the probe connection. The forward reset is necessary because the probe connection and inner connection use the same IP addresses and ports and both map to the same firewall connection. The reset clears the probe connection created by the Steelhead appliance and allows for the full transparent inner connection to traverse the firewall. Both the client-side and server-side Steelhead appliances must be running RiOS v6.0 or later. **Notes:** • For details on configuring WAN visibility and its implications, see the *Riverbed Deployment Guide*. • WAN visibility works with auto-discover in-path rules only. It does not work with fixed-target rules or server-side out-of-path Steelhead appliance configurations. • To turn full transparency on globally by default, create an in-path auto-discover rule, select Full, and place it above the default in-path rule and after the Secure, Interactive, and RBT-Proto rules. • You can configure a Steelhead appliance for WAN visibility even if the server-side Steelhead appliance does not support it, but the connection is not transparent. • You can enable full transparency for servers in a specific IP address range and you can enable port transparency on a specific server. For details, see the *Riverbed Deployment Guide*. • The Top Talkers report displays statistics on the most active, heaviest users of WAN bandwidth, providing some WAN visibility without enabling a WAN Visibility Mode.
Description	Describe the rule to facilitate administration.
Add	Adds the rule to the list. The Management Console redisplays the In-Path Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected Rules	Select the check box next to the name and click **Remove Selected Rules**.
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

(7 of 7)

Tip: The default rule, Auto, which optimizes all remaining traffic that has not been selected by another rule, cannot be removed and is always listed last.

To edit an in-path rule

1. Choose Configure > Optimization > In-Path Rules to display the In-Path Rules page.

2. Select the rule number in the rule list.

3. Edit the rule.

4. Click **Save** to save your settings permanently.

After the Management Console has applied your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details, see "Managing Configuration Files" on page 358.

Related Topics

- "About In-Path Rules" on page 27
- "About Default In-Path Rules" on page 28
- "Configuring General Service Settings" on page 60
- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring Port Labels" on page 94
- "Configuring HTTP Optimization" on page 111
- "Secure Inner Channel Overview" on page 231
- "Viewing Current Connections" on page 385
- "Viewing Connection History" on page 400

CHAPTER 3

Modifying Host and Network Interface Settings

This chapter describes how to configure host and network interface settings. You initially set these properties when you ran the installation wizard. This section describes how you can view and modify these settings, if needed. It includes the following sections:

- "Modifying General Host Settings" on page 39
- "Modifying Base Interfaces" on page 43
- "Modifying In-Path Interfaces" on page 50

Modifying General Host Settings

You can view and modify general host settings in the Configure > Networking > Host Settings page.

When you initially ran the installation wizard, you set required network host settings for the Steelhead appliance. Use the following groups of controls on this page only if modifications or additional configuration is required:

- **Name** - Modify the hostname only if your deployment requires it.
- **DNS Settings** - Riverbed recommends you use DNS resolution.
- **Hosts** - If you do not use DNS resolution, or if the host does not have a DNS entry, you can create a host-IP address resolution map.
- **Proxies** - Configure proxy addresses for Web or FTP proxy access to the Steelhead appliance.
- **Date and Time** - Riverbed recommends you configure NTP time synchronization.

To modify general host settings

- Choose Configure > Networking > Host Settings to display the Host Settings page.

Figure 3-1. **Host Settings Page**

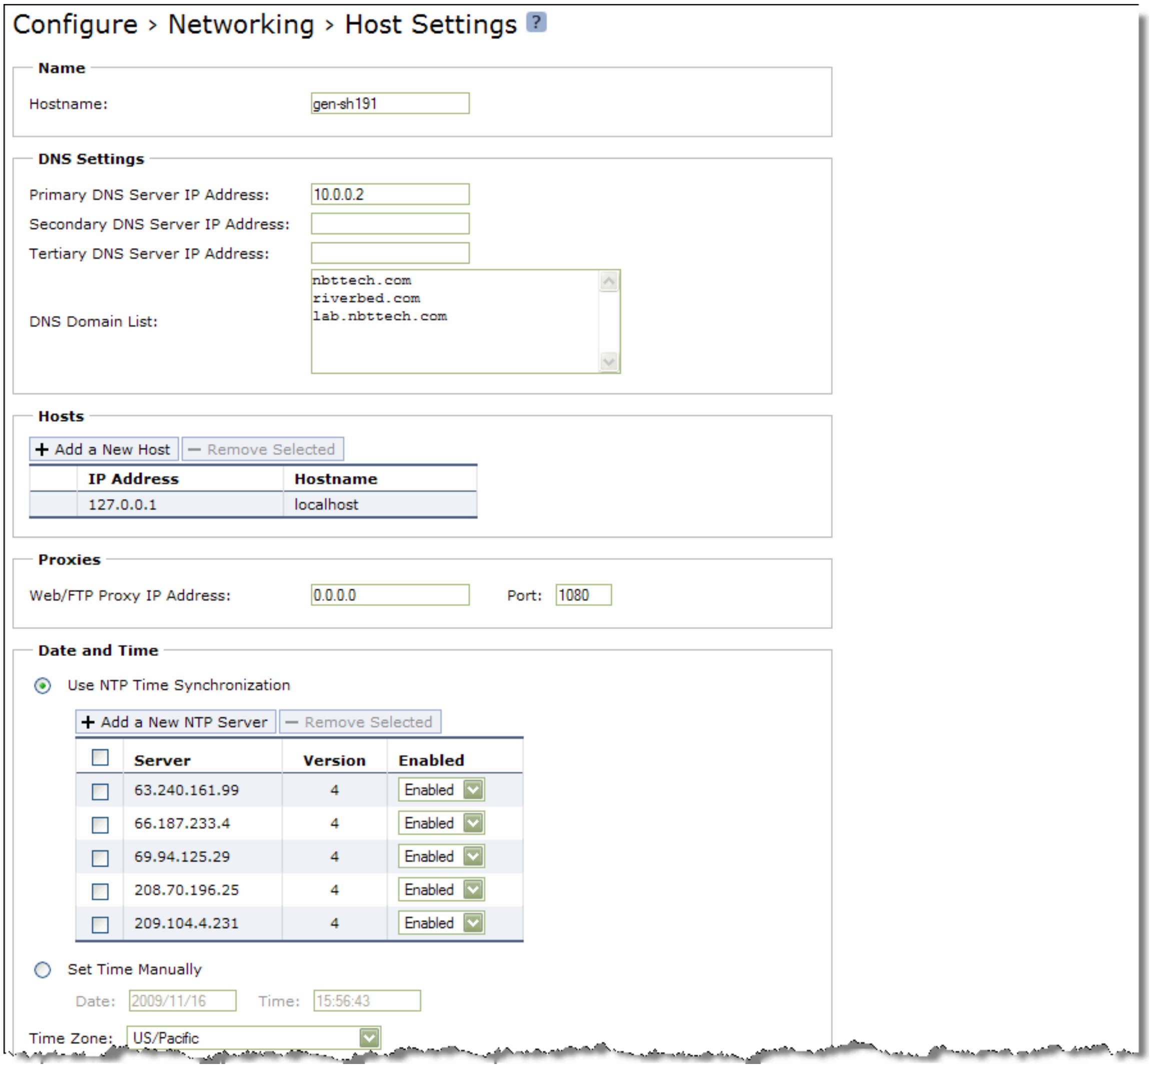

Configure > Networking > Host Settings ?

Name

Hostname: gen-sh191

DNS Settings

Primary DNS Server IP Address: 10.0.0.2

Secondary DNS Server IP Address:

Tertiary DNS Server IP Address:

DNS Domain List:
nbttech.com
riverbed.com
lab.nbttech.com

Hosts

IP Address	Hostname
127.0.0.1	localhost

Proxies

Web/FTP Proxy IP Address: 0.0.0.0 Port: 1080

Date and Time

● Use NTP Time Synchronization

□	Server	Version	Enabled
□	63.240.161.99	4	Enabled
□	66.187.233.4	4	Enabled
□	69.94.125.29	4	Enabled
□	208.70.196.25	4	Enabled
□	209.104.4.231	4	Enabled

○ Set Time Manually

Date: 2009/11/16 Time: 15:56:43

Time Zone: US/Pacific

To change the hostname

1. Choose Configure > Networking > Host Settings to display the Host Settings page.

2. Under Name, modify the value in the Hostname field.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To specify DNS settings

1. Choose Configure > Networking > Host Settings to display the Host Settings page.

2. Under DNS Settings, complete the configuration as described in the following table.

Control	Description
Primary DNS Server IP Address	Specify the IP address for the primary name server.
Secondary DNS Server IP Address	Optionally, specify the IP address for the secondary name server.
Tertiary DNS Server IP Address	Optionally, specify the IP address for the tertiary name server.
DNS Domain List	Specify an ordered list of domain names. If you specify domains the system automatically finds the appropriate domain for each of the hosts that you specify in the system.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To add a new host

1. Choose Configure > Networking > Host Settings to display the Host Settings page.

2. Under Hosts, complete the configuration as described in the following table.

Control	Description
Add a New Host	Displays the controls for adding a new host.
IP Address	Specify the IP address for the host.
Hostname	Specify a hostname.
Add	Adds the host.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To set a Web proxy

1. Choose Configure > Networking > Host Settings to display the Host Settings page.

2. Under Proxies, complete the configuration as described in the following table.

Control	Description
Web/FTP Proxy IP Address	Specify the IP address for the Web/FTP proxy.
Port	Specify the port for the Web/FTP proxy.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To use NTP Time Synchronization

1. Choose Configure > Networking > Host Settings to display the Host Settings page.
2. Under Date and Time, click **Use NTP Time Synchronization**.
3. As a best practice, you should configure your own internal NTP servers; however, if you want to use the Riverbed-provided NTP server, the hard-coded IP address that is pre-configured into every Steelhead appliance is 208.70.196.25. This IP address appears in the NTP server list.

   To add a new NTP server, complete the configuration as described in the following table.

Control	Description
Add a New NTP Server	Displays the controls to add a server.
Hostname or IP Address	Specify the hostname or IP address for the NTP server.
Version	Select the NTP server version from the drop-down list: 3 or 4.
Enabled/Disabled	Enables or disables the connection to the NTP server.
Add	Adds the NTP server to the table list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

4. Click **Save** to save your settings permanently.

Tip: To modify server properties, select the server name in the server table row.

To set the time manually

1. Choose Configure > Networking > Host Settings to display the Host Settings page.
2. Under Date and Time, click **Set Time Manually**.
3. Complete the configuration as described in the following table.

Control	Description
Date	Specify the date in the following format: YYYY/MM/DD
Time	Specify military time in the following format: HH:MM:SS
Time Zone	Select a time zone from the drop-down list. The default value is GMT. **Note:** If you change the time zone, log messages retain the old time zone until you reboot the Steelhead appliance.

4. Click **Apply** to apply your changes to the running configuration.
5. Click **Save** to save your settings permanently.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Modifying Base Interfaces

You can view and modify settings for the appliance primary and auxiliary interfaces in the Configure > Networking > Base Interfaces page.

When you initially ran the Configuration wizard, you set required settings for the base interfaces for the Steelhead appliance. Only use the following groups of controls on this page if modifications or additional configuration is required:

- **Primary Interface** - On the appliance, the primary interface is the port you connect to the LAN switch. The primary interface is the appliance management interface. You connect to the primary interface to use the Web UI or the CLI.
- **Auxiliary Interface** - On the appliance, the auxiliary interface is an optional port you can use to connect the appliance to a non-Riverbed network management device. The IP address for the auxiliary interface must be on a subnet different from the primary interface subnet.
- **Main Routing Table** - Displays a summary of the main routing table for the appliance. If necessary, you can add static routes that might be required for out-of-path deployments or particular device management subnets.

To display and modify the configuration for base interfaces

1. Choose Configure > Networking > Base Interfaces to display the Base Interfaces page.

Figure 3-2. **Base Interfaces Page**

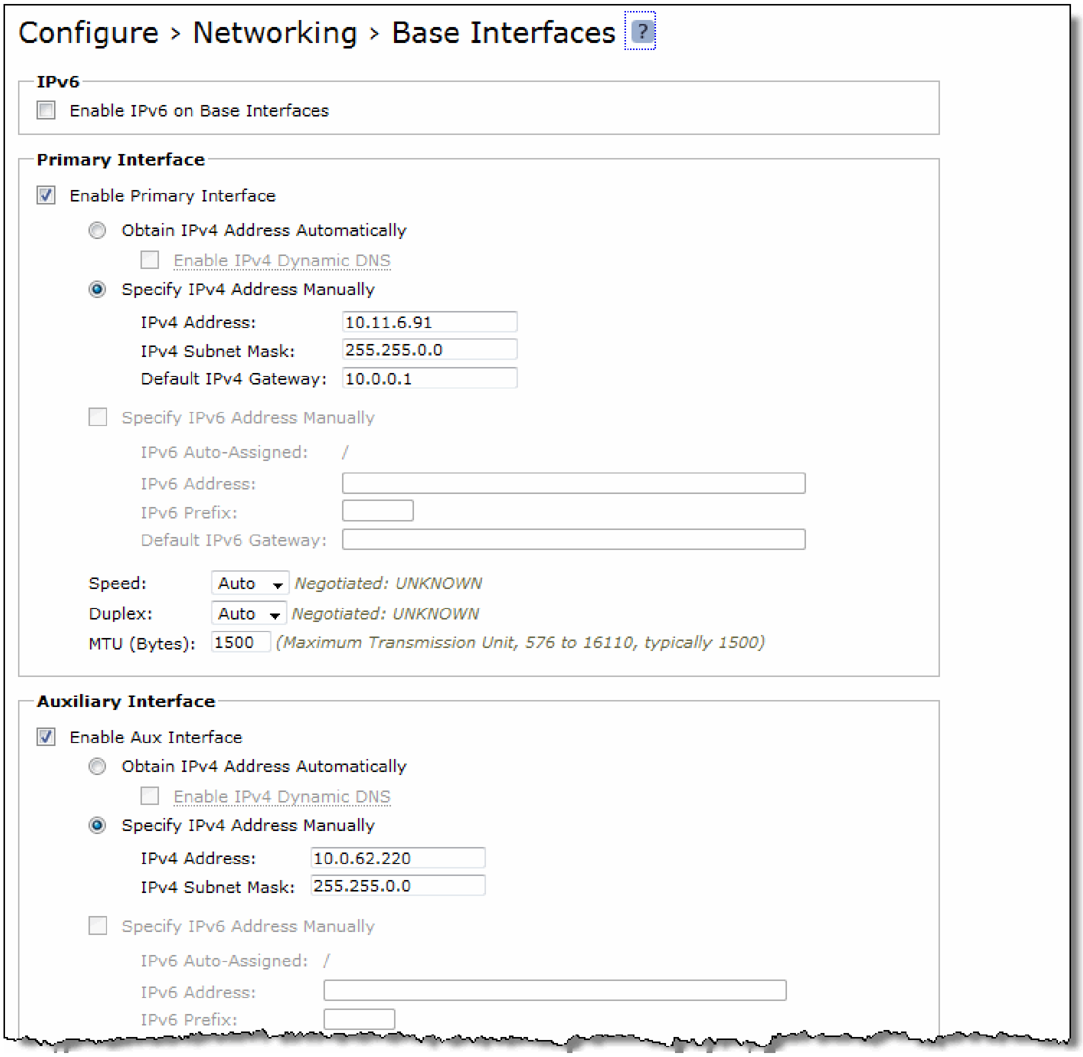

Configure > Networking > Base Interfaces ?

IPv6

☐ Enable IPv6 on Base Interfaces

Primary Interface

☑ Enable Primary Interface

○ Obtain IPv4 Address Automatically

☐ Enable IPv4 Dynamic DNS

● Specify IPv4 Address Manually

IPv4 Address:	10.11.6.91
IPv4 Subnet Mask:	255.255.0.0
Default IPv4 Gateway:	10.0.0.1

☐ Specify IPv6 Address Manually

IPv6 Auto-Assigned:	/
IPv6 Address:	
IPv6 Prefix:	
Default IPv6 Gateway:	

Speed:	Auto	*Negotiated: UNKNOWN*
Duplex:	Auto	*Negotiated: UNKNOWN*
MTU (Bytes):	1500	*(Maximum Transmission Unit, 576 to 16110, typically 1500)*

Auxiliary Interface

☑ Enable Aux Interface

○ Obtain IPv4 Address Automatically

☐ Enable IPv4 Dynamic DNS

● Specify IPv4 Address Manually

IPv4 Address:	10.0.62.220
IPv4 Subnet Mask:	255.255.0.0

☐ Specify IPv6 Address Manually

IPv6 Auto-Assigned:	/
IPv6 Address:	
IPv6 Prefix:	

2. RiOS v6.5 provides the option to enable IPv6 on base interfaces. To enable IPv6, complete the configuration as described in the following table.

Control	Description

Enable IPv6 on Base Interfaces	Enables configuration of IPv6 addresses on the primary and auxiliary interfaces. After enabling IPv6 and specifying the IPv6 addresses address and appropriate routing, you can log in to the Steelhead Management Console and Riverbed Command-Line Interface (CLI) using an IPv6 address and perform maintenance tasks from an IPv6-enabled node.
	Note: When you enable IPv6, RiOS automatically generates a link-local IPv6 address for the Primary and Auxilary interfaces. This link-local address appears in the Management Console. You can also display it by entering the **show interface primary** and **show interface aux** CLI commands.
	By default, IPv6 is disabled.
	The basic steps are:
	1. Choose Configure > Networking > Base Interfaces and enable IPv6 on base interfaces.
	2. Save the configuration.
	3. Reboot the Steelhead appliance.
	4. Choose Configure > Networking > Base Interfaces.
	5. Under either the Primary or Auxiliary Interface, select Specify IPv6 Address Manually, and specify the IPv6 addresses (one per interface).
	6. If necessary, configure additional IPv6 routes.
	When the configuration is complete, you can:
	• Access the Steelhead Management Console and CLI through the configured IPv6 address.
	• Resolve IPv6 addressed hostnames.
	• Log in through an IPv6 interface with AAA (as long as the AAA server is configured with an IPv4 address).
	• Use the following commands to check connectivity with another IPv6-enabled system, such as a Steelhead appliance or a Windows or Linux box:
	`ping6`
	`traceroute6`
	To disable IPv6, clear the Enable IPv6 on Base Interfaces check box. Save the configuration and reboot the Steelhead appliance.
	Notes:
	• Because the IPv6 addresses are limited to the management interfaces, network interfaces related to optimization have no knowledge of IPv6.
	• You can configure only one IPv6 address for each management network interface. You can use IPv4 addresses on the same interface.
	• You cannot configure IPv6 addresses on a management in-path interface.
	• Steelhead appliances do not support auto-configuration.
	• You can only use IPv6 addresses on the management interfaces for management functions. Features like out-of-path optimization and datastore synchronization on management interfaces must use IPv4 addresses.

- Under Primary Interface, complete the configuration as described in the following table.

Control	Description
Enable Primary Interface	Enables the appliance management interface, which can be used for both managing the Steelhead appliance and serving data for a server-side out-of-path (OOP) configuration.
Obtain IPv4 Address Automatically	Select this option to automatically obtain the IP address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it. **Important:** The primary and in-path interfaces can share the same network subnet. The primary and auxiliary interfaces cannot share the same network subnet.
Enable IPv4 Dynamic DNS	Select this option to send the hostname with the DHCP request for registration with Dynamic DNS. The hostname is specified in the Configure > Networking > Host Settings page.
Specify IPv4 Address Manually	Select this option if you do not use a DHCP server to set the IPv4 address. Specify the following settings: • **IPv4 Address** - Specify an IP address. • **IPv4 Subnet Mask** - Specify a subnet mask. • **Default IPv4 Gateway** - Specify the default gateway IPv4 address. The default gateway must be in the same network as the primary interface. You must set the default gateway for in-path configurations.

Specify IPv6 Address Manually	Select this option and specify the following settings to set an IPv6 address.
	• **IPv6 Auto-Assigned** - Displays the link-local address that is automatically generated when IPv6 is enabled on the base interfaces.
	• **IPv6 Address** - Specify an IP address using the following format: eight 16-bit hex strings separated by colons, 128-bits. For example:
	`2001:38dc:0052:0000:0000:e9a4:00c5:6282`
	You do not need to include leading zeros. For example:
	`2001:38dc:52:0:0:e9a4:c5:6282`
	You can replace consecutive zero strings with double colons (::). For example:
	`2001:38dc:52::e9a4:c5:6282`
	• **IPv6 Prefix** - Specify a prefix. The prefix length is 0 to 128, separated from the address by a forward slash (/). In the following example, 60 is the prefix:
	`2001:38dc:52::e9a4:c5:6282/60`
	• **Default IPv6 Gateway** - Specify the default gateway IP address. The default gateway must be in the same network as the primary interface.
	Note: You cannot set an IPv6 address dynamically using a DHCP server.

Control	Description
Speed and Duplex	**Speed** - Select a speed from the drop-down list. The default value is Auto. **Duplex** - Select Auto, Full, or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them manually. The speed and duplex must match (LAN and WAN) in an in-path configuration. If they do not match, you might have a large number of errors on the interface when it is in bypass mode, because the switch and the router are not set with the same duplex settings.
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.

- Under Auxiliary Interface, complete the configuration as described in the following table.

Control	Description
Enable Aux Interface	Enables an auxiliary interface, which can only be used for managing the Steelhead appliance. It cannot be used for an out-of-path (OOP) Steelhead appliance data service. Typically this is used for device-management networks.
Obtain IPv4 Address Automatically	Select this option to automatically obtain the IP address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it. **Important:** The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.
Enable IPv4 Dynamic DNS	Select this option to send the hostname with the DHCP request for registration with Dynamic DNS. The hostname is specified in the Configure > Networking > Host Settings page.
Specify IPv4 Address Manually	Select this option if you do not use a DHCP server to set the IPv4 address. Specify the following settings: • **IPv4 Address** - Specify an IP address. • **IPv4 Subnet Mask** - Specify a subnet mask.

Specify IPv6 Address Manually	Select this option and specify the following settings to set an IPv6 address.
	• **IPv6 Auto-Assigned** - Displays the link-local address that is automatically generated when IPv6 is enabled on the base interfaces.
	• **IPv6 Address** - Specify an IP address, using the following format: eight 16-bit hex strings separated by colons, 128-bits. For example:
	`2001:38dc:0052:0000:0000:e9a4:00c5:6282`
	You do not need to include leading zeros. For example:
	`2001:38dc:52:0:0:e9a4:c5:6282`
	You can replace consecutive zero strings with double colons (::). For example:
	`2001:38dc:52::e9a4:c5:6282`
	• **IPv6 Prefix** - Specify a prefix. The prefix length is 0 to 128, separated from the address by a forward slash (/). In the following example, 60 is the prefix:
	`2001:38dc:52::e9a4:c5:6282/60`
	Note: You cannot set an IPv6 address dynamically using a DHCP server.

Control	Description
Speed and Duplex	**Speed** - Select the speed from the drop-down list. The default value is Auto. **Duplex** - Select Auto, Full or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them on the device manually. The speed and duplex must match (LAN and WAN) in an in-path configuration. To avoid a speed and duplex mismatch, configure your LAN external pair to match the WAN external pair.
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.

7. Click **Apply** to apply your changes to the running configuration.

8. Click **Save** to save your changes permanently.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

To configure routes for IPv4

- Under the Main IPv4 Routing Table you can configure a static routing in the main routing table for out-of-path deployments or if your device management network requires static routes.

  You can add or remove routes from the table list as described in the following table.

Control	Description
Add a New Route	Displays the controls for adding a new route.
Destination IPv4 Address	Specify the destination IP address for the out-of-path appliance or network management device.
IPv4 Subnet Mask	Specify the subnet mask.
Gateway IPv4 Address	Specify the IP address for the gateway. The gateway must be in the same network as the primary or auxiliary interface you are configuring.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

The Management Console writes your configuration changes to memory.

To configure routes for IPv6

- Under the Main IPv6 Routing Table you can configure static routing in the main routing table if your device management network requires static routes.

You can add or remove routes from the table list as described in the following table.

Control	Description
Add a New Route	Displays the controls for adding a new route.
Destination IPv6 Address	Specify the destination IP address.
IPv6 Prefix	Specify a prefix. The prefix length is 0 to 128, separated from the address by a forward slash (/).
Gateway IPv6 Address	Specify the IP address for the gateway. The gateway must be in the same network as the primary or auxiliary interface you are configuring.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

The Management Console writes your configuration changes to memory.

Important: You can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Modifying In-Path Interfaces

You can view and modify settings for the appliance in-path interfaces in the Configure > Networking > In-Path Interfaces page. You can also enable a management in-path interface on this page.

You configure in-path interfaces for deployments where the Steelhead appliance is in the direct path (the same subnet) as the client and the server in your network. You also set the in-path gateway (WAN router).

Note: In the Riverbed system, appliances have a unique in-path interface for each pair of LAN/WAN ports. For each appliance, the Management Console detects LAN/WAN pairs, including those added through bypass cards, and identifies them according to slot (for example, inpath0_0, inpath0_1, inpath1_0, inpath1_1, and so on.)

To display and modify the configuration for in-path interfaces

1. Choose Configure > Networking > In-Path Interfaces to display the In-Path Interfaces page.

Figure 3-3. In-Path Interfaces Page

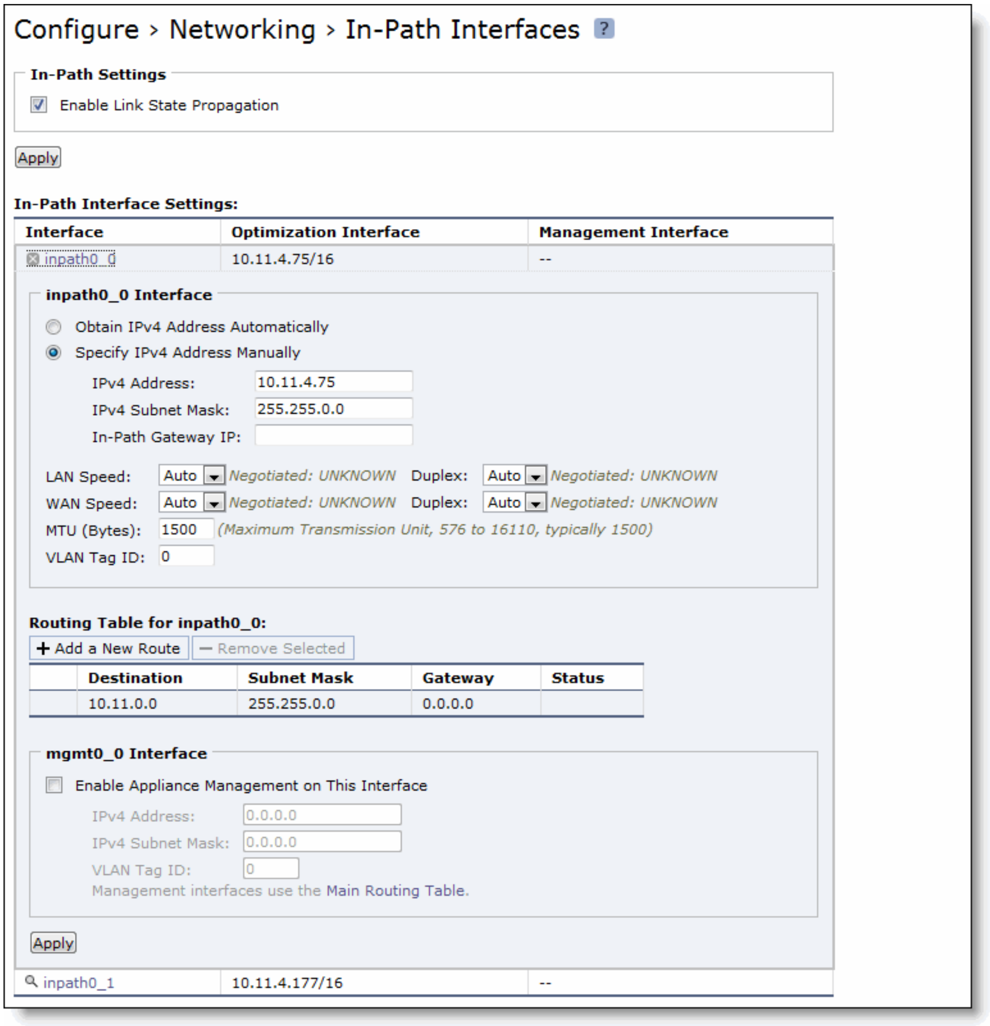

Configure > Networking > In-Path Interfaces ?

In-Path Settings

☑ Enable Link State Propagation

Apply

In-Path Interface Settings:

Interface	Optimization Interface	Management Interface
☒ inpath0_0	10.11.4.75/16	--

inpath0_0 Interface

○ Obtain IPv4 Address Automatically

● Specify IPv4 Address Manually

IPv4 Address: 10.11.4.75

IPv4 Subnet Mask: 255.255.0.0

In-Path Gateway IP:

LAN Speed: Auto *Negotiated: UNKNOWN* Duplex: Auto *Negotiated: UNKNOWN*

WAN Speed: Auto *Negotiated: UNKNOWN* Duplex: Auto *Negotiated: UNKNOWN*

MTU (Bytes): 1500 *(Maximum Transmission Unit, 576 to 16110, typically 1500)*

VLAN Tag ID: 0

Routing Table for inpath0_0:

+ Add a New Route — Remove Selected

	Destination	Subnet Mask	Gateway	Status
	10.11.0.0	255.255.0.0	0.0.0.0	

mgmt0_0 Interface

□ Enable Appliance Management on This Interface

IPv4 Address: 0.0.0.0

IPv4 Subnet Mask: 0.0.0.0

VLAN Tag ID: 0

Management interfaces use the Main Routing Table.

Apply

Interface	Optimization Interface	Management Interface
inpath0_1	10.11.4.177/16	--

2. To enable Link State Propagation, under In-Path Settings, complete the configuration as described in the following table.

Control	Description
Enable Link State Propagation	Enables Link State Propagation (LSP). With LSP enabled, if the LAN interface drops the link, the WAN also drops the link. LSP is enabled by default. If you require a Steelhead appliance to fail-to-wire (bypass) when the LAN or WAN ports become disconnected, enable this feature. This feature is similar to what ISPs do to follow the state of a link. **Note:** You cannot reach a MIP interface when Link State Propagation (LSP) is also enabled and the corresponding in-path interface fails. In physical in-path deployments, LSP shortens the recovery time of a link failure. LSP communicates link status between the devices connected to the Steelhead appliance and is enabled by default in RiOS v6.0 and later. CSH VSH Virtual Steelhead and Cloud Steelhead models do not support LSP.

3. Under In-Path Interface Settings, select the interface name and complete the configuration as described in the following table.

Control	Description
Obtain IPv4 Address Automatically	Specify this option to automatically obtain the IP address from a DHCP server. (A DHCP server must be available so that the Steelhead appliance can request the IP address from it.) **Important:** The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.
Specify IPv4 Address Manually	Specify the following settings if you do not use a DHCP server to set the IP address: • **IPv4 Address** - Specify an IP address. This IP address is the in-path main interface. • **IPv4 Subnet Mask** - Specify the subnet mask. • **In-Path Gateway IP** - Specify the IP address for the in-path gateway. If you have a router (or a Layer-3 switch) on the LAN side of your network, specify this device as the in-path gateway. **Important:** If there is a routed network on the LAN-side of the in-path appliance, the router that is the default gateway for the appliance must not have the ACL configured to drop packets from the remote hosts as its source. The in-path appliance uses IP masquerading to appear as the remote server.
LAN Speed and Duplex WAN Speed and Duplex	**Speed** - Select Auto, 1000, 100, or 10 from the drop-down list. The default value is Auto. **Duplex** - Select Auto, Full, or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them on the device manually. The speed and duplex must match (LAN and WAN) in an in-path configuration. To avoid a speed and duplex mismatch, configure your LAN external pair to match the WAN external pair. **Note:** Speed and duplex mismatches can easily occur in a network. For example, if one end of the link is set at half or full-duplex and the other end of the link is configured to auto negotiate (auto), the link defaults to half-duplex, regardless of the duplex setting on the non-auto-negotiated end. This duplex mismatch passes traffic, but it causes interface errors and results in degraded optimization. The following guidelines can help you avoid speed and duplex mismatches when configuring the Steelhead appliance: • Routers are often configured with fixed speed and duplex settings. Check your router configuration and set it to match the Steelhead appliance WAN and LAN settings. Make sure your switch has the correct setting. • After you finish configuring the Steelhead appliance, check for speed and duplex error messages (crc or frame errors) in the System Log page of the Management Console. • If there is a serious problem with the Steelhead appliance and it goes into bypass mode (that is, it automatically continues to pass traffic through your network), a speed and duplex mismatch might occur when you reboot the Steelhead appliance. To avoid a speed and duplex mismatch, configure your LAN external pair to match the WAN external pair.

Control	Description
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. Applies to optimized traffic only. The default value is 1500.
VLAN Tag ID	Specify a numeric VLAN Tag ID. When you specify the VLAN Tag ID for the MIP interface, all packets originating from the Steelhead appliance are tagged with that identification number. Specify the VLAN tag that the appliance uses to communicate with other Steelhead appliances in your network. The VLAN Tag ID might be the same value or a different value than the VLAN tag used on the client. A zero (0) value specifies non-tagged (or native VLAN) and is the correct setting if there are no VLANs present. For example, if the in-path interface is 192.168.1.1 in VLAN 200, you would specify tag 200. **Note:** When the Steelhead appliance communicates with a client or a server it uses the same VLAN tag as the client or the server. If the Steelhead appliance cannot determine which VLAN the client or server is in, it uses its own VLAN until it is able to determine that information. You must also define in-path rules to apply to your VLANs.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

4. Under Routing Table for <interface name>, you can configure a static routing table for in-path interfaces. You can add or remove routes from the table list.

Control	Description
Add a New Route	Displays the controls to add a route.
Destination IP Address	Specify the destination IP address.
Subnet Mask	Specify the subnet mask.
Gateway IP Address	Specify the IP address for the gateway. The gateway must be in the same network as the in-path interface.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

5. Click **Apply** to apply your changes to the running configuration.

6. Click **Save** to save your settings permanently.

Configuring a Management In-Path Interface

You can configure a Management In-Path (MIP) interface in the Configure > Networking > InPath <slot> page.

In a typical in-path deployment, optimized and pass-through traffic flows through the Steelhead appliance LAN and WAN interfaces and Riverbed network management traffic flows through the auxiliary interface. You can also use the auxiliary interface to connect the appliance to a non-Riverbed network management device. Some deployments do not allow access to the auxiliary management interface when plugged into a private subnet with a separate IP address space. In this type of deployment you cannot use the auxiliary interface to manage the Steelhead appliance.

RiOS v6.1 and later provides a way to configure a secondary MIP interface that you can reach through the physical in-path LAN and WAN interfaces. Configuring a secondary MIP interface is a way to manage Steelhead appliances from a private network while maintaining a logical separation of network traffic. This configuration eliminates the need to deploy a switch or borrow a switchport. You can configure one MIP interface for each LAN and WAN interface pair.

A MIP interface is accessible from both the LAN and WAN side and you can reach it even when:

- the primary interface is unavailable.
- the optimization service is not running.
- the (logical) in-path interface fails.

A MIP interface is not accessible if the (physical) LAN and WAN interfaces fail.

Note: You cannot configure IPv6 addresses on a Management In-Path interface.

Figure 3-4. Management In-Path Interface Deployment

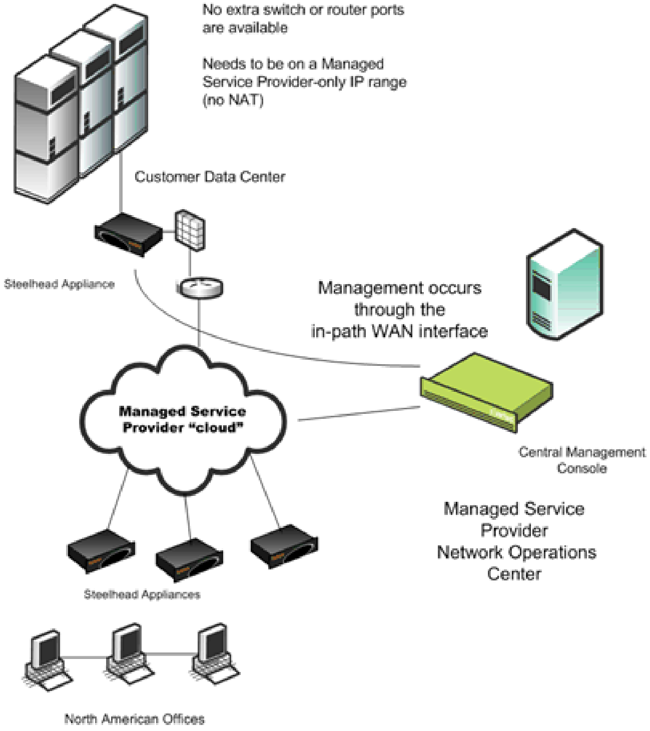
No extra switch or router ports are available
Needs to be on a Managed Service Provider-only IP range (no NAT)
Customer Data Center
Steelhead Appliance
Management occurs through the in-path WAN interface
Managed Service Provider "cloud"
Central Management Console
Managed Service Provider Network Operations Center
Steelhead Appliances
North American Offices

Dependencies

- Any connections destined to a MIP interface are not optimized by that Steelhead appliance and do not appear in the Current Connections report.
- A MIP interface cannot reside in the same subnet as the Primary or Auxiliary interfaces.
- A MIP interface must reside in its own subnet. It cannot share the same subnet with any other interfaces on the Steelhead appliance.
- You cannot enable a MIP interface after fail-to-block has been enabled and the corresponding in-path interface fails. When fail-to-block is enabled, in the event of a failure or loss of power, the Steelhead appliance LAN and WAN interfaces completely lose link status. The failed Steelhead appliance blocks traffic along its path, forcing traffic to be re-routed onto other paths (where the remaining Steelhead appliances are deployed). For details on fail-to-block, see the *Riverbed Deployment Guide*.
- You cannot enable a MIP interface when RSP is enabled and vice versa. For details, see "Installing and Configuring RSP" on page 180.

- You cannot reach a MIP interface when Link State Propagation (LSP) is also enabled and the corresponding in-path interface fails. In physical in-path deployments, LSP shortens the recovery time of a link failure. LSP communicates link status between the devices connected to the Steelhead appliance and is enabled by default in RiOS v6.0 and later. To disable LSP, enter the **no in-path lsp enable** CLI command at the system prompt.
- This feature supports 802.1Q VLAN.
- A MIP interface uses the main routing table.

Enabling a MIP Interface

Use the following controls on this page when you need to enable a MIP interface or the interface requires additional configuration.

To configure a management in-path interface

1. Choose Configure > Networking > In-Path <slot> to display the In-Path <slot> page.
2. Under Management <interface name>, complete the configuration as described in the following table.

Control	Description
Enable Appliance Management on This Interface	Enables a secondary MIP interface that you can reach through the physical in-path LAN and WAN interfaces. Configuring a secondary MIP interface allows management of Steelhead appliances from a private network while maintaining a logical separation of network traffic. **Note:** If LSP or fail-to-block is enabled, a message reminds you to disable the feature before enabling the MIP interface.
IPv4 Address	Specify the IP address for the MIP interface.
IPv4 Subnet Mask	Specify the subnet mask.
VLAN Tag ID	Specifies a numeric VLAN Tag ID. When you specify the VLAN Tag ID for the MIP interface, all packets originating from the Steelhead appliance are tagged with that identification number. Specify the VLAN tag that the appliance uses to communicate with other Steelhead appliances in your network. The VLAN Tag ID might be the same value or a different value than the in-path interface VLAN tag ID. The MIP interface could be un-tagged and in-path interface could be tagged and vice versa. A zero (0) value specifies non-tagged (or native VLAN) and is the correct setting if there are no VLANs present. For example, if the MIP interface is 192.168.1.1 in VLAN 200, you would specify tag 200.

3. Click **Apply** to apply your changes to the running configuration.
4. Click **Save** to save your settings permanently.

Tip: After you apply your settings, choose Reports > Networking > Interface Counters to view MIP interface statistics.

Note: You can remove MIP interfaces from the main routing table in the Configure > Networking > Base Interfaces page.

Related Topics

- "Configuring General Service Settings" on page 60
- "Configuring In-Path Rules" on page 28
- "Modifying Base Interfaces" on page 43

CHAPTER 4 Configuring Optimization Features

This chapter describes how to enable and configure optimization features. It includes the following sections:

Configuring General Service Settings

You can configure general optimization service settings in the Configure > Optimization > General Service Settings page.

Enabling Basic Deployment Options

General Service Settings include controls to enable or disable in-path, out-of-path, failover support, and to set connection limits and the maximum connection pooling size.

If you have a Steelhead appliance that contains multiple bypass cards, the Management Console displays options to enable in-path support for these ports. The number of these interface options depends on the number of pairs of LAN and WAN ports that you have enabled in your Steelhead appliance.

The properties and values you set on this page depend on your deployment. For example, the following deployment types would require different choices:

- **Physical In-Path** - The Steelhead appliance is physically in the direct path between the client and the server. The clients and servers continue to see client and server IP addresses. Physical in-path configurations are suitable for any location where the total bandwidth is within the limits of the installed Steelhead appliance.
- **Virtual In-Path** - The Steelhead appliance is virtually in the path between the client and the server. This differs from a physical in-path in that a packet redirection mechanism is used to direct packets to Steelhead appliances that are not in the physical path. Redirection mechanisms include WCCP, Layer-4 switches, and PBR. In this configuration, clients and servers continue to see client and server IP addresses.
- **Out-of-Path** - The Steelhead appliance is not in the direct path between the client and the server. Servers see the IP address of the server-side Steelhead appliance rather than the client IP address, which might impact security policies. An out-of-path configuration is suitable for data center locations where physically in-path or virtually in-path configurations are not possible.

For an overview of in-path and out-of-path deployment options, see the *Riverbed Deployment Guide*.

Enabling Failover

In the event of appliance failure, the Steelhead appliance enters bypass mode to avoid becoming a single point of failure in your network. If you want optimization to continue in the event of appliance failure, you can deploy redundant appliances as failover buddies.

For details on failover redundancy, see the *Riverbed Deployment Guide*.

Physical In-Path Failover Deployment

For a physical in-path failover deployment, you configure a pair of Steelhead appliances: one as a master and the other as a backup. The master Steelhead appliance in the pair (usually the Steelhead appliance closest to the LAN) is active and the backup Steelhead appliance is passive. The master Steelhead appliance is active unless it fails for some reason. The backup is passive while the master is active and becomes active if either the master fails or the master reaches its connection limit and enters *admission control* status. A backup Steelhead appliance does not intercept traffic while the master appliance is active. It pings the master Steelhead appliance to make sure that it is alive and processing data. If the master Steelhead appliance fails, the backup takes over and starts processing all of the connections. When the master Steelhead appliance comes back up, it sends a message to the backup that it has recovered. The backup Steelhead appliance stops processing new connections (but continues to serve old ones until they end).

Out-of-Path Failover Deployment

For an out-of-path failover deployment, you deploy two server-side Steelhead appliances and add a fixed-target rule to the client-side Steelhead appliance to define the master and backup target appliances. When both the master and backup Steelhead appliances are functioning properly, the connections traverse the master appliance. If the master Steelhead appliance fails, subsequent connections traverse the backup Steelhead appliance.

The master Steelhead appliance uses an Out-of-Band (OOB) connection. The OOB connection is a single, unique TCP connection that communicates internal information only; it does not contain optimized data. If the master Steelhead appliance becomes unavailable, it loses this OOB connection and the OOB connection times out in approximately 40-45 seconds. Once the OOB connection times out, the client-side Steelhead appliance declares the master Steelhead appliance unavailable and connects to the backup Steelhead appliance.

During the 40-45 second delay before the client-side Steelhead appliance declares a peer unavailable, it passes through any incoming new connections; they are not blackholed.

While the client-side Steelhead appliance is using the backup Steelhead appliance for optimization, it attempts to connect to the master Steelhead appliance every 30 seconds. If the connection succeeds, the client-side Steelhead appliance reconnects to the master Steelhead appliance for any new connections. Existing connections remain on the backup Steelhead appliance for their duration. This is the only time, (immediately after a recovery from a master failure), that connections are optimized by both the master Steelhead appliance and the backup.

If both the master and backup Steelhead appliances become unreachable, the client-side Steelhead appliance tries to connect to both appliances every 30 seconds. Any new connections are passed through the network unoptimized.

Synchronizing Master and Backup Failover Pairs

In addition to enabling failover and configuring buddy peering, you must synchronize the datastores for the master-backup pairs to ensure optimal use of SDR for *warm* data transfer. With warm transfers, only new or modified data is sent, dramatically increasing the rate of data transfer over the WAN. For information on synchronizing datastores for master-backup pairs, see “Synchronizing Peer Datastores” on page 78.

Configuring Connection Limits

In the General Service Settings page, you can also modify default settings for the maximum half-opened connections from a single source IP address and the connection pool size. For details, pay careful attention to the configuration descriptions included in the following procedure.

To configure general optimization service settings

1. Choose Configure > Optimization > General Service Settings to display the General Service Settings page.

Figure 4-1. General Service Settings Page

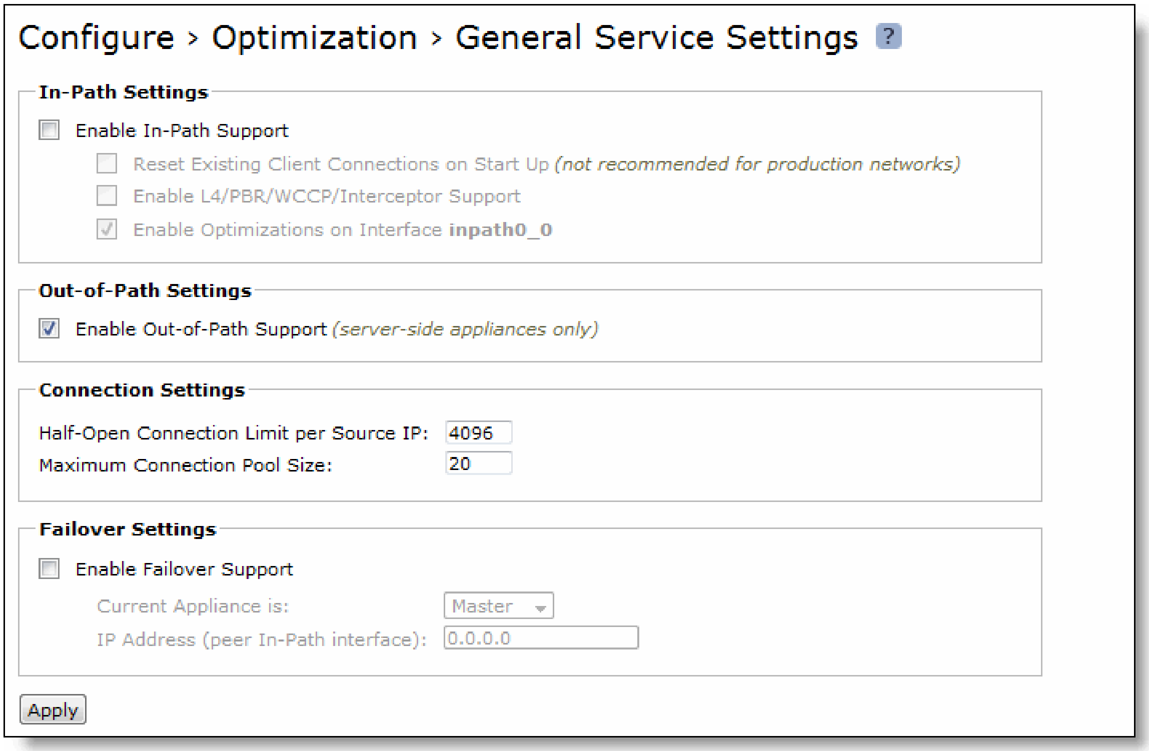

Configure > Optimization > General Service Settings ?

In-Path Settings

□ Enable In-Path Support

□ Reset Existing Client Connections on Start Up *(not recommended for production networks)*

□ Enable L4/PBR/WCCP/Interceptor Support

☑ Enable Optimizations on Interface **inpath0_0**

Out-of-Path Settings

☑ Enable Out-of-Path Support *(server-side appliances only)*

Connection Settings

Half-Open Connection Limit per Source IP: 4096

Maximum Connection Pool Size: 20

Failover Settings

□ Enable Failover Support

Current Appliance is: Master

IP Address (peer In-Path interface): 0.0.0.0

2. Under In-Path Settings, complete the configuration as described in the following table.

Control	Description
Enable In-Path Support	Enables optimization on traffic that is in the direct path of the client, server, and Steelhead appliance.
Reset Existing Client Connections on Start Up	Enables *kickoff* globally. If you enable kickoff, connections that exist when the Steelhead service is started and restarted are disconnected. When the connections are retried they are optimized. Generally, connections are short lived and kickoff is not necessary. It is suitable for very challenging remote environments. In a remote branch-office with a T1 and 35 ms round-trip time, you would want connections to migrate to optimization gracefully, rather than risk interruption with kickoff. **Note:** RiOS v6.1 and later provides a way to reset pre-existing connections that match an in-path rule and the rule has kickoff enabled. You can also reset a single pass-through or optimized connection in the Current Connections report, one connection at a time. **Note:** Do not enable kickoff for in-path Steelhead appliances that use auto-discover or if you do not have a Steelhead appliance on the remote side of the network. If you do not set any in-path rules the default behavior is to auto-discover all connections. If kickoff is enabled, all connections that existed before the Steelhead appliance started are reset.

Control	Description
Enable L4/PBR/WCCP Interceptor Support	Enables optional, virtual in-path support on all the interfaces for networks that use Layer-4 switches, PBR, WCCP, and Interceptor. External traffic redirection is supported only on the first in-path interface. The following redirection methods are available: • **Layer-4 Switch** - You enable Layer-4 switch support when you have multiple Steelhead appliances in your network, so that you can manage large bandwidth requirements. • **Policy-Based Routing (PBR)** - PBR allows you to define policies to route packets instead of relying on routing protocols. You enable PBR to redirect traffic that you want optimized by a Steelhead appliance that is not in the direct physical path between the client and server. • **Web Cache Communication Protocol (WCCP)** - If your network design requires you to use WCCP, a packet redirection mechanism directs packets to RiOS appliances that are not in the direct physical path to ensure that they are optimized. For details about configuring Layer-4 switch, PBR, and WCCP deployments, see the *Riverbed Deployment Guide*. CSH The AWS Cloud Steelhead does not support L4/PBR/WCCP and Interceptor, but the ESX Cloud Steelhead supports it.
CSH Enable Agent-Intercept This feature is only supported by the Cloud Steelhead.	Select this checkbox to enable configuration of the transparency mode in the Cloud Steelhead and transmit it to the Discovery Agent. The Discovery Agent in the server provides the following transparency modes for client connections: • **Restricted transparent** - All client connections are transparent with the following restrictions: – If the client connection is from a NATed network, the application server sees the private IP address of the client. – You can use this mode only if there is no conflict between the private IP address ranges (there are no duplicate IP addresses) and ports. This is the default mode. • **Safe transparent** - If the client is behind a NAT device, the client connection to the application server is non-transparent—the application server sees the connection as a connection from the Cloud Steelhead IP address and not the client IP address. All connections from a client that is not behind a NAT device are transparent and the server sees the connections from the client IP address instead of the Cloud Steelhead IP address. • **Non-transparent** - All client connections are non-transparent—the application server sees the connections from the server-side Steelhead IP address and not the client IP address. Riverbed recommends that you use this mode as the last option.
Enable Optimizations on Interface <interface_name>	Enables in-path support for additional bypass cards. If you have an appliance that contains multiple two-port, four-port, or six-port bypass cards, the Management Console displays options to enable in-path support for these ports. The number of these interface options depends on the number of pairs of LAN and WAN ports that you have enabled in your Steelhead appliance. The interface names for the bypass cards are a combination of the slot number and the port pairs (inpath<slot>_<pair>, inpath<slot>_<pair>). For example, if a four-port bypass card is located in slot 0 of your appliance, the interface names are: inpath0_0 and inpath0_1. Alternatively, if the bypass card is located in slot 1 of your appliance, the interface names are: inpath1_0 and inpath1_1. For details about installing additional bypass cards, see the *Network Interface Card Installation Guide*.

3. Under Out-of-Path Settings, complete the configuration as described in the following table.

Control	Description
Enable Out-of-Path Support	Enables out-of-path support on a server-side Steelhead appliance, where only a Steelhead appliance primary interface connects to the network. The Steelhead appliance can be connected anywhere in the LAN. There is no redirecting device in an out-of-path Steelhead appliance deployment. You configure fixed-target in-path rules for the client-side Steelhead appliance. The fixed-target in-path rules point to the primary IP address of the out-of-path Steelhead appliance. The out-of-path Steelhead appliance uses its primary IP address when communicating to the server. The remote Steelhead appliance must be deployed either in a physical or virtual in-path mode. If you set up an out-of-path configuration with failover support, you must set fixed-target rules that specify the master and backup Steelhead appliances.

4. Under Connection Settings, complete the configuration as described in the following table.

Control	Description
Half-Open Connection Limit per Source IP	Restricts half-opened connections on a source IP address initiating connections (that is, the client machine). Set this feature to block a source IP address that is opening multiple connections to invalid hosts or ports simultaneously (for example, a virus or a port scanner). This feature does not prevent a source IP address from connecting to valid hosts at a normal rate. Thus, a source IP address could have more established connections than the limit. The default value is 4096. The appliance counts the number of half-opened connections for a source IP address (connections that check if a server connection can be established before accepting the client connection). If the count is above the limit, new connections from the source IP address are passed through unoptimized. **Note:** If you have a client connecting to valid hosts or ports at a very high rate, some of its connections might be passed through even though all of the connections are valid.
Maximum Connection Pool Size	Specify the maximum number of TCP connections in a connection pool. Connection pooling enhances network performance by reusing active connections instead of creating a new connection for every request. Connection pooling is useful for protocols which create a large number of short-lived TCP connections, such as HTTP. To optimize such protocols, a connection pool manager maintains a pool of idle TCP connections, up to the maximum pool size. When a client requests a new connection to a previously visited server, the pool manager checks the pool for unused connections and returns one if available. Thus, the client and the Steelhead appliance do not have to wait for a three-way TCP handshake to finish across the WAN. If all connections currently in the pool are busy and the maximum pool size has not been reached, the new connection is created and added to the pool. When the pool reaches its maximum size, all new connection requests are queued until a connection in the pool becomes available or the connection attempt times out. The default value is 20. A value of 0 specifies no connection pool. **Important:** You must restart the Steelhead appliance after changing this setting. **Tip:** Viewing the Connection Pooling report can help determine whether to modify the default setting. If the report indicates an unacceptably low ratio of pool hits per total connection requests, increase the pool size.

5. Under Failover Settings, complete the configuration as described in the following table.

Control	Description
Enable Failover Support	Configures a failover deployment on either a master or backup Steelhead appliance. In the event of a failure in the master appliance, the backup appliance takes its place with a hot RiOS datastore, and can begin delivering fully-optimized performance immediately. The master and backup Steelhead appliances must be the same hardware model.

Control	Description
Current Appliance is	Select Master or Backup from the drop-down list. A master Steelhead appliance is the primary appliance; the backup Steelhead appliance is the appliance that automatically optimizes traffic if the master appliance fails.
IP Address (peer in-path interface)	Specify the IP address for the master or backup Steelhead appliance. You must specify the in-path IP address (inpath0_0) for the Steelhead appliance, not the primary interface IP address. **Important:** You must specify the inpath0_0 interface as the other appliance's in-path IP Address.

CSH The Cloud Steelhead does not support configuring Failover Settings.

6. Click **Apply** to apply your settings.

7. Click **Save** to save your settings permanently.

Tip: After applying the settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Related Topics

- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring In-Path Rules" on page 28
- "Configuring the Datastore" on page 76
- "Configuring Service Ports" on page 93
- "Configuring Connection Forwarding Features" on page 255
- "Modifying In-Path Interfaces" on page 50
- "Configuring Subnet Side Rules" on page 260

Enabling Peering and Configuring Peering Rules

This section describes how to enable peering and configure peering rules. It includes the following sections:

- "About Regular and Enhanced Auto-Discovery" on page 67
- "Configuring Peering" on page 69

About Regular and Enhanced Auto-Discovery

With enhanced auto-discovery, the Steelhead appliance automatically finds the furthest Steelhead appliance peer in a network and optimization occurs there. By default, auto-discovery is enabled. When auto-discovery is disabled, the Steelhead appliance uses regular auto-discovery. With regular auto-discovery, the Steelhead appliance finds the next appliance in the group and optimization occurs there.

In some deployments, enhanced auto-discovery can simplify configuration and make your deployments more scalable. When enhanced auto-discovery is enabled, the Steelhead appliance automatically finds the furthest Steelhead appliance in a network and optimization occurs there. For example, if you had a deployment with four Steelhead appliance (A, B, C, D), where D represents the appliance that is furthest from A, the Steelhead appliance automatically finds D. This simplifies configuration and makes your deployment more scalable.

CSH The Cloud Steelhead does not use automatic peering. When you run a server in the cloud, you deploy the Cloud Steelhead to be the furthest Steelhead in the network because the Discovery Client on the server is configured to use the Cloud Steelhead automatically. When you run a client in the cloud, and there are multiple Steelheads in the path to the server, the Cloud Steelhead is selected for optimization first. You can enable automatic peering on the remote Steelheads to make the Cloud Steelhead peer with the furthest Steelhead in the network.

Enhanced auto-discovery is recommended for the deployments described in the following table.

Deployment Type	Description
Serial Cascade Deployments	Cascade configurations enable optimal multi-site deployments where connections between the client and the server might pass through intermediate Steelhead appliances to reach their final destination. Enhanced auto-discovery for cascading Steelhead appliances detects when more than two Steelhead appliances are present between the client and the server and automatically chooses the two outside Steelhead appliances, optimizing all traffic in between.
Serial Cluster Deployments	You can provide increased optimization by deploying two or more Steelhead appliances back-to-back in an in-path configuration to create a serial cluster. Appliances in a serial cluster process the peering rules you specify in a spill-over fashion. When the maximum number of TCP connections for a Steelhead appliance is reached, that appliance stops intercepting new connections. This allows the next Steelhead appliance in the cluster the opportunity to intercept the new connection, if it has not reached its maximum number of connections. The in-path peering rules and in-path rules tell the Steelhead appliance in a cluster not to intercept connections between themselves. You configure peering rules that define what to do when a Steelhead appliance receives an auto-discovery probe from another Steelhead appliance. You can deploy serial clusters on the client or server-side of the network. **Supported models** Two-appliance serial clusters are supported for all Steelhead appliance *xx*20 and *xx*50 models, except the 250 model. The Steelhead appliances must be the same model running RiOS v5.5.3 or later or RiOS v5.0.8. The following Steelhead appliance models support serial clusters: 550 series, 1050 series, 2050, 5050, 6050, 7050, 1020, 2020, 3020, 3520, 5000, 5010, 5520, and 6020 These models can reach their specifications even while potentially passing through the LAN-side traffic for optimized connections for the other Steelhead appliance in the cluster. When running a RiOS software version earlier than v5.5.3, models 5520, 6020, and 6120 are qualified by Riverbed for serial clusters. **Important:** For environments that want to optimize MAPI or FTP traffic which require all connections from a client to be optimized by one Steelhead appliance, Riverbed strongly recommends using the master and backup redundancy configuration instead of a serial cluster. For larger environments that require multi-appliance scalability and high availability, Riverbed recommends using the Interceptor to build multi-appliance clusters. For details, see the *Riverbed Deployment Guide* and the *Interceptor Appliance User's Guide*. **Note:** A serial cluster has the same bandwidth specification as the Steelhead appliance model deployed in the cluster. The bandwidth capability does not increase because the cluster contains more than one Steelhead appliance. For example, a serial cluster comprised of two Steelhead appliance 2050M models with a bandwidth specification of 20 Mbps has a bandwidth specification of 20 Mbps. **Note:** If the active Steelhead appliance in the cluster enters a degraded state because the CPU load is too high, it continues to accept new connections.

For details on these deployment types, see the *Riverbed Deployment Guide*.

Extending the Number of Peers

RiOS supports a large number of peers (up to 20,000) per Steelhead appliance. This feature is available only on Steelhead appliance models 5050, 5520, 6020, 6050, 6120, and 7050. Riverbed recommends enabling the extended peer table if you have more than 4,000 peers. After enabling extended peer table support, you must clear the datastore and stop and restart the service. See "Configuring Peering" on page 69.

Configuring Peering

You can display, add, and modify auto-discovery peering settings in the Configure > Optimization > Peering Rules page. You can also enable extended peer table support.

To enable enhanced auto-discovery

1. Choose Configure > Optimization > Peering Rules to display the Peering Rules page.

Figure 4-2. Peering Rules Page

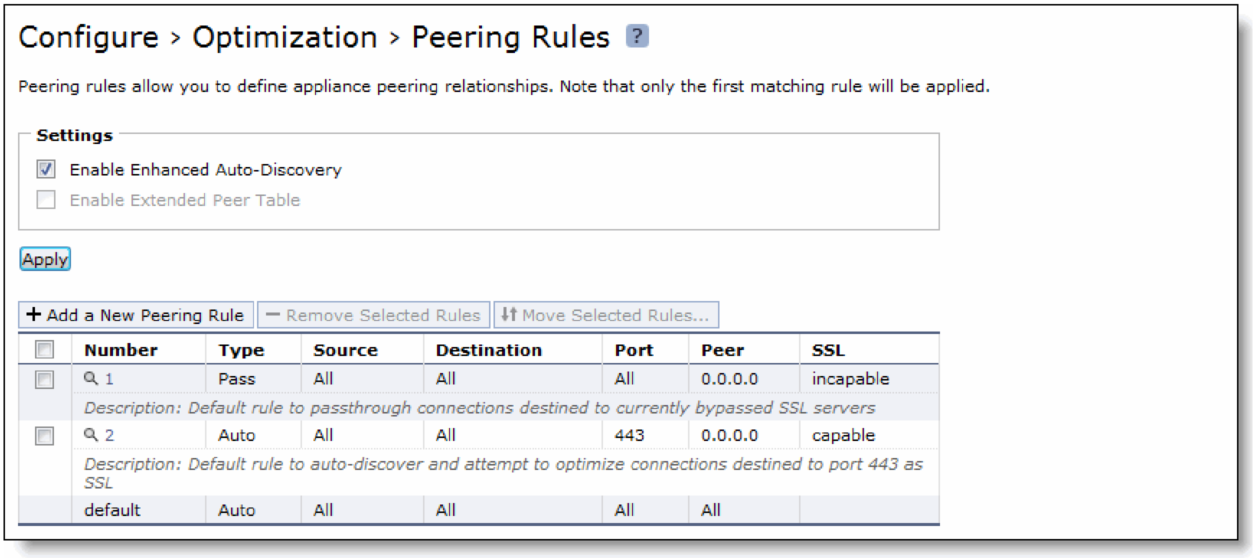

Configure > Optimization > Peering Rules

Peering rules allow you to define appliance peering relationships. Note that only the first matching rule will be applied.

Settings

☑ Enable Enhanced Auto-Discovery

☐ Enable Extended Peer Table

☐	Number	Type	Source	Destination	Port	Peer	SSL
☐	1	Pass	All	All	All	0.0.0.0	incapable
	Description: Default rule to passthrough connections destined to currently bypassed SSL servers						
☐	2	Auto	All	All	443	0.0.0.0	capable
	Description: Default rule to auto-discover and attempt to optimize connections destined to port 443 as SSL						
	default	Auto	All	All	All	All	

2. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable Enhanced Auto-Discovery	Enables enhanced auto-discovery. With enhanced auto-discovery, the Steelhead appliance automatically finds the furthest Steelhead appliance along the connection path of the TCP connection and optimization occurs there. For example, in a deployment with four Steelhead appliances (A, B, C, D), where D represents the appliance that is furthest from A, the Steelhead appliance automatically finds D. This simplifies configuration and makes your deployment more scalable. By default, enhanced auto-discovery peering is enabled. If you do not enable enhanced auto-discovery, the Steelhead appliance uses regular auto-discovery. With regular auto-discovery, the Steelhead appliance finds the first remote Steelhead appliance along the connection path of the TCP connection and optimization occurs there. For example, if you had a deployment with four Steelhead appliances (A, B, C, D) where D represents the appliance that is furthest from A, the Steelhead appliance automatically finds B, then C, and finally D and optimization takes place in each. For detailed information about deployments that require enhanced auto-discovery peering, see the *Riverbed Deployment Guide*.
Enable Extended Peer Table	Enables support for up to 20,000 peers on high-end server-side Steelhead appliances (models 5520, 6020, 6050, and 6120) to accommodate large Steelhead client deployments. The datastore maintains the peers in groups of 1,024 in the global peer table. Riverbed recommends enabling the extended peer table if you have more than 4,000 peers. By default, this option is disabled and it is unavailable on Steelhead appliance models that do not support it. After enabling this option you must clear the datastore and stop and restart the service. **Important:** Before enabling this feature you must have a thorough understanding of performance and scaling issues. When deciding whether to use extended peer table support, you need to compare it with a serial cluster deployment. For details on serial clusters, see the *Riverbed Deployment Guide*. **Important:** After enabling extended peer table support, you cannot install a RiOS software version earlier than v5.5 without first clearing the datastore.

3. Click **Apply** to apply your settings. If you have enabled Extended Peer Table Support, a message tells you to clear the datastore and restart the service.

4. Click **Save** to save your settings permanently.

Peering Rules

Peering rules control Steelhead appliance behavior when it sees probe queries.

Peering rules are an ordered list of fields a Steelhead appliance uses to match with incoming SYN packet fields (for example, source or destination subnet, IP address, VLAN, or TCP port) as well as the IP address of the probing Steelhead appliance. This is especially useful in complex networks.

The Peering Rules List

The Peering Rules page displays a list of peering rules. The list contains the default peering rules and any peering rules you add.

The system evaluates the rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied. If the conditions set in the rule do not match, then the rule is not applied and the system moves on to the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted.

The Rule Type of a matching rule determines which action the Steelhead appliance takes on the connection.

Figure 4-3. The Default Peering Rules

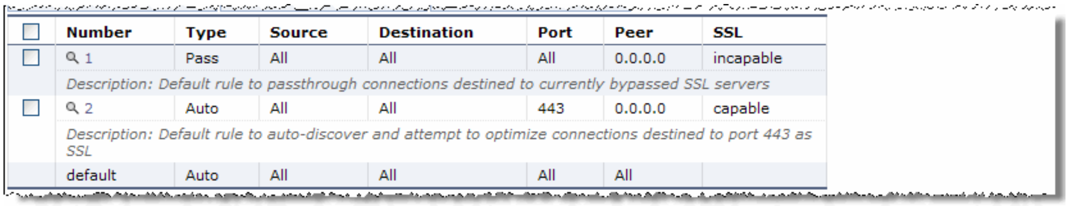

□	Number	Type	Source	Destination	Port	Peer	SSL
□	1	Pass	All	All	All	0.0.0.0	incapable
	Description: Default rule to passthrough connections destined to currently bypassed SSL servers						
□	2	Auto	All	All	443	0.0.0.0	capable
	Description: Default rule to auto-discover and attempt to optimize connections destined to port 443 as SSL						
	default	Auto	All	All	All	All	

About the Default Peering Rules

The default peering rules are adequate for typical network configurations, such as in-path configurations. However, you might need to add peering rules for complex network configurations. For details on deployment cases requiring peering rules, see the *Riverbed Deployment Guide*.

Note: Riverbed recommends using in-path rules to optimize SSL connections on destination ports other than the default port 443. For details, see "Configuring In-Path Rules" on page 28.

- The default peering rule number 1 with the SSL incapable flag matches any SSL connection whose IP address and destination port appear in the list of bypassed servers in the Configure > Optimization > SSL Main Settings page. The bypassed list includes the IP addresses and port numbers of SSL servers that the Steelhead appliance is bypassing because it could not match the common name of the server's certificate with one in its certificate pool. The list also includes servers whose IP address and port combination have experienced an SSL handshake failure. For example, a handshake failure occurs when the Steelhead appliance cannot find the issuer of a server certificate on its list of trusted certificate authorities.

  Note: Once a server appears in the bypassed servers list, follow-on connections to the same destination IP and port number always match rule number 1.

- The default peering rule number 2 with the SSL capable flag matches connections on port 443 that did not match default peering rule number 1. RiOS versions prior to v6.0 required a valid SSL server certificate to have been installed and the specific IP address and port associated with them to be configured. In RiOS v6.0 or later, the Steelhead appliance attempts to automatically discover certificate matches for servers answering on port 443. For all connections that match, the Steelhead appliance performs both enhanced auto-discovery (finding the nearest and farthest Steelhead appliance pair) and SSL optimization.

To configure a peering rule

1. To add, move, or remove a peering rule, complete the configuration as described in the following table.

Control	Description
Add a New Peering Rule	Displays the controls for adding a new peering rule.
Rule Type	Determines which action the Steelhead appliance takes on the connection. Select one of the following rule types from the drop-down list: • **Auto** - Allows built-in functionality to determine the response for peering requests (performs the best peering possible). If the receiving Steelhead appliance is not using automatic auto-discovery, this has the same effect as the **Accept** peering rule action. If automatic auto-discovery is enabled, the Steelhead appliance only becomes the optimization peer if it is the last Steelhead appliance in the path to the server. • **Accept** - Accepts peering requests that match the source-destination-port pattern. The receiving Steelhead appliance responds to the probing Steelhead appliance and becomes the remote-side Steelhead appliance (that is, the peer Steelhead appliance) for the optimized connection. • **Passthrough** - Allows pass-through peering requests that match the source and destination port pattern. The receiving Steelhead appliance does not respond to the probing Steelhead appliance, and allows the SYN+probe packet to continue through the network.
Insert Rule At	Determines the order in which the system evaluates the rule. Select Start, End, or a rule number from the drop-down list. The system evaluates rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied and the system moves on to the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted. The Rule Type of a matching rule determines which action the Steelhead appliance takes on the connection.
Source Subnet	Specify an IP address and mask for the traffic source, or you can specify all or 0.0.0.0/0 as the wildcard for all traffic. Use the following format: XXX.XXX.XXX.XXX/XX
Destination Subnet	Specify an IP address and mask pattern for the traffic destination, or you can specify all or 0.0.0.0/0 as the wildcard for all traffic. Use the following format: XXX.XXX.XXX.XXX/XX
	Port - Specify the destination port number, port label, or all.
Peer IP Address	Specify the in-path IP address of the probing Steelhead appliance. If more than one in-path interface is present on the probing Steelhead appliance, apply multiple peering rules, one for each in-path interface.

Control	Description
SSL Capability	Enables an SSL Capability flag, which specifies a criteria for matching an incoming connection with one of the rules in the peering rules table. This flag is typically set on a server-side Steelhead appliance. Select one of the following options from the drop-down list to determine how to process attempts to create secure SSL connections: • **No Check** - The peering rule does not determine whether the server Steelhead appliance is present for the particular destination IP address and port combination. • **Capable** - The peering rule determines that the connection is SSL-capable if the destination port is 443 (irrespective of the destination port value on the rule), and the destination IP and port do not appear on the bypassed servers list. The Steelhead appliance accepts the condition and, assuming all other proper configurations and that the peering rule is the best match for the incoming connection, optimizes SSL. • **Incapable** - The peering rule determines that the connection is SSL-incapable if the destination IP and port appear in the bypassed servers list. The service adds a server to the bypassed servers list when there is no SSL certificate for the server or for any other SSL handshake failure. The Steelhead appliance passes the connection through unoptimized without affecting connection counts. **Note:** Riverbed recommends that you use in-path rules to optimize SSL connections on non-443 destination port configurations.
Description	Specify a description to help you identify the peering relationship.
Add	Adds a peering rule to the list. The Management Console redisplays the Peering Rules table and applies your modifications to the running configuration, which is stored in memory.
Cancel	Cancels your actions.
Remove Selected Rules	Select the check box next to the name and click **Remove Selected Rules**.
Move Selected Rules	Select the check box next to the rule and click **Move Selected Rules**. Click the arrow next to the desired rule position; the rule moves to the new position.

2. Click **Save** to save your settings permanently.

Preventing an Unknown (or Unwanted) Steelhead Appliance from Peering

Enhanced auto-discovery greatly reduces the complexities and time it takes to deploy Steelhead appliances. It works so seamlessly that occasionally it has the undesirable effect of peering with Steelheads on the Internet that are not in your organization's management domain or your corporate business unit. When an unknown (or unwanted) Steelhead appears connected to your network, you can create a peering rule to prevent it from peering and remove it from your list of connected appliances. The peering rule defines what to do when a Steelhead appliance receives an auto-discovery probe from the unknown Steelhead appliance.

To prevent an unknown Steelhead from peering

1. Choose Configure > Optimization > Peering Rules.

2. Click **Add a New Peering Rule**.

3. Select Passthrough as the rule type.

4. Specify the source and destination subnets. The source subnet is the remote location network subnet (in the format XXX.XXX.XXX.XXX/XX). The destination subnet is your local network subnet (in the format XXX.XXX.XXX.XXX/XX).

5. Click **Add**.

   In this example, the peering rule passes through traffic from the unknown Steelhead in the remote location.

Note: When you use this method and add a new remote location in the future, you need to create a new peering rule that accepts traffic from the remote location. Place this new Accept rule before the Pass-through rule.

   If you do not know the network subnet for the remote location, there is another option: you can create a peering rule that allows peering from your corporate network subnet and denies it otherwise. For example, create a peering rule that accepts peering from your corporate network subnet and place it as the first rule in the list. Next, create a second peering rule to pass-through all other traffic. In this example, when the local Steelhead receives an auto-discovery probe, it checks the peering rules first (from top to bottom). If it matches the first Accept rule, the local Steelhead peers with the other Steelhead. If it does not match the first Accept rule, the local Steelhead checks the next peering rule, which is the pass-through rule for all other traffic. In this case, the local Steelhead appliance just passes through the traffic, and does not peer with the other Steelhead appliance.

After you add the peering rule, the unknown Steelhead appliance appears in the Current Connections report as a Connected Appliance until the connection times out. After the connection becomes inactive, it appears dimmed. To remove the unknown appliance completely, restart the optimization service.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Configuring General Service Settings" on page 60
- "Configuring Port Labels" on page 94
- "Secure Inner Channel Overview" on page 231
- "Viewing Current Connections" on page 385

Configuring NAT IP Address Mapping

CSH This feature is only supported by the Cloud Steelhead.

You configure NAT IP address mapping for the Cloud Steelhead in the Configure > Optimization > NAT IP Address Mapping page.

To configure NAT IP address mapping

1. Choose Configure > Optimization > NAT IP Address Mapping to display the NAT IP Address Mapping page.

Figure 4-4. NAT IP Address Mapping Page

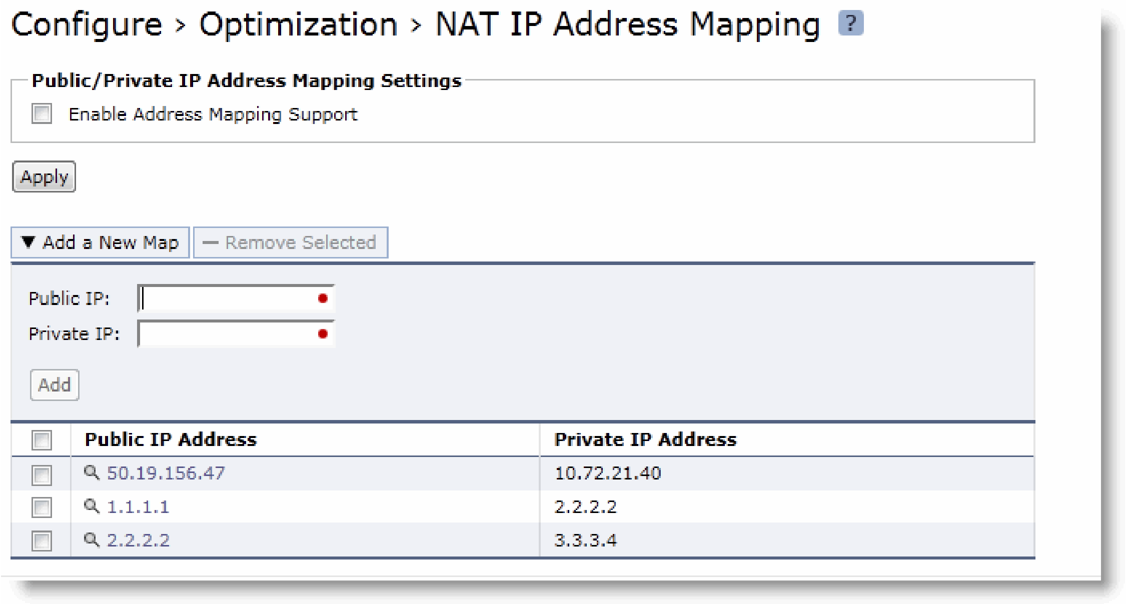

Configure > Optimization > NAT IP Address Mapping

Public/Private IP Address Mapping Settings

□ Enable Address Mapping Support

□	Public IP Address	Private IP Address
□	50.19.156.47	10.72.21.40
□	1.1.1.1	2.2.2.2
□	2.2.2.2	3.3.3.4

2. Under Public/Private IP Address Mapping Settings, select the Enable Address Mapping Support checkbox to enable the Cloud Steelhead to support public or private IP address mapping.

3. Click **Apply** to apply your settings to the running configuration.

4. Complete the configuration as described in the following table.

Control	Description
Add a New Map	Displays the controls to add a new IP address map.
Remove Selected	Select the checkbox next to the IP address and click Remove Selected to delete it from the system.
Public IP	Type the current public IP address of the appliance.
Private IP	Type the private IP address (cloud vendor-assigned) of the appliance.
Add	Adds the public IP address and private IP address of the appliance to the system.

Configuring Discovery Service

CSH This feature is only supported by the Cloud Steelhead.

You configure discovery service in the Cloud Steelhead in the Configure > Optimization > Discovery Service page. Discovery service enables the Cloud Steelhead to find and propagate the public and private IP address of the Cloud Steelhead.

To configure discovery service

1. Choose Configure > Optimization > Discovery Service to display the Discovery Service page.

Figure 4-5. Discovery Service Page

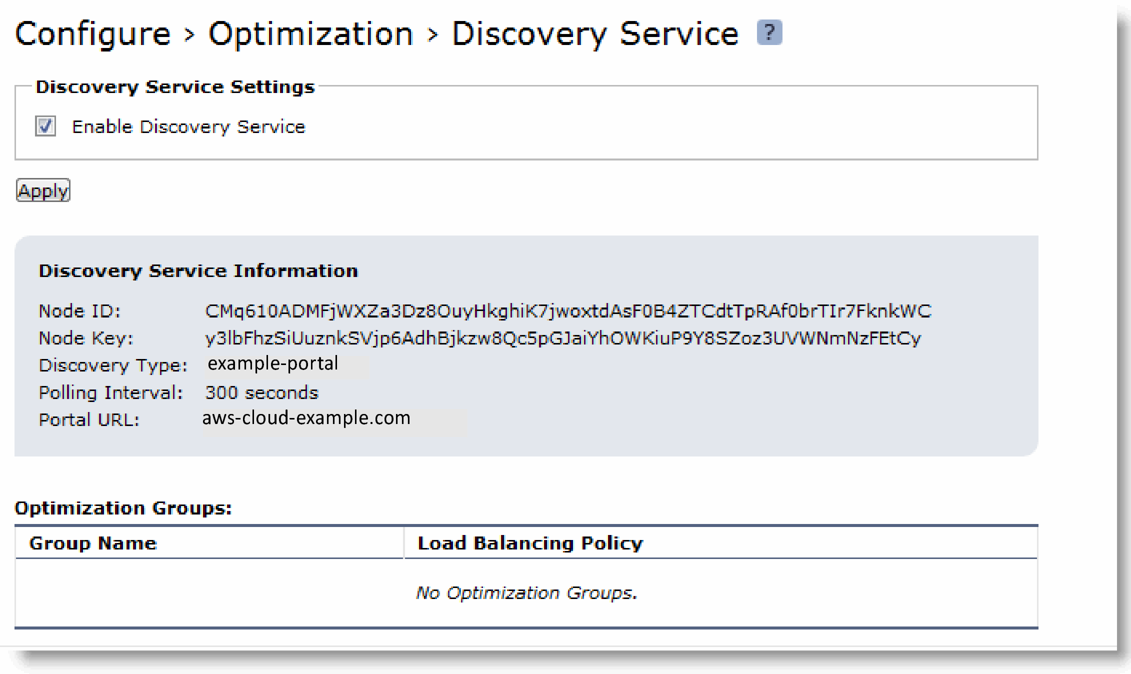

Configure > Optimization > Discovery Service

Discovery Service Settings

☑ Enable Discovery Service

Discovery Service Information

Node ID:	CMq610ADMFjWXZa3Dz8OuyHkghiK7jwoxtdAsF0B4ZTCdtTpRAf0brTIr7FknkWC
Node Key:	y3lbFhzSiUuznkSVjp6AdhBjkzw8Qc5pGJaiYhOWKiuP9Y8SZoz3UVWNmNzFEtCy
Discovery Type:	example-portal
Polling Interval:	300 seconds
Portal URL:	aws-cloud-example.com

Optimization Groups:

Group Name	Load Balancing Policy
	No Optimization Groups.

2. Under Discovery Service Settings, check the Enable Discovery Service check box to enable discovery service. This option is selected by default.

The Cloud Steelhead displays the following discovery service information: node ID, node key, discovery type, polling interval, and portal URL.

The Optimization Groups table displays the group name and the load balancing policy of the optimization groups that you configured in the Riverbed Cloud Portal. Click the magnifying glass icon before the group name to display more information about the list of nodes in each group. Click the magnifying glass icon before the node to display more information about the node such as the load balancing policy, node ID, public interfaces, and local interfaces.

Configuring the Datastore

This section describes how to configure datastore settings. It includes the following sections:

- "Encrypting the Datastore" on page 77
- "Synchronizing Peer Datastores" on page 78
- "Clearing the Datastore" on page 80
- "Warming Branch Steelhead Mobile Clients" on page 81

You can display and modify datastore settings in the Configure > Optimization > Data Store page. This page is typically used to turn on datastore encryption and synchronization.

Steelhead appliances transparently intercept and analyze all of your WAN traffic. TCP traffic is segmented, indexed, and stored as *segments* of data, and the *references* representing that data is stored on the datastore within Steelhead appliances on both sides of your WAN. After the data has been indexed, it is compared to data already on the disk. Segments of data that have been seen before are not transferred across the WAN again; instead a reference is sent in its place that can index arbitrarily large amounts of data, thereby massively reducing the amount of data that needs to be transmitted. One small reference can refer to megabytes of existing data that has been transferred over the WAN before.

Encrypting the Datastore

You enable datastore encryption in the Configure > Optimization > Data Store page.

Encrypting the datastore significantly limits the exposure of sensitive data in the event an appliance is compromised by loss, theft, or a security violation. The secure data is difficult for a third party to retrieve.

Before you encrypt the datastore, the secure vault must be unlocked. The encryption key is stored in the secure vault. For details, see "Unlocking the Secure Vault" on page 370.

Important: Encrypting the datastore *and* enabling SSL optimization provides maximum security. For details, see "Configuring SSL Server Certificates and Certificate Authorities" on page 213.

Note: Datastore synchronization traffic is not encrypted.

Encryption Strengths

Encrypting the datastore can have performance implications; generally, higher security means less performance. Several encryption strengths are available to provide the right amount of security while maintaining the desired performance level. When selecting an encryption type, you must evaluate the network structure, the type of data that travels over it, and how much of a performance trade-off is worth the extra security.

Encrypted Datastore Downgrade Limitations

The Steelhead appliance cannot use an encrypted datastore with an earlier RiOS software version, unless the release is an update (v4.x.x). For example, an encrypted datastore created in v4.1.4 would work with v4.1.2, but not with v4.0.x.

Before downgrading to an earlier software version, you must select none as the encryption type, clear the datastore, and restart the service. After you clear the datastore, the data is removed from persistent storage and cannot be recovered.

If you return to a previous software version and there is a mismatch with the encrypted datastore, the status bar indicates that the datastore is corrupt. You can either:

- Use the backup software version after clearing the datastore and rebooting the service.

—or—

- Return to the software version in use when the datastore was encrypted, and continue using it.

To encrypt the datastore

1. Choose Configure > Optimization > Data Store to display the Data Store page.

Figure 4-6. Data Store Page

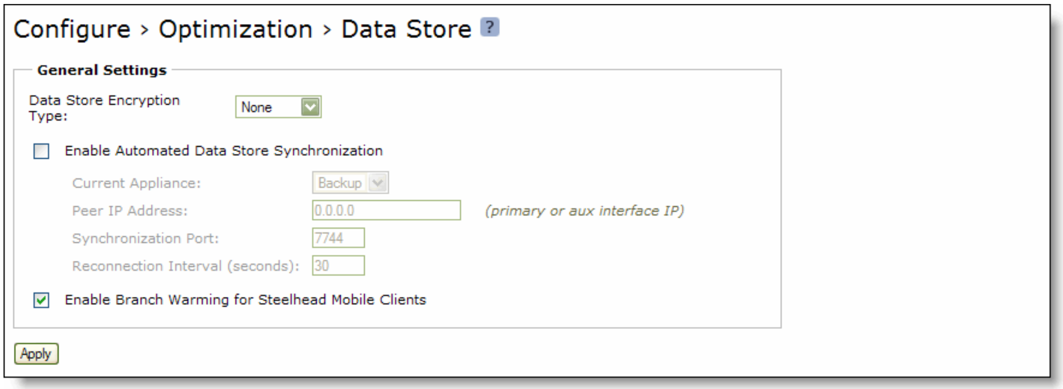

Configure > Optimization > Data Store

General Settings

Data Store Encryption Type: None

□ Enable Automated Data Store Synchronization

Current Appliance: Backup

Peer IP Address: 0.0.0.0 (primary or aux interface IP)

Synchronization Port: 7744

Reconnection Interval (seconds): 30

☑ Enable Branch Warming for Steelhead Mobile Clients

2. Under General Settings, complete the configuration as described in the following table.

Control	Description
Data Store Encryption Type	Select one of the following encryption types from the drop-down list. The encryption types are listed from the least to the most secure. • **None** - Turns off data encryption. • **AES_128** - Encrypts data using the AES cryptographic key length of 128 bits. • **AES_192** - Encrypts data using the AES cryptographic key length of 192 bits. • **AES_256** - Encrypts data using the AES cryptographic key length of 256 bits.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

5. Select **Clear the Data Store on Reboot** and reboot the Steelhead appliance as described in "Rebooting and Shutting Down the Steelhead Appliance" on page 349.

Important: You must clear the datastore and reboot the Steelhead service on the Steelhead appliance after turning on, changing, or turning off the encryption type. After you clear the datastore, the data cannot be recovered. If you do not want to clear the datastore, reselect your previous encryption type and reboot the service. The Steelhead appliance uses the previous encryption type and encrypted datastore. For details, see "Rebooting and Shutting Down the Steelhead Appliance" on page 349.

Synchronizing Peer Datastores

For deployments requiring the highest levels of redundancy and performance, RiOS supports *warm* standby between designated master and backup devices. Using automated datastore synchronization, the data segments and the references created via data streamlining are automatically copied from the master to the backup appliance. In the event of a failure in the master appliance, the backup appliance takes its place with a warm datastore, and can begin delivering fully-optimized performance immediately. Warm data transfers send only new or modified data, dramatically increasing the rate of data transfer over the WAN.

RiOS supports active-active configurations, in which each appliance is serving both as a master for some traffic and as a backup for the other appliance, with full datastore synchronization. Automatic synchronization can include appliances in a serial or WCCP cluster, and appliances using connection forwarding.

Note: Synchronization takes place over the primary or auxiliary port only.

Failover is not required for datastore synchronization. Although the failover and synchronization features are typically enabled together, you can enable datastore synchronization independently of standard failover.

Note: In most implementations in which both failover and synchronization are enabled, the same Steelhead appliance serves as the master for both failover and datastore synchronization. However, if you enable failover and synchronization, the failover master and the synchronization master do not have to be the same Steelhead appliance.

You configure two Steelhead appliances to enable synchronization, one as a server (the synchronization master) and the other as a backup. The synchronization master and its backup:

- must be on the same LAN.
- do not have to be in the same physical location. If they are in different physical locations, they must be connected via a fast, reliable LAN connection with minimal latency.
- must be running the same version of the RiOS software.
- must have the same hardware model.
- must be configured on the primary or auxiliary interface.

When you have configured the master and backup appliances, you must restart the Steelhead service on the backup Steelhead appliance. The master restarts automatically.

After you have enabled and configured synchronization, the datastores are actively kept synchronized. For details on how synchronized appliances replicate data and how datastore synchronization is commonly used in high availability designs, see the *Riverbed Deployment Guide*.

Note: If one of the synchronized Steelhead appliances is under high load, some data might not be copied. For details, see the *Riverbed Deployment Guide*.

Note: If datastore synchronization is interrupted for any reason (such as a network interruption or if one of the Steelhead appliances is taken out of service), the Steelhead appliances continue other operations without disruption. When the interruption is resolved, datastore synchronization resumes without risk of data corruption.

To synchronize the datastore

1. Choose one Steelhead appliance to be the master and one to be the backup. The backup has its datastore overwritten by the master datastore.

2. Make sure there is a network connection between the two Steelhead appliances.

3. Connect to the Management Console on the Steelhead appliance you have chosen to be the master appliance.

4. Choose Configure > Optimization > Data Store to display the Data Store page.

5. Under General Settings, complete the configuration as described in the following table.

Control	Description
Enable Automated Data Store Synchronization	Enables automated data store synchronization. Data store synchronization ensures that each data store in your network has *warm* data for maximum optimization. All operations occur in the background and do not disrupt operations on any of the systems.
Current Appliance	Select Master or Backup from the drop-down list.
Peer IP Address	Specify the IP address for the peer appliance. You must specify either the IP address for the primary or auxiliary interface (if you use the auxiliary interface in place of the primary).
Synchronization Port	Specify the destination TCP port number used when establishing a connection to synchronize data. The default value is 7744.
Reconnection Interval	Specify the number of seconds to wait for reconnection attempts. The default value is 30.

6. Click **Apply** to apply your settings.

7. Click **Save** to save your settings permanently.

8. Choose Configure > Maintenance > Services to display the Services page.

9. Select Clear the Data Store and click **Restart** to restart the service on the Steelhead appliance.

Note: When redeploying a synchronized pair, you must clear the datastore. For details, see "Clearing the Datastore" on page 80.

Clearing the Datastore

The appliance continues to write data references to the datastore until it reaches capacity. In certain situations, you might need to clear the datastore. For example, you need to clear the datastore:

- after turning encryption on or off, or changing the encryption type.
- before downgrading to an earlier software version.
- to redeploy an active-active synchronization pair.
- after testing or evaluating the appliance.
- after receiving a "data store is corrupt" message.

For details on clearing the datastore, see “Rebooting and Shutting Down the Steelhead Appliance” on page 349.

Note: After clearing the datastore and rebooting the service, the data transfers are cold. Performance improves with subsequent warm data transfers over the WAN.

Warming Branch Steelhead Mobile Clients

You enable branch warming for Steelhead Mobile Clients in the Configure > Optimization > Data Store page. By default, branch warming is enabled.

Branch warming keeps track of data segments created while a Steelhead Mobile user is in a Steelhead appliance-enabled branch office and sends the new data back to the Steelhead Mobile user’s laptop. When the user goes back on the road, they receive warm performance.

Branch warming co-operates with and optimizes transfers for a server-side Steelhead appliance. New data transfers between the client and server are populated in the Steelhead Mobile datastore, the branch Steelhead appliance datastore, and the server-side Steelhead appliance datastore.

When the server downloads data, the server-side Steelhead appliance checks if either the Steelhead Mobile Client or the branch Steelhead appliance has the data in their datastore. If either device already has the data segments, the server-side Steelhead appliance sends only references to the data. The Mobile Client and the branch Steelhead appliance communicate with each other to resolve the references.

Other clients at a branch office benefit from branch warming as well, because data transferred by one client at a branch also populates the branch Steelhead appliance datastore. Performance improves with all clients at the branch because they receive warm performance for that data.

Note: For details, see the *Riverbed Deployment Guide*.

Requirements

The following requirements must be met for branch warming to work:

- Enable latency-based location awareness and branch warming on the Steelhead Mobile Controller.
- Enable branch warming on both the client-side and server-side Steelhead appliances.
- Both the client-side and server-side Steelhead appliances must be deployed in-path.
- Enable enhanced auto-discovery on both the client-side and server-side Steelhead appliances.
- The Steelhead Mobile Controller appliance must be running RiOS v3.0 or later.
- The Steelhead appliances must be running RiOS v6.0 or later.
- The Steelhead Mobile Client must be running RiOS v3.0 or later.

Branch Warming does not improve performance for configurations using:

- SSL connections
- Out-of-path with fixed-target rules

- Steelhead Mobile Clients that communicate with multiple server-side appliances in different scenarios. For example, if a Steelhead Mobile Client home user peers with one server-side Steelhead appliance after logging in through a VPN network and peers with a different server-side Steelhead appliance after logging in from the branch office, branch warming does not improve performance.

To enable branch warming

1. On both the client-side and the server-side Steelhead appliances, choose Configure > Optimization > Data Store to display the Data Store page.

Figure 4-7. Data Store Page

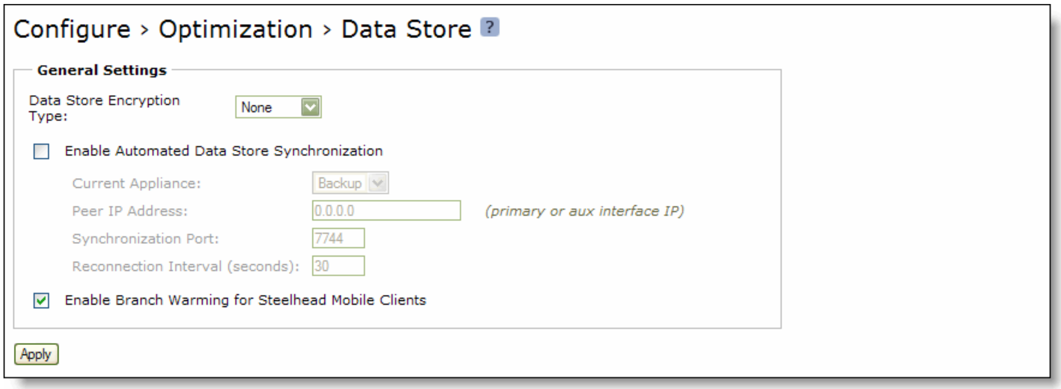

Configure > Optimization > Data Store

General Settings

Data Store Encryption Type: None

□ Enable Automated Data Store Synchronization

Current Appliance: Backup

Peer IP Address: 0.0.0.0 *(primary or aux interface IP)*

Synchronization Port: 7744

Reconnection Interval (seconds): 30

☑ Enable Branch Warming for Steelhead Mobile Clients

2. Under General Settings, select Enable Branch Warming for Steelhead Mobile Clients.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

5. You must restart the optimization service. For details, see “Starting and Stopping the Optimization Service” on page 345.

Related Topics

- “Enabling Failover” on page 60
- “Improving Performance” on page 82
- “Unlocking the Secure Vault” on page 370
- “Viewing Data Store Status Reports” on page 436

Improving Performance

You can enable settings to improve network and datastore performance in the Configure > Optimization > Performance page. The following sections describe the default settings and the cases in which you might consider changing the default values.

Selecting a Datastore Segment Replacement Policy

The datastore segment replacement policy selects the technique used to replace the data in the datastore. While the default setting works best for most Steelhead appliances, occasionally Riverbed Support recommends changing the policy to improve performance.

Note: The segment replacement policy should match on both the client-side and server-side Steelhead appliances.

To select a datastore segment replacement policy

1. Choose Configure > Optimization > Performance to display the Performance page.

2. Under Data Store, select one of the following replacement algorithms from the drop-down list.

Control	Description
Segment Replacement Policy	• **Riverbed LRU** - Replaces the least recently used data in the datastore, which improves hit rates when the data in the datastore are not equally used. This is the default setting. • **FIFO** - Replaces data in the order received (first in, first out).

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

5. Restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Note: Upgrading from RiOS v5.0.x to v5.5 changes the default datastore segment replacement policy from FIFO to Riverbed LRU.

Optimizing the Datastore for High-Throughput Environments

You can optimize the datastore for high-throughput Data Replication (DR) or data center workloads in the Configure > Optimization > Performance page.

You might benefit from changing the performance settings if your environment uses a high-bandwidth WAN. DR and SAN replication workloads at these high throughputs might benefit from the settings that enhance datastore performance while still receiving data reduction benefits from SDR.

To maintain consistent levels of performance, Riverbed recommends using separate Steelhead appliances for DR workloads than for optimization of other application traffic.

Setting an Adaptive Streamlining Mode

The adaptive data streamlining mode monitors and controls the different resources available on the Steelhead appliance and adapts the utilization of these system resources to optimize LAN throughput. Changing the default setting is *optional*; Riverbed recommends you select another setting only with guidance from Riverbed Support or the Riverbed Sales Team.

Generally, the default setting provides the most data reduction. When choosing an adaptive streamlining mode for your network, contact Riverbed Support to help you evaluate the setting based on:

- the amount of data replication your Steelhead appliance is processing.
- the type of data being processed and its effects on disk throughput on the Steelhead appliances.
- your primary goal for the project, which could be maximum data reduction or maximum throughput. Even when your primary goal is maximum throughput you can still achieve high data reduction.

To select an adaptive data streamlining mode

1. Choose Configure > Optimization > Performance to display the Performance page.

2. Under Adaptive Data Streamlining Modes, select one of the following settings.

Setting	Description
Default	This setting is enabled by default and works for most implementations. The default setting: • Provides the most data reduction. • Reduces random disk seeks and improves disk throughput by discarding very small data margin segments that are no longer necessary. This Margin Segment Elimination (MSE) process provides network-based disk defragmentation. • Writes large page clusters. • Monitors the disk write I/O response time to provide more throughput.

Setting	Description
SDR-Adaptive	**Legacy** - Includes the default settings and also: • Balances writes and reads. • Monitors both read and write disk I/O response and, based on statistical trends, can employ a blend of disk-based and non-disk-based data reduction techniques to enable sustained throughput during periods of high disk-intensive workloads. **Important:** Use caution with the SDR-Adaptive Legacy setting, particularly when you are optimizing CIFS or NFS with prepopulation. Please contact Riverbed Support for more information. **Advanced** - Maximizes LAN-side throughput dynamically under different data work loads. This switching mechanism is governed with a throughput and bandwidth reduction goal using the available WAN bandwidth. Both Steelheads must be running RiOS v6.0.x or later. **Upgrade notes**: If you have enabled SDR-Adaptive prior to upgrading to RiOS v6.0, the default setting is SDR-Adaptive Legacy. If you did not change the SDR-Adaptive setting prior to upgrading to RiOS v6.0, the default setting is SDR-Adaptive Advanced.
SDR-M	Performs data reduction entirely in memory, which prevents the Steelhead appliance from reading and writing to and from the disk. Enabling this option can yield high LAN-side throughput because it eliminates all disk latency. This is typically the preferred configuration mode for SAN replication environments. SDR-M is most efficient when used between two identical high-end Steelhead appliance models; for example, 6050 - 6050. When used between two different Steelhead appliance models, the smaller model limits the performance. After enabling SDR-M on both the client-side and the server-side Steelhead appliances, restart both Steelheads to avoid performance degradation. **Important:** You cannot use peer data store synchronization with SDR-M.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

5. If you have selected a new adaptive data streamlining mode, you must restart the Steelhead service on the client-side and server-side Steelhead appliances. For details, see "Starting and Stopping the Optimization Service" on page 345.

Note: If you select SDR-M as the adaptive data streamlining mode, the Clear the Datastore option is not available when you restart the Steelhead service because the SDR-M mode has no effect on the datastore disk.

Tip: After changing the datastore adaptive streamlining setting, you can verify whether changes have had the desired effect by reviewing the Optimized Throughput report. From the menu bar, choose Reports > Optimization > Optimized Throughput.

Configuring CPU Settings

Use the CPU settings to balance throughput with the amount of data reduction and balance the connection load. The CPU settings are useful with high-traffic loads to scale back compression, increase throughput, and maximize Long Fat Network (LFN) utilization.

To configure the CPU settings

1. Choose Configure > Optimization > Performance to display the Performance page.

2. Under CPU Settings, complete the configuration as described in the following table.

Setting	Description
Compression Level	Specifies the relative trade-off of data compression for LAN throughput speed. Generally, a lower number provides faster throughput and slightly less data reduction. Select a datastore compression value of 1 (minimum compression, uses less CPU) through 9 (maximum compression, uses more CPU) from the drop-down list. The default value corresponds to level 6. Riverbed recommends setting the compression level to 1 in high-throughput environments such as data center to data center replication.
Adaptive Compression	Detects LZ data compression performance for a connection dynamically and turns it off (sets the compression level to 0) momentarily if it is not achieving optimal results. Improves end-to-end throughput over the LAN by maximizing the WAN throughput. By default, this setting is disabled.
Multi-Core Balancing	Enables multi-core balancing which ensures better distribution of workload across all CPUs, thereby maximizing throughput by keeping all CPUs busy. Core balancing is useful when handling a small number of high-throughput connections (approximately 25 or less). By default, this setting is disabled.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Synchronizing Peer Datastores" on page 78

Configuring CIFS Prepopulation

You can enable prepopulation and add, modify, and delete prepopulation shares in the Configure > Optimization > CIFS Prepopulation page.

The prepopulation operation effectively performs the first Steelhead appliance read of the data on the prepopulation share. Subsequently, the Steelhead appliance handles read and write requests as effectively as with a warm data transfer. With warm transfers, only new or modified data is sent, dramatically increasing the rate of data transfer over the WAN.

Note: Riverbed does not support prepopulation with Windows Domain Controller servers with SMB signing set to Required. If your network environment requires SMB signing, use the RCU to prepopulate your shares. You can obtain the RCU from the Riverbed Support site at https://support.riverbed.com.

CSH The AWS Cloud Steelhead does not support CIFS Prepopulation. The ESX Cloud Steelhead supports CIFS Prepopulation if it is deployed with WCCP or PBR (not with the Discovery Agent).

To enable CIFS prepopulation and add, modify, or delete a prepopulation share

1. Choose Configure > Optimization > CIFS Prepopulation to display the CIFS Prepopulation page.

Figure 4-8. CIFS Prepopulation Page

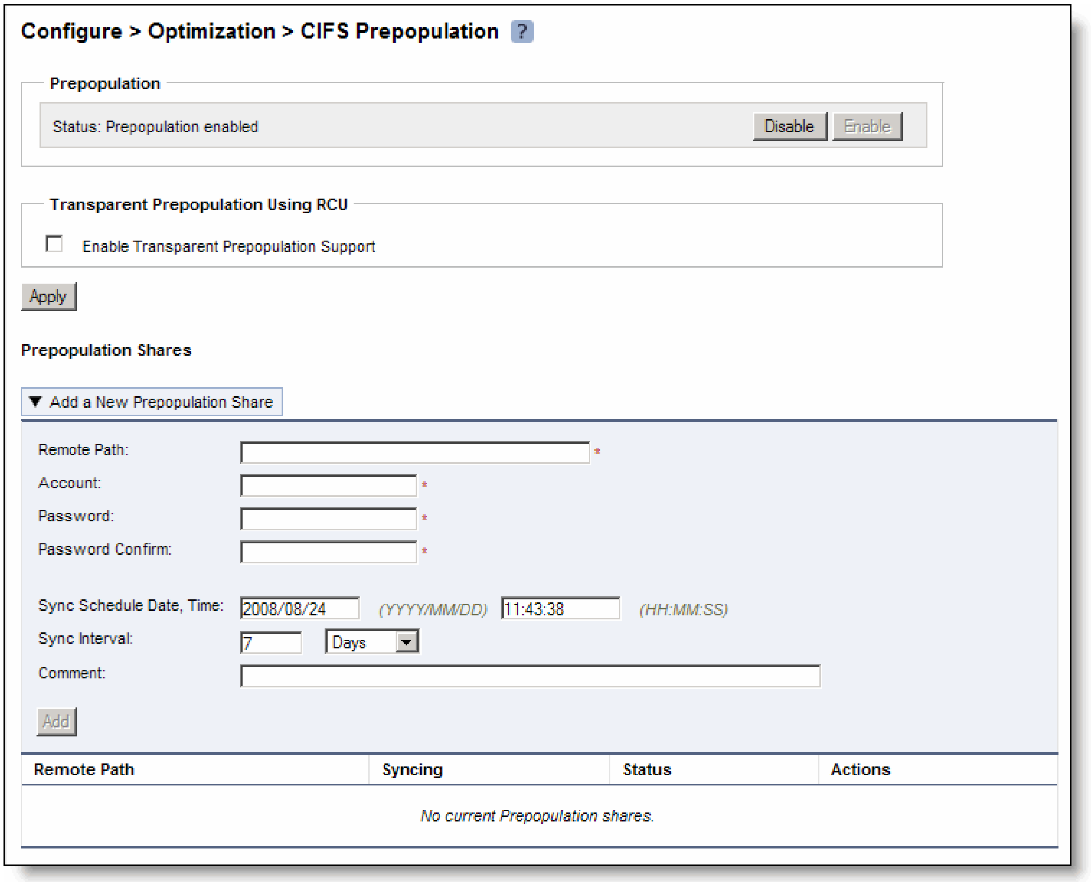

Configure > Optimization > CIFS Prepopulation

Prepopulation

Status: Prepopulation enabled

Transparent Prepopulation Using RCU

□ Enable Transparent Prepopulation Support

Prepopulation Shares

▼ Add a New Prepopulation Share

Remote Path:	*
Account:	*
Password:	*
Password Confirm:	*
Sync Schedule Date, Time:	2008/08/24 (YYYY/MM/DD) 11:43:38 (HH:MM:SS)
Sync Interval:	7 Days
Comment:	

Remote Path	Syncing	Status	Actions
No current Prepopulation shares.			

2. Under Prepopulation, complete the configuration as described in the following table.

Control	Description
Enable/Disable buttons	Enables or disables CIFS prepopulation, which pre-warms the Steelhead data store. In this setup, the Primary interface of the Steelhead acts as a client and pre-requests data from the share you want to use to warm the data store. When data is requested again by a client on the local LAN, only new or modified data is sent over the WAN which dramatically increases the rate of data transfers. When CIFS prepopulation is enabled, the data request is generated from the primary interface (acting as a client). This request goes through the LAN interface to WAN interface out to the server-side Steelhead, causing the in-path interface to see the data as a normal client request.

3. Under Transparent Prepopulation Using RCU, complete the configuration as described in the following table.

Control	Description
Enable Transparent Prepopulation Support	Enables transparent prepopulation using the RCU to prepopulate your shares. (Use this option if your network environment requires SMB signing.)

4. Click **Apply** to apply your settings.

5. When prepopulation is enabled, you can add and remove shares (or modify properties of shares) as described in the following table.

Control	Description
Add a New Prepopulation Share	Displays the controls for adding a new prepopulation share.
Remote Path	Specify the path to the data on the origin server or the UNC path of a share to which you want to make available for prepopulation. Set up the prepopulation share on the remote box pointing to the actual share in the headend data center server. For example: \\<origin-file-server>\<local-name> **Important:** The share and the origin-server share names must not use any characters other than letters, numbers, underscore, space, or backslash (directory separator). The names cannot contain any of the following characters: < > * ? | / + = ; : " , &
Account	Specify the account used to access the prepopulation shares. For example: <Domain>\<username>
Password/Password Confirm	Specify and confirm the password for the local administrator account.
Comment	Optionally, include a comment to help you administer the share in the future. Comments must not use any characters other than letters, numbers, underscore, space, or backslash (directory separator).

Control	Description
Sync Schedule, Date and Time	Specify a date and time to perform the initial synchronization job. The first synchronization, or the initial copy, retrieves data from the origin file server and copies it to the local disk on the Steelhead appliance. Subsequent synchronizations are based on the synchronization interval. Date Format: YYYY/MM/DD Time Format: HH:MM:SS
Sync Interval	Specify the interval for subsequent synchronization jobs. After the initial synchronization, the Steelhead appliance retrieves data from the server at every synchronization interval. In these subsequent synchronizations, only new data that was modified or created after the previous synchronization is sent from the origin-file server to Steelhead appliance. In the text box, specify a number. Select a time unit from the drop-down list: Minutes, Hours, Days, or Disabled.
Add	Adds the share to the Prepopulations Share list.

Tip: To modify share properties or delete a share, select the remote path for the share in the share table row.

6. Click **Save** to save your settings permanently.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Viewing CIFS Prepopulation Share Logs

After adding a CIFS prepopulation share, you can view CIFS prepopulation share logs from the Edit Prepopulation Share page. View the prepopulation share log to see more detail regarding the initial copy of the share or the last share synchronization.

To view CIFS prepopulation share logs

1. Choose Configure > Optimization > CIFS Prepopulation to display the CIFS Prepopulation page.

2. Select the remote path for the share.

3. Use the controls to customize the report as described in the following table.

Field	Description
Initial-copy log	Displays the date and time the initial share copy started and completed. The log includes how many directories, files, and bytes were received and how long it took to receive them. The log also lists any errors or deletions.
Last-sync log	Displays the date and time the last share synchronization started and completed. The log includes how many directories, files, and bytes were received and how long it took to receive them. The log also lists any errors or deletions.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Related Topics

- "Configuring CIFS Optimization" on page 97
- "Viewing CIFS Prepopulation Share Log Reports" on page 427

Configuring TCP and High-Speed TCP

This section describes how to configure TCP settings. It includes the following section:

- "TCP and High-Speed TCP Optimization" on page 91

You configure TCP and high-speed TCP in the Configure > Optimization > Transport Settings page.

TCP and High-Speed TCP Optimization

The high-speed TCP feature provides acceleration and high throughput for high-bandwidth links (also known as Long Fat Networks, or LFNs) where the WAN pipe is large but latency is high. High-speed TCP is activated for all connections that have a BDP larger than 100 packets.

Note: For details on using HS-TCP in data protection scenarios, see the *Riverbed Deployment Guide*.

HS-TCP Basic Steps

The following table describes the basic steps needed to configure high-speed TCP.

Task	Reference
1. Enable high-speed TCP support.	"To enable TCP and high-speed TCP optimization" on page 91.
2. Increase the WAN buffers to 2 * Bandwidth Delay Product (BDP). You can calculate the BDP WAN buffer size: Buffer size in bytes = 2 * bandwidth (in bits per sec) * delay (in sec) / 8 (bits per byte) Example: For a link of 155 Mbps and 100 ms round-trip delay. Bandwidth = 155 Mbps = 155000000 bps Delay = 100 ms = 0.1 sec BDP = 155 000 000 * 0.1 / 8 = 1937500 bytes Buffer size in bytes = 2 * BDP = 2 * 1937500 = 3 875 000 bytes. If this number is greater than the default (256 KB), enable HS-TCP with the correct buffer size.	"To configure buffer settings" on page 92.
3. Increase the LAN buffers to 1 MB.	"To configure buffer settings" on page 92.
4. Enable in-path support.	"Configuring General Service Settings" on page 60.

To enable TCP and high-speed TCP optimization

1. Choose Configure > Optimization > Transport Settings to display the Transport Settings page.

2. Under Transport Optimization, complete the configuration as described in the following table.

Control	Description
Enable HighSpeed TCP	Enables high-speed TCP for more complete use of long fat pipes (high-bandwidth, high-delay networks). Riverbed recommends that you enable HS-TCP only after you have carefully evaluated whether it will benefit your network environment. For details about the trade-offs of enabling HS-TCP, see **tcp highspeed enable** in the *Riverbed Command-Line Interface Reference Manual*.
Use Default Steelhead TCP Optimization	Optimizes TCP connections by applying data and transport streamlining for TCP traffic over the WAN. For details on data and transport streamling, see the *Riverbed Deployment Guide*. This is the default setting.

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

If you change the default TCP optimization setting, you must restart the Steelhead service on the server-side Steelhead appliance. For details, see “Starting and Stopping the Optimization Service” on page 285.

Configuring Buffer Settings

The buffer settings in the Transport Settings page support HS-TCP and are also used in data protection scenarios to improve performance. For details on data protection deployments, see the *Riverbed Deployment Guide*.

To configure buffer settings

1. Choose Configure > Optimization > Transport Settings to display the Transport Settings page.

2. Under Buffer Settings, complete the configuration as described in the following table.

Control	Description
LAN Send Buffer Size	Specify the send buffer size used to send data out of the LAN. The default value is 81920.
LAN Receive Buffer Size	Specify the receive buffer size used to receive data from the LAN. The default value is 32768.
WAN Default Send Buffer Size	Specify the send buffer size used to send data out of the WAN. The default value is 262140.
WAN Default Receive Buffer Size	Specify the receive buffer size used to receive data from the WAN. The default value is 262140.

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

Configuring Service Ports

You can configure service port settings in the Configure > Optimization > Service Ports page.

Service ports are the ports used for inner connections between Steelhead appliances.

You can configure multiple service ports on the server-side of the network for multiple QoS mappings. You define a new service port and then map destination ports to that port, so that QoS configuration settings on the router are applied to that service port.

Configuring service port settings is *optional*.

To set a service port

1. Choose Configure > Optimization > Service Ports to display the Service Ports page.

Figure 4-9. Service Ports Page

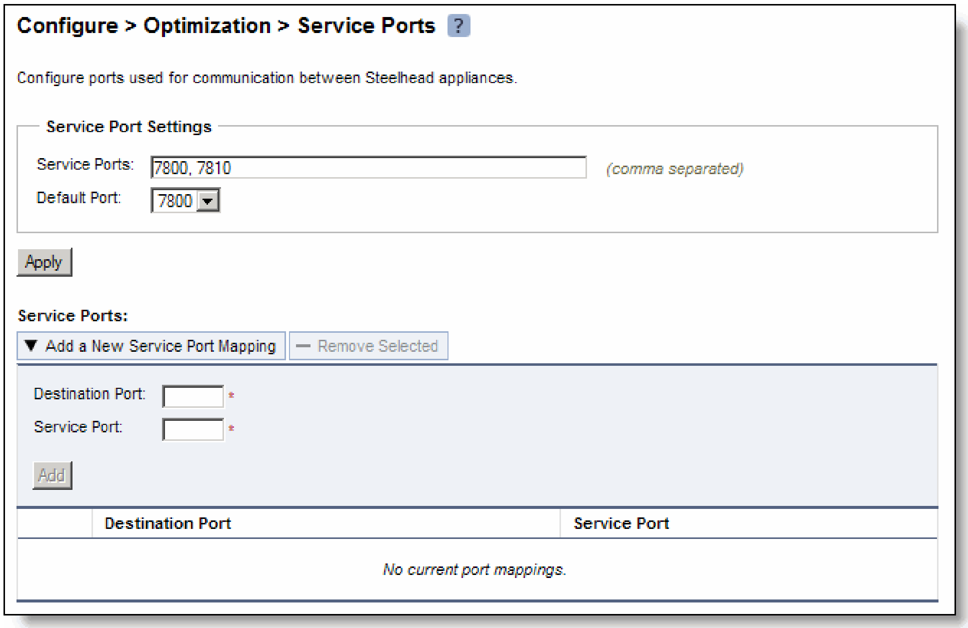

Configure > Optimization > Service Ports

Configure ports used for communication between Steelhead appliances.

Service Port Settings

Service Ports: 7800, 7810 (comma separated)

Default Port: 7800

	Destination Port	Service Port
	No current port mappings.	

2. Under Service Port Settings, complete the configuration as described in the following table.

Control	Description
Service Ports	Specify ports in a comma-separated list. The default service ports are 7800 and 7810.
Default Port	Select the default service port from the drop-down list. The default service ports are 7800 and 7810.

3. Click **Apply** to apply your settings.

To add a service port

1. Under Service Ports, complete the configuration as described in the following table.

Control	Description
Add a New Service Port Mapping	Displays the controls to add a new mapping.
Destination Port	Specify a destination port number.
Service Port	Specify a port number.
Add	Adds the port numbers.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Save** to save your settings permanently.

Related Topic

- "Configuring General Service Settings" on page 60

Configuring Port Labels

You create port labels in the Port Labels page. Port labels are names given to sets of port numbers. You use port labels when configuring in-path rules. For example, you can use port labels to define a set of ports for which the same in-path, peering, QoS classification, and QoS marking rules apply.

The following table summarizes the port labels that are provided by default.

Port Type	Description and Ports
Interactive	Use this port label to automatically pass-through traffic on interactive ports (for example, Telnet, TCP ECHO, remote logging, and shell).
RBT-Proto	Use this port label to automatically pass-through traffic on ports used by the system: 7744 (datastore synchronization), 7800-7801 (in-path), 7810 (out-of-path), 7820 (failover), 7850 (connection forwarding), 7860 (Interceptor appliance), 7870 (Steelhead Mobile Controller).
Secure	Use this port label to automatically pass-through traffic on commonly secure ports (for example, ssh, https, and smtps).

If you do not want to automatically forward traffic on interactive or secure ports, you must delete the Interactive and Secure in-path rules. For details, see "About In-Path Rules" on page 27.

For information on common port assignments, see "Steelhead Appliance Ports" on page 491.

This feature is *optional*.

To create a port label

1. Choose Configure > Networking > Port Labels to display the Port Labels page.

Figure 4-10. Port Labels Page

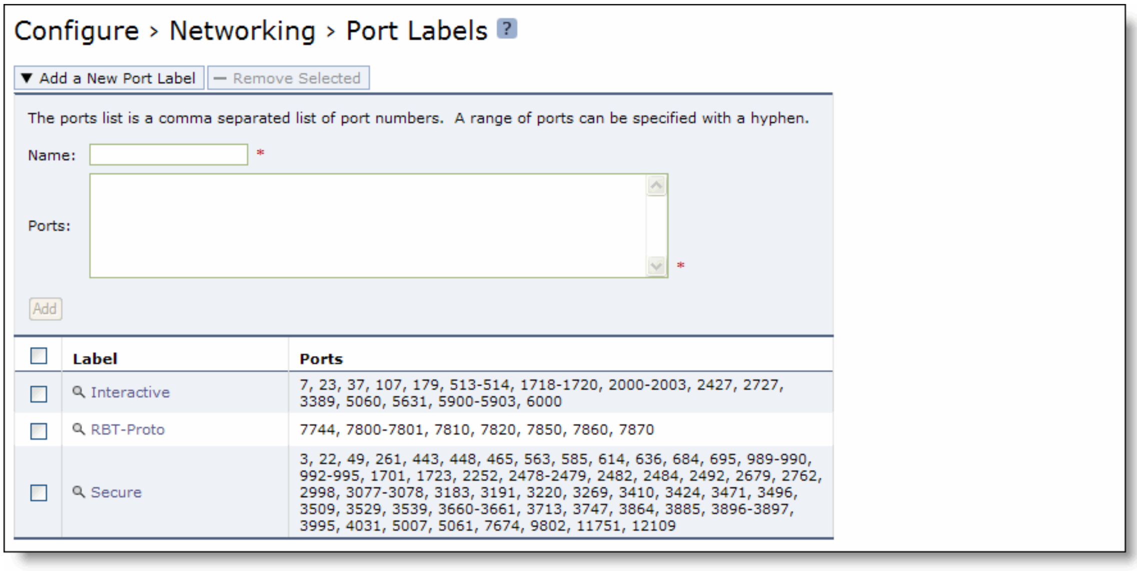

Configure > Networking > Port Labels

The ports list is a comma separated list of port numbers. A range of ports can be specified with a hyphen.

Name: *

Ports: *

□	Label	Ports
□	Interactive	7, 23, 37, 107, 179, 513-514, 1718-1720, 2000-2003, 2427, 2727, 3389, 5060, 5631, 5900-5903, 6000
□	RBT-Proto	7744, 7800-7801, 7810, 7820, 7850, 7860, 7870
□	Secure	3, 22, 49, 261, 443, 448, 465, 563, 585, 614, 636, 684, 695, 989-990, 992-995, 1701, 1723, 2252, 2478-2479, 2482, 2484, 2492, 2679, 2762, 2998, 3077-3078, 3183, 3191, 3220, 3269, 3410, 3424, 3471, 3496, 3509, 3529, 3539, 3660-3661, 3713, 3747, 3864, 3885, 3896-3897, 3995, 4031, 5007, 5061, 7674, 9802, 11751, 12109

2. To add a port label, complete the configuration as described in the following table.

Control	Description
Add a New Port Label	Displays the controls to add a new port label.
Name	Specify the label name. The following rules apply: • Port labels are not case sensitive and can be any string consisting of letters, the underscore (_), or the hyphen (-). There cannot be spaces in port labels. • The fields in the various rule pages of the Management Console that take a physical port number also take a port label. • To avoid confusion, do not use a number for a port label. • Port labels that are used in in-path and other rules, such as QoS and peering rules, cannot be deleted. • Port label changes (that is, adding and removing ports inside a label) are applied immediately by the rules that use the port labels that you have modified.
Ports	Specify a comma-separated list of ports.
Remove Selected	Select the check box next to the name and click **Remove Selected**.
Add	Adds the port label.

3. Click **Save** to save your settings permanently.

Modifying Ports in a Port Label

You can add or delete ports associated with a port label in the Port Label: <Port Label Name> page.

To modify ports in a port label

1. Choose Configure > Networking > Port Labels to display the Port Labels page.

2. Select the port label name in the Port Labels list to display the Editing Port Labels Interactive group.

Figure 4-11. Editing Port Labels Page

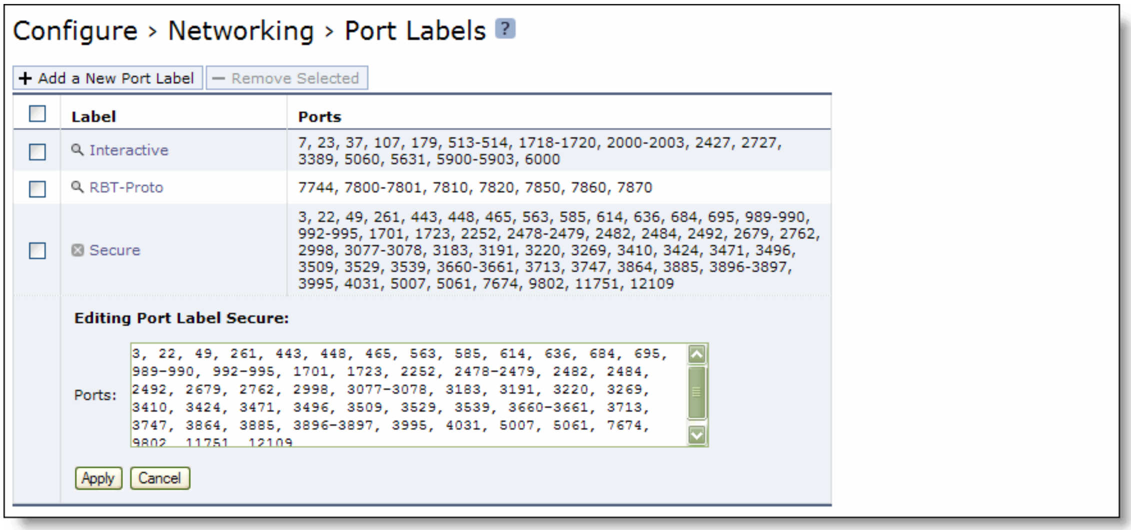

□	Label	Ports
□	Interactive	7, 23, 37, 107, 179, 513-514, 1718-1720, 2000-2003, 2427, 2727, 3389, 5060, 5631, 5900-5903, 6000
□	RBT-Proto	7744, 7800-7801, 7810, 7820, 7850, 7860, 7870
□	☒ Secure	3, 22, 49, 261, 443, 448, 465, 563, 585, 614, 636, 684, 695, 989-990, 992-995, 1701, 1723, 2252, 2478-2479, 2482, 2484, 2492, 2679, 2762, 2998, 3077-3078, 3183, 3191, 3220, 3269, 3410, 3424, 3471, 3496, 3509, 3529, 3539, 3660-3661, 3713, 3747, 3864, 3885, 3896-3897, 3995, 4031, 5007, 5061, 7674, 9802, 11751, 12109

Editing Port Label Secure:

Ports: 3, 22, 49, 261, 443, 448, 465, 563, 585, 614, 636, 684, 695, 989-990, 992-995, 1701, 1723, 2252, 2478-2479, 2482, 2484, 2492, 2679, 2762, 2998, 3077-3078, 3183, 3191, 3220, 3269, 3410, 3424, 3471, 3496, 3509, 3529, 3539, 3660-3661, 3713, 3747, 3864, 3885, 3896-3897, 3995, 4031, 5007, 5061, 7674, 9802, 11751, 12109

3. Under Editing Port Label <port label name>, add or delete ports in the Ports text box.

4. Click **Apply** to save your settings to the running configuration; click **Cancel** to cancel your changes.

5. Click **Save** to save your settings permanently.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring Citrix ICA Optimization" on page 137
- "Creating QoS Classes" on page 288
- "Configuring QoS Marking" on page 300

Configuring CIFS Optimization

This section describes how to optimize CIFS. It includes the following sections:

- "Optimizing CIFS SMB1" on page 97
- "Optimizing SMB2" on page 101
- "Configuring SMB Signing" on page 103

You can display and modify CIFS optimization and SMB Signing settings in the Configure > Optimization > CIFS (SMB1) page and the Configure > Optimization > SMB2 pages.

RiOS v5.5x and later includes settings to optimize Microsoft Office and CIFS traffic with SMB signing enabled.

RiOS v6.0 and later supports CIFS latency optimization and SMB Signing settings for Mac OS X 10.5.x and later clients.

RiOS v6.5 supports SMB2 latency optimization.

CIFS latency optimization does not require a separate license and is enabled by default.

Typically, you disable CIFS optimizations only to troubleshoot the system.

Optimizing CIFS SMB1

CIFS SMB1 optimization performs latency and SDR optimizations on SMB1 traffic. Without this feature, Steelhead appliances perform only SDR optimization without improving CIFS latency.

Important: You must restart the client Steelhead service after enabling the SMB1 latency optimization.

To display CIFS optimization settings for SMB1

1. Choose Configure > Optimization > CIFS (SMB1) to display the CIFS (SMB1) page.

Figure 4-12. CIFS SMB1 Page

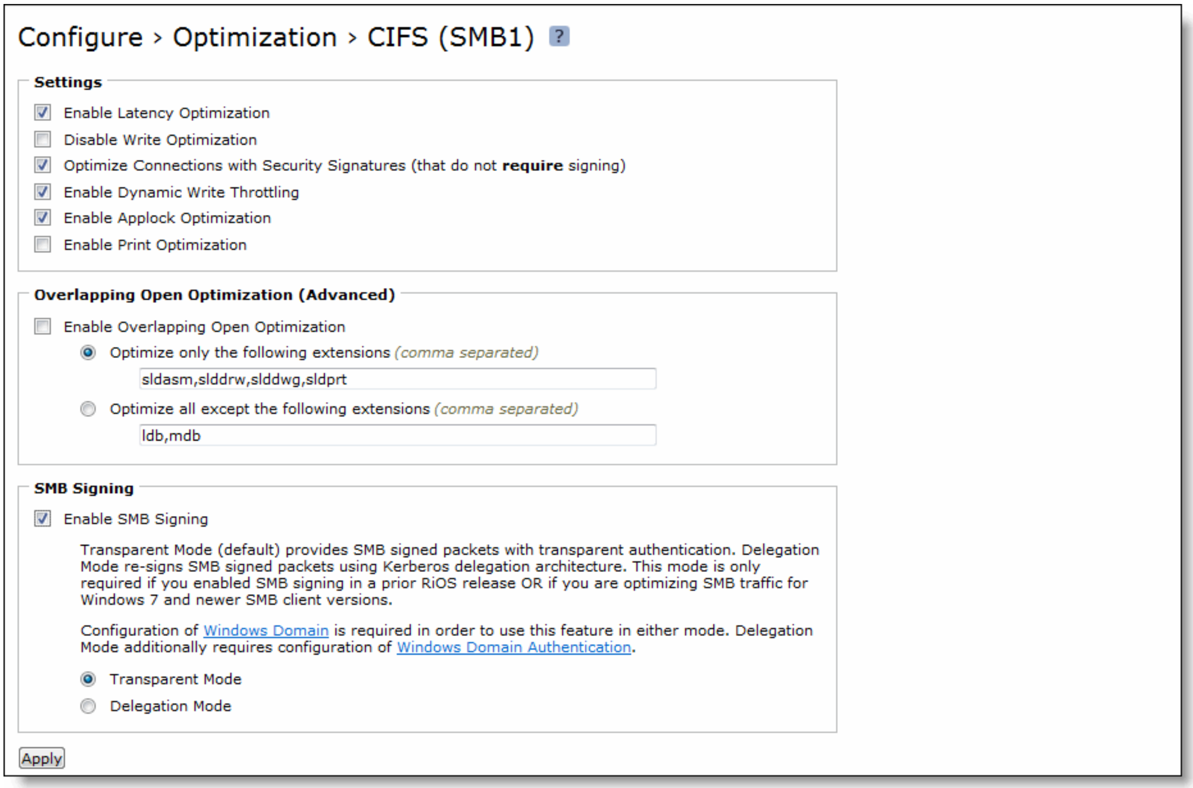

Configure > Optimization > CIFS (SMB1)

Settings

☑ Enable Latency Optimization

□ Disable Write Optimization

☑ Optimize Connections with Security Signatures (that do not **require** signing)

☑ Enable Dynamic Write Throttling

☑ Enable Applock Optimization

□ Enable Print Optimization

Overlapping Open Optimization (Advanced)

□ Enable Overlapping Open Optimization

● Optimize only the following extensions *(comma separated)*

sldasm,slddrw,slddwg,sldprt

○ Optimize all except the following extensions *(comma separated)*

ldb,mdb

SMB Signing

☑ Enable SMB Signing

Transparent Mode (default) provides SMB signed packets with transparent authentication. Delegation Mode re-signs SMB signed packets using Kerberos delegation architecture. This mode is only required if you enabled SMB signing in a prior RiOS release OR if you are optimizing SMB traffic for Windows 7 and newer SMB client versions.

Configuration of Windows Domain is required in order to use this feature in either mode. Delegation Mode additionally requires configuration of Windows Domain Authentication.

● Transparent Mode

○ Delegation Mode

2. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable Latency Optimization	Enables SMB1 optimized connections for file opens and reads. This is the fundamental component of the CIFS module and is required for base optimized connections for file opens and reads. Although latency optimization incorporates several hundred individual optimized connection types, the most frequent type of file opens is where exclusive opportunistic locks have been granted, and read-ahead operations are initiated on the file data. RiOS optimizes the bandwidth used to transfer the read-ahead data from the server side to the client side. Only clear this check box if you want to disable latency optimization. Typically, you disable latency optimization to troubleshoot problems with the system. **Important:** Latency optimization must be enabled (or disabled) on both Steelhead appliances.
Disable Write Optimization	Disables write optimization. Disable write optimization only if you have applications that assume and require write-through in the network. If you disable write optimization, the Steelhead appliance still provides optimization for CIFS reads and for other protocols, but you might experience a slight decrease in overall optimization. Most applications operate safely with write optimization because CIFS allows you to explicitly specify write-through on each write operation. However, if you have an application that does not support explicit write-through operations, you must disable it in the Steelhead appliance. If you do not disable write-through, the Steelhead appliance acknowledges writes before they are fully committed to disk, to speed up the write operation. The Steelhead appliance does not acknowledge the file close until the file is safely written.
Optimize Connections with Security Signatures (that do not require signing)	Prevents Windows SMB signing. This is the default setting. This feature automatically stops Windows SMB signing. SMB signing prevents the Steelhead appliance from applying full optimization on CIFS connections and significantly reduces the performance gain from a Steelhead deployment. Because many enterprises already take additional security precautions (such as firewalls, internal-only reachable servers, and so on), SMB signing adds little additional security, at a significant performance cost (even without Steelhead appliances). Before you enable this feature, consider the following factors: • If the client-side machine has Required signing, enabling this feature prevents the client from connecting to the server. • If the server-side machine has Required signing, the client and the server connect but you cannot perform full latency optimization with the Steelhead appliance. Domain Controllers default to Required. **Important:** If your deployment requires SMB signing, you can optimize signed CIFS messages using the Enable SMB Signing feature. For details about SMB signing and the performance cost associated with it, see the *Riverbed Deployment Guide*.
Enable Dynamic Write Throttling	Enables CIFS dynamic throttling mechanism which replaces the current static buffer scheme. If you enable CIFS dynamic throttling, it is activated only when there are sub-optimal conditions on the server-side causing a backlog of write messages; it does not have a negative effect under normal network conditions.

Control	Description
Enable Applock Optimization	Enables CIFS latency optimizations to improve read and write performance for Microsoft Word (.doc) and Excel (.xls) documents when multiple users have the file open. This setting is enabled by default in v6.0 and later. This feature enhances the Enable Overlapping Open Optimization feature by identifying and obtaining locks on read write access at the application level. The overlapping open optimization feature handles locks at the file level. **Note:** Enable the applock optimization feature on the client-side Steelhead appliance. The client-side Steelhead appliance must be running RiOS v5.5 or later.
Enable Print Optimization	Improves centralized print traffic performance. For example, when the print server is located in the data center and the printer is located in the branch office, enabling this option speeds the transfer of a print job spooled across the WAN to the server and back again to the printer. By default, this setting is disabled. Enabling this option requires an optimization service restart. This option supports Windows XP (client), Vista (client), Windows 2003 (server), and Windows 2008 (server). Both the client and server-side Steelhead appliance must be running RiOS v6.0 or later. **Note:** This feature does not improve optimization for a Windows Vista client printing over a Windows 2008 server, because this client and server pair uses a different print protocol.

3. Click **Apply** to apply your settings to the current configuration.

4. Click **Save** to save your settings permanently.

5. If you enabled print optimization, you must restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Note: For details on SMB signing, see "Configuring SMB Signing" on page 103.

To enable Overlapping Open Optimization

1. On the client-side Steelhead appliance, under Overlapping Open Optimization (Advanced), complete the configuration as described in the following table.

Control	Description
Enable Overlapping Open Optimization	Enables overlapping opens to obtain better performance with applications that perform multiple opens on the same file (for example, CAD applications). By default, this setting is disabled. **Note:** Enable this setting on the client-side Steelhead appliance. With overlapping opens enabled the Steelhead appliance optimizes data where exclusive access is available (in other words, when locks are granted). When an oplock is not available, the Steelhead appliance does not perform application-level latency optimizations but still performs SDR and compression on the data as well as TCP optimizations. **Note:** If a remote user opens a file that is optimized using the overlapping opens feature and a second user opens the same file, they might receive an error if the file fails to go through a v3.x.x or later Steelhead appliance or if it does not go through a Steelhead appliance (for example, certain applications that are sent over the LAN). If this occurs, disable overlapping opens for those applications. Use the radio buttons to set either an include list or exclude list of file types subject to overlapping opens optimization.
Optimize only the following extensions	Specify a list of extensions you want to include in overlapping open optimization.
Optimize all except the following extensions	Specify a list of extensions you do not want to include; for example, specify any file extensions that Enable Applock Optimization is being used for.

2. Click **Apply** to apply your settings to the current configuration.

3. Click **Save** to save your settings permanently.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Optimizing SMB2

RiOS v6.5 includes support for SMB2 traffic latency optimization for native SMB2 clients and servers. SMB2 allows more efficient access across disparate networks. It is the default mode of communication between Windows Vista and Windows Server 2008. Microsoft has subsequently modified SMB2 again (to SMB v2.1) for Windows 7 and Windows Server 2008 R2.

SMB2 brought a number of improvements, including but not limited to:

- A vastly reduced set of opcodes (a total of only 18); in contrast SMBv1 has over 70 separate opcodes. Note that use of SMB2 does not result in lost functionality (most of the SMB1 opcodes were redundant).
- General mechanisms for data pipelining and lease-based flow control.
- Request compounding which allows multiple SMB requests to be sent as a single network request.

- Larger reads and writes provide for more efficient use of networks with high latency.
- Caching of folder and file properties, where clients keep local copies of folders and files.
- Improved scalability for file sharing (number of users, shares and open files per server greatly increased).

To display optimization settings for SMB2

1. Choose Configure > Optimization > SMB2 to display the SMB2 page.

Figure 4-13. SMB2 Page

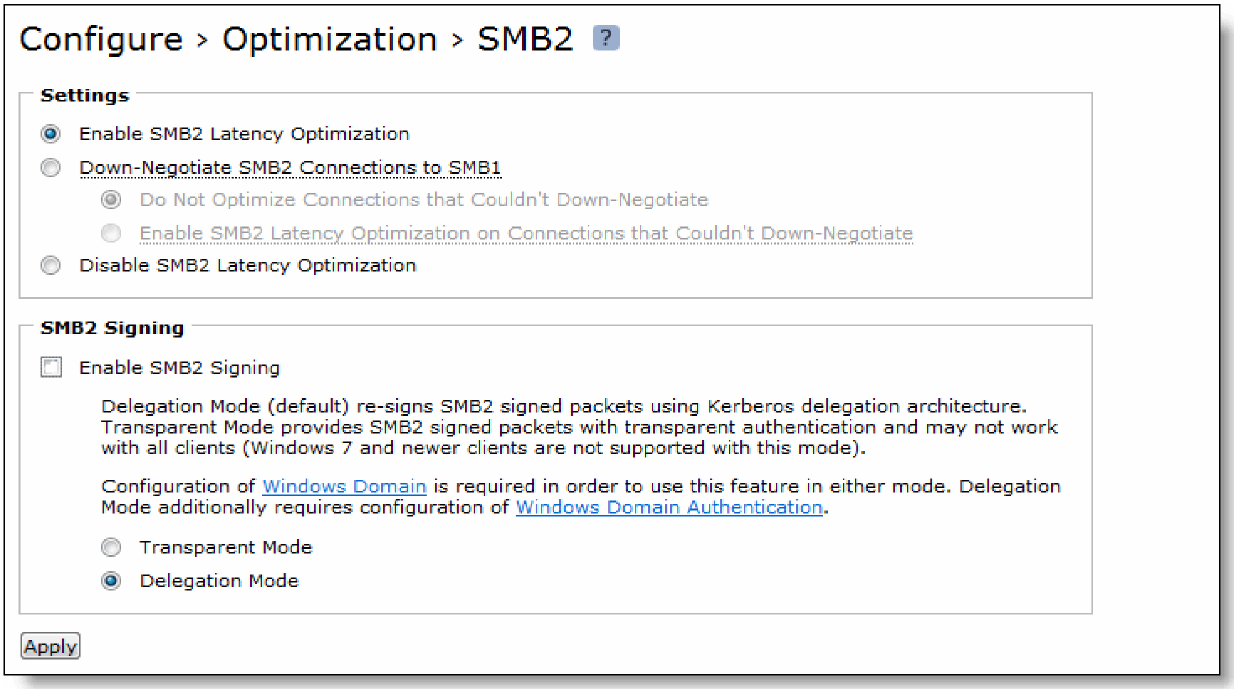

Configure > Optimization > SMB2

Settings

● Enable SMB2 Latency Optimization

○ Down-Negotiate SMB2 Connections to SMB1

   ● Do Not Optimize Connections that Couldn't Down-Negotiate

   ○ Enable SMB2 Latency Optimization on Connections that Couldn't Down-Negotiate

○ Disable SMB2 Latency Optimization

SMB2 Signing

□ Enable SMB2 Signing

Delegation Mode (default) re-signs SMB2 signed packets using Kerberos delegation architecture. Transparent Mode provides SMB2 signed packets with transparent authentication and may not work with all clients (Windows 7 and newer clients are not supported with this mode).

Configuration of Windows Domain is required in order to use this feature in either mode. Delegation Mode additionally requires configuration of Windows Domain Authentication.

○ Transparent Mode

● Delegation Mode

2. Under Settings, complete the configuration on both the client-side and server-side Steelhead appliances as described in the following table.

Control	Description
Enable SMB2 Latency Optimization	Performs SMB2 latency optimization in addition to the existing bandwidth optimization features. These optimizations include cross-connection caching, read-ahead, write-behind, and batch prediction among several other techniques to ensure low latency transfers. RiOS maintains the data integrity and the client always receives data directly from the servers. By default, SMB2 optimization is disabled. **Important:** You must enable (or disable) SMB2 latency optimization on both the client-side and server-side Steelhead appliances and both Steelheads must be running RiOS v6.5. After enabling SMB2 optimization, you must restart the optimization service.
Down-negotiate SMB2 connections to SMB1	Enable on the client-side Steelhead appliance. Optimizes connections that are successfully negotiated down to SMB1 according to the settings on the Configure > Optimization > CIFS (SMB1) page. Down negotiation is bypassed when the client or the server is configured to only use SMB2 or the client has already established an SMB2 connection with the server. If the client already has a connection with the server, you need to restart the client.
Do Not Optimize Connections that Couldn't Down Negotiate	Specifies that the Steelhead appliance does not optimize the connection when it is unable to negotiate down to SMB1.
Enable SMB2 Latency Optimization on Connections that Couldn't Down Negotiate	Enable to use SMB1 latency optimization when possible, but use SMB2 latency optimization when the Steelhead appliance is unable to negotiate down to SMB1.
Disable SMB2 Latency Optimization	Disables SMB2 latency optimization. **Important:** You must enable (or disable) SMB2 latency optimization on both the client-side and server-side Steelhead appliances and both Steelheads must be running RiOS v6.5. After enabling SMB2 optimization, you must restart the optimization service.

3. Click **Apply** to apply your settings to the current configuration.

Related Topic

- "Configuring CIFS Prepopulation" on page 86

Configuring SMB Signing

You can display and modify SMB signing settings in the Configure > Optimization > CIFS page.

When sharing files, Windows provides the ability to sign CIFS messages to prevent man-in-the-middle attacks. Each CIFS message has a unique signature which prevents the message from being tampered with. This security feature is called SMB signing. Prior to v5.5, RiOS did not provide latency optimization for signed traffic.

You can enable the RiOS SMB signing feature on a server-side Steelhead appliance to alleviate latency in file access with CIFS acceleration while maintaining message security signatures. With SMB signing on, the Steelhead appliance optimizes CIFS traffic by providing bandwidth optimizations (SDR and LZ), TCP optimizations, and CIFS latency optimizations—even when the CIFS messages are signed.

RiOS v6.5 includes support for optimizing SMB2-signed traffic for native SMB2 clients and servers. SMB2 signing support includes:

- Windows domain integration, including domain join and domain-level support
- Authentication using transparent mode and delegation mode. Delegation mode is the default for SMB2. Transparent mode works out-of-the-box with Windows Vista (but not Windows 7). For details, see “Authentication” on page 104.
- Secure inner-channel SSL support. For details, see “Configuring Secure Peers” on page 231.

By default, RiOS SMB2 signing is disabled.

Domain Security

The RiOS SMB signing feature works with Windows domain security and is fully compliant with the Microsoft SMB-signing v1 and v2 protocols. RiOS v6.1 and later supports domain security in both native and mixed modes for:

- Windows 2000
- Windows 2003 R2
- Windows 2008
- Windows 2008 R2

The server-side Steelhead appliance in the path of the signed CIFS traffic becomes part of the Windows trust domain. The Windows domain is either the same as the domain of the user or has a trust relationship with the domain of the user. The trust relationship can be either a parent-child relationship or an unrelated trust relationship.

RiOS v6.0 and later optimizes signed CIFS traffic even when the logged-in user or client machine and the target server belong to different domains, provided these domains have a trust relationship with the domain the Steelhead appliance has joined. RiOS v6.1 and later supports delegation for users that are in domains trusted by the server's domain. The trust relationships include:

- A basic parent and child domain relationship. Users from the child domain try to access CIFS/MAPI servers in the parent domain. For example, users in ENG.RVBD.COM accessing servers in RVBD.COM.
- A grandparent and child domain relationship. Users from grandparent domain access resources from the child domain. For example, users from RVBD.COM accessing resources in DEV.ENG.RVBD.COM.
- A sibling domain relationship. For example, users from ENG.RVBD.COM try to access resources in MARKETING.RVBD.COM.

Authentication

The process RiOS uses to authenticate domain users depends upon the release version.

RiOS v5.5.x uses Kerberos authentication between the server-side Steelhead appliance and any configured servers participating in the signed session. It uses NTLM authentication between the client-side and server-side Steelhead appliances.

RiOS v6.0 and later features two authentication modes:

- **Delegation mode** - uses Kerberos delegation architecture to authenticate signed packets between the server-side Steelhead appliance and any configured servers participating in the signed session. NTLM is used between the client-side and server-side Steelhead appliance. This is the default mode for SMB2. SMB2 delegation mode in RiOS v6.5 supports Windows 7 and Samba 4 clients. Delegation mode requires additional configuration of Windows Domain Authentication.

- **Transparent mode** - uses NTLM authentication end-to-end between the client-side and server-side Steelhead appliances and the server-side Steelhead and the server. This is the default mode for SMB1. Transparent mode in RiOS v6.1 and later supports all Windows servers, including Windows 2008 R2, that have NTLM enabled. It is easier to configure.

Transparent mode in RiOS v6.1 and later does *not* support:

- Windows 7 clients.
- Windows 2008 R2 domains which have NTLM disabled.
- Windows servers that are in domains with NTLM disabled.

In RiOS v6.0 and later, you can enable extra security using the secure inner channel. The peer Steelhead appliances using the secure channel encrypt signed CIFS traffic over the WAN. For details, see "Configuring Secure Peers" on page 231.

Prerequisites

- With RiOS SMB signing enabled, Steelhead appliances sign the traffic between the client and the client-side Steelhead appliance and between the server and the server-side Steelhead appliance. The traffic is not signed between the Steelhead appliances, but the Steelheads implement their own integrity mechanisms. Whether Steelhead appliances are used or not, SMB-signed traffic is only signed, not encrypted. For maximum security, Riverbed recommends that you configure the Steelhead appliances as SSL peers and use the secure inner channel to secure the traffic between them. For details, see "Configuring Secure Peers" on page 231.
- When upgrading from RiOS v6.1 to v6.5, you might already have a delegate user and be joined to a domain. If so, enabling SMB2 signing will work when enabled with no additional configuration.
- SMB signing requires joining a Windows domain. It is vital to set the correct time zone for joining a domain. The most common reason for failing to join a domain is a significant difference in the system time on the Windows Domain Controller and the Steelhead appliance. When the time on the Domain Controller and the Steelhead appliance do not match, the following error message appears:

  ```
  lt-kinit: krb5_get_init_creds: Clock skew too great
  ```

  Riverbed recommends using NTP time synchronization to synchronize the client and server clocks. It is critical that the Steelhead appliance time is the same as on the Active Directory controller. Sometimes an NTP server is down or inaccessible, in which case there can be a time difference. You can also disable NTP if it is not being used and manually set the time. You must also verify that the time zone is correct. For details, see "Modifying General Host Settings" on page 39. For more troubleshooting, see "Troubleshooting a Domain Join Failure" on page 309.
- Both the client and the server must support SMB2 to use RiOS SMB2 signing.

Verifying the Domain Functional Level and Host Settings

This section describes how to verify the domain and DNS settings before joining the Windows domain and enabling SMB signing.

To verify the domain functional level (Delegation mode only)

1. If you are using delegation mode, verify that the Windows domain functionality is at the Windows 2003 level or higher. In Windows, open Active Directory Users and Computers on the Domain Controller, choose Domain Name, right-click, and select Raise Domain functionality level. If the domain is not

already at the Windows 2003 level or higher, manually raise the domain functionality. For details on delegation mode, see “Enabling SMB Signing” on page 108.

Note: Once you raise the domain level, you cannot lower it.

Figure 4-14. Verifying the Domain Level Before Enabling SMB Signing

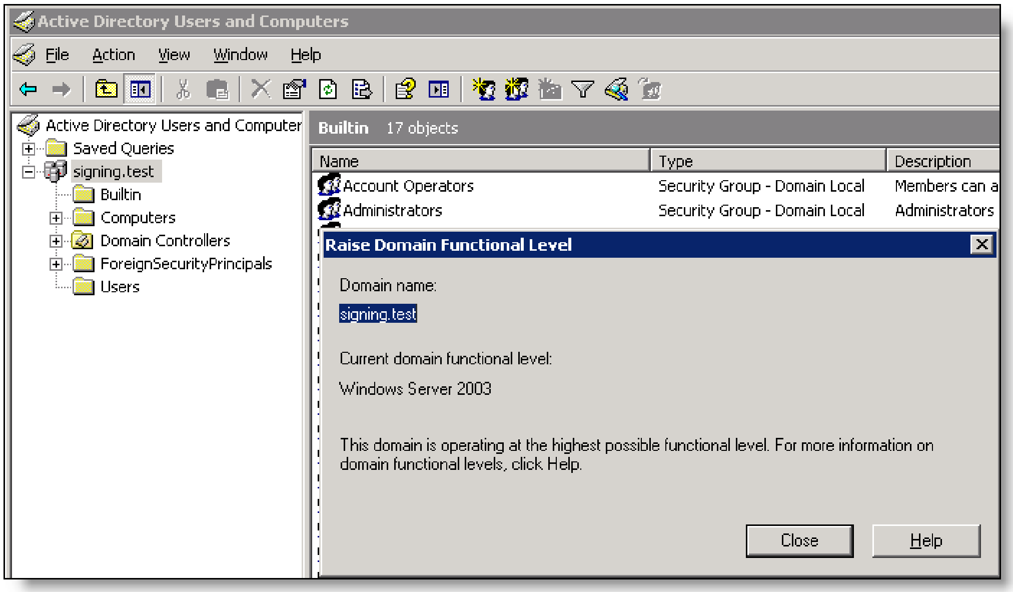
Active Directory Users and Computers
File Action View Window Help
Active Directory Users and Computer
Saved Queries
signing.test
Builtin
Computers
Domain Controllers
ForeignSecurityPrincipals
Users
Builtin 17 objects
Raise Domain Functional Level
Close
Help

Name	Type	Description
Account Operators	Security Group - Domain Local	Members can a
Administrators	Security Group - Domain Local	Administrators

Domain name:

signing.test

Current domain functional level:

Windows Server 2003

This domain is operating at the highest possible functional level. For more information on domain functional levels, click Help.

For details, see the Microsoft Windows Server 2003 Active Directory documentation

http://www.microsoft.com/windowsserver2003technologies/directory/activedirectory/default.mspx

2. Identify the full domain name, which must be the same as DNS. You need to specify this name when you join the server-side Steelhead appliance to the domain.

3. Identify the short (NetBIOS) domain name by pressing Ctrl+Alt+Del on any member server. You need to explicitly specify the short domain name when the Steelhead appliance joins the domain if it does not match the leftmost portion of the fully-qualified domain name.

4. Make sure that the primary or auxiliary interface for the server-side Steelhead appliance is routable to the DNS and the Domain Controller.

5. Verify the DNS settings.

   You must be able to ping the server-side Steelhead appliance, by name, from a CIFS server joined to the same domain that the server-side Steelhead appliance joins. If you cannot, you must manually create an entry in the DNS server for the server-side Steelhead appliance and perform a DNS replication prior to joining the Windows Domain. The Steelhead appliance does not automatically register the required DNS entry with the Windows Domain Controller.

You must be able to ping the Domain Controller, by name, whose domain the server-side Steelhead appliance joins. If you cannot, choose Configure > Networking > Host Settings to configure the DNS settings.

Figure 4-15. Verifying the DNS Settings for SMB Signing

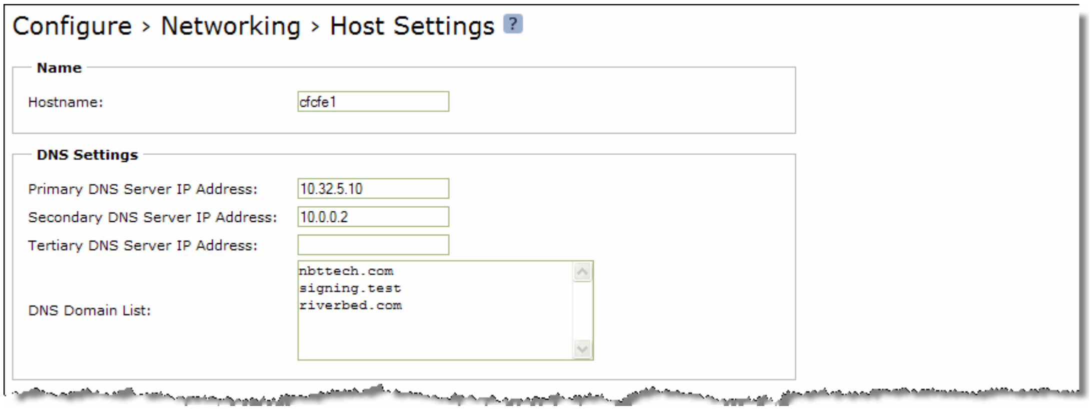

Configure > Networking > Host Settings

Name

Hostname: cfcfe1

DNS Settings

Primary DNS Server IP Address: 10.32.5.10

Secondary DNS Server IP Address: 10.0.0.2

Tertiary DNS Server IP Address:

DNS Domain List:

```
nbttech.com
signing.test
riverbed.com
```

For details, see “Modifying General Host Settings” on page 39.

The next step is to join a Windows Domain.

To join a Windows domain

- Choose Configure > Networking > Windows Domain on the server-side Steelhead appliance and join the domain.

Figure 4-16. Windows Domain Page

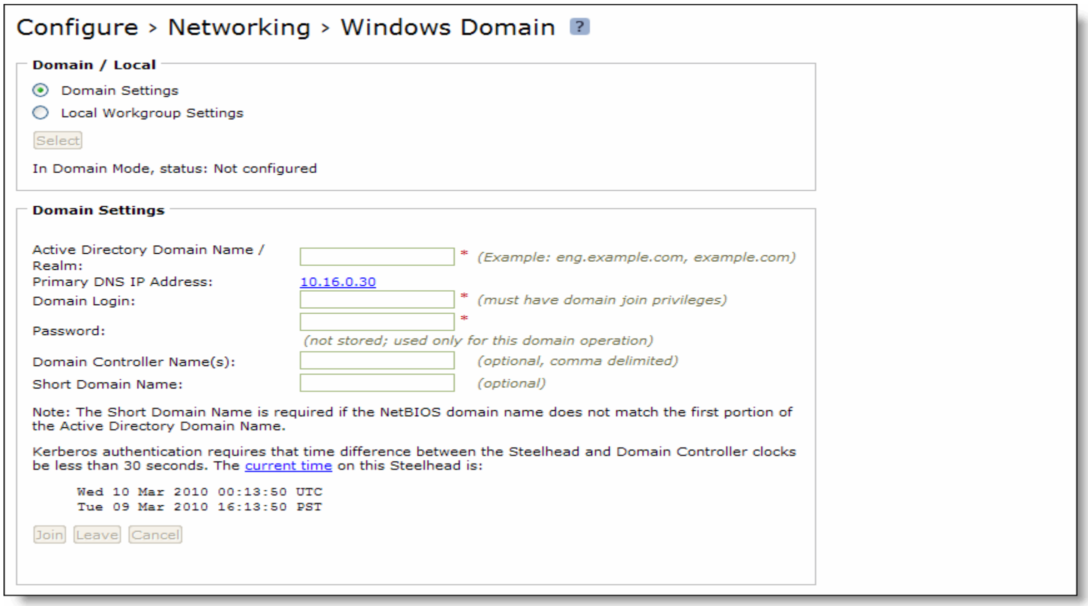

Configure > Networking > Windows Domain

Domain / Local

● Domain Settings

○ Local Workgroup Settings

Select

In Domain Mode, status: Not configured

Domain Settings

Active Directory Domain Name / Realm: * (Example: eng.example.com, example.com)

Primary DNS IP Address: 10.16.0.30

Domain Login: * (must have domain join privileges)

Password: * (not stored; used only for this domain operation)

Domain Controller Name(s): (optional, comma delimited)

Short Domain Name: (optional)

Note: The Short Domain Name is required if the NetBIOS domain name does not match the first portion of the Active Directory Domain Name.

Kerberos authentication requires that time difference between the Steelhead and Domain Controller clocks be less than 30 seconds. The current time on this Steelhead is:

```
Wed 10 Mar 2010 00:13:50 UTC
Tue 09 Mar 2010 16:13:50 PST
```

For details, see “Joining a Windows Domain or Workgroup” on page 305. After you have joined the domain, the next step is to enable SMB signing.

Enabling SMB Signing

Now that you have joined a Windows domain you can enable SMB signing.

Important: When SMB signing is set to Enabled for both the client and server-side SMB component (but not set to Required), and the RiOS Optimize Connections with Security Signatures feature is enabled, it takes priority and prevents SMB signing. You can resolve this by disabling the Optimize Connections with Security Signatures feature and restarting the Steelhead appliance *before* enabling this feature.

Important: The RiOS Optimize Connections with Security Signatures feature can lead to unintended consequences in the scenario when SMB signing is required on the client, but set to Enabled on the server. With this feature enabled the client concludes that the server does not support signing and may terminate the connection with the server as a result. You can resolve this by using one of the following procedures *before* enabling this feature:

1. Disable the Optimize Connections with Security Signatures feature and restart the Steelhead appliance.

—or—

2. Apply a Microsoft Service pack update to the clients (recommended). You can download the update from the Microsoft Download Center: http://support.microsoft.com/kb/916846

To enable SMB1 signing

1. On the server-side Steelhead appliance, choose Configure > Optimization > CIFS (SMB1) to display the CIFS page.

Figure 4-17. CIFS SMB1 Page

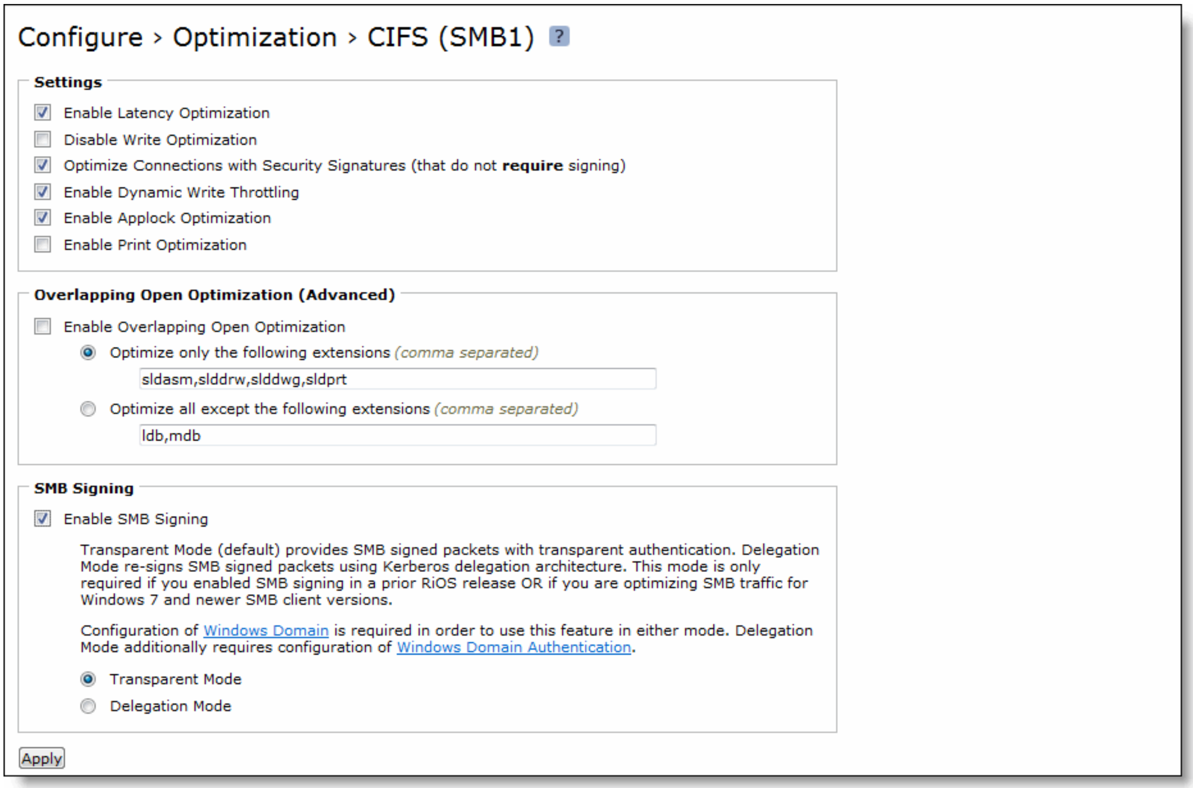

Configure > Optimization > CIFS (SMB1) ?

Settings

☑ Enable Latency Optimization
□ Disable Write Optimization
☑ Optimize Connections with Security Signatures (that do not **require** signing)
☑ Enable Dynamic Write Throttling
☑ Enable Applock Optimization
□ Enable Print Optimization

Overlapping Open Optimization (Advanced)

□ Enable Overlapping Open Optimization
● Optimize only the following extensions *(comma separated)*
sldasm,slddrw,slddwg,sldprt
○ Optimize all except the following extensions *(comma separated)*
ldb,mdb

SMB Signing

☑ Enable SMB Signing

Transparent Mode (default) provides SMB signed packets with transparent authentication. Delegation Mode re-signs SMB signed packets using Kerberos delegation architecture. This mode is only required if you enabled SMB signing in a prior RiOS release OR if you are optimizing SMB traffic for Windows 7 and newer SMB client versions.

Configuration of Windows Domain is required in order to use this feature in either mode. Delegation Mode additionally requires configuration of Windows Domain Authentication.

● Transparent Mode
○ Delegation Mode

2. Under SMB Signing, complete the configuration as described in the following table.

Control	Description
Enable SMB Signing	Enables CIFS traffic optimization by providing bandwidth optimizations (SDR and LZ), TCP optimizations, and CIFS latency optimizations even when the CIFS messages are signed. By default, this setting is disabled. You must enable this feature on the server-side Steelhead appliance. **Note:** If you enable this feature without first joining a Windows Domain, a message tells you that the Steelhead appliance must join a domain before it can support SMB signing.
Transparent Mode	Provides SMB1 signing with transparent authentication. The server-side Steelhead uses NTLM to authenticate users. Select transparent mode with Vista for the simplest configuration.
Delegation Mode	Re-signs SMB signed packets using the Kerberos delegation facility. This setting is enabled by default when you enable SMB signing. Delegation mode is required for Windows 7, but works with all clients. Delegation mode requires additional configuration. Choose Configure > Optimization > Windows Domain Authentication or click the link provided in the CIFS Optimization page.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

To enable SMB2 signing

1. On the server-side Steelhead appliance, choose Configure > Optimization > SMB2 to display the SMB2 page.

Figure 4-18. CIFS Page for SMB2 Signing

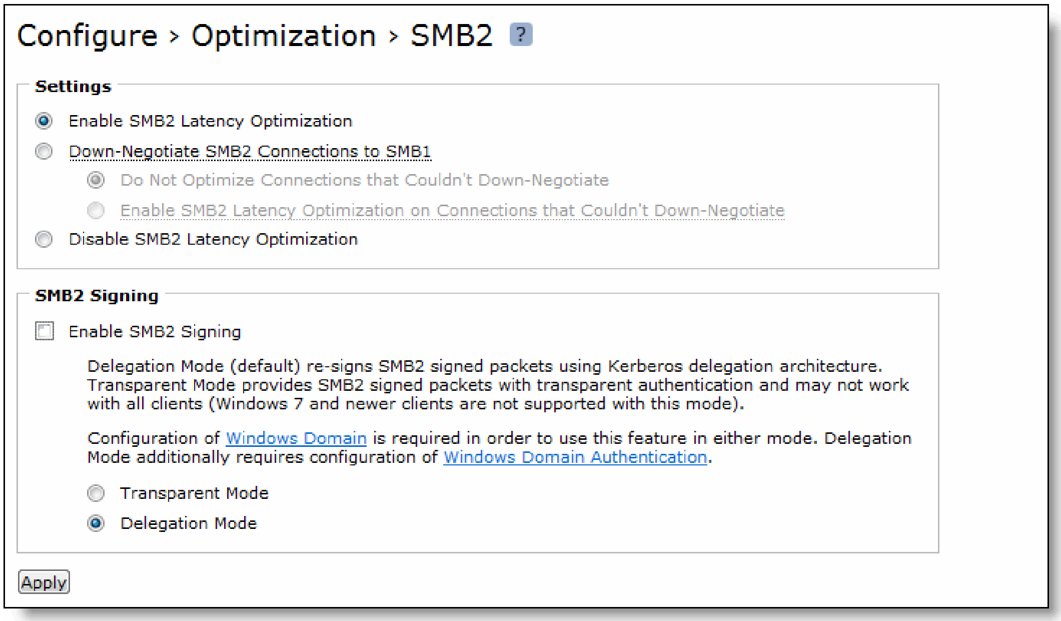

Configure > Optimization > SMB2 ?

Settings

● Enable SMB2 Latency Optimization
○ Down-Negotiate SMB2 Connections to SMB1
  ● Do Not Optimize Connections that Couldn't Down-Negotiate
  ○ Enable SMB2 Latency Optimization on Connections that Couldn't Down-Negotiate
○ Disable SMB2 Latency Optimization

SMB2 Signing

□ Enable SMB2 Signing

Delegation Mode (default) re-signs SMB2 signed packets using Kerberos delegation architecture. Transparent Mode provides SMB2 signed packets with transparent authentication and may not work with all clients (Windows 7 and newer clients are not supported with this mode).

Configuration of Windows Domain is required in order to use this feature in either mode. Delegation Mode additionally requires configuration of Windows Domain Authentication.

○ Transparent Mode
● Delegation Mode

2. Under SMB2 Signing, complete the configuration as described in the following table.

Control	Description
Enable SMB2 Signing	Enables SMB2 traffic optimization by providing bandwidth optimizations (SDR and LZ), TCP optimizations, and SMB2 latency optimizations even when the SMB2 messages are signed. By default, this setting is disabled. You must enable this feature on the server-side Steelhead appliance. **Important:** If you are upgrading from RiOS v6.1 to v6.5, you might already have a delegate user and be joined to a domain. If so, enabling SMB2 signing will work when enabled with no additional configuration. **Note:** If you enable this feature without first joining a Windows Domain, a message tells you that the Steelhead appliance must join a domain before it can support SMB2 signing.
Transparent Mode	Provides SMB2 signing with transparent authentication. The server-side Steelhead uses NTLM to authenticate users. Select transparent mode with Vista for the simplest configuration.
Delegation Mode	Re-signs SMB2 signed packets using the Kerberos delegation facility. This setting is enabled by default when you enable SMB2 signing. Delegation mode is required for Windows 7, but works with all clients. Delegation mode requires additional configuration. Choose Configure > Optimization > Windows Domain Authentication or click the link in the CIFS Optimization page.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. If you change the SMB2 mode, you must restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Related Topics

- "Configuring CIFS Prepopulation" on page 86
- "Windows Domain Authentication" on page 148
- "Joining a Windows Domain or Workgroup" on page 305
- "Viewing Current Connections" on page 385

Configuring HTTP Optimization

This section describes how to configure HTTP optimization features. It includes the following sections:

- "About HTTP Optimization" on page 111
- "Configuring HTTP Optimization Feature Settings" on page 114

About HTTP Optimization

A typical Web page is not a single file that is downloaded all at once. Instead, Web pages are composed of dozens of separate objects—including .jpg and .gif images, JavaScript code, cascading style sheets, and more—each of which must be requested and retrieved separately, one after the other. Given the presence of latency, this behavior is highly detrimental to the performance of Web-based applications over the WAN. The higher the latency, the longer it takes to fetch each individual object and, ultimately, to display the entire page.

HTTP optimization works for most HTTP and HTTPS applications, including SAP, Customer Relationship Management, Enterprise Resource Planning, Financials, Document Management, and Intranet portals.

The RiOS HTTP latency optimizations include features that target different types of Web applications. The following features can be used individually or in combination with each other.

- **URL Learning** - The Steelhead appliance learns associations between a base request and a follow-on request. This feature is most effective for Web applications with large amounts of static content, for example, images, style sheets, and so on. Instead of saving each object transaction, the Steelhead appliance saves only the request URL of object transactions in a Knowledge Base and then generates related transactions from the list. This feature uses the Referer header field to generate relationships between object requests and the base HTML page that referenced them and to group embedded objects. This information is stored in an internal HTTP database. The following objects are retrieved by default: .gif, .jpg, .css, .js, .png. You can add more object types to be retrieved.
- **Parse and Prefetch** - The Steelhead appliance includes a specialized algorithm that determines which objects are going to be requested for a given Web page and prefetches them so that they are readily available when the client makes its requests. This feature complements the URL Learning feature by handling dynamically generated pages and URLs that include state information.

Parse and Prefetch reads a page, finds HTML tags that it recognizes as containing a prefetchable object, and sends out prefetch requests for those objects. Typically, a client would need to request the base page, parse it, and then send out requests for each of these objects. This still occurs, but with Parse and Prefetch the Steelhead appliance has quietly perused the page before the client receives it and has already sent out the requests. This allows it to serve the objects as soon as the client requests them, rather than forcing the client to wait on a slow WAN link.

For example, when an HTML page contains the tag <img src=”my_picture.gif”>, the Steelhead appliance prefetches the image my_picture.gif because it parses an img tag with an attribute of src by default. The HTML tags that are prefetched by default are base/href, body/background, img/src, link/href, and script/src. You can add additional object types to be prefetched.

- **Removal of Unfetchable Objects** - The Steelhead appliance removes unfetchable objects from the URL Learning Knowledge Base.

- **Object Prefetch Table** - The Steelhead appliance stores object prefetches from HTTP GET requests for cascading style sheets, static images, and JavaScript files. This helps the client-side Steelhead appliance respond to If-Modified-Since (IMS) requests and regular requests from the client, thus cutting back on round trips across the WAN. This feature is useful for applications that use a lot of cacheable content.

- **Persistent Connections** - The Steelhead appliance uses an existing TCP connection between a client and a server to prefetch objects from the Web server that it determines are about to be requested by the client. Many Web browsers open multiple TCP connections to the Web server when requesting embedded objects. Typically, each of these TCP connections go through a lengthy authentication dialog before the browser can request and receive objects from the Web server on that connection. NTLM is a Microsoft authentication protocol which employs a challenge-response mechanism for authentication, in which clients are required to prove their identities without sending a password to a server. NTLM requires the transmission of three messages between the client (wanting to authenticate) and the server (requesting authentication).

  Because these authentication dialogs are time consuming, if your Web servers require NTLM authentication you can configure your Steelhead appliance to reuse existing NTLM authenticated connections to avoid unnecessarily authenticating extra connections.

All HTTP optimization features are driven by the client-side Steelhead appliance. The client-side Steelhead appliance sends the prefetched information to the server-side Steelhead appliance. Prefetched data and object prefetches are served from the client-side Steelhead appliance upon request from the browser.

You can set up an optimization scheme that applies to all HTTP traffic, or create individual schemes for each server subnet. Therefore, you can configure an optimization scheme that includes your choice of prefetch optimizations for one range of server addresses, with that range encompassing as large a network as you need, from a single address to all possible addresses.

The following situations might affect HTTP optimization:

- **Fat Client** - Not all applications accessed through a Web browser use the HTTP protocol. This is especially true for fat clients that run inside a Web browser which might use proprietary protocols to communicate with a server. HTTP optimization does not improve performance in such cases.

- **Digest for Authentication** - Some Web servers might require users to authenticate themselves before allowing them access to certain Web content. Digest Authentication is one of the less popular Authentication schemes, although it is still supported by most Web servers and browsers. Digest Authentication requires the browser to include a secret value which only the browser and server know how to generate and decode. Because the Steelhead appliance cannot generate these secret values, it cannot prefetch objects protected by Digest Authentication.

- **Object Authentication** - It is uncommon for Web servers to require separate authentication for each object requested by the client, but occasionally Web servers are configured to use *per object authentication*. In such cases, the HTTP prefetch may provide limited performance improvement.

Comparing the HTTP Optimization Features

The following table compares the HTTP optimization features.

	URL Learning	Parse and Prefetch	Object Prefetch Table
The application includes dynamic URLs.	Not effective	Good results	Good results
Is there a learning phase (first user transaction)?	Yes	No	Yes
When does the prefetch occur?	With the base request, after the learning phase	After one Round-Trip Time	N/A
Does the application include embedded object requests from JavaScript and CSS?	Yes	No	Yes

Note: HTTP optimization has been tested on the following browsers: Internet Explorer v6.0 or later, and Firefox v2 or later. HTTP optimization has been tested on the following servers: Apache v1.3, Apache v2.2, Microsoft IIS v5.0 and v6.0, Microsoft Sharepoint, ASP.net, Microsoft Internet Security and Acceleration Server (ISA).

Basic Steps

The following table summarizes the basic steps for configuring HTTP optimization, followed by detailed procedures.

Task	Reference
1. Enable HTTP optimization for prefetching Web objects. This is the default setting.	"Configuring HTTP Optimization Feature Settings" on page 114
2. Specify object prefetch extensions that represent prefetched objects for URL Learning. By default, the Steelhead appliance prefetches .jpg, .gif, .js, .png, and .css objects.	"Configuring HTTP Optimization Feature Settings" on page 114
3. Optionally, specify which HTML tags to prefetch for Parse and Prefetch. By default, the Steelhead appliance prefetches base/href, body/background, img/src, link/href, and script/src HTML tags.	"To prefetch HTML tags" on page 116
4. Optionally, set an HTTP optimization scheme for each server subnet. For example, an optimization scheme can include a combination of the URL Learning, Parse and Prefetch, or Object Prefetch features. The default setting is URL Learning only. RiOS v6.1 and later supports authorization optimizations and basic tuning for server subnets. Riverbed recommends that you enable: • **Strip compression** - removes the Accept-Encoding lines from the HTTP headers that contain gzip or deflate. These Accept-Encoding directives allow Web browsers and servers to send and receive compressed content rather than raw HTML. • **Insert cookie** - tracks repeat requests from the client. • **Insert Keep Alive** - maintains persistent connections. Often this feature is turned off even though the Web server can support it. This is especially true for Apache Web servers that serve HTTPS to Microsoft Internet Explorer browsers.	"Adding a Server Subnet" on page 116
5. If necessary, define in-path rules that specify when to apply HTTP optimization and whether to enable HTTP latency support for HTTPS.	"Configuring In-Path Rules" on page 28

Note: In order for the Steelhead appliance to optimize HTTPS traffic (HTTP over SSL), you must configure a specific in-path rule that enables both SSL optimization and HTTP optimization.

Configuring HTTP Optimization Feature Settings

You can display and modify HTTP optimization feature settings in the Configure > Optimization > HTTP page. For an overview of the HTTP optimization features and basic deployment considerations, see "Configuring HTTP Optimization" on page 111.

Note: All of the HTTP optimization features operate on the client-side Steelhead appliance. As long as the server-side Steelhead appliance is running v4.0.x or later, you configure HTTP optimizations only on the client-side Steelhead appliance.

To display HTTP optimization feature settings or to modify them

1. Choose Configure > Optimization > HTTP to display the HTTP page.

Figure 4-19. HTTP Page

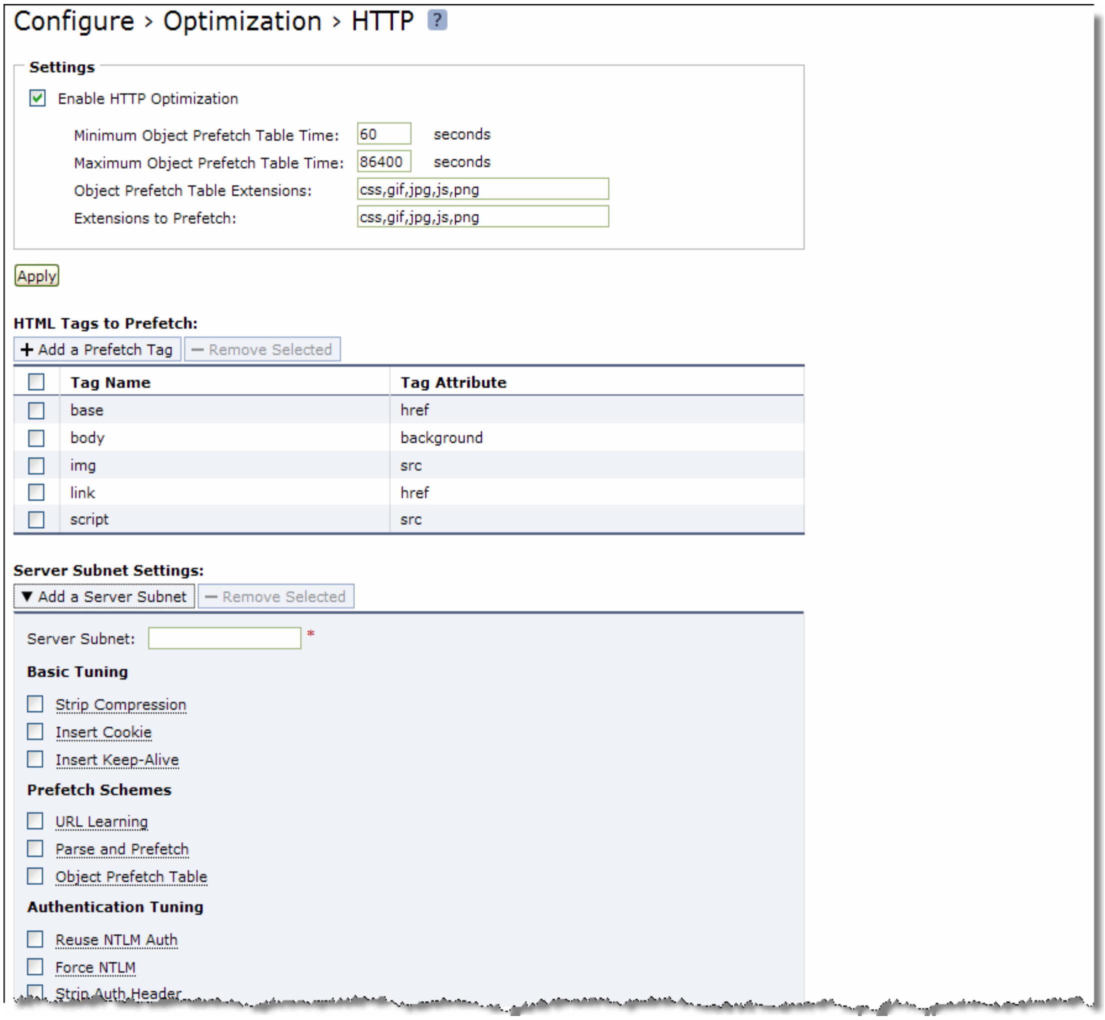

Configure > Optimization > HTTP ?

Settings

☑ Enable HTTP Optimization

Minimum Object Prefetch Table Time:	60	seconds
Maximum Object Prefetch Table Time:	86400	seconds
Object Prefetch Table Extensions:	css,gif,jpg,js,png	
Extensions to Prefetch:	css,gif,jpg,js,png	

Apply

HTML Tags to Prefetch:

+ Add a Prefetch Tag — Remove Selected

□	Tag Name	Tag Attribute
□	base	href
□	body	background
□	img	src
□	link	href
□	script	src

Server Subnet Settings:

▼ Add a Server Subnet — Remove Selected

Server Subnet: *

Basic Tuning

□ Strip Compression

□ Insert Cookie

□ Insert Keep-Alive

Prefetch Schemes

□ URL Learning

□ Parse and Prefetch

□ Object Prefetch Table

Authentication Tuning

□ Reuse NTLM Auth

□ Force NTLM

□ Strip Auth Header

2. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable HTTP Optimization	Enables HTTP acceleration, which prefetches and stores objects embedded in Web pages to improve HTTP traffic performance. By default, HTTP optimization is enabled.
Minimum Object Prefetch Table Time	Specify this option to set the minimum number of seconds the objects are stored in the local object prefetch table. The default is 60 seconds. This setting specifies the minimum lifetime of the stored object. During this lifetime, any qualified If-Modified-Since (IMS) request or regular request from the client receives an HTTP 304 response, indicating that the resource for the requested object has not changed since stored.

Control	Description
Maximum Object Prefetch Table Time	Specify this option to set the maximum number of seconds the objects are stored in the local object prefetch table. The default is 86,400 seconds. This setting specifies the maximum lifetime of the stored object. During this lifetime, any qualified If-Modified-Since (IMS) request or regular request from the client receives an HTTP 304 response, indicating that the resource for the requested object has not changed since stored.
Object Prefetch Table Extensions	Specify the object extensions to store, separated by commas. By default the Steelhead appliance stores .jpg, .gif, .js, .png, and .css object extensions. **Note:** These extensions are only for objects stored in the object prefetch table and do not affect other prefetch types.
Extensions to Prefetch	Specify object extensions to prefetch, separated by commas. By default the Steelhead appliance prefetches .jpg, .gif, .js, .png, and .css object extensions. **Note:** These extensions are only for URL Learning and do not affect other prefetch types.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

To prefetch HTML tags

1. Under HTML Tags to Prefetch, select which HTML tags to prefetch. By default, the following tags are prefetched: base/href, body/background, img/src, link/href, and script/src.

Note: These tags are for the Parse and Prefetch feature only and do not affect other prefetch types, such as object extensions.

2. To add a new tag, complete the configuration as described in the following table.

Control	Description
Add a Prefetch Tag	Displays the controls to add an HTML tag.
Tag Name	Specify the tag name.
Attribute	Specify the tag attribute.
Add	Adds the tag.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see “Managing Configuration Files” on page 358.

Adding a Server Subnet

Under Server Subnet Settings, you can enable URL Learning, Parse and Prefetch, and Object Prefetch Table in any combination for any server subnet. You can also enable authorization optimization in RiOS v6.1 and later to tune a particular subnet *dynamically,* with no service restart required.

The default setting is URL Learning only for all traffic. The default setting applies when HTTP optimization is enabled, regardless of whether there is an entry in the Server Subnet list. In the case of overlapping subnets, specific list entries override any default settings.

For example, suppose the majority of your Web servers have dynamic content applications but you also have several static content application servers. You could configure your entire server subnet to disable URL Learning and enable Parse and Prefetch and Object Prefetch Table, optimizing HTTP for the majority of your Web servers. Next, you could configure your static content servers to use URL Learning only, disabling Parse and Prefetch and Object Prefetch Table.

To add or configure a server subnet

1. On the client-side Steelhead appliance, under Server Subnet Settings, complete the configuration as described in the following table.

Control	Description
Add a Server Subnet	Displays the controls for adding a server subnet. The server must support keep-alive.
Server Subnet	Specify an IP address and mask pattern for the server subnet on which to set up the HTTP optimization scheme. Use the format: XXX.XXX.XXX.XXX/XX
Basic Tuning	
Strip Compression	Removes the accept-encoding lines from the HTTP compression header. An accept-encoding directive compresses content rather than using raw HTML. Enabling this option improves the performance of the Steelhead appliance data reduction algorithms. By default, strip compression is enabled.
Insert Cookie	Adds a cookie to HTTP applications that do not already have one. HTTP applications frequently use cookies to keep track of sessions. The Steelhead appliance uses cookies to distinguish one user session from another. If an HTTP application does not use cookies, the client Steelhead appliance inserts one so that it can track requests from the same client. By default, this setting is disabled.
Insert Keep Alive	Uses the same TCP connection to send and receive multiple HTTP requests and responses, as opposed to opening a new one for every single request and response. Specify this option when using the URL Learning or Parse and Prefetch features with HTTP v1.0 or HTTP v1.1 applications using the Connection Close method. By default, this setting is disabled.
Prefetch Schemes	
URL Learning	Enables URL Learning, which learns associations between a base URL request and a follow-on request. Stores information about which URLs have been requested and which URLs have generated a 200 OK response from the server. This option fetches the URLs embedded in style sheets or any JavaScript associated with the base page and located on the same host as the base URL. URL Learning works best with non-dynamic content that does not contain session-specific information. URL Learning is enabled by default. Your system must support cookies and persistent connections to benefit from URL Learning. If your system has cookies turned off and depends on URL rewriting for HTTP state management, or is using HTTP v1.0 (with no keep-alives), you can force the use of cookies using the Add Cookie option and force the use of persistent connections using the Insert Keep Alive option.
Parse and Prefetch	Enables Parse and Prefetch, which parses the base HTML page received from the server and prefetches any embedded objects to the client-side Steelhead appliance. This option complements URL Learning by handling dynamically generated pages and URLs that include state information. When the browser requests an embedded object, the Steelhead appliance serves the request from the prefetched results, eliminating the round-trip delay to the server. The prefetched objects contained in the base HTML page can be images, style sheets, or any Java scripts associated with the base page and located on the same host as the base URL. Parse and Prefetch requires cookies. If the application does not use cookies, you can insert one using the Insert Cookie option.

Control	Description
Object Prefetch Table	Enables the Object Prefetch Table, which stores HTTP object prefetches from HTTP GET requests for cascading style sheets, static images, and Java scripts in the Object Prefetch Table. When the browser performs If-Modified-Since (IMS) checks for cached content or sends regular HTTP requests, the client-side Steelhead appliance responds to these IMS checks and HTTP requests, cutting back on round trips across the WAN.
Authentication Tuning	
Reuse Auth	Allows an unauthenticated connection to serve prefetched objects, as long as the connection belongs to a session whose base connection is already authenticated. This option is most effective when the Web server is configured to use per-connection NTLM or Kerberos authentication.
Force NTLM	In the case of negotiated Kerberos and NTLM authentication, forces NTLM. Kerberos is less efficient over the WAN because the client must contact the Domain Controller to answer the server authentication challenge and tends to be employed on a per-request basis. Riverbed recommends enabling Strip Auth Header along with this option.
Strip Auth Header	Removes all credentials from the request on an already authenticated connection. This works around Internet Explorer behavior that re-authorizes connections that have previously been authorized. This option is most effective when the Web server is configured to use per-connection NTLM authentication. **Important:** If the Web server is configured to use per-request NTLM authentication, enabling this option might cause authentication failure.
Gratuitous 401	Prevents a WAN round trip by issuing the first 401 containing the realm choices from the client-side Steelhead appliance. Riverbed recommends enabling Strip Auth Header along with this option. This option is most effective when the Web server is configured to use per-connection NTLM authentication or per-request Kerberos authentication. **Important:** If the Web server is configured to use per-connection Kerberos authentication, enabling this option might cause additional delay.
Add	Adds the subnet.

Note: To modify subnet configuration properties, use the drop-down lists in the table row for the configuration.

2. Click **Apply** to apply your settings to the running configuration.

3. Click **Save** to save your settings permanently.

Tip: To modify server properties, use the drop-down list in the table row for the server.

Related Topic

- "Viewing CIFS Prepopulation Share Log Reports" on page 427

Configuring Oracle Forms Optimization

You can display and modify Oracle Forms optimization settings in the Configure > Optimization > Oracle Forms page.

Oracle Forms is a platform for developing user interface applications to interact with an Oracle database. It uses a Java applet to interact with the database in either native, HTTP, or HTTPS mode. The Steelhead appliance decrypts, optimizes, and then re-encrypts the Oracle Forms traffic.

You can configure Oracle Forms optimization in the following modes:

- **Native** - The Java applet communicates with the backend server, typically over port 9000. Native mode is also known as socket mode.
- **HTTP** - The Java applet tunnels the traffic to the Oracle Forms server over HTTP, typically over port 8000.
- **HTTPS** - The Java applet tunnels the traffic to the Oracle Forms server over HTTPS, typically over port 443. HTTPS mode is also known as SSL mode.

Use Oracle Forms optimization to improve Oracle Forms traffic performance. RiOS v5.5.x and later supports 6i, which comes with Oracle Applications 11i. RiOS v6.0 and later supports 10gR2, which comes with Oracle E-Business Suite R12.

This feature does not need a separate license and is enabled by default. However, you must also set an in-path rule to enable this feature.

Note: Optionally, you can enable IPSec encryption to protect Oracle Forms traffic between two Steelhead appliances over the WAN or use the Secure Inner Channel on all traffic.

Determining the Deployment Mode

Before enabling Oracle Forms optimization, you need to know the mode in which Oracle Forms is running at your organization.

To determine the Oracle Forms deployment mode

1. Start the Oracle application that uses Oracle Forms.
2. Click a link in the base HTML page to download the Java applet to your browser.
3. On the Windows taskbar, right-click the Java icon (a coffee cup) to access the Java console.
4. Choose Show Console (JInitiator) or Open <*version*> Console (Sun JRE).

5. Locate the "connectMode=" message in the Java Console window. This message indicates the Oracle Forms deployment mode at your organization, for example:

```
connectMode=HTTP, native
connectMode=Socket
connectMode=HTTPS, native
```

Enabling Oracle Forms Optimization

This section describes how to enable Oracle Forms optimization for the deployment mode your organization uses.

To enable the Oracle Forms optimization feature in native and HTTP modes

1. Choose Configure > Optimization > Oracle Forms to display the Oracle Forms page.

Figure 4-20. Oracle Forms Page

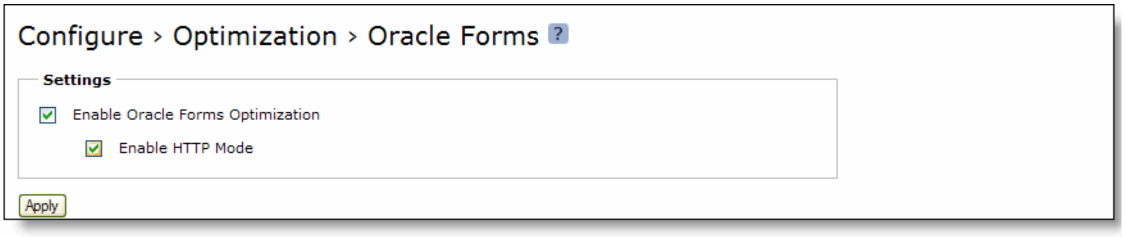

Configure > Optimization > Oracle Forms

Settings

☑ Enable Oracle Forms Optimization

☑ Enable HTTP Mode

2. On the client-side and server-side Steelhead appliances, under Settings, complete the configuration as described in the following table.

Control	Description
Enable Oracle Forms Optimization	Enables Oracle Forms optimization in native mode, also known as socket mode. Oracle Forms native mode optimization is enabled by default. Disable this option only to turn off Oracle Forms optimization; for example, if your network users do not use Oracle applications.
Enable HTTP Mode	Enables Oracle Forms optimization in HTTP mode. All internal messaging between the forms server and the Java client is encapsulated in HTTP packets. In RiOS v6.0 and later, HTTP mode is enabled by default. You must also click the Enable Oracle Forms Optimization check box to enable HTTP mode.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. If you change the Oracle Forms setting, you must restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

6. If you have not already done so, Choose Configure > Optimization > In-path Rules and click **Add a New In-path Rule**. Add an in-path rule with the following properties.

Property	Value
Type	Auto-discover or Fixed-target.
Destination Subnet/Port	Specify the server IP address (for example, 10.11.41.14/32), and a port number: • **9000** - Native mode, using the default forms server. • **8000** - HTTP mode.
Preoptimization Policy	Oracle Forms.
Optimization Policy	Normal.
Latency Optimization Policy	**HTTP** - Select this policy to separate any non-Oracle Forms HTTP traffic from the standard Oracle Forms traffic. This policy applies HTTP latency optimization to the HTTP traffic to improve performance. Both the client-side and server-side Steelhead appliances must be running RiOS v6.0 or later.
Neural Framing Mode	Always.
WAN Visibility	Correct Addressing.

To enable the Oracle Forms optimization feature in HTTPS mode

1. Configure and enable SSL optimization *before* enabling the Oracle Forms support. For details, see "Configuring SSL Server Certificates and Certificate Authorities" on page 213.

2. Choose Configure > Optimization > Oracle Forms to display the Oracle Forms page.

Figure 4-21. Oracle Forms Page

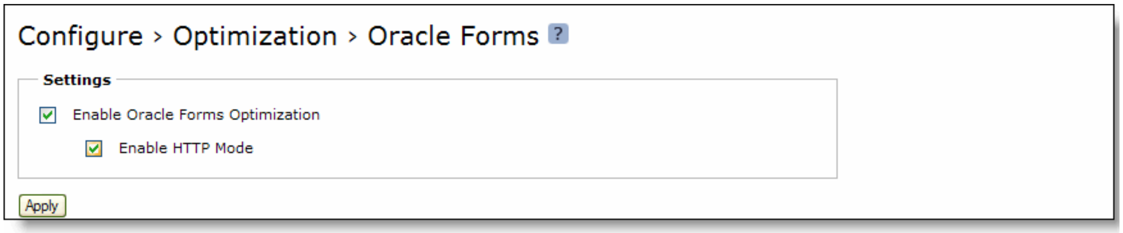

Configure > Optimization > Oracle Forms ?

Settings

☑ Enable Oracle Forms Optimization

☑ Enable HTTP Mode

3. Under Settings, select both check boxes as described in the following table.

Control	Description
Enable Oracle Forms Optimization	Enables Oracle Forms optimization in native mode, also known as socket mode. Oracle Forms native mode optimization is enabled by default. Disable this option only to turn off Oracle Forms optimization; for example, if your network users do not use Oracle applications.
Enable HTTP Mode	Enables Oracle Forms optimization in HTTP mode. All internal messaging between the forms server and the Java client is encapsulated in HTTP packets. In RiOS v6.0 and later, HTTP mode is enabled by default. You must also click the Enable Oracle Forms Optimization check box to enable HTTP mode.

4. Click **Apply** to apply your settings to the running configuration.

5. Click **Save** to save your settings permanently.

6. If you change the Oracle Forms setting, you must restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

7. Choose Configure > Optimization > In-path Rules and click **Add a New In-path Rule**. Use the following in-path rule settings.

Property	Value
Type	Auto-discover or Fixed-target.
Destination Subnet/Port	Specify the server IP address (for example, 10.11.41.14/32), and a port number (for example, 443).
Preoptimization Policy	Oracle Forms over SSL. **Note:** If you upgrade a Steelhead appliance running RiOS v5.5.x or earlier to v6.0 and were using the Oracle Forms preoptimization policy with v5.5.x, you must change it to Oracle Forms over SSL for v6.0 and later.
Optimization Policy	Normal.
Latency Optimization Policy	**HTTP** - Select this policy to separate any non-Oracle Forms HTTP traffic from the standard Oracle Forms traffic. This policy applies HTTP latency optimization to the HTTP traffic to improve performance. Both the client-side and server-side Steelhead appliances must be running RiOS v6.0 or later.
Neural Framing Mode	Always.
WAN Visibility	Correct Addressing.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Configuring HTTP Optimization Feature Settings" on page 114
- "Configuring SSL Server Certificates and Certificate Authorities" on page 213

Configuring MAPI Optimization

You can display and modify MAPI optimization settings in the Configure > Optimization > MAPI page.

MAPI optimization does not require a separate license and is enabled by default.

RiOS v6.0 and later uses the Steelhead secure inner channel to ensure all MAPI traffic sent between the client-side and the server-side Steelhead appliances are secure.

You must enable MAPI optimization on all Steelhead appliances optimizing MAPI in your network, not just the client-side Steelhead appliance.

To configure MAPI optimization features

1. Set up secure peering between the client-side and server-side Steelhead appliances and enable inner channel SSL with secure protocols. For details, see "Configuring Secure Peers" on page 231.

2. Choose Configure > Optimization > MAPI to display the MAPI page.

Figure 4-22. MAPI Page

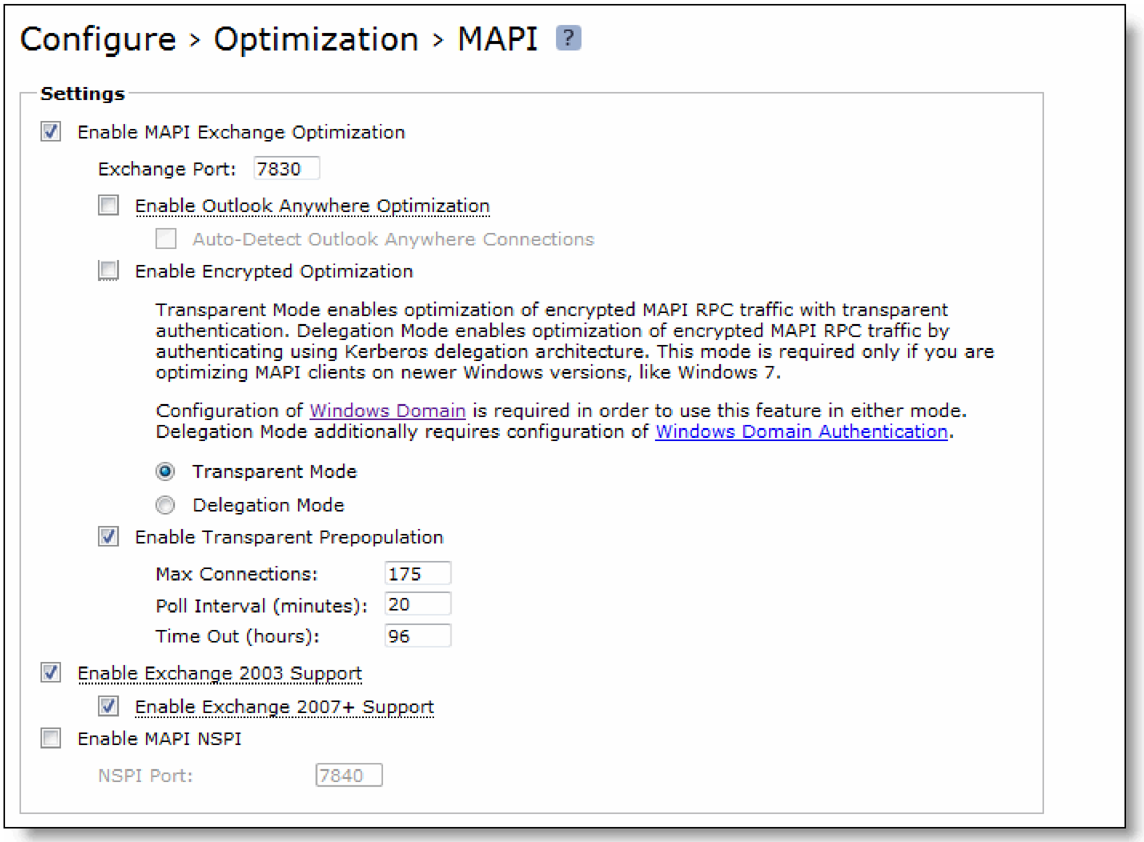

Configure > Optimization > MAPI

Settings

☑ Enable MAPI Exchange Optimization

Exchange Port: 7830

□ Enable Outlook Anywhere Optimization

□ Auto-Detect Outlook Anywhere Connections

□ Enable Encrypted Optimization

Transparent Mode enables optimization of encrypted MAPI RPC traffic with transparent authentication. Delegation Mode enables optimization of encrypted MAPI RPC traffic by authenticating using Kerberos delegation architecture. This mode is required only if you are optimizing MAPI clients on newer Windows versions, like Windows 7.

Configuration of Windows Domain is required in order to use this feature in either mode. Delegation Mode additionally requires configuration of Windows Domain Authentication.

● Transparent Mode

○ Delegation Mode

☑ Enable Transparent Prepopulation

Max Connections: 175

Poll Interval (minutes): 20

Time Out (hours): 96

☑ Enable Exchange 2003 Support

☑ Enable Exchange 2007+ Support

□ Enable MAPI NSPI

NSPI Port: 7840

3. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable MAPI Exchange Optimization	Enables the fundamental component of the MAPI optimization module, which includes optimization for Read, Write (Receive, Send), and Sync operations. By default, MAPI Exchange optimization is enabled. Only clear this check box to disable MAPI optimization. Typically, you disable MAPI optimization to troubleshoot problems with the system. For example, if you are experiencing problems with Outlook clients connecting with Exchange, you can disable MAPI latency acceleration (while continuing to optimize with SDR for MAPI).
Exchange Port	Specify the MAPI Exchange port for optimization. Typically, you do not need to modify the default value, 7830. If you have changed the MEISI port in your Exchange Server environment, change port 7830 to the static port number you have configured in your Exchange environment. For more details about changing (MEISI) ports, see the Microsoft Exchange Information Store Interface at: https://support.microsoft.com/kb/270836/en-us

Control	Description
Enable Outlook Anywhere Optimization	Enables Outlook Anywhere latency optimization. Outlook Anywhere is a feature of Microsoft Exchange Server 2003, 2007, and 2010 that allows Microsoft Office Outlook 2003, 2007, and 2010 clients to connect to their Exchange servers over the Internet using the Microsoft RPC tunneling protocol. Outlook Anywhere allows for a VPN-less connection as the MAPI RPC protocol is tunneled over HTTP or HTTPS. RPC over HTTP can transport regular or encrypted MAPI. If you use encrypted MAPI, the server-side Steelhead appliance must be a member of the Windows domain. Enable this feature on the client-side and server-side Steelheads. Both Steelheads must be running RiOS v6.5. By default, this feature is disabled. To use this feature, you must also enable HTTP Optimization on the client-side and server-side Steelheads (HTTP optimization is enabled by default). If you are using Outlook Anywhere over HTTPS, you must enable SSL and the IIS certificate must be installed on the server-side Steelhead: • When using HTTP, Outlook can only use NTLM proxy authentication. • When using HTTPS, Outlook can use NTLM or Basic proxy authentication. • When using encrypted MAPI with HTTP or HTTPS, you must enable and configure encrypted MAPI in addition to this feature. **Note:** Outlook Anywhere optimized connections cannot start MAPI prepopulation. After you apply your settings, you can verify that the connections appear in the Current Connections report as a MAPI-OA or an eMAPI-OA (encrypted MAPI) application. The Outlook Anywhere connection entries appear in the system log with an RPCH prefix. **Important:** Outlook Anywhere creates twice as many connections on the Steelhead than regular MAPI. This results in the Steelhead entering admission control twice as fast with Outlook Anywhere than with regular MAPI. For details, see Appendix B - Steelhead Appliance MIB. For details and troubleshooting information, see the *Riverbed Deployment Guide*. For details about enabling Outlook Anywhere, see http://technet.microsoft.com/en-us/library/bb123513(EXCHG.80).aspx

Control	Description
Auto-Detect Outlook Anywhere Connections	Automatically detects the RPC over HTTPS protocol used by Outlook Anywhere. This feature is dimmed and unavailable until you enable Outlook Anywhere optimization. You can enable automatic detection of RPC over HTTPS using this option or you can set in-path rules. Auto-detect is best for simple Steelhead configurations with only a single Steelhead at each site and when the IIS server is also handling Web sites. If the IIS server is only used as RPC Proxy, and for configurations with asymmetric routing, connection forwarding or Interceptor installations, add in-path rules that identify the RPC Proxy server IP addresses and select the Outlook Anywhere latency optimization policy. After adding the in-path rule, disable the auto-detect option. On an Interceptor, add load-balancing rules to direct traffic for RPC Proxy to the same Steelhead appliance. In-path rules interact with auto-detect as follows: • When auto-detect is enabled and the in-path rule does not match, RiOS optimizes Outlook Anywhere if it detects the RPC over HTTPS protocol. • When auto-detect is not enabled and the in-path rule does not match, RiOS does not optimize Outlook Anywhere. • When auto-detect is enabled and the in-path rule matches with HTTP only, RiOS does not optimize Outlook Anywhere (even if it detects the RPC over HTTPS protocol). • When auto-detect is not enabled and the in-path rule does not match with HTTP only, RiOS does not optimize Outlook Anywhere. • When auto-detect is enabled and the in-path rule matches with an Outlook Anywhere latency optimization policy, RiOS optimizes Outlook Anywhere (even if it does not detect the RPC over HTTPS protocol). • When auto-detect is not enabled and the in-path rule matches with Outlook Anywhere, RiOS optimizes Outlook Anywhere.
Enable Encrypted Optimization	Enables encrypted MAPI RPC traffic optimization between Outlook and Exchange. By default, this option is disabled. The basic steps to enable encrypted optimization are: **1.** Choose Configure > Networking > Windows Domain and join the server-side Steelhead appliance to the same Windows Domain that the Exchange server belongs to and operates as a member server. **2.** Verify that Outlook is encrypting traffic. **3.** Enable this option on all Steelheads involved in optimizing MAPI encrypted traffic. **4.** Windows 7 MAPI clients must use Delegation mode. Delegation mode is the default in RiOS v6.5. Use Transparent mode for all other clients. **5.** Make sure that both Enable MAPI Exchange 2003 Acceleration and Enable MAPI Exchange 2007 Acceleration are enabled. Both options are enabled by default. **6.** Restart the service on all Steelheads that have this option enabled. **Note:** Both the server-side and client-side Steelheads must be running RiOS v5.5.x or later. **Note:** When this option is enabled and Enable MAPI Exchange 2007 Acceleration is disabled on either Steelhead appliance, MAPI Exchange 2007 acceleration remains in effect for unencrypted connections.

Control	Description
Transparent Mode	Provides encrypted MAPI with transparent NTLM authentication. By default, this setting is enabled with encrypted MAPI optimization. Transparent mode supports all Windows servers, including Windows 2008 R2 (assuming they are not in domains with NTLM disabled). Transparent mode does *not* support Windows 7 clients or Windows 2008 R2 domains with NTLM disabled. Windows 7 clients must use Delegation mode. In RiOS v6.1 and later, transparent mode includes support for trusted domains, wherein users are joined to a different domain from the Exchange server being accessed.
Delegation Mode	Provides encrypted MAPI optimization using the Kerberos delegation facility. Select this mode if you are encrypting MAPI traffic for Windows 7 or earlier client versions. Both the server-side and client-side Steelhead appliances must be running RiOS v6.1 or later. **Note:** CIFS SMB Signing and Encrypted MAPI optimization share the delegate user account. If you enable Delegation mode for both features, the delegate user account must have delegation privileges for both features as well. If you are upgrading from RiOS v6.0, a delegation account might already be in place for CIFS SMB Signing. In RiOS v6.1 and later, Delegation mode includes support for trusted domains, wherein users are joined to a different domain from the filer being accessed. Delegation mode requires additional configuration. To configure Delegation mode, choose Configure > Optimization > Windows Domain Authentication.

Control	Description
Enable Transparent Prepopulation	Enables MAPI transparent prepopulation. You must enable this feature on the server-side and client-side Steelhead appliance. By default, MAPI transparent prepopulation is enabled. Transparent prepopulation provides a mechanism for sustaining Microsoft Exchange MAPI connections between the client and server even after the Outlook client has shut down. This allows mail data to be delivered between the Exchange server and the client-side Steelhead appliance while the Outlook client is offline or inactive. When a user logs into their Outlook client, the mail data is already prepopulated on the client-side Steelhead appliance. This accelerates the first access of the client's e-mail. Transparent prepopulation creates virtual MAPI connections to the Exchange server for Outlook clients that are offline. When the remote Steelhead appliance detects that an Outlook client has shut down, the virtual MAPI connections are triggered. The remote Steelhead appliance uses these virtual connections to pull mail data from the Exchange server over the WAN link. Enable this feature to allow email data to be delivered between the Exchange server and the client-side Steelhead appliance while the Outlook client is offline. When a user logs in to their MAPI client, the mail has already been seen by the client-side Steelhead appliance and is retrieved with LAN-like performance. MAPI prepopulation does not use any additional Client Access Licenses (CALs). The Steelhead appliance holds open a existing authenticated MAPI connection after Outlook is shut down. No user credentials are used or saved by the Steelhead appliance when performing prepopulation. In RiOS v6.5, the client-side Steelhead appliance controls MAPI prepopulation v2. This allows for a higher rate of prepopulated session, and enables the MAPI prepopulation to take advantage of the read-ahead feature in the MAPI optimization blade. MAPI prepopulation v2 is supported in RiOS v6.0.4 or later, v6.1.2 or later, and v6.5. The client-side and server-side Steelhead appliance can be running any of these code train levels and provide prepopulation v2 capabilities. For example, a client-side Steelhead appliance running RiOS v6.0.4 connecting to a server-side Steelhead appliance running RiOS v6.5 provides prepopulation v2 capabilities. In contrast, a 6.0.1a client-side Steelhead appliance connecting to a RiOS v6.5 server-side Steelhead appliance supports prepopulation v1, but does not provide prepopulation v2. If a user starts a new Outlook session, the MAPI prepopulation session terminates. If for some reason the MAPI prepopulation session does not terminate (for example, the user starts a new session in a location that is different than the Steelhead appliance that has the MAPI prepopulation session active), the MAPI prepopulation session eventually times-out per the configuration setting. **Note:** MAPI transparent prepopulation is not started with Outlook Anywhere connections.
Max Connections	Specify the maximum number of virtual MAPI connections to the Exchange server for Outlook clients that have shut down. Setting the maximum connections limits the aggregate load on all Exchange servers through the configured Steelhead appliance. The default value varies by model; for example, on a 5520 the default is 3750. You must configure the maximum connections on both the client and server-side of the network.
Poll Interval (minutes)	Sets the number of minutes you want the appliance to check the Exchange server for newly-arrived email for each of its virtual connections. The default value is 20.
Time Out (hours)	Specify the number of hours after which to time-out virtual MAPI connections. When this threshold is reached, the virtual MAPI connection is terminated. The time-out is enforced on a per-connection basis. Time-out prevents a buildup of stale or unused virtual connections over time. The default value is 96.

Control	Description
Enable Exchange 2003 Support	Enables MAPI 2003 support. By default, this option is enabled. This feature increases optimization of traffic between Exchange 2003 and Outlook 2003. Do not disable when moving to a later version of MAPI in your network; for example, if you are running Exchange 2007 with Outlook 2007 clients, do not disable the Exchange 2003 option. You must enable MAPI Exchange Optimization to optimize outbound traffic from Outlook 2003 to the Exchange server. Regardless of the MAPI Exchange Optimization setting, you must configure this option identically on the client-side and server-side Steelheads. For out-of-path deployments, to optimize MAPI Exchange 2003, you must define fixed-target, in-path rules that specify the following ports on the client-side Steelhead appliance: the Microsoft end-point mapper port: 135; the Steelhead appliance port for Exchange traffic: 7830; the Steelhead appliance port for Exchange Directory NSPI traffic: 7840.
Enable Exchange 2007+ Support	Enables native MAPI 2007 support. By default, this option is enabled. If you have Outlook 2007 and Exchange 2003 or 2007 in your environment, this option increases optimization of traffic between Exchange and Outlook 2007. Sharing calendars between Outlook 2007 and Exchange 2007 increases the number of connections (anywhere from 1 to 2 extra connections per each user sharing calendars). The connections are persistent and remain even when users are not actively checking other user's calendars. Enabling this option helps keep connection counts at sustained, low levels, thereby increasing optimization. You must enable MAPI Exchange Optimization to optimize outbound traffic from Outlook 2007 and later to the Exchange server. Regardless of the MAPI Exchange Optimization setting, you must configure this option identically on the client-side and server-side Steelheads.
Enable MAPI NSPI	Performs latency optimization for MAPI Name Service Provider Interface (NSPI) connections when using the Exchange 2000 Server or when the client is not using Cached Exchange mode.By default, NSPI optimization is disabled. NSPI is the address book subcomponent of the Exchange protocol. Enable this feature to perform latency optimization for the connection when using the Exchange 2000 Server or when the client is not using Cached Exchange mode.
NSPI Port	Specify the NSPI port. The default value is 7840.

7. Click **Apply** to apply your settings to the running configuration.

8. Click **Save** to save your settings permanently.

Tip: When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Optimizing MAPI Exchange in Out-of-Path Deployments

In out-of-path deployments, if you want to optimize MAPI Exchange by destination port, you must define a fixed-target, in-path rule that specifies the following ports on the client-side appliance:

- **Port 135** - The Microsoft end-point mapper port.
- **Port 7830** - The Steelhead appliance port used for Exchange traffic.
- **Port 7840** - The Steelhead appliance port used for Exchange Directory NSPI traffic.

For details on defining in-path rules, see "Configuring In-Path Rules" on page 27.

Configuring MS-SQL Optimization

You can display and modify MS-SQL optimization settings in the Configure > Optimization > MS-SQL page.

Enabling MS-SQL optimization applies default rules to increase optimization for Microsoft Project (MS Project).

By default, Riverbed provides MS-SQL optimizations only for Microsoft Project Enterprise 2003. Each application interacts with the database differently and customizations are needed before the MS-SQL feature can be used for any other application. To optimize all other SQL applications with the MS-SQL Application acceleration module, contact Riverbed Professional Services.

To configure MS-SQL optimization features

1. Choose Configure > Optimization > MS-SQL to display the MS-SQL page.

Figure 4-23. MS-SQL Page

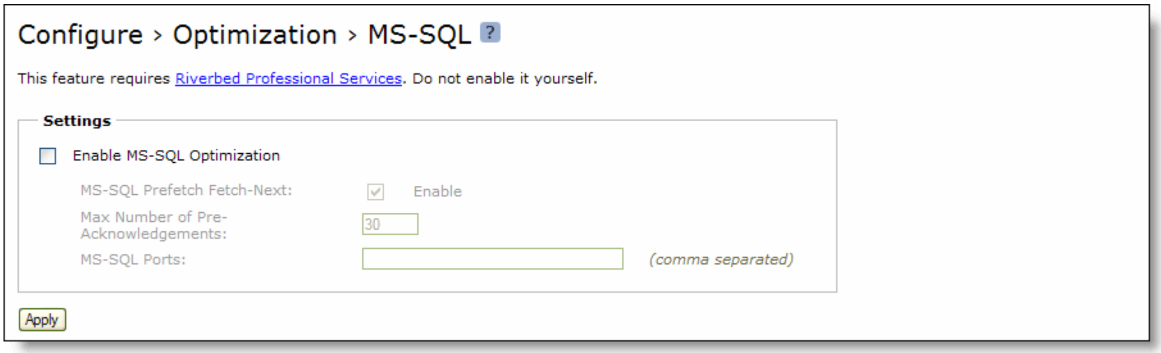

Configure > Optimization > MS-SQL

This feature requires Riverbed Professional Services. Do not enable it yourself.

Settings

□ Enable MS-SQL Optimization

MS-SQL Prefetch Fetch-Next: ☑ Enable

Max Number of Pre-Acknowledgements: 30

MS-SQL Ports: (comma separated)

2. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable MS-SQL Optimization	Increases optimization for Microsoft Project. The MS-SQL feature also optimizes other database applications, but you must define SQL rules to obtain maximum optimization. If you are interested in enabling the MS-SQL feature for other database applications, contact Riverbed Professional Services.
MS-SQL Prefetch Fetch-Next	Enables prefetching requests to request the next row in MS Project. This feature is enabled by default. The server-side Steelhead appliance prefetches sequential row results and the client-side Steelhead appliance caches them.
Max Number of Pre-Acknowledgements	Specify the number of requests to pre-acknowledge before waiting for a server response to be returned. The default value is 30.
MS-SQL Ports	Specify a comma-separated list of port numbers for MS-SQL servers. By default, 1433 is optimized; if you specify other ports they are optimized instead.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

Tip: After you apply your settings, you can verify whether changes have had the desired effect by reviewing related reports. When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Configuring NFS Optimization

You can display and modify NFS optimization settings in the Configure > Optimization > NFS page.

NFS optimization provides latency optimization improvements for NFS operations by prefetching data, storing it on the client Steelhead appliance for a short amount of time, and using it to respond to client requests. You enable NFS optimization in high-latency environments.

You can configure NFS settings globally for all servers and volumes or you can configure NFS settings that are specific to particular servers or volumes. When you configure NFS settings for a server, the settings are applied to all volumes on that server unless you override settings for specific volumes.

Important: NFS optimization is not supported in an out-of-path deployment.

Note: NFS optimization is only supported for NFS v3. When a transaction using NFS version 2 or 4 is optimized, the NFS latency module cannot be used and an alarm is triggered. Bandwidth optimization, SDR and LZ compression will still apply.

To configure NFS optimization

1. Choose Configure > Optimization > NFS to display the NFS page.

Figure 4-24. NFS Page

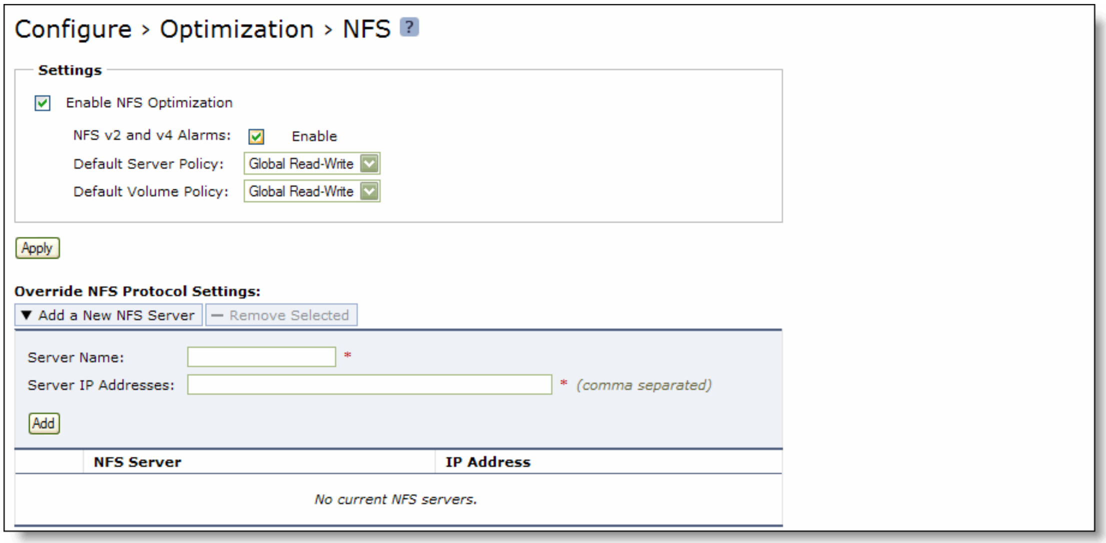

Configure > Optimization > NFS

Settings

☑ Enable NFS Optimization

NFS v2 and v4 Alarms: ☑ Enable

Default Server Policy: Global Read-Write

Default Volume Policy: Global Read-Write

Override NFS Protocol Settings:

Server Name: *

Server IP Addresses: * (comma separated)

	NFS Server	IP Address
	No current NFS servers.	

2. Under Settings, complete the configuration as described in the following table.

Control	Description
Enable NFS Optimization	Enables NFS optimization. You enable NFS optimization where NFS performance over the WAN is impacted by a high-latency environment. By default, this feature is enabled.
NFS v2 and v4 Alarms	Enables alarm notification when NFS v2 and NFS v4 traffic is detected. When triggered, the Steelhead appliance displays the Needs Attention health state. The alarm provides a link to this page and a button to reset the alarm.
Default Server Policy	Select one of the following server policies for NFS servers: • **Global Read-Write** - Specifies a policy that provides data consistency rather than performance. All of the data can be accessed from any client, including LAN-based NFS clients (which do not go through the Steelhead appliances) and clients using other file protocols such as CIFS. This option severely restricts the optimization that can be applied without introducing consistency problems. This is the default configuration. • **Custom** - Specifies a custom policy for the NFS server. • **Read-only** - Specifies that the clients can read the data from the NFS server or volume but cannot make changes. The default server policy is used to configure any connection to a server which does not have a policy.
Default Volume Policy	Select one of the following volume policies for NFS volumes: • **Global Read-Write** - Specifies a policy that provides data consistency rather than performance. All of the data can be accessed from any client, including LAN-based NFS clients (which do not go through the Steelhead appliances) and clients using other file protocols such as CIFS. This option severely restricts the optimization that can be applied without introducing consistency problems. This is the default configuration. • **Custom** - Specifies a custom policy for the NFS volume. • **Read-only** - Specifies that the clients can read the data from the NFS server or volume but cannot make changes. The default volume policy is used to configure a volume that does not have a policy.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

You can add server configurations to override your default settings. You can also modify or remove these configuration overrides. If you do not override settings for a server or volume, the Steelhead appliance uses the global NFS settings.

To override NFS settings for a server or volume

1. Choose Configure > Optimization > NFS to display the NFS page.

Figure 4-25. Partial NFS Page

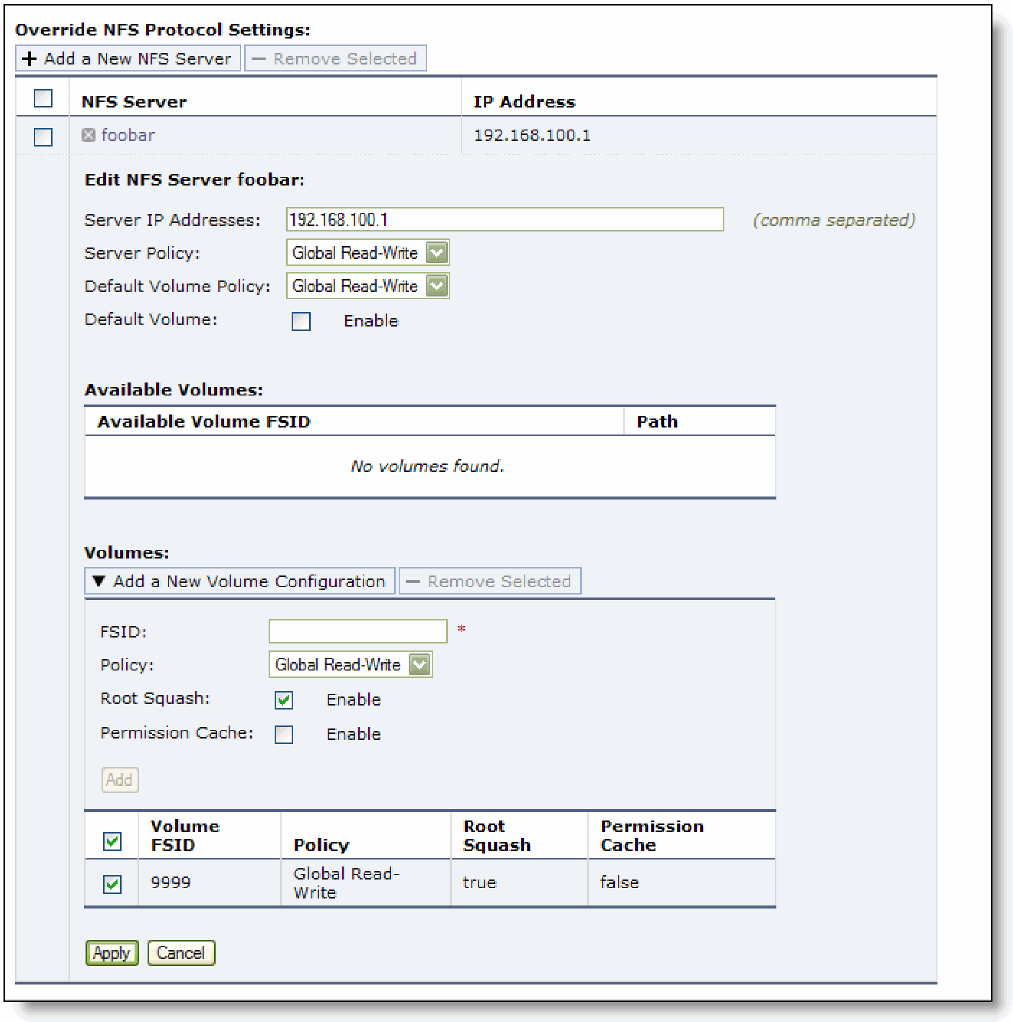

Override NFS Protocol Settings:

+ Add a New NFS Server — Remove Selected

□	NFS Server	IP Address
□	☒ foobar	192.168.100.1

Edit NFS Server foobar:

Server IP Addresses:	192.168.100.1	(comma separated)
Server Policy:	Global Read-Write	
Default Volume Policy:	Global Read-Write	
Default Volume:	□ Enable	

Available Volumes:

Available Volume FSID	Path
No volumes found.	

Volumes:

▼ Add a New Volume Configuration — Remove Selected

FSID:	*
Policy:	Global Read-Write
Root Squash:	☑ Enable
Permission Cache:	□ Enable

Add

☑	Volume FSID	Policy	Root Squash	Permission Cache
☑	9999	Global Read-Write	true	false

Apply Cancel

2. Under Override NFS Protocol Settings, complete the configuration as described in the following table.

Control	Description
Add a New NFS Server	Displays the controls to add an NFS server configuration.
Server Name	Specify the name of the server.
Server IP Addresses	Specify the IP addresses of the servers, separated by commas, and click **Add Server**.
Add	Adds the configuration to the NFS Servers list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

To modify the properties for an NFS server

1. Select the NFS server name in the table and complete the configuration as described in the following table.

Control	Description
Server IP Addresses	Specify the server IP addresses, separated by commas.
Server Policy	Select one of the following server policies for this NFS server configuration from the drop-down list: • **Global Read-Write** - Choose this policy when the data on the NFS server can be accessed from any client, including LAN clients and clients using other file protocols. This policy ensures data consistency but does not allow for the most aggressive data optimization. This is the default value. • **Custom** - Create a custom policy for the NFS server. • **Read-only** - Any client can read the data on the NFS server or volume but cannot make changes.
Default Volume Policy	Select one of the following default volume configurations for this server from the drop-down list: • **Global Read-Write** - Choose this policy when the data on the NFS volume can be accessed from any client, including LAN clients and clients using other file protocols. This policy ensures data consistency but does not allow for the most aggressive data optimization. This is the default value. • **Custom** - Create a custom policy for the NFS server. • **Read-only** - Any client can read the data on the NFS server or volume but cannot make changes.
Default Volume	Enables the default volume configuration for this server.
Apply	Applies the changes.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Save** to save your settings permanently.

After you add a server, the NFS page includes options to configure volume policies. The Available Volumes table provides an uneditable list of NFS volumes that are available for the current NFS server. You can use the NFS volume information listed in this table to facilitate adding new NFS volumes.

To add an NFS volume configuration for this server

1. Select the NFS server name in the table and complete the configuration as described in the following table.

Control	Description
Add a New Volume Configuration	Displays the controls to add a new volume.
FSID	Specify the volume File System ID. An FSID is a number NFS uses to distinguish mount points on the same physical file system. Because two mount points on the same physical file system have the same FSID, more than one volume can have the same FSID.
Policy	Optionally, choose one of the following default volume configurations for this server from the drop-down list: • **Global Read-Write** - Choose this policy when the data on the NFS volume can be accessed from any client, including LAN clients and clients using other file protocols. This policy ensures data consistency but does not allow for the most aggressive data optimization. This is the default value. • **Custom** - Create a custom policy for the NFS server. • **Read-only** - Any client can read the data on the NFS server or volume but cannot make changes.
Root Squash	Enables the root squash feature for NFS volumes from this server. This feature turns off Steelhead optimizations for the root user on NFS clients. When the root user accesses an NFS share, its ID is *squashed* (mapped) to another user (most commonly "nobody") on the server. This is for security reasons, as it prevents clients from giving themselves access to the server file system.
Permission Cache	Enables the permission cache, where the Steelhead appliance stores file read data and uses it to respond to client requests. For example, if a user downloads data and another user tries to access that data, the Steelhead ensures that the second user has permission to read the data before releasing it.
Default Volume	Enables the default volume configuration for this server.
Add	Adds the volume.
Remove Selected	Select the check box next to the volume FSID and click **Remove Selected**.

2. Click **Save** to save your settings permanently.

To reset the NFS alarm

1. Choose Configure > Optimization > NFS to display the NFS page. The option to reset the NFS alarm appears only after the service triggers the NFS v2 and v4 alarm. The alarm remains triggered until you manually reset it.

2. Under Reset NFS Alarm, click **Reset NFS Alarm**.

3. Click **Save** to save your settings permanently.

Related Topic

- "Viewing NFS Reports" on page 432

Configuring Lotus Notes Optimization

You can enable and modify Lotus Notes optimization settings in the Configure > Optimization > Lotus Notes page.

Lotus Notes is a client-server collaborative application that provides email, instant messaging, calendar, resource, and file sharing. RiOS provides latency and bandwidth optimization for Lotus Notes v6.0 and later traffic across the WAN, accelerating email attachment transfers and server-to-server or client-to-server replications.

RiOS saves bandwidth by automatically disabling socket compression (which makes SDR more effective), and by decompressing Huffman-compressed attachments and LZ-compressed attachments when they are sent or received and recompressing them on the other side. This allows SDR to recognize attachments which have previously been sent in other ways, that is; over CIFS, HTTP, or other protocols, and also allows SDR to optimize the sending and receiving of attachments that are slightly changed from previous sends and receives.

To use this feature both the client-side and server-side Steelhead appliances must be running RiOS v5.5.x or later.

Enabling Lotus Notes provides latency optimization regardless of the compression type (Huffman, LZ, or none).

Before enabling Lotus Notes optimization, be aware that:

- Riverbed cannot optimize encrypted Lotus Notes connections.
- Lotus Notes Optimization automatically disables socket level compression for connections going through Steelheads that have this feature enabled.

To configure Lotus Notes optimization

1. Choose Configure > Optimization > Lotus Notes to display the Lotus Notes page.

Figure 4-26. Lotus Notes Page

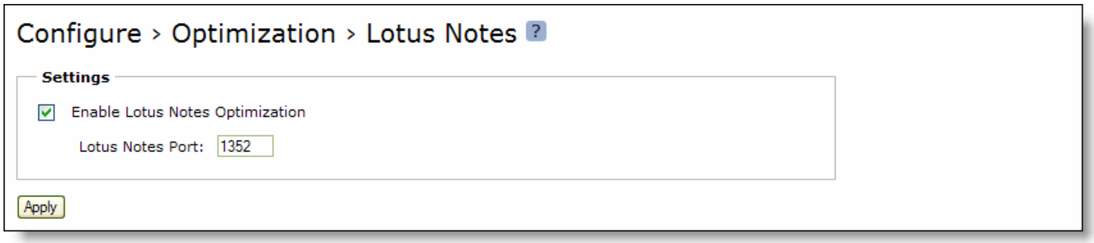

Configure > Optimization > Lotus Notes

Settings

☑ Enable Lotus Notes Optimization

Lotus Notes Port: 1352

2. Under Settings, complete the configuration on the client-side Steelhead appliance as described in the following table.

Control	Description
Enable Lotus Notes Optimization	Provides latency and bandwidth optimization for Lotus Notes v6.0 and later traffic across the WAN, accelerating email attachment transfers and server-to-server or client-to-server replications. By default, Lotus Notes optimization is disabled.
Lotus Notes Port	Specify the Lotus Notes port for optimization. Typically, you do not need to modify the default value 1352.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. If you have enabled or disabled Lotus Notes or changed the port, you need to restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Configuring Citrix ICA Optimization

You can enable and modify Citrix ICA optimization settings in the Configure > Optimization > Citrix ICA page.

To consolidate operations, some organizations install thin clients in their branch offices and install a Citrix Presentation Server in the data center to front-end the applications. The proprietary protocol that Citrix uses to move updates between the client and the server is called ICA (Independent Computing Architecture). The thin clients at the branch offices have a Citrix ICA client accessing the services at the data center which are front-ended by a Citrix Presentation Server (also called Citrix Metaframe Server in earlier versions).

RiOS v6.0 and later provides the following ways to recognize, prioritize, encrypt and optimize Citrix traffic:

- Optimize the native ICA traffic bandwidth.
- Classify and shape Citrix ICA traffic using QoS.

Citrix Version Support

RiOS v6.0 and later provides support for the following Citrix software versions on the client side.

Citrix software running on an ICA Client or Receiver:

- Version 9 (starting in RiOS v6.0.4 and v6.1.2)
- Version 10 (RiOS v6.0.0 and later)
- Version 11 (RiOS v6.0.0 and later)
- Version 12 (RiOS v6.1.2a and later)
- Wyse V10L and S10 Thin clients (RiOS v6.0.2 and later)

On XenDesktop:

- XenDesktop version 4 (RiOS v6.1.2a and later)
- XenDesktop version 5 (RiOS v6.1.2a and later)

RiOS v6.0 and later provides support for the following Citrix software versions on the server side:

On XenApp:

- RiOS can automatically negotiate session encryption and compression for basic and secure ICA, and can create QoS classes from Citrix virtual channels. Currently, RiOS does not provide the capacity to add pass-through Citrix traffic into the ICA channel or support latency optimization for ICA over SSL and optimization for client drive mapping (RiOS does provide data reduction).
- Presentation Server version 4.5 (RiOS v6.0.0 and later)
- XenApp Server version 5.0 (RiOS v6.0.0 and later)
- XenApp Server version 6.0 (RiOS 6.1.2a and later)

Basic Steps

The following table describes the basic steps needed to configure Citrix Optimization, followed by detailed procedures.

Task	Reference
1. Remove ports 1494 and 2598 from the Interactive Ports label.	"Configuring Port Labels" on page 94
2. Enable Citrix optimization on the client-side and server-side Steelhead appliance.	"To configure Citrix ICA optimization" on page 138
3. Optionally, encrypt the ICA protocol.	"To configure Citrix ICA optimization" on page 138
4. Optionally, use a QoS rule to prioritize Citrix traffic.	"Citrix ICA QoS Default Rule" on page 274
5. If you have changed the ICA or Session Reliability port, restart the optimization service.	"Starting and Stopping the Optimization Service" on page 345

To configure Citrix ICA optimization

1. Choose Configure > Networking > Ports Labels to display the Ports Labels page.

2. Select the Interactive port label in the Port Labels list to display the Editing Port Labels Interactive group.

Figure 4-27. Editing Port Labels Page

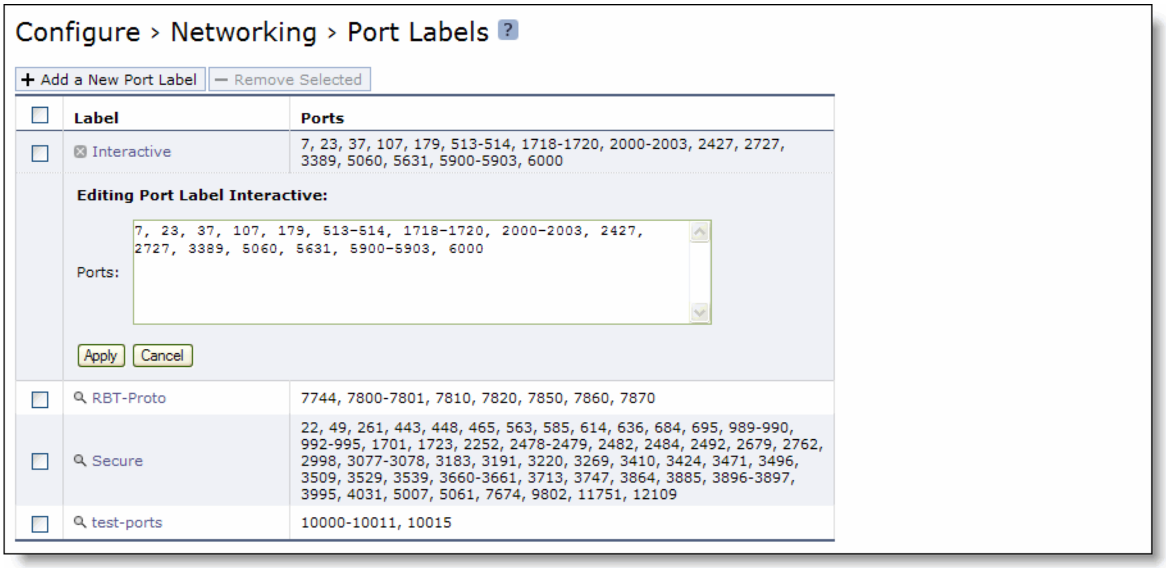

□	Label	Ports
□	☒ Interactive	7, 23, 37, 107, 179, 513-514, 1718-1720, 2000-2003, 2427, 2727, 3389, 5060, 5631, 5900-5903, 6000
□	🔍 RBT-Proto	7744, 7800-7801, 7810, 7820, 7850, 7860, 7870
□	🔍 Secure	22, 49, 261, 443, 448, 465, 563, 585, 614, 636, 684, 695, 989-990, 992-995, 1701, 1723, 2252, 2478-2479, 2482, 2484, 2492, 2679, 2762, 2998, 3077-3078, 3183, 3191, 3220, 3269, 3410, 3424, 3471, 3496, 3509, 3529, 3539, 3660-3661, 3713, 3747, 3864, 3885, 3896-3897, 3995, 4031, 5007, 5061, 7674, 9802, 11751, 12109
□	🔍 test-ports	10000-10011, 10015

Editing Port Label Interactive:

Ports: 7, 23, 37, 107, 179, 513-514, 1718-1720, 2000-2003, 2427, 2727, 3389, 5060, 5631, 5900-5903, 6000

3. Under Editing Port Label Interactive, remove Citrix ICA ports 1494 and 2598 from the Ports text box.

4. Click **Apply** to save your settings to the running configuration.

5. Choose Configure > Optimization > Citrix ICA to display the Citrix ICA page.

Figure 4-28. Citrix ICA Page

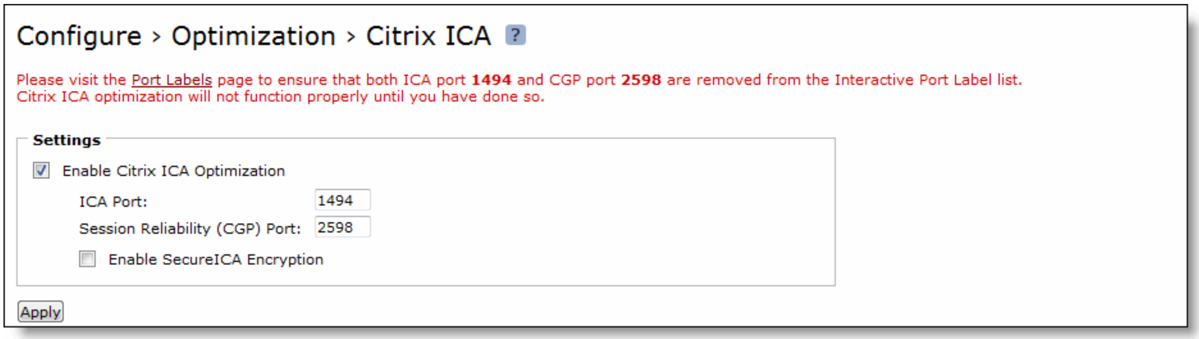

Configure > Optimization > Citrix ICA ?

Please visit the Port Labels page to ensure that both ICA port **1494** and CGP port **2598** are removed from the Interactive Port Label list.
Citrix ICA optimization will not function properly until you have done so.

Settings

☑ Enable Citrix ICA Optimization

ICA Port: 1494

Session Reliability (CGP) Port: 2598

□ Enable SecureICA Encryption

6. Under Settings, complete the configuration on the client-side and server-side Steelhead appliances as described in the following table.

Control	Description
Enable Citrix ICA Optimization	Optimizes the native Citrix traffic bandwidth. By default, Citrix optimization is disabled.
ICA Port	Specify the port on the Presentation Server for inbound traffic. The default port is 1494.
Session Reliability (CGP) Port	Specify the port number for Common Gateway Protocol (CGP) connections. CGP uses the session reliability port to keep the session window open even if there is an interruption on the network connection to the server. The default port is 2598.
Enable SecureICA Encryption	Uses the RC5 algorithm to encrypt the ICA protocol, securing communication sent between a MetaFrame Presentation Server and a client.

7. Click **Apply** to apply your settings to the running configuration.

8. Click **Save** to save your settings permanently.

9. If you have enabled or disabled Citrix optimization or changed the port, you need to restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Related Topics

- "Configuring Port Labels" on page 94
- "Creating QoS Classes" on page 288

Configuring FCIP Optimization

You can enable and modify FCIP storage optimization module settings in the Configure > Optimization > FCIP page.

Fibre Channel over TCP/IP (FCIP) is a transparent Fibre Channel (FC) tunneling protocol that transmits FC information between FC storage facilities over IP networks. FCIP is designed to overcome the distance limitations of FC.

RiOS v6.1 and later FCIP storage optimization provides support for environments using storage technology that originates traffic as FC and then uses either a Cisco Multilayer Director Switch (MDS) or a Brocade 7500 FCIP gateway.

To increase the data reduction LAN-to-WAN ratio with either equal or greater data throughput in environments with FCIP traffic, RiOS separates the FCIP headers from the application data workload written to storage. The FCIP headers contain changing protocol state information, such as sequence numbers. These headers interrupt the network stream and reduce the ability of SDR to match large, contiguous data patterns. After isolating the header data, the Steelhead appliance performs SDR network deduplication on the larger, uninterrupted storage data workload and LZ compression on the headers. RiOS then optimizes, reassembles, and delivers the data to the TCP consumer without compromising data integrity.

Note: Environments with Symmetrix Remote Data Facility (SRDF) traffic originated through Symmetrix FC ports (RF ports) only require configuration of the RiOS FCIP storage optimization module. Traffic originated through Symmetrix GigE ports (RE ports) require configuration of the RiOS SRDF storage optimization module. For details, see "Configuring SRDF Optimization" on page 144. For details on storage technologies that originate traffic through FC, see the *Riverbed Deployment Guide*.

You configure the RiOS FCIP storage optimization module on the Steelhead appliance closest to the FCIP gateway that opens the FCIP TCP connection by sending the initial SYN packet. This can vary by environment. If you are unsure which gateway initiates the SYN, enable FCIP on both the client-side and server-side Steelhead appliances.

By default, FCIP optimization is disabled.

For details on data replication deployments, see the *Riverbed Deployment Guide*.

To configure FCIP optimization

1. Choose Configure > Optimization > FCIP to display the FCIP page.

Figure 4-29. FCIP Page

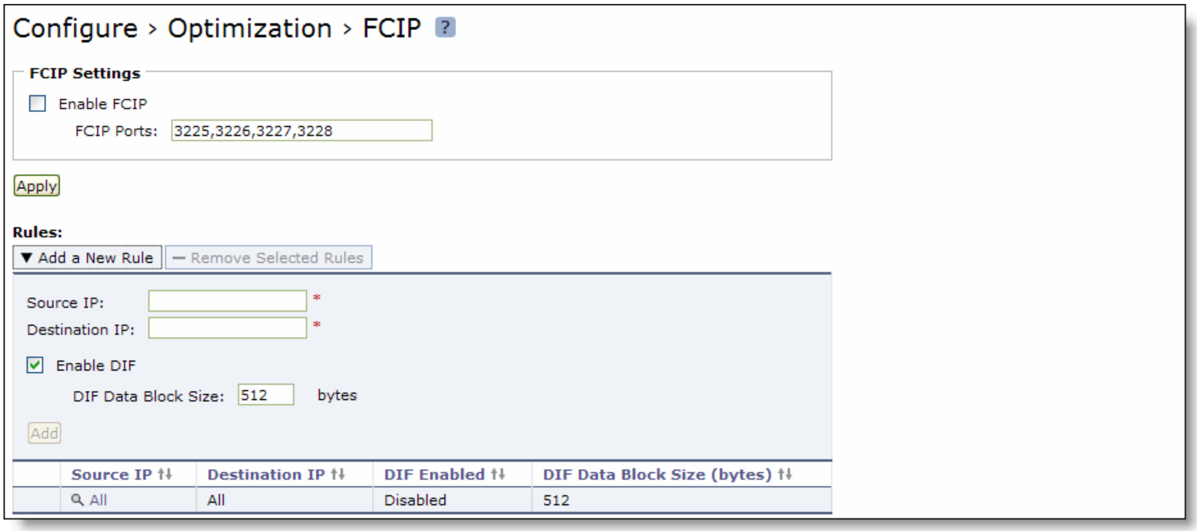

Configure > Optimization > FCIP

FCIP Settings

□ Enable FCIP

FCIP Ports: 3225,3226,3227,3228

Source IP: *

Destination IP: *

☑ Enable DIF

DIF Data Block Size: 512 bytes

	Source IP	Destination IP	DIF Enabled	DIF Data Block Size (bytes)
	All	All	Disabled	512

2. Under FCIP Settings, select Enable FCIP. By default, RiOS directs all traffic on the standard ports 3225, 3226, 3227, and 3228 through the FCIP optimization module. For most environments, the configuration is complete and you can skip to step 4.

   Environments with RF-originated SRDF traffic between VMAX arrays need additional configuration to isolate and optimize the DIFs embedded within the headers of the FCIP data payload. For details, see "FCIP Rules (VMAX-to-VMAX Traffic Only)" on page 142.

3. Optionally, you can add FCIP port numbers separated by commas or remove a port number. Do not specify a port range.

Note: The FCIP ports field must always contain at least one FCIP port.

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

6. If you have enabled or disabled FCIP optimization or changed a port, you need to restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Viewing FCIP Connections

After completing the FCIP configuration on both Steelhead appliances and restarting the optimization service, you can view the FCIP connections in the Current Connections report. Choose Reports > Networking > Current Connections. In the list of optimized connections, look for the FCIP connection in the Application column. Verify that the FCIP connection appears in the list without a Protocol Error icon:

- If the report lists a connection as TCP instead of FCIP, the module is not optimizing the connection. You need to verify the configuration. For example, make sure that the peer Steelhead appliances are running RiOS v6.1 or later.
- If the report lists a connection as FCIP but a red protocol error icon appears in the Notes column, click the magnifying glass to view the reason for the error.

For details, see "Viewing Current Connections" on page 385.

Note: You can view combined throughput and reduction statistics for two or more FCIP tunnel ports by entering the following command from the Command-Line Interface:

```
protocol fcip stat-port <num>
```

For details, see the *Riverbed Command-Line Interface Reference Manual*.

FCIP Rules (VMAX-to-VMAX Traffic Only)

Environments with RF-originated SRDF traffic between VMAX arrays need additional configuration beyond enabling FCIP optimization. You need to add FCIP rules to isolate the Data Integrity Field (DIF) headers within the FCIP data stream. These DIF headers further interrupt the data stream. You can add or remove FCIP rules by defining a match for source or destination IP traffic.

The FCIP Default Rule

The default rule optimizes all remaining traffic that has not been selected by another rule. It always appears as the last in the list. You cannot remove the default rule; however, you can change its DIF setting. The default rule uses 0.0.0.0 in the source and destination IP address fields, specifying all IP addresses. You cannot specify 0.0.0.0 as the source or destination IP address for any other rule.

To add an FCIP rule

1. Choose Configure > Optimization > FCIP to display the FCIP page.

2. Under Rules, complete the configuration as described in the following table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Source IP	Specify the connection source IP address of the FCIP gateway tunnel endpoints. **Note:** The source IP address cannot be the same as the destination IP address.

Control	Description
Destination IP	Specify the connection destination IP address of the FCIP gateway tunnel endpoints.
Enable DIF	Isolates and optimizes the DIFs embedded within the FCIP data workload.
DIF Data Block Size	Specify the size of a standard block of storage data, in bytes, after which a DIF header begins. The valid range is from 1 - 2048 bytes. The default value is 512, which is a standard block size for Open System environments. When you enable DIF, RiOS FCIP optimization looks for a DIF header after every 512 bytes of storage data unless you change the default setting. Open System environments (such as Windows, UNIX, and Linux) inject the DIF header into the data stream after every 512 bytes of storage data. IBM iSeries AS/400 host environments inject the DIF header into the data stream after every 520 bytes. This field is required when you enable DIF.
Add	Adds the rule to the list. The Management Console redisplays the Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected	Select the check box next to the name and click **Remove Selected**.
Move Selected	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. You need to restart the optimization service after adding or removing a FCIP rule. For details, see "Starting and Stopping the Optimization Service" on page 345.

To edit an FCIP rule

1. Choose Configure > Optimization > FCIP to display the FCIP page.

2. Select the rule number in the rule list.

3. Edit the rule.

4. Click **Save** to save your settings permanently.

Example—Adding an FCIP rule to isolate DIF headers on the FCIP tunnel carrying the VMAX-to-VMAX SRDF traffic.

Suppose your environment consists mostly of regular FCIP traffic without DIF headers that has some RF-originated SRDF between a pair of VMAX arrays. A pair of FCIP gateways uses a tunnel to carry the traffic between these VMAX arrays. The source IP address of the tunnel is 10.0.0.1 and the destination IP is 10.5.5.1. The pre-existing default rule does not look for DIF headers on FCIP traffic. It handles all of the non-VMAX FCIP traffic. To isolate the DIF headers on the FCIP tunnel carrying the VMAX-to-VMAX SRDF traffic, add the following rule.

1. Choose Configure > Optimization > FCIP to display the FCIP page.

2. Click **Add a New Rule**.

3. Specify the following properties for the FCIP rule.

Control	Setting
Source IP	10.0.0.1.
Destination IP	10.5.5.1
Enable DIF	Select the check box.
DIF Data Block Size	Leave the default setting 512.

4. Click **Add**.

Related Topic

- "Configuring SRDF Optimization" on page 144

Configuring SRDF Optimization

You can enable and modify SRDF storage module optimization settings in the Configure > Optimization > SRDF page.

EMC's Symmetrix Remote Data Facility/Asynchronous (SRDF/A) is a SAN replication product. It carries out the data replication over GigE (instead of the Fibre Channel) using gateways that implement the SRDF protocol.

RiOS v6.1 and later SRDF storage optimization provides support for environments using storage technology that originates traffic through Symmetrix GigE ports. For details on storage technologies that originate traffic through GigE RE ports, see the *Riverbed Deployment Guide*.

To increase the data reduction LAN-to-WAN ratio with either equal or greater data throughput in environments with SRDF traffic, RiOS separates the SRDF headers from the application data workload written to storage. The SRDF headers contain changing protocol state information, such as sequence numbers. These headers interrupt the network stream and reduce the ability of SDR to match large, contiguous data patterns. After isolating the header data, the Steelhead appliance performs SDR network deduplication on the larger, uninterrupted storage data workload and LZ compression on the headers. RiOS then optimizes, reassembles, and delivers the data to the TCP consumer without compromising data integrity.

Note: Traffic originated through Symmetrix GigE ports (RE ports) require configuration of the RiOS SRDF storage optimization module. Environments with SRDF traffic originated through Symmetrix FC ports (RF ports) require configuration of the RiOS FCIP storage optimization module. For details, see "Configuring FCIP Optimization" on page 139.

You configure the SRDF storage optimization module on the Steelhead appliance closest to the Symmetrix array that opens the SRDF TCP connection by sending the initial SYN packet. This can vary by environment. If you are unsure which array initiates the SYN, configure SRDF on both the client-side and server-side Steelhead appliances.

By default, SRDF optimization is disabled.

For details on data replication deployments, see the *Riverbed Deployment Guide*.

To configure SRDF optimization

1. Choose Configure > Optimization > SRDF to display the SRDF page.

Figure 4-30. SRDF Page

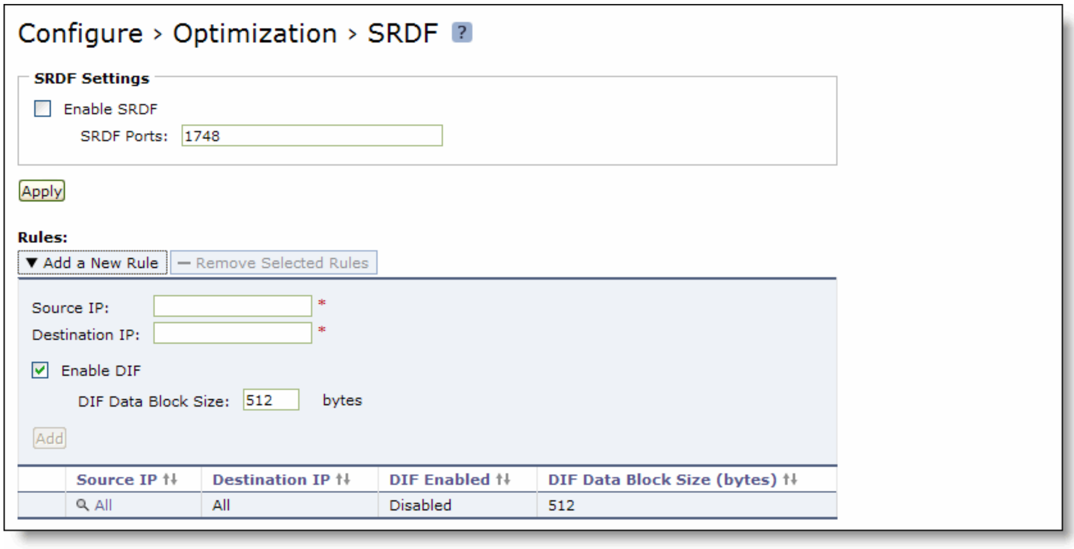

Configure > Optimization > SRDF

SRDF Settings

□ Enable SRDF

SRDF Ports: 1748

Apply

Rules:

▼ Add a New Rule — Remove Selected Rules

Source IP: *

Destination IP: *

☑ Enable DIF

DIF Data Block Size: 512 bytes

Add

	Source IP	Destination IP	DIF Enabled	DIF Data Block Size (bytes)
	All	All	Disabled	512

2. Under SRDF Settings, select Enable SRDF. By default, RiOS directs all traffic on the standard port 1748 through the SRDF module for enhanced SRDF header isolation. For most environments, the configuration is complete and you can skip to step 4.

   Environments with RE-originated SRDF traffic between VMAX arrays need additional configuration to isolate and optimize the DIFs embedded within the data payload. For details, see “SRDF Rules (VMAX-to-VMAX Traffic Only)” on page 146.

3. Optionally, specify non-standard individual SRDF port numbers separated by commas. Do not specify a port range.

Note: The SRDF ports field must always contain at least one port.

---

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

6. If you have enabled or disabled SRDF optimization or changed a port, you need to restart the optimization service. For details, see “Starting and Stopping the Optimization Service” on page 345.

Viewing SRDF Connections

After completing the SRDF configuration on both Steelhead appliances and restarting the optimization service, you can view the SRDF connections in the Current Connections report.

- If the report lists a connection as TCP instead of SRDF, the module is not optimizing the connection. You need to verify the configuration. For example, make sure that the peer Steelhead appliances are running RiOS v6.1 or later.
- If the report lists a connection as SRDF but a red protocol error icon appears in the Notes column, click the magnifying glass to view the reason for the error.

For details, see "Viewing Current Connections" on page 385.

SRDF Rules (VMAX-to-VMAX Traffic Only)

Environments with GigE-based RE-originated SRDF traffic between VMAX arrays need to isolate DIF headers within the data stream in addition to enabling SRDF optimization. These DIF headers further interrupt the data stream. To isolate the DIF headers, you add SRDF rules by defining a match for source or destination IP traffic.

The SRDF Default Rule

The default rule optimizes all remaining traffic that has not been selected by another rule. It always appears as the last in the list. You cannot remove the default rule; however, you can change the DIF setting of the default rule. The default rule uses 0.0.0.0 in the source and destination IP address fields, specifying all IP addresses. You cannot specify 0.0.0.0 as the source or destination IP address for any other rule.

To add a SRDF rule

1. Choose Configure > Optimization > SRDF to display the SRDF page.

2. Under Rules, complete the configuration as described in the following table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Source IP	Specify the connection source IP address of the Symmetrix or VMAX hosts handling the replication. **Note:** The source IP address cannot be the same as the destination IP address.
Destination IP	Specify the connection destination IP address of the Symmetrix or VMAX hosts receiving the replication.
Enable DIF	Isolates and optimizes the Data Integrity Fields embedded within the SRDF data workload.

Control	Description
DIF Data Block Size	Specify the size of a standard block of storage data, in bytes, after which a DIF header begins. The valid range is from 1 - 2048 bytes. The default value is 512, which is a standard block size for Open System environments. When you enable DIF, RiOS SRDF optimization looks for a DIF header after every 512 bytes of storage data unless you change the default setting. Open System environments (such as Windows, UNIX, and Linux) inject the DIF header into the data stream after every 512 bytes of storage data. IBM iSeries (AS/400) host environments inject the DIF header into the data stream after every 520 bytes. Do not add a module rule isolating DIF headers in mainframe environments, as SRDF environments that replicate mainframe traffic do not currently include DIF headers. This field is required when you enable DIF.
Add	Adds the rule to the list. The Management Console redisplays the Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected	Select the check box next to the name and click **Remove Selected**.
Move Selected	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. You need to restart the optimization service after adding or removing a SRDF rule. For details, see "Starting and Stopping the Optimization Service" on page 345.

To edit a SRDF rule

1. Choose Configure > Optimization > SRDF to display the SRDF page.

2. Select the rule number in the rule list.

3. Edit the rule.

4. Click **Save** to save your settings permanently.

Related Topic

- "Configuring FCIP Optimization" on page 139

Windows Domain Authentication

This section describes how to configure a user account that is trusted for delegation to the CIFS or MAPI service on target servers and add delegate users to a Windows Domain. Follow the procedures in this section to set up Delegation Mode for the following features

- RiOS SMB signing (earlier than RiOS v6.0)
- SMB2 signing (RiOS v6.5)
- Encrypted MAPI for Windows 7 clients (RiOS v6.1 and later).

Important: If you are upgrading from RiOS v6.1 to v6.5, you might already have a delegate user and be joined to a domain. If so, enabling SMB2 signing will work when enabled with no additional configuration.

For delegation mode, you only have to update the server-side Steelhead appliance.

Follow the procedures in this section after joining a Windows Domain and enabling the SMB Signing or MAPI optimization features. For details, see "Joining a Windows Domain or Workgroup" on page 305, "Configuring SMB Signing" on page 103, or "Configuring MAPI Optimization" on page 123.

Important: RiOS v6.1 and later supports constrained delegation for users that are in domains trusted by the server's domain.

Important: In RiOS v6.0 and later, transparent authentication replaces the delegation trust authentication for SMB signing, eliminating the need to set up delegate users. However, Windows 7 clients and SMB2 signing must use delegation trust authentication.

RiOS v6.1 and later supports Windows 7 clients, Windows 2008 R2 servers, and Windows 2008 R2 Domains (in both native and mixed-mode environments).

The SMB signing feature in RiOS v5.5.x and the Encrypted MAPI for Windows 7 feature in RiOS v6.1 and later use the Kerberos constrained delegation facility. This facility refers to the identity of a delegation user from end-to-end. A delegate user is a user with the privilege to impersonate other users and computers for specific services.

To use the Kerberos delegation facility, you add a user name to trust for delegation. When RiOS trusts a user for delegation, that user can impersonate an incoming client user. You need to create one delegate user. The credentials for the delegate user are stored on the server-side Steelhead.

The following procedures are *required* before enabling RiOS SMB1 signing in v5.5.x, SMB2 signing, and Encrypted MAPI for Windows 7 clients in v6.1 and later.

Delegation

Delegation mode in RiOS v5.5 or later requires that you manually grant the user access to delegate. A delegate user is required in each of the domains where a server is going to be optimized.

To set up the Domain Controller

- In Windows, create a user in the domain controller responsible for the domain of which the CIFS or MAPI server is a member. Choose Active Directory Users and Computers > Domain Name > Users and create the user (for example, with the name *delegate_user*). Alternatively, you can select an existing user.

  The next step is to create a Service Principal Name for the delegate user.

To create the Service Principal Name (SPN)

- In Windows, create an SPN for the user using the setspn.exe command-line tool. The Windows Server 2003 SP1 Support Tools product CD includes this tool or you can download it from the Microsoft Download Center. The SPN:

  - must be unique because the DC assigns the Kerberos ticket for it.
  - cannot be used by another service.
  - cannot be cifs/<hostname of domain controller>, or mapi/<hostname of domain controller>, which are used by the CIFS and MAPI services.

  For example:

  ```
   c:\> setspn.exe -A cifs/delegate delegate_user
  ```

  where

  ```
   -A
  ```

  adds the specified SPN to the specified account

  ```
  cifs/delegate
  ```

  is the name of the SPN, and

  ```
  delegate_user
  ```

  is the name of the delegate user account.

Note: For details on SPN, go to
http://msdn.microsoft.com/en-us/library/ms677949%28VS.85%29.aspx

---

The next step is to grant the user access to delegate for the CIFS or MAPI service in Windows. You must perform the following procedure for every server on which you want to enable RiOS SMB signing or Encrypted MAPI.

To grant the user access to delegate

1. Open Active Directory Users and Computers, and select the Delegate User > Properties > Delegation tab.

Note: If the Delegation tab does not appear, raise the Windows Domain functionality to the Windows 2003 level or higher and create a Service Principal Name for the delegate user.

---

2. Select Trust this user for delegation to specified services only and Use any authentication protocol.

Figure 4-31. Granting User Access to Delegate

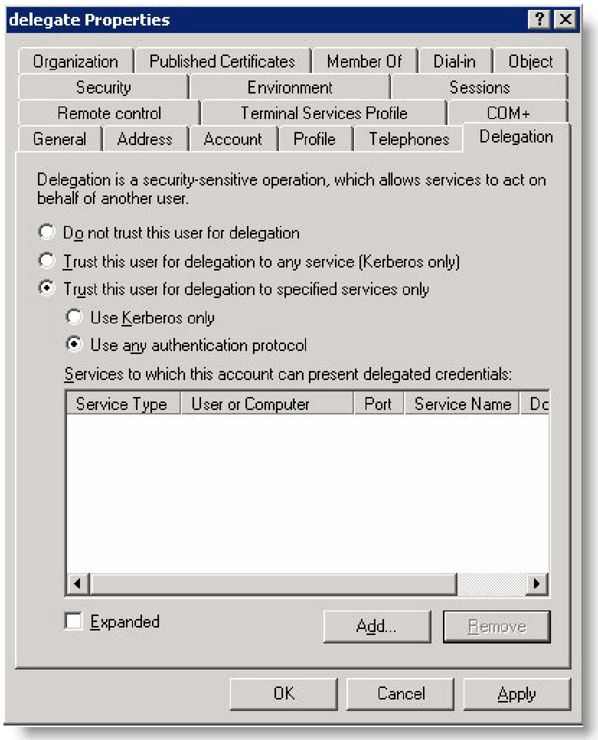

Delegation is a security-sensitive operation, which allows services to act on behalf of another user.

○ Do not trust this user for delegation

○ Trust this user for delegation to any service (Kerberos only)

● Trust this user for delegation to specified services only

○ Use Kerberos only

● Use any authentication protocol

Services to which this account can present delegated credentials:

□ Expanded

3. Click **Add**.

Figure 4-32. Adding the Server to be Delegated

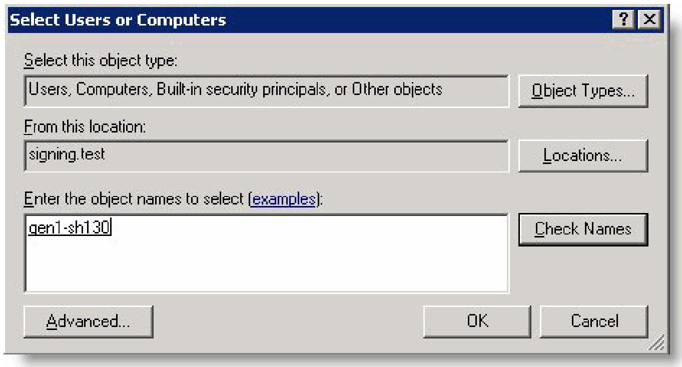

Select this object type:
Users, Computers, Built-in security principals, or Other objects

From this location:
signing.test

Enter the object names to select (examples):
gen1-sh130

4. Click Users or Computers.

5. In the Select Users or Computers dialog box, enter the CIFS or MAPI server as the local hostname and click **OK**.

Figure 4-33. Entering the Server Name

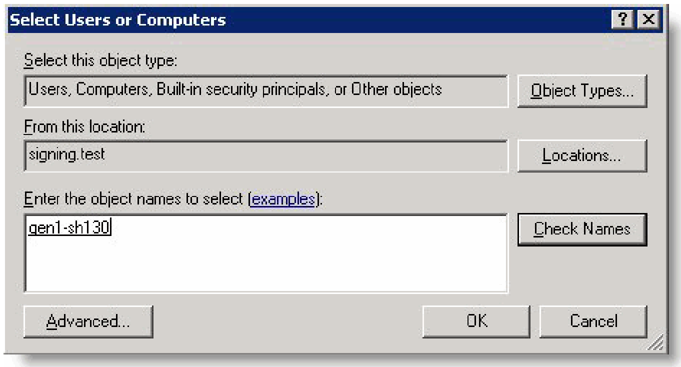

Select Users or Computers

Select this object type:
Users, Computers, Built-in security principals, or Other objects

From this location:
signing.test

Enter the object names to select (examples):
gen1-sh130

6. In the Add Services dialog box, select either the CIFS or exchangeMDB service (MAPI) type for delegation and click **OK**.

Figure 4-34. Adding the Service for Delegation

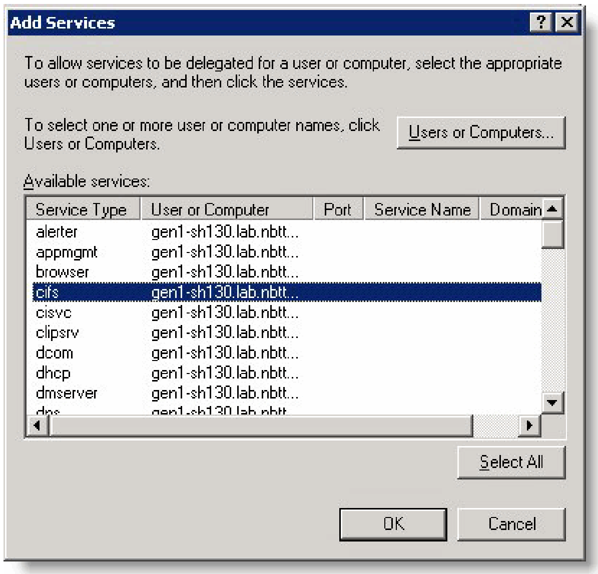

Add Services

To allow services to be delegated for a user or computer, select the appropriate users or computers, and then click the services.

To select one or more user or computer names, click Users or Computers.

Available services:

Service Type	User or Computer	Port	Service Name	Domain
alerter	gen1-sh130.lab.nbtt...			
appmgmt	gen1-sh130.lab.nbtt...			
browser	gen1-sh130.lab.nbtt...			
cifs	gen1-sh130.lab.nbtt...			
cisvc	gen1-sh130.lab.nbtt...			
clipsrv	gen1-sh130.lab.nbtt...			
dcom	gen1-sh130.lab.nbtt...			
dhcp	gen1-sh130.lab.nbtt...			
dmserver	gen1-sh130.lab.nbtt...			
dns	gen1-sh130.lab.nbtt			

Once you have performed steps 1 through 6 for every server on which you want to enable RiOS SMB signing or Encrypted MAPI, the next step is to add delegate users to the server-side Steelhead appliance.

Note: For automatic delegation mode, you do not need to perform steps 1 through 6 for all servers but you must still configure one CIFS or exchangeMDB service, as this is required by the Active Directory interface. Also continue with the Steelhead delegate user configuration steps that follow.

To add delegate users on the Steelhead appliance

1. On the server-side Steelhead appliance, choose Configure > Optimization > Windows Domain Auth to display the Windows Domain Auth page.

Figure 4-35. Adding a New Delegate User

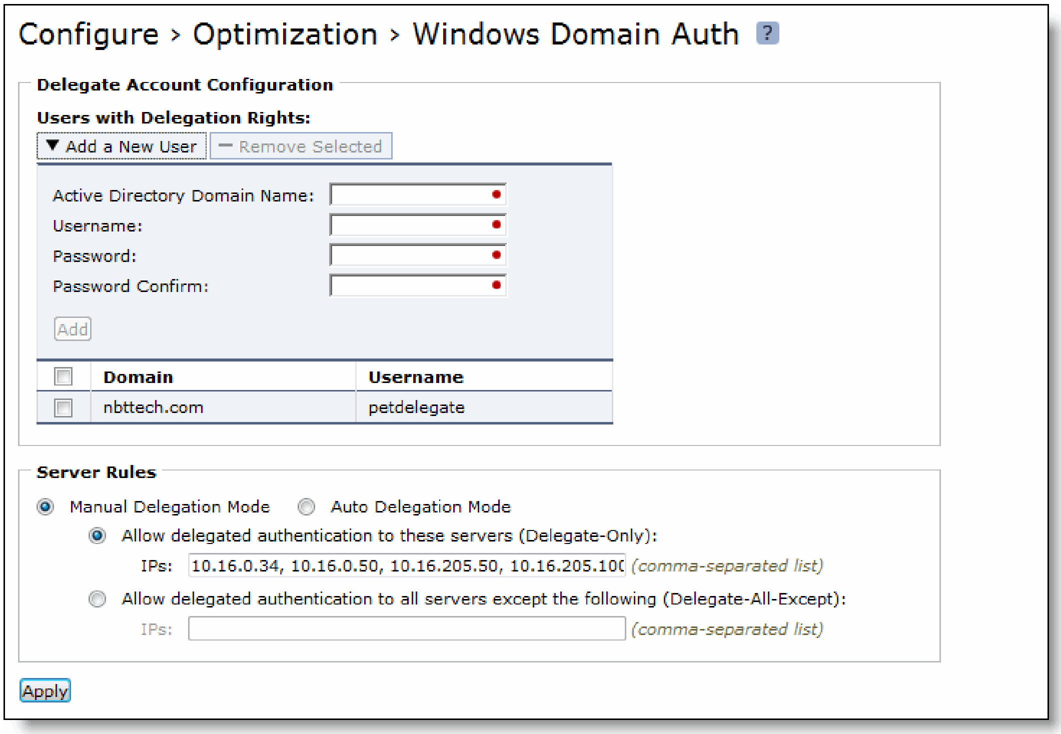

Configure > Optimization > Windows Domain Auth

Delegate Account Configuration

Users with Delegation Rights:

▼ Add a New User | — Remove Selected

Active Directory Domain Name:

Username:

Password:

Password Confirm:

Add

□	Domain	Username
□	nbttech.com	petdelegate

Server Rules

● Manual Delegation Mode ○ Auto Delegation Mode

● Allow delegated authentication to these servers (Delegate-Only):

IPs: 10.16.0.34, 10.16.0.50, 10.16.205.50, 10.16.205.10(*(comma-separated list)*

○ Allow delegated authentication to all servers except the following (Delegate-All-Except):

IPs: *(comma-separated list)*

Apply

2. Under Users with Delegation Rights, complete the configuration as described in the following table.

Control	Description
Add a New User	Displays the controls to add a user with trusted delegation rights to a domain. **Important:** You can only add one delegate user per domain. A delegate user is required in each of the domains where a server is going to be optimized.
Active Directory Domain Name	Specify the delegation domain in which you want to make the delegate user a trusted member, for example: `SIGNING.TEST` **Note:** You cannot specify a single-label domain name (a name without anything after the dot), as in riverbed instead of riverbed.com.
Username	Specify the delegate user name. The maximum length is 20 characters. The username cannot contain any of the following characters: / \ [ ] : ; | = , + * ? < > @ " **Note:** The system translates the user name into uppercase to match the registered server realm information.
Password	Specify the user account password.
Password Confirm	Confirm the user account password.
Add	Adds the user.

3. Click **Apply** to apply your settings to the running configuration.

To set up manual delegation (specifying each server allowed to delegate), continue to the next procedure.

To set up auto-delegation (automatic server detection), see "Auto-Delegation Mode" on page 154.

To specify manual delegation mode and allowed servers

1. On the server-side Steelhead appliance, choose Configure > Optimization > Windows Domain Auth to display the Windows Domain Auth page.

2. Under Server Rules, complete the configuration as described in the following table.

Control	Description
Manual Delegation Mode	Enables transparent authentication using NTLM and provides more control to specify the exact servers to perform optimization for. When you select this mode, you need to specify each server on which to delegate and sign for each domain using the Delegate-Only and Delegate-All-Except controls. This is the default setting in RiOS v6.0 and later.
Auto Delegation Mode	Enables delegate user authentication and automatically discovers the servers on which to delegate and sign. This eliminates the need to set up the servers on which to delegate and sign for each domain. This mode requires additional configuration. For details, see Auto-Delegation mode. A delegate user is required in each of the domains where a server is going to be optimized. **Note:** If you switch between manual and automatic delegation you must restart the optimization service.
Allow delegated authentication to these servers (Delegate-Only)	Click to intercept the connections destined for the servers in this list. By default, this setting is enabled. Specify the file server IP addresses for SMB signed or MAPI encrypted traffic in the text box, separated by commas. **Tip:** You can switch between the Delegate-Only and Delegate-All-Except controls without losing the list of IP addresses for the control. Only one list is active at a time.
Allow delegated authentication to all servers except the following (Delegate-All-Except)	Click to intercept all of the connections except those destined for the servers in this list. Specify the file server IP addresses that do *not* require SMB signing or MAPI encryption in the text box, separated by commas. By default, this setting is disabled. Only the file servers that do *not* appear in the list are signed or encrypted. **Note:** You must register any servers not on this list with the Domain Controller or be using Auto-Delegation Mode.

3. Click **Apply** to apply your settings to the running configuration.

4. Click **Save** to save your settings permanently.

5. If you change the delegation mode, you must restart the optimization service.

Note: A delegate user with access to the CIFS and exchangeMDB (MAPI) service does not have log on privileges.

Auto-Delegation Mode

Delegation mode in RiOS v6.1 and later automatically updates the delegate user in Active Directory with delegation rights to servers. The service updates the user in real-time, eliminating the need to grant the user access to delegate on every server. Auto-delegation mode also updates the server IP address if it changes.

This section describes how to grant special privileges to the delegate user so they have automatic delegation rights to servers. The first step is to create a Delegate User with a Service Principal Name (SPN). The procedure to create a delegate user with an SPN is the same for both Windows DC R2 2003 and Windows DC R2 2008. Next, you need to grant the delegate user the right to delegate on the Domain Controller. Because the procedures to grant the delegate user rights on the DC is different for Windows DC R2 2003 and Windows DC R2 2008, the procedures to do so are separate.

Note: A delegate user that is an Administrator already has the correct delegation rights for auto-delegation mode.

Note: A delegate user is required in each of the domains where a server is going to be optimized.

Note: If you update the password for the delegate user in Active Directory, you must also update the account information on the Steelhead appliance. To do this, delete the old account and add a new one with the updated password.

To create a delegate user with an SPN

- In Windows, create an SPN for the user using the setspn.exe command-line tool. The Windows Server 2003 SP1 Support Tools product CD includes this tool or you can download it from the Microsoft Download Center. The SPN:
  - must be unique because the DC assigns the Kerberos ticket for it.
  - cannot be used by another service.
  - cannot be cifs/<hostname of domain controller>, or mapi/<hostname of domain controller>, which are used by the CIFS and MAPI services.

  For example:

  ```
   c:\> setspn.exe -A cifs/delegate delegate_user
  ```

  where

  ```
   -A
  ```

  adds the specified SPN to the specified account

  ```
  cifs/delegate
  ```

  is the name of the SPN, and

  ```
  delegate_user
  ```

is the name of the delegate user account.

Note: For details on SPN, go to
http://msdn.microsoft.com/en-us/library/ms677949%28VS.85%29.aspx

---

To grant the delegate user rights in the Controlling Security Group Policy Object (GPO) for Windows DC R2 2008

1. Choose Start > Administrative Tools > Group Policy Management to display the Group Policy Management viewer.

2. Locate the Domain Controller Security Policy and choose Edit to display the Group Policy Management editor.

Figure 4-36. Editing the Domain Controller Security Policy

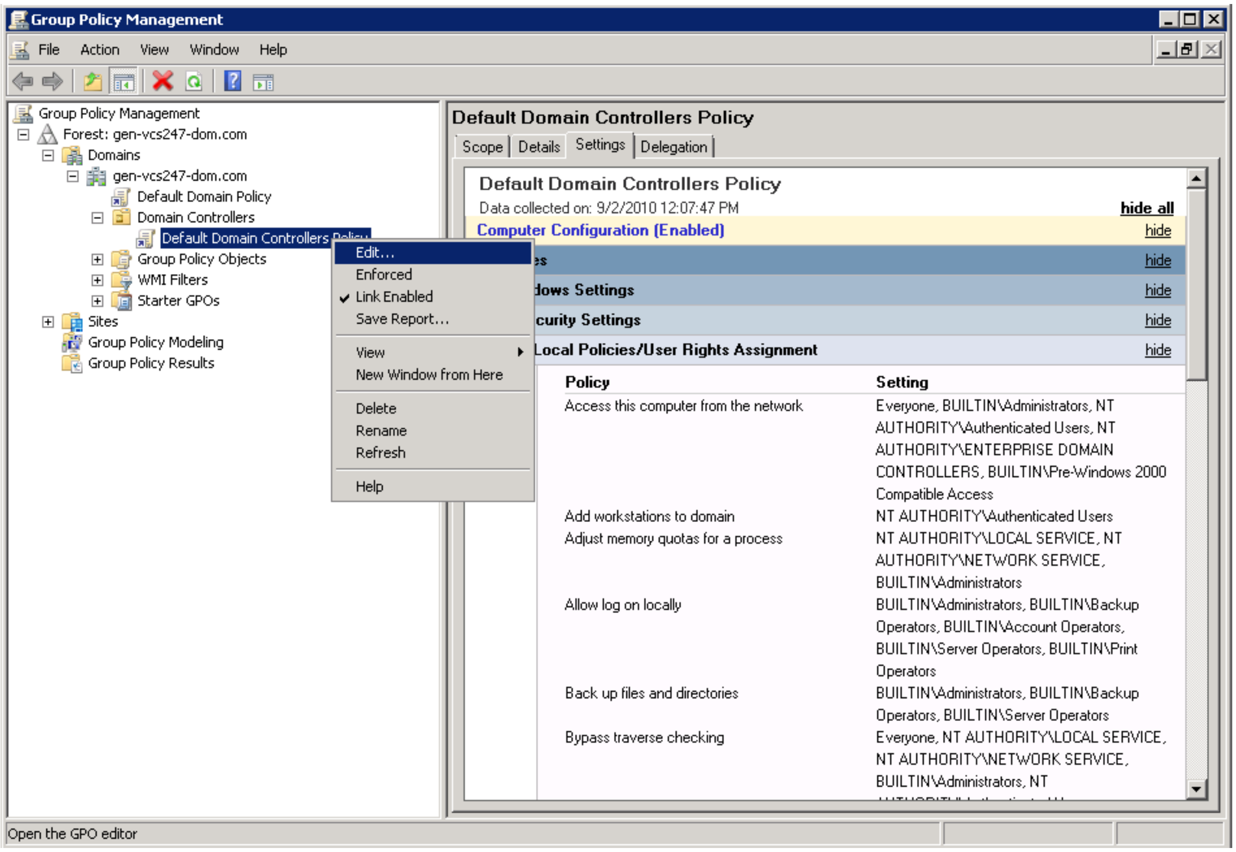

Default Domain Controllers Policy

Default Domain Controllers Policy
Data collected on: 9/2/2010 12:07:47 PM

Computer Configuration (Enabled)

...es

...dows Settings

...curity Settings

Local Policies/User Rights Assignment

Policy	Setting
Access this computer from the network	Everyone, BUILTIN\Administrators, NT AUTHORITY\Authenticated Users, NT AUTHORITY\ENTERPRISE DOMAIN CONTROLLERS, BUILTIN\Pre-Windows 2000 Compatible Access
Add workstations to domain	NT AUTHORITY\Authenticated Users
Adjust memory quotas for a process	NT AUTHORITY\LOCAL SERVICE, NT AUTHORITY\NETWORK SERVICE, BUILTIN\Administrators
Allow log on locally	BUILTIN\Administrators, BUILTIN\Backup Operators, BUILTIN\Account Operators, BUILTIN\Server Operators, BUILTIN\Print Operators
Back up files and directories	BUILTIN\Administrators, BUILTIN\Backup Operators, BUILTIN\Server Operators
Bypass traverse checking	Everyone, NT AUTHORITY\LOCAL SERVICE, NT AUTHORITY\NETWORK SERVICE, BUILTIN\Administrators, NT

3. Choose Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment.

Figure 4-37. Navigating to the User Rights Assignment

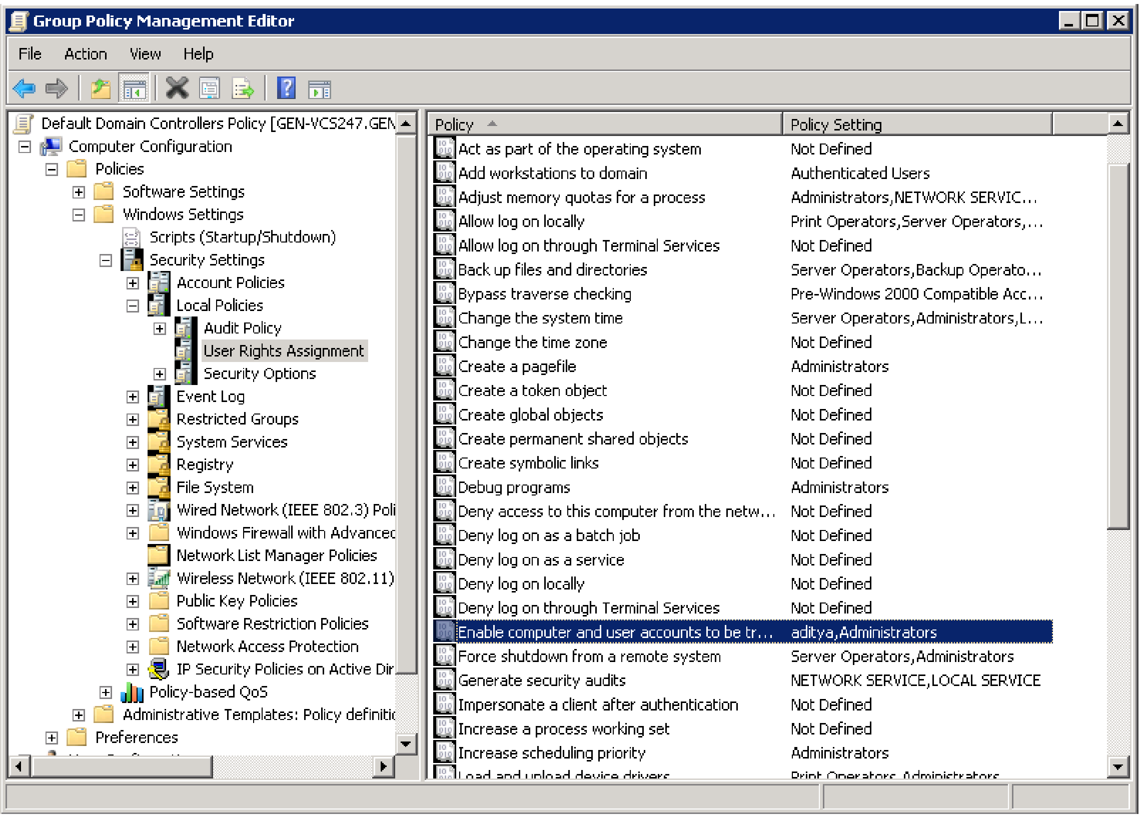
Group Policy Management Editor
File
Action
View
Help
Default Domain Controllers Policy [GEN-VCS247.GEN
Computer Configuration
Policies
Software Settings
Windows Settings
Scripts (Startup/Shutdown)
Security Settings
Account Policies
Local Policies
Audit Policy
User Rights Assignment
Security Options
Event Log
Restricted Groups
System Services
Registry
File System
Wired Network (IEEE 802.3) Poli
Windows Firewall with Advancec
Network List Manager Policies
Wireless Network (IEEE 802.11)
Public Key Policies
Software Restriction Policies
Network Access Protection
IP Security Policies on Active Dir
Policy-based QoS
Administrative Templates: Policy definitic
Preferences

Policy	Policy Setting
Act as part of the operating system	Not Defined
Add workstations to domain	Authenticated Users
Adjust memory quotas for a process	Administrators,NETWORK SERVIC...
Allow log on locally	Print Operators,Server Operators,...
Allow log on through Terminal Services	Not Defined
Back up files and directories	Server Operators,Backup Operato...
Bypass traverse checking	Pre-Windows 2000 Compatible Acc...
Change the system time	Server Operators,Administrators,L...
Change the time zone	Not Defined
Create a pagefile	Administrators
Create a token object	Not Defined
Create global objects	Not Defined
Create permanent shared objects	Not Defined
Create symbolic links	Not Defined
Debug programs	Administrators
Deny access to this computer from the netw...	Not Defined
Deny log on as a batch job	Not Defined
Deny log on as a service	Not Defined
Deny log on locally	Not Defined
Deny log on through Terminal Services	Not Defined
Enable computer and user accounts to be tr...	aditya,Administrators
Force shutdown from a remote system	Server Operators,Administrators
Generate security audits	NETWORK SERVICE,LOCAL SERVICE
Impersonate a client after authentication	Not Defined
Increase a process working set	Not Defined
Increase scheduling priority	Administrators
Load and unload device drivers	Print Operators,Administrators

4. In the right pane, Under Policy, right-click Enable computer and user accounts to be trusted for delegation policy.

5. Click **Add** to display the Add User or Group dialog box.

Figure 4-38. Adding a Delegate User to the Group Policy Management in Windows 2008

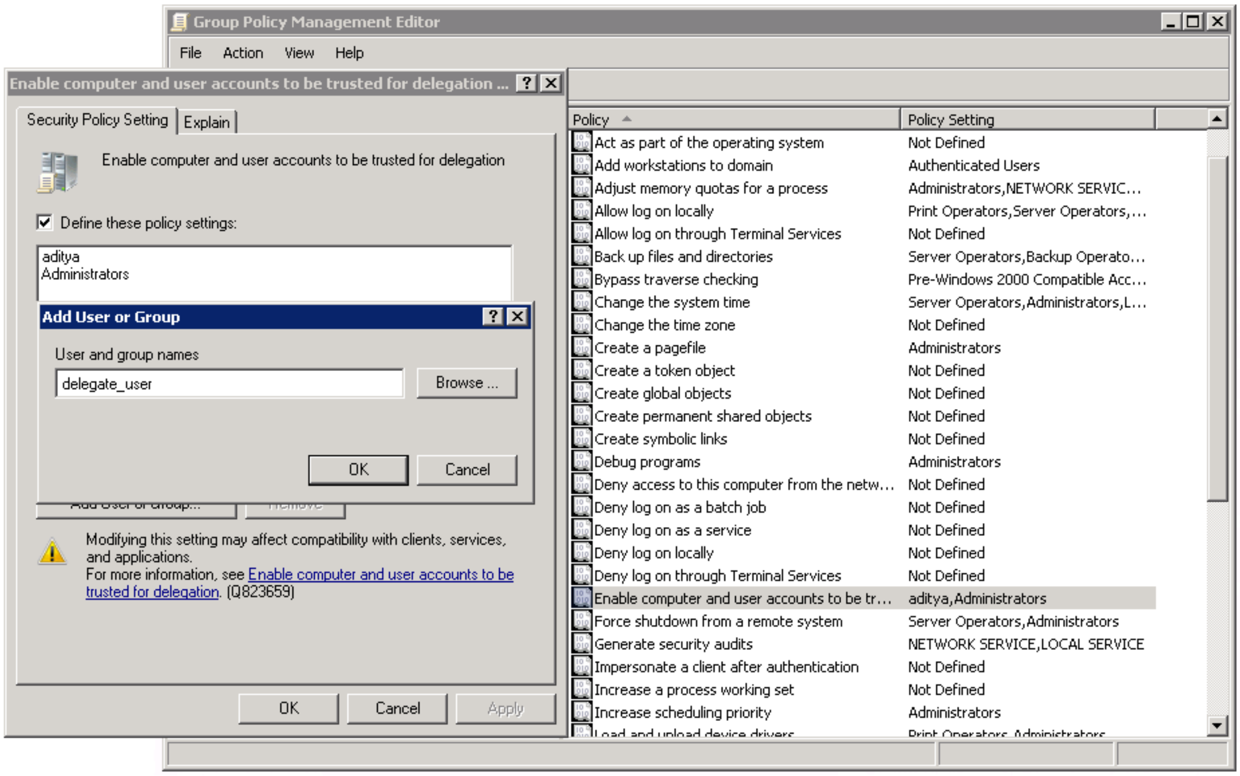

Enable computer and user accounts to be trusted for delegation

☑ Define these policy settings:

aditya
Administrators

User and group names: delegate_user

Modifying this setting may affect compatibility with clients, services, and applications.
For more information, see Enable computer and user accounts to be trusted for delegation. (Q823659)

Policy	Policy Setting
Act as part of the operating system	Not Defined
Add workstations to domain	Authenticated Users
Adjust memory quotas for a process	Administrators,NETWORK SERVIC...
Allow log on locally	Print Operators,Server Operators,...
Allow log on through Terminal Services	Not Defined
Back up files and directories	Server Operators,Backup Operato...
Bypass traverse checking	Pre-Windows 2000 Compatible Acc...
Change the system time	Server Operators,Administrators,L...
Change the time zone	Not Defined
Create a pagefile	Administrators
Create a token object	Not Defined
Create global objects	Not Defined
Create permanent shared objects	Not Defined
Create symbolic links	Not Defined
Debug programs	Administrators
Deny access to this computer from the netw...	Not Defined
Deny log on as a batch job	Not Defined
Deny log on as a service	Not Defined
Deny log on locally	Not Defined
Deny log on through Terminal Services	Not Defined
Enable computer and user accounts to be tr...	aditya,Administrators
Force shutdown from a remote system	Server Operators,Administrators
Generate security audits	NETWORK SERVICE,LOCAL SERVICE
Impersonate a client after authentication	Not Defined
Increase a process working set	Not Defined
Increase scheduling priority	Administrators
Load and unload device drivers	Print Operators,Administrators

6. Specify the delegate user name and click **OK**.

7. Click **OK** to close the Group Policy Management editor.

Now that the delegate user has rights in the Windows 2008 GPO, you need to grant the delegate user more privileges to use auto-delegation mode. For details, see “To grant the delegate user rights to modify the msDS-AllowedToDelegateTo Active Directory attribute on itself” on page 158.

To grant the delegate user rights in the Controlling Security Group Policy Object (GPO) in Windows 2003

1. Choose Start > Administrative Tools > Domain Controller Security Policy.

   The GPO viewer appears.

2. Choose Security Settings > Local Policies > User Rights Assignment.

3. Right-click Enable computer and user accounts to be trusted for delegation policy.

Figure 4-39. Adding a Delegate User to the Group Policy Management in Windows 2003

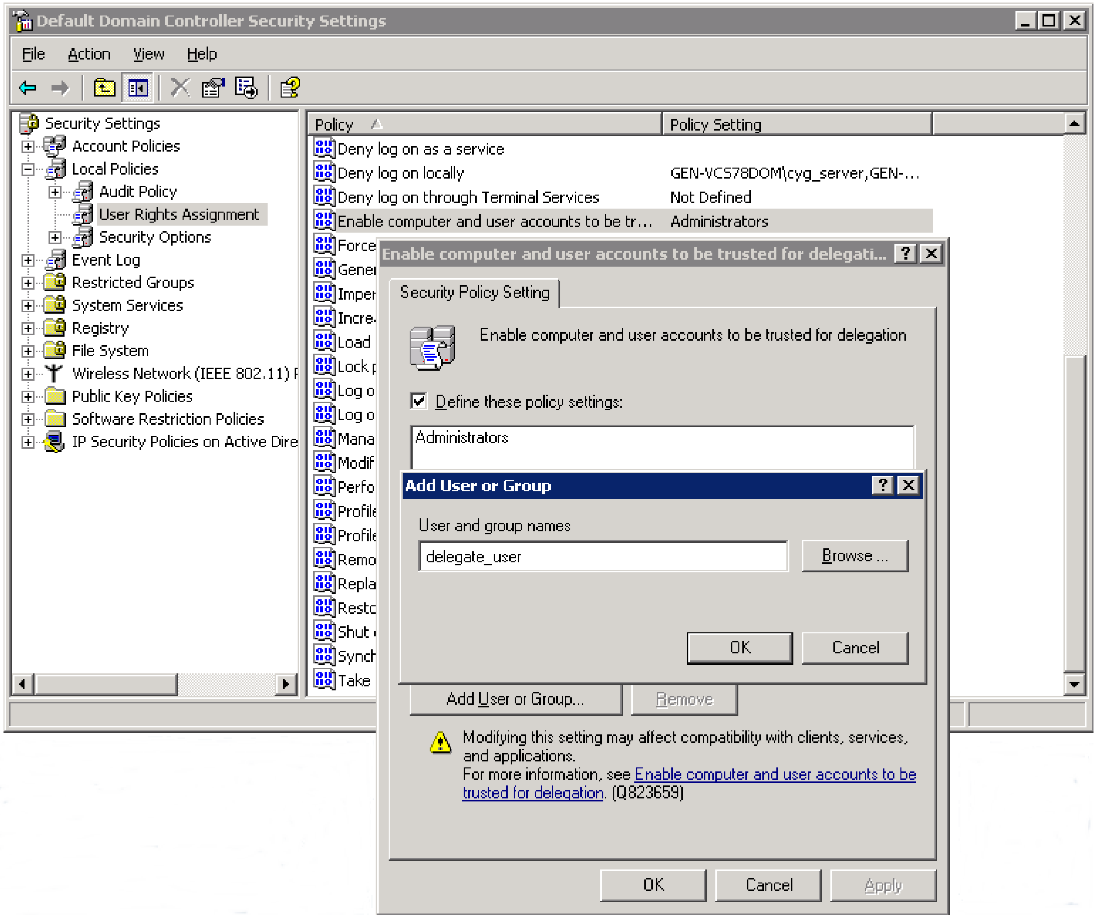

Policy	Policy Setting
Deny log on as a service	
Deny log on locally	GEN-VCS78DOM\cyg_server,GEN-...
Deny log on through Terminal Services	Not Defined
Enable computer and user accounts to be tr...	Administrators

Enable computer and user accounts to be trusted for delegation

☑ Define these policy settings:

Administrators

User and group names

delegate_user

Modifying this setting may affect compatibility with clients, services, and applications.
For more information, see Enable computer and user accounts to be trusted for delegation. (Q823659)

4. Click **Properties**.

5. Specify the delegate user name.

6. Click **OK** to add the user name.

7. Close the Group Policy Management editor.

8. Click **OK** to close the Group Policy Management viewer.

Now that the delegate user has rights in the Windows 2008 GPO, you need to grant the delegate user more privileges to use auto-delegation mode. For details, see "To grant the delegate user rights to modify the msDS-AllowedToDelegateTo Active Directory attribute on itself" on page 158.

To grant the delegate user rights to modify the msDS-AllowedToDelegateTo Active Directory attribute on itself

1. Open the ADSI Edit utility.

2. Choose Start > Run, and open adsiedit.msc.

3. Select Default naming context > *Domain DN* > CN=Users > CN=<Delegate User>.

4. Right-click CN=<Delegate User> and select Properties.

5. Select the Security tab, select Advanced, then click **Add**.

6. Specify the delegate user name and click **OK**.

7. Select the Properties tab in the Permission Entry dialog box.

8. Select the Allow check box next to the following:
   - Read msDS-AllowedToDelegateTo
   - Write msDS-AllowedToDelegateTo

9. Click **OK**.

10. On the server-side Steelhead appliance, choose Configure > Optimization > Windows Domain Auth to display the Windows Domain Auth page.

Figure 4-40. Selecting Auto-Delegation Mode After Granting Delegate User Privileges

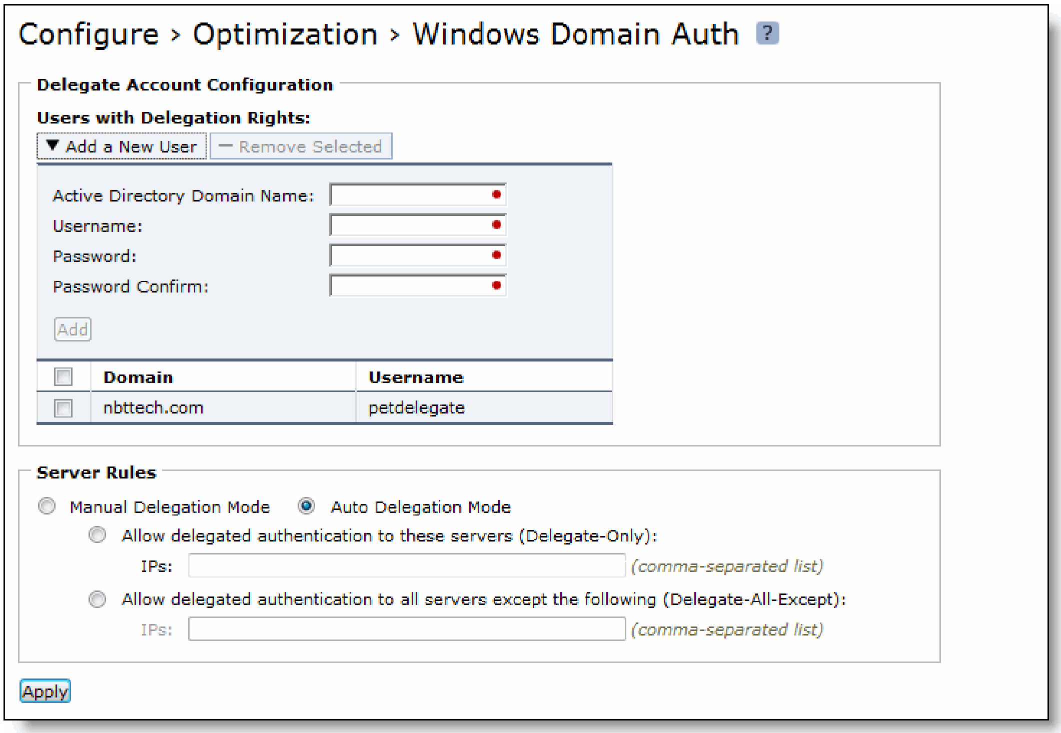

Configure > Optimization > Windows Domain Auth

Delegate Account Configuration

Users with Delegation Rights:

Active Directory Domain Name:

Username:

Password:

Password Confirm:

□	Domain	Username
□	nbttech.com	petdelegate

Server Rules

○ Manual Delegation Mode ● Auto Delegation Mode

○ Allow delegated authentication to these servers (Delegate-Only):

IPs: *(comma-separated list)*

○ Allow delegated authentication to all servers except the following (Delegate-All-Except):

IPs: *(comma-separated list)*

11. Under Server Rules, select Auto-Delegation Mode.

12. Click **Apply** to apply your settings to the running configuration.

13. Click **Save** to save your settings permanently.

14. Click **Restart** to restart the optimization service.

Troubleshooting Delegate Users

This section provides information on troubleshooting the delegate user set up, if necessary.

- When the CIFS or exchangeMDB service (MAPI) cannot obtain a delegate user's credentials, the following message appears:

  ```
  kinit: krb5_get_init_creds: Clients credentials have been revoked
  ```

  This message indicates that Login Denied is set for the delegate user for the entire day. To verify when the delegate user has permission to log in, select the Account tab in the Delegate User Properties dialog box and click Logon Hours.

- When the CIFS or exchangeMDB service cannot obtain permissions to access certain required user account attributes, the following message appears:

  ```
  kgetcred: krb5_get_creds: Client (delegate@SIGNING.TEST) unknown
  ```

  Add the delegate user to the Windows Authorization Access group. For details, see

  http://support.microsoft.com/kb/331951

- For details on constrained delegation, see

  http://technet.microsoft.com/en-us/library/cc739587(WS.10).aspx

Related Topics

- "Configuring CIFS Optimization" on page 97
- "Configuring MAPI Optimization" on page 123
- "Joining a Windows Domain or Workgroup" on page 305
- "Viewing Current Connections" on page 385

CHAPTER 5 Configuring Branch Services

This chapter describes how to enable and configure a local DNS name server, the Proxy File Service (PFS), and install and configure Riverbed Services Platform (RSP) for the Steelhead appliance. It includes the following sections:

Configuring PFS

You can enable and configure PFS support in the Configure > Branch Services > PFS Settings page.

PFS is an integrated virtual file server that allows you to store copies of files on the Steelhead appliance with Windows file access, creating several options for transmitting data between remote offices and centralized locations with improved performance. Data is configured into file shares and the shares are periodically synchronized transparently in the background, over the optimized connection of the Steelhead appliance. PFS leverages the integrated disk capacity of the Steelhead appliance to store file-based data in a format that allows it to be retrieved by NAS clients.

Important: Do not configure both RSP and PFS on the same Steelhead appliance. Riverbed does not support this configuration because PFS has no disk boundaries and can overwrite the space allocated to RSP.

Note: PFS is supported on Steelhead appliance models 250, 520, 550, 1010, 1020, 1050, 1520, 2010, 2011, 2020, 2050, 2510, 2511, 3010, 3020, 3030, 3510, 3520, and 5010.

CSH VSH Cloud Steelhead and Virtual Steelhead models do not support PFS.

When to Use PFS

Before you configure PFS, evaluate whether it is suitable for your network needs. Advantages of using PFS are:

- **LAN access to data residing across the WAN** - File access performance is improved between central and remote locations. PFS creates an integrated file server, enabling clients to access data directly from the proxy filer on the LAN as opposed to the WAN. In the background, data on the proxy filer is transparently synchronized with data from the origin-file server over the WAN.
- **Continuous access to files in the event of WAN disruption** - PFS supports disconnected operations. In the event of a network disruption that prevents access over the WAN to the origin-file server, files can still be accessed on the local Steelhead appliance.
- **Simple Branch Infrastructure and Backup Architectures** - PFS consolidates file servers and local tape backup from the branch office into the data center. PFS reduces the number and the size of backup windows running in complex backup architectures.
- **Automatic Content Distribution** - PFS automatically distributes new and changed content throughout a network.

If any of these advantages can benefit your environment, then enabling PFS in the Steelhead appliance is appropriate.

However, PFS requires pre-identification of files and is not appropriate in environments where there is concurrent read-write access to data from multiple sites:

- **Pre-identification of PFS files** - PFS requires that files accessed over the WAN are identified in advance. If the data set accessed by the remote users is larger than the specified capacity of your Steelhead appliance model, or if it cannot be identified in advance, then you must have end-users access the origin-file server directly through the Steelhead appliance without PFS. (This configuration is also known as Global mode.)
- **Concurrent Read-Write Data Access from Multiple Sites** - In a network environment where users from multiple branch offices update a common set of centralized files and records over the WAN, the Steelhead appliance without PFS is the most appropriate solution because file locking is directed between the client and the server. The Steelhead appliance always consults the origin-file server in response to a client request; it never provides a proxy response or data from its datastore without consulting the origin-file server.

Prerequisites and Tips

This section describes prerequisites and tips for using PFS:

- Before you enable PFS, configure the Steelhead appliance to use NTP to synchronize the time. To use PFS, the Steelhead appliance and Domain Controller (DC) clocks must be synchronized. For details on setting the date and time in the Steelhead appliance, see "Modifying General Host Settings" on page 39.
- The PFS Steelhead appliance must run the same version of the Steelhead appliance software as the server-side Steelhead appliance.
- PFS traffic to and from the Steelhead appliance travels through the primary interface. PFS requires that the primary interface is connected to the same switch as the LAN interface.
- The PFS share and origin-server share names cannot contain Unicode characters.

Upgrading Version 2 PFS Shares

By default, when you configure PFS shares with Steelhead appliance software v3.x, you create Version 3 PFS shares. PFS shares configured with Steelhead appliance software RiOS v2.x are Version 2 shares. Version 2 shares are not upgraded when you upgrade Steelhead appliance software.

If you have shares created with RiOS v2.x software, you must upgrade them to Version 3 shares in the Management Console. If you upgrade any Version 2 shares, you must upgrade all of them. After you have upgraded shares to Version 3, you must only create Version 3 shares.

If you do not upgrade your Version 2 shares:

- you must not create Version 3 shares.
- you must install and start the RCU on the origin-file server or on a separate Windows host with write-access to the data PFS uses. The account that starts the RCU must have write permission to the folder on the origin-file server that contains the data PFS uses. You can download the RCU from the Riverbed Support site at https://support.riverbed.com.

  In RiOS v3.x or later, you do not need to install the RCU service on the server for synchronization purposes. All RCU functionality has been moved to the Steelhead appliance.
- you must configure domain, not workgroup, settings. Domain mode supports Version 2 PFS shares but Workgroup mode does not.

Domain and Local Workgroup Settings

When you configure your PFS Steelhead appliance, set either domain or local workgroup settings. For details, see "Joining a Windows Domain or Workgroup" on page 305.

PFS Share Operating Modes

PFS provides Windows file service in the Steelhead appliance at a remote site. When you configure PFS, you specify an operating mode for each individual file share on the Steelhead appliance. The proxy-file server can export data volumes in Local mode, Broadcast mode, and Stand-Alone mode. After the Steelhead appliance receives the initial copy of the data and ACLs, shares can be made available to local clients. In Broadcast and Local mode only, shares on the Steelhead appliance are periodically synchronized with the origin-file server at intervals you specify, or manually if you choose. During the synchronization process, the Steelhead appliance optimizes this traffic across the WAN.

- **Broadcast Mode** - Use Broadcast mode for environments seeking to broadcast a set of read-only files to many users at different sites. Broadcast mode quickly transmits a read-only copy of the files from the origin-file server to your remote offices.

- **Local Mode** - Use Local mode for environments that need to efficiently and transparently copy data created at a remote site to a central data center, perhaps where tape archival resources are available to back up the data. Local mode enables read-write access at remote offices to update files on the origin-file server.
- **Stand-Alone Mode** - Use Stand-Alone mode for network environments where it is more effective to maintain a separate copy of files that are accessed locally by the clients at the remote site. The PFS share also creates more storage space.

  The PFS share on the Steelhead appliance is a one-time, working copy of data copied from the origin-file server. You can specify a remote path to a directory on the origin-file server, creating a copy at the branch office. Users at the branch office can read from or write to stand-alone shares, but there is no synchronization back to the origin-file server because a stand-alone share is an initial and one-time-only synchronization.

Figure 5-1. PFS Deployment

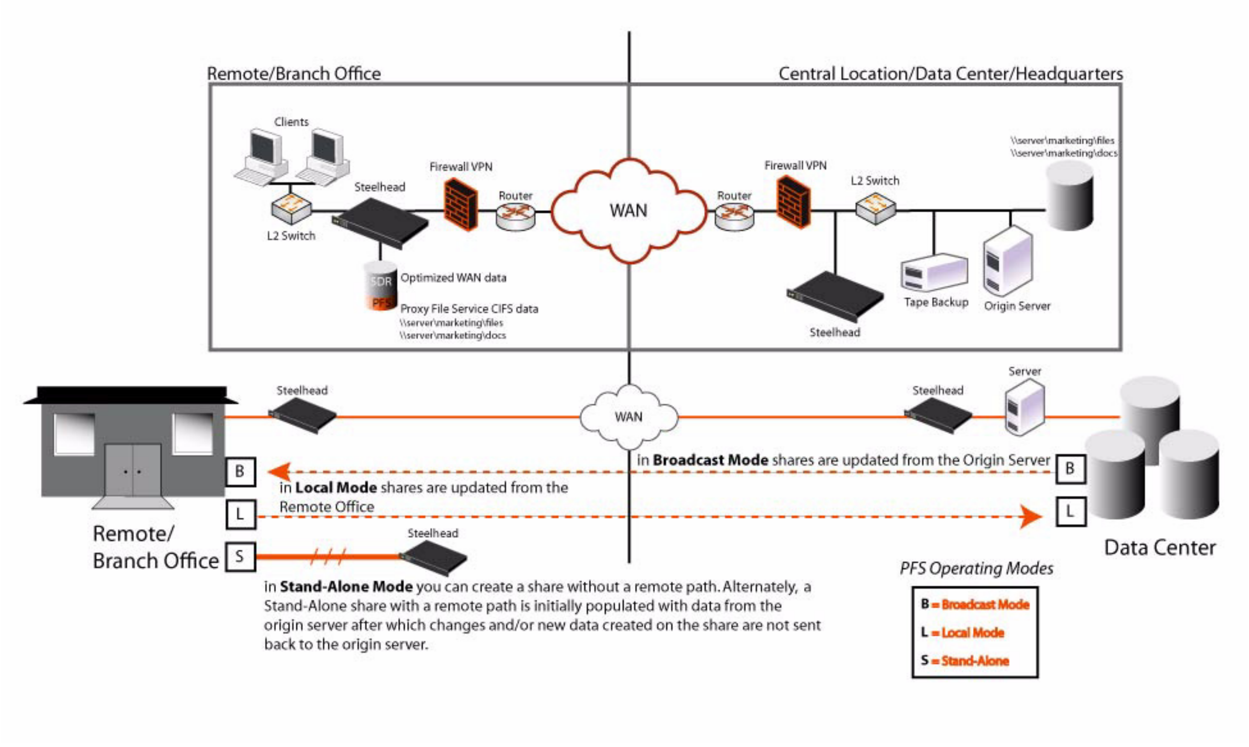
Remote/Branch Office
Central Location/Data Center/Headquarters
Clients
Steelhead
Firewall VPN
Router
L2 Switch
WAN
SDR
PFS
Optimized WAN data
Proxy File Service CIFS data
\\server\marketing\files
\\server\marketing\docs
Router
Firewall VPN
L2 Switch
\\server\marketing\files
\\server\marketing\docs
Steelhead
Tape Backup
Origin Server
Steelhead
WAN
Steelhead
Server
B
in **Broadcast Mode** shares are updated from the Origin Server
B
in **Local Mode** shares are updated from the Remote Office
L
L
Remote/
Branch Office
S
Steelhead
Data Center
in **Stand-Alone Mode** you can create a share without a remote path. Alternately, a Stand-Alone share with a remote path is initially populated with data from the origin server after which changes and/or new data created on the share are not sent back to the origin server.
PFS Operating Modes
B = Broadcast Mode
L = Local Mode
S = Stand-Alone

Important: If you set up a PFS share on a NetApp filer, the filer allows all users access regardless of the permissions set on the NetApp share. For example, if you set "No Access" for a user for a share, the NetApp filer does not translate it into the appropriate ACL (Access Control List) entry on the folder. When a PFS share is created from this origin share, the user is allowed access to the share because there is not a deny entry present in the ACL

Lock Files

When you configure a Version 3 Local mode share or any Version 2 share (except a Stand-Alone share in which you do not specify a remote path to a directory on the origin-file server), a text file (._rbt_share_lock. txt) keeps track of which Steelhead appliance owns the share created on the origin-file server. Do not remove this file. If you remove the ._rbt_share_lock. txt file on the origin-file server, PFS does not function properly. (Version 3 Broadcast and Stand-Alone shares do not create this file.)

For details on PFS and when to enable it, see the *Riverbed Deployment Guide*.

Enabling PFS support is *optional*.

To enable PFS on the client-side Steelhead appliance

Note: For the server-side Steelhead appliance, you need only verify that it is intercepting and optimizing connections. No configuration is required for the server-side Steelhead appliance.

1. Choose Configure > Branch Services > PFS Settings to display the PFS Settings page.

Figure 5-2. PFS Settings Page

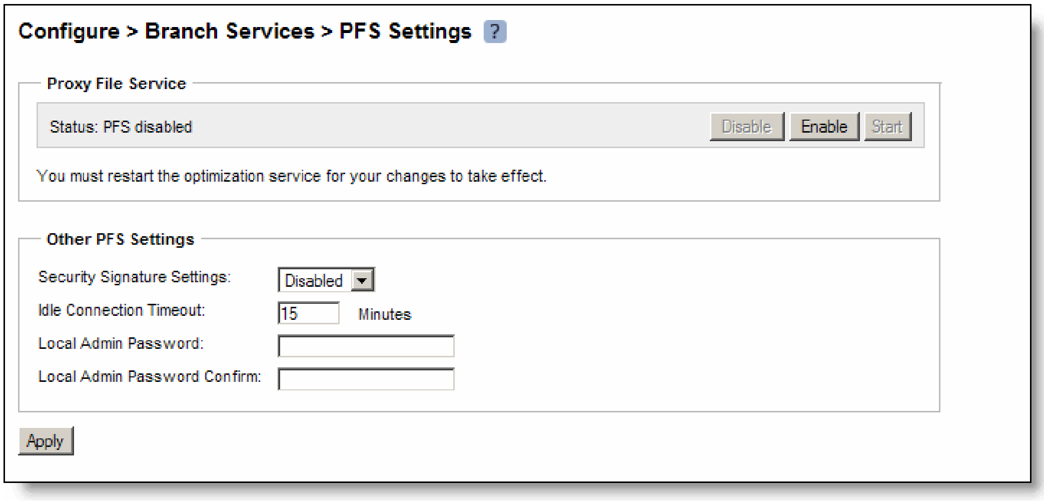

Status: PFS disabled

You must restart the optimization service for your changes to take effect.

Security Signature Settings: Disabled

Idle Connection Timeout: 15 Minutes

Local Admin Password:

Local Admin Password Confirm:

2. Under Proxy File Service, complete the configuration as described in the following table.

Control	Description
Enable/Disable	Enables or disables PFS to improve performance by deferring current requests or permitting certain operations even if WAN connectivity is disrupted. You must enable PFS before you configure it. After you enable PFS, the page expands to display Domain and Local Workgroup controls. **Note:** You must restart the Steelhead appliance each time you enable or disable PFS.
Start	Starts the PFS service.
Stop	Appears when PFS is started. Stops the PFS service.

3. To configure PFS in Local Workgroup mode, select Configure > Networking > Windows Domain.

   Under Domain/Local Workgroup Settings, click **Local Workgroup Settings** and complete the configuration as described in "Domain and Local Workgroup Settings" on page 305.

4. To configure PFS in Domain mode, select Configure > Networking > Windows Domain.

Under Domain/Local Workgroup Settings, click **Domain Settings** and complete the configuration as described in "Domain and Local Workgroup Settings" on page 305.

Note: For troubleshooting errors while joining a Windows Domain, see "Troubleshooting a Domain Join Failure" on page 309.

5. To configure other PFS settings, under Other PFS Settings, complete the configuration as described in the following table.

Control	Description
Security Signature Settings	Select one of the following options from the drop-down list: • **Enabled** - This setting supports any type of security signature setting requested by the client machine. • **Disabled** - This is the default setting. In this setting, PFS does not support clients with security signatures set to Required. • **Required** - In this setting, PFS supports clients with security signatures set to Enabled or Required. For details on SMB signing and security signatures, see the *Steelhead Appliance Installation and Configuration Guide*.
Idle Connection Timeout	Specify the number of minutes after which idle connections are timed-out. If there is no read or write activity on a mapped PFS share on a client machine, then the TCP connection times out according to the value set and the client has to re-map the share.
Local Admin Password/ Confirm	Specify and confirm the local administrator password.

6. Click **Apply** to apply your settings to the running configuration.

7. Click **Save** to save your settings permanently.

Related Topics

- "Adding PFS Shares" on page 166
- "Modifying General Host Settings" on page 39
- "Joining a Windows Domain or Workgroup" on page 305
- "Viewing PFS Share Logs" on page 449
- "Viewing PFS Data Reports" on page 450

Adding PFS Shares

You create and manage PFS shares in the Configure > Branch Services > PFS Shares page.

A share is the data volume exported by the proxy file server.

To set PFS share parameters

1. Choose Configure > Branch Services > PFS Shares to display the PFS Shares page.

Figure 5-3. PFS Shares Page

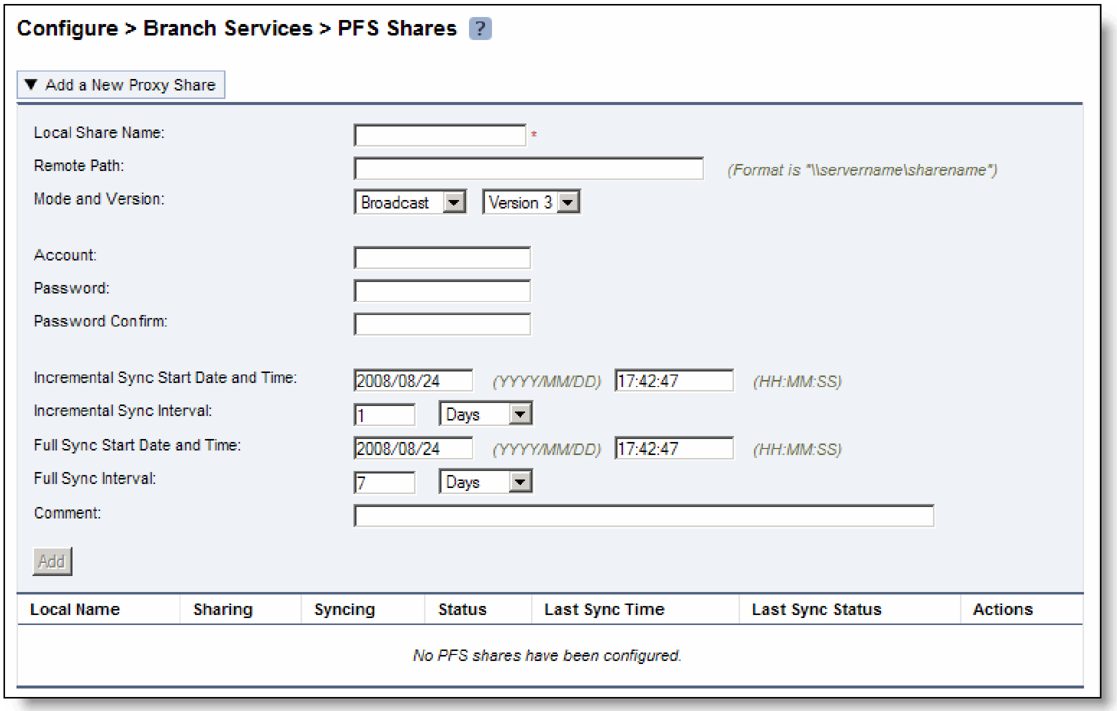

Configure > Branch Services > PFS Shares ?

▼ Add a New Proxy Share

Local Share Name:	*	
Remote Path:		(Format is "\\servername\sharename")
Mode and Version:	Broadcast ▼ Version 3 ▼	
Account:		
Password:		
Password Confirm:		
Incremental Sync Start Date and Time:	2008/08/24 (YYYY/MM/DD) 17:42:47	(HH:MM:SS)
Incremental Sync Interval:	1 Days ▼	
Full Sync Start Date and Time:	2008/08/24 (YYYY/MM/DD) 17:42:47	(HH:MM:SS)
Full Sync Interval:	7 Days ▼	
Comment:		

Add

Local Name	Sharing	Syncing	Status	Last Sync Time	Last Sync Status	Actions

No PFS shares have been configured.

2. Complete the configuration as described in the following table.

Control	Description
Add a New Proxy Share	Displays the controls to add a new proxy share.
Local Share Name	Specify a name for the share on the Steelhead appliance. This is the name by which users identify and map a share. The maximum length is 80 characters. **Important**: Do not use any characters in the share name other than letters, numbers, underscore, space, or backslash (directory separator). The PFS share and the origin-server share name cannot contain Unicode characters. The Management Console does not support Unicode characters.

Control	Description
Remote Path	Specify the path to the data on the origin-file server or the UNC path of a share you want to make available to PFS. You must use UNC for the remote path for Version 3 shares. For example, \\<origin-file-server>\<share> **Important**: Do not use any characters in the share name other than letters, numbers, underscore, space, or backslash (directory separator). The PFS share and the origin-server share name cannot contain Unicode characters. The Management Console does not support Unicode characters. **Notes**: • For a Version 2 share, the remote path is relative to the RCU server running on the origin Windows server. If the origin share is on the Windows server running RCU, the remote path is c:\folder. If the origin share is a shared folder on a computer different than the Windows server running RCU, the remote path is in UNC format. • For a Version 3 share, the remote path is always in UNC format.
Mode	Select one of the following options from the drop-down list: • **Broadcast** - Use Broadcast mode for environments seeking to broadcast a set of read-only files to many users at different sites. Broadcast mode quickly transmits a read-only copy of the files from the origin-file server to your remote offices. The PFS share on the Steelhead appliance contains read-only copies of files on the origin-file server. The PFS share is synchronized from the origin-file server according to parameters you specify. However, files deleted on the origin-file server are not deleted on the Steelhead appliance until you perform a full synchronization. Additionally, if you perform directory moves on the origin-file server (for example, move .\dir1\dir2 .\dir3\dir2) regularly, incremental synchronization does not reflect these directory changes. In this case, you must perform a full synchronization frequently to keep the PFS shares in synchronization with the origin-file server. • **Local** - Use Local mode for environments that need to efficiently and transparently copy data created at a remote site to a central data center, perhaps where tape archival resources are available to back up the data. Local mode enables read-write access at remote offices to update files on the origin-file server. After the PFS share on the Steelhead appliance receives the initial copy from the origin-file server, the PFS share copy of the data becomes the master copy. New data generated by clients is synchronized from the Steelhead appliance copy to the origin-file server based on parameters you specify when you configure the share. The folder on the origin-file server essentially becomes a back-up folder of the share on the Steelhead appliance. If you use Local mode, users must not directly write to the corresponding folder on the origin-file server. **Caution**: In Local mode, the Steelhead appliance copy of the data is the master copy; do not make changes to the shared files from the origin-file server while in Local mode. Changes are propagated from the remote office hosting the share to the origin-file server. **Important**: Riverbed recommends that you do not use Windows file shortcuts if you use PFS. For details, contact Riverbed Support at https://support.riverbed.com. • **Stand-Alone** - Use Stand-Alone mode for network environments where it is more effective to maintain a separate copy of files that are accessed locally by the clients at the remote site. The PFS share also creates more storage space. The PFS share on the Steelhead appliance is a one-time, working copy of data copied from the origin-file server. You can specify a remote path to a directory on the origin-file server, creating a copy at the branch office. Users at the branch office can read from or write to stand-alone shares but there is no synchronization back to the origin-file server since a stand-alone share is an initial and one-time only synchronization.

Control	Description
Version	Select one of the following software versions from the drop-down list. The controls change according to the version you choose. This value represents the version of the share that you want to create. • **Version 2** - Specify the server name and remote path for the share folder on the origin-file server. With Version 2, you must have the RCU service running on a Windows server—this can be the origin-file server or a separate server. **Important**: You must convert your Version 2 shares to Version 3 shares. Riverbed recommends you do not configure a mixed system of PFS shares; that is, Version 2 shares and Version 3 shares. For details, see "Upgrading Shares from Version 2 to Version 3" on page 171. **Important**: If you have shares that were created with RiOS v2.x.x of the Steelhead appliance software, the account that starts the RCU must have full permissions to the folder on the origin-file server. Also, the log-in user for the RCU server (which is used for Version 2 shares) and the share creation user for Version 3 shares must be a member of the Administrators group, either locally on the origin-file server (the local Administrators group) or globally in the domain (the Domain Administrator group). • **Version 3** - Specify the remote path, login, and password used to access the share folder on the origin-file server. With Version 3, the RCU runs on the Steelhead appliance—you do not need to install the RCU service on a Windows server. **Important**: Make sure the users are members of the Administrators group on the remote share server, either locally on the origin-file server (the local Administrators group) or globally in the domain (the Domain Administrator group).
Version 2	**Server Name** - Specify the server located in the data center which hosts the origin data volumes (folders).
	Port - Specify the port for the server.
Version 3	**Account** - Specify the fully qualified Windows login (including domain) to be used to access the shares on the origin-file server. For example, <Domain>\Administrator **Important**: Make sure the user is a member of the Administrators group on the remote share server, either locally on the origin-file server (the local Administrators group) or globally in the domain (the Domain Administrator group).
	Password/Password Confirm - Specify and confirm the password to be used to access the share on the origin-file server.
Incremental Sync Start Date and Time	Specify the date and time that you want incremental synchronization (updates) to start. The first synchronization, or the initial copy, retrieves data from the origin-file server and copies it to the local disk on the Steelhead appliance. Subsequent synchronizations are based on the synchronization interval. In incremental synchronization, only new and changed data is sent between the proxy-file server and the origin-file server. **Important**: For Local mode, changes are synchronized from the Steelhead appliance to the origin-file server; Broadcast mode changes are synchronized from the origin-file server to the Steelhead appliance. **Important:** For Broadcast mode: if you are performing directory moves regularly (for example, mv ./dir1/dir2 ./dir3/dir2), incremental synchronization does not reflect these directory changes. You must perform a full synchronization more frequently to keep the PFS shares in synchronization with the remote site.

Control	Description
Incremental Sync Interval	Specify a number and select a synchronization interval of Minutes, Hours, Days, or Disabled from the drop-down list.
Full Sync Start Date and Time	Specify the date and time that you want full synchronization (updates) to start. Use full synchronization if performance is not an issue. The first synchronization, or the initial copy, retrieves data from the origin-file server and copies it to the local disk on the Steelhead appliance. Subsequent synchronizations are based on the synchronization interval. In full synchronization, the system performs a full directory comparison and sends all changes since the last full synchronization between the proxy-file server and the origin-file server. **Important**: In Local mode, the system synchronizes changes from the Steelhead appliance to the origin-file server; in Broadcast mode the system synchronizes changes from the origin-file server to the Steelhead appliance. **Important**: In Broadcast mode, if you are performing directory moves regularly (for example, mv ./dir1/dir2 ./dir3/dir2), incremental synchronization does not reflect these directory changes. You must perform a full synchronization more frequently to keep the PFS shares in synchronization with the remote site.
Full Sync Interval	Specify a number and select a synchronization interval of Minutes, Hours, Days, or Disabled from the drop-down list.
Add	Adds the share.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to write your changes to disk.

   After you save your share settings the share is added to the Shares list.

Note: The PFS service must be started to add a share. To start the PFS service, click **Enable** in the Configure > Branch Services > PFS Settings page. For details, see "To enable PFS on the client-side Steelhead appliance" on page 165.

Enabling and Synchronizing Shares

After you have configured your PFS shares, you must perform the initial synchronization and enable your shares in the Configure > Branch Services > PFS Shares page.

When you perform the initial synchronization of the share, a copy of the data is downloaded from the origin-file server to the Steelhead appliance. The Steelhead appliance also configures the share for automatic synchronization according to the parameters you specified previously.

When you enable sharing for the first time, the share is made available to users for mounting. Users can map the mounted share using standard Windows mapping procedures. For example, map a network drive using the following format:

\\<appliance name or primary interface IP address>\<name of share>

To initialize and enable a share

1. Choose Configure > Branch Services > PFS Shares to display the PFS Shares page.

2. To enable synching, edit the share, select the Synching check box, and save the changes.

3. To make the share available to end users for mounting, edit the share, select the Sharing check box, and save the changes. End users are able to read the share by mapping to the mounted share (for example, \\Steelhead\share1).

Note: When performing the initial synchronization, or when changing large amounts of data, your bandwidth utilization and other graphs might show pockets of inactivity. This is by design.

Upgrading Shares from Version 2 to Version 3

When you upgrade to RiOS v3.x.x or later software, all of your existing shares run as Version 2 shares. You must manually upgrade your Version 2 shares.

In RiOS v3.x.x or later, you no longer need to install the RCU service on the server for synchronization purposes—all RCU functionality has been moved to the Steelhead appliance.

If you have legacy shares; that is, shares created with RiOS v2.x.x software, you must upgrade your Version 2 shares to Version 3 shares in the Management Console.

Important: You must convert your Version 2 shares to Version 3 shares. Riverbed recommends you do not configure a mixed system of PFS shares; that is, Version 2 shares and Version 3 shares.

If the remote path is in UNC format, you simply copy that path to the New Remote Path field.

To upgrade your share

1. Choose Configure > Branch Services > PFS Shares to display the PFS Shares page.

2. In the Shares list select the share name that you want to upgrade.

3. Click **Upgrade to Version 3 on Save**. The page refreshes with more controls necessary to upgrade the share.

4. Use the controls to complete the configuration, as described in the following table.

Control	Description
Account	Specify the fully qualified Windows login (including domain) to be used to access the share on the origin-file server. For example, <Domain>\Administrator. **Important**: Make sure the user is a member of the Administrators group on the remote share server, either locally on the origin-file server (the local Administrators group) or globally in the domain (the Domain Administrator group).

Control	Description
Password/Password Confirm	Specify and confirm the password to be used to access the share on the origin-file server.
New Remote Path	Specify the remote path of the origin-file server where the share resides. You must use UNC for the remote path for Version 3 shares. For example, \\<origin-file-server>\<share> **Important**: Do not use any characters in the share name other than letters, numbers, underscore, space, or backslash (directory separator). The PFS share and the origin-server share name cannot contain Unicode characters. The Management Console does not support Unicode characters.

5. Click **Apply** to apply your changes to the running configuration.

6. Click **Save** to write your changes to disk.

Modifying Share Settings

You can modify your share settings in the Configure > Branch Services > PFS Shares page.

To modify share settings

1. Choose Configure > Branch Services > PFS Shares to display the PFS Shares page.

2. In the Shares list select the share name that you want to modify to expand the page.

3. Under Edit Share, use the controls to modify the configuration as described in the following table.

Control	Description
Mode	Select one of the following options from the drop-down list: • **Broadcast** - Use Broadcast mode for environments seeking to broadcast a set of read-only files to many users at different sites. Broadcast mode quickly transmits a read-only copy of the files from the origin-file server to your remote offices. The PFS share on the Steelhead appliance contains read-only copies of files on the origin-file server. The PFS share is synchronized from the origin-file server according to parameters you specify. However, files deleted on the origin-file server are not deleted on the Steelhead appliance until you perform a full synchronization. Additionally, if you perform directory moves on the origin-file server regularly (for example, move .\dir1\dir2 .\dir3\dir2), incremental synchronization does not reflect these directory changes. In this case, you must perform a full synchronization frequently to keep the PFS shares in synchronization with the origin-file server. • **Local** - Use Local mode for environments that need to efficiently and transparently copy data created at a remote site to a central data center, perhaps where tape archival resources are available to back up the data. Local mode enables read-write access at remote offices to update files on the origin-file server. After the PFS share on the Steelhead appliance receives the initial copy from the origin-file server, the PFS share copy of the data becomes the master copy. New data generated by clients is synchronized from the Steelhead appliance copy to the origin-file server based on parameters you specify when you configure the share. The folder on the origin-file server essentially becomes a back-up folder of the share on the Steelhead appliance. If you use Local mode, users must not directly write to the corresponding folder on the origin-file server. **Caution**: In Local mode, the Steelhead appliance copy of the data is the master copy; do not make changes to the shared files from the origin-file server while in Local mode. The system propagates the changes from the remote office hosting the share to the origin-file server. **Important**: Riverbed recommends that you do not use Windows file shortcuts if you use PFS. For details, contact Riverbed Support at https://support.riverbed.com. • **Stand-Alone** - Use Stand-Alone mode for network environments where it is more effective to maintain a separate copy of files that are accessed locally by the clients at the remote site. The PFS share also creates more storage space. The PFS share on the Steelhead appliance is a one-time, working copy of data copied from the origin-file server. You can specify a remote path to a directory on the origin-file server, creating a copy at the branch office. Users at the branch office can read from or write to stand-alone shares but there is no synchronization back to the origin-file server since a stand-alone share is an initial and one-time only synchronization.
Sharing	Enables sharing.
Syncing	Enables syncing.
	Port (Version 2) - Specify the port for the server.

Control	Description
	Account - (Version 3) Specify the fully qualified Windows login (including domain) to be used to access the share on the origin-file server. For example, <Domain>\Administrator. **Important**: Make sure the user is a member of the Administrators group on the remote share server, either locally on the origin-file server (the local Administrators group) or globally in the domain (the Domain Administrator group).
	Password/Password Confirm - Specify and confirm the password to be used to access the share on the origin-file server.
Incremental Sync Start Date and Time	Specify the date and time that you want incremental synchronization (updates) to start. The first synchronization, or the initial copy, retrieves data from the origin-file server and copies it to the local disk on the Steelhead appliance. The system bases subsequent synchronizations on the synchronization interval. **Important**: In Local mode, the system synchronizes changes from the Steelhead appliance to the origin-file server; In Broadcast mode the system synchronizes changes from the origin-file server to the Steelhead appliance.
Incremental Sync Interval	Specify a number and select a synchronization interval of Minutes, Hours, Days, or Disabled from the drop-down list.
Full Sync Start Date and Time	Specify the date and time that you want full synchronization (updates) to start. Use full synchronization if performance is not an issue. The first synchronization, or the initial copy, retrieves data from origin-file server and copies it to the local disk on the Steelhead appliance. The system bases subsequent synchronizations on the synchronization interval. **Important**: In Local mode, the system synchronizes changes from the Steelhead appliance to the file server; In Broadcast mode the system synchronizes changes from the origin-file server to the Steelhead appliance.
Full Sync Interval	Specify a number and select a synchronization interval of Minutes, Hours, Days, or Disabled from the drop-down list.

4. Click **Apply** to apply your changes to the running configuration.

5. Click **Save** to write your changes to disk.

Performing Manual Actions on Shares

You can verify a share, perform a full synchronization, cancel an operation, or delete a share in the Shares list. The shares list appears in the PFS Shares page.

To perform manual actions on shares

1. Choose Configure > Branch Services > PFS Shares to display the PFS Shares page.

2. Select one of the following actions for the share, as described in the following table.

Control	Description
Actions	Select one of the following actions from the drop-down list: • **Start Verify** - Generates a list of the differences between the share on the Steelhead appliance and the origin-file server. A list of differences is available in the PFS Shares report. • **Manual Sync** - Allows you to immediately synchronize the share and its corresponding remote share on the origin-file server. • **Cancel Operation** - Cancels the current share action. • **Delete Share** - Deletes the selected share.

Related Topics

- "Joining a Windows Domain or Workgroup" on page 305
- "Modifying General Host Settings" on page 39
- "Viewing PFS Share Reports" on page 448
- "Viewing PFS Share Logs" on page 449

Enabling DNS Caching

You configure a local DNS name server for caching in the Configure > Branch Services > Caching DNS page. By default, the DNS cache is disabled.

A DNS name server resolves hostnames to IP addresses and stores them locally in a single Steelhead appliance. Any time your browser requests a URL, it first looks in the local cache to see if it is there before querying the external name server. If it finds the resolved URL locally, it uses that IP address.

This is a non-transparent DNS caching service. Any client machine must point to the client-side Steelhead appliance as their DNS server.

Hosting the DNS name server function provides:

- Improved performance for applications by saving the round trips previously needed to resolve names. Whenever the name server receives address information for another host or domain, it stores that information for a specified period of time. That way, if it receives another name resolution request for that host or domain, the name server has the address information ready, and does not need to send another request across the WAN.
- Improved performance for services by saving round trips previously required for updates.
- Continuous DNS service locally when the WAN is disconnected, with no local administration needed, eliminating the need for DNS servers at branch offices.

A cache holds the resolved address entries information. For information on DNS Statistics, see "Viewing DNS Cache Hits" on page 453.

To enable the DNS name server

1. Choose Configure > Branch Services > Caching DNS to display the Caching DNS page.

Figure 5-4. Caching DNS Page

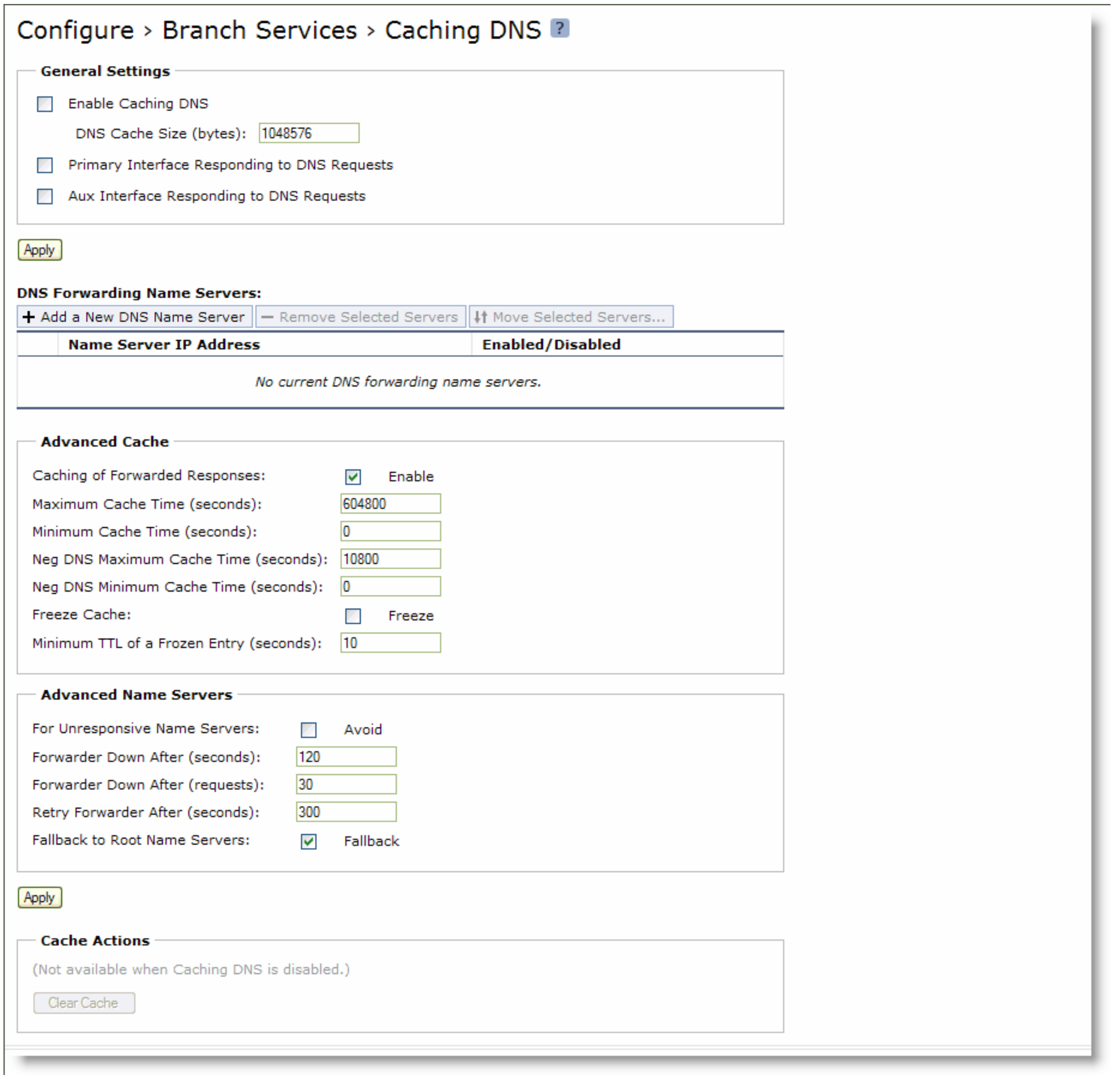

Configure > Branch Services > Caching DNS ?

General Settings

□ Enable Caching DNS

DNS Cache Size (bytes): 1048576

□ Primary Interface Responding to DNS Requests

□ Aux Interface Responding to DNS Requests

Apply

DNS Forwarding Name Servers:

+ Add a New DNS Name Server | − Remove Selected Servers | ⇅ Move Selected Servers...

	Name Server IP Address	Enabled/Disabled
	No current DNS forwarding name servers.	

Advanced Cache

Caching of Forwarded Responses: ☑ Enable

Maximum Cache Time (seconds): 604800

Minimum Cache Time (seconds): 0

Neg DNS Maximum Cache Time (seconds): 10800

Neg DNS Minimum Cache Time (seconds): 0

Freeze Cache: □ Freeze

Minimum TTL of a Frozen Entry (seconds): 10

Advanced Name Servers

For Unresponsive Name Servers: □ Avoid

Forwarder Down After (seconds): 120

Forwarder Down After (requests): 30

Retry Forwarder After (seconds): 300

Fallback to Root Name Servers: ☑ Fallback

Apply

Cache Actions

(Not available when Caching DNS is disabled.)

Clear Cache

2. Under General Settings, complete the configuration as described in the following table.

Control	Description
Enable Caching DNS	**Enabled** - Forwards name resolution requests to a DNS name server, then stores the address information locally in the Steelhead appliance. By default, the requests go to the root name server, unless you specify another name server. **Disabled** - Stops the Steelhead appliance from acting as the DNS name server.
DNS Cache Size (bytes)	Specifies the cache size, in bytes. The default value is 1048576. The range is 524288 to 2097152.

Control	Description
Primary Interface Responding to DNS Requests	**Enabled** - Enables the name server to listen for name resolution requests on the primary interface. **Disabled** - Stops the name server from using the primary interface.
Aux Interface Responding to DNS Requests	**Enabled** - Enables the name server to listen for name resolution requests on the auxiliary interface. **Disabled** - Stops the name server from using the auxiliary interface.

Tip: To move the position of a name server in the DNS Forwarding name server list, select the name server IP address and click **Move Selected Servers**.

Tip: To remove a name server from the list, select the name server IP address and click **Remove Selected Servers**. You cannot remove the last name server in the list unless the root name server is enabled.

3. Click **Apply** to apply your changes to the running configuration.

4. Under DNS Forwarding Name Servers, complete the configuration as described in the following table.

Control	Description
Add a New DNS Name Server	Displays the controls to add a DNS name server to which the Steelhead appliance forwards requests to cache responses. By default, the Steelhead appliance only forwards requests to the Internet root name servers when you enable caching DNS without specifying any name servers to forward requests to. You can add multiple name servers to use; the Steelhead appliance uses failover to these if one name server is not responding.
Name Server IP Address	Specify an IP address for the name server.
Position	Specify the order in which the name servers are queried (when using more than one). If the first name server, or *forwarder*, does not respond, the Steelhead appliance queries each remaining forwarder in sequence until it receives an answer or until it exhausts the list.
Add	Adds the name server.
Remove Selected	Select the check box next to the name and click **Remove Selected Servers**.
Move Selected	Select the check box next to the name and click **Move Selected Servers**.

5. Under Advanced Cache, complete the configuration as described in the following table.

Control	Description
Caching of Forwarded Responses	Enables the cache that holds the resolved address entries. The cache is enabled by default; however nothing is actually cached until you select the General Setting Enable Caching DNS.
Maximum Cache Time (seconds)	Specify the maximum number of seconds the name server stores the address information. The default setting is one week (604,800 seconds). The minimum is 2 seconds and the maximum is thirty days (2,592,000 seconds). You can adjust this setting to reflect how long the cached addresses remain up-to-date and valid. **Note:** Changes to this setting affect new address information and do not change responses already in the cache.
Minimum Cache Time (seconds)	Specify the minimum number of seconds that the name server stores the address entries. The default value is 0. The maximum value is the current value of Maximum Cache Time. Typically there is no need to adjust this setting. **Note:** Changes to this setting affect new responses and do not change any responses already in the cache.
Neg DNS Maximum Cache Time (seconds)	Specify the maximum number of seconds that an unresolved negative address is cached. The valid range is from two seconds to thirty days (2,592,000 seconds). The default value is 10,800 seconds. A negative entry occurs when a DNS request fails and the address remains unresolved. When a negative entry is in the cache, the appliance does not request it again until the cache expires, the maximum cache time is reached, or the cache is cleared.
Neg DNS Minimum Cache Time (seconds)	Specify the TTL for a negative entry, which is always this value or above, even if the server returns a smaller TTL value. For example, when this value is set to 300 seconds and the client queries aksdfjh.com, the DNS service returns a negative answer with a TTL of 100 seconds, but the DNS cache stores the entry as having a TTL of 300 seconds. The default value is 0, which specifies that the Steelhead appliance still caches negative responses; it does not place a lower bound on what the TTL value for the entry can be.
Freeze Cache	Freezes the cache contents. When the cache is frozen, entries do not automatically expire from the cache. They are still returned in response to DNS queries. This is useful to keep local services available when the WAN is disconnected. By default, this setting is disabled. **Note:** When the cache is frozen and full, entries can still be pushed out of the cache by newer entries.
Minimum TTL of a Frozen Entry (seconds)	Specify the minimum TTL in seconds that a response from a frozen cache has when sent to a branch office client. The default value is 10. For example, suppose this value is set to 60 seconds. At the time the cache is frozen, the cache entry for riverbed.com has a TTL of 300 seconds. For subsequent client requests for riverbed.com, the service responds with a TTL of 300 seconds minus however much time has elapsed since the cache freeze. After 240 seconds have elapsed, the service responds to all subsequent requests with a TTL of 60 seconds regardless of how much time elapses, until the cache is unfrozen.

6. Under Advanced Name Servers, complete the configuration as described in the following table.

Control	Description
For Unresponsive Name Servers	Detects when one of the name servers is not responding and sends requests to a responsive name server instead.
Forwarder Down After (seconds)	Specify how many seconds can pass without a response from a name server until the appliance considers it unresponsive. The default value is 120. When the name server receives a request but does not respond within this time *and* does not respond after the specified number of failed requests, the appliance determines that it is down. It then queries each remaining forwarder in sequence until it receives an answer or it exhausts the list. When the list is exhausted and the request is still unresolved, you can specify that the Steelhead appliance try the root name server.
Forwarder Down After (requests)	Specify how many requests a name server can ignore before the appliance considers it unresponsive. The default value is 30. When the name server does not respond to this many requests *and* does not respond within the specified amount of time, the appliance determines that it is down. It then queries each remaining forwarder in sequence until it receives an answer or it exhausts the list. When the list is exhausted and the request is still unresolved, you can specify that the Steelhead appliance try the root name server.
Retry Forwarder After (seconds)	Specify the time limit, in seconds, that the appliance forwards the name resolution requests to name servers that are responding instead of name servers that are down. The appliance also sends a single query to name servers that are down using this time period. If they respond, the appliance considers them back up again. The default value is 300. The single query occurs at intervals of this value if the value is set to 300. A request is allowed to go to a forwarder considered down about every 300 seconds until it responds to one.
Fallback to Root Name Servers	Forwards the request to a root name server when all other name servers have not responded to a request. This is the default setting; either this option must be enabled or a server must be present. When the fallback to root name servers option is disabled, the Steelhead appliance only forwards a request to the forwarding name servers listed above. If it exhausts these name servers and does not get a response, it does not forward the request to a root name server and returns a server failure. **Note:** If the name servers used by the Steelhead appliance are internal name servers; that is, they can resolve hostnames that external name servers like the Internet DNS root servers cannot, you must disable this option. Otherwise, if the name servers all fail, the root name servers might inform the Steelhead appliance that a host visible only to internal name servers does not exist, might cache that response, and return it to clients until it expires. This prolongs the period of time until service comes back up after name servers are down.

7. Click **Apply** to apply your changes to the running configuration.

8. Click **Save** to save your settings permanently.

Clearing the Cache

To clear the cache

- Under Cache Actions, complete the configuration as described in the following table.

Control	Description
Clear Cache	Removes entries from the cache, even if it is frozen. All cached data expires. **Note:** A small amount of data remains in the cache for internal use only.

Related Topics

- "Configuring HTTP Optimization" on page 111
- "Viewing DNS Cache Hits" on page 453
- "Viewing DNS Cache Utilization" on page 455

Installing and Configuring RSP

You can install, start, stop, and restart Riverbed Services Platform (RSP) in the Configure > Branch Services > RSP Service page.

RSP offers branch-office-in-a-box services with the following benefits:

- A VMware-based virtualization platform provides the benefits of the most commonly deployed and advanced virtualization tool set.
- Support for running up to five different additional services simultaneously on a single Steelhead appliance, depending on the service and Steelhead appliance model.
- Support for more services and types of services. These include in-band packages that sit in-line with optimization such as the Universal Threat Management (UTM) security services, proxy solutions such as video or network monitoring services, and improved support for out-of-band packages like Windows Active Directory, DNS and DHCP management software, and print services.
- A comprehensive, integrated, user management interface that provides granular control of RSP including setup, reporting, and the definition of the data flow between services.

Note: RSP is supported on Models 250, 520, 550, 1020, 1050, 1520, 2020, 2050, 3020, 3520, 5050, 6050, and 7050.

CSH VSH Cloud Steelhead and Virtual Steelhead models do not support RSP.

Important: Do not configure both RSP and PFS on the same Steelhead appliance. Riverbed does not support this configuration because PFS has no disk boundaries and can overwrite the space allocated to RSP.

RSP uses VMware Server 2.0 as the virtualization platform. Both 32 and 64-bit versions of the RSP installation image are available.

Important: 64-bit guest VMs (such as Windows Server 2008 R2) are not supported on Steelhead appliance models 250, 550, and the 1U xx20s because these models do not incorporate Virtual Technology (VT) support. For details, see "Riverbed Services Platform 32-Bit and 64-Bit Support" on page 15.

Figure 5-5 shows the RSP setup on the Steelhead appliance.

Figure 5-5. The RSP Platform

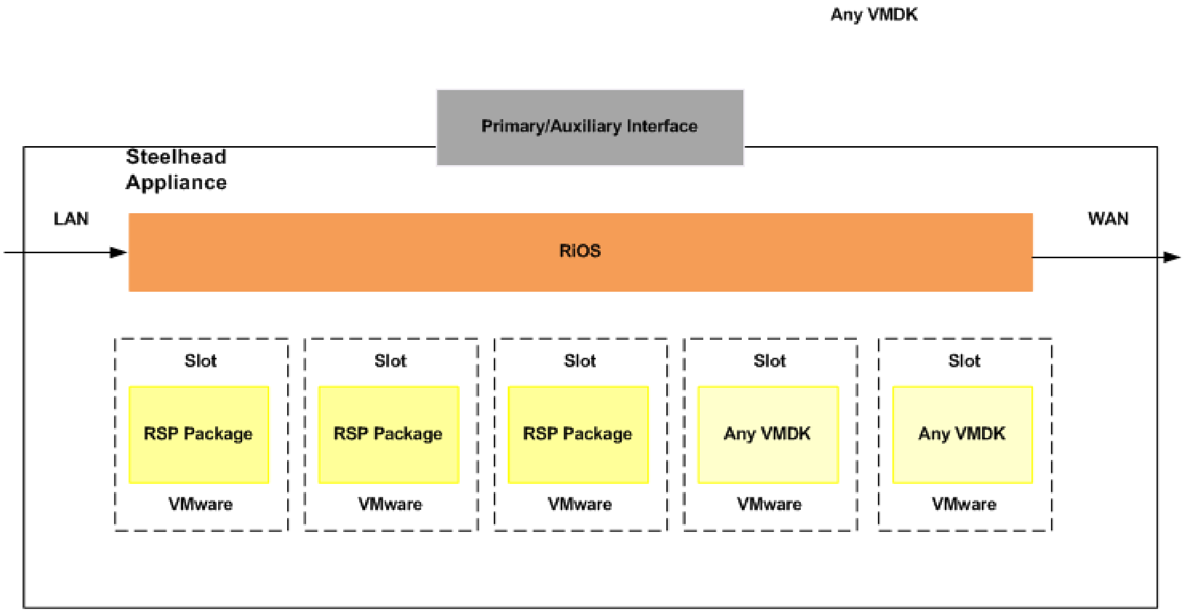

Note: For details on RSP, see the *RSP User's Guide*.

After installing the RSP installation image, you can add packages to run more services and applications. RSP includes configuration options that allow you to determine the data flow to and from a package, and the ability to chain packages together.

RSP Support for Virtual-In Path Deployments

RiOS v6.0 and later supports Virtual in-path RSP (VRSP). VRSP extends the Riverbed Services Platform to allow RSP to operate with a single connection to the network while presenting an in-path view of the network to the optimization and virtual services running on the platform.

VRSP supports the following types of connections:

- **Virtual-in-Path connections** - connections going from client to the server, neither of which is running on the RSP platform. Routers redirect such connections to VRSP using one of the following mechanisms:
  - L2 redirect.
  - Policy-Based Routing (PBR) allows you to define policies to route packets instead of relying on routing protocols. You enable PBR to redirect traffic that you want optimized by a Steelhead appliance that is not in the direct physical path between the client and server.
  - Generic Routing Encapsulation (GRE) with Web Cache Communication Protocol (WCCP).
- **Out-of-Path connections** - connections that originate from a system outside of RSP to a service running on the RSP platform (inbound) or connections that originate from an RSP service to an external system (outbound).

- **Redirected Out-of-Path connections** - this is a variation of Out-of-Path connections where traffic to RSP is redirected to by a router using L2, PBR, or GRE with WCCP.

VRSP, being an extension of RSP, inherits configuration/management mechanisms from RSP. The same facilities are still present and operating as they were in RSP, with the following exceptions:

- VNI types
- VNI rules

For details on configuring VRSP, see “Configuring Subnet Side Rules” on page 260 and “To add a DNAT rule to a virtual in-path VNI” on page 211, and “Configuring WCCP” on page 312.

Basic Steps for Installing and Configuring RSP

This section provides an overview of the basic steps to install and configure RSP, followed by detailed procedures.

Important: Follow these steps in the order given to avoid traffic disruption.

Task	Reference
1. Verify that your Steelhead appliance supports RSP requirements.	*RSP User's Guide*
2. Download and install the RSP image, which contains the VMware server and the RSP service. The Steelhead appliance RiOS image does not include the RSP image. You must install the RSP image separately. RSP is pre-installed on newly manufactured Steelhead appliances if you ordered RSP.	“Installing the RSP Image” on page 183
3. Start the RSP service on the Steelhead appliance.	“Installing and Configuring RSP” on page 180
4. Obtain an RSP package: • Use an existing package from Riverbed, a third-party vendor, or from within your organization. • Create your own package. **Important:** The package files can be quite large and can take up several GBs of disk space.	“Adding RSP Packages” on page 184 For details about creating an RSP Package for your application or service, see the *RSP User's Guide.*
5. Install a package in a slot.	“Installing a Package in a Slot” on page 186
6. Enable the slot.	“Enabling, Disabling, and Restarting Slots” on page 189
7. Specify VM settings, specify watchdog settings, create and manage virtual disks, manage Virtual Network Interfaces (VNIs), and perform operations such as uninstalling a slot or restoring RSP backup.	“Enabling, Disabling, and Restarting Slots” on page 189
8. Place the slotted package optimization VNI into the data flow. **Note**: This step is not required for out-of-band packages.	“Viewing RSP Statistics Reports” on page 457

Task	Reference
9. Create and add data flow rules to the VNI. **Note**: This step is not required if you use the default rules for the package.	"Adding Rules to an Optimization VNI" on page 208
10. Configure virtual in-path support for RSP.	"Configuring Subnet Side Rules" on page 260
11. Verify your configuration.	"Viewing RSP Statistics Reports" on page 457

Installing the RSP Image

This section describes the prerequisites and the basic steps to install the RSP installation image using the Management Console.

Prerequisites and Tips

- RSP requires at least 2 GB of additional memory on the Steelhead appliance.
- You must have role-based permission for RSP to install RSP. For details, see "Viewing Permissions" on page 358.
- Before installing a new RSP image, you must stop the RSP service currently running on the Steelhead appliance.
- If you have previously installed RSP for RiOS v5.0.x, you must reinstall the RSP image for RiOS v5.5 and later. RSP for RiOS v5.0.x is not compatible with RSP for RiOS v5.5.
- Installing a new RSP image replaces the previously installed image (the RSP packages and installed VMs are unmodified).
- You cannot enable RSP when a Management In-Path (MIP) interface is enabled and vice versa. For details, see "Configuring a Management In-Path Interface" on page 54.
- For details on Steelhead appliance RSP support for guest operating systems, see the product specification sheets at: http://www.riverbed.com/products/appliances/

Note: When you upgrade from Steelhead models 5050M to 5050H, 1050L to 1050_LR, 1050_M to 1050_MR, and 1050H to 1050HR, the upgrade process deletes the RSP partition and data. Before you install the new image, back up the RSP slots and after you install the image restore the slots from the backup. For information on backing up, see the *RSP User's Guide*.

To install the RSP image

1. Choose Configure > Branch Services > RSP Service to display the RSP Service page.

Figure 5-6. RSP Service Page

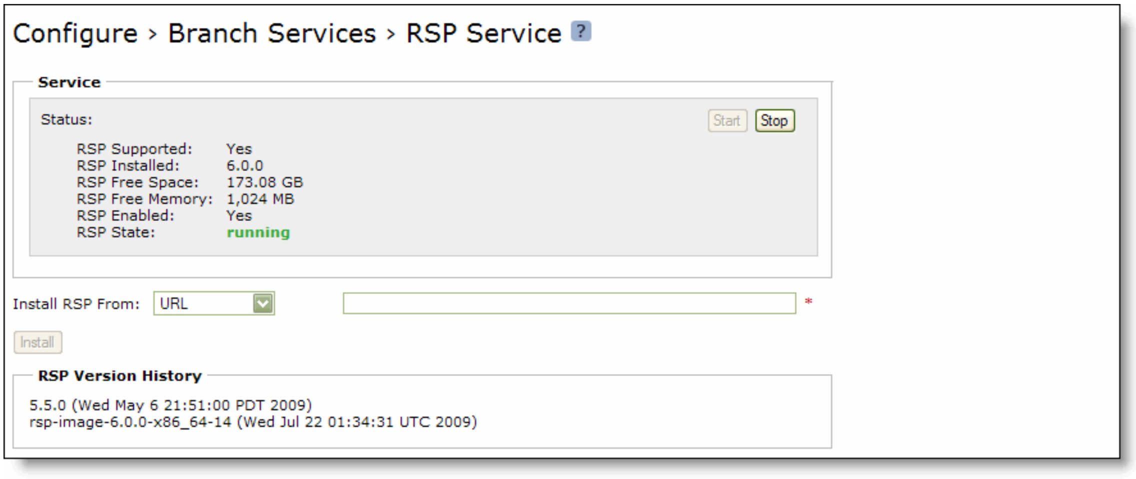

Configure > Branch Services > RSP Service

Service

Status:

```
RSP Supported:     Yes
RSP Installed:     6.0.0
RSP Free Space:    173.08 GB
RSP Free Memory:   1,024 MB
RSP Enabled:       Yes
RSP State:         running
```

Install RSP From: URL

RSP Version History

5.5.0 (Wed May 6 21:51:00 PDT 2009)
rsp-image-6.0.0-x86_64-14 (Wed Jul 22 01:34:31 UTC 2009)

2. Select the Install **RSP** From drop-down menu and complete the configuration as described in the following table.

Control	Description
Install RSP From:	Select one of the following options from the drop-down list: • **URL** - Select this option and specify the URL to the RSP image in the text box next to the drop-down menu. • **Local File** - Select this option and specify the pathname in the text box next to the drop-down menu, or click **Browse** to go to the local file directory. • **Existing Image** - Select an image that you uploaded to the Steelhead appliance using the Riverbed CLI. For details, see the *Riverbed Command-Line Interface Reference Manual*.

3. Click **Install** to download and install RSP on your system.

4. Click **Start** to start the RSP service.

Adding RSP Packages

You add RSP packages in the Configure > Branch Services > RSP Packages page.

The RSP package contains the service or applications in the virtual machine and RSP-specific configuration files. RSP packages contain a service or application running in a virtual machine. RSP packages also contain Riverbed configuration files specifying the package VNIs, and how the package interacts with the Steelhead appliance, a .vmx file, and one or more .vmdk files. Do not open or modify any of the files in the package. The package files can be quite large, and can take up several GBs of disk space.

Before installing a package, you must install the RSP installation image. For details, "Installing the RSP Image" on page 183. After installing the RSP image, you must start the RSP service and then you can download and install packages.

A package can be:

- a VM created by a third-party vendor that also contains configuration files specific to the Riverbed RSP implementation.
- a VM created by Riverbed.
- a VM created internally within your organization.

You can download any number of packages to the Steelhead appliance, but you can only run up to five packages at a time. The exact number depends on the package size, the amount of resources available, and your Steelhead appliance model.

For the Steelhead xx50 model appliances, Riverbed provides an option to purchase fully-licensed OEM Microsoft Windows 2008 Standard package. If you purchase this option, the RSP package is already installed on the Steelhead appliance with a Certificate of Authenticity (COA) sticker containing your license key attached to the appliance. To obtain the package, simply browse for it and the file ms_windows_2008_standard<version>.pkg appears. The package name varies based on the version installed.

Note: For details on RSP, see the *RSP User's Guide*.

When you install an RSP package you must select an RSP *slot*. A slot is a directory on disk. When you install a package into a slot, the system unpacks the VM into the directory. When you uninstall a VM from a slot, the system removes the VM and deletes the files from the slot.

For details about creating an RSP package for a Windows Server, see the *RSP Package Creation Guide*.

Note: If an out-of-band connection uses WCCP or PBR redirected to the Steelhead, configure the redirected rule and the NAT so that the redirected packets reach the VM (even if the packet destination is the VM IP address).

To add an RSP package

1. Choose Configure > Branch Services > RSP Packages to display the RSP Packages page.

Figure 5-7. RSP Packages Page

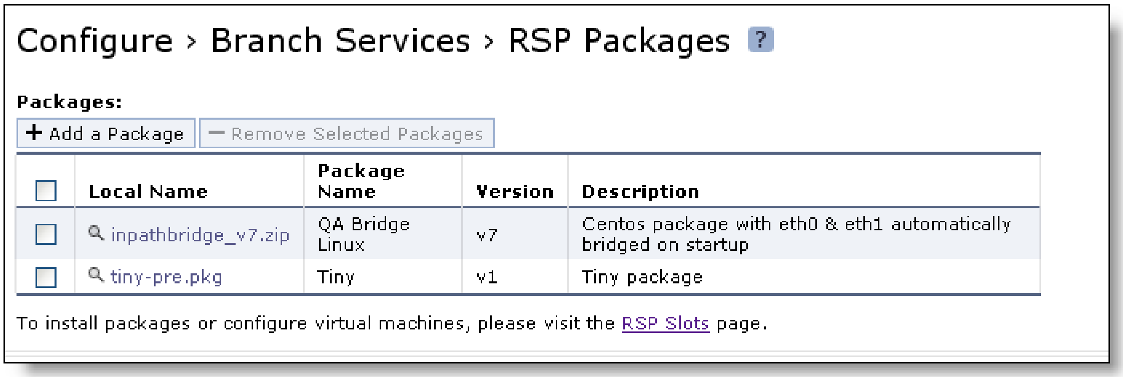

Configure > Branch Services > RSP Packages ?
Packages:
+ Add a Package
− Remove Selected Packages

□	Local Name	Package Name	Version	Description
□	inpathbridge_v7.zip	QA Bridge Linux	v7	Centos package with eth0 & eth1 automatically bridged on startup
□	tiny-pre.pkg	Tiny	v1	Tiny package

To install packages or configure virtual machines, please visit the RSP Slots page.

2. Under Packages, complete the configuration as described in the following table.

Control	Description
Add a Package	Displays the controls to obtain an RSP package.
Name	Optionally, specify a descriptive name for the package (up to eight characters).
From URL	Select this option and specify the URL to the RSP package. For third-party packages, get the URL directly from the vendor.
From Local File (for packages less than 2 GB in size)	Select this option and specify the path or click **Browse** to go to the local file directory. **Important:** You cannot use this option to upload a package file that is larger than 2 GB. If the file is larger than 2 GB you can use SCP or FTP to transfer it using the CLI. Alternatively, you can push the package to the Steelhead using scp and upload the package to the /rsp/packages directory on the Steelhead. For example: `scp windows.pkg <Steelhead username>:<Steelhead password>@<Steelhead IP>:/rsp/packages`
Add Package	Downloads the RSP package to your system.
Remove Selected Packages	Select the check box next to the name and click **Remove Selected Packages**.

Tip: To view the VNIs and watchdog settings for a package, select the local name of the package.

Installing a Package in a Slot

You can install one package per slot. By default, the RSP slots are numbered 1 to 5. However, Riverbed recommends changing slot names to meaningful, descriptive terms because RSP-enabled Steelhead appliances may be remotely configured by the Central Management Console (CMC). Meaningful names reduce the potential for confusion. Riverbed also recommends you give slots with identical VMs identical names to facilitate batch management.

Note: Installing a package into a slot does not affect the RSP package file itself, as it is a copy of the files within the package that are installed into the slot. The package remains unmodified and can be installed into other slots as needed.

To install an RSP package in a slot

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

To install an RSP package in an occupied slot, you must first uninstall the package in that slot. Installing a package into a slot and uninstalling that particular slot affects only the slot directory, not the package itself.

Figure 5-8. Installing a Package in a Slot

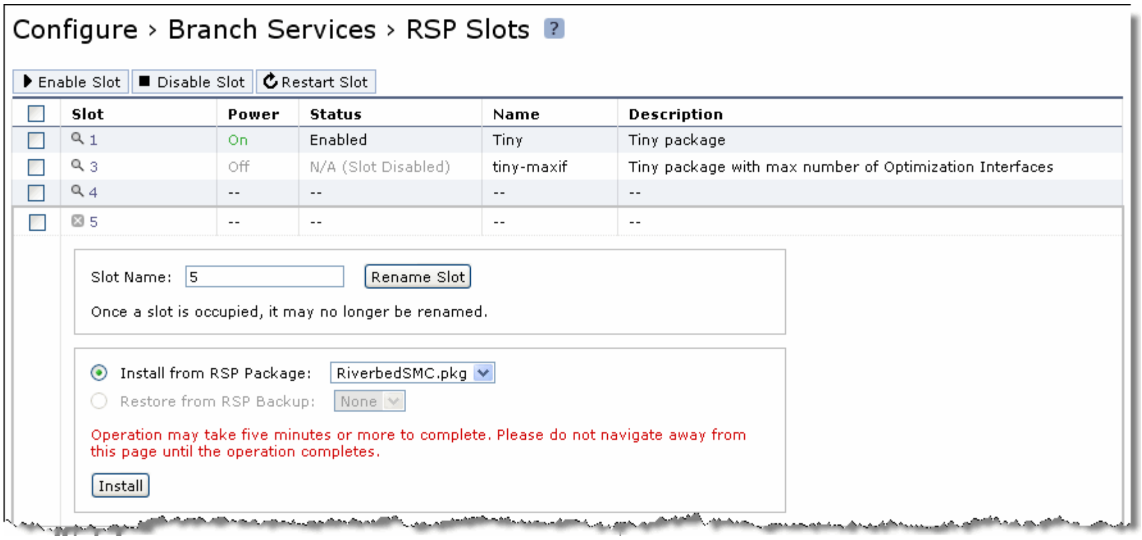

□	Slot	Power	Status	Name	Description
□	1	On	Enabled	Tiny	Tiny package
□	3	Off	N/A (Slot Disabled)	tiny-maxif	Tiny package with max number of Optimization Interfaces
□	4	--	--	--	--
□	5	--	--	--	--

Slot Name: 5

Once a slot is occupied, it may no longer be renamed.

● Install from RSP Package: RiverbedSMC.pkg

○ Restore from RSP Backup: None

Operation may take five minutes or more to complete. Please do not navigate away from this page until the operation completes.

2. Under Slot, select the name of an empty slot and complete the configuration as described in the following table.

Control	Description
Slot Name	Optionally, specify a unique descriptive name for the slot (up to eight characters). **Note:** The slot must be empty before you can change its name.
Install From RSP Package	Select the package name from the drop-down list of packages you added.
Restore from RSP Backup	Select the package name from the drop-down list of packages in the backup.
Install	Installs the package to the slot and updates the configuration. **Note**: This step can take five minutes or longer to complete.

3. Click **Save** to save your settings permanently.

Related Topics

- "Enabling, Disabling, and Restarting Slots" on page 189
- "Specifying VM Settings" on page 190
- "Specifying Watchdog Settings" on page 191
- "Managing Virtual Disks" on page 193
- "Viewing RSP Statistics Reports" on page 457
- "Capturing and Uploading TCP Dumps" on page 482

Viewing Slot Status

You can view slot status in the RSP Slots page.

To view slot status

- Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

Figure 5-9. RSP Slots Page

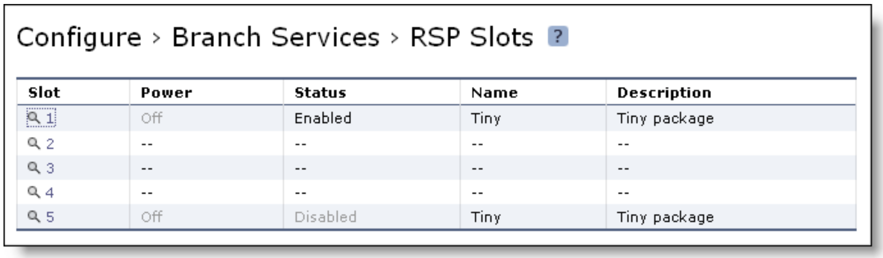

Configure > Branch Services > RSP Slots ?

Slot	Power	Status	Name	Description
1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--
4	--	--	--	--
5	Off	Disabled	Tiny	Tiny package

The list at the top of this page displays the following information:

Control	Description
Slot	Displays the slot name.
Power	Displays one of the following power states: • **On** - The VM is powered on. • **Off** - The VM is powered off. • **N/A** - RSP service is not running.
Status	Displays one of the following status messages: • **Enabled** - The slot is enabled but the watchdog is not monitoring it. • **Disabled** - The slot is disabled but the watchdog is not monitoring it. This status can also indicate that the RSP service is not running. • **Healthy** - The slot is enabled and watchdog is monitoring it. • **Blocked** - The slot is enabled, and the watchdog has triggered block-on-failure mode. • **Bypass** - The slot is enabled, and the watchdog has triggered bypass-on-failure mode. • **Startup Grace Period** - The slot is enabled, and is still in the power-on grace period (watchdog has not yet begun monitoring). • **N/A (RSP service not running)** - You have not started the RSP service. For details, see "Installing the RSP Image" on page 183.
Name	Displays the name of the package. You specify the name in the package creator when you create the package. For details, see the *RSP User's Guide*
Description	Displays a description about the package in the slot. You specify the description in the package creator when you create the package. For details, see the *RSP User's Guide*.

Enabling, Disabling, and Restarting Slots

This section describes how to enable, disable, or restart a slot in the RSP Slots page. It includes the following tasks:

- "Specifying VM Settings" on page 190
- "Specifying Watchdog Settings" on page 191
- "Managing Virtual Disks" on page 193
- "Managing Virtual Network Interfaces" on page 198
- "Performing RSP Operations" on page 199

To enable, disable, or restart a slot

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

Figure 5-10. RSP Slots Page - Slot Operations

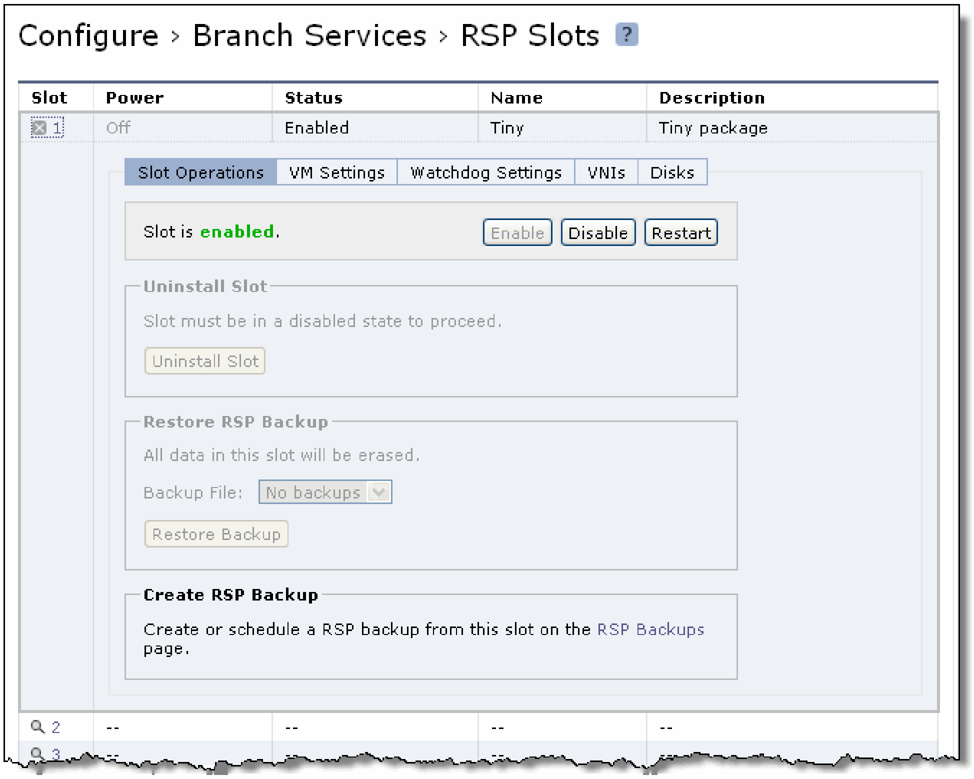

Configure > Branch Services > RSP Slots ?

Slot	Power	Status	Name	Description
☒ 1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--

Slot is **enabled**.

Uninstall Slot

Slot must be in a disabled state to proceed.

Restore RSP Backup

All data in this slot will be erased.

Backup File: No backups

Create RSP Backup

Create or schedule a RSP backup from this slot on the RSP Backups page.

2. Click the magnifying glass next to the slot in the Slot column.

3. To enable or disable a slot, complete the configuration as described in the following table.

Control	Description
Enable	Enables the slot (makes it active).
Disable	Disables the slot (makes it inactive).
Restart	Shuts down and restarts the slot.

4. Click **Save** to save your settings permanently.

Specifying VM Settings

This section describes how to change the memory footprint for a VM in a slot and how to open a virtual machine (VM) console for an RSP package.

To specify VM settings

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

Figure 5-11. Slots Page - VM Settings

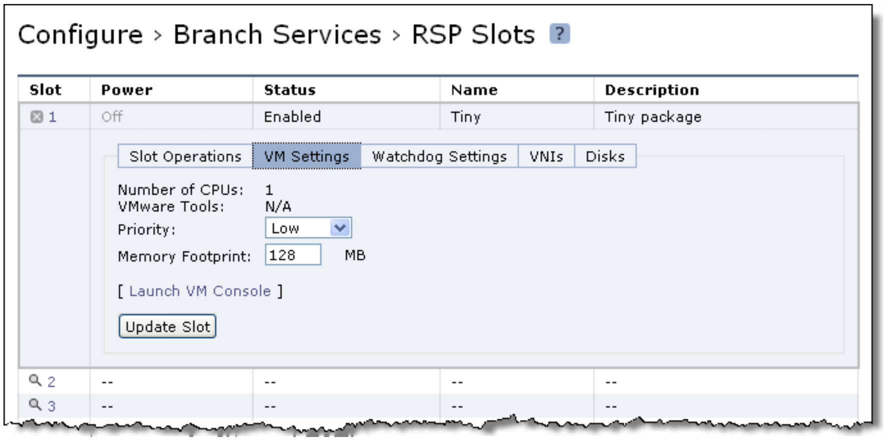

Configure > Branch Services > RSP Slots ?

Slot	Power	Status	Name	Description
1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--

Slot Operations | VM Settings | Watchdog Settings | VNIs | Disks

Number of CPUs: 1
VMware Tools: N/A
Priority: Low
Memory Footprint: 128 MB

[Launch VM Console]

Update Slot

2. Select the slot name in the Slot column.

3. Select VM Settings and complete the configuration as described in the following table.

Control	Description
Priority	Specify the priority level for the VM processing: • **Low** - Specifies low priority relative to other VMs. • **Normal** - Specifies normal priority relative to other VMs. • **High** - Specifies a high priority relative to other VMs. **Note:** The default priority setting is Normal. In the event of CPU contention, CPU resources are allocated to the VMs according to the priority specified. VMs with the same priority level receive equal access to the CPU.
Memory Footprint	Specify how many megabytes of memory to allocate to the VM. The value must be a multiple of 4. **Note:** Although you can change the memory footprint for a VM in the slot using this option, Riverbed recommends using the default value provided by the vendor.
Launch VM Console	Opens the VMware console. You can log in to the VM console and open the VM connection to the package.
Update Slot	Updates the slot information.

4. Click **Restart Slot** to restart the slot and the VM.

Specifying Watchdog Settings

You can configure an RSP watchdog using ping monitoring or heartbeat monitoring (or both). If either mechanism fails to respond within the time-out period, the watchdog assumes the slot has failed. For details on RSP watchdog, see the *RSP User's Guide*.

Configuring the Heartbeat Watchdog

A guest application inside the VM opens a virtual serial port and transmits a signal at regular intervals out of this port to the Steelhead appliance. This signal is called the *slot heartbeat*.

On the host side the RSP system opens a UNIX pipe, created by the VMware Server, corresponding to the other end of this virtual serial port.

If the RSP system does not detect the slot heartbeat in a certain time period, it assumes that the VM has failed and applies the failure policy to the slot. Other than the designated heartbeat characters, the VM ignores all data received over the virtual serial port.

A package sending a heartbeat through a virtual serial port must also accept any input sent through the same port by the Steelhead appliance.

The watchdog continues to monitor the slot for the heartbeat. If and when it re-appears, the watchdog assumes the VM is fully operational and directs the network traffic through the slot.

To configure an RSP watchdog

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-12. RSP Slots Page - Watchdog Settings

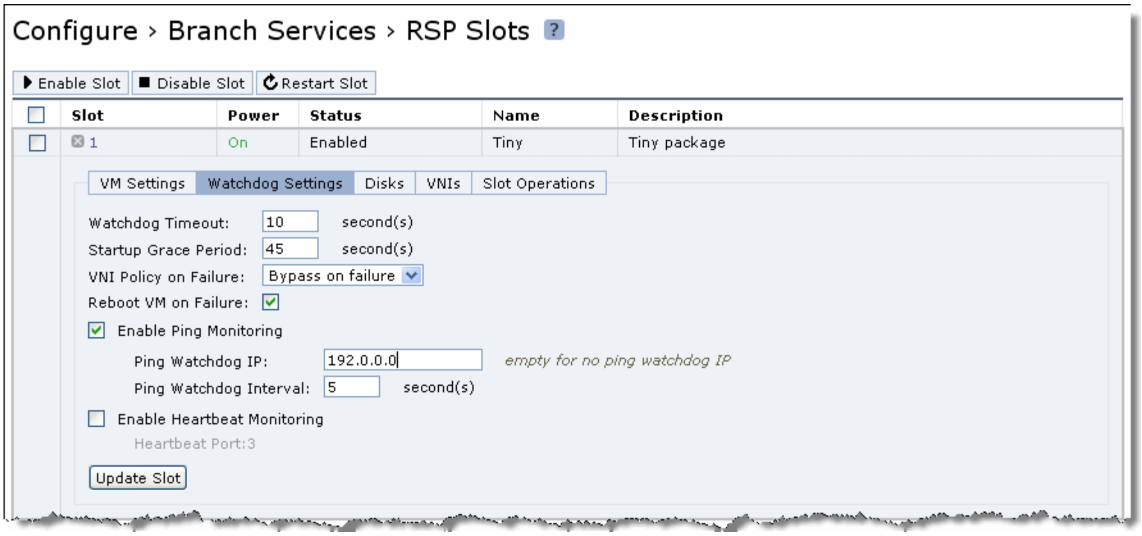

□	Slot	Power	Status	Name	Description
□	1	On	Enabled	Tiny	Tiny package

Watchdog Timeout: 10 second(s)

Startup Grace Period: 45 second(s)

VNI Policy on Failure: Bypass on failure

Reboot VM on Failure: ☑

☑ Enable Ping Monitoring

Ping Watchdog IP: 192.0.0.0 *empty for no ping watchdog IP*

Ping Watchdog Interval: 5 second(s)

□ Enable Heartbeat Monitoring

Heartbeat Port:3

2. Select the slot name in the Slot column.

3. Select Watchdog Settings and complete the configuration as described in the following table.

Control	Description
Watchdog Timeout	Specify the number of seconds to wait for a response from the package. You must enter a positive integer. **Note:** If you enable ping monitoring, specify a larger number of seconds for the watchdog time-out than for the ping interval. The time should be greater than or equal to the ping interval.
Startup Grace Period	Specify the number of seconds to wait during the slot startup time before watchdog monitoring begins.

Control	Description
VNI Policy on Failure	Select one of the following options from the drop-down list: • **Block on failure** - Specifies that the watchdog blocks traffic if the VM in the slot fails. • **Bypass on failure** - Specifies that the watchdog bypasses traffic if the VM in the slot fails.
Reboot VM on Failure	Reboot the VM if it fails (stops running).
Enable Ping Monitoring	Enables ping monitoring that monitors the package by simply pinging a network interface in a virtual machine. The RSP package must have a management VNI before you can configure ping monitoring. For details, see "Managing Virtual Network Interfaces" on page 198. Specify the following options: • **Ping Watchdog IP** - Specify an IP address of the management VNI to ping. • **Ping Watchdog Interval** - Specify the number of seconds between pings to determine whether the package is responsive (for example, 60). You must enter a positive integer.
Enable Heartbeat Monitoring	Enables a watchdog process that operates as a heartbeat to monitor each package's health, and if necessary, automatically restart the virtual machines gracefully. For details, see "Configuring the Heartbeat Watchdog" on page 192.
Update Slot	Applies your configuration to the slot.

Managing Virtual Disks

You can add, extend, or delete virtual disks for an installed VM inside the RSP system.

This section describes the following:

- "Creating or Deleting a Virtual Disk" on page 193
- "Attaching a Virtual Disk to a VM" on page 195
- "Extending a Virtual Disk" on page 196

Creating or Deleting a Virtual Disk

You create a virtual disk in the RSP Slots page.

To create a virtual disk

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

Important: Before creating a virtual disk, choose Configure > Branch Services > RSP Service to display the RSP Service Page. Confirm that there is sufficient RSP Free Space available for the virtual disk.

Figure 5-13. RSP Slots Page - Disks

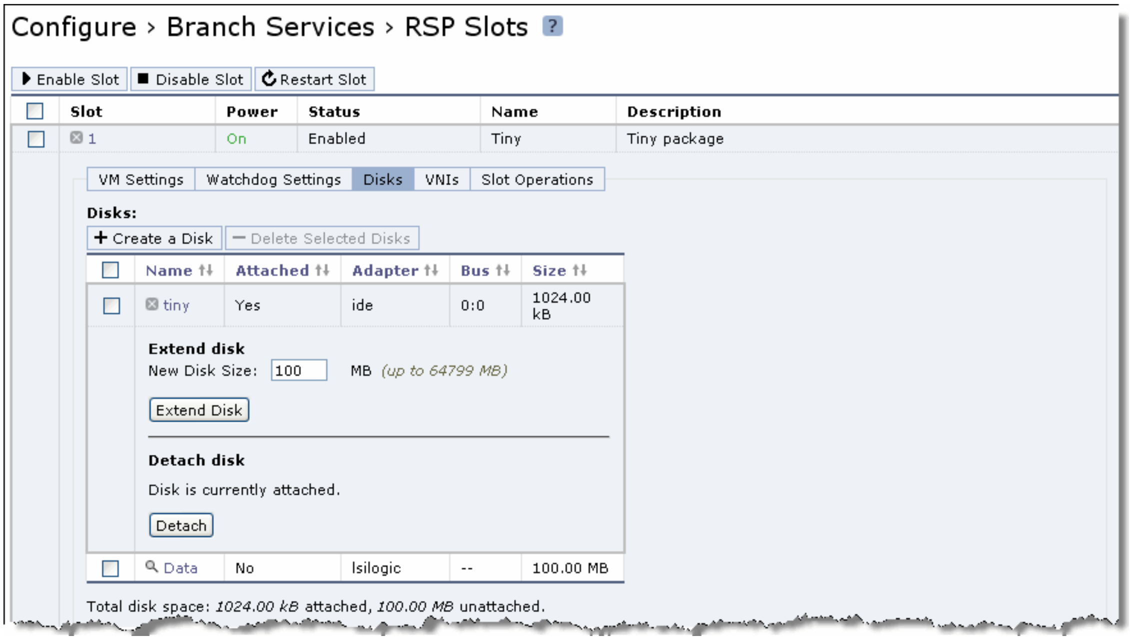

□	Slot	Power	Status	Name	Description
□	1	On	Enabled	Tiny	Tiny package

Disks:

□	Name	Attached	Adapter	Bus	Size
□	tiny	Yes	ide	0:0	1024.00 kB
□	Data	No	lsilogic	--	100.00 MB

Extend disk

New Disk Size: 100 MB *(up to 64799 MB)*

Detach disk

Disk is currently attached.

Total disk space: *1024.00 kB* attached, *100.00 MB* unattached.

2. Select the slot name in the Slot column.

3. Click **Disks**.

   A list of each disk name, whether it is attached to the VM, the adapter, bus, and size information appears.

4. Click **Create Disk** and complete the configuration as described in the following table.

Control	Description
Create Disk	Displays the controls to create a new virtual disk.
Disk Name	Specify a descriptive name for the virtual disk.
Disk Size	Specify the virtual disk size in MB. The maximum value of disk size is limited to the RSP free space displayed in the RSP Service page.

Control	Description
Disk Adapter	Specify the adapter type. The adapters connect the disk to the system. Select an adapter from the drop-down menu: • IDE adapter • LSI Logic SCSI adapter • Bus Logic SCSI adapter The LSI Logic and Bus Logic adapters are Small Computer Systems Interface (SCSI) adapters.
Delete Selected Disks	Select the check boxes next to the names of the virtual disks you want to delete from the VM, and click **Delete Selected Disks**. When you delete a virtual disk, the system deletes the files in the disk and returns the free space to the RSP system. **Important:** You can delete a virtual disk if it is detached from the VM. For details, see "Detaching a Virtual Disk from a VM" on page 197.

Attaching a Virtual Disk to a VM

After you create a virtual disk, you can attach it to a VM. Ensure that the VM is powered off before you attach the virtual disk to it because the system does not attach a virtual disk to a running VM.

To attach a virtual disk to a VM

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-14. RSP Slots Page - Disks - Attach Disk

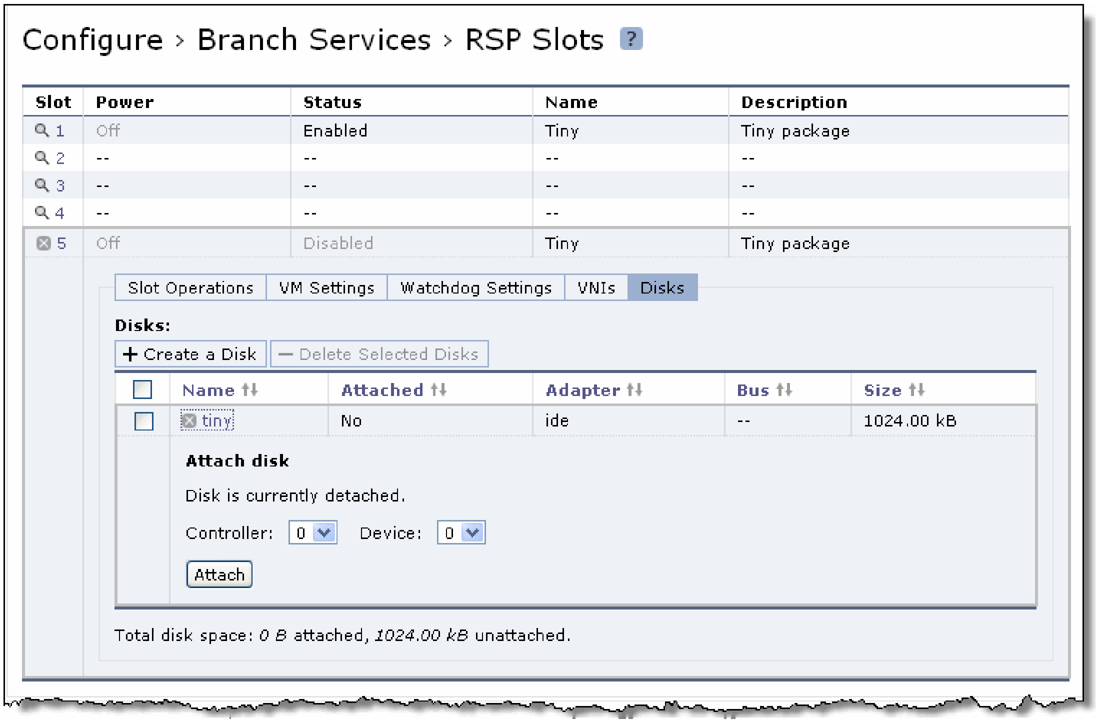

Configure › Branch Services › RSP Slots ?

Slot	Power	Status	Name	Description
1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--
4	--	--	--	--
5	Off	Disabled	Tiny	Tiny package

Disks:

□	Name	Attached	Adapter	Bus	Size
□	tiny	No	ide	--	1024.00 kB

Attach disk

Disk is currently detached.

Controller: 0 Device: 0

Total disk space: *0 B* attached, *1024.00 kB* unattached.

2. Select the slot name in the Slot column.

3. Click **Disks**.

A list of each disk name, whether it is attached to the VM, the adapter, bus, and size information appears.

4. Select the name of a detached disk and complete the configuration as described in the following table.

Control	Description
Controller	Select a value for the controller ID from the drop-down list. The controller enables the CPU to communicate with the disk. For SCSI adapters (LSI Logic and Bus Logic), you can select a controller value from 0 to 3. For the IDE adapter, you can select a controller value 0 or 1.
Device	Select a value for the device ID from the drop-down menu. For SCSI adapters, the device has a value from 0 to 15. For the IDE adapter, the device is 0 or 1.
Attach	Attaches the virtual disk to the VM.

Extending a Virtual Disk

You can increase the size of (extend) a virtual disk attached to a VM.

Before extending a disk:

- Ensure that the VM is powered off.
- Check that your system meets RSP requirements and has enough free space to accommodate the new disk size. For details, see the *RSP User's Guide*.
- Create a backup of your disk.

To extend a virtual disk

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-15. RSP Slots Page - Disks - Extend Disk

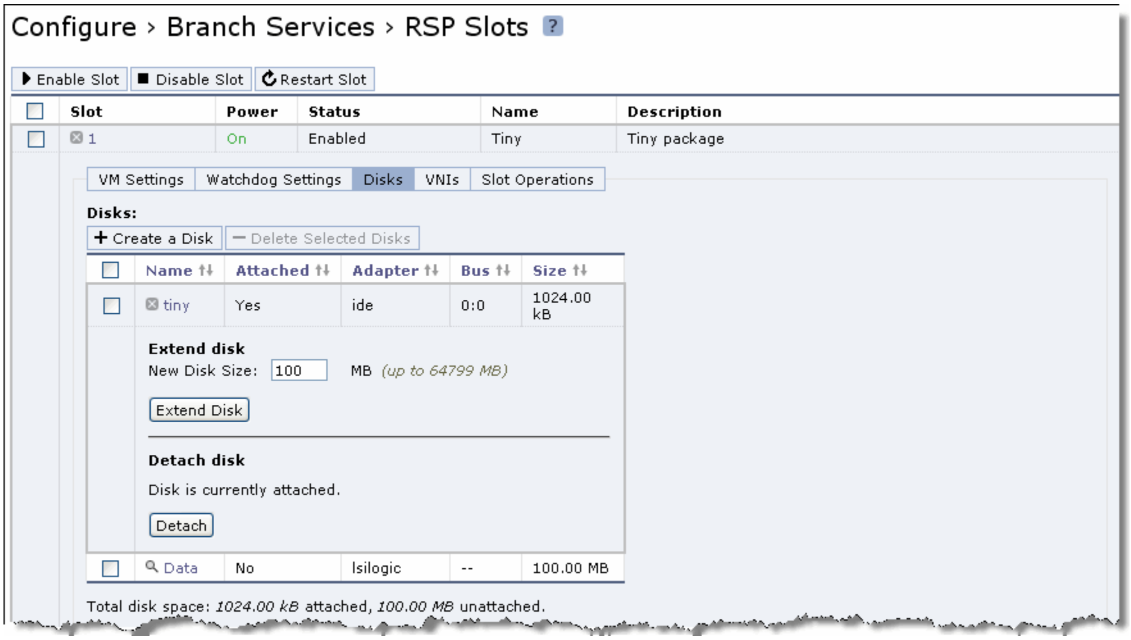

Configure › Branch Services › RSP Slots

□	Slot	Power	Status	Name	Description
□	1	On	Enabled	Tiny	Tiny package

Disks:

□	Name	Attached	Adapter	Bus	Size
□	tiny	Yes	ide	0:0	1024.00 kB
□	Data	No	lsilogic	--	100.00 MB

Extend disk

New Disk Size: 100 MB *(up to 64799 MB)*

Detach disk

Disk is currently attached.

Total disk space: *1024.00 kB* attached, *100.00 MB* unattached.

2. Select the slot name in the Slot column.

3. Click **Disks**.

   A list of each disk name, whether it is attached to the VM, the adapter, bus, and size information appears.

4. Select the name of an attached disk and complete the configuration as described in the following table.

Control	Description
New Disk Size	Specify the new disk size in MB. The maximum size is limited to the RSP free space.
Extend Disk	Enlarges the virtual disk size to the new value you specify.

Detaching a Virtual Disk from a VM

You can detach an existing disk from a VM. If the VM is powered on, the disk is detached only after you restart the VM. You can detach any disk from the VM, including the disk in the source RSP package.

After you detach a disk from the VM, the files corresponding to the disk remain in the VM slot directory.

To remove these files, you must delete the disk. For details, see "Creating or Deleting a Virtual Disk" on page 193.

To detach a virtual disk

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-16. RSP Slots Page - Disks - Detach Disk

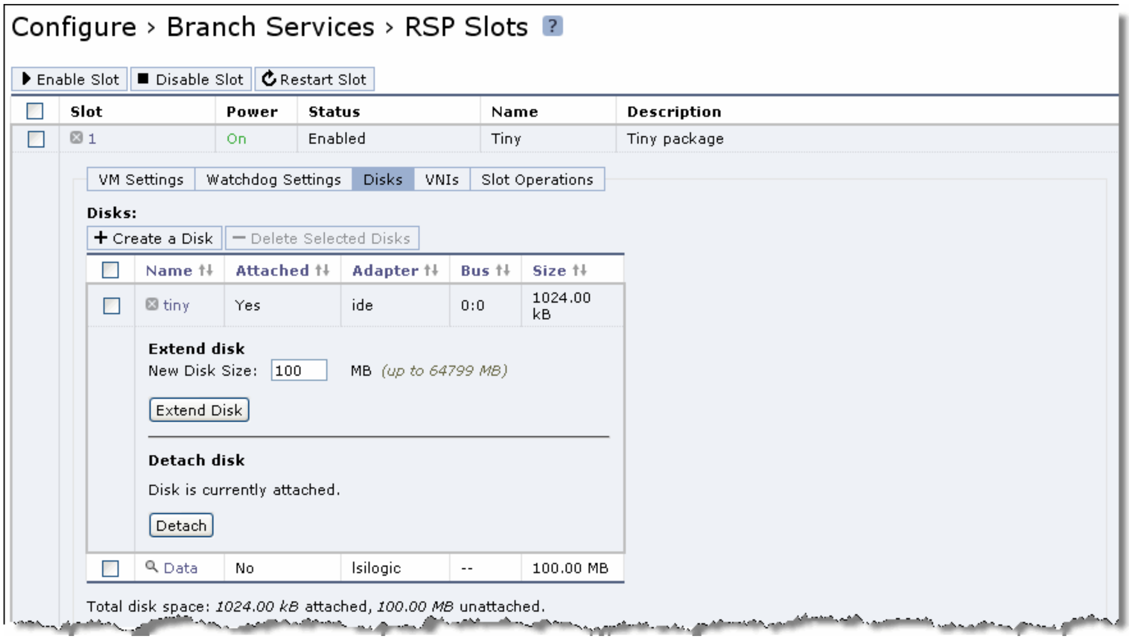

☐	Slot	Power	Status	Name	Description
☐	1	On	Enabled	Tiny	Tiny package

Disks:

☐	Name	Attached	Adapter	Bus	Size
☐	tiny	Yes	ide	0:0	1024.00 kB
☐	Data	No	lsilogic	--	100.00 MB

Extend disk

New Disk Size: 100 MB *(up to 64799 MB)*

Detach disk

Disk is currently attached.

Total disk space: *1024.00 kB* attached, *100.00 MB* unattached.

2. Select the slot name in the Slot column.

3. Click **Disks**.

   A list of each disk name, whether it is attached to the VM, the adapter, bus, and size information appears.

4. Select the name of an attached disk and click **Detach** to detach the disk from the VM.

Managing Virtual Network Interfaces

After you install and configure RSP, you can add a VNI to your configuration. For an overview of VNIs, see the *RSP User's Guide*.

This section describes how to update VLAN settings in the optimization VNIs and how to bridge a management VNI to an interface.

To update VLAN settings

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-17. RSP Slots Page - VNI Settings

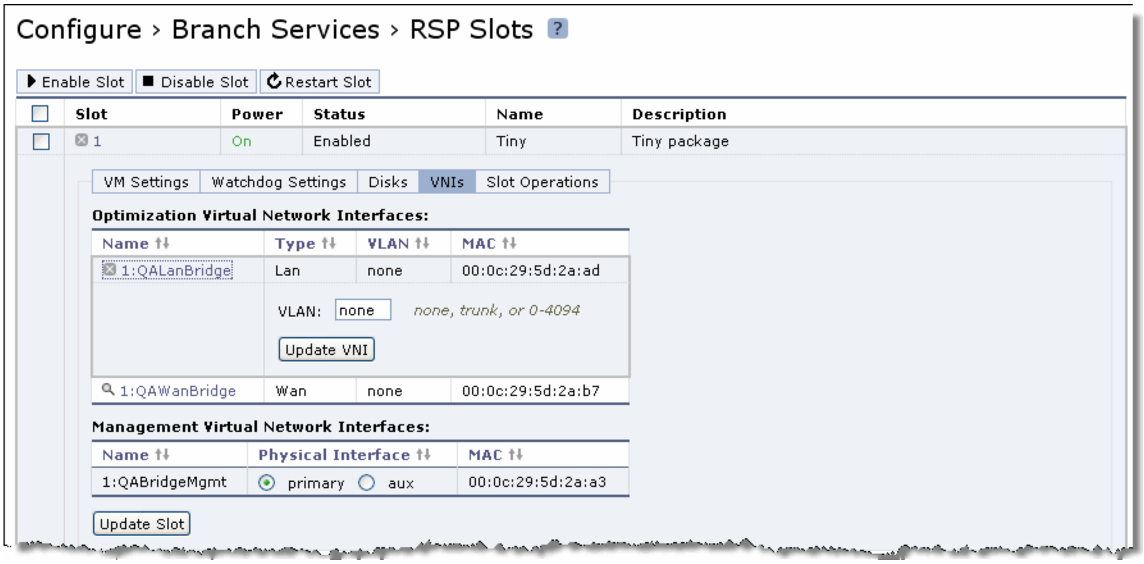

Configure > Branch Services > RSP Slots ?

☐	Slot	Power	Status	Name	Description
☐	☒ 1	On	Enabled	Tiny	Tiny package

Optimization Virtual Network Interfaces:

Name	Type	VLAN	MAC
☒ 1:QALanBridge	Lan	none	00:0c:29:5d:2a:ad
	VLAN: none	*none, trunk, or 0-4094*	
	Update VNI		
1:QAWanBridge	Wan	none	00:0c:29:5d:2a:b7

Management Virtual Network Interfaces:

Name	Physical Interface	MAC
1:QABridgeMgmt	● primary ○ aux	00:0c:29:5d:2a:a3

Update Slot

2. Select the slot name in the Slot column.

3. Select the VNIs tab.

   The optimization VNIs and management VNIs appear.

4. Under Optimization Virtual Network Interfaces, select the name of a VNI.

5. Specify trunk, none, or a VLAN tag identification number from 0 to 4094 (available VLAN numbers).

6. Click **Update VNI**.

To bridge a management VNI to an interface

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

2. Select the slot name in the Slot column.

3. Select the VNIs tab.

   The optimization VNIs and management VNIs appear.

4. Under Management Virtual Network Interfaces, after the VNI name, click **primary** or **aux** as the physical interface.

5. Click **Update Slot**.

The change takes effect the next time the slot is powered on.

Performing RSP Operations

This section describes how to perform the following RSP operations:

- "Uninstalling a Slot" on page 199
- "Restoring an RSP Backup" on page 200

Uninstalling a Slot

Before you uninstall an RSP package from a slot, disable the slot in which the RSP package resides. For details, see "To enable, disable, or restart a slot" on page 190.

To uninstall a slot

1. Choose Configure > Branch Services > RSP Slots to display the RSP slots page.

Figure 5-18. RSP Slots Page - Operations - Uninstall Slot

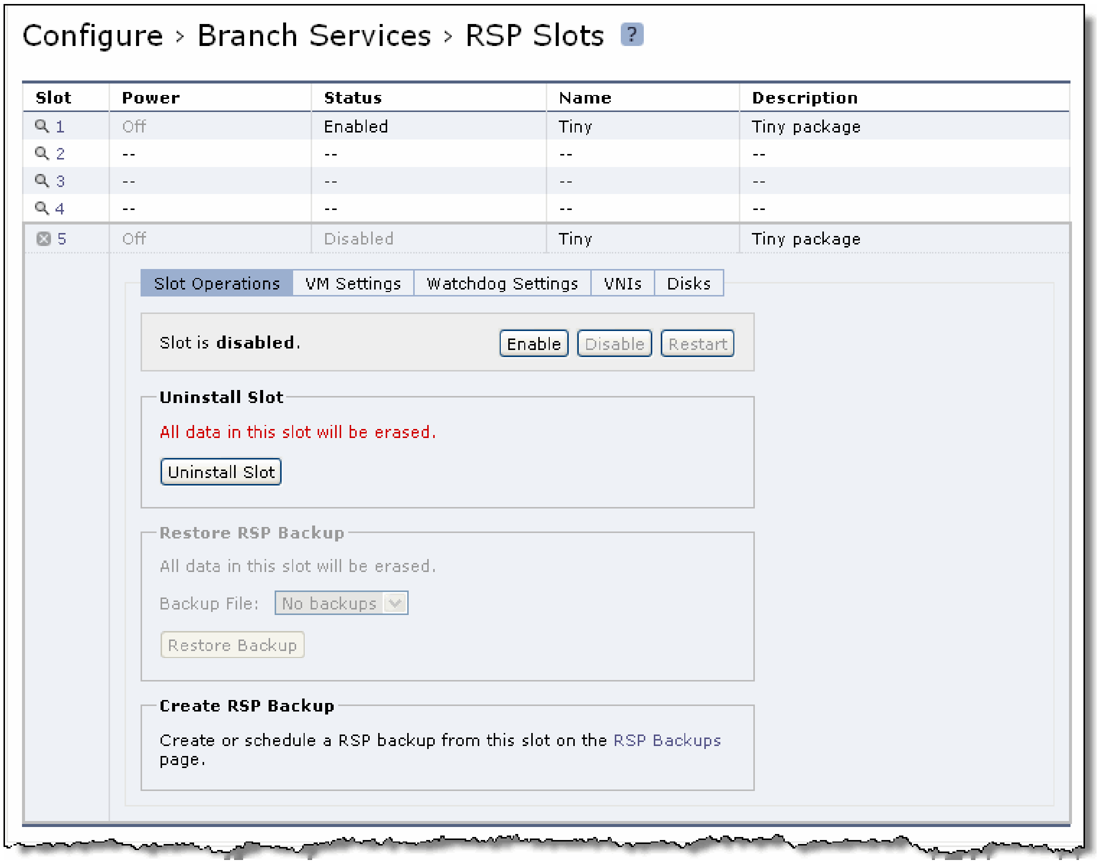

Configure > Branch Services > RSP Slots

Slot	Power	Status	Name	Description
1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--
4	--	--	--	--
5	Off	Disabled	Tiny	Tiny package

Slot is **disabled**.

Uninstall Slot

All data in this slot will be erased.

Restore RSP Backup

All data in this slot will be erased.

Backup File: No backups

Create RSP Backup

Create or schedule a RSP backup from this slot on the RSP Backups page.

2. Select the slot name in the Slot column.

3. Select the Slot Operations tab.

   If the slot is powered on, the controls to restore an RSP backup appear. If the slot is powered off, the controls to uninstall the slot and to restore an RSP backup appear.

4. Click **Uninstall Slot**.

   This deletes all data in the slot.

Restoring an RSP Backup

You use the RSP Slots - Operations page to restore the RSP data in case the Steelhead appliance fails. You create the RSP backup file in the RSP Backups page. For details, see “Configuring RSP Backups” on page 201.

You must restore an RSP backup into the same slot in which it was previously installed on the same Steelhead appliance model.

To restore an RSP backup

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.

Figure 5-19. RSP Slots Page - Operations - Restore Backup

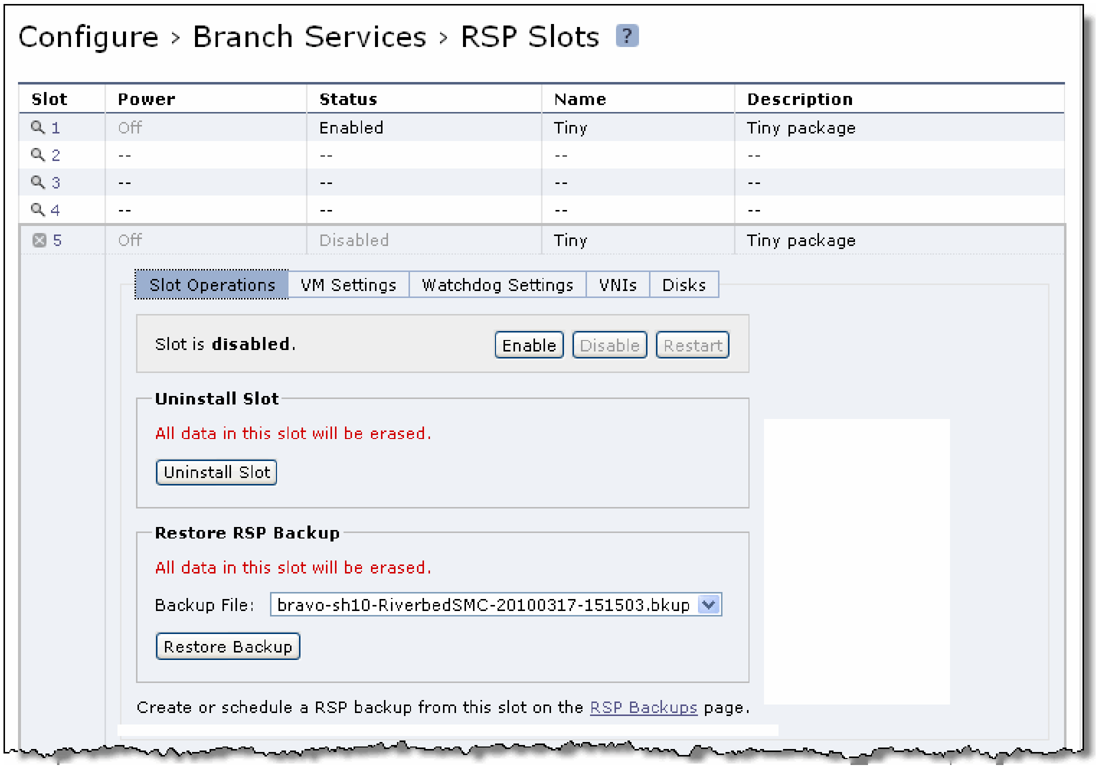

Configure > Branch Services > RSP Slots

Slot	Power	Status	Name	Description
1	Off	Enabled	Tiny	Tiny package
2	--	--	--	--
3	--	--	--	--
4	--	--	--	--
5	Off	Disabled	Tiny	Tiny package

Slot is **disabled**.

Uninstall Slot

All data in this slot will be erased.

Restore RSP Backup

All data in this slot will be erased.

Backup File: bravo-sh10-RiverbedSMC-20100317-151503.bkup

Create or schedule a RSP backup from this slot on the RSP Backups page.

2. Select the slot name in the Slot column.

3. Select the Slot Operations tab.

   If the slot is powered on, the controls to restore RSP backup appears. If the slot is powered off, the controls to uninstall the slot and to restore RSP backup appears.

4. Under Restore RSP Backup, select a backup file from the drop-down menu.

5. Click **Restore Backup**.

   This restores the data in the slot.

Configuring RSP Backups

You configure RSP backups in the RSP Backups page. Use RSP backups to save and restore the RSP data in case the Steelhead appliance fails. The backup operation generates a backup file with a .bkup file extension. The default backup filename is <steelheadname>-<slotname>-<date>.bkup.

Do not shut down the VM before you create a backup.

You must restore an RSP backup into the same slot in which it was previously installed.

Note: Before you create an RSP backup file, ensure that the RSP partition has free space that is greater than or equal to the size of the RSP slot (total size of all VMDKs) + Slot RAM size + 20 MB. For example, an RSP slot with 30 GB VMDK and 1 GB of RAM must have at least 31 GB of free disk space on the RSP partition for the RSP backup operation to succeed.

RSP Backup Limitation

You can restore live backups (backups of virtual machines or slots that are powered on) only on the same model of appliance on which they were created. For example, if a live backup was created on a model 550 Steelhead appliance, you can restore it only on a model 550 Steelhead appliance. If you want to move the backup to a different model Steelhead appliance, power off the slot before performing the backup (architecture limitations still apply—you cannot start a 64-bit virtual machine on a 32-bit appliance).

To configure an RSP backup

1. Choose Configure > Branch Services > RSP Backups to display the RSP Backups page.

Figure 5-20. RSP Backups Page

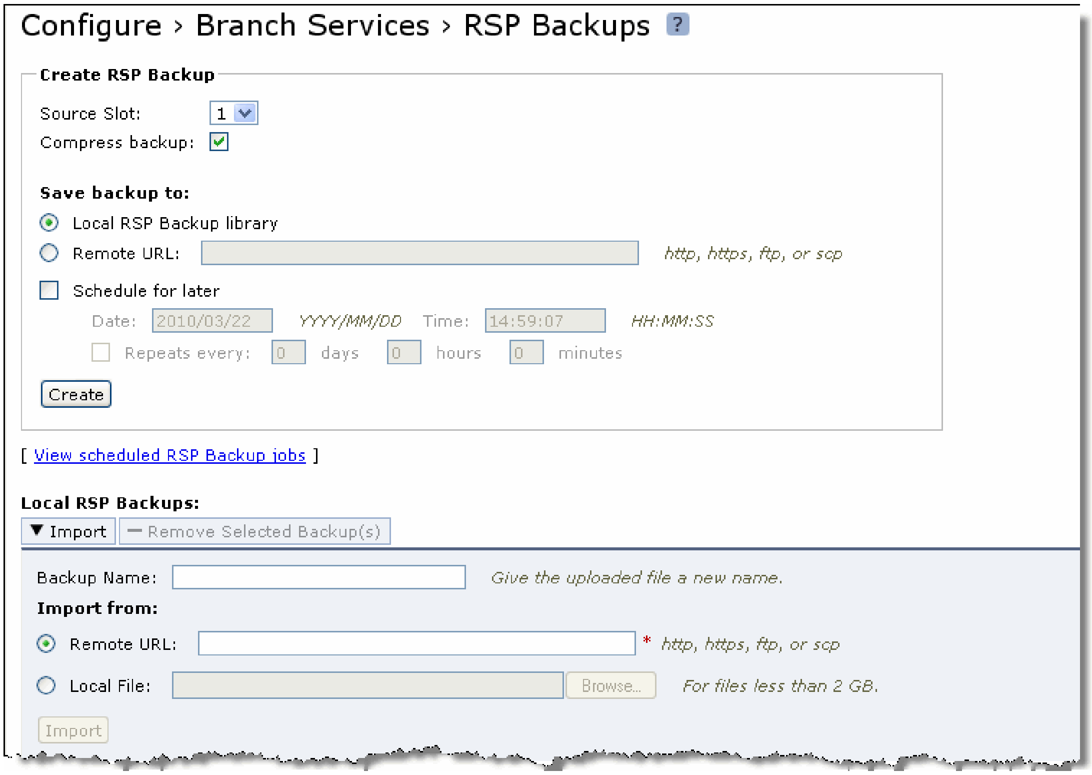

Configure > Branch Services > RSP Backups

Create RSP Backup

Source Slot: 1
Compress backup: ☑

Save backup to:
● Local RSP Backup library
○ Remote URL: *http, https, ftp, or scp*

□ Schedule for later
Date: 2010/03/22 *YYYY/MM/DD* Time: 14:59:07 *HH:MM:SS*
□ Repeats every: 0 days 0 hours 0 minutes

Create
[View scheduled RSP Backup jobs]

Local RSP Backups:

▼ Import
— Remove Selected Backup(s)

Backup Name: *Give the uploaded file a new name.*
Import from:
● Remote URL: * *http, https, ftp, or scp*
○ Local File: Browse... *For files less than 2 GB.*

2. Under Create RSP Backup, complete the configuration as described in the following table.

Control	Description
Source Slot	Select the slot for which you create a backup file from the drop-down menu.
Compress backup	Compresses the RSP backup file. This option is enabled by default. In most cases, Riverbed recommends that you compress RSP backups. However, if you are transferring—or expect to transfer—an RSP backup over a slow WAN, you might want to leave it uncompressed. The Steelhead appliance transfers an uncompressed RSP backup faster, assuming there is a corresponding Steelhead appliance on the other side of the transfer.

3. Under Save backup to, select one of the following options:

- **Local RSP Backup library** - Select this option to save the backup file to the local RSP backup library on the Steelhead appliance file system.
- **Remote URL** - Select this option and specify the remote URL to store the RSP backup file using FTP or SCP.

4. Under Schedule for later, complete the configuration as described in the following table.

Control	Description
Date	Specify the date (in yyyy/mm/dd format) on which the RSP backup file should be created.
Time	Specify the time (in hh:mm:ss format) at which the RSP backup file should be created.
Repeats every	Specify the number of days, hours, and minutes for which the RSP backup operation should be repeated.

5. Click **Create** to create the RSP backup file.

6. Under Local RSP Backups, complete the configuration as described in the following table.

Control	Description
Backup Name	Specify a name for the local backup file.
Import from	• **Remote URL** - Specify the remote URL from which the RSP backup file should be imported using ftp or scp. • **Local file** - Click **Browse** and select the local backup file. Use this option only if the local RSP backup file is less than 2 GB.
Import	Imports the RSP backup file from the location you specify.
Remove Selected Backups	Deletes the selected backup files.

Configuring RSP HA

You configure RSP High Availability (HA) in the RSP HA page.

RSP HA enables specific RSP slots on a Steelhead appliance (source) to be copied automatically to another Steelhead appliance (target). If the source Steelhead appliance ever fails, you can replace its slots with the slot functionality using the cloned slots on the target Steelhead appliance.

There are two types of RSP HA transfers—incoming and outgoing. Incoming HA transfers are the replicated RSP slots copied into the target Steelhead appliance. Outgoing HA transfers are replicated RSP slots copied out of the source Steelhead appliance.

You can schedule the RSP slots on the source Steelhead appliance to be transferred automatically to the target Steelhead appliance at a regular time period.

You must set up the following before the source Steelhead appliance accepts incoming RSP slot copies:

- Ensure that the packages installed on a given slot on the source Steelhead appliance match the packages on the target Steelhead appliance. For example, if slot 1 on the source Steelhead appliance has the package smc_pkg v3.0 installed on it, slot 1 on the target Steelhead appliance must also have the same package (smc_pkg v3.0) installed on it. You must disable the slot on the target Steelhead appliance throughout the process for incoming transfers to work.
- Specify an RSP HA password on the target Steelhead appliance. The source Steelhead appliance administrator uses this password to copy slots and transfer slots to the target Steelhead appliance.
- Enable incoming RSP HA transfers on the target Steelhead appliance.

Important: You can perform HA data synchronization only between Steelhead appliances of the same model. RSP HA is only supported on Steelhead appliance models x50 and xx50.

To configure RSP HA

1. Choose Configure > Branch Services > RSP HA to display the RSP HA page.

Figure 5-21. RSP HA Page

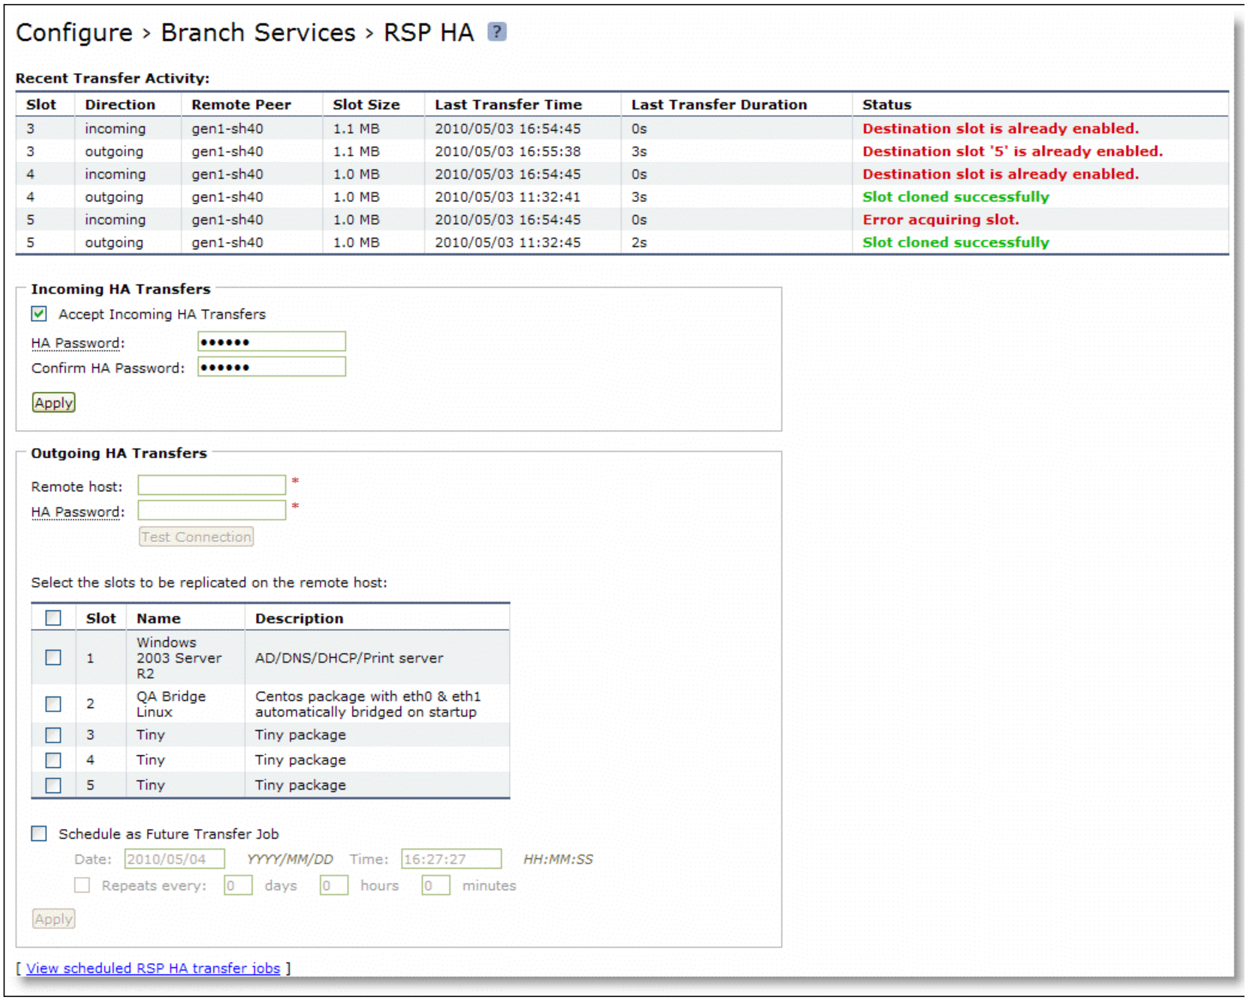

Configure > Branch Services > RSP HA

Recent Transfer Activity:

Slot	Direction	Remote Peer	Slot Size	Last Transfer Time	Last Transfer Duration	Status
3	incoming	gen1-sh40	1.1 MB	2010/05/03 16:54:45	0s	Destination slot is already enabled.
3	outgoing	gen1-sh40	1.1 MB	2010/05/03 16:55:38	3s	Destination slot '5' is already enabled.
4	incoming	gen1-sh40	1.0 MB	2010/05/03 16:54:45	0s	Destination slot is already enabled.
4	outgoing	gen1-sh40	1.0 MB	2010/05/03 11:32:41	3s	Slot cloned successfully
5	incoming	gen1-sh40	1.0 MB	2010/05/03 16:54:45	0s	Error acquiring slot.
5	outgoing	gen1-sh40	1.0 MB	2010/05/03 11:32:45	2s	Slot cloned successfully

Incoming HA Transfers

☑ Accept Incoming HA Transfers

HA Password: ••••••

Confirm HA Password: ••••••

Apply

Outgoing HA Transfers

Remote host: *

HA Password: *

Test Connection

Select the slots to be replicated on the remote host:

☐	Slot	Name	Description
☐	1	Windows 2003 Server R2	AD/DNS/DHCP/Print server
☐	2	QA Bridge Linux	Centos package with eth0 & eth1 automatically bridged on startup
☐	3	Tiny	Tiny package
☐	4	Tiny	Tiny package
☐	5	Tiny	Tiny package

☐ Schedule as Future Transfer Job

Date: 2010/05/04 YYYY/MM/DD Time: 16:27:27 HH:MM:SS

☐ Repeats every: 0 days 0 hours 0 minutes

Apply

[View scheduled RSP HA transfer jobs]

The Recent Transfer Activity table lists the slot name, the direction of transfer (incoming or outgoing), the remote peer name, slot size, last transfer date (in yyyy/mm/dd format) and time (in hh:mm:ss format), last transfer duration, and status of the HA transfer.

2. Under Incoming HA Transfers, complete the configuration as described in the following table.

Control	Description
Accept Incoming HA Transfers	Enables the Steelhead appliance to accept incoming HA transfers.
HA Password	Specify a password for the incoming HA transfer. This password decrypts slot images that are received from HA peers. Remote hosts must supply this password to initiate HA transfers. Do not use your Steelhead appliance account password in this field.
Confirm HA Password	Specify the HA password again to confirm.
Apply	Applies your configuration changes.

3. Under Outgoing HA Transfers, complete the configuration as described in the following table.

Control	Description
Remote host	Specify the remote hostname or IP address.
HA Password	Specify a password for the outgoing HA transfer. This password encrypts slot images that are sent to HA peers. This password must match the incoming HA password of the remote hosts. Do not use your Steelhead appliance account password in this field.
Select the slots to be replicated on the remote host	Select the check box before a slot name to select the slot. The system makes a copy of the slot on the remote host you specify.
Schedule as Future Transfer Job	Select to schedule the outgoing HA transfer at the date and time you specify. **Date -** Specify the date (in yyyy/mm/dd format) on which RSP should make the outgoing HA transfer. **Time** - Specify the time (in hh:mm:ss format) at which RSP should make the outgoing HA transfer. **Repeats every -** Specify the days, hours, and minutes at which RSP should repeat the outgoing HA transfer.
Apply	Click to apply your configuration.

Configuring RSP Data Flow

You can configure Virtual Network Interfaces (VNIs) for an in-path RSP package in the Configure > Branch Services > RSP Data Flow page.

RSP packages use their own virtual network interfaces (VNIs), equivalent to VMware network interfaces, to communicate with the network. These virtual network interfaces are connected to the physical network interfaces on the Steelhead appliance, or taps into the optimization data flow. VNIs are network taps that enable data flow in and out of the RSP packages. VNIs are the connection points for the LAN, WAN, primary, and auxiliary interfaces of the Steelhead appliance.

Each package can have up to ten RSP network interfaces, which enables each package to support ten VNIs. VNIs provide a great deal of configuration flexibility. VNIs also provide the basis for controlling the flow of data through packages in a multiple package scenario.

You can use the following types of VNIs:

- **Optimization VNIs** - Optimization VNIs are used with in-band packages. Optimization VNIs are part of the optimized data flow on either the LAN- or WAN-side of RiOS. There are several types of optimization VNIs:
    - **In-path** - In-path VNIs are used for packages such as security packages. The following types of in-path optimization VNIs are available:
    - – **LAN** - LAN VNIs forward packets from the LAN-side to the virtual machine, to the WAN-side, or both. LAN VNIs unconditionally forward packets from the virtual machine to the LAN-side for RSP. LAN VNIs cannot receive packets from the WAN-side. For VRSP, packets from LAN or WAN VNIs can go in either direction, depending on the subnet side rules. For details, see “Configuring Subnet Side Rules” on page 260.
    - – **WAN** - WAN VNIs forward packets from the WAN-side to the virtual machine, to the LAN-side, or both. WAN VNIs unconditionally forward packets from the virtual machine to the WAN-side. WAN VNIs cannot receive packets from the LAN-side.

- **Virtual-in-path** - Virtual-in-path VNIs are used for packages that need redirection to intercept traffic. The types of virtual-in-path VNIs are:
  - **DNAT** - Destination Network Address Translation (DNAT) VNIs are used for proxy-based solutions, such as video proxies.
  - **Mirror** - Mirror VNIs are used with network monitoring-based solutions and act like a SPAN port.
- **Management VNIs** - Management VNIs are bridged to the Steelhead appliance primary or auxiliary port. Management VNIs are used as a management interface for in-band packages. Management VNIs are the primary communication path for out-of-band packages.

Adding a VNI to the Data Flow

After you install and configure RSP, you can add VNI to your configuration.

To add a VNI to the data flow

1. Choose Configure > Branch Services > RSP Data Flow to display the RSP Data Flow page.

Figure 5-22. RSP Data Flow Page

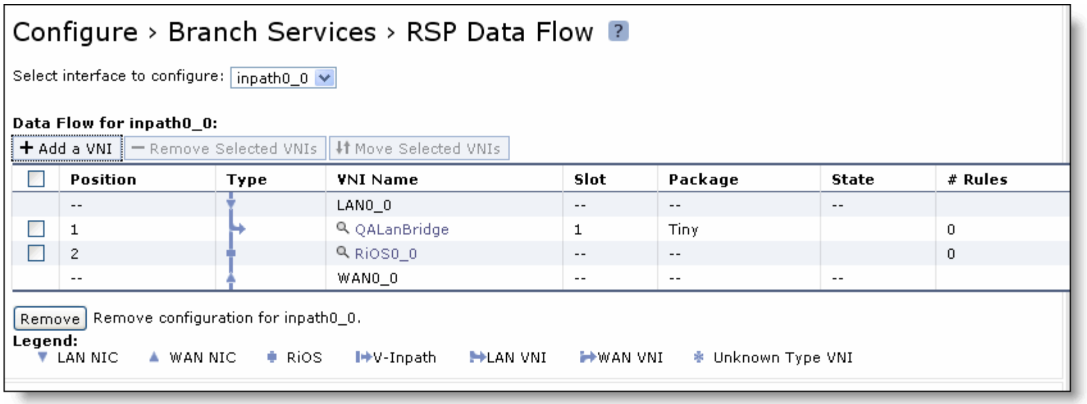

Configure › Branch Services › RSP Data Flow ?

Select interface to configure: inpath0_0

Data Flow for inpath0_0:

□	Position	Type	VNI Name	Slot	Package	State	# Rules
	--		LAN0_0	--	--	--	
□	1		QALanBridge	1	Tiny		0
□	2		RiOS0_0	--	--		0
	--		WAN0_0	--	--	--	

Remove Remove configuration for inpath0_0.

Legend:

▼ LAN NIC ▲ WAN NIC ● RiOS V-Inpath LAN VNI WAN VNI Unknown Type VNI

2. Under Data Flow for the selected, complete the configuration as described in the following table.

Control	Description
Add a VNI	Displays the controls to add a VNI.
Interface	Select an in-path interface from the drop-down list.
Data Flow Position	Select one of the following from the drop-down list. • **Start** - Locates the VNI next to the LAN. A packet coming from the Steelhead appliance LAN interface goes to this VNI first. • **End** - Locates the VNI next to the WAN. A packet coming from the Steelhead appliance WAN interface goes to this VNI first. • **Order number** - Specifies the VNI order number. A lower number locates the VNI closer to the LAN. A higher number locates the VNI closer to the WAN.
Add	Adds the VNI to the data flow.

Control	Description
Removed Selected VNIs	Select the check box next to the name and click **Remove Selected VNIs**.
Move Selected VNIs	Moves the selected VNIs. Click the arrow next to the desired VNI position; the VNI moves to the new position.

The next step is to add rules to configure the traffic flow, unless you use the default rules for the package. For more details on the package data flow, refer to the package documentation.

Important: To route optimized traffic through a Steelhead appliance that is using QoS and RSP simultaneously, RiOS must be positioned last in the RSP traffic data flow. At the bottom of the RSP Data Flow page, verify that RiOS appears as the last item in the data flow table, directly above the WAN interface. If it does not appear last, select the check box next to RiOS, click **Move Selected VNIs**, and then click the arrow next to the position directly above WAN. RiOS moves to the last position.

Adding Rules to an Optimization VNI

VNI rules determine what the VNI does with the traffic it receives. After you install a package and add it to a slot, you need to add rules to configure the data flow for the package unless you use the default rules for the package. For a WAN VNI, you add WAN-to-LAN rules to direct traffic. For a LAN VNI, you add LAN-to-WAN rules to direct traffic.

Rules are used with in-path and virtual in-path optimization VNIs. You do not need to add rules to management VNIs.

VNI rules determine what the VNI does with the traffic it receives. You can control the redirection by rules based on IP or port. VNI rules can perform one of the following actions:

- Redirect the packets to the VM.
- Pass the packets along the data flow to the next VNI.
- Pass the packets along the data flow and send a copy of the packets to the VM.

Using VNI Rules to Chain Packages Together

For example, if you installed a video streaming package, a security package, and a VPN package on the Steelhead appliance, you can define rules to invoke the following data path:

1. Flash video traffic coming from the LAN Steelhead appliance is redirected to a video proxy solution.

2. All other traffic goes directly to RiOS and is optimized.

3. After RiOS optimizes the traffic, it is redirected to the security package on the WAN-side of the Steelhead appliance that checks the data (or, if it is a VPN solution, encrypts it), and sends it back out of the WAN. You can control the data redirection using rules based on IP or port.

DNAT Rules

Destination Network Address Translation (DNAT) rules are used for in-path proxy-based solutions. You can only add DNAT rules for a Virtual In-path optimization VNIs.

By default, DNAT is disabled. When DNAT is enabled, it translates the network address of packets that match the source and destination IP and the port (or port range) to the target IP and, optionally, the target port. It then routes them to the correct device, host, or network.

For example, you can install an RSP package for live video streaming and add a DNAT rule (using the IP address, port number, or both) that transparently proxies all traffic redirected to the local RSP video instance.

Consequently, the local RSP video instance responds to the local clients on behalf of the original server, simultaneously communicating with the original server in the background over the WAN. This process streamlines the number of requests over the WAN, resulting in time and bandwidth savings.

The RSP rule that determines which traffic is network address translated is provided in the data flow rules for the Virtual In-path VNI.

Usage Notes

Use the following tips when you create data flow rules:

- Place the VNI in the data flow only for in-band packages
- You can either use the default IP rules or non-IP rules per VNI (customizable by each package vendor).
- Add rules to control traffic (such as in-path rules) per VNI and perform the following actions:
  - Redirect the packets to the VM.
  - Pass the packets along the data flow to the next VNI.
  - Pass the packets along the data flow and send a copy of the packets to the VM.
- Each given data flow rule pertains to one direction only, incoming or outgoing.
- For a WAN VNI, only WAN-to-LAN rules apply.
- For a LAN VNI, only LAN-to-WAN rules apply.
- You can only add DNAT rules for a virtual in-path VNI.
- You can specify a target port range with DNAT rules.

  The following table shows some typical rule actions that you can use to control the data flow for the various VNI types.

Optimization VNI type	Typical Data Flow Rule Actions
LAN	Pass traffic around a VM or redirect it to the VM.
WAN	Pass traffic around a VM or redirect it to the VM.
Virtual In-Path DNAT	Pass traffic to the target IP or redirect it to a target IP.
Virtual In-Path Mirror	Pass traffic along the data flow and copy it for monitoring.

To add a LAN-to-WAN or WAN-to-LAN rule to a VNI

1. Choose Configure > Branch Services > RSP Data Flow to display the RSP Data Flow page.

2. In the VNI list, select the VNI name.

3. Complete the rule properties under LAN to WAN Rules or WAN to LAN Rules, as described in the following table.

Control	Description
Add a Rule	Displays the controls to add a rule.
Rule Number	Optionally, specify a number (0 - 65535) to replace the default rule number.
Protocol	Optionally, select All, TCP, UDP, or specify a protocol number (1-254). The default setting is All.
Source Network	Optionally, specify the source subnet of the packet; for example, 1.2.3.0/24, or leave it blank to specify all subnets.
Source Port	Optionally, specify a single source port value or a port range of port1-port2, where port1 must be less than port2. Leave it blank to specify all ports. **Note:** You can only specify a port when the protocol is either TCP or UDP.
Destination Network	Optionally, specify the destination network or leave it blank to specify all destination networks.
Destination Port	Optionally, specify the destination port of the packet, either a single port value or a port range of port1-port2, where port1 must be less than port2. **Note:** You can only specify a port when the protocol is either TCP or UDP.
Vlan	Optionally, specify a VLAN identification number or leave it blank to specify all VLANs.
Rule Action	Select one of the following from the drop-down list: • **Redirect Traffic to Slot** - Click to redirect the packet to a VM. • **Pass Traffic Around Slot** - Click to pass the packet along the data flow, bypassing the VM. • **Copy Traffic to Slot** - Click to copy the packet to the VM and also pass it along the data flow.
Add	Adds the rule to the VNI.
Removed Selected Rules	Removes the selected rules. Select the check box next to the name and click **Remove Selected Rules**.
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

Changing the Default VNI Rules

The LAN-to-WAN and WAN-to-LAN rule tables include default rules to direct traffic that does not match any other rules. The tables always list these default rules last and you cannot delete them.

Riverbed supplies the following default rules to direct traffic for optimization VNIs.

Optimization VNI Type	Data Flow Action on IP Traffic	Data Flow Action on Non-IP Traffic
LAN	Pass	Pass
WAN	Pass	Pass
Virtual in-path	L2 switch	L2 switch

To edit the default VNI rule settings

1. Choose Configure > Branch Services > RSP Data Flow to display the RSP Data Flow page.

2. In the VNI list, select the VNI name.

3. Under either Default IP Policy or Default Non-IP Policy, select one of the following data flow options:
   - **Pass Traffic Around Slot** - Redirects the packet to a VM in a slot.
   - **Redirect Traffic to Slot**- Passes the packet along the data flow, bypassing the VM in a slot.
   - **Copy Traffic to Slot** - Copies the packet to the slot and also pass it along the data flow.
   - **Use L2 switching** - Uses L2 switching.

4. Click **Apply**.

5. Click **Save** to save your changes permanently.

To add a DNAT rule to a virtual in-path VNI

1. Choose Configure > Branch Services > RSP Data Flow to display the RSP Data Flow page.

2. In the VNI list, select the VNI name.

3. Complete the rule properties under Destination NAT Rules, as described in the following table.

Control	Description
Add a Rule	Displays the controls to add a rule.
Rule Number	Optionally, specify a number (0 - 65535) to replace the default rule number.
Protocol	Optionally, select All, TCP, UDP, or specify a protocol number (1-254). The default setting is All.
Source Network	Optionally, specify the source subnet of the packet; for example, 1.2.3.0/24, or leave it blank to specify all subnets.
Source Port or Range	Optionally, specify a single source port value or a port range of port1-port2, where port1 must be less than port2. Leave it blank to specify all ports.
Original Destination Network	Optionally, specify the destination network or leave it blank to specify all destination networks.
Original Destination Port or Range	Optionally, specify the destination port of the packet, either a single port value or a port range of port1-port2, where port1 must be less than port2.
Target Destination Address	Optionally, specify the target address or leave it blank to specify all target addresses.
Target Destination Port or Range	Specify the target destination port of the packet, either a single port value or a port range of port1-port2, where port1 must be less than port2. Leave it blank to specify all target ports.
Add	Adds the rule to the VNI.

Control	Description
Removed Selected Rules	Removes the selected rules. Select the check box next to the name and click **Remove Selected Rules**.
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

Tip: To remove the data flow configuration for an interface, click **Remove Configuration**.

Bridging a Management VNI to an Interface

If your package has a management VNI, you can bridge it to a primary or auxiliary interface and connect to the respective physical Ethernet adapter on the Steelhead appliance. This provides network connectivity.

If you want the RSP watchdog to use ping monitoring, you must bridge a management VNI on the Steelhead appliance to the virtual machine interface. For details, see “Specifying Watchdog Settings” on page 191.

To bridge a management VNI to an interface

1. Choose Configure > Branch Services > RSP Slots to display the RSP Slots page.
2. Under Slots, select the package slot number.
3. Under Management Virtual Network Interfaces, after the VNI name, click **Primary** or **Aux**.
4. Click **Save** to save your settings permanently.

CHAPTER 6 Configuring SSL and a Secure Inner Channel

This chapter describes how to configure SSL support. It includes the following sections:

Configuring SSL Server Certificates and Certificate Authorities

The following section provides an overview of SSL support and describes how to configure SSL server certificates and Certificate Authorities. It includes the following sections:

SSL is a cryptographic protocol which provides secure communications between two parties over the Internet.

Typically in a Web-based application, it is the client that authenticates the server. To identify itself, an SSL certificate is installed on a Web server and the client checks the credentials of the certificate to make sure it is valid and signed by a trusted third party. Trusted third parties that sign SSL certificates are called Certificate Authorities (CA).

How Does SSL Work?

With Riverbed SSL, Steelhead appliances are configured to have a trust relationship, so they can exchange information securely over an SSL connection. SSL clients and servers communicate with each other exactly as they do without Steelhead appliances; no changes are required for the client and server application, nor are they required for the configuration of proxies. RiOS splits up the SSL handshake, the sequence of message exchanges at the start of an SSL connection.

In an ordinary SSL handshake, the client and server first establish identity using public-key cryptography, and then negotiate a symmetric session key to be used for data transfer. With Riverbed SSL acceleration, the initial SSL message exchanges take place between the client and the server-side Steelhead appliance.

Figure 6-1. Riverbed SSL

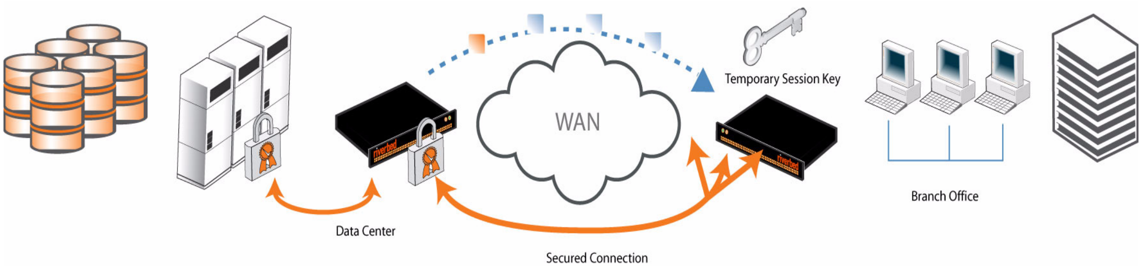

Prior to RiOS v6.0, the SSL handshakes from the client are always handled by the server-side Steelhead appliance. RiOS v6.0 and later provides an alternative handshake, called distributed termination, which terminates full handshakes on the client-side Steelhead appliance. The master secret containing information that allows the computation of the session key for reusing the session is transported to the session cache of the client-side Steelhead appliance. The subsequent handshakes are reused and the client's SSL connection is physically and logically terminated on the client-side Steelhead appliance.

Distributed termination improves performance by lessening the CPU load because it eliminates expensive asymmetric key operations. It also shortens the key negotiation process by avoiding WAN roundtrips to the server. You can find the setting to reuse a client-side session for distributed termination in the Configure > Optimization > Advanced Settings page. See "Setting Advanced SSL Options" on page 241.

In RiOS v6.1 and earlier, SSL optimization intercepts and optimizes SSL connections where only the SSL server uses a certificate. RiOS v6.5 provides client-side authentication, used to optimize SSL connections where the SSL server challenges the SSL client to present its own certificate, in addition to authenticating servers using SSL certificates. See "Configuring Advanced and SSL Cipher Settings" on page 240.

The Steelhead appliance also contains a secure vault which stores all SSL server settings, other certificates (that is, the CA, peering trusts, and peering certificates), and the peering private key. The secure vault protects your SSL private keys and certificates when the Steelhead appliance is not powered on. You set a password for the secure vault which is used to unlock it when the Steelhead appliance is powered on. After rebooting the Steelhead appliance, SSL traffic is not optimized until the secure vault is unlocked with the correct password. See "Unlocking the Secure Vault" on page 370.

Prerequisite Tasks

Complete the following prerequisite tasks before you begin SSL configuration:

1. Connect to the Management Console using HTTPS to protect your SSL private keys and certificates.

2. On the client and server-side Steelhead appliance, make sure you have a valid Enhanced Cryptography License Key. To verify your license, see "Managing Licenses and Model Upgrades" on page 349. If you

do not have a valid Enhanced Cryptography License Key file, go to https://sslcert.riverbed.com and follow the procedures documented there.

Note: The SSL License in RiOS v5.0.x and later is called the Enhanced Cryptographic License Key, because it also activates datastore encryption and creates secure channels while optimizing encrypted MAPI and SMB-signed traffic (even if the Steelhead appliances are not configured for optimizing SSL traffic).

3. Back up your private keys and the CA-signed certificates before you begin the SSL configuration process.

Basic Steps

This section provides an overview of the basic steps to configure SSL, followed by detailed procedures.

Task	Reference
1. Enable SSL support on the server-side and client-side Steelhead appliances.	"Configuring SSL Main Settings" on page 218
2. Set the SSL secure vault password on the client and server-side Steelhead appliance.	"Unlocking the Secure Vault" on page 370
3. Optionally, enable the Steelhead appliance to reuse the client-side SSL session. This is a client-side setting that improves connection setup performance. Both the client-side Steelhead appliance and the server-side Steelhead appliance must be running RiOS v6.0 or later. Enabling this option requires an optimization service restart.	"Setting Advanced SSL Options" on page 241
4. On the server-side Steelhead appliance, configure a proxy certificate and private key for the SSL back-end server. This step enables the server-side Steelhead appliance to act as a proxy for the back-end server, which is necessary to intercept the SSL connection and to optimize it.	"Configuring SSL Server Certificates" on page 220
5. Create a new in-path rule for the client-side Steelhead appliance. **In-path configurations** - Create a client-side in-path rule with the Preoptimization Policy = SSL. If you want to enable the HTTP latency optimization module for connections to this server, you add a corresponding in-path rule with Latency Optimization Policy = HTTP. **Out-of-path configurations** - On the client-side Steelhead appliance, add a new in-path rule to identify which connections are to be intercepted and applied to SSL optimization. Use the following property values: • Type - Fixed target • Destination Subnet/Port - Riverbed recommends you specify the exact SSL server IP address (for example, 10.11.41.14/32) and the default SSL port 443. • VLAN Tag - All • Preoptimization Policy - SSL • Optimization Policy - Normal • Latency Optimization Policy - HTTP Note: Latency optimization is not always HTTP, especially for applications that use the SSL protocol but are not HTTP based. In such cases, specify None for the latency optimization. • Neural Framing Mode - Always	"Configuring In-Path Rules" on page 28

Task	Reference
6. Configure mutual peering trusts so the server-side Steelhead appliance trusts the client-side Steelhead appliance and vice versa. Use one of the following approaches: **Use the secure inner channel and peering lists**: • Configure the inner channel SSL settings as described in "Configuring Secure Peers" on page 231. Both the client-side and server-side Steelhead appliances must be running RiOS v5.0 or later. • To automatically discover Steelhead appliances using self-signed certificates, open your secure application to send some traffic through the Steelhead appliances. The connection will be passed through to the server without optimization, but the Steelheads will automatically discover the peers and place them in the self-signed peer *gray* list. • Manually move the peers from the gray list to the trusted white list by simply marking them as trusted. The connections are not optimized until after you move the peers to the white list. • Reopen your secure application. —or— **Add CA-signed peer certificates**: • Add the PEM certificate of the designated CA as a new trusted entity to the peering trust list for each Steelhead appliance. • For production networks with multiple Steelhead appliances, use the CMC or the bulk import and export feature to simplify configuring trusted peer relationships. For details, see the *Steelhead Central Management Console User's Guide* or "Performing Bulk Imports and Exports" on page 248. **Tip**: Your organization can choose to replace all of the default self-signed identity certificates and keys on their Steelhead appliances with those certificates signed by another CA (either internal to your organization or an external well-known CA). In such cases, every Steelhead appliance must simply have the certificate of the designated CA (that signed all those Steelhead appliance identity certificates) added as a new trusted entity.	"Configuring Secure Peers" on page 231
7. If your organization uses internal CAs to sign their SSL server certificates you must import each of the certificates (in the chain) on to the server-side Steelhead appliance. You must perform this step if you use internal CAs because the Steelhead appliance default list of well-known CAs (trusted by our server-side Steelhead appliance) does not include your internal CA certificate. To identify the certificate of your internal CA (in some cases, the chain of certificate authorities) go to your Web browser repository of trusted-root or intermediate CAs. (For example, Internet Explorer > Tools > Internet Options > Certificates.)	"Configuring SSL Certificate Authorities" on page 223
8. On the client and server-side Steelhead appliance, restart the Steelhead service.	"Starting and Stopping the Optimization Service" on page 345

Verifying SSL and Secure Inner Channel Optimization

Use the following tools to verify that you have configured SSL support correctly.

- **SSL Optimization** - After completing the SSL configuration on both Steelhead appliances and restarting the optimization service, access the secure server from the Web browser. The following events take place in a successful optimization:

  - If you specified a self-signed proxy certificate for the server on the server-side Steelhead appliance, a pop-up window appears on the Web browser. View the certificate details to ensure that it is the same as the certificate on the server-side Steelhead appliance.
  - In the Management Console, the Current Connections report lists the new connection as optimized without a Protocol Error flag (new statistics appear every 60 seconds).
  - In the Management Console, the Traffic Summary report displays encrypted traffic (typically, HTTPS).
  - Verify that the back-end server IP appears in the SSL Discovered Server Table (Optimizable) in the SSL Main Settings page.

Note: Because all the SSL handshake operations are processed by the server-side Steelhead appliance, all the SSL statistics are reported on the server-side Steelhead appliance. No SSL statistics are reported on the client-side Steelhead appliance.

- **Monitoring SSL Connections** - Use the following tools to verify SSL optimization and to monitor SSL progress:
  - On the client Web browser, click the **Lock** icon to obtain certificate details. The certificate must match the proxy certificate installed on server-side Steelhead appliance.
  - In the Current Connections report, verify the destination IP address, port 443, the Connection Count as Established (three yellow arrows on the left side of the table), SDR Enabled (three cascading yellow squares on the right side of the table), and that there is no Protocol Error (a red triangle on the right side of the table).
  - In the SSL Statistics report (on the server-side Steelhead appliance only) look for connection requests (established and failed connections), connection establishment rate, and concurrent connections.
- **Monitoring Secure Inner Channel Connections** - Use the following tools to verify that secure inner channels are in use for the selected application traffic types:
  - Choose Reports > Networking > Current Connections. Look for the **Lock** icon and three yellow arrows, which indicate the connection is encrypted and optimized. If the Lock icon is not visible or is dimmed, click the magnifying glass to view a failure reason that explains why the Steelhead appliance is not using the secure inner channel to encrypt the connection. If there is a red protocol error, click the magnifying glass to view the reason for the error.
  - Search the client-side and server-side Steelhead appliance logs for ERR and WARN.
  - Check that both Steelhead appliances appear in the white peering trust list on the client-side and server-side Steelhead appliances, indicating that they trust each other.

For details on the secure inner channel, see "Secure Inner Channel Overview" on page 231.

- **SSL Issues with Internet Explorer 6 and Oracle R12** - Previously, RiOS fixed a vulnerability found in CBC-based ciphers prior to versions 0.9.6e by inserting an empty frame on the wire to avoid a Chosen Plaintext Attack on cipher-block chaining (CBC) ciphers. Some versions of client and server applications do not understand the insertion of empty frames into the encrypted stream and close the connection when they detect these frames. Therefore, RiOS no longer inserts empty frames by default. Examples of applications that close the connection when they detect these empty frames are IE6 and Oracle R12. Sharepoint under IIS has also exhibited this behavior.

The failure occurs when the SSL application fails to understand the data payload when either the client or server is using a CBC mode as the chosen cipher. This can be with DES, AES, or 3DES using CBC. Note that when Steelhead appliances are deployed, the chosen cipher can be different than when the client is negotiating directly with the SSL server.

Important: Because current Web browsers do not protect themselves from this vulnerability, Steelhead appliances are no less secure than other vendor's appliances. From a security perspective, fixing this vulnerability is the responsibility of a server, not a patched client.

To determine if the Steelhead appliances are inserting empty frames to avoid an attack, capture TCP dumps on the server-side Steelhead LAN interface and look at the Server Hello message where it displays the selected cipher. Verify that DES, AES, or 3DES is the cipher. Also, check for the existence of 32-byte length SSL application data (this is the empty frame) on the LAN traces followed by an SSL Alert.

To change the default and insert empty frames, enter the CLI command **no protocol ssl bug-work-around dnt-insrt-empty**.

Note: For details on the vulnerability, see
http://www.openssl.org/~bodo/tls-cbc.txt.

Configuring SSL Main Settings

You can configure SSL optimization in the Configure > Optimization > SSL Main Settings page. Enabling SSL allows you to accelerate encrypted traffic (for example, HTTPS).

The Steelhead appliance securely decrypts, optimizes, and then re-encrypts SSL traffic. To configure SSL support, you do not need to make configuration changes on the client and the server—clients continue connecting to the same server name or IP address.

To enable SSL

1. Choose Configure > Optimization > SSL Main Settings to display the SSL Main Settings page.

Figure 6-2. **SSL Main Settings Page**

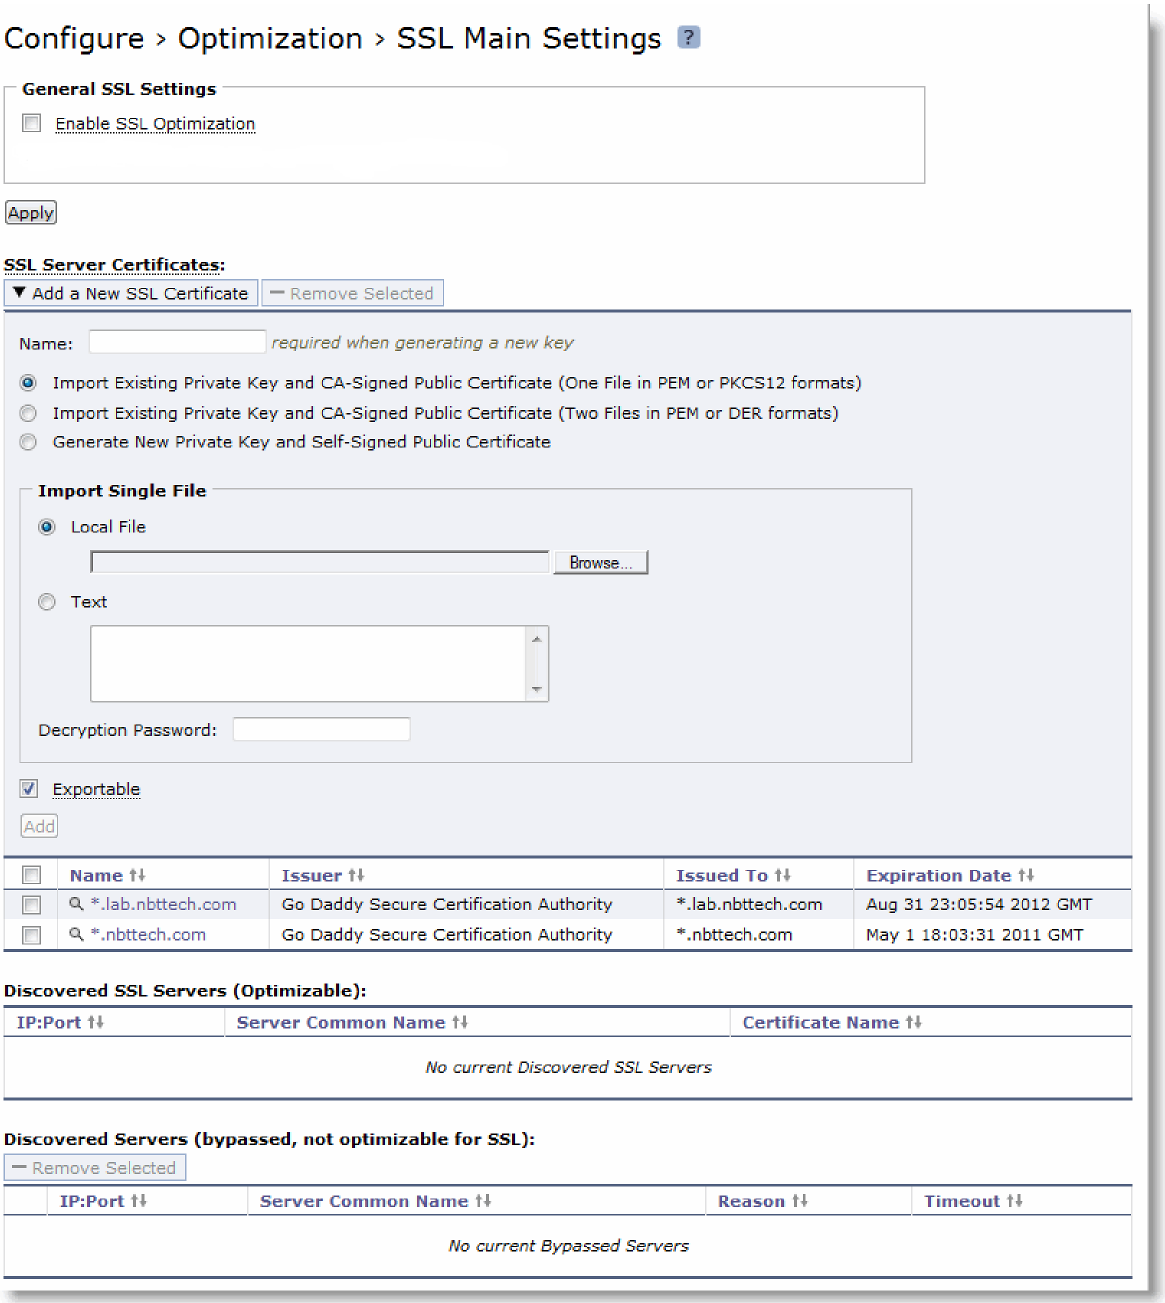

Configure > Optimization > SSL Main Settings ?
General SSL Settings
Apply
SSL Server Certificates:
▼ Add a New SSL Certificate
— Remove Selected
Import Single File
Browse...
Add
Discovered SSL Servers (Optimizable):
Discovered Servers (bypassed, not optimizable for SSL):
— Remove Selected

□ Enable SSL Optimization

Name: *required when generating a new key*

● Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)
○ Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)
○ Generate New Private Key and Self-Signed Public Certificate

● Local File
○ Text

Decryption Password:

☑ Exportable

□	Name	Issuer	Issued To	Expiration Date
□	*.lab.nbttech.com	Go Daddy Secure Certification Authority	*.lab.nbttech.com	Aug 31 23:05:54 2012 GMT
□	*.nbttech.com	Go Daddy Secure Certification Authority	*.nbttech.com	May 1 18:03:31 2011 GMT

IP:Port	Server Common Name	Certificate Name
No current Discovered SSL Servers		

	IP:Port	Server Common Name	Reason	Timeout
	No current Bypassed Servers			

2. Under General SSL Settings, complete the configuration on both the client-side and server-side Steelhead appliances as described in the following table.

Control	Description
Enable SSL Optimization	Enables SSL optimization, which accelerates applications that use SSL to encrypt traffic. By default, this option is disabled. You can choose to enable SSL optimization only on certain sessions (based on source and destination addresses, subnets, and ports), or on all SSL sessions, or on no SSL sessions at all. An SSL session that is not optimized simply passes through the Steelhead appliance unmodified.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

5. You must restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Configuring SSL Server Certificates

You configure an SSL Server Certificate in the Configure > Optimization > SSL Main Settings page.

RiOS v6.0 or later simplifies the SSL configuration process because it eliminates the need to add each server certificate individually. Prior to v6.0, you need to provide an IP address, port, and certificate to enable SSL optimization for a server. In RiOS v6.0 or later, you need only add unique certificates to a Certificate Pool on the server-side Steelhead appliance. When a client initiates an SSL connection with a server, the Steelhead appliance matches the common name of the server's certificate with one in its certificate pool. If it finds a match, it adds the server name to the list of Discovered servers that are optimizable and all subsequent connections to that server are optimized.

If it does not find a match, it adds the server name to the list of Bypassed servers and all subsequent connections to that server are not optimized. The Discovered and Bypassed server lists appear in the SSL Main Settings page.

The Steelhead appliance supports RSA private keys for peers and SSL servers.

Important: Optimization does not occur for a particular server IP address and port unless a suitable proxy server certificate is configured on the server-side Steelhead appliance.

When you configure the back-end server proxy certificate and key on the server-side Steelhead appliance, if you choose not to use the actual certificate for the back-end server and key, you can use a self-signed certificate and key or another CA-signed certificate and key. If you have a CA-signed certificate and key, import it.

If you do not have a CA-signed certificate and key, you can add the proxy server configuration with a self-signed certificate and key, back up the private key, generate a CSR, have it signed by a CA, and import the newly CA-signed certificate and the backed up private key.

Tip: To back up a single pair of certificate and key (that is, the peering certificate and key pair and a single certificate and key for the server) use the **Export (in PEM format only)** option. Make sure you check **Include Private Key** and enter the encryption password. Save the exported file that contains the certificate and the encrypted private key. For details, see "Configuring Secure Peers" on page 231.

You can also simply use the generated self-signed certificate and key, but it might be undesirable because the clients by default do not trust it, requiring action from the end-users.

To add an SSL server certificate

1. Choose Configure > Optimization > SSL Main Settings to display the SSL Main Settings page.

Figure 6-3. SSL Main Settings Page

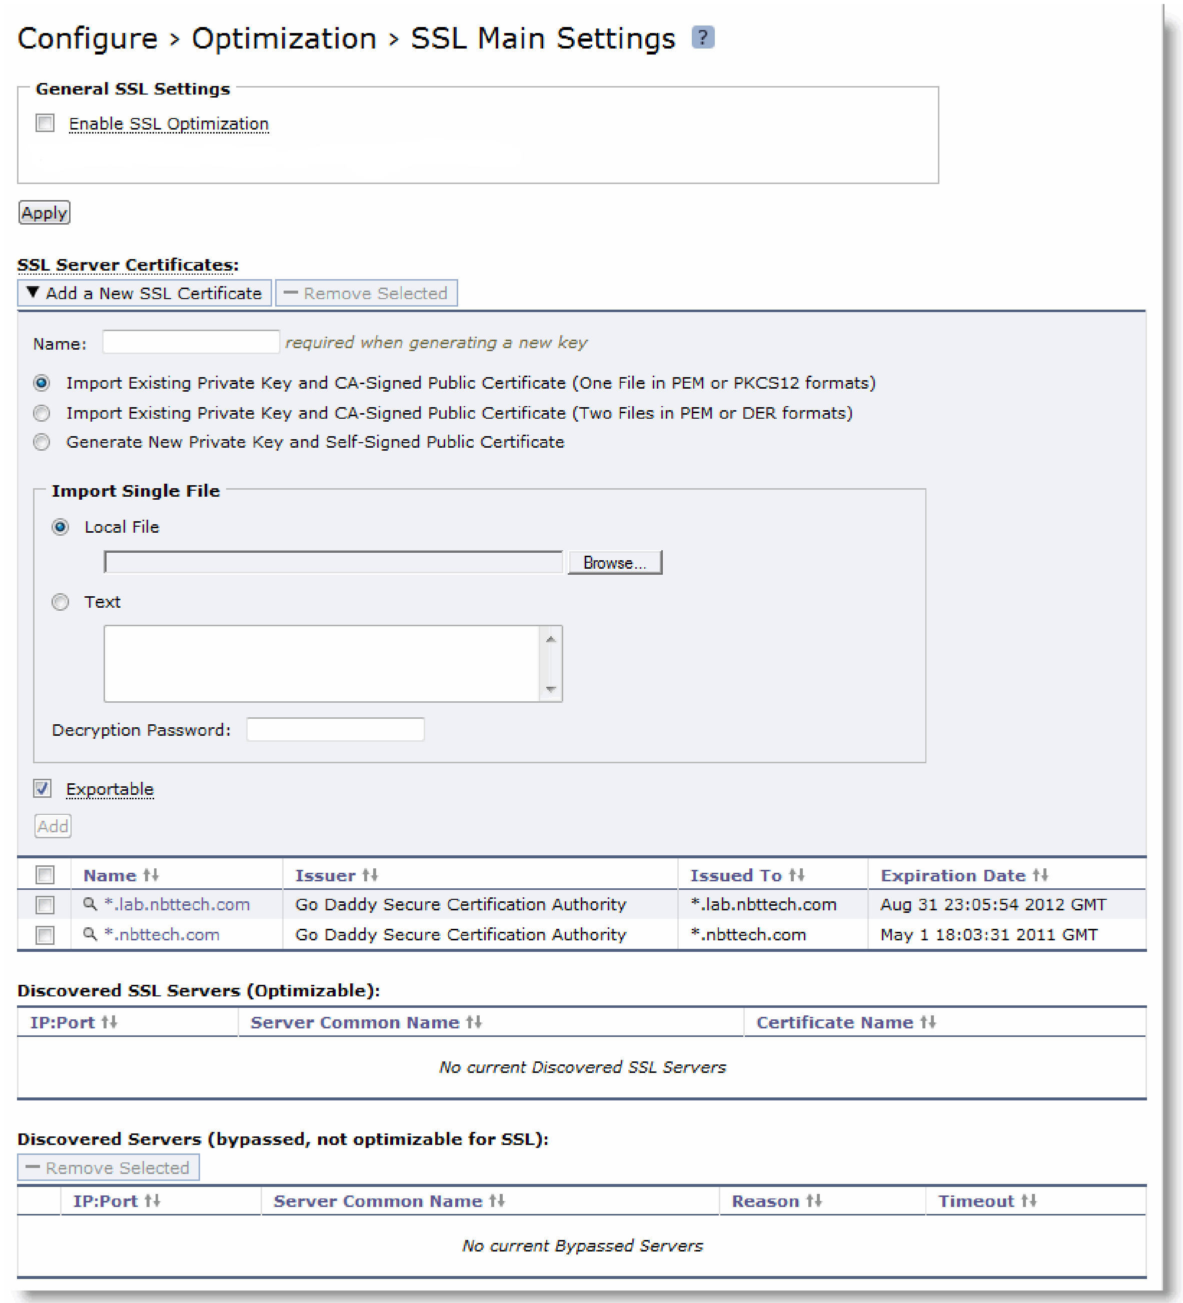

Configure > Optimization > SSL Main Settings ?

General SSL Settings

□ Enable SSL Optimization

Apply

SSL Server Certificates:

▼ Add a New SSL Certificate — Remove Selected

Name: ________ *required when generating a new key*

● Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)

○ Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)

○ Generate New Private Key and Self-Signed Public Certificate

Import Single File

● Local File

________ Browse...

○ Text

Decryption Password: ________

☑ Exportable

Add

□	Name	Issuer	Issued To	Expiration Date
□	*.lab.nbttech.com	Go Daddy Secure Certification Authority	*.lab.nbttech.com	Aug 31 23:05:54 2012 GMT
□	*.nbttech.com	Go Daddy Secure Certification Authority	*.nbttech.com	May 1 18:03:31 2011 GMT

Discovered SSL Servers (Optimizable):

IP:Port	Server Common Name	Certificate Name

No current Discovered SSL Servers

Discovered Servers (bypassed, not optimizable for SSL):

— Remove Selected

	IP:Port	Server Common Name	Reason	Timeout

No current Bypassed Servers

2. On the server-side Steelhead appliance, under SSL Server Certificates, complete the configuration as described in the following table.

Control	Description
Add a New SSL Certificate	Displays the controls to add a new server certificate.
Name	Specify a name for the proxy certificate (required when generating a certificate, leave blank when importing a certificate).
Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)	Imports the key and certificate. Select this option if the existing private key and CA-signed certificate are located in one file. The page expands displaying Private Key and CA-Signed Public Certificate controls for browsing to the key and certificate files or a text box for copying and pasting the key and certificate. The private key is required regardless of whether you are adding or updating.
	Local File - Browse to the local file.
	Text - Paste the contents of the file.
	Decryption Password - Specify the decryption password, if necessary.
	Exportable - Allows the certificate and server key to be exported. This is the default setting. Disable this setting to make sure the private key does not leave the Steelhead appliance.
	Add - Adds the server.
Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)	Imports the key and certificate. Select this option if the existing private key and CA-signed certificate are located in two files. The page expands displaying Private Key and CA-Signed Public Certificate controls for browsing to the key and certificate files or text boxes for copying and pasting the keys and certificates.
Generate New Private Key and Self-Signed Public Certificate	Select this option to generate a new private key and self-signed public certificate.
	Cipher Bits - Select the key length from the drop-down list.
	Common Name - Specify the common name of a certificate. To facilitate configuration, you can use wildcards in the name; for example, *.nbttech.com. If you have three origin servers using different certificates such as webmail.nbttech.com, internal.nbttech.com, and marketingweb.nbttech.com, on the server-side Steelhead appliances, all three server configurations may use the same certificate name *.nbttech.com.
	Organization Name - Specify the organization name (for example, the company).
	Organization Unit Name - Specify the organization unit name (for example, the section or department).
	Locality - Specify the city.
	State (no abbreviations) - Specify the state.
	Country (2-letter code) - Specify the country (2-letter code only).
	Email Address - Specify the email address of the contact person.
	Validity Period (Days) - Specify how many days the certificate is valid.

Control	Description
Add	Adds the server certificate.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

Configuring SSL Certificate Authorities

You add SSL Certificate Authorities (CA) in the Configure > Optimization > Certificate Authorities page.

A CA is a third-party entity in a network which issues digital certificates and manages security credentials and public keys for message encryption. A CA issues a public key certificate which states that the CA attests that the public key contained in the certificate belongs to the person, organization, server, or other entity noted in the certificate. The CA verifies applicant credentials, so that relying parties can trust the information in the CA certificates. If you trust the CA and can verify the CA signature, then you can also verify that a certain public key does indeed belong to whomever is identified in the certificate.

Important: Before adding a CA, it is critical to verify that it is genuine; a malicious CA can compromise network security by signing fake certificates.

You might need to add a new CA in the following situations:

- Your organization has an internal CA that signs the certificates or peering certificates for the back-end server.
- The server certificates are signed by an intermediate or root CA unknown to the Steelhead appliance (perhaps external to the organization).
- The CA certificate included in the trusted list of the Steelhead appliance has expired or has been revoked and needs replacing.

To add SSL certificate authorities

1. On the server-side Steelhead appliance, choose Configure > Optimization > Certificate Authorities to display the Certificate Authorities page.

Figure 6-4. **Certificate Authorities Page**

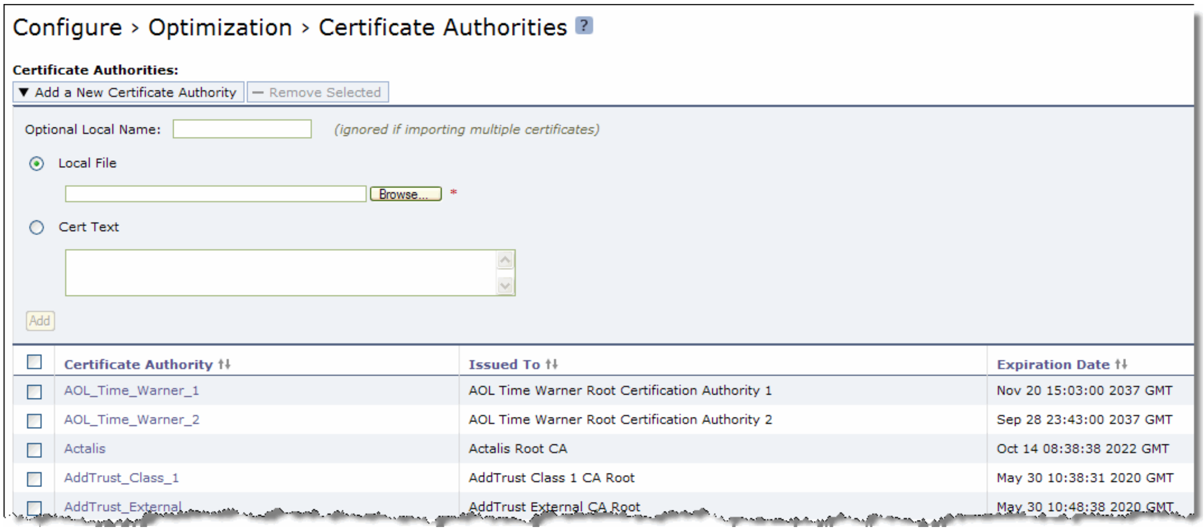

Configure > Optimization > Certificate Authorities

Certificate Authorities:

Optional Local Name: (ignored if importing multiple certificates)

● Local File

○ Cert Text

□	Certificate Authority	Issued To	Expiration Date
□	AOL_Time_Warner_1	AOL Time Warner Root Certification Authority 1	Nov 20 15:03:00 2037 GMT
□	AOL_Time_Warner_2	AOL Time Warner Root Certification Authority 2	Sep 28 23:43:00 2037 GMT
□	Actalis	Actalis Root CA	Oct 14 08:38:38 2022 GMT
□	AddTrust_Class_1	AddTrust Class 1 CA Root	May 30 10:38:31 2020 GMT
□	AddTrust_External	AddTrust External CA Root	May 30 10:48:38 2020 GMT

2. Under Certificate Authorities, complete the configuration as described in the following table.

Control	Description
Add a New Certificate Authority	**Optional Local Name (ignored if importing multiple certificates)** - Specify the local name.
	Local File - Browse to the local certificate authority file.
	Cert Text - Paste the certificate authority into the text box and click **Add**.
Add	Adds the certificate authority.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

Tip: Select the Certificate Authority name to display details.

Modifying SSL Server Certificate Settings

After initial SSL server configuration, you can modify server certificate settings in the Configure > Optimization > SSL Main Settings page. You can remove a server certificate, view the server certificate details, change the server certificate and private key, export a certificate, or generate a CSR.

For details on initial SSL server configuration, see "Configuring SSL Main Settings" on page 218.

Note: After initial configuration, you might need to generate a Certificate Signing Request and import a Certificate Authority-signed certificate before activating the SSL server for optimization.

Removing or Changing an SSL Server Certificate

The following section describes how to remove or change an existing SSL server certificate.

To remove a server certificate

1. Choose Configure > Optimization > SSL Main SSL Settings to display the SSL Main Settings page.

2. Under Bypassed SSL Servers, select the certificate name you want to remove and click **Remove Selected**.

To change an SSL server certificate

1. Choose Configure > Optimization > SSL Main Settings to display the SSL Main Settings page.

2. Under SSL Server Certificates, select the certificate name.

3. Click **Modify**.

4. Complete the configuration as described in the following table.

Control	Description
Rename Certificate	Displays the controls to rename the certificate. **Name** - Specify the new certificate name. **Change** - Changes the certificate name.
Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)	Select this option if the existing private key and CA-signed certificate are located in one file. The page expands displaying Private Key and CA-Signed Public Certificate controls for browsing to the key and certificate files or a text box for copying and pasting the key and certificate. The private key is required regardless of whether you are adding or updating.
	Local File - Browse to the local file.
	Text - Paste the content of the file.
	Decryption Password - Specify the password used to decrypt, if necessary.
	Set - Changes the settings.
Import Existing Private Keys and CA-Signed Public Certificate (Two Files in PEM or DER formats)	Select this option if the existing private key and CA-signed certificate are located in two files. The page expands displaying Private Key and CA-Signed Public Certificate controls for browsing to the key and certificate files or text boxes for copying and pasting the keys and certificates. A private key is optional for existing server configurations.
Private Key	**Private Key Local File** - Browse to the local file containing the private key.
	Private Key Text - Paste the private key text.
CA-Signed Public Certificate	**Local File** - Browse to the local file.
	Cert Text - Paste the content of the certificate text file.
	Decryption Password - Specify the password used to decrypt, if necessary.
	Change - Changes the settings.
Generate New Private Key and Self-Signed Public Certificate	Select this option to generate a new private key and self-signed public certificate.
	Cipher Bits - Select the key length from the drop-down list. The default value is 1024.
	Common Name - Specify the domain name of the server.
	Organization Name - Specify the organization name (for example, the company).
	Organization Unit Name - Specify the organization unit name (for example, the section or department).
	Locality - Specify the city.
	State (no abbreviations) - Specify the state.
	Country (2-letter code) - Specify the country (2-letter code only).
	Email Address - Specify the email address of the contact person.
	Validity Period (Days) - Specify how many days the certificate is valid.
	Change - Changes the settings.

Exporting an SSL Server Certificate

The following section describes how to export an existing certificate for an SSL server.

To export an SSL server certificate

1. Choose Configure > Optimization > Main SSL Settings to display the Main SSL Settings page.

2. Under SSL Server Certificates, select the certificate name.

3. To export an existing certificate, click **Export** and complete the configuration as described in the following table.

Control	Description
Include Private Key	Includes the private key in the export.
Password/Password Confirm	Specify and confirm the encrypted password if you are including the private key (required if including the key). The password must be at least 4 characters.
Export	Exports the Steelhead appliance peering certificate and key.

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

Generating a CSR

This section describes how to generate a Certificate Signing Request (CSR) for an existing SSL server off the current private key.

To generate a CSR

1. Choose Configure > Optimization > Main SSL Settings to display the Main SSL Settings page.

2. Under SSL Server Certificates, select the certificate name.

3. Click **Generate CSR** and complete the configuration as described in the following table.

Control	Description
Common Name (required)	Specify the common name (hostname) of the peer.
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).
Locality	Specify the city.
State	Specify the state. Do not abbreviate.
Country (2-letter code)	Specify the country (2-letter code only).
Email Address	Specify the email address of the contact person.
Generate CSR	Generates the Certificate Signing Request.

4. Click **Save** to save the settings permanently.

Adding a Chain Certificate

This section describes how to add or remove a chain certificate for an existing SSL server.

To add a chain certificate

1. Choose Configure > Optimization > Main SSL Settings to display the Main SSL Settings page.

2. Under SSL Server Certificates, select the certificate name.

3. Click **Chain** and complete the configuration as described in the following table.

Control	Description
Add a New Chain Certificate	Displays the controls to add a chain certificate.
Use Existing CA	Select to use an existing certificate authority, and then select the certificate authority from the drop-down list.
Use New Certificate(s) PEM or DER formats	Select to use a new certificate.
Optional Local Name	Optionally, specify a local name for the certificate.
Local File	Browse to the local file.
Cert Text	Paste the contents of the certificate text file into the text box.
Add	Adds the chain certificate to the chain certificate list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

4. Click **Save** to save the settings permanently.

Related Topics

- “Configuring In-Path Rules” on page 28
- “Enabling Peering and Configuring Peering Rules” on page 67
- “Configuring HTTP Optimization” on page 111
- “Secure Inner Channel Overview” on page 231
- “Unlocking the Secure Vault” on page 370
- “Viewing SSL Reports” on page 434
- “Viewing the System Dumps List” on page 480

Configuring CRL Management

RiOS v6.5 provides a way to configure Certificate Revocation Lists (CRLs) for an automatically discovered CA using the Management Console. CRLs allow CAs to revoke issued certificates (for example, when the private key of the certificate has been compromised). By default, CRLs are not used in the Steelhead appliance.

A CRL is a database that contains a list of digital certificates that have been invalidated before their expiration date, including the reasons for the revocation and the names of the issuing certificate signing authorities. The CRL is issued by the CA which issues the corresponding certificates. All CRLs have a lifetime during which they are valid (often 24 hours or less).

CRLs are used when a:

- server-side Steelhead appliance verifies the certificate presented by the server in the SSL handshake between the server-side Steelhead appliance and the server.
- server-side Steelhead appliance verifies the certificate presented by the client-side Steelhead appliance in the handshake between the two Steelhead appliances for establishing a secure inner channel over the WAN.
- client-side Steelhead appliance verifies the certificate presented by the server-side Steelhead appliance in the handshake between the two Steelhead appliances for establishing a secure inner channel over the WAN.

The two types of CAs issuing CRLs are:

- Conventional CAs, which are listed in the Certificate Authorities page.
- Peering CAs, which are listed in the Trusted Entities list in the Secure Peering page.

You configure each type of CA separately.

Note: Currently, the Steelhead appliance only supports downloading CRLs from Lightweight Directory Access Protocol (LDAP) servers.

To enable CRL management

1. On the server-side Steelhead appliance, choose Configure > Optimization > CRL Management to display the CRL Management page.

Figure 6-5. CRL Management Page

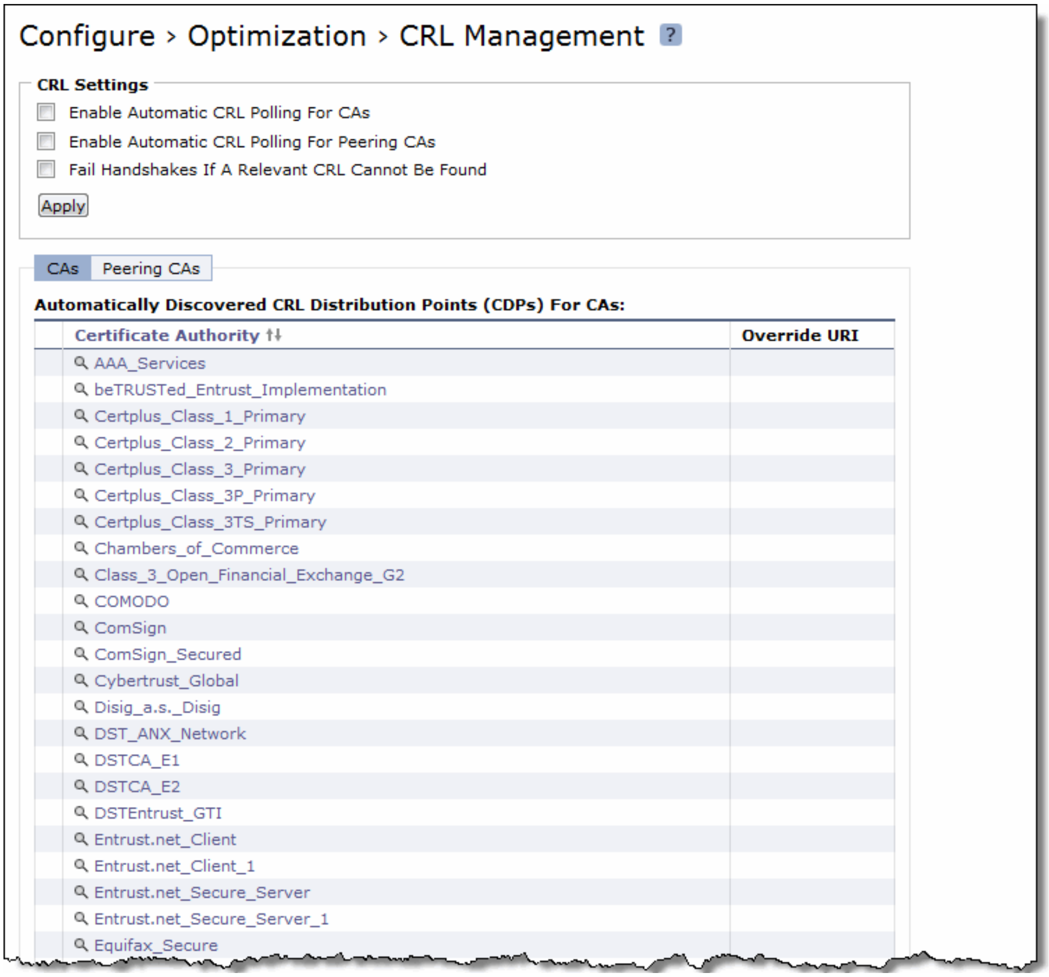

Configure > Optimization > CRL Management

CRL Settings

□ Enable Automatic CRL Polling For CAs

□ Enable Automatic CRL Polling For Peering CAs

□ Fail Handshakes If A Relevant CRL Cannot Be Found

Automatically Discovered CRL Distribution Points (CDPs) For CAs:

Certificate Authority	Override URI
AAA_Services	
beTRUSTed_Entrust_Implementation	
Certplus_Class_1_Primary	
Certplus_Class_2_Primary	
Certplus_Class_3_Primary	
Certplus_Class_3P_Primary	
Certplus_Class_3TS_Primary	
Chambers_of_Commerce	
Class_3_Open_Financial_Exchange_G2	
COMODO	
ComSign	
ComSign_Secured	
Cybertrust_Global	
Disig_a.s._Disig	
DST_ANX_Network	
DSTCA_E1	
DSTCA_E2	
DSTEntrust_GTI	
Entrust.net_Client	
Entrust.net_Client_1	
Entrust.net_Secure_Server	
Entrust.net_Secure_Server_1	
Equifax_Secure	

2. Under CRL Settings, complete the configuration as described in the following table.

Control	Description
Enable Automatic CRL Polling for CAs	Enables CRL polling and use of a CRL in handshake verifications of CA certificates. Currently, the Steelhead only supports downloading CRLs from Lightweight Directory Access Protocol (LDAP) servers.
Enable Automatic CRL Polling for Peering CAs	Configures a CRL for an automatically discovered peering CA.
Fail Handshakes If A Relevant CRL Cannot Be Found	Configures handshake behavior for a CRL. Fails the handshake verification if a relevant CRL for either a peering or server certificate cannot be found.

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

Managing CRL Distribution Points (CDPs)

You can view, override, or remove CRL distribution points (CDPs) for CAs in the Configure > Optimization > CRL Management page.

A CDP is a field within a certificate containing information that describes where to retrieve the CRL for the certificate.

To view a list of CDPs for a CA

1. On the server-side Steelhead appliance, choose Configure > Optimization > CRL Management to display the CRL Management page.

2. Select the CAs tab to view conventional CAs or the Peering CAs tab to view secure CAs.

   The Automatically Discovered CRL Distribution Points table displays a list of CAs for which CDPs have been automatically discovered. Because not all CAs have CDPs, this list is a subset of the main CA list in the Configure > Certificate Authorities page or a subset of the CAs in the Peering Trust table in the Configure > Optimization > Secure Peering page.

Figure 6-6. **CRL Management Page - Automatically Discovered CDPs**

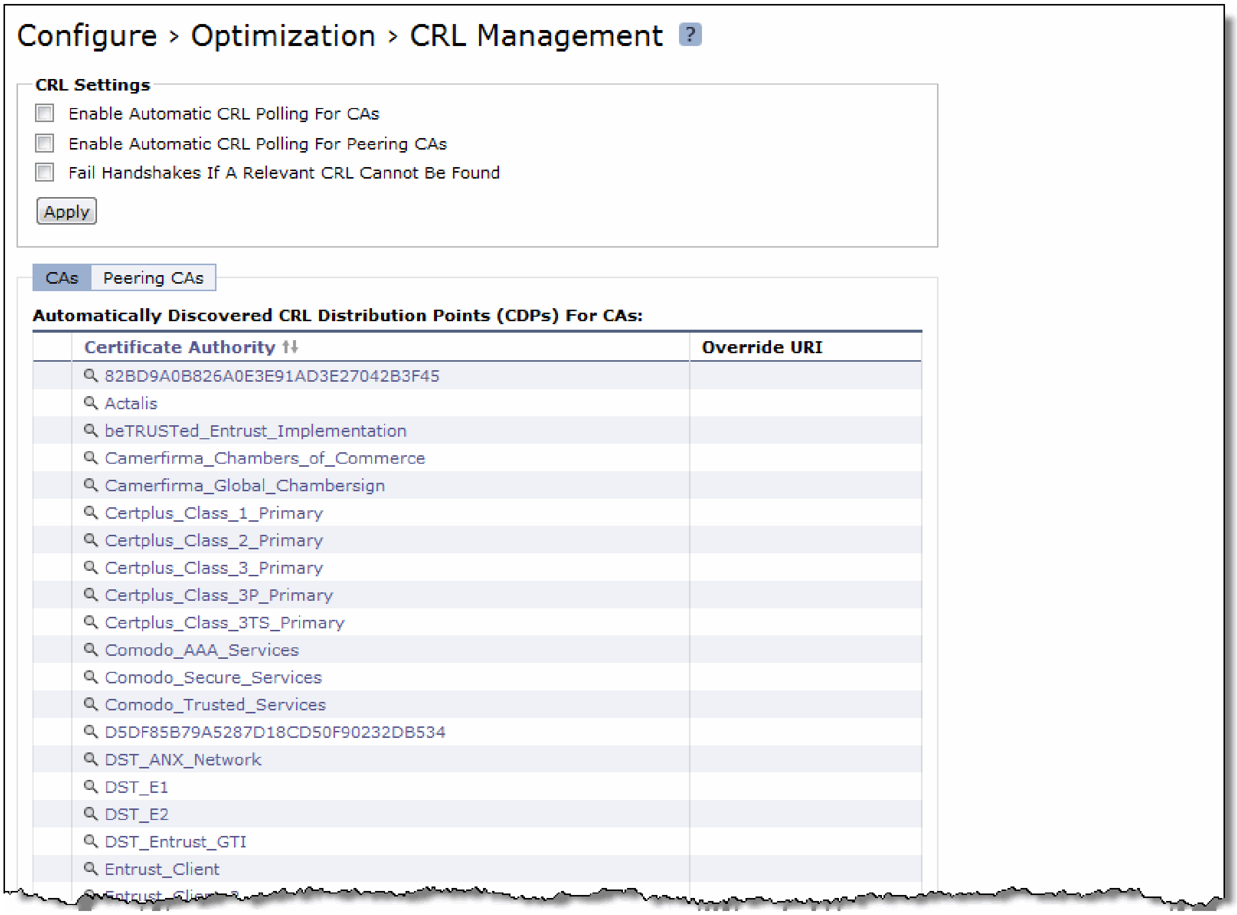

Configure > Optimization > CRL Management ?

CRL Settings

□ Enable Automatic CRL Polling For CAs

□ Enable Automatic CRL Polling For Peering CAs

□ Fail Handshakes If A Relevant CRL Cannot Be Found

Apply

CAs | Peering CAs

Automatically Discovered CRL Distribution Points (CDPs) For CAs:

Certificate Authority ⇅	Override URI
82BD9A0B826A0E3E91AD3E27042B3F45	
Actalis	
beTRUSTed_Entrust_Implementation	
Camerfirma_Chambers_of_Commerce	
Camerfirma_Global_Chambersign	
Certplus_Class_1_Primary	
Certplus_Class_2_Primary	
Certplus_Class_3_Primary	
Certplus_Class_3P_Primary	
Certplus_Class_3TS_Primary	
Comodo_AAA_Services	
Comodo_Secure_Services	
Comodo_Trusted_Services	
D5DF85B79A5287D18CD50F90232DB534	
DST_ANX_Network	
DST_E1	
DST_E2	
DST_Entrust_GTI	
Entrust_Client	

3. Select the CA name in the Automatically Discovered CRL Distribution Points table.

   If a CDP has been manually overridden for the CA, it appears in the override column.

To view CDP details and access history

1. Click the **Expand** icon next to the CDP name.
2. To see the CDP access points, select the CDP Details tab.

   Use the scroll bar to view the entire address.
3. To see the Certificate Revocation List, select the View CRL tab.

   The display includes a CRL Access History list.
4. Click **Check for Update** to refresh the display.

To override an existing CDP

Perform this task to manually override the existing CDP for a certificate with an LDAP server specification.

1. Click **Add Manual Override**.
2. Select a CA name from the drop-down list.
3. Specify the CDP Uniform Resource Indicator (URI) for an LDAP server; for example,

   ```
   http://ca.actalis.it/crl/root/getCRL
   ```
4. Click **Add**.

Configuring Secure Peers

You configure secure peers in the Configure > Optimization > Secure Peering (SSL) page. In RiOS v6.0 and later, secure, encrypted peering extends beyond traditional SSL traffic encryption. In addition to SSL-based traffic like HTTPS that always needs a secure connection between the client-side and the server-side Steelhead appliance, you can also secure other types of traffic such as:

- MAPI-encrypted, SMB1, and SMB2-signed traffic.
- all other traffic that inherently does not require a secure connection.

Important: To encrypt and optimize the connection for *non-SSL* traffic, both the client-side and server-side Steelhead appliances must be running RiOS v6.0 or later. You must configure secure peering on both the client-side and the server-side Steelhead appliances and the settings must match on both sides. RiOS v4.0 and later encrypts and optimizes SSL traffic.

Secure Inner Channel Overview

Each Steelhead appliance is manufactured with its own self-signed certificate and private key which uniquely identify that Steelhead. The secure inner channel setup process begins with the peer Steelhead appliances authenticating each other by exchanging certificates and negotiating a separate encryption key for each intercepted connection. Next, the Steelhead appliances create corresponding inner connections for all outer connections between the client and the client-side Steelhead appliance and between the server and the server-side Steelhead appliance.

Peers are detected the first time a client-side Steelhead appliance attempts to connect to the server. The optimization service bypasses this initial connection and does not perform data reduction, but rather uses it to detect peers and populate the peer entry tables. On both Steelhead appliances, an entry appears in a peering list with the certificate of the other peer and identifying information such as IP address and hostname. You can then accept or decline the trust relationship with each Steelhead appliance requesting a secure inner channel.

Once the appliances trust each other, they send encrypted data between themselves over secure inner connections matching the outer connections of the selected traffic types. The trust relationship between the Steelheads is bi-directional; the client-side Steelhead appliance trusts the server-side Steelhead appliance, and vice versa.

Riverbed recommends using the secure inner channel in place of IPSec encryption to secure traffic.

Enabling Secure Peers

This section describes how to secure traffic between client-side and server-side Steelhead appliances.

Tip: You rarely need to replace a self-signed certificate on a Steelhead; however, if you do, add the CA's certificate to the peering trust section so each Steelhead can verify the peer certificate for its peers. For details, see "Configuring Peer Trust" on page 234.

To enable secure peering

1. If you are securing SMB-signed traffic or encrypted MAPI traffic, enable the CIFS or MAPI protocol as follows.

   - Choose Configure > Optimization > CIFS and select Enable SMB Signing.

   —or—

   - Choose Configure > Optimization > MAPI and select Enable Encrypted Optimization.

2. On both the server-side and client-side Steelhead appliances, choose Configure > Optimization > Secure Peering (SSL) to display the Secure Peering (SSL) page.

Figure 6-7. **Secure Peering (SSL) Page**

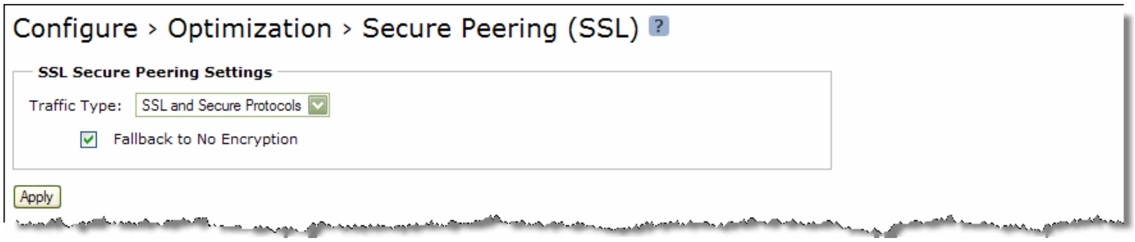

Configure > Optimization > Secure Peering (SSL)

SSL Secure Peering Settings

Traffic Type: SSL and Secure Protocols

☑ Fallback to No Encryption

3. Under SSL Secure Peering Settings, complete the configuration as described in the following table.

Control	Description
Traffic Type	Select one of the following traffic types from the drop-down list: • **SSL Only** - The peer client-side Steelhead appliance and the server-side Steelhead appliance authenticate each other and then encrypt and optimize all SSL traffic; for example, HTTPS traffic on port 443. This is the default setting. • **SSL and Secure Protocols** - The peer client-side Steelhead appliance and the server-side Steelhead appliance authenticate each other and then encrypt and optimize all traffic traveling over the following secure protocols: SSL, SMB-signed, and encrypted MAPI. When you select this traffic type, SMB-signing and MAPI encryption must be enabled. Enabling this option requires an optimization service restart. • **All** - The peer client-side Steelhead appliance and the server-side Steelhead appliance authenticate each other and then encrypt and optimize all traffic. Only the optimized traffic is secure; pass-through traffic is not. Enabling this option requires an optimization service restart.
Fallback to No Encryption	Specifies that the Steelhead appliance optimizes but does not encrypt the connection when it is unable to negotiate a secure, encrypted inner channel connection with the peer. This is the default setting. Enabling this option requires an optimization service restart. **Important:** Riverbed strongly recommends enabling this setting on both the client-side and the server-side Steelhead appliances, especially in mixed deployments where one Steelhead appliance is running RiOS v6.0 and the other Steelhead is running an earlier RiOS version. This option applies only to non-SSL traffic and is unavailable when you select SSL Only as the traffic type. Clear the check box to pass through connections that do not have a secure encrypted inner channel connection with the peer. Use caution when disabling this setting, as doing so specifies that you strictly do not want traffic optimized between non-secure Steelheads. Consequently, when this setting is disabled connections might be dropped. For example, consider a configuration with a client-side Steelhead running RiOS v5.5.x or earlier and a server-side Steelhead running RiOS v6.0. When this setting is disabled on the server-side Steelhead and All is selected as the traffic type, it will not optimize the connection when a secure channel is unavailable, and might drop it.

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

6. If you have changed an encryption setting, you need to restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Important: The Steelhead appliance supports RSA private keys for peers and SSL servers.

Configuring Peer Trust

The first time a client-side Steelhead appliance attempts to connect to the server, the optimization service detects peers and populates the peer entry tables. On both Steelhead appliances, an entry appears in a peering list with the information and certificate of the other peer. A peer list provides you with the option of accepting or declining the trust relationship with each Steelhead appliance requesting a secure inner channel. The self-signed peer lists are designated by the following color categories:

- **White** - Lists all trusted Steelhead appliances. When you select Trust for a peer in a black or gray list, the public key of the Steelhead appliance peer is copied into the white list of the local Steelhead appliance trusted host. The list includes the peer expiration date, IP address, and hostname.
- **Black** - Lists all untrusted Steelhead appliances. When you select Do Not Trust for a peer in a white or gray list, the public key of the Steelhead appliance peer is copied into the black list of the local Steelhead appliance untrusted host. The list includes the peer expiration date, IP address, and hostname.
- **Gray** - Lists all Steelhead appliances of unknown status. This list serves as a temporary holding place for all discovered peer Steelhead appliances that are attempting to establish a secure inner channel. The list includes the peer expiration date, IP address, and hostname. You can select one of the following actions to change the status of the peer and move it to the white or black lists: Trust, Do Not Trust, or Remove.

Note: When a self-signed peer has already been added to a peering trust list manually, the SSL server recognizes it upon the first connection from that peer and automatically places it in the white list (without action by the administrator). The certificate that was previously copied and pasted (or imported) into the trusted list is not removed.

The Configure > Optimization > Secure Peering (SSL) page also provides you with the following options for configuring peer certificates and Mobile Controller trust:

- **Peering Trust** - Add and view the following types of entities:
  - Certificates of trusted peers.
  - Certificates of trusted Certificate Authorities (CAs) that may sign certificates for peers.
- **SCEP Peering Trust** - Add and view trusted SCEP entities.
- **Mobile Trust** - Add and view trusted Steelhead Mobile Controller entities that may sign certificates for Steelhead Mobile Clients.

To configure SSL peers

1. Choose Configure > Optimization > Secure Peering (SSL) to display the Secure Peering (SSL) page.

Figure 6-8. **Secure Peering (SSL) Page**

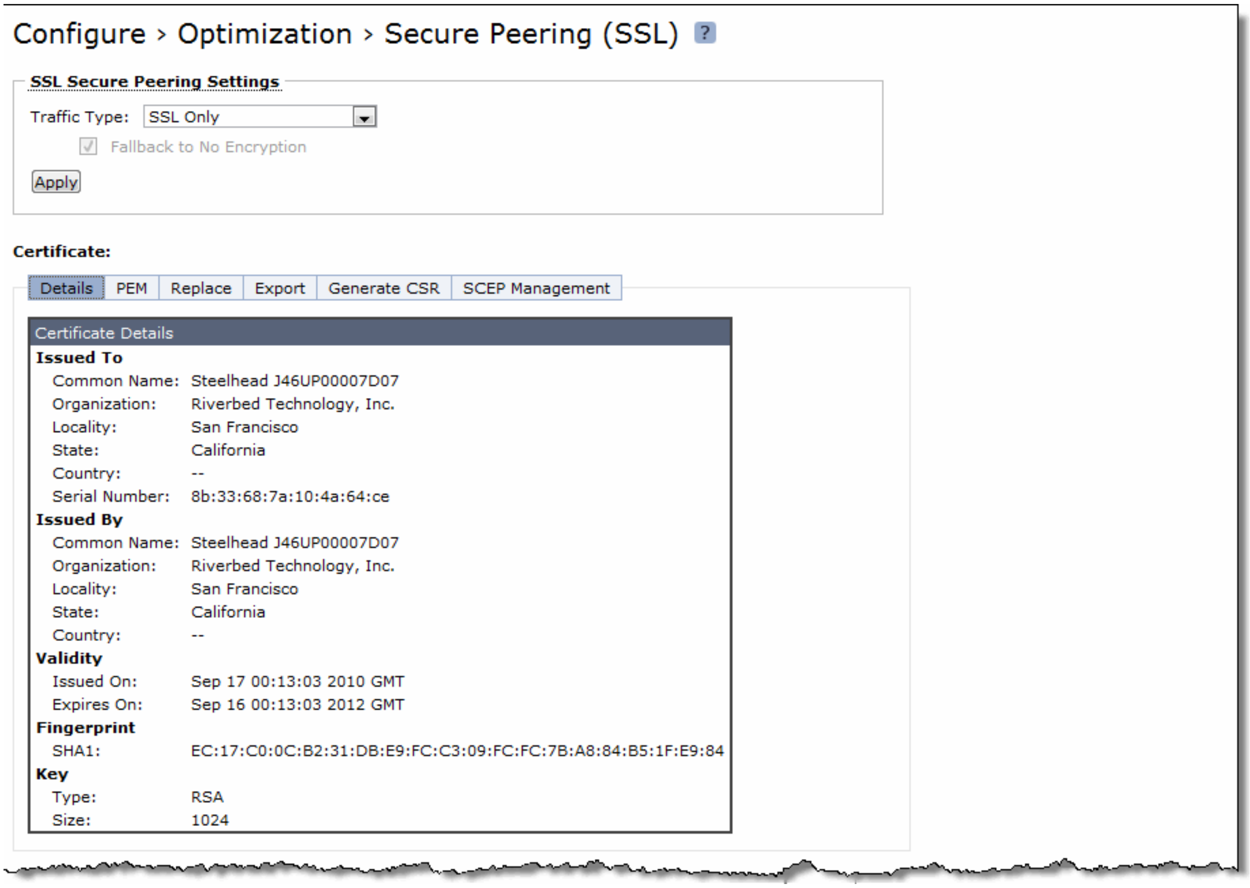

Configure > Optimization > Secure Peering (SSL)

SSL Secure Peering Settings

Traffic Type: SSL Only

☑ Fallback to No Encryption

Certificate Details

Issued To

Common Name:	Steelhead J46UP00007D07
Organization:	Riverbed Technology, Inc.
Locality:	San Francisco
State:	California
Country:	--
Serial Number:	8b:33:68:7a:10:4a:64:ce

Issued By

Common Name:	Steelhead J46UP00007D07
Organization:	Riverbed Technology, Inc.
Locality:	San Francisco
State:	California
Country:	--

Validity

Issued On:	Sep 17 00:13:03 2010 GMT
Expires On:	Sep 16 00:13:03 2012 GMT

Fingerprint

SHA1:	EC:17:C0:0C:B2:31:DB:E9:FC:C3:09:FC:FC:7B:A8:84:B5:1F:E9:84

Key

Type:	RSA
Size:	1024

The Steelhead identity certificate details appear, as described in the following table.

Control	Description
Issued To/Issued By	**Common Name** - Specifies the common name of the certificate authority.
	Organization - Specifies the organization name (for example, the company).
	Organization Unit - Specifies the organization unit name (for example, section or department).
	Locality - Specifies the city.
	State - Specifies the state.
	Country - Specifies the country.
	Serial Number - Specifies the serial number (Issued To, only).
Validity	**Issued On** - Specifies the date the certificate was issued.
	Expires On - Specifies the date the certificate expires.
Fingerprint	Specifies the SSL fingerprint.

Control	Description
Key	**Type** - Specifies the key type.
	Size - Specifies the size in bytes.

2. To replace an existing certificate, Under Certificate, click **Replace** and complete the configuration as described in the following table.

Control	Description
Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)	Click this option if the existing private key and CA-signed certificate are located in one file. The page displays a Private Key and CA-Signed Public Certificate control for browsing to the key and certificate files or a text box for copying and pasting the key and certificate. **Note**: The private key is required.
	Local File - Browse to the local file.
	Text - Paste the text content of the file into the text box.
	Decryption Password - Specify the decryption password, if necessary.
	Set - Sets the peer.
Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)	Select this option if the existing private key and CA-signed certificate are located in two files. The page displays a Private Key and CA-Signed Public Certificate control for browsing to the key and certificate, or a text box for copying and pasting the key and certificate. **Note:** Importing the private key is optional.
Generate New Private Key and Self-Signed Public Certificate	Select to generate a new private key and self-signed public certificate.
	Cipher Bits - Select the key length from the drop-down list. The default value is 1024.
	Common Name (required) - Specify the hostname of the peer.
	Organization Name - Specify the organization name (for example, the company).
	Organization Unit Name - Specify the organization unit name (for example, the section or department).
	Locality - Specify the city.
	State (no abbreviations) - Specify the state.
	Country (2-letter code) - Specify the country (2-letter code only).
	Email Address - Specify the email address of the contact person.
	Validity Period (Days) - Specify how many days the certificate is valid. The default value is 730.

3. To export an existing certificate, under Certificate, click **Export** and complete the configuration as described in the following table.

Control	Description
Password/Password Confirm	Specify and confirm the encrypted password if you are including the private key (required if including key). The password must be at least 4 characters long.
Include Private Key	Includes the private key in the export.
Export	Exports the Steelhead appliance peering certificate and key.

4. To generate a CSR, under Certificate, click **Generate CSR** and complete the configuration as described in the following table.

Control	Description
Common Name (required)	Specify the common name (hostname) of the peer.
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).
Locality	Specify the city.
State	Specify the state. Do not abbreviate.
Country (2-letter code)	Specify the country (2-letter code only).
Email Address	Specify the email address of the contact person.
Generate CSR	Generates the Certificate Signing Request.

5. To use SCEP to manage the certificate, under Certificate, click **SCEP Management** and complete the configuration as described in the following table.

Control	Description
URL	Specify the URL of the SCEP responder. Use the following format: http://host[:port]/path[/to/service] For example: `http:examplehost:1212/pathtoservice`
Maximum Number of Polls	Specify the maximum number of polls before the Steelhead appliance cancels the enrollment. The peering certificate is not modified. The default value is 5. A poll is a request to the server for an enrolled certificate by the Steelhead appliance. The Steelhead appliance polls only if the server responds with pending. If the server responds with fa**il** then the Steelhead appliance does not poll.
Poll Period	Specify the poll frequency in minutes. The default value is 5.
Challenge Passphrase	Specify the challenge password phrase.

Control	Description
Enable Auto Enrollment	Enables automatic re-enrollment of a certificate to be signed by a CA using SCEP. • **Expiration Threshold** - Specify the amount of time (in days) to schedule re-enrollment before the certificate expires. The range is 1 to 60 days. The default value is 30 days.
Update SCEP Settings	Updates the SCEP settings.

6. To add or remove a Trusted entity, under Peering Trust, complete the configuration as described in the following table.

Control	Description
Add a New Trusted Entity	Displays the controls for adding trusted entities.
Trust Existing CA	Select an existing CA from the drop-down list.
Trust New Certificate	Adds a new CA or peer certificate. The Steelhead appliance supports RSA and DSA for peering trust entities.
Optional Local Name	Optionally, specify a local name for the entity (for example, the fully qualified domain name).
Local File	Browse to the local file.
Cert Text	Paste the content of the certificate text file into the text box.
Add	Adds the trusted entity (or peer) to the trusted peers list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

7. To add or remove a SCEP Trusted entity, under SCEP Peering Trust, complete the configuration as described in the following table.

Control	Description
Add a New SCEP Entity	Displays the controls for adding a trusted SCEP entity.
Peering Trust	Select a peering trust from the drop-down list.
Add	Adds the trusted entity (or peer) to the trusted peers list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

8. To add or remove a Steelhead Mobile Controller trusted entity, under Mobile Trust, complete the configuration as described in the following table.

Control	Description
Add a New Mobile Entity	Displays the controls for adding a trusted Steelhead Mobile Controller entity.
Optional Local Name	Optionally, specify a local name for the entity (for example, the fully qualified domain name).
Local File	Browse to the local file.
Cert Text	Paste the content of the certificate text file into the text box.

Control	Description
Add	Adds the trusted entity (or peer) to the trusted peers list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

9. To change the trust status of a self-signed peer and move it to another list, or to remove a peer from a list, click the down arrow in the **Actions** drop-down list and complete the configuration as described in the following table.

   The white, gray, and black peering lists sort the peers by IP address.

Figure 6-9. A self-signed peering white list.

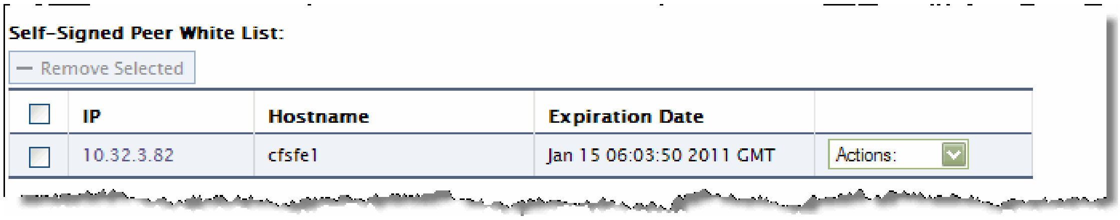

Self-Signed Peer White List:

— Remove Selected

□	IP	Hostname	Expiration Date	
□	10.32.3.82	cfsfe1	Jan 15 06:03:50 2011 GMT	Actions:

Important: Before moving a peer from the gray list to the trusted peers white list, it is critical to verify that the certificate fingerprint does indeed belong to a peer Steelhead appliance, particularly to avoid the potential risk of a man-in-the-middle attack.

Control	Description
Trust	Changes the peer Steelhead appliance to a trusted entity. The Steelhead appliance automatically finds all Steelhead appliances in your deployment and lists them in the gray list. When a self-signed peer becomes a trusted entity it moves to the white list.
Do Not Trust	Changes the self-signed peer from a trusted entity to an untrusted entity. The Steelhead appliance automatically finds all Steelhead appliances in your deployment and lists them by IP address in the gray list. When a self-signed peer becomes an untrusted entity it moves to the black list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

Note: When the same certificate appears in both the trusted entity and a self-signed peer list, deleting the certificate from one list automatically deletes it from the other.

10. Click **Apply** to save your settings to the running configuration.

11. Click **Save** to save your settings permanently.

12. Restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Verifying the Secure Inner Channel Connections

This section describes what happens when a secure inner channel cannot be established for traffic between Steelhead appliances running v6.0 or later and how to verify whether connections are using a secure inner channel.

When the Steelhead appliances are configured to use secure inner channels for **S**SL traffic only or All optimized traffic:

- The first connection that runs into a failure is passed through without optimization. This connection appears as established in the Current Connections report, but it is flagged with a protocol error.
- For up to five minutes all follow-on or subsequent connections are passed through. These follow-on connections appear as pass-through in the Current Connections report. You can click the magnifying glass for details on the pass-through reason.

When the Steelhead appliances are configured to use secure inner channels for SSL and Secure Protocols:

- The first secure protocol connection (either encrypted MAPI or SMB Signed) that runs into a failure is passed through without optimization if Fallback to No Encryption is disabled. See "Enabling Secure Peers" on page 232.
- The first SSL connection that runs into a failure is passed through without optimization. This connection appears as established in the Current Connections report, but it is flagged with a protocol error.
- For up to five minutes all follow-on or subsequent connections are passed-through.

To verify that the secure inner channel is encrypting and optimizing traffic, choose Reports > Networking > Current Connections. Look for the **Lock** icon and three yellow arrows, which indicate the connection is encrypted and optimized. If the **Lock** icon is not visible, or is dimmed, click the magnifying glass to view a failure reason that explains why the Steelhead appliance is not encrypting the connection. If there is a red protocol error, click the magnifying glass to view the reason for the error. For details, see "Viewing Current Connections" on page 385 and "Verifying SSL and Secure Inner Channel Optimization" on page 216.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring CIFS Optimization" on page 97
- "Configuring MAPI Optimization" on page 123
- "Unlocking the Secure Vault" on page 370
- "Viewing SSL Reports" on page 434
- "Viewing the System Dumps List" on page 480

Configuring Advanced and SSL Cipher Settings

This section describes the SSL advanced settings you can use to expedite SSL configurations, improve performance for short-lived SSL connections, and configure SSL cipher settings.

It includes the following sections

"Setting Advanced SSL Options" on page 241

"Configuring SSL Cipher Settings" on page 245

“Performing Bulk Imports and Exports” on page 248

Setting Advanced SSL Options

You can synchronize the SSL chain certificate configuration, configure Steelhead Mobile for SSL, improve performance for SSL connection establishment, and enable client certificate authentication in the Configure > Optimization > Advanced Settings page.

To set Advanced SSL options

1. Choose Configure > Optimization > Advanced Settings to display the Advanced Settings page.

Figure 6-10. Advanced Settings Page

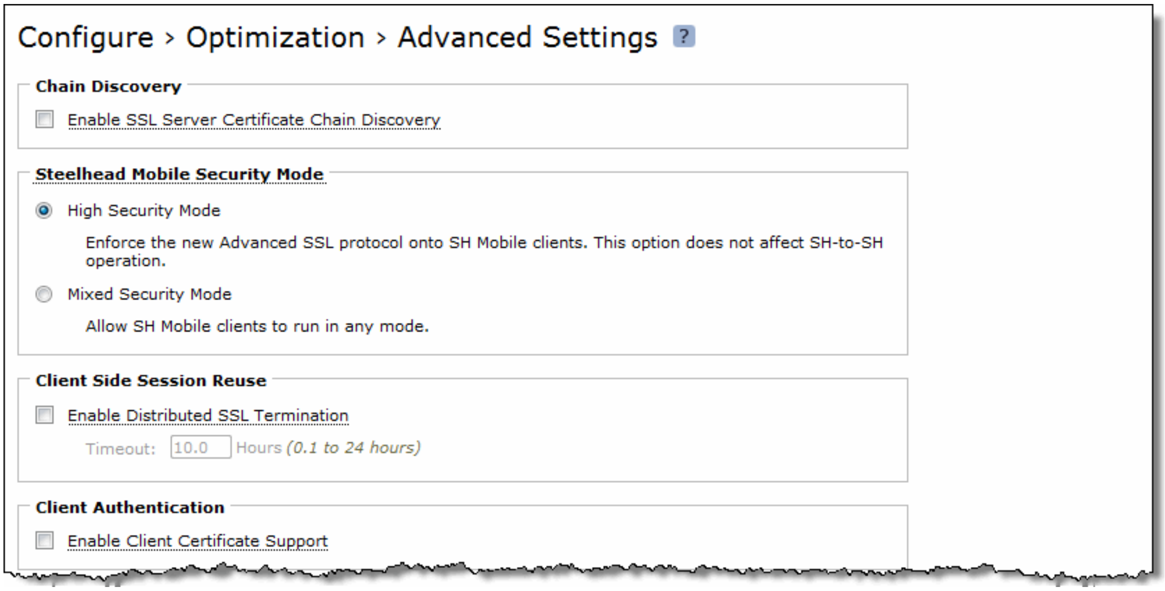

Configure > Optimization > Advanced Settings ?

Chain Discovery

□ Enable SSL Server Certificate Chain Discovery

Steelhead Mobile Security Mode

● High Security Mode

Enforce the new Advanced SSL protocol onto SH Mobile clients. This option does not affect SH-to-SH operation.

○ Mixed Security Mode

Allow SH Mobile clients to run in any mode.

Client Side Session Reuse

□ Enable Distributed SSL Termination

Timeout: 10.0 Hours (0.1 to 24 hours)

Client Authentication

□ Enable Client Certificate Support

2. Complete the configuration as described in the following table.

Control	Description
Enable SSL Server Certificate Chain Discovery	Synchronizes the chain certificate configuration on the server-side Steelhead appliance with the chain certificate configuration on the back-end server. The synchronization occurs after a handshake fails between the client-side and server-side Steelhead appliance. By default, this option is disabled. Enable this option when you replace an existing chain certificate on the back-end server with a new chain to ensure that the certificate chain remains in sync on both the server-side Steelhead appliance and the back-end server. **Note:** This option never replaces the server certificate. It updates the chain containing the intermediate certificates and the root certificate in the client context.
Steelhead Mobile Security Mode	On the server-side Steelhead appliance, select one of the following security modes: • **High Security Mode** - Enforces the advanced SSL protocol on the Steelhead Mobile Clients for increased security (v5.5.x or later). • **Mixed Security Mode** - Allows Steelhead Mobile clients to run in any SSL mode. **Note:** This option does not affect Steelhead appliance-to-Steelhead appliance operation.

Control	Description
Enable Distributed SSL Termination	Enables reuse of the original session on a client-side Steelhead appliance when the client reconnects to an SSL server. Reusing the session provides two benefits: it lessens the CPU load because it eliminates expensive asymmetric key operations and it shortens the key negotiation process by avoiding WAN roundtrips to the server. By default, this option is disabled. Both the client-side and server-side Steelheads must be configured to optimize SSL traffic. • **Timeout** - Specify the amount of time the client can reuse a session with an SSL server after the initial connection ends. The range is 6 minutes to 24 hours. The default value is 10 hours. Enabling this option requires an optimization service restart.

Control	Description
Enable Client Certificate Support	Enables acceleration of SSL traffic to those SSL servers that authenticate SSL clients. The SSL server verifies the SSL client certificate. In the client authentication SSL handshake, each client has a unique client certificate and the SSL server, in most cases, maintains the state that is specific to each client when answering the client's requests. The SSL server must receive exactly the same certificate that is originally issued for a client on all the connections between the client and the server. Typically the client's unique certificate and private key are stored on a smart card, such as a Common Access Card (CAC), or on a similar location that is inaccessible to other devices on the network. Enabling the client authentication feature allows Steelhead appliances to compute the encryption key while the SSL server continues to authenticate the original SSL client exactly as it would without the Steelhead appliances. The server-side Steelhead appliance observes the SSL handshake messages as they go back and forth. With access to the SSL server's private key, the Steelhead appliance computes the session key exactly as the SSL server does. The SSL server continues to perform the actual verification of the client, so any dependencies on the uniqueness of the client certificate for correct operation of the application are met. Because the Steelhead appliance does not modify any of the certificates (or the handshake messages) exchanged between the client and the server, there is no change to their trust model. The client and server continue to trust the same set of Certificate Authorities as they did without the Steelhead appliances accelerating their traffic. **Note:** If the data center has a mixed environment with a few SSL servers that authenticate clients along with those that do not authenticate clients, Riverbed recommends enabling client authentication. **Requirements** • Both the client-side and the server-side Steelhead appliance must be running RiOS v6.5. • Enable client certificate support on the server-side Steelhead appliance. • The server-side Steelhead appliance must have access to the exact private key used by the SSL server. • The SSL server must be configured to ask for client certificates. • The Steelhead appliance must have a compatible cipher chosen by the server. • SSL sessions that reuse previous secrets that are unknown to the Steelhead appliance cannot be decrypted. • Client-side certificates with renegotiation handshakes are not supported. • Client certificate supports the RSA key exchange only. It does not support the Diffie-Hellman key exchange. **Basic Steps** The basic steps to enable client authentication are: **1.** Perform the basic steps to enable SSL optimization (described in Configuring SSL Server Certificates and Certificate Authorities). **2.** On the server-side Steelhead appliance, choose Configure > Optimization > Advanced Settings, select Enable Client Certificate Support, and click **Apply**. **3.** Choose Configure > Optimization> SSL Main Settings, import the private key and certificate used by the SSL server to the server-side Steelhead appliance, and click **Save** to save the configuration. You do not need to restart the optimization service. **Verification** To verify client authentication, on the server-side Steelhead appliance, check the Discovered Server (Optimizable) table in the Configure > Optimization > SSL Main Settings page. Optimizable servers that are using client authentication appear as optimizable. For servers that are not using client authentication, the server appears in the Discovered Server (bypassed, not optimizable) table with the reason “No proxy certificate configured for the server."

4. Click **Apply** to apply your settings.

5. Click **Save** to save your settings permanently.

6. If you have enabled Client Side Session Reuse you need to restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Note: For details on Steelhead Mobile security mode and client-side session reuse, see the *Riverbed Deployment Guide*.

Configuring SSL Cipher Settings

You configure SSL cipher settings in the Configure > Optimization > Advanced Settings page.

Note: Unless you have specific organizational requirements, typically you do not need to change SSL cipher settings.

In cryptography, a cipher is an algorithm for performing encryption and decryption. In RiOS, the types of ciphers are:

- **Server ciphers** - communicate with the server on the segment between the server-side Steelhead appliance and the SSL server.
- **Client ciphers** - communicate with the client on the segment between the client-side Steelhead appliance and the SSL client. Although this segment does not include the server-side Steelhead appliance, you must configure the client ciphers on the server-side Steelhead appliance, because the server-side Steelhead appliance actually handles the SSL handshake with the SSL client.
- **Peer ciphers** - communicate between the two Steelhead appliances.

The default cipher setting is DEFAULT which represents a variety of high-strength ciphers that allow for compatibility with many browsers and servers.

Use the default cipher configuration to limit the possible ciphers which are negotiated on the three parts of the secure inner channel connection (the client-to-Steelhead appliance, the server-to-Steelhead appliance, and Steelhead appliance-to-Steelhead appliance).

To configure SSL ciphers

1. Choose Configure > Optimization > Advanced Settings to display the Advanced Settings page.

Figure 6-11. **Partial Advanced Settings Page**

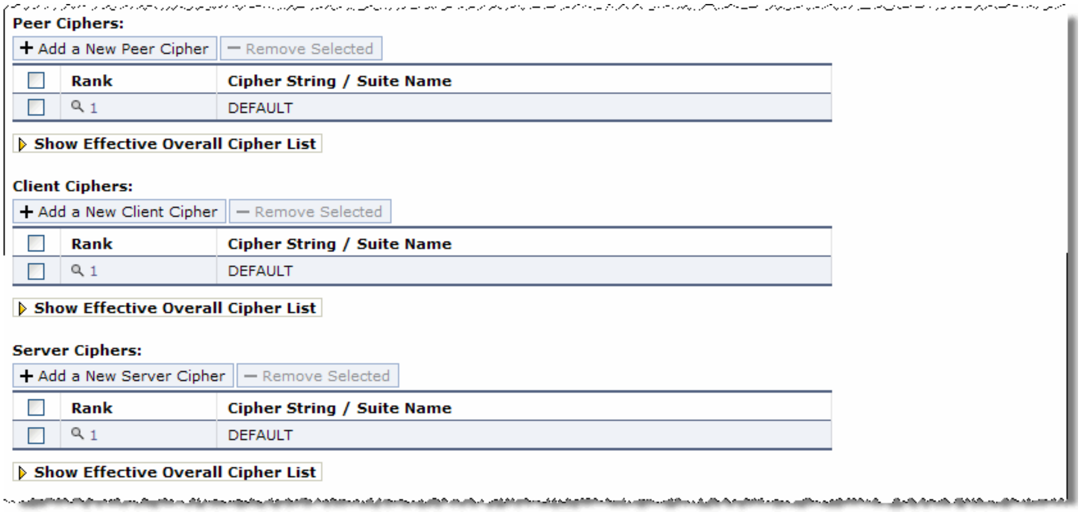

Peer Ciphers:

+ Add a New Peer Cipher | — Remove Selected

□	Rank	Cipher String / Suite Name
□	1	DEFAULT

▷ Show Effective Overall Cipher List

Client Ciphers:

+ Add a New Client Cipher | — Remove Selected

□	Rank	Cipher String / Suite Name
□	1	DEFAULT

▷ Show Effective Overall Cipher List

Server Ciphers:

+ Add a New Server Cipher | — Remove Selected

□	Rank	Cipher String / Suite Name
□	1	DEFAULT

▷ Show Effective Overall Cipher List

2. Under Peer Ciphers, complete the configuration on both the server-side and client-side Steelhead appliances, as described in the following table.

Control	Description
Add a New Peer Cipher	Displays the controls for adding a new peer cipher.
Cipher	Select the cipher type for communicating with peers from the drop-down list. The Hint text box displays information about the cipher. You must specify at least one cipher for peers, clients, and servers for SSL to function properly. The default cipher setting is DEFAULT which represents a variety of high-strength ciphers that allow for compatibility with many browsers and servers.
Insert Cipher At	Select Start, End, or the cipher number from the drop-down list. The default cipher, if used, must be rule number 1.
Add	Adds the cipher to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. On the server-side Steelhead appliance, under Client Ciphers, you can add or remove a client or peer cipher by completing the configuration as described in the following tables.

Control	Description
Add a New Client Cipher	Displays the controls for adding a new client cipher.
Cipher	Select the cipher type for communicating with clients from the drop-down list. The Hint text box displays information about the cipher. You must specify at least one cipher for peers, clients, and servers for SSL to function properly. The default cipher setting is DEFAULT which represents a variety of high-strength ciphers that allow for compatibility with many browsers and servers.
Insert Cipher At	Select Start, End, or a cipher number from the drop-down list. The default cipher, if used, must be rule number 1.
Add	Adds the cipher to the list.
Cancel	Cancels your settings.
Removed Selected	Select the check box next to the name and click **Remove Selected**.

Control	Description
Add a New Peer Cipher	Displays the controls for adding a new peer cipher.
Cipher	Select the cipher type for communicating with peers from the drop-down list. The Hint text box displays information about the cipher. You must specify at least one cipher for peers, clients, and servers for SSL to function properly. The default cipher setting is DEFAULT which represents a variety of high-strength ciphers that allow for compatibility with many browsers and servers.
Insert Cipher At	Select Start, End, or the cipher number from the drop-down list. The default cipher, if used, must be rule number 1.
Add	Adds the cipher to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

On the server-side Steelhead appliance, you can add or remove a server cipher by completing the configuration as described in the following table.

Control	Description
Add a New Server Cipher	Displays the controls for adding a new server cipher.
Cipher	Select the cipher type for communicating with servers from the drop-down list. The Hint text box displays information about the cipher. You must specify at least one cipher for peers, clients, and servers for SSL to function properly. The default cipher setting is DEFAULT which represents a variety of high-strength ciphers that are compatible with many browsers and servers.
Insert Cipher At	Select Start, End, or a cipher number from the drop-down list. The default cipher, if used, must be rule number 1.
Add	Adds the cipher to the list.

Control	Description
Cancel	Cancels your settings.
Removed Selected	Select the check box next to the name and click **Remove Selected**.

4. Click **Show Effective Overall Cipher List** to display a list of ciphers.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring HTTP Optimization" on page 111
- "Viewing SSL Reports" on page 434
- "Viewing the System Dumps List" on page 480
- "Unlocking the Secure Vault" on page 370

Performing Bulk Imports and Exports

You can perform bulk import and export operations in the Configure > Optimization > Advanced Settings page.

The following import and export features expedite configuring backup and peer trust relationships:

- **Backup** - You can use the bulk export feature to back up your SSL configurations, including your server configurations and private keys.

Important: To protect your server private keys, you can choose to not include your **Server Configurations and Private Keys** when performing bulk exports of trusted peers.

- **Peer Trust** - If you use self-signed peering certificates and have multiple Steelhead appliances (including multiple server-side appliances), you can use the bulk import feature to avoid configuring each peering trust relationship between the pairs of Steelhead appliances.

The bulk data that you import contains the serial number of the exporting Steelhead appliance. The Steelhead appliance importing the data compares its own serial number with the serial number contained in the bulk data.

The following rules apply to bulk data when importing and exporting the data:

- **Peering Certificate and Key Data** - If the serial numbers match, the Steelhead appliance importing the bulk data overwrites its existing peering certificates and keys with that bulk data. If the serial numbers do not match, the Steelhead appliance importing the bulk data does not overwrite its peering certificate and key.

- **Certificate Authority, Peering Trust, and SSL Server Configuration Data** - For all other configuration data such as certificate authorities, peering trusts, and server configurations (if included), if there is a conflict, the imported configuration data takes precedence (that is, the imported configuration data overwrites any existing configurations).

Note: Bulk data importing operations do not delete configurations; they can only add or overwrite them.

Note: Bulk importing does not require a Steelhead service restart.

To perform bulk export operations

1. Select one Steelhead appliance (A) and trust all the Steelhead appliances peering certificates. Make sure you include the peering certificate for Steelhead appliance A. For details on configuring trusted peers, see "Configuring Secure Peers" on page 231.

2. Choose Configure > Optimization > Advanced Settings to display the Advanced Settings page.

3. Under Bulk Export, complete the configuration as described in the following table.

Control	Description
Include Server Certificates and Private Keys	Includes the server certificates and keys in the export file. **Important:** To protect your server private keys, do not select when performing bulk exports of trusted peers.
Include SCEP/CRL Configuration	Includes the SCEP and CRL configurations with the export file.
Password	Specify and confirm the password used for the export file.
Export	Exports your SSL configuration and optionally your server private keys and certificates.

4. Click **Save** to save your settings permanently.

To perform bulk import operations

1. Choose Configure > Optimization > Advanced Settings to display the Advanced Settings page.

Figure 6-12. **Advanced Settings Page**

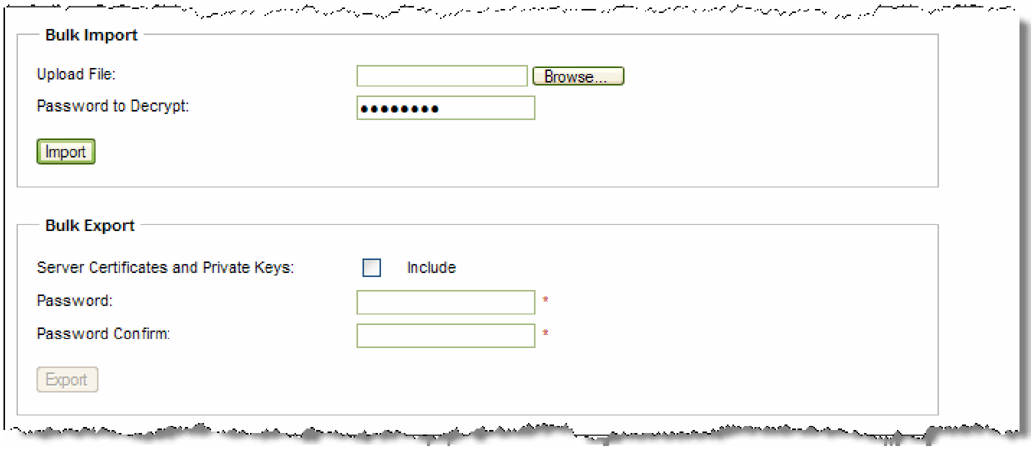

Bulk Import

Upload File: [ ] Browse...

Password to Decrypt: ••••••••

Import

Bulk Export

Server Certificates and Private Keys: □ Include

Password: [ ] *

Password Confirm: [ ] *

Export

2. Under Bulk Import, complete the configuration as described in the following table.

Control	Description
Upload File	Browse to the previously exported bulk file that contains the certificates and keys.
Password to Decrypt	Specify the password used to decrypt the file.
Import	Imports your SSL configuration, keys, and certificates, so that all of the Steelhead appliances trust one another as peers.

3. Click **Save** to save your settings permanently.

Related Topics

- "Configuring In-Path Rules" on page 28
- "Enabling Peering and Configuring Peering Rules" on page 67
- "Configuring HTTP Optimization" on page 111
- "Unlocking the Secure Vault" on page 370
- "Viewing SSL Reports" on page 434
- "Viewing the System Dumps List" on page 480

CHAPTER 7

Configuring Network Integration Features

This chapter describes how to configure advanced features such as asymmetric routing, connection forwarding, encryption, flow export, QoS, joining a Windows domain, simplified routing, and WCCP.

This chapter includes the following topics:

For details on basic and advanced deployment types, see the *Riverbed Deployment Guide*.

Configuring Asymmetric Routing Features

You enable asymmetric route detection in the Configure > Networking > Asymmetric Routing page.

Asymmetric route detection automatically detects and reports asymmetric routing conditions and caches this information to avoid losing connectivity between a client and a server.

Asymmetric routing is when a packet takes one path to the destination and takes another path when returning to the source. Asymmetric routing is common within most networks; the larger the network, the more likely there is asymmetric routing in the network.

Asymmetric routing is undesirable for many network devices including, firewalls, VPNs, and Steelhead appliances. These devices all rely on seeing every packet to function properly. When Steelhead appliances are deployed in a network, all TCP traffic must flow through the same Steelhead appliances in the forward and reverse directions. If traffic flows through a Steelhead appliance in one direction and not the other, then TCP clients are unable to make connections to TCP servers. When deploying Steelhead appliances into redundant networks, there is a possibility of traffic taking different forward and return paths so that traffic in one direction goes through Steelhead appliances but traffic in the reverse direction does not.

Asymmetric auto-detection enables Steelhead appliances to detect the presence of asymmetry within the network. Asymmetry is detected by the client-side Steelhead appliances. Once detected, the Steelhead appliance passes through asymmetric traffic unoptimized allowing the TCP connections to continue to work. The first TCP connection for a pair of addresses might be dropped because during the detection process the Steelhead appliances have no way of knowing that the connection is asymmetric.

If asymmetric routing is detected, an entry is placed in the asymmetric routing table and any subsequent connections from that IP-address pair is passed through unoptimized. Further connections between these hosts are not optimized until that particular asymmetric routing cache entry times out.

The Configure > Networking > Asymmetric Routing page displays the asymmetric routing table. The following table describes the different types of asymmetry.

Type	Description	Asymmetric Routing Table and Log Entries
Complete Asymmetry	Packets traverse both Steelhead appliances going from the client to the server but bypass both Steelhead appliances on the return path.	• Asymmetric Routing Table: bad RST • Log: Sep 5 11:16:38 gen-sh102 kernel: [intercept.WARN] asymmetric routing between 10.11.111.19 and 10.11.25.23 detected (bad RST)
Server-Side Asymmetry	Packets traverse both Steelhead appliances going from the client to the server but bypass the server-side Steelhead appliance on the return path.	• Asymmetric Routing Table: bad SYN/ACK • Log: Sep 7 16:17:25 gen-sh102 kernel: [intercept.WARN] asymmetric routing between 10.11.25.23:5001 and 10.11.111.19:33261 detected (bad SYN/ACK)
Client-Side Asymmetry	Packets traverse both Steelhead appliances going from the client to the server but bypass the client-side Steelhead appliance on the return path.	• Asymmetric Routing Table: no SYN/ACK • Log: Sep 7 16:41:45 gen-sh102 kernel: [intercept.WARN] asymmetric routing between 10.11.111.19:33262 and 10.11.25.23:5001 detected (no SYN/ACK)
Multi-SYN Retransmit	The types of Multi-SYN Retransmits are: • Probe-filtered occurs when the client-side Steelhead appliance sends out multiple SYN+ frames and does not get a response. • SYN-remit occurs when the client-side Steelhead appliance receives multiple SYN retransmits from a client and does not see a SYN/ACK packet from the destination server.	• Asymmetric Routing Table: probe-filtered(not-AR) • Log: Sep 13 20:59:16 gen-sh102 kernel: [intercept.WARN] it appears as though probes from 10.11.111.19 to 10.11.25.23 are being filtered. Passing through connections between these two hosts.

Detecting and caching asymmetric routes does not optimize these packets. If you want to optimize asymmetric routed packets you must make sure that packets going to the WAN always go through a Steelhead appliance either by using a multi-port Steelhead appliance, connection forwarding, or using external ways to redirect packets, such as WCCP or PBR.

For details, see "Configuring Connection Forwarding Features" on page 255 or the *Riverbed Deployment Guide*.

Troubleshooting Asymmetric Routes

You can use the following tools to detect and analyze asymmetric routes:

- **TCP Dump** - Run a TCP dump diagnostic report on the client-side Steelhead appliance to verify the packet sequence that is causing the asymmetric route detection. You can take traces on the LAN and WAN ports of the Steelhead appliance and, based on the packet maps, look for the packet sequence that is expected for the type of warning message that was in the log.

  For example, to obtain information about all packets on the WAN interface sourced from or destined to 10.0.0.1, and with a source and destination TCP port of 80:

  1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
  2. Click **Add a New TCP Dump**.
  3. Select the WAN interface.
  4. Specify 10.0.0.1 as the source and destination address.
  5. Specify TCP port 80 as the source and destination port.
  6. Select the Schedule Dump check box and specify the date and time to initiate the dump.
  7. Specify any other options such as the capture filename or duration.
  8. Click **Add**.

For details, see “Capturing and Uploading TCP Dumps” on page 482.

- **Trace Route**. From the CLI, run the **traceroute** tool to discover what path a packet is taking from the client to the server and from the server to the client. You access the client and run the **traceroute** command with the IP address of the server, then run the **traceroute** command from the server with the IP address of the client. For example, for a Cisco router:

```
#Client’s Address: 10.1.0.2
#Server’s Address: 10.0.0.4
client# traceroute 10.0.0.4 Type escape sequence to abort.
Tracing the route to 10.0.0.4
1 10.1.0.1 4 msec 0 msec 4 msec
2 10.0.0.2 4 msec 4 msec 0 msec
3 10.0.0.3 4 msec 4 msec 0 msec
4 10.0.0.4 4 msec 4 msec 0 msec
server# traceroute 10.1.0.2 Type escape sequence to abort.
Tracing the route to 10.1.0.2
1 10.0.0.6 4 msec 0 msec 4 msec
2 10.0.0.5 4 msec 4 msec 0 msec
3 10.1.0.1 4 msec 4 msec 0 msec
4 10.1.0.2 4 msec 4 msec 0 msec
```

For details, see the *Riverbed Command-Line Interface Reference Manual* or the *Riverbed Deployment Guide*.

To automatically detect asymmetric routing

1. Choose Configure > Networking > Asymmetric Routing to display the Asymmetric Routing page.

Figure 7-1. Asymmetric Routing Page

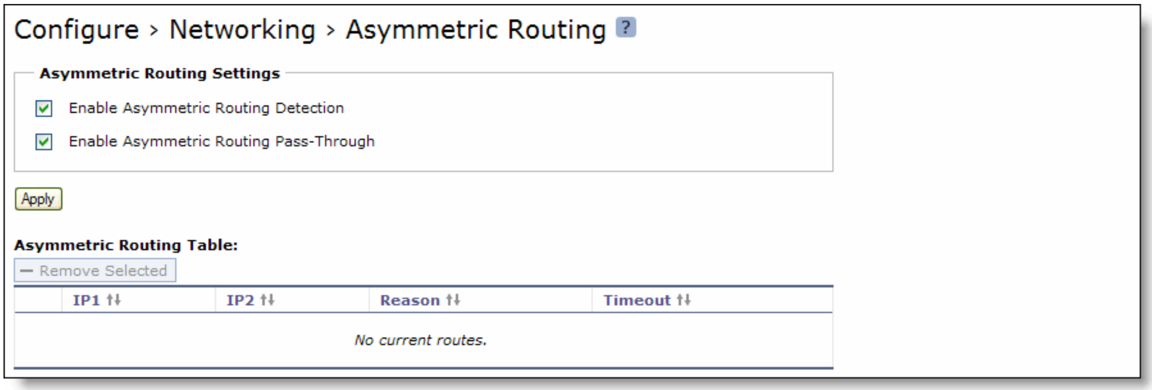

Asymmetric Routing Settings

☑ Enable Asymmetric Routing Detection

☑ Enable Asymmetric Routing Pass-Through

Asymmetric Routing Table:

	IP1	IP2	Reason	Timeout
		No current routes.		

2. Under Asymmetric Routing Settings, complete the configuration as described in the following table.

Control	Description
Enable Asymmetric Routing Detection	Detects asymmetric routes in your network.
Enable Asymmetric Routing Pass-Through	Enables pass-through traffic if asymmetric routing is detected. If asymmetric routing is detected, the pair of IP addresses, defined by the client and server addresses of this connection, is cached on the Steelhead appliance. Further connections between these hosts are passed through unoptimized until that particular asymmetric routing cache entry times out. Detecting and caching asymmetric routes does not optimize these packets. If you want to optimize asymmetric routed packets you must make sure that the packets going to the WAN always go through a Steelhead appliance either by using a multi-port Steelhead appliance, connection forwarding, or using external ways to redirect packets, such as WCCP or PBR. For details, see the *Riverbed Deployment Guide*.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Save** to save your settings permanently.

Related Topics

- "Configuring Connection Forwarding Features" on page 255
- "Viewing the System Dumps List" on page 480
- "Viewing Process Dumps" on page 481

Configuring Connection Forwarding Features

You configure connection forwarding for a network with multiple paths from the server in the Configure > Networking > Connection Forwarding page.

CSH The AWS Cloud Steelhead does not support connection forwarding; however, the ESX Cloud Steelhead supports it.

You enable connection forwarding only in asymmetric networks; that is, networks in which a client request traverses a different network path than the server response. The default port for connection forwarding is 7850.

For virtual in-path deployments with multiple Steelhead appliances, including WCCP clusters and connection forwarding, you must allow in-path neighbor failure. This is because certain events, such as network failures, and router or Steelhead appliance cluster changes, can cause routers to change the destination Steelhead appliance for TCP connection packets. When this happens, Steelhead appliances must be able to redirect traffic to each other to ensure that optimization continues.

To optimize connections in asymmetric networks, packets traveling in both directions must pass through the same client-side and server-side Steelhead appliance. If you have one path from the client to the server and a different path from the server to the client, you need to enable in-path connection forwarding and configure the Steelhead appliances to communicate with each other. These Steelhead appliances are called neighbors and exchange connection information to redirect packets to each other.

In RiOS v6.5 and later, you must enable connection forwarding in a WCCP cluster. With connection forwarding enabled, the WCCP load balancing algorithm considers the total number of in-path interfaces of all neighbors in the service group when balancing the load across the interfaces. If you do not enable connection forwarding, the Steelhead with the lowest IP address assigns all traffic flows to itself. For details, see the *Riverbed Deployment Guide*.

Figure 7-2. Asymmetric Network

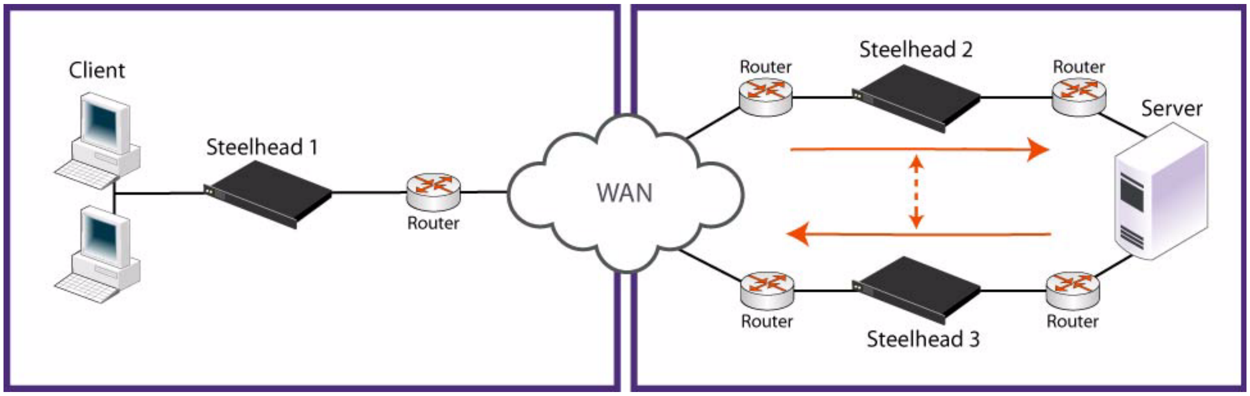

Neighbors can be placed in the same physical site or in different sites, but the latency between them should be small because the packets travelling between them are not optimized.

Important: When you define a neighbor, you specify the Steelhead appliance in-path IP address, not the primary IP address.

If there are more than two possible paths, additional Steelhead appliances must be installed on each path and configured as neighbors. Neighbors are notified in parallel so that the delay introduced at the connection setup is equal to the time it takes to get an acknowledgement from the furthest neighbor.

Important: Connection forwarding neighbors must use the same WAN visibility mode. For details, see "Configuring In-Path Rules" on page 28.

For details on connection forwarding, see the *Riverbed Deployment Guide*.

To enable connection forwarding

1. Choose Configure > Networking > Connection Forwarding to display the Connection Forwarding page.

Figure 7-3. Connection Forwarding Page

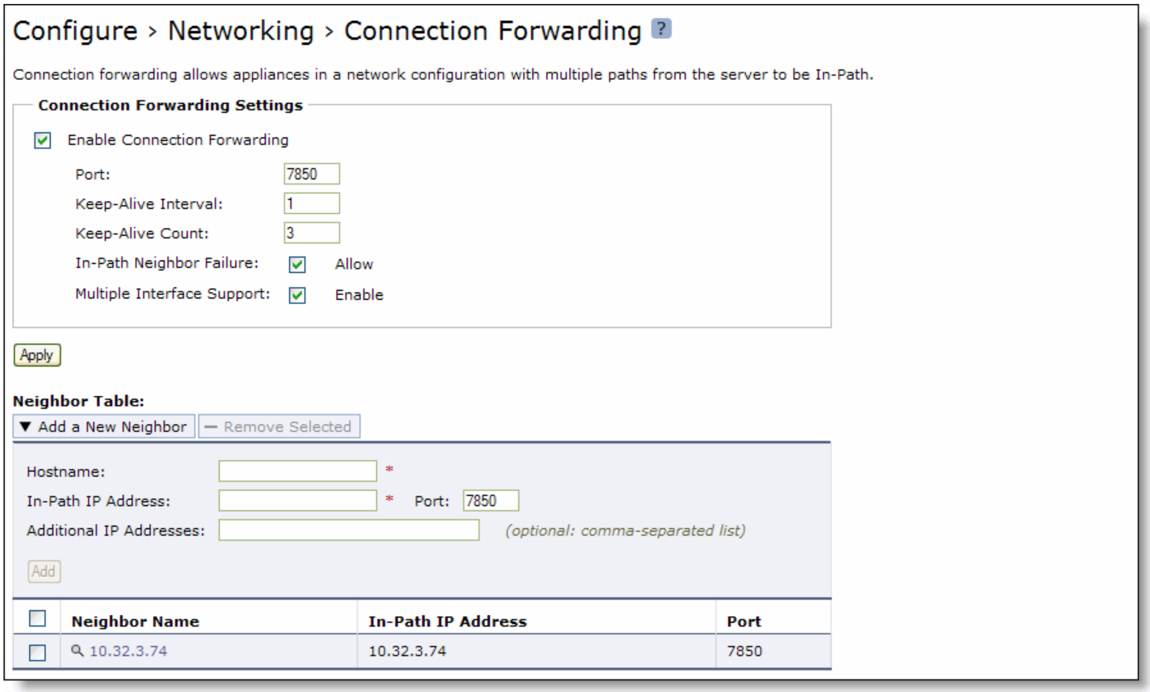

Configure > Networking > Connection Forwarding ?

Connection forwarding allows appliances in a network configuration with multiple paths from the server to be In-Path.

Connection Forwarding Settings

☑ Enable Connection Forwarding

Port:	7850	
Keep-Alive Interval:	1	
Keep-Alive Count:	3	
In-Path Neighbor Failure:	☑	Allow
Multiple Interface Support:	☑	Enable

Apply

Neighbor Table:

▼ Add a New Neighbor — Remove Selected

Hostname: *
In-Path IP Address: * Port: 7850
Additional IP Addresses: *(optional: comma-separated list)*

Add

□	Neighbor Name	In-Path IP Address	Port
□	10.32.3.74	10.32.3.74	7850

2. Under Connection Forwarding Settings, complete the configuration as described in the following table.

Control	Description
Enable Connection Forwarding	Enables connection forwarding by default on all neighbors added to the peer list. The default value is 7850.
Port	Specify the port number to use as the default for the neighbor Steelhead appliance in-path port. The default value is 7850.
Keep-Alive Interval	Specify the number of seconds to use as the default interval for ping commands between neighbor Steelhead appliances.
Keep-Alive Count	Specify the number of tries to use as the default number of failed ping attempts before an appliance terminates a connection with a neighbor. The default value is 3.

Control	Description
In-Path Neighbor Failure	Uses the neighbor appliance to optimize new connections if the appliance fails. For in-path deployments that use connection forwarding with WCCP, enabling this option ensures that if one appliance fails, the neighbor appliance continues to optimize new connections. For in-path deployments that use connection forwarding without WCCP, enabling this option ensures that a Steelhead appliance attempts to optimize new connections that are symmetrically routed, even after all of the neighbor Steelhead appliances on another network path failed. New asymmetrically routed connections are not optimized but passed through.
Multiple Interface Support	Enables high availability on Steelhead appliances configured with multiple in-path interfaces and using connection forwarding with another multi-port Steelhead appliance. This option makes all neighbor in-path interface IP addresses visible to each peer to ensure proper neighbor communication if the in-path0_0 interface fails. RiOS v6.5 and later requires connection forwarding in a WCCP cluster.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

To add a new neighbor

1. Under Neighbor Table, complete the configuration as described in the following table.

Control	Description
Add a New Neighbor	Displays the controls to add a new neighbor.
Hostname	Specify a hostname.
In-Path IP Address	Specify the in-path IP address for the neighbor Steelhead appliance. When you define a neighbor, you must specify the appliance in-path IP address, not the primary IP address.
Port	Specify the in-path port for the neighbor Steelhead appliance. The default port is 7850.
Additional IP Addresses	Adds a neighbor Steelhead appliance to the neighbor list.
Add	Adds a new neighbor.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Apply** to apply your settings.

3. Click **Save** to save your settings permanently.

Tip: To modify the neighbor properties, select the IP address of the neighbor and complete the configuration.

Related Topics

- "Configuring General Service Settings" on page 60
- "Configuring Asymmetric Routing Features" on page 251

Configuring IPSec Encryption

You configure IPSec encryption to allow data to be communicated securely between peer Steelhead appliances in the Configure > Optimization > Secure Peering (IPSEC) page.

Enabling IPSec encryption makes it difficult for a third party to view your data or pose as a computer you expect to receive data from. To enable IPSec, you must specify at least one encryption and authentication algorithm. Only optimized data is protected, pass-through traffic is not.

Enabling IPSec support is *optional*.

RiOS v6.0 and later also provides support for SSL peering beyond traditional HTTPS traffic. For details, see "Configuring Secure Peers" on page 231.

Important: You must set IPSec support on each peer Steelhead appliance in your network for which you want to establish a secure connection. You must also specify a shared secret on each peer Steelhead appliance.

Note: If you NAT traffic between Steelhead appliances, you cannot use the IPSec channel between the Steelhead appliances because the NAT changes the packet headers, causing IPSec to reject them.

To enable IPSec encryption

1. Choose Configure > Optimization > Secure Peering (IPSEC) to display the Secure Peering (IPSEC) page.

Figure 7-4. Secure Peering (IPSEC) Page

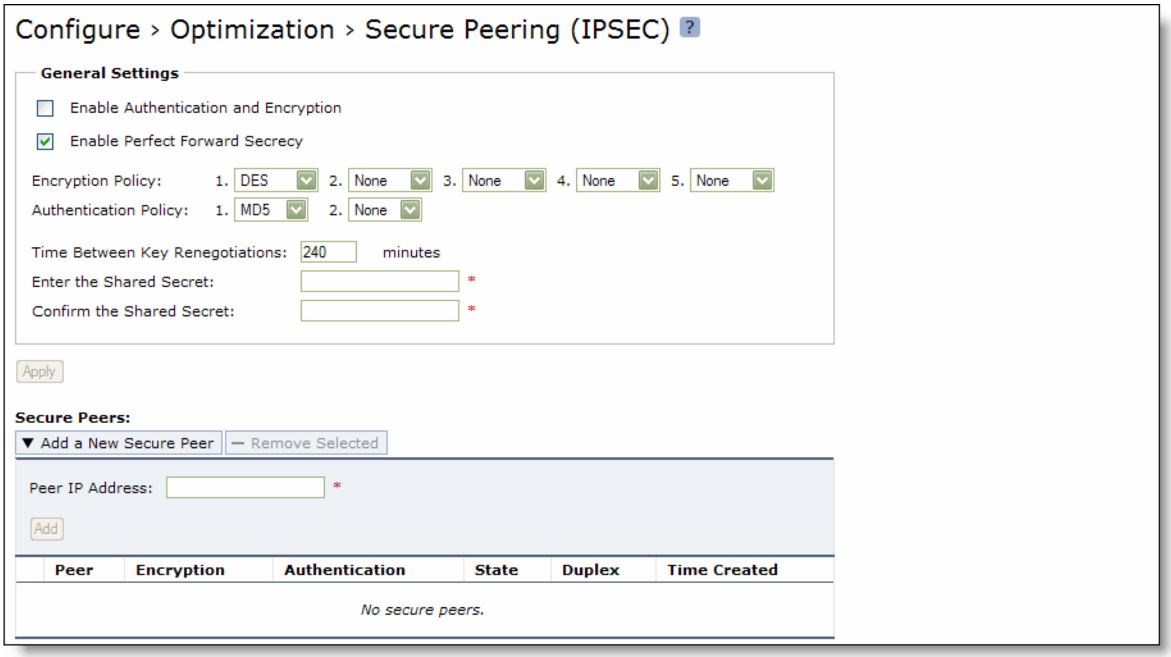

Configure > Optimization > Secure Peering (IPSEC) ?

General Settings

□ Enable Authentication and Encryption

☑ Enable Perfect Forward Secrecy

Encryption Policy: 1. DES 2. None 3. None 4. None 5. None

Authentication Policy: 1. MD5 2. None

Time Between Key Renegotiations: 240 minutes

Enter the Shared Secret: *

Confirm the Shared Secret: *

Apply

Secure Peers:

▼ Add a New Secure Peer — Remove Selected

Peer IP Address: *

Add

Peer	Encryption	Authentication	State	Duplex	Time Created
No secure peers.					

2. Under General Settings, complete the configuration as described in the following table.

Control	Description
Enable Authentication and Encryption	Enables authentication between Steelhead appliances. By default, this option is disabled.
Enable Perfect Forward Secrecy	Enables additional security by renegotiating keys at specified intervals. If one key is compromised, subsequent keys are secure because they are not derived from previous keys. By default, this option is enabled.
Encryption Policy	Select one of the following encryption methods from the drop-down list: • **DES** - Encrypts data using the Data Encryption Standard algorithm. DES is the default value. • **NULL** - Specifies the null encryption algorithm. • **None** - Does not apply an encryption policy. • **3DES** - Appears when a valid Enhanced Cryptography License Key is installed on the appliance. Encrypts data using the Triple Digital Encryption Standard with a 168-bit key length. This standard is supported for environments where AES has not been approved, but is both slower and less secure than AES. • **AES** - Appears when a valid Enhanced Cryptography License Key is installed on the appliance. Encrypts data using the Advanced Encryption Standard (AES) cryptographic key length of 128 bits. • **AES256** - Appears when a valid Enhanced Cryptography License Key is installed. Encrypts data using the Advanced Encryption Standard (AES) cryptographic key length of 256 bits. Provides the highest security. Optionally, select an algorithm from the method 2, 3, 4, or 5 drop-down lists to create a prioritized list of encryption policies for negotiating between peers. **Note:** Peer Steelhead appliances must both have a valid Enhanced Cryptography License Key installed to use 3DES, AES, or AES256. When a Steelhead appliance has the valid Enhanced Cryptography License Key installed and an IPSec encryption level is set to 3DES or AES, and a peer Steelhead appliance does not have a valid Enhanced Cryptography License Key installed, the appliances uses the highest encryption level set on the appliance without the key.
Authentication Policy	Select one of the following authentication methods from the drop-down list: • **MD5** - Specifies the Message-Digest 5 algorithm, a widely-used cryptographic hash function with a 128-bit hash value. This is the default value. • **SHA-1** - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA-1 is considered to be the successor to MD5. Optionally, select an algorithm from the method 2 drop-down list to create a secondary policy for negotiating the authentication method to use between peers. If the first authentication policy negotiation fails, the peer Steelhead appliances use the secondary policy to negotiate authentication.
Time Between Key Renegotiations	Specify the number of minutes between quick-mode renegotiation of keys using the Internet Key Exchange (IKE) protocol. IKE uses public key cryptography to provide the secure transmission of a secret key to a recipient so that the encrypted data can be decrypted at the other end. The default value is 240 minutes.
Enter the Shared Secret/Confirm the Shared Secret	Specify and confirm the shared secret. All the Steelhead appliances in a network for which you want to use IPSec must have the same shared secret.

Control	Description
Add a New Secure Peer	Displays the controls to add a new secure peer. • **Peer IP Address** - Specify the IP address for the peer Steelhead appliance (in-path interface) for which you want to make a secure connection.
Add	Adds the peer specified in the Peer IP Address text box. If a connection has not been established between the two Steelhead appliances that are configured to use IPSec security, the peers list does not display the peer Steelhead appliance status as mature. **Note:** Adding a peer causes a short service disruption (3-4 seconds) to the peer that is configured to use IPSec security.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Save** to save your settings permanently.

4. If you have changed an IPSec encryption setting, you need to restart the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Note: The peered Steelhead appliances do not establish the IPSec channel until they are optimizing traffic.

About the Secure Peers List

The Secure Peers list displays the peers with the encryption and authentication policies. The state of the peer is listed as one of the following:

- **Mature** - The IPSec connection is established and usable.
- **Larval** - The IPSec connection is being established.
- **Disconnected** - The IPSec connection is not yet established or is not usable.

Configuring Subnet Side Rules

You need to configure subnet side rules to support RSP and Flow Export on a virtual in-path deployment in the Configure > Networking > Subnet Side Rules page.

Subnet side rules let you configure subnets as LAN-side subnets or WAN-side subnets for a virtual in-path Steelhead appliance. The subnet side rules determine whether traffic originated from the LAN or the WAN-side of the Steelhead appliance based on the source subnet. You must configure subnets on each Steelhead appliance in a virtual in-path configuration, as the subnets for each will likely be unique.

With subnet side rules in place, RiOS can send incoming packets to the correct RSP VNIs for VRSP, and a virtual in-path Steelhead can use flow export collectors such as NetFlow to analyze non-optimized or passed through traffic correctly. Otherwise, the Steelhead appliance cannot discern whether the traffic is traveling from the LAN to the WAN or in the opposite direction. This can result in over-reporting traffic in a particular direction or for a particular interface.

Note: FakeIndex is necessary for correct optimized traffic reporting. For details, see the *Riverbed Deployment Guide*.

Note: Before you use Virtual RSP, you must disable simplified routing. For details, see "Configuring Simplified Routing Features" on page 311.

To add subnet side rules

1. Choose Configure > Networking > Subnet Side Rules to display the Subnet Side Rules page.

Figure 7-5. Subnet Side Rules Page

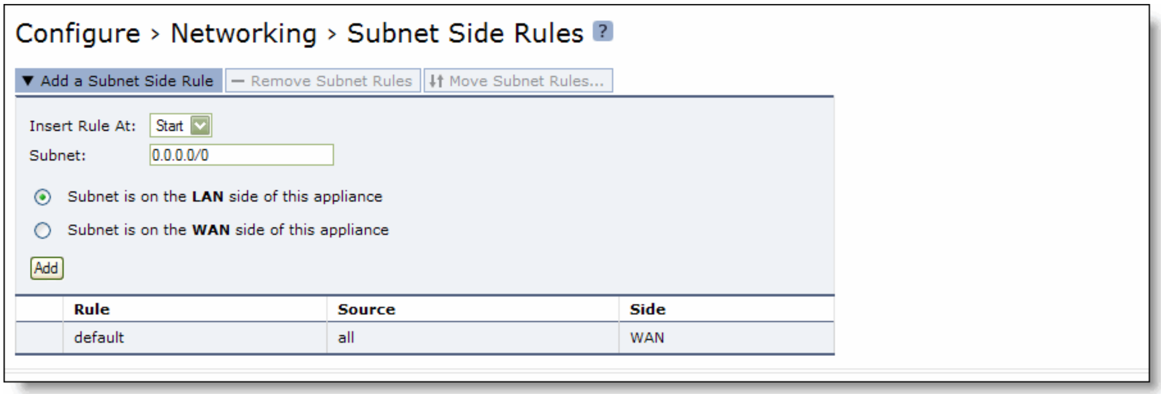

Insert Rule At: Start

Subnet: 0.0.0.0/0

● Subnet is on the **LAN** side of this appliance

○ Subnet is on the **WAN** side of this appliance

	Rule	Source	Side
	default	all	WAN

2. Complete the configuration as described in the following table.

Control	Description
Add a Subnet Side Rule	Displays the controls to create a subnet side rule.
Insert Rule At	Select Start, End, or a rule number from the drop-down list. Steelhead appliances evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted.
Subnet	Specify the subnet. Use the following format: <IP address>/<subnet mask>
Subnet is on the LAN side of this appliance	In virtual in-path configurations, all traffic is flowing in and out of one physical interface. Select to specify that the subnet is on the LAN side of the device.
Subnet is on the WAN side of this appliance	In virtual in-path configurations, all traffic is flowing in and out of one physical interface. Select to specify that the subnet is on the WAN side of the device.

Control	Description
Add	Adds the rule to the subnet map table. The Management Console redisplays the subnet map table and applies your changes to the running configuration, which is stored in memory.
Remove Subnet Rules	Select the check box next to the name and click **Remove Subnet Rules**.
Move Subnet Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

Tip: You cannot delete the default rule, Default, which optimizes all remaining WAN-side traffic that has not been selected by another rule. This rule is always listed last.

Related Topic

- "Configuring Flow Export" on page 262

Configuring Flow Export

You enable and configure flow export and Top Talker report settings in the Configure > Networking > Flow Export page. By default, flow export is disabled.

Flow export enables you to export network statistics to external collectors that provide information about network data flows such as the top users, peak usage times, traffic accounting, security, and traffic routing. You can export pre-optimization and post-optimization data to an external collector. The Top Talkers feature enables a report that details the hosts, applications, and host and application pairs that are either sending or receiving the most data on the network. Top Talkers does not use a NetFlow Collector.

Important: Steelhead appliances support NetFlow v5.0, CascadeFlow, NetFlow v9, and CascadeFlow-compatible.

Flow export requires the following components:

- **Exporter** - When you enable flow export support, the Steelhead appliance exports data about the individual flows that it sees as they traverse the network.
- **Collector** - A server or appliance designed to aggregate data sent to it by the Steelhead appliance and other exporters.
- **Analyzer** - A collection of tools used to analyze the data and provide relevant data summaries and graphs. NetFlow analyzers are available for free or from commercial sources. Analyzers are often provided in conjunction with the collectors.

Before you enable flow export in your network, consider the following:

- Flow data typically consumes less than 1% of link bandwidth. Take care with low bandwidth links to ensure that flow export does not consume too much bandwidth and thereby impacting application performance.
- You can reduce the amount of bandwidth consumption by applying filters that only export the most critical information needed for your reports.

Flow Export in Virtual In-Path Deployments

For virtual in-path deployments such as WCCP or PBR, because the traffic is arriving and leaving from the same WAN interface, when the Steelhead appliance exports data to a flow export collector, all traffic has the WAN interface index. This is the correct behavior because the input interface is the same as the output interface.

For details on configuring flow export in a virtual in-path deployment, see "Configuring Subnet Side Rules" on page 260.

To distinguish between LAN-to-WAN and WAN-to-LAN traffic in virtual in-path deployments, see the *Riverbed Deployment Guide*.

Troubleshooting

To troubleshoot your flow export settings:

- Make sure the port configuration matches on the Steelhead appliance and the listening port of the collector.
- Ensure that you can reach the collector from the Steelhead appliance (for example, ping 1.1.1.1 where 1.1.1.1 is the NetFlow collector).
- Verify that your capture settings are on the correct interface and that traffic is flowing through it.

To enable flow export and Top Talker report settings

1. Choose Configure > Networking > Flow Export to display the Flow Export page.

Figure 7-6. Flow Export Page

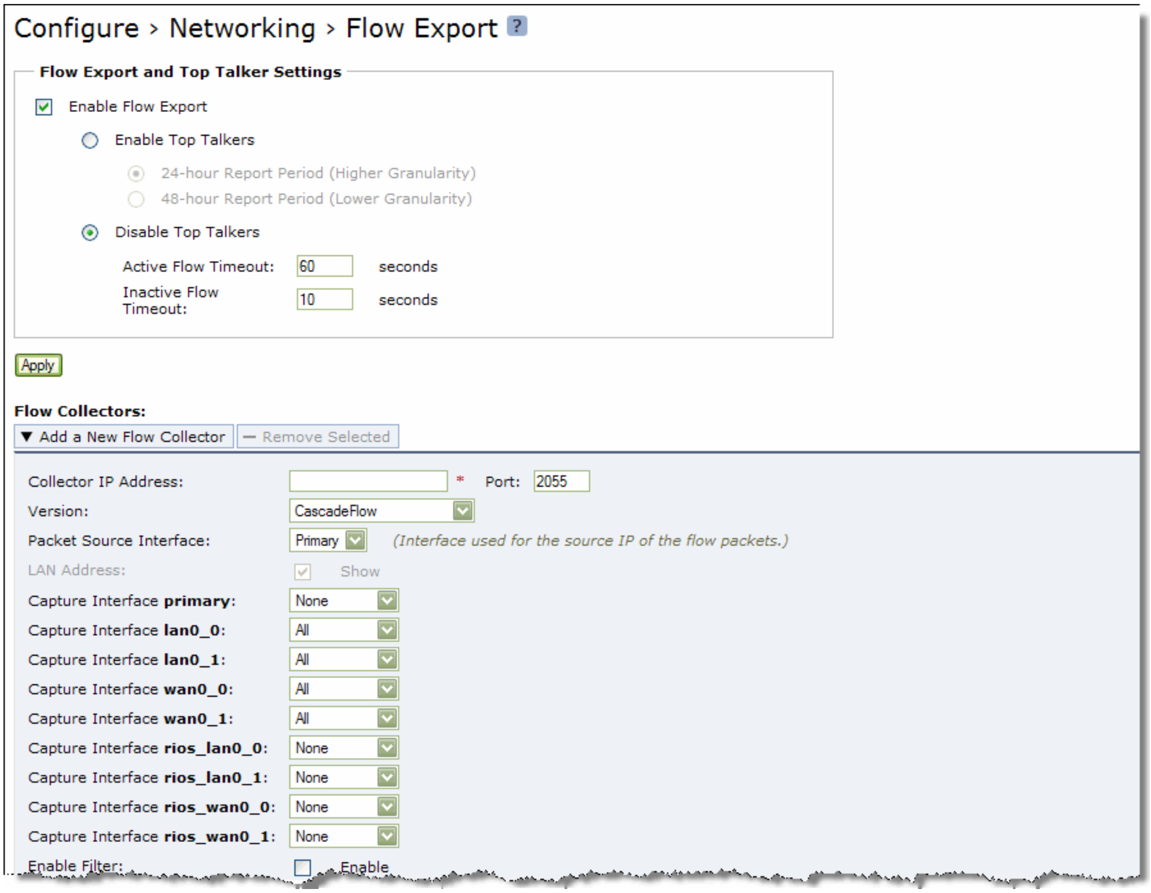

Configure > Networking > Flow Export ?

Flow Export and Top Talker Settings

☑ Enable Flow Export

○ Enable Top Talkers

● 24-hour Report Period (Higher Granularity)

○ 48-hour Report Period (Lower Granularity)

● Disable Top Talkers

Active Flow Timeout: 60 seconds

Inactive Flow Timeout: 10 seconds

Apply

Flow Collectors:

▼ Add a New Flow Collector — Remove Selected

Collector IP Address: * Port: 2055

Version: CascadeFlow

Packet Source Interface: Primary *(Interface used for the source IP of the flow packets.)*

LAN Address: ☑ Show

Capture Interface **primary**: None

Capture Interface **lan0_0**: All

Capture Interface **lan0_1**: All

Capture Interface **wan0_0**: All

Capture Interface **wan0_1**: All

Capture Interface **rios_lan0_0**: None

Capture Interface **rios_lan0_1**: None

Capture Interface **rios_wan0_0**: None

Capture Interface **rios_wan0_1**: None

Enable Filter: □ Enable

2. Under Flow Export and Top Talker Settings, complete the configuration as described in the following table.

Control	Description
Enable Flow Export	Enables the Steelhead appliance to export network statistics about the individual flows that it sees as they traverse the network. By default, this setting is disabled.
Enable Top Talkers	Continuously collects statistics for the most active traffic flows. A traffic flow consists of data sent and received from a single source IP address and port number to a single destination IP address and port number over the same protocol. The most active, heaviest users of WAN bandwidth are called the *Top Talkers*. A flow collector identifies the top consumers of the available WAN capacity (the top 50 by default) and displays them in the Top Talkers report. Collecting statistics on the Top Talkers provides visibility into WAN traffic without applying an in-path rule to enable a WAN visibility mode. You can analyze the Top Talkers for accounting, security, troubleshooting, and capacity planning purposes. You can also export the complete list in CSV format. The collector gathers statistics on the Top Talkers based on the proportion of WAN bandwidth consumed by the top hosts, applications, and host and application pair conversations. The statistics track pass-through or optimized traffic, or both. Data includes TCP or UDP traffic, or both (configurable in the Top Talkers report page). You must enable Flow Export before you enable Top Talkers. A NetFlow collector is not required for this feature. Enabling Top Talkers automatically sets the Active Flow Timeout to **60** seconds. Optionally, select a time period to adjust the collection interval: • **24-hour Report Period** - For a five-minute granularity (the default setting). • **48-hour Report Period** - For a ten-minute granularity.
Disable Top Talkers	Stops collecting statistics on the most active, heaviest users of WAN bandwidth.
Active Flow Timeout	Optionally, specify the amount of time, in seconds, the collector retains the list of active traffic flows. The default value is 1800 seconds. Enabling Top Talkers automatically sets the time-out period to 60 seconds and disables this option.
Inactive Flow Timeout	Optionally, specify the amount of time, in seconds, the collector retains the list of inactive traffic flows. The default value is 15 seconds.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

Related Topics

- "Configuring Subnet Side Rules" on page 260
- "Viewing Top Talkers Reports" on page 409

To add a Flow collector

1. Under Flow Collectors, complete the configuration as described in the following table.

Control	Description
Add a New Flow Collector	Displays the controls to add a Flow collector.
Collector IP Address	Specify the IP address for the Flow collector.
Port	Specify the UDP port the Flow collector is listening on. The default value is 2055.
Version	Select one of the following versions from the drop-down list: • **CascadeFlow** - Use with Cascade v8.4 or later. • **CascadeFlow-compatible** - Use with Cascade v8.34 or earlier. • **NetFlow v5** - Enables ingress flow records. • **NetFlow v9** - Enables both ingress and egress flow records. For details on using NetFlow records with Cascade, see the *Riverbed Deployment Guide*. CascadeFlow and CascadeFlow-compatible are enhanced versions of flow export to Riverbed Cascade. These versions allow automatic discovery and interface grouping for Steelhead appliances in the Riverbed Cascade Profiler or Cascade Gateway and support WAN and optimization reports in Cascade. For details, see the *Cascade Profiler User Manual* and the *Cascade Gateway User Manual*.
Packet Source Interface	Select the interface to use as the source IP address of the flow packets (Primary, Aux, or MIP) from the drop-down list. NetFlow records sent from the Steelhead appliance appear to be sent from the IP address of the selected interface.
LAN Address	Causes the TCP/IP addresses and ports reported for optimized flows to contain the original client and server IP addresses and not those of the Steelhead appliance. The default setting displays the IP addresses of the original client and server without the IP address of the Steelhead appliances. This setting is unavailable with NetFlow v9, because the optimized flows are always sent out with both the original client server IP addresses and the IP addresses used by the Steelhead appliance.
Capture Interface/Type	Specify the traffic type to export to the flow collector. Select one of the following types from the drop-down list: • **All** - Exports both optimized and non-optimized traffic. • **Optimized** - Exports optimized traffic. • **Optimized** - Exports optimized LAN or WAN traffic when WCCP is enabled. • **Passthrough** - Exports pass-through traffic. • **None** - Disables traffic flow export. The default is All for LAN and WAN interfaces, for all four collectors. The default for the other interfaces (Primary, rios_lan, and rios_wan) is None. You cannot select a MIP interface.
Enable Filter	(CascadeFlow and NetFlow v9 only) Filter flow reports by IP and subnets or IP:ports included in the Filter list. When disabled, reports include all IP addresses and subnets.
Filter	(CascadeFlow and NetFlow v9 only) Specify the IP and subnet or IP:port to include in the report, one entry per line, up to 25 filters maximum.

Control	Description
Add	Adds the collector to the Collector list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Apply** to apply your settings.

3. Click **Save** to save your settings permanently.

Applying QoS Policies

This section describes how to set Riverbed Quality of Service (QoS) policies. It includes the following sections:

- "QoS Overview" on page 267
- "Configuring Basic QoS" on page 275
- "Configuring Advanced QoS" on page 286
- "Setting ToS DSCP Values" on page 301
- "Creating a QoS Map List" on page 303

You apply Riverbed QoS policies in the Configure > Networking > Basic QoS or Advanced QoS pages. This section describes how Steelhead appliances use Riverbed QoS policies to allocate bandwidth and latency priorities, and provides specific examples for setting policies for FTP and Citrix ICA traffic.

Note: For details on QoS, including integrating Steelhead appliances into an existing QoS implementation, see the *Riverbed Deployment Guide*. The *Riverbed Deployment Guide* also includes configuration examples and Riverbed QoS best practices.

QoS Overview

QoS is a reservation system for network traffic in which you use QoS sites or classes to distribute network resources. RiOS v6.5 improves QoS with the following benefits:

- **Simplified setup** - A basic QoS configuration page streamlines setup for networks that require minimal configuration of network traffic.
- **AppFlow Engine (AFE)** - Allows advanced classification and shaping of network traffic. The Steelhead inspects classification rules for information within the TCP/UDP payload in addition to packet headers to distinguish between different traffic types.

  With AFE, QoS can identify applications accurately and differentiate applications that use the same port on the same server. For example, Sharepoint and Microsoft Background Intelligent Transfer Service (BITS) can use port 80 on the same server. Once an application is identified, you can place into different classes for QoS enforcement. AFE identification is similar to deep packet inspection (DPI) because it identifies applications based on patterns. This approach allows you to more accurately identify modern applications than signature-based DPI methods by being aware of the more complex ways they communicate and the dependencies between multiple flows.

The QoS classifier in RiOS v6.5 accommodates multiple types of traffic monitoring, including header-based and third-party protocol matching rules. A protocol matching rule contains a combination of header and Layer-7 information to identify applications accurately. By adjusting a global list of applications or class bandwidth allocations, you can use QoS in v6.5 to create endless combinations of Layer-7 applications.

To view the predefined global application list, go to

http://www.riverbed.com/us/products/technology/riverbed_classification_qos_engine.php

In addition to supporting many well-known applications, you can use AFE to add signatures to identify custom applications. For example, you can identify a new HTTP application based on a specific domain name or relative path.

You can also use AFE to classify encrypted applications, such as HTTPS. You do not need the public and private key pair in order to use AppFlow Engine.

AFE works with both pass-through and optimized traffic and is available in both the basic QoS and advanced QoS modes.

- **Port label handling** - Allows you to specify a port range for more efficient port handling.
- **Connection tracking for pass-through traffic** - Eliminates per-packet inspection of flow oriented traffic, thereby enhancing performance.
- **Rule hierarchy** - Allows you to increase the number of rules per site, up to 2000.
- **Site awareness** - Improves performance and scalability in multi-site configurations.

QoS classes are based on traffic importance, bandwidth needs, and delay-sensitivity. You allocate network resources to each of the classes. Traffic flows according to the network resources allocated to its class.

You configure QoS on client-side and server-side Steelhead appliances to control the prioritization of different types of network traffic and to ensure that Steelhead appliances give certain network traffic (for example, VoIP) higher priority over other network traffic.

Traffic Classification

QoS allows you to specify priorities for particular classes of traffic and properly distribute excess bandwidth among classes. The QoS classification algorithm provides mechanisms for link sharing and priority services while decoupling delay and bandwidth allocation.

Many QoS implementations use some form of Packet Fair Queueing (PFQ), such as Weighted Fair Queueing or Class-Based Weighted Fair Queueing. As long as high-bandwidth traffic requires a high priority (or vice-versa), PFQ systems perform adequately. However, problems arise for PFQ systems when the traffic mix includes high-priority, low-bandwidth traffic, or high-bandwidth traffic that does not require a high priority, particularly when both of these traffic types occur together. Features such as low-latency queueing (LLQ) attempt to address these concerns by introducing a separate system of strict priority queueing that is used for high-priority traffic. However, LLQ is not an effective way of handling bandwidth and latency trade-offs. LLQ is a separate queueing mechanism meant as a work around for PFQ limitations.

The Riverbed QoS system is not based on PFQ, but rather on Hierarchical Fair Service Curve (HFSC). HFSC delivers low latency to traffic without wasting bandwidth and delivers high bandwidth to delay-insensitive traffic without disrupting delay-sensitive traffic. The Riverbed QoS system achieves the benefits of LLQ without the complexity and potential configuration errors of a separate queueing mechanism.

The Steelhead appliance HFSC-based QoS enforcement system provides the flexibility needed to simultaneously support varying degrees of delay requirements and bandwidth usage. For example, you can enforce a mix of high-priority, low-bandwidth traffic patterns (for example, SSH, Telnet, Citrix, RDP, and CRM systems) with lower priority, high-bandwidth traffic (for example, FTP, backup, and replication). RiOS QoS allows you to protect delay-sensitive traffic such as VoIP, as well as other delay-sensitive traffic such as RDP and Citrix. You can do this without having to reserve large amounts of bandwidth for their traffic classes.

QoS classification occurs during connection setup for optimized traffic, before optimization and compression. QoS shaping and enforcement occurs after optimization and compression.

By design, QoS is applied to both pass-through and optimized traffic; however, you can choose to classify either pass-through or optimized traffic. QoS is implemented in the operating system; it is not a part of the optimization service. When the optimization service is disabled, all the traffic is pass-through and is still shaped by QoS.

Basic or Advanced QoS

RiOS v6.5 provides two types of QoS configurations: basic and advanced. The QoS configuration you implement depends on how much classification and shaping your network traffic requires and whether you are migrating from a previous RiOS version or configuring QoS on a Steelhead for the first time.

Advanced QoS supports different bandwidths for different interfaces; basic QoS does not, but you can specify the remote site absolute bandwidth.

After upgrading a Steelhead to RiOS v6.5, the default is:

- Basic QoS on new and upgraded Steelheads that do not have a QoS configuration.
- Advanced QoS on Steelheads that have a existing QoS configuration. The Steelhead preserves the configuration.

Use basic QoS when you:

- currently do not have RiOS QoS configured.
- are currently using RiOS v6.1.x or earlier QoS but are willing to consolidate and reconfigure your existing rules. The existing configuration is lost when you switch from advanced to basic QoS.
- do not need more granular control and can use the default settings.

Use advanced QoS when you:

- are currently using RiOS v6.1.x or earlier QoS and do not want to reconfigure your existing rules. The Steelhead preserves the configuration.
- need to use the MX-TCP queue. For details, see "Enabling MX-TCP Queue Policies (Advanced QoS only)" on page 299.
- need to set application priorities for Citrix ICA traffic (this requires packet-order queue).
- have WAN links with different bandwidth (basic QoS assumes all links of the same size). For example, you might have a 2 Mbps MPLS link with a 1 Mbps ADSL backup.

QoS Classes

QoS classes set priorities and bandwidths. Basic QoS comes with six predefined classes, and you cannot add or delete classes. In advanced QoS, you can create multiple QoS classes. There is no requirement that QoS classes represent applications, traffic to remote sites, or any other particular aggregation.

The QoS classes that are always present on the Steelhead appliance in Advanced QoS mode are:

- **Root class** - The root class is used to constrain the total outbound rate of traffic leaving the Steelhead appliance to the configured, per-link WAN bandwidth. This class is not configured directly, but is created when you enable QoS classification and enforcement on the Steelhead appliance.
- **Built-in default class** - The QoS scheduler applies the built-in default class constraints and parameters on traffic not placed in a class by the configured QoS rules.

QoS classes are configured in one of two different modes: *flat* or *hierarchical*. The difference between the two modes primarily consists of how QoS classes are created.

Note: For details on QoS classes, see the *Riverbed Deployment Guide*.

Hierarchical Mode

In hierarchical mode, you create QoS classes as children of QoS classes other than the root class. This allows you to create overall parameters for a certain traffic type, and specify parameters for subtypes of that traffic. There is no enforced limit to the number of QoS class levels you can create.

In hierarchical mode, the following relationships exist between QoS classes:

- **Sibling classes** - Classes that share the same parent class.
- **Leaf classes** - Classes at the bottom of the class hierarchy.
- **Inner classes** - Classes that are neither the root class nor leaf classes.

In hierarchical mode, QoS rules can only specify leaf classes as targets for traffic.

Riverbed QoS controls the traffic of hierarchical QoS classes in the following manner:

- QoS rules assign active traffic to leaf classes.
- The QoS scheduler:
  - applies active leaf class parameters to the traffic.
  - applies parameters to inner classes that have active leaf class children.

Flat Mode (Advanced QoS only)

In flat mode, all of the QoS classes you create must have the root class as their parent. Accordingly, all of the QoS classes you create are siblings.

The QoS scheduler treats QoS classes in flat mode the same way that it does in hierarchical mode. However, only a single class level is defined. QoS rules place active traffic into the leaf classes. Each active class has their own QoS rule parameters which the QoS scheduler applies to traffic.

Note: You can use the CMC to enable QoS and to configure and apply QoS policies centrally to Steelhead appliances. For details, see the *Steelhead Central Management Console User's Guide* and the *Riverbed Deployment Guide*.

Selecting a QoS Enforcement System

Selecting the appropriate QoS enforcement system depends on the location of WAN bottlenecks for traffic leaving the site.

Use the following guidelines when implementing advanced QoS:

- A site that acts as a data server for other locations, such as a data center or regional hub, typically uses hierarchical mode. The first level of classes represents remote sites, and those remote site classes have child classes that either represent application types, or are indirectly connected remote sites.
- A site that typically receives data from other locations, such as a branch site, typically uses flat mode. The classes represent different application types.

For example, suppose you have a network with ten locations, and you want to choose the correct mode for site 1. Traffic from site 1 normally goes to two other sites: sites 9 and 10. If the WAN links at sites 9 and 10 are at a higher bandwidth than the link at site 1, the WAN bottleneck rate for site 1 is always the link speed for site 1. In this case, you can use flat mode to enforce QoS at site 1, because the bottleneck that needs to be managed is the link at site 1. In flat mode, the parent class for all created classes is the root class that represents the WAN link at site 1.

In the same network, site 10 sends traffic to sites 1 through 8. Sites 1 through 8 have slower bandwidth links than site 10. Because the traffic from site 10 faces multiple WAN bottlenecks (one at each remote site), you configure hierarchical mode for site 10.

Note: For details on configuring QoS for a branch office and data center, see the *Riverbed Deployment Guide*.

QoS Classification for the FTP Data Channel

When configuring QoS classification for FTP, the QoS rules differ depending on whether the FTP data channel is using *active* or *passive* FTP. Active versus passive FTP determines whether the FTP client or the FTP server select the port connection for use with the data channel, which has implications for QoS classification.

Active FTP Classification

With active FTP, the FTP client logs in and enters the PORT command, informing the server which port it must use to connect to the client for the FTP data channel. Next, the FTP server initiates the connection towards the client. From a TCP perspective, the server and the client swap roles. The FTP server becomes the client because it sends the SYN packet, and the FTP client becomes the server because it receives the SYN packet.

Although not defined in the RFC, most FTP servers use source port 20 for the active FTP data channel.

For active FTP, configure a QoS rule on the server-side Steelhead appliance to match source port 20. On the client-side Steelhead appliance, configure a QoS rule to match destination port 20.

You can also use AFE to classify active FTP data.

Passive FTP Classification

With passive FTP, the FTP client initiates both connections to the server. First, it requests passive mode by entering the PASV command after logging in. Next, it requests a port number for use with the data channel from the FTP server. The server agrees to this mode, selects a random port number, and returns it to the client. Once the client has this information, it initiates a new TCP connection for the data channel to the server-assigned port. Unlike active FTP, there is no role swapping and the FTP client initiates the SYN packet for the data channel.

The FTP client receives a random port number from the FTP server. Because the FTP server cannot return a consistent port number to use with the FTP data channel, RiOS does not support QoS Classification for passive FTP in versions earlier than RiOS v4.1.8, v5.0.6, or v5.5.1. Later RiOS releases support passive FTP and the QoS Classification configuration for passive FTP is the same as active FTP.

When configuring QoS Classification for passive FTP, port 20 on both the server and client-side Steelhead appliances means the port number used by the data channel for passive FTP, as opposed to the literal meaning of source or destination port 20.

Note: The Steelhead appliance must intercept the FTP control channel (port 21), regardless of whether the FTP data channel is using active or passive FTP.

Figure 7-7. Active and Passive FTP

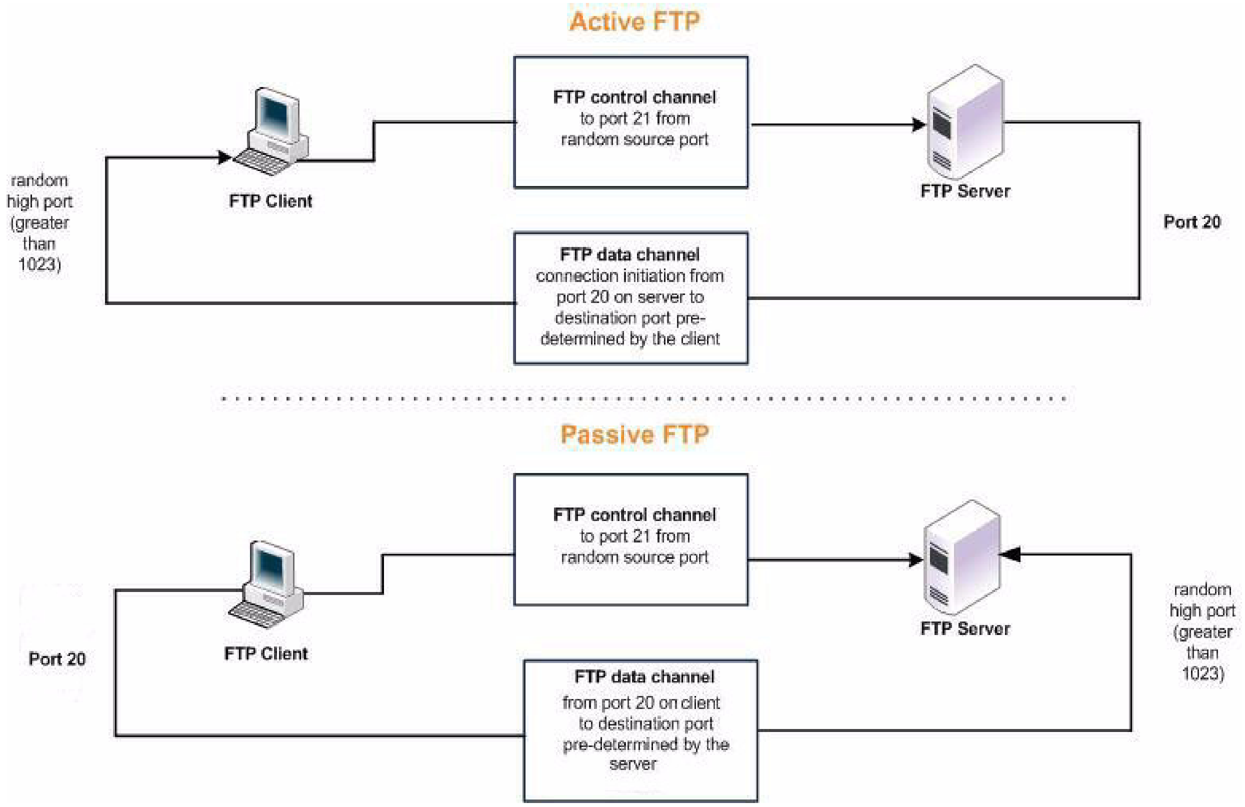
Active FTP
FTP control channel to port 21 from random source port
FTP Client
FTP Server
random high port (greater than 1023)
Port 20
FTP data channel connection initiation from port 20 on server to destination port pre-determined by the client
Passive FTP
FTP control channel to port 21 from random source port
FTP Client
FTP Server
Port 20
random high port (greater than 1023)
FTP data channel from port 20 on client to destination port pre-determined by the server

For details, see "QoS Marking with the FTP Data Channel" on page 300.

Using QoS with RSP

To route optimized traffic through a Steelhead appliance that is using QoS and RSP simultaneously, RiOS must be positioned last in the RSP data flow. To verify that RiOS is positioned correctly, choose Configure > Branch Services > RSP Data Flow. In the RSP Data Flow table at the bottom of the page, check that RiOS appears as the last item in the data flow, directly above the WAN interface. For details, see "Configuring RSP Data Flow" on page 206.

QoS Classification for Citrix Traffic

RiOS v6.0 and later provides a way to classify Citrix traffic using QoS to differentiate between different traffic types within a Citrix session. QoS classification for Citrix traffic is beneficial in mixed-use environments where Citrix users perform printing and use drive-mapping features. Using QoS to classify Citrix traffic in a mixed-use environment provides optimal network performance for end users. Note that if the Citrix sessions in your environment carry only interactive traffic, you can use basic QoS.

Citrix QoS classification provides support for Presentation Server v4.5, XenApp v5.0 and v6.0, and v10.x, v11.x, and v12.x clients.

The essential RiOS capabilities that ensure optimal delivery of Citrix traffic over the network are:

- **Latency priority** - The Citrix traffic application priority affects traffic latency. This allows you to assign interactive traffic a higher priority than print or drive-mapping traffic. A typical application priority for interactive Citrix sessions, such as screen updates, is real-time or interactive. Keep in mind that priority is relative to other classes in your QoS configuration. You must use advanced QoS.
- **Bandwidth allocation** (also known as traffic shaping) - When configuring QoS for Citrix traffic, it is important to allocate the correct amount of bandwidth for each QoS traffic class. The amount you specify reserves a pre-determined amount of bandwidth for each traffic class. Bandwidth allocation is important for ensuring that a given class of traffic cannot consume more bandwidth than it is allowed. It is also important to ensure that a given class of traffic has a minimum amount of bandwidth available for delivery of data through the network.

RiOS v6.5 provides a default rule for Citrix ICA traffic. You can simply use the default rule or edit it to suit your configuration in either basic or advanced QoS.

The default ports for the Citrix service are 1494 (native ICA traffic) and 2598 (session reliability). To use session reliability, you must enable Citrix optimization on the Steelhead appliance in order to classify the traffic correctly. You can enable and modify Citrix ICA optimization settings in the Configure > Optimization > Citrix ICA page. For details, see "Configuring Citrix ICA Optimization" on page 137.

You can use session reliability with optimized traffic only. Session reliability with RiOS QoS does not support pass-through traffic. For details about disabling session reliability, go to http://support.citrix.com/proddocs/index.jsp?topic=/xenapp5fp-w2k8/ps-sessions-sess-rel.html

---

Important: If you upgrade from a previous RiOS version with an existing Citrix QoS configuration, the upgrade automatically combines the five pre-existing Citrix rules into one.

Note: For QoS configuration examples, see the *Riverbed Deployment Guide*.

Figure 7-8. Citrix ICA QoS Default Rule

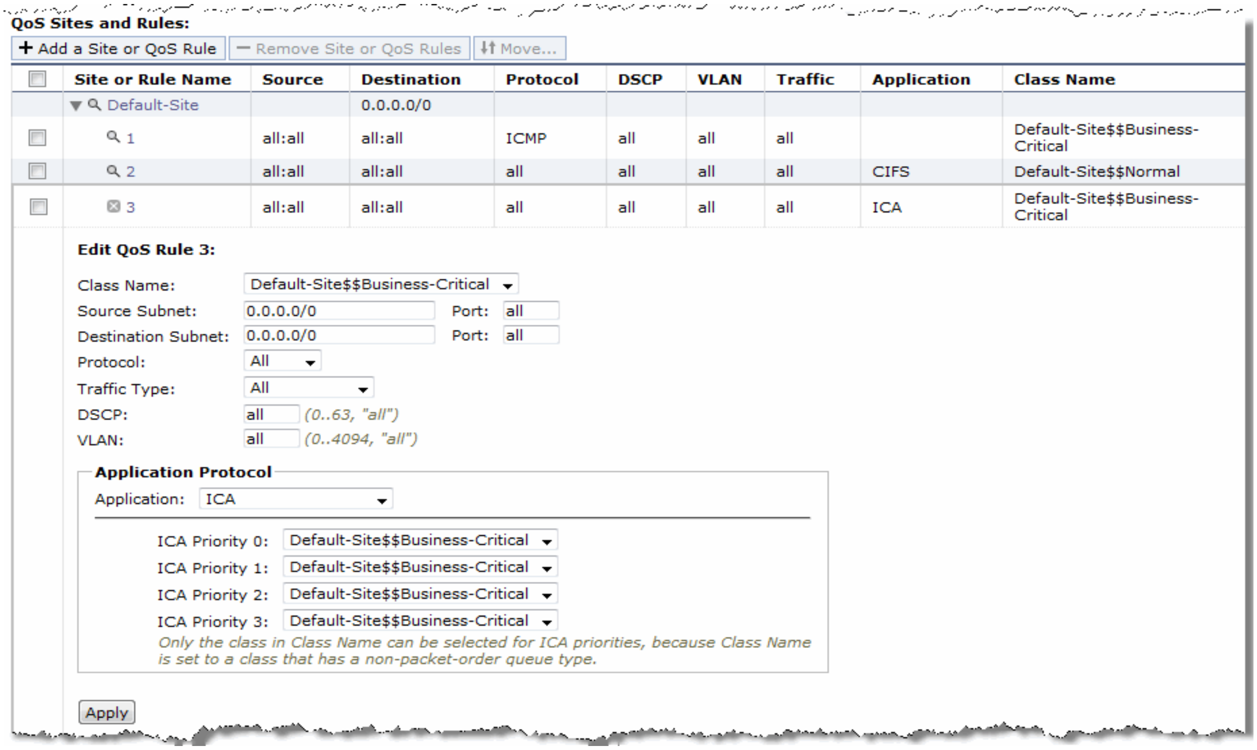

QoS Sites and Rules:

□	Site or Rule Name	Source	Destination	Protocol	DSCP	VLAN	Traffic	Application	Class Name
	▼ Default-Site		0.0.0.0/0						
□	1	all:all	all:all	ICMP	all	all	all		Default-Site$$Business-Critical
□	2	all:all	all:all	all	all	all	all	CIFS	Default-Site$$Normal
□	3	all:all	all:all	all	all	all	all	ICA	Default-Site$$Business-Critical

Edit QoS Rule 3:

Class Name: Default-Site$$Business-Critical

Source Subnet: 0.0.0.0/0 Port: all

Destination Subnet: 0.0.0.0/0 Port: all

Protocol: All

Traffic Type: All

DSCP: all *(0..63, "all")*

VLAN: all *(0..4094, "all")*

Application Protocol

Application: ICA

ICA Priority 0: Default-Site$$Business-Critical

ICA Priority 1: Default-Site$$Business-Critical

ICA Priority 2: Default-Site$$Business-Critical

ICA Priority 3: Default-Site$$Business-Critical

Only the class in Class Name can be selected for ICA priorities, because Class Name is set to a class that has a non-packet-order queue type.

Configuring Basic QoS

This section describes how to configure basic QoS. It contains the following sections:

- "Overview" on page 275
- "Adding a Remote Site" on page 278
- "Adding an Application" on page 280
- "Adding a Service Policy" on page 283

For a QoS overview, see "Applying QoS Policies" on page 267. For information on whether to deploy basic or advanced QoS, see "Basic or Advanced QoS" on page 269.

Overview

Basic QoS simplifies QoS configuration by accurately identifying business applications and classifying traffic according to priorities. The Steelhead uses this information to control the amount of WAN resources that each application can use. This ensures that your important applications are prioritized and removes the guesswork from protecting performance of key applications. In addition, basic QoS prevents recreational applications from interfering with business applications.

Basic QoS comes with a predefined set of six classes, a list of global applications, and a predefined set of policies. All interfaces have the same link rate.

To view the predefined global application list, go to

http://www.riverbed.com/us/products/technology/riverbed_classification_qos_engine.php

Basic QoS includes a default site that is tied to the predefined service policy Medium Office. The bandwidth for the default site is automatically set to the same bandwidth as the interface's WAN throughput value. You can edit the bandwidth for the default site but you cannot edit the subnet.

You cannot add or delete classes in basic QoS.

Enabling Local WAN Oversubscription

Basic QoS includes an optional local WAN oversubscription feature that allows the sum of remote site bandwidths to exceed the WAN uplink speed. Riverbed recommends enabling this option when your network includes remote site bandwidths that collectively exceed the available bandwidth of the local WAN uplink interface speed when a subset of remote sites are active at once. This feature is only available in basic QoS.

Figure 7-9. Bandwidth Oversubscription Feature in Basic QoS

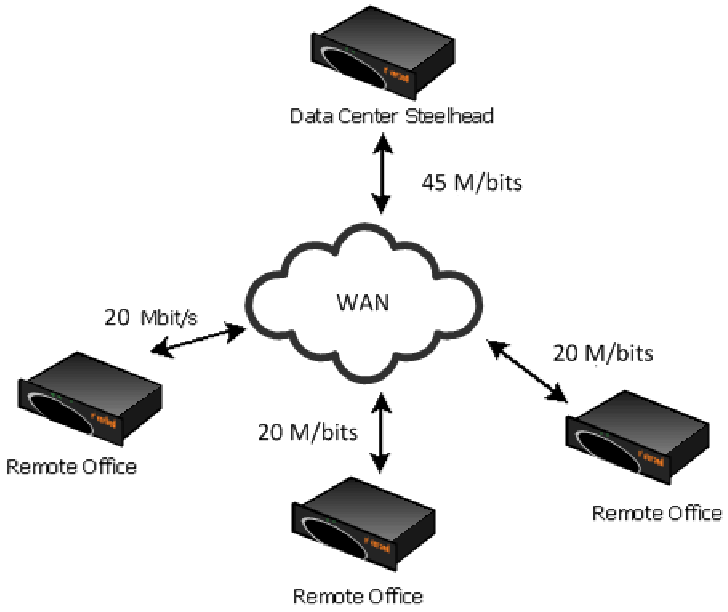

To enable basic QoS

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

Figure 7-10. Basic QoS Page

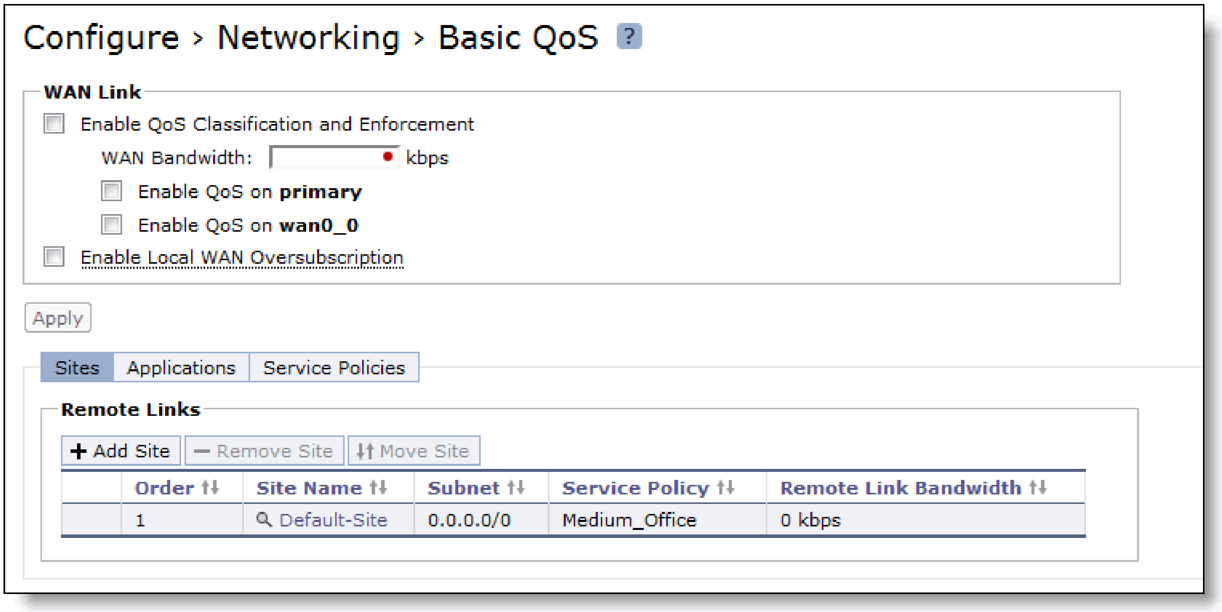

Configure > Networking > Basic QoS

WAN Link

□ Enable QoS Classification and Enforcement

WAN Bandwidth: ______ kbps

□ Enable QoS on **primary**

□ Enable QoS on **wan0_0**

□ Enable Local WAN Oversubscription

	Order	Site Name	Subnet	Service Policy	Remote Link Bandwidth
	1	Default-Site	0.0.0.0/0	Medium_Office	0 kbps

2. Under WAN Link, complete the configuration as described in the following table.

Control	Description
Enable QoS Classification and Enforcement	Enables QoS classification to control the prioritization of different types of network traffic and to ensure that the Steelhead gives certain network traffic (for example, Voice Over IP) higher priority than other network traffic. Traffic is not classified until at least one WAN interface is enabled. To disable QoS, clear this check box and restart the optimization service.
WAN Bandwidth (kbps)	Specify the interface bandwidth link rate in kbps. The Steelhead automatically sets the bandwidth for the default site to this value. The link rate is the *bottleneck* WAN bandwidth, not the interface speed out of the WAN interface into the router or switch. For example, if your Steelhead connects to a router with a 100 Mbps link, do not specify this value—specify the actual WAN bandwidth (for example, T1, T3). **Important:** Different WAN interfaces can have different WAN bandwidths; you must enter the bandwidth link rate correctly for QoS to function properly.

Control	Description
Enable QoS on <interface>	Specify a WAN interface <XXXX-X> to enable.
Enable Local WAN Oversubscription	Optionally, select to allow the sum of remote site bandwidths to exceed the WAN uplink speed. Bandwidth oversubscription shares the bandwidth fairly when the network includes remote site bandwidths that collectively exceed the available bandwidth of the local WAN uplink interface speed. The link sharing provides bandwidth guarantees when some of the sites are partially or fully inactive. For example, your data center uplink might be 45 Mbit/s with three remote office sites each with 20 Mbit/s uplinks. When disabled, you can only allocate bandwidth for the remote sites such that the total bandwidth does not exceed the bandwidth of any of the interfaces on which QoS is enabled. **Note:** Enabling this option can degrade latency guarantees when the remote sites are fully active.

3. Click **Apply** to apply your settings.

   A message tells you the WAN throughput for the default site has been set, and the throughput appears in the Remote sites table.

4. Click **Save** to save your settings permanently.

5. You can optionally customize QoS further by adjusting the global application list or adjusting the class bandwidth allocations as described in the following sections. When you finish configuring basic QoS, select the Applications tab to make sure the applications belong to the desired class, and restart the optimization service.

Adding a Remote Site

The Sites tab provides you with the ability to optionally add a remote site. A site is a logical grouping of subnets. Sites represent the physical and logical topology of a site type. You can classify traffic for each site using network addresses. Site types are typically data center, small, medium and large branch office, and so on. Each site uses a service policy, and the sites have an order. Traffic is matched to the first matching site.

The overal maximum number of rules is 2000. The maximum number of sites is 100.

The default site is a catch-all site that has a subnet of 0.0.0.0/0. You do not need to add a remote site if you only have one remote site and the default site is suitable.

To add a remote site

1. On the client-side and server-side Steelheads, choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Select the Sites tab.

Figure 7-11. Basic QoS Page - Sites

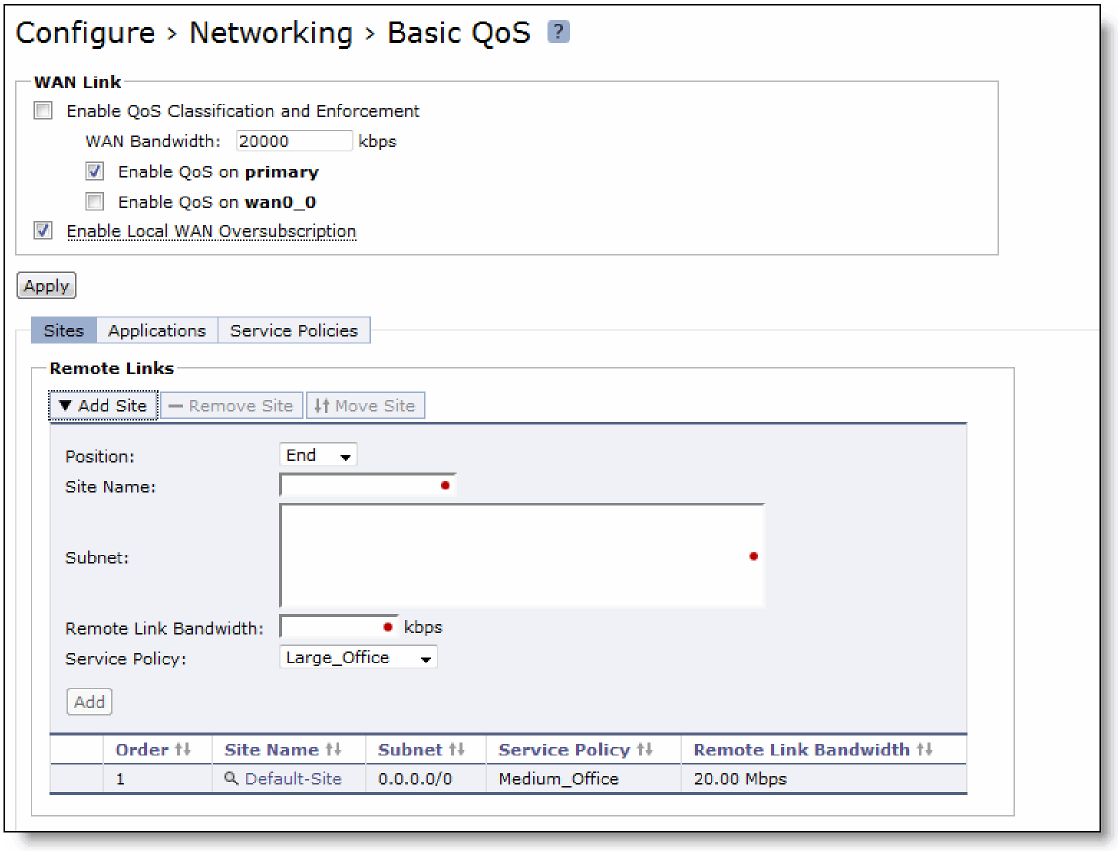

WAN Link

□ Enable QoS Classification and Enforcement

WAN Bandwidth: 20000 kbps

☑ Enable QoS on **primary**

□ Enable QoS on **wan0_0**

☑ Enable Local WAN Oversubscription

Remote Links

Position: End

Site Name:

Subnet:

Remote Link Bandwidth: kbps

Service Policy: Large_Office

	Order	Site Name	Subnet	Service Policy	Remote Link Bandwidth
	1	Default-Site	0.0.0.0/0	Medium_Office	20.00 Mbps

3. Under Remote Links, complete the configuration as described in the following table.

Control	Description
Add Site	Displays the controls to define a remote site.
Position	Select Start or End from the drop-down list. Steelhead appliances evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted. **Note:** The default site, which is tied to the Medium Office policy, cannot be removed and is always listed last.
Site Name	Specify the site name. For example, data center.
Subnet	Specify a maximum of five destination subnets that represent individual sites. **Note:** You cannot edit the subnet for the default site.
Remote Link Bandwidth	Specify the maximum WAN bandwidth in kbps.

Control	Description
Service Policy	Optionally, select a service policy from the drop-down list. The default policy is Large Office.
Add	Adds the site to the list. The Management Console redisplays the Sites table and applies your modifications to the running configuration, which is stored in memory. This button is dimmed and unavailable until you enter the WAN bandwidth.
Remove Site	Select the check box next to the name and click **Remove Site**.
Move Site	Moves the selected sites. Click the arrow next to the desired rule position; the site moves to the new position.

4. Click **Apply** to apply your settings.

Adding an Application

An application definition determines the global performance rules for an application, including latency priority. The Applications tab provides the ability to map classification parameters (for example, name and header) to a predefined service class (latency priority) and the ability to specify a rule order for the mappings.

You can select an application protocol definition from a predefined global application list or you can add a custom application.

To view the predefined global application list, go to

http://www.riverbed.com/us/products/technology/riverbed_classification_qos_engine.php

To define custom applications or edit existing application definitions

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Select the Applications tab.

Figure 7-12. Basic QoS Page - Applications

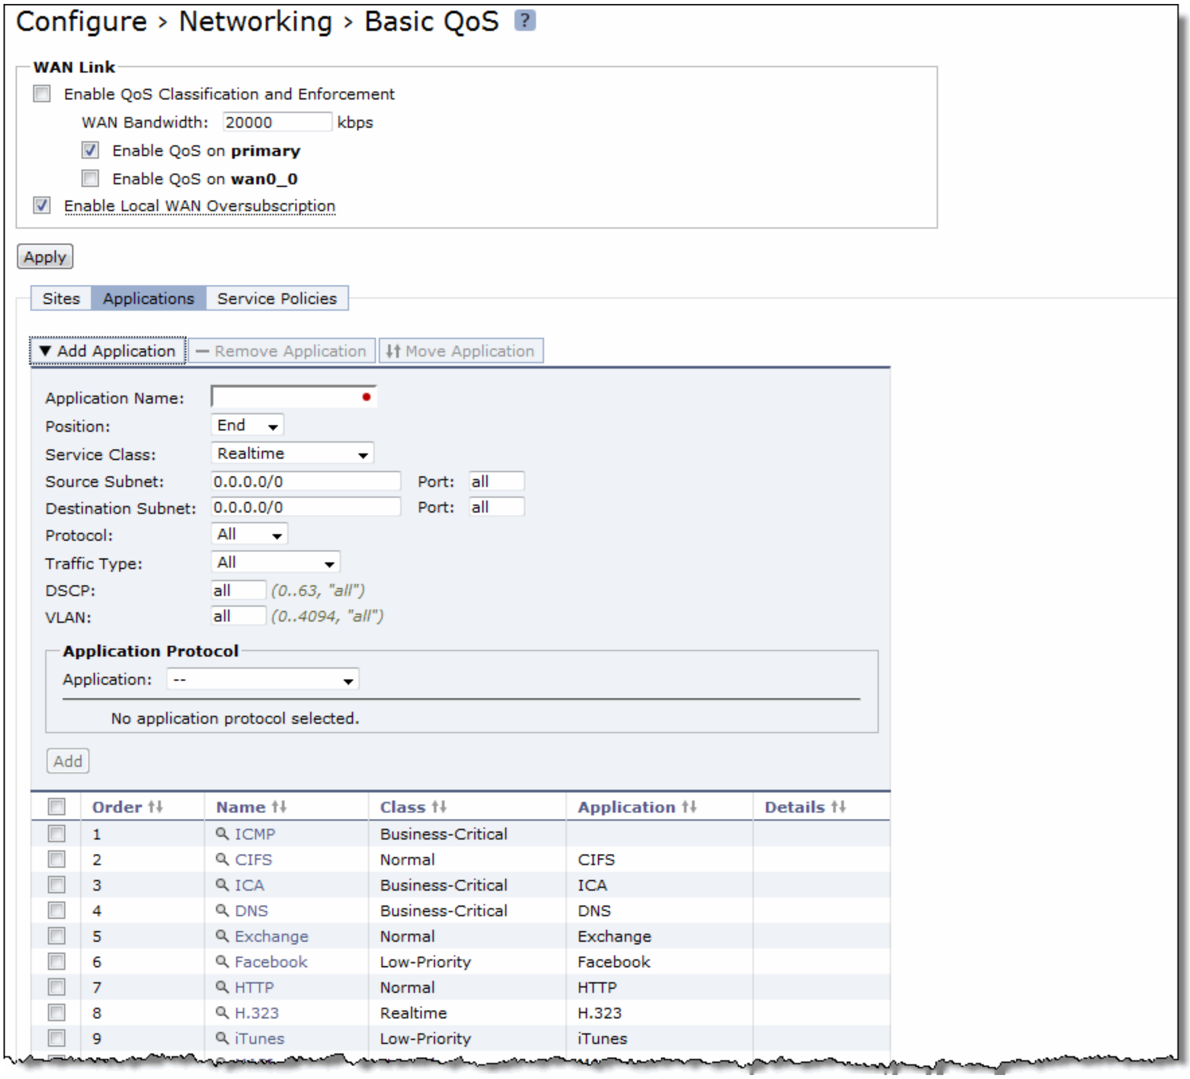

Configure > Networking > Basic QoS

WAN Link

☐ Enable QoS Classification and Enforcement

WAN Bandwidth: 20000 kbps

☑ Enable QoS on **primary**

☐ Enable QoS on **wan0_0**

☑ Enable Local WAN Oversubscription

Application Name:	
Position:	End
Service Class:	Realtime
Source Subnet:	0.0.0.0/0 Port: all
Destination Subnet:	0.0.0.0/0 Port: all
Protocol:	All
Traffic Type:	All
DSCP:	all *(0..63, "all")*
VLAN:	all *(0..4094, "all")*

Application Protocol

Application: --

No application protocol selected.

☐	Order	Name	Class	Application	Details
☐	1	ICMP	Business-Critical		
☐	2	CIFS	Normal	CIFS	
☐	3	ICA	Business-Critical	ICA	
☐	4	DNS	Business-Critical	DNS	
☐	5	Exchange	Normal	Exchange	
☐	6	Facebook	Low-Priority	Facebook	
☐	7	HTTP	Normal	HTTP	
☐	8	H.323	Realtime	H.323	
☐	9	iTunes	Low-Priority	iTunes	

3. To define a custom application and add it to the application list, complete the configuration as described in the following table.

Control	Description
Add Application	Displays the controls to define an application.
Application Name	Optionally, specify the application name, for example, Outlook Anywhere.
Position	Select Start, End, or a rule number from the drop-down list. Steelhead appliances evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted.
Source Subnet	Specify an IP address and mask for the traffic source, or you can specify all or 0.0.0.0/0 as the wildcard for all traffic. Use the following format: XXX.XXX.XXX.XXX/XX
Source Port	Optionally, specify all source ports, a single source port value or a port range of port1-port2, where port1 must be less than port2. The default setting is all ports.
Destination Subnet	Specify an IP address and mask pattern for the traffic destination, or you can specify all or 0.0.0.0/0 as the wildcard for all traffic. Use the following format: XXX.XXX.XXX.XXX/XX
Destination Port	Optionally, specify all destination ports, a single source port value or a port range of port1-port2, where port1 must be less than port2. The default setting is all ports.
Protocol	Select All, TCP, UDP GRE, ICMP, or IPsec from the drop-down list. The default setting is All.
Traffic Type	Select Optimized, Passthrough, or All from the drop-down list. The default setting is All.
DSCP	Optionally, specify a DSCP value from 0 to 63, or all to use all DSCP values.
VLAN	Optionally, specify a VLAN tag as follows: • Specify a numeric VLAN tag identification number from 0 to 4094. • Specify all to specify the rule applies to all VLANs. • Specify none to specify the rule applies to untagged connections. RiOS supports VLAN v802.1Q. To configure VLAN tagging, configure transport rules to apply to all VLANs or to a specific VLAN. By default, rules apply to all VLAN values unless you specify a particular VLAN ID. Pass-through traffic maintains any pre-existing VLAN tagging between the LAN and WAN interfaces.

Control	Description
Service Class	The service class indicates how delay-sensitive a traffic class is to the QoS scheduler. Select a service class for the application from the drop-down list (highest priority to lowest): • **Real-Time** - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, or video conferencing. • **Interactive** - Specifies an interactive traffic class. For example, Citrix, RDP, telnet and ssh. • **Business Critical** - Specifies the high priority traffic class. For example, Thick Client Applications, ERPs, and CRMs. • **Normal Priority** - Specifies a normal priority traffic class. For example, Internet browsing, file sharing, and email. • **Low Priority** - Specifies a low priority traffic class. For example, FTP, backup, replication, other high-throughput data transfers, and recreational applications such as audio file sharing. • **Best Effort** - Specifies the lowest priority. These are minimum service class guarantees; if better service is available, it is provided. For example, if a class is specified as low priority and the higher priority classes are not active, then the low priority class receives the highest possible available priority for the current traffic conditions. This parameter controls the priority of the class relative to the other classes. **Important:** The service class describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how *important* the traffic is compared to other classes. Typically you configure low priority for high-throughput, non-packet delay sensitive applications like FTP, backup, and replication.
Application	Select an application from the drop-down list. Selecting HTTP expands the controls to include the Domain Name and Relative Path controls. The relative path is the part of the URL that follows the domain name.
Add	Adds the rule to the list. The Management Console redisplays the Applications table and applies your modifications to the running configuration, which is stored in memory.
Remove Application	Select the check box next to the name and click **Remove Application**.
Move Application	Moves the selected applications. Click the arrow next to the desired rule position; the application moves to the new position.

Adding a Service Policy

You can use the default policies or you can optionally add a service policy to allocate a bandwidth percentage for any of the six predefined service classes. When you create a service policy, you are configuring a template for the service classes to use preconfigured priorities.

The default policies appear in the policies list.

To add a service policy

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Select the Service Policies tab.

Figure 7-13. Basic QoS Page - Service Policies

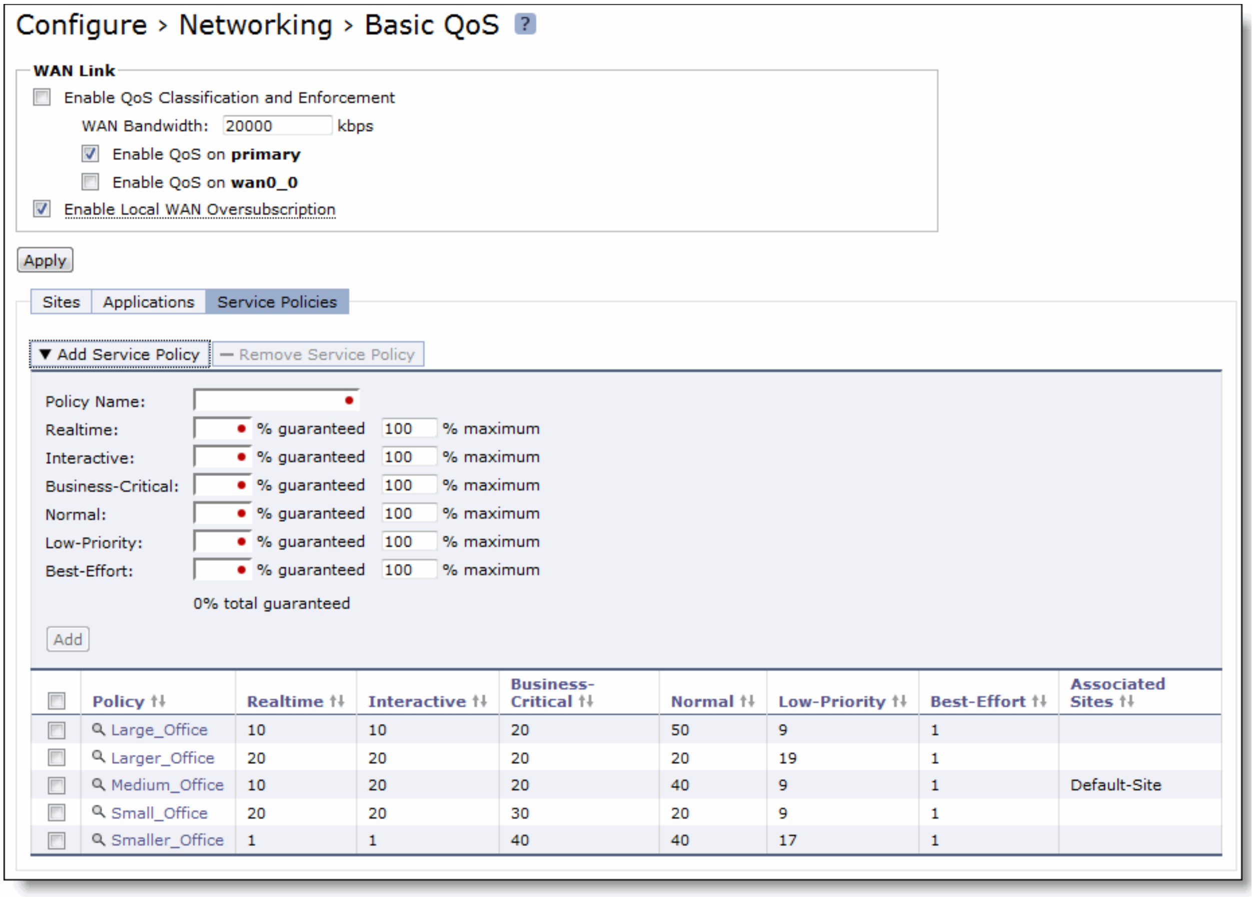

Configure > Networking > Basic QoS ?

WAN Link

□ Enable QoS Classification and Enforcement

WAN Bandwidth: 20000 kbps

☑ Enable QoS on **primary**

□ Enable QoS on **wan0_0**

☑ Enable Local WAN Oversubscription

Policy Name: ●

Realtime: ● % guaranteed 100 % maximum

Interactive: ● % guaranteed 100 % maximum

Business-Critical: ● % guaranteed 100 % maximum

Normal: ● % guaranteed 100 % maximum

Low-Priority: ● % guaranteed 100 % maximum

Best-Effort: ● % guaranteed 100 % maximum

0% total guaranteed

□	Policy	Realtime	Interactive	Business-Critical	Normal	Low-Priority	Best-Effort	Associated Sites
□	Large_Office	10	10	20	50	9	1	
□	Larger_Office	20	20	20	20	19	1	
□	Medium_Office	10	20	20	40	9	1	Default-Site
□	Small_Office	20	20	30	20	9	1	
□	Smaller_Office	1	1	40	40	17	1	

3. Complete the configuration as described in the following table.

Control	Description
Add Service Policy	Displays the controls to add a service policy.
Policy Name	Specify the policy name. For example, New York Office.
Realtime	Specify the percentage to allocate for the guaranteed and maximum bandwidth. The guaranteed bandwidth is the percentage of the bandwidth that is guaranteed to be allocated to the applications in the traffic class. A lower value indicates that the traffic in the class is more likely to be delayed. The maximum bandwidth is the maximum percentage of the bandwidth that can be allocated to the applications in the traffic class. A zero indicates that all traffic in the class is dropped.
Interactive	Specify the percentage to allocate for the guaranteed and maximum bandwidth.
Business-Critical	Specify the percentage to allocate for the guaranteed and maximum bandwidth.
Normal	Specify the percentage to allocate for the guaranteed and maximum bandwidth.
Low-Priority	This is the default service policy; specify the percentage to allocate for the guaranteed and maximum bandwidth.
Best Effort	Specify the percentage to allocate for the guaranteed and maximum bandwidth.
Add	Adds the service policy to the list. The Management Console redisplays the Policies table and applies your modifications to the running configuration, which is stored in memory.
Remove Service Policy	Select the check box next to the name and click **Remove Service Policy**.

To modify the maximum bandwidth and bandwidth guarantees for a service policy

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Click the magnifying glass next to a policy name in the policy list and specify the guaranteed and maximum bandwidth percentage.

Figure 7-14. Basic QoS Page - Modifying a Service Policy

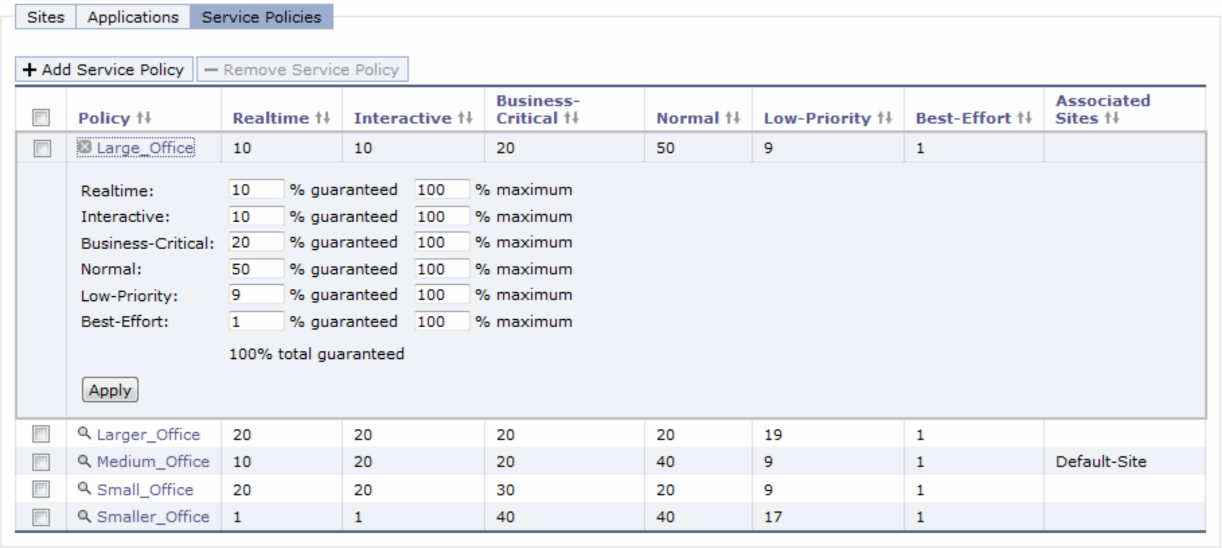

	Policy	Realtime	Interactive	Business-Critical	Normal	Low-Priority	Best-Effort	Associated Sites
□	☒ Large_Office	10	10	20	50	9	1	

Realtime: 10 % guaranteed 100 % maximum
Interactive: 10 % guaranteed 100 % maximum
Business-Critical: 20 % guaranteed 100 % maximum
Normal: 50 % guaranteed 100 % maximum
Low-Priority: 9 % guaranteed 100 % maximum
Best-Effort: 1 % guaranteed 100 % maximum

100% total guaranteed

Apply

□	Larger_Office	20	20	20	20	19	1	
□	Medium_Office	10	20	20	40	9	1	Default-Site
□	Small_Office	20	20	30	20	9	1	
□	Smaller_Office	1	1	40	40	17	1	

3. Click **Apply** to apply your settings.

Configuring Advanced QoS

You configure advanced QoS in the Configure > Networking > Advanced QoS page. Advanced QoS provides a greater degree of configurability than basic QoS; for example, you can separate rules by sites and you can perform AppFlow Engine matching.

- If you are configuring QoS for the first time, you need to migrate from basic to advanced QoS. For details, see "Migrating from Basic to Advanced QoS" on page 286.
- If you are upgrading a Steelhead with an existing QoS configuration running RiOS v6.1.x or earlier, the system automatically upgrades to advanced QoS. For details, see "Advanced QoS Steps" on page 287.

Migrating from Basic to Advanced QoS

After upgrading a Steelhead with no QoS configuration running RiOS v6.1.x or earlier to RiOS v6.5, you must migrate from basic to advanced QoS on both the client-side and server-side Steelhead appliances before configuring advanced QoS.

You might also want to migrate from basic to advanced QoS after configuring basic and finding you need more control.

To migrate from basic to advanced QoS

1. Choose Configure > Networking > Advanced QoS to display the Advanced QoS page.

Figure 7-15. Advanced QoS - Migrate Page

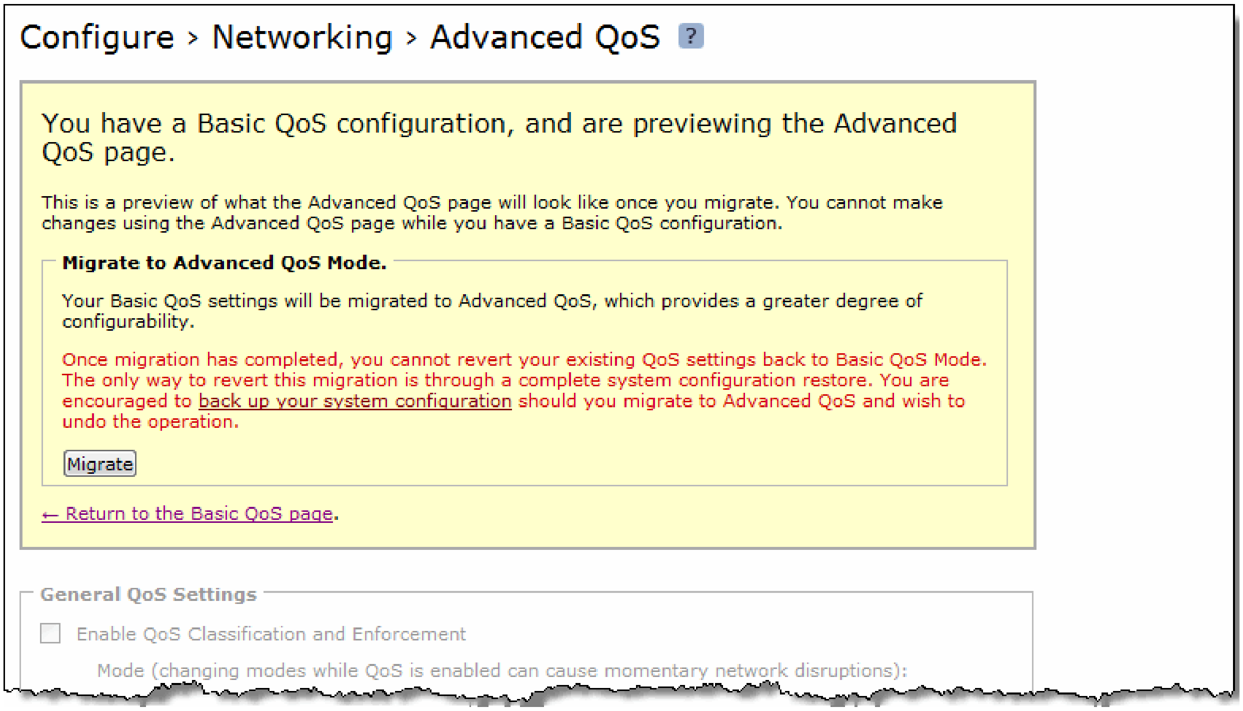

Configure > Networking > Advanced QoS

You have a Basic QoS configuration, and are previewing the Advanced QoS page.

This is a preview of what the Advanced QoS page will look like once you migrate. You cannot make changes using the Advanced QoS page while you have a Basic QoS configuration.

Migrate to Advanced QoS Mode.

Your Basic QoS settings will be migrated to Advanced QoS, which provides a greater degree of configurability.

Once migration has completed, you cannot revert your existing QoS settings back to Basic QoS Mode. The only way to revert this migration is through a complete system configuration restore. You are encouraged to back up your system configuration should you migrate to Advanced QoS and wish to undo the operation.

General QoS Settings

□ Enable QoS Classification and Enforcement

Mode (changing modes while QoS is enabled can cause momentary network disruptions):

2. Riverbed recommends that you back up your current system configuration. Optionally, click **back up your system configuration**.

3. In the Configure > Configurations page, under Save Current Configuration, specify a filename and click **Save**.

4. Choose Configure > Networking > Advanced QoS to return to the Advanced QoS page.

5. Click **Migrate**.

6. Click **OK**.

   A message confirms that the migration was successful. You can now configure Advanced QoS.

Advanced QoS Steps

The following table describes the steps for configuring advanced QoS, followed by detailed procedures.

Task	Reference
1. Select each WAN interface and define the bandwidth link rate for each interface.	"To enable advanced QoS" on page 287
2. Select the Enable QoS Classification and Enforcement check box.	"To enable advanced QoS" on page 287
3. Select either Flat or Hierarchical QoS.	"To enable advanced QoS" on page 287
4. Define the QoS classes for each traffic flow.	"To add a QoS class" on page 289
5. Add sites and define rules for each class or subclass.	"Adding a QoS Site or Rule (Advanced QoS)" on page 294
6. Restart the optimization service.	"Starting and Stopping the Optimization Service" on page 345

Important: If you delete or add new rules, the existing connections are not effected; the changes only affect new connections.

To enable advanced QoS

1. Choose Configure > Networking > Advanced QoS to display the Advanced QoS page.

Figure 7-16. Advanced QoS Page

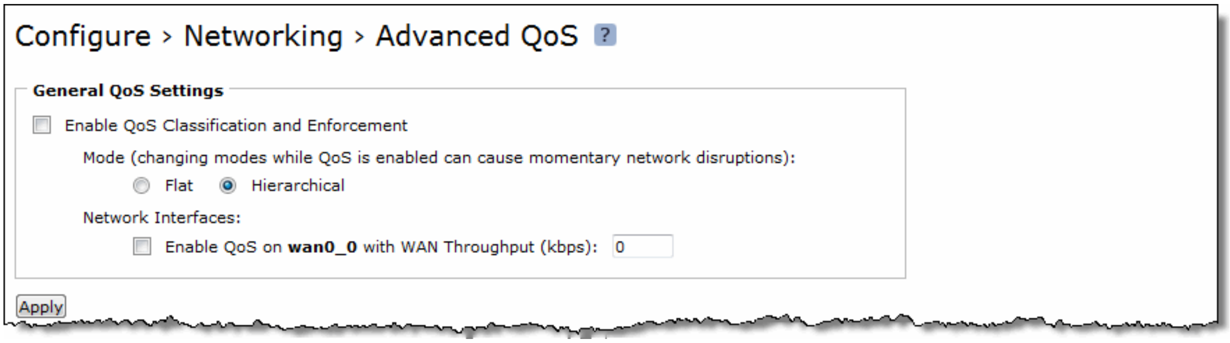

Configure > Networking > Advanced QoS ?

General QoS Settings

□ Enable QoS Classification and Enforcement

Mode (changing modes while QoS is enabled can cause momentary network disruptions):

○ Flat ● Hierarchical

Network Interfaces:

□ Enable QoS on **wan0_0** with WAN Throughput (kbps): 0

2. Under General QoS Settings, complete the configuration as described in the following table.

Control	Description
Enable QoS Classification and Enforcement	Enables QoS classification to control the prioritization of different types of network traffic and to ensure that the Steelhead gives certain network traffic (for example, Voice Over IP) higher priority than other network traffic. Traffic is not classified until at least one WAN interface is enabled. To disable QoS, clear this check box and restart the optimization service.
Mode	Specify a QoS structure: • Flat mode creates all classes at the same level. • Hierarchical mode creates a tree structure that can contain children of class parents. This is the default setting. Use this setting to segregate traffic based on flow source or destination and apply different shaping rules to each child. Use a hierarchical structure to effectively manage and support remote sites with different bandwidth characteristics. **Note:** Selecting a QoS mode does not enable QoS traffic classification. You must select the Enable QoS Classification and Enforcement check box and set a bandwidth link rate before traffic optimization begins. **Important:** Changing the QoS enforcement mode while QoS is enabled can cause a momentary service disruption to traffic flowing through the Steelhead appliance. Riverbed recommends that you configure QoS while the QoS functionality is disabled and only enable it after you are ready for the changes to take effect.
Enable QoS on <interface> with WAN bandwidth (kbps)	Enables a WAN interface <XXXX-X>. Specify its bandwidth link rate in kbps. The bandwidth for the default site is automatically set to this value. The link rate is the *bottleneck* WAN bandwidth, not the interface speed out of the WAN interface into the router or switch. For example, if your Steelhead appliance connects to a router with a 100 Mbps link, do not specify this value—specify the actual WAN bandwidth (for example, T1, T3). **Important:** Different WAN interfaces can have different WAN bandwidths; you must enter the bandwidth link rate correctly for QoS to function properly.

3. Click **Apply** to apply your settings.

   A message tells you the WAN throughput for the default site has been set, and the throughput appears in the Remote sites table.

4. Click **Save** to save your settings permanently.

5. If you are finished configuring advanced QoS, restart the optimization service.

Creating QoS Classes

Priorities and bandwidths are set by QoS class. You can create multiple classes.

Note: For details on QoS, see the *Riverbed Deployment Guide*.

To add a QoS class

1. Choose Configure > Networking > Advanced QoS to display the Advanced QoS page.

Figure 7-17. Default Advanced QoS Classes

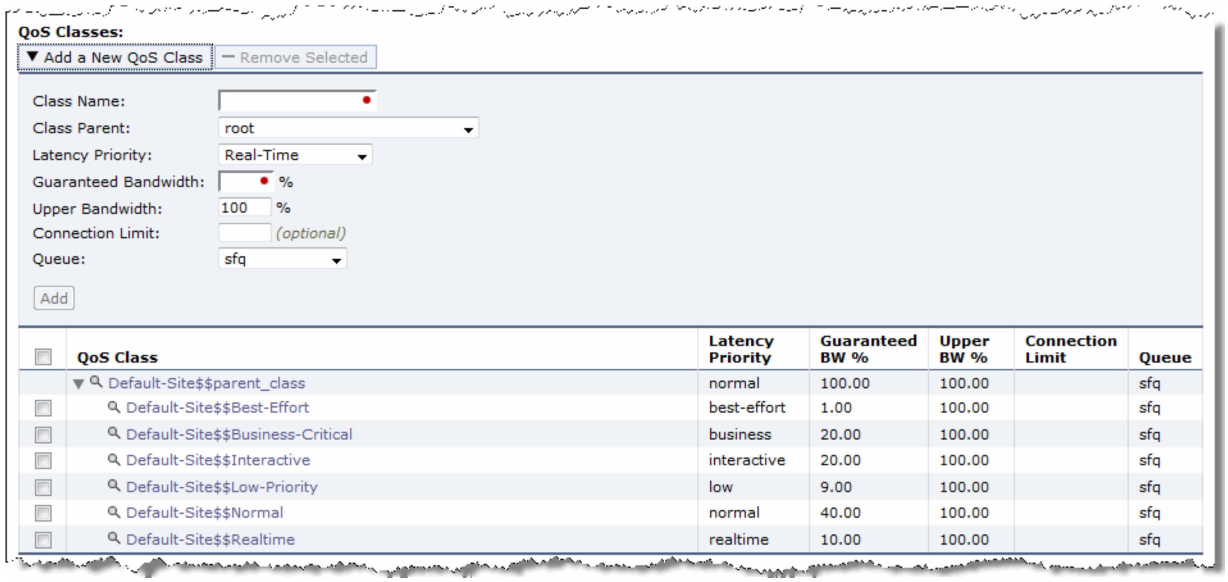

QoS Classes:

Class Name:	
Class Parent:	root
Latency Priority:	Real-Time
Guaranteed Bandwidth:	%
Upper Bandwidth:	100 %
Connection Limit:	(optional)
Queue:	sfq

□	QoS Class	Latency Priority	Guaranteed BW %	Upper BW %	Connection Limit	Queue
	▼ Default-Site$$parent_class	normal	100.00	100.00		sfq
□	Default-Site$$Best-Effort	best-effort	1.00	100.00		sfq
□	Default-Site$$Business-Critical	business	20.00	100.00		sfq
□	Default-Site$$Interactive	interactive	20.00	100.00		sfq
□	Default-Site$$Low-Priority	low	9.00	100.00		sfq
□	Default-Site$$Normal	normal	40.00	100.00		sfq
□	Default-Site$$Realtime	realtime	10.00	100.00		sfq

2. Under QoS Classes, complete the configuration as described in the following table.

Control	Description
Add a New QoS Class	Displays the controls for adding a class.
Class Name	Specify a name for the QoS class.
Class Parent	Appears only when a QoS hierarchy is enabled. Select the parent for a child class. The class inherits the parent's definitions. For example, if the parent class has a business critical latency priority, and its child has a real-time latency priority, the child inherits the business critical priority from its parent, and uses a real-time priority only with respect to its siblings. Select a class parent from the drop-down list.

Control	Description
Latency Priority	Indicates how delay-sensitive a traffic class is to the QoS scheduler. Select the latency priority for the class from the drop-down list (highest priority to lowest): • **Real-Time** - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, video conferencing. • **Interactive** - Specifies an interactive traffic class. For example, Citrix, RDP, telnet and ssh. • **Business Critical** - Specifies the high priority traffic class. For example, Thick Client Applications, ERPs, and CRMs. • **Normal Priority** - Specifies a normal priority traffic class. For example, Internet browsing, file sharing, and email. • **Low Priority** - Specifies a low priority traffic class for all traffic that does not fall into any other service class. For example, FTP, backup, replication, other high-throughput data transfers, and recreational applications such as audio file sharing. • **Best Effort** - Specifies the lowest priority. These are minimum priority guarantees; if better service is available, it is provided. For example, if a class is specified as low priority and the higher priority classes are not active, then the low priority class receives the highest possible available priority for the current traffic conditions. This parameter controls the priority of the class relative to the other classes. **Important:** The latency priority describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how *important* the traffic is compared to other classes. Typically, you configure low latency priority for high-throughput, non-packet delay sensitive applications like FTP, backup, and replication.
Guaranteed Bandwidth	Specify the minimum amount of bandwidth (as a percentage) to guarantee to a traffic class when there is bandwidth contention. All of the classes combined cannot exceed 100%. During contention for bandwidth, the class is guaranteed the amount of bandwidth specified. The class receives more bandwidth if there is unused bandwidth remaining. The guaranteed bandwidth must fall within the bandwidth limit for the Steelhead appliance. In hierarchical mode, excess bandwidth is allocated based on the relative ratios of guaranteed bandwidth. The total minimum guaranteed bandwidth of all QoS classes must be less than or equal to 100% of the parent class. A default class is automatically created with guaranteed bandwidth of 10%. Traffic that does not match any of the rules is put into the default class. Riverbed recommends that you change the guaranteed bandwidth of the default class to the appropriate value. The guaranteed bandwidth calculated based on this percentage must be no less than 1 kbps.
Link Share Weight	Specify the weight for the class. Applies to flat mode only. The link share weight determines how the excess bandwidth is allocated among sibling classes. Link share does not depend on the minimum guaranteed bandwidth. By default, all the link shares are equal. Classes with a larger weight are allocated more of the excess bandwidth than classes with a lower link share weight. You cannot specify a Link Share Weight in Hierarchical QoS. In Hierarchical QoS, the link share weight is the same proportion as the guaranteed bandwidth of the class. The Link Share Weight does not apply to MX-TCP queues.

Control	Description
Upper Bandwidth	Specify the maximum allowed bandwidth (as a percentage) a class receives as a percentage of the parent class guaranteed bandwidth. The limit is applied even if there is excess bandwidth available. Upper Bandwidth does not apply to MX-TCP queues.
Connection Limit	Optionally, specify the maximum number of optimized connections for the class. When the limit is reached, all new connections are passed through unoptimized. In hierarchical mode, a parent class connection limit does not affect its child. Each child class optimized connection is limited by the connection limit specified for their class. For example, if B is a child of A, and the connection limit for A is set to 5, while the connection limit for B is set to 10, the connection limit for B is 10. Connection Limit is supported only in in-path configurations. It is not supported in out-of-path or virtual-in-path configurations. Connection Limit does not apply to the packet-order queue or Citrix ICA traffic. RiOS does not support a connection limit assigned to any QoS class that is associated with a QoS rule with an AFE component. An AFE component consists of a Layer-7 protocol specification. RiOS cannot honor the class connection limit because the QoS scheduler may subsequently re-classify the traffic flow after applying a more precise match using AFE identification.

Control	Description
Queue	Optionally, select one of the following queue methods for the leaf class from the drop-down list (the queue does not apply to the inner class): • **SFQ** - Shared Fair Queueing (SFQ) is the default queue for all classes. Determines Steelhead appliance behavior when the number of packets in a QoS class outbound queue exceeds the configured queue length. When SFQ is used, packets are dropped from within the queue in a round-robin fashion, among the present traffic flows. SFQ ensures that each flow within the QoS class receives a fair share of output bandwidth relative to each other, preventing bursty flows from starving other flows within the QoS class. • **FIFO** - Transmits all flows in the order that they are received (first in, first out). Bursty sources can cause long delays in delivering time-sensitive application traffic and potentially to network control and signaling messages. • **MXTCP** - Has very different use cases than the other queue parameters. MX-TCP also has secondary effects that you need to understand before configuring: – When optimized traffic is mapped into a QoS class with the MX-TCP queuing parameter, the TCP congestion control mechanism for that traffic is altered on the Steelhead appliance. The normal TCP behavior of reducing the outbound sending rate when detecting congestion or packet loss is disabled, and the outbound rate is made to match the minimum guaranteed bandwidth configured on the QoS class. – You can use MX-TCP to achieve high-throughput rates even when the physical medium carrying the traffic has high loss rates. For example, MX-TCP is commonly used for ensuring high throughput on satellite connections where a lower-layer-loss recovery technique is not in use. – Another usage of MX-TCP is to achieve high throughput over high-bandwidth, high-latency links, especially when intermediate routers do not have properly tuned interface buffers. Improperly tuned router buffers cause TCP to perceive congestion in the network, resulting in unnecessarily dropped packets, even when the network can support high-throughput rates. MX-TCP is incompatible with AFE identification. A traffic flow cannot be classified as MX-TCP and then subsequently classified in a different queue. This re-classification can occur if there is a more exact match of the traffic using AFE identification. You must ensure the following when you enable MX-TCP: • The QoS rule for MX-TCP is at the top of QoS rules list. • The rule does not use AFE identification. • You only use MX-TCP for optimized traffic. MX-TCP does not work for unoptimized traffic. Use caution when specifying MX-TCP. The outbound rate for the optimized traffic in the configured QoS class immediately increases to the specified bandwidth, and does not decrease in the presence of network congestion. The Steelhead appliance always tries to transmit traffic at the specified rate. If no QoS mechanism (either parent classes on the Steelhead appliance, or another QoS mechanism in the WAN or WAN infrastructure) is in use to protect other traffic, that other traffic might be impacted by MX-TCP not backing off to fairly share bandwidth. When MX-TCP is configured as the queue parameter for a QoS class, the following parameters for that class are also affected: – **Link share weight**. The link share weight parameter has no effect on a QoS class configured with MX-TCP. – **Upper limit**. The upper limit parameter has no effect on a QoS class configured with MX-TCP.

Control	Description
	• **Packet-order** - Protects the TCP stream order by keeping track of flows that are currently inside the packet-shaping infrastructure. Packet-order protection allows only one packet from each flow into the HFSC traffic shaper at a time. The backlog for each flow stores the packets from the flow in order until the packet inside the HFSC infrastructure is dequeued for delivery to the network interface. This packet order priority protection works for both TCP and UDP streams. For best performance, select this queue with Citrix real-time latency priority traffic.
Add	Adds the QoS class.
Remove Selected	Select the check box next to the name and click **Remove Selected**. To remove a parent class, delete all rules for the corresponding child classes first. When a parent class has rules or children, the check box for the parent class is unavailable.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

Tip: The QoS classes appear in the QoS class table. To display QoS rules associated with the class, click the magnifying glass. To hide the rules associated with a QoS class, click the **close** icon.

Switching from Hierarchical QoS to Flat QoS

In certain situations, it might be necessary to switch from hierarchical QoS to flat QoS. For example, you might need to use link share weights, which are not supported in hierarchical QoS. Before changing from hierarchical QoS to flat QoS, you must delete all rules and classes in the hierarchy greater than one level deep.

Important: Changing the QoS enforcement mode while QoS is enabled can cause a momentary service disruption to traffic flowing through the Steelhead appliance.

To switch from hierarchical QoS to flat QoS

1. Start with a blank Basic QoS configuration. If necessary, migrate to Advanced QoS mode, return to Basic QoS mode and press **Clear**.

2. Migrate to Advanced QoS.

3. Choose Configure > Networking > Advanced QoS to display the Advanced QoS page.

4. Select all rules.

5. Click **Remove Selected**.

6. Select all child classes in the hierarchy greater than one level deep.

7. Click **Remove Selected**.

8. Under the default site, add a new class that is the child of the root class.

9. Change the default rule to use the new class you just added. You might need to adjust the minimum bandwidths on the existing classes.

10. Select all of the classes except the class you just added.

11. Click **Remove Selected**.

12. In the WAN Link section, select **Flat** mode.

13. Click **Apply**.

Adding a QoS Site or Rule (Advanced QoS)

Each rule maps a type of network traffic to a QoS class. You can create more than one QoS rule for a class. When more than one QoS rule is created for a class, the rules are followed in the order in which they are shown in the Advanced QoS page and only the first matching rule is applied to the class. Steelhead appliances support up to 2000 rules and up to 100 sites. When a port label is used to add a QoS rule, the range of ports cannot be more than 2000 ports.

In hierarchical QoS, only child classes can have rules.

Note: In RiOS v5.5 and earlier, the DSCP field in a QoS classification rule matches the DSCP value *before* DSCP marking rules are applied. In RiOS 6.0.x and v6.1.x, the DSCP field in a QoS classification rule matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value.

In RiOS v6.5 and later, the DSCP field in a QoS classification rule for pass-through traffic matches the DSCP value *before* DSCP marking rules are applied. The DSCP field in a QoS classification rule for optimized traffic matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value.

To add a QoS site or rule in Advanced QoS

1. Choose Configure > Networking > Advanced QoS to display the Advanced QoS page.

Figure 7-18. Advanced QoS Sites and Rules

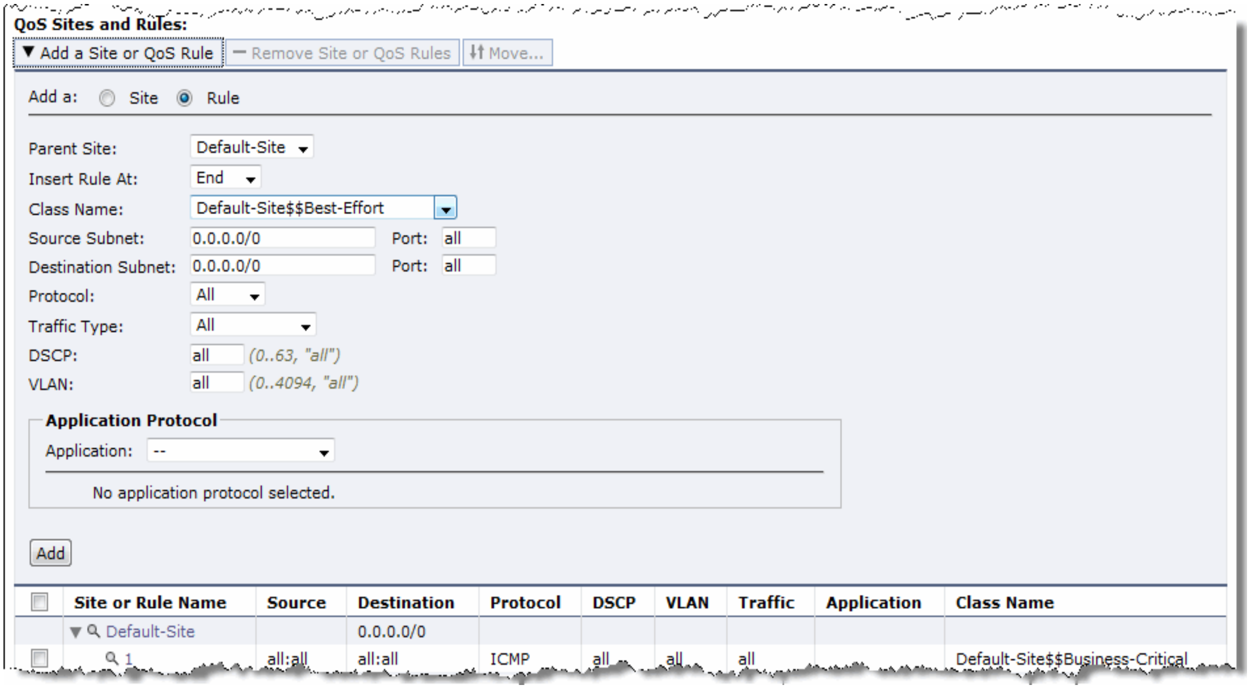

Add a: ○ Site ● Rule

Parent Site: Default-Site

Insert Rule At: End

Class Name: Default-Site$$Best-Effort

Source Subnet: 0.0.0.0/0 Port: all

Destination Subnet: 0.0.0.0/0 Port: all

Protocol: All

Traffic Type: All

DSCP: all (0..63, "all")

VLAN: all (0..4094, "all")

Application Protocol

Application: --

No application protocol selected.

□	Site or Rule Name	Source	Destination	Protocol	DSCP	VLAN	Traffic	Application	Class Name
	▼ Default-Site		0.0.0.0/0						
□	1	all:all	all:all	ICMP	all	all	all		Default-Site$$Business-Critical

2. Under QoS Sites and Rules, complete the configuration as described in the following table.

Control	Description
Add Site or QoS Rule	Displays the controls to add a QoS site or rule.
Add a	Select either Site or Rule. The default is rule.
Parent Site	Appears in hierarchical mode only. Select a parent site from the drop-down list. The default value is Default-site.
Insert Rule At	Inserts a QoS rule for a QoS class. Select Start, End, or a rule number from the drop-down list. Steelhead appliances evaluate rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted.
Class Name	Select a class name from the drop-down list. If the rule matches, the specified rule sends the packet to this class.
Source Subnet	Specify the IP address for the source network. Use the following format: XXX.XXX.XXX.XXX/XX
Port	Specify the port or port label for the source subnet. The default value is All. **Tip:** Rules support port labels for source and destination ports.

Control	Description
Destination Subnet	Specify the IP address for the destination network. Use the following format: XXX.XXX.XXX.XXX/XX
Port	Specify the port or port label for the destination subnet. The default value is All. **Tip:** Rules support port labels for source and destination ports.
Protocol	Select All, TCP, GRE, UDP, ICMP, or IPSec from the drop-down list. All specifies all TCP and UDP-based protocols.
Traffic Type	Select All, Optimized, or Pass-Through from the drop-down list. The system applies the QoS rules to optimized and pass-through (egress only) traffic. **Note:** Session reliability (port 2598) is not supported with pass-through Citrix traffic.
DSCP	Optionally, select a DSCP level from the drop-down list. **Note:** In RiOS v5.5 and earlier, the DSCP field in a QoS classification rule matches the DSCP value *before* DSCP marking rules are applied. In RiOS 6.0.x and v6.1.x, the DSCP field in a QoS classification rule matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value. In RiOS v6.5, the DSCP field in a QoS classification rule for pass-through traffic matches the DSCP value *before* DSCP marking rules are applied. The DSCP field in a QoS classification rule for optimized traffic matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value.
VLAN	Optionally, specify the VLAN tag for the rule.
Application Protocols	Select an application from the drop-down list of global applications. You can define and add any applications that do not appear in the list. Selecting HTTP expands the control to include the Domain Name and Relative Path controls. Enter the domain name and relative path. The relative path is the part of the URL that follows the domain name. Selecting ICA expands the control to include priorities 0 - 3. Select a priority for the Citrix application to separate low-priority traffic (such as print jobs), from high-priority traffic (such as interactive screen updates). Citrix classification using a priority supports optimized and pass-through traffic. You must select the packet-order queue when using ICA priorities.
Add	Adds a rule or site to the QoS rule or site list.
Remove Site or QoS Rules	Removes the selected sites or rules.
Move Site or QoS Rules	Select the box next to the name and click **Move QoS Rules**. Click the arrow next to the desired rule position. The rule or sites moves to the new position.

3. Click **Apply** to apply your settings.

4. Click **Save** to save your settings permanently.

Tip: To display the QoS rules associated with a site, click the magnifying glass in the QoS Site table. To hide the rules associated with a QoS site, click the **close** icon.

Tip: To modify a QoS rule, delete it and add a new rule.

Verifying and Saving a QoS Configuration

After you apply your settings, you can verify whether the traffic is categorized in the correct class by choosing Reports > Networking > QoS Stats Sent and viewing the report. For example, if you have configured VoIP traffic as real-time, check the Real-time class and verify that the other classes are not receiving VoIP traffic.

You can verify whether the configuration is honoring the bandwidth allocations by reviewing the QoS Stats Sent and QoS Stats Dropped reports.

When you have verified appropriate changes, you can write the active configuration that is stored in memory to the active configuration file (or you can save it as any filename you choose). For details on saving configurations, see "Managing Configuration Files" on page 358.

Related Topics

- "Configuring Port Labels" on page 94
- "Managing Configuration Files" on page 358
- "Viewing QoS Stats Dropped Reports" on page 405
- "Viewing QoS Stats Sent Reports" on page 407

Modifying QoS Classes

You can modify QoS classes in the Advanced QoS page.

To modify a QoS class

1. Choose Configure > Networking > QoS Classification to display the Advanced QoS page.

2. Select the class name in the list to display the Editing QoS Group page.

Figure 7-19. Editing QoS Class Page

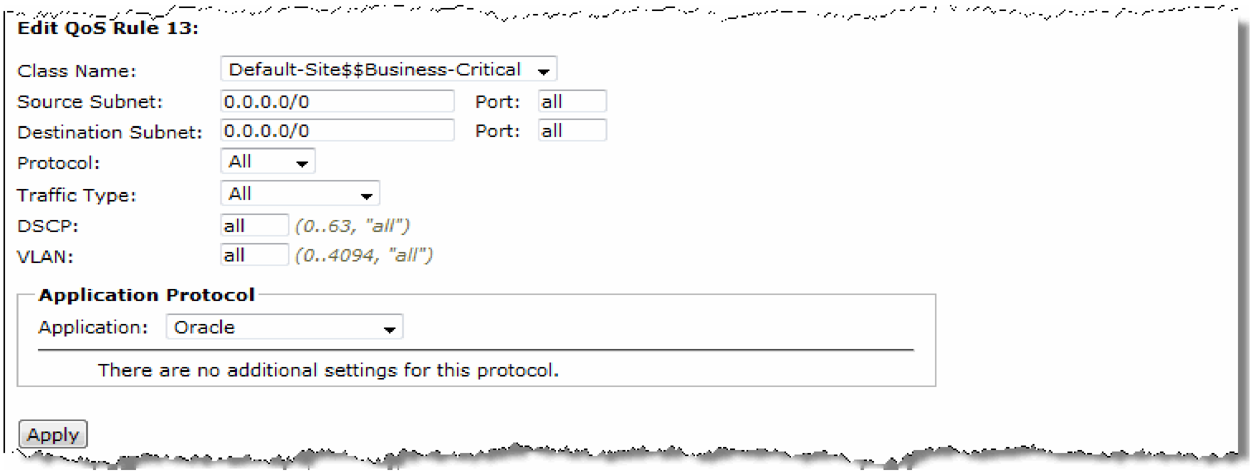

Edit QoS Rule 13:

Class Name: Default-Site$$Business-Critical

Source Subnet: 0.0.0.0/0 Port: all

Destination Subnet: 0.0.0.0/0 Port: all

Protocol: All

Traffic Type: All

DSCP: all (0..63, "all")

VLAN: all (0..4094, "all")

Application Protocol

Application: Oracle

There are no additional settings for this protocol.

3. Under Editing QoS Class, modify the settings.

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

Clearing an Advanced QoS Configuration to Return to Basic QoS

In certain situations, it might be necessary to revert from advanced to basic QoS. You can either revert to a saved configuration or start over with a blank basic QoS configuration.

You can only revert to a previous basic QoS configuration if you backed up your configuration before you migrated to advanced QoS. Reverting to a previously saved configuration restores your entire Steelhead configuration.

Reverting to basic QoS without using a previous configuration deletes all your current QoS settings.

To clear an advanced QoS configuration and return to a blank basic QoS configuration

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Click **Clear**.

Figure 7-20. Basic QoS - Clear Confirmation

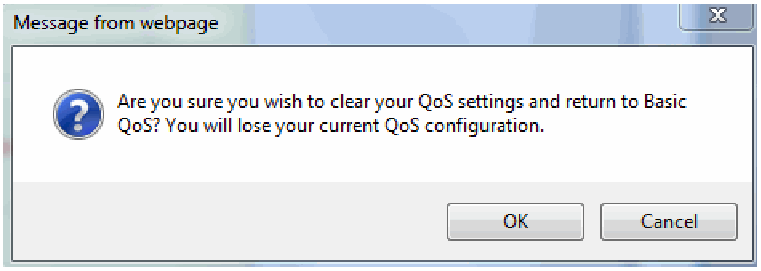

Message from webpage

Are you sure you wish to clear your QoS settings and return to Basic QoS? You will lose your current QoS configuration.

3. Click **OK**.

The process takes approximately two minutes but can take longer depending on the existing configuration. When the system returns to basic QoS, the Basic QoS page appears.

You can now configure basic QoS. For details, see "Configuring Basic QoS" on page 275.

4. Click **Save** to save your settings permanently.

To revert from advanced QoS to a previously saved basic QoS configuration

1. Choose Configure > Networking > Basic QoS to display the Basic QoS page.

2. Click **revert to a prior system configuration** to display the Configurations page.

3. Under Change Active Configuration, select the previous configuration for basic QoS from the drop-down list.

4. Click **Activate**.

   Reverting takes approximately two minutes but can take longer depending on the configuration.

5. Click **Restart** to restart the optimization service.

6. Click **Save** to save your settings permanently.

Enabling MX-TCP Queue Policies (Advanced QoS only)

When you define a QoS class, you can enable an MX-TCP queue policy, which prioritizes TCP/IP traffic to provide more throughput for high loss links or links that have large bandwidth and high latency LFNs. For example:

- **Data-Intensive Applications** - Many large, data-intensive applications running across the WAN can negatively impact performance due to latency, packet loss, and jitter. MX-TCP enables you to maximize your TCP throughput for data intensive applications.
- **High Loss Links** - TCP does not work well on misconfigured links (for example, an under-sized bottleneck queue) or links with even a small amount of loss, which leads to link under-utilization. If you have dedicated point-to-point links and want those links to function at predefined rates, configure the Steelhead appliance to prioritize TCP traffic.
- **Privately Owned Links** - If your network includes privately-owned links dedicated to rate-based TCP, configure the Steelhead appliance to prioritize TCP traffic.

After enabling the MX-TCP queue to forward TCP traffic regardless of congestion or packet loss, you can assign QoS rules that incorporate this policy only to links where TCP is of exclusive importance.

The following exceptions to QoS classes apply to MX-TCP queues:

- The Link Share Weight and Upper BW limit parameters do not apply to MX-TCP queues.
- MX-TCP queues apply only to optimized traffic (that is, no pass-through traffic).
- MX-TCP queues cannot be configured to contain more bandwidth than the license limit.

MX-TCP is incompatible with the AppFlow Engine. A traffic flow cannot be classified as MX-TCP and then subsequently classified in a different queue. This re-classification can happen if there is a more exact match of the traffic.

When enabling MX-TCP, ensure that:

- the QoS rule is at the top of QoS rules list.
- the rule does not use AppFlow Engine identification.

Basic Steps for MX-TCP

The following table describes the basic steps to configure MX-TCP, followed by detailed procedures. Enabling this feature is *optional*.

Task	Reference
1. Select either Flat or Hierarchical mode. **Note:** Selecting a mode does *not* enable QoS traffic classification. The Enable QoS Classification and Enforcement check box must be selected and a bandwidth link rate must be set for each WAN interface before traffic optimization begins.	"Selecting a QoS Enforcement System" on page 270 "To enable basic QoS" on page 277
2. Select each WAN interface and define the bandwidth link rate for each interface.	"To enable basic QoS" on page 277
3. Add an MX-TCP class for the traffic flow. Make sure you specify MX-TCP as your queue.	"To add a QoS class" on page 289
4. Define QoS rules to point to the MX-TCP class.	"Adding a QoS Site or Rule (Advanced QoS)" on page 294
5. Select the Enable QoS Classification and Enforcement check box. Your changes take effect immediately.	"To enable basic QoS" on page 277

Task	Reference
6. Optionally, to test a single connection, change the WAN socket buffer size (to at least the BDP). You must set this parameter on both the client-side and the server-side Steelhead appliance.	"Configuring Buffer Settings" on page 92
7. Check and locate the inner connection.	"Viewing Alarm Status Reports" on page 459
8. Check the throughput.	"Viewing Current Connections" on page 385

Configuring QoS Marking

You set QoS marking in the Configure > Networking > QoS Marking page.

This section describes how to use Steelhead appliance QoS marking when integrating Steelhead appliances into an existing QoS architecture. Steelhead appliances can retain or alter the DSCP or IP ToS value of both pass-through traffic and optimized traffic. To alter the DSCP or IP ToS value of optimized or pass-through traffic, you create a list that maps which traffic receives a certain DSCP value. The first matching mapping is applied.

This section includes the following topics:

- "QoS Marking Default Setting" on page 300
- "Setting ToS DSCP Values" on page 301
- "Creating a QoS Map List" on page 303

For details on QoS marking, see the *Riverbed Deployment Guide*.

Note: RiOS does not support QoS Marking using AppFlow Engine identification.

QoS Marking Default Setting

By default, Steelhead appliances reflect the DSCP or IP ToS value found on pass-through traffic and optimized connections. This means that the DSCP or IP ToS value on pass-through traffic is unchanged when it passes through the Steelhead appliance.

After you map a source-destination-port pattern and a DSCP level, every packet corresponding to the connection with that destination port has the DSCP field set to that value in the forward and backward direction. On the WAN side of the Steelhead appliance, you configure a network router or a traffic shaper to prioritize packets according to the value in the DSCP field before they are sent across the WAN.

Enabling these features is *optional*.

QoS Marking with the FTP Data Channel

The method you use to configure QoS for FTP depends on whether the data channel is using *active* or *passive* FTP traffic. For details on active versus passive FTP, see "QoS Classification for the FTP Data Channel" on page 271.

Active FTP Marking

The procedure you use to configure QoS for active FTP traffic depends on the RiOS version:

- **RiOS versions prior to v5.0.7 and v5.5.2** - Configure a QoS map on the server-side Steelhead appliance to match the *destination* port 20. This might seem counter-intuitive as active FTP uses source port 20 and not destination port 20. This is because QoS marking does not support the creation of QoS maps based on the source port for optimized traffic in RiOS releases prior to 5.0.7 and 5.5.2. It is not necessary to create a QoS rule on the client-side Steelhead appliance because the default behavior is to automatically reflect the DSCP value.
- **RiOS versions v5.0.7, v5.5.2, and later** - For active FTP, configure a QoS map on the server-side Steelhead appliance to match the *source* port 20. It is not necessary to create a QoS map on the client-side Steelhead appliance because the default behavior is to automatically reflect the DSCP value.

For details, see "To add an optimized QoS map" on page 303.

Passive FTP Marking

For passive FTP, specify destination port 20 on the client-side Steelhead appliance when adding an optimized QoS map. This might seem counter-intuitive as passive FTP does not use destination port 20, but rather some random port number. However, the Steelhead appliance has specific intelligence built-in so that it knows which port number passive FTP is using as its destination port number. Consequently, for QoS marking with passive FTP, destination port 20 on the client-side Steelhead appliance simply means the port number being used by the data channel for passive FTP, as opposed to the literal meaning of destination port 20. You do not need to create an optimized QoS map on the server-side Steelhead appliance because the default behavior is to automatically reflect the DSCP value.

Note: The Steelhead appliance must intercept the FTP control channel (port 21), regardless of whether FTP is active or passive.

Setting ToS DSCP Values

The ToS DSCP level corresponds to the DiffServ DSCP field in the IP packets header. The ToS precedence values (0 to 7) use the upper three bits of the DiffServ field; DSCP values (0 to 63) use the upper six bits.

To set a ToS DSCP value

1. Choose Configure > Networking > QoS Marking to display the QoS Marking page.

Figure 7-21. QoS Marking Page

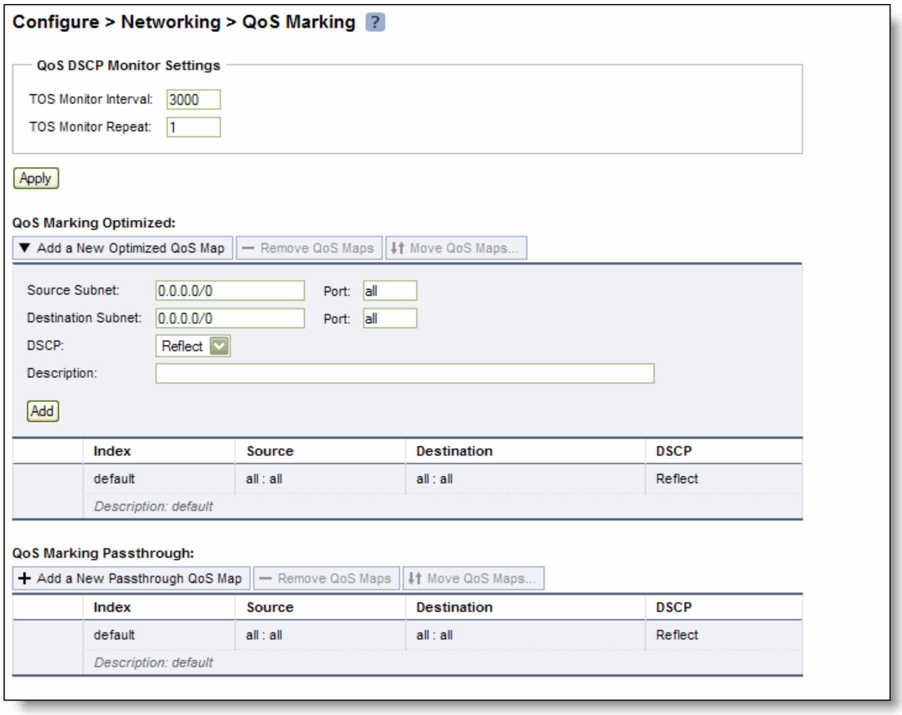

Configure > Networking > QoS Marking

QoS DSCP Monitor Settings

TOS Monitor Interval: 3000

TOS Monitor Repeat: 1

QoS Marking Optimized:

Source Subnet: 0.0.0.0/0 Port: all

Destination Subnet: 0.0.0.0/0 Port: all

DSCP: Reflect

Description:

	Index	Source	Destination	DSCP
	default	all : all	all : all	Reflect
	Description: default			

QoS Marking Passthrough:

	Index	Source	Destination	DSCP
	default	all : all	all : all	Reflect
	Description: default			

2. Under QoS DSCP Monitor Settings, complete the configuration as described in the following table.

Control	Description
TOS Monitor Interval	Specify how many TCP bytes the client Steelhead appliance receives on the upstream connection before sending packets that reflect the same DSCP value. The default value is 3000. For example, after the TCP connection has received 3000 bytes of data, the Steelhead appliance checks the DSCP value received in the last packet for that connection and uses that value to mark packets on the next hop. The DSCP value in packets received from the server is used in packets sent from the server-side Steelhead appliance to the client-side Steelhead appliance. As soon as the server sends data back, the DSCP value is sent for packets in the reverse direction. This also applies to packets sent from a server-side Steelhead appliance to the server. If you set the interval to 1, the connection setup packets (SYN/SYN-ACK/ACK) are not marked, but the next packets are marked, because the server-side Steelhead appliance sends data to the server only after it receives data from the client-side Steelhead appliance.
TOS Monitor Repeat	Specify how often the client-side Steelhead appliance rechecks the DSCP value of the traffic. The default value is 1. Change this value when you expect the DSCP value to change during the duration of the connection and you want to use the most recent value. If you want to check indefinitely, set the repeat interval to negative 1 (-1).

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

Creating a QoS Map List

You can create separate map lists for optimized and pass-through traffic.

Note: Optimized traffic is marked in both directions, but pass-through traffic is marked only on the egress traffic.

Note: Only the first matching mapping is applied.

To add an optimized QoS map

1. Under QoS Marking Optimized, complete the configuration as described in the following table.

Control	Description
Add a New Optimized QoS Map	Displays the controls to add an optimized QoS map.
Source Subnet	Specify the IP address for the source network. Use the following format: XXX.XXX.XXX.XXX/XX
Source Port	Specify the source port number, port label, or all. A port label is a label that you assign to a set of ports so that you can reduce the number of configuration rules in your system. For the MAPI data channel, specify port 7830 and the corresponding DSCP level. The method you use to configure QoS for active FTP depends on the RiOS version: • **RiOS versions v5.0.7 and v5.5.2** - For the FTP data channel, specify source port 20 and the corresponding DSCP level on the Steelhead appliance closest to the FTP server (assuming the FTP server initiates the data channel on port 20). Setting QoS for port 20 on the server-side Steelhead appliance affects *active* FTP. • **RiOS versions prior to v5.0.7 and v5.5.2** - For the FTP data channel, configure a QoS map on the server-side Steelhead appliance to match the *destination* port 20, because RiOS versions prior to v5.0.7 and v5.5.2 do not support the creation of QoS maps based on the source port for optimized traffic.
Destination Subnet	Specify the IP address for the destination subnet. Use the following format: XXX.XXX.XXX.XXX/XX
Destination Port	Specify the destination port number, port label, or all. A port label is a label that you assign to a set of ports so that you can reduce the number of configuration rules in your system. For the MAPI data channel, specify port 7830 and the corresponding DSCP level. For the FTP data channel, specify destination port 20 and the corresponding DSCP level. Setting QoS for port 20 on the server-side Steelhead appliance affects *passive* FTP.

Control	Description
DSCP	Optionally, select a DSCP level (0-63) or Reflect (the default setting) from the drop-down list. Reflect specifies that the DSCP level or IP ToS value found on optimized traffic is unchanged when it passes through the Steelhead appliance. **Important:** If your connections already have a DSCP level and you do not define one on the client-side Steelhead appliance, the Steelhead appliance uses the existing DSCP level for the connection between the Steelhead appliances. If you define a DSCP level on the client-side Steelhead appliance, the Steelhead appliance overrides the existing DSCP level and the value that you defined is applied. **Note:** Optimized traffic is marked in both directions, but pass-through traffic is marked only on the egress traffic. **Note:** In RiOS v5.5 and earlier, the DSCP field in a QoS classification rule matches the DSCP value *before* DSCP marking rules are applied. In RiOS 6.0.x and v6.1.x, the DSCP field in a QoS classification rule matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value. In RiOS v6.5, the DSCP field in a QoS classification rule for pass-through traffic matches the DSCP value *before* DSCP marking rules are applied. The DSCP field in a QoS classification rule for optimized traffic matches the DSCP value *after* DSCP marking rules are applied; that is, it matches the post-marking DSCP value.
Description	Optionally, specify a description to identify the rule.
Add	Adds the rule to the optimized QoS map list.
Remove QoS Maps	Removes the selected map configurations.
Move QoS Maps	Reorders the selected maps in the list.

2. Click **Save** to save your settings permanently.

To add a pass-through map

1. Under QoS Marking Passthrough, complete the configuration as described in the following table.

Control	Description
Add a New Passthrough QoS Map	Displays the controls to add a pass-through QoS map.
Source Subnet	Specify the IP address for the source network. Use the following format: XXX.XXX.XXX.XXX/XX
Source Port	Specify the source port number, port label, or all. A port label is a label that you assign to a set of ports so that you can reduce the number of configuration rules in your system. For the MAPI data channel, specify port 7830 and the corresponding DSCP level. You cannot optimize a pass-through FTP data channel connection.
Destination Subnet	Specify the IP address for the destination subnet. Use the following format: XXX.XXX.XXX.XXX/XX
Destination Port	Specify the destination port number, port label, or all. A port label is a label that you assign to a set of ports so that you can reduce the number of configuration rules in your system. For the MAPI data channel, specify port 7830 and the corresponding DSCP level. You cannot optimize a pass-through FTP data channel connection.

Control	Description
DSCP	Optionally, select a DSCP level (**0-63**) or **Reflect** (the default setting) from the drop-down list. Reflect specifies that the DSCP level or IP ToS value found on pass-through traffic is unchanged when it passes through the Steelhead appliance. **Important:** If your connections already have a DSCP level and you do not define one in the Management Console, the Steelhead appliance uses the existing DSCP level for the connection between the Steelhead appliances. If you define a DSCP level in the Management Console, the Steelhead appliance overrides the existing DSCP level and the value that you defined is applied. **Note:** Optimized traffic is marked in both directions, but pass-through traffic is marked only on the egress traffic.
Description	Optionally, specify a description to help you identify the map.
Add	Adds the map to the pass-through QoS map list.
Remove QoS Maps	Removes the selected map configurations.
Move QoS Maps	Reorders the selected maps in the list.

2. Click **Save** to save your settings permanently.

Related Topics

- "Configuring Port Labels" on page 94
- "Viewing QoS Stats Dropped Reports" on page 405
- "Viewing QoS Stats Sent Reports" on page 407

Joining a Windows Domain or Workgroup

A server-side Steelhead appliance can join a Windows domain or workgroup in the Configure > Networking > Windows Domain page. This page provides a central place for a Steelhead appliance to join a Windows Domain or workgroup for the following RiOS features:

- SMB signing delegation trust for CIFS optimizations. For details, see "Configuring SMB Signing" on page 103.
- MAPI 2007 encrypted traffic optimization authentication. For details, see "Configuring MAPI Optimization" on page 123.

Workgroup mode is provided for PFS and does not support these features. For details, see "Configuring PFS" on page 161.

Domain and Local Workgroup Settings

You can choose between two user authentication modes: domain or local workgroup. Creating a local workgroup eliminates the need to join a Windows domain and simplifies the configuration process, but a workgroup does *not* support SMB signing or MAPI 2007 encrypted traffic optimization authentication.

Domain Mode

In Domain mode, you configure the Steelhead appliance to join a Windows domain (typically, the domain of your company). When you configure the Steelhead appliance to join a Windows domain, you do not have to manage local accounts in the branch office, as you do in Local Workgroup mode.

Domain mode allows a Domain Controller (DC) to authenticate users accessing its file shares. The DC can be located at the remote site or over the WAN at the main data center. The Steelhead appliance must be configured as a Member Server in the Windows 2000, or later, Active Directory Services (ADS) domain. Domain users are allowed to access the PFS shares, use the Kerberos delegation trust facility and/or NTLM environments for MAPI 2007 encryption or SMB signing, based on the access permission settings provided for each user.

Data volumes at the data center are configured explicitly on the proxy-file server and are served locally by the Steelhead appliance. As part of the configuration, the data volume and ACLs from the origin-file server are copied to the Steelhead appliance. RiOS allocates a portion of the Steelhead appliance datastore for users to access as a network file system.

Before enabling Domain mode make sure you:

- configure the DNS server correctly. The configured DNS server must be the same DNS server to which all the Windows client computers point. To use SMB signing, the server-side Steelhead appliance must be in DNS. For details, see “To specify DNS settings” on page 40.
- have a fully-qualified domain name. This domain name must be the domain name for which all the Windows desktop computers are configured.
- set the owner of all files and folders in all remote paths to a domain account and not a local account.

Using PFS in Domain Mode

PFS does not support local user and group accounts. These accounts reside only on the host where they are created. During an initial copy from the origin file server to the PFS Steelhead appliance, if PFS encounters a file or folder with permissions for both domain and local accounts, the Steelhead appliance preserves only the domain account permissions. If your DC is across the WAN, in the event of a WAN outage, you cannot perform user authentication. To prevent this, you either need a local DC (perhaps running in RSP), or you can switch to Local Workgroup mode, which requires you to configure local usernames and passwords or use shares that are open to everyone. For details, see “Local Workgroup Mode” on page 306.

Regarding the user account required to join the Steelhead to the domain:

- This account does *not* need to be a domain admin account. Any account that has sufficient privileges to join a machine to Active Directory works (that is; if you have created a non-domain Admin account that has permission to add machine accounts, and it works for regular Windows computers).
- Regardless of what account is entered, RiOS does *not* store the account information on the Steelhead appliance. RiOS uses it for a one-time attempt to join the domain.
- If you ever need to rejoin the computer (for example, if the account was deleted from the Active Directory), you need to re-enter your credentials.

For details on how the ACLS are propagated from the origin-file server to a PFS share, refer to the Riverbed Support site at https://support.riverbed.com.

Local Workgroup Mode

In Local Workgroup mode, you define a workgroup and add individual users that have access to the Steelhead appliance. The Steelhead appliance does not join a Windows domain.

Use Local Workgroup mode in environments where you do not want the Steelhead appliance to be a part of a Windows domain. Creating a workgroup eliminates the need to join a Windows domain and simplifies the configuration process.

Note: If you use Local Workgroup mode you must manage the accounts and permissions for the branch office on the Steelhead appliance. The Local Workgroup account permissions might not match the permissions on the origin-file server.

To configure a Windows domain in Local Workgroup mode

1. Select Configure > Networking > Windows Domain to display the Windows Domain page.

Figure 7-22. Windows Domain Page

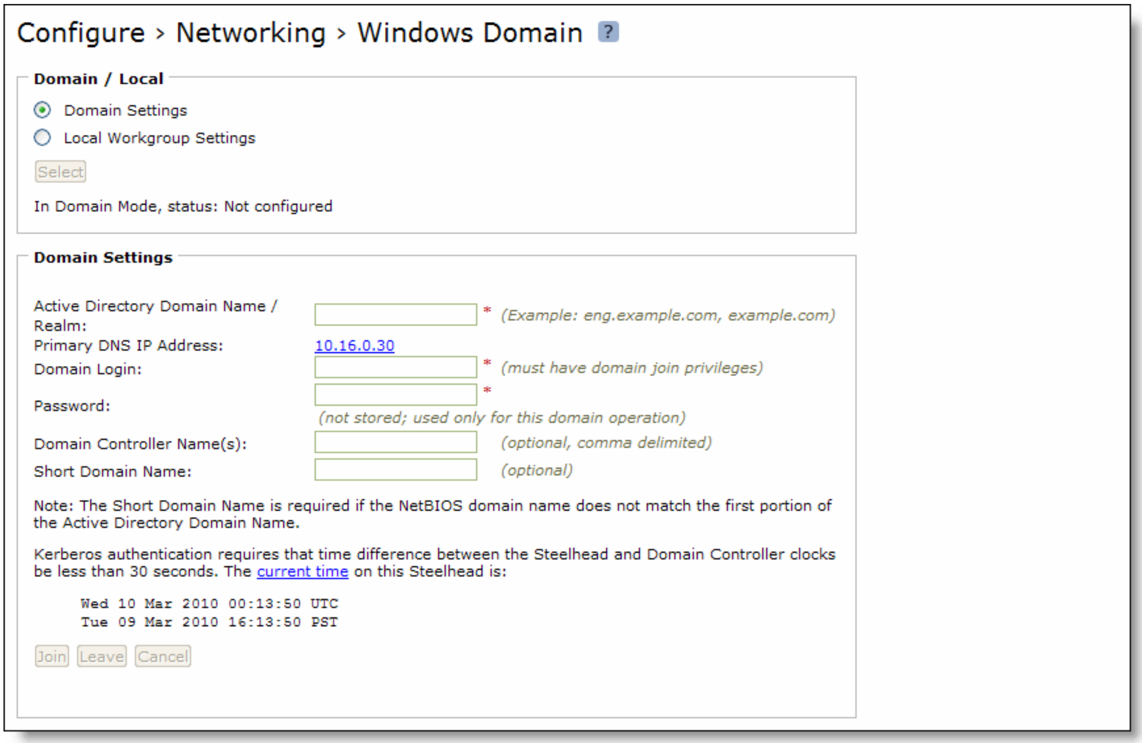

Configure > Networking > Windows Domain

Domain / Local

● Domain Settings

○ Local Workgroup Settings

Select

In Domain Mode, status: Not configured

Domain Settings

Active Directory Domain Name / Realm:	*	(Example: eng.example.com, example.com)
Primary DNS IP Address:	10.16.0.30	
Domain Login:	*	(must have domain join privileges)
Password:	* (not stored; used only for this domain operation)	
Domain Controller Name(s):		(optional, comma delimited)
Short Domain Name:		(optional)

Note: The Short Domain Name is required if the NetBIOS domain name does not match the first portion of the Active Directory Domain Name.

Kerberos authentication requires that time difference between the Steelhead and Domain Controller clocks be less than 30 seconds. The current time on this Steelhead is:

```
Wed 10 Mar 2010 00:13:50 UTC
Tue 09 Mar 2010 16:13:50 PST
```

2. Under Domain/Local Workgroup Settings, select Local Workgroup Settings, click **Select**, and then click **OK** when a dialog asks if you really want to change the setting.

3. Complete the configuration as described in the following table.

Control	Description
Workgroup Name	Specify a local workgroup name. If you configure in Local Workgroup mode the Steelhead appliance does not need to join a domain. Local Workgroup accounts are used by clients when they connect to the Steelhead appliance. **Note:** PFS, MAPI 2007, or SMB signing must be enabled and Local Workgroup Settings must be selected before you can set the Workgroup Name. After you have set a Workgroup Name, click **Join**.
Add a New User	Displays the controls to add a new user to the local workgroup.

Control	Description
User	Specify the login to create a local workgroup account so that users can connect to the Steelhead appliance.
Password/Password Confirm	Specify and confirm the user account password.
Add	Adds users to the local workgroup.
Remove Selected	Removes the selected names.

4. Click **Apply** to apply your settings to the running configuration.

5. Click **Save** to save your settings permanently.

To configure a Windows domain in Domain mode

1. Select Configure > Networking > Windows Domain to display the Windows Domain page.

2. Under Domain/Local Workgroup Settings, click **Domain Settings**, click **Select**, and then click **OK** when a dialog asks if you really want to change the setting.

3. Complete the configuration as described in the following table.

Control	Description
Active Directory Domain Name/Realm	Specify the domain in which to make the Steelhead appliance a member. Typically, this is your company domain name. RiOS v5.5 and later supports Windows 2000 or later domains. RiOS does not support non-domain accounts other than administrator accounts. If you create Local mode shares on a non-administrator account, your security permissions for the share are not preserved on the origin-file server.
Primary DNS IP Address	By default, this field displays the primary DNS IP set in the DNS Settings page. To modify this entry, click the IP address.
Domain Login	Specify the login name, which must have domain join privileges. Domain administrator credentials are not strictly required.
Password	Specify the password. This control is case-sensitive.
Domain Controller Name(s)	Optionally, specify the hosts that provide user login service in the domain. (Typically, with Windows 2000 Active Directory Service domains, given a domain name, the system automatically retrieves the DC name.) This control is case-sensitive. **Note:** Riverbed recommends specifying the domain controller name in high latency situations, as it reduces the time to join the domain significantly.
Short Domain Name	Specify the short domain (NetBIOS) name if it does not match the first portion of the Active Directory Domain name. Case matters; NBTTECH is not the same as nbttech.
Join/Leave	Joins the domain or leaves the domain. **Important:** If you are in Domain mode and have joined a domain, you cannot change to Local Workgroup mode until you leave the domain.
Rejoin	Rejoins the domain.
Cancel	Cancels any current domain action that is in progress, such as joining or leaving a domain.

4. Click **Apply** to apply your settings to the running configuration.

5. Click **Save** to save your settings permanently.

   When you have successfully joined the domain, the status updates to **In a Domain**.

Troubleshooting a Domain Join Failure

This section describes common problems that can occur when joining a Windows domain.

System Time Mismatch

The number one cause of failing to join a domain is a significant difference in the system time on the Windows domain controller and the Steelhead appliance. When the time on the domain controller and the Steelhead appliance do not match, the following error message appears:

```
lt-kinit: krb5_get_init_creds: Clock skew too great
```

Riverbed recommends using NTP time synchronization to synchronize the client and server clocks. It is critical that the Steelhead appliance time is the same as on the Active Directory controller. Sometimes an NTP server is down or inaccessible, in which case there can be a time difference. You can also disable NTP if it is not being used and manually set the time. You must also verify that the time zone is correct. For details, see "Modifying General Host Settings" on page 39.

Tip: Select the Primary DNS IP address to view the Configure > Networking > Host Settings page.

Invalid Domain Controller IP

A domain join can fail when the DNS server returns an invalid IP address for the Domain Controller. When a DNS misconfiguration occurs during an attempt to join a domain, the following error messages appear:

```
Failed to join domain: failed to find DC for domain <domain name>
Failed to join domain : No Logon Servers
```

Additionally, the Domain Join alarm triggers and messages similar to the following appear in the logs:

```
Oct 13 14:47:06 bravo-sh81 rcud[10014]: [rcud/main/.ERR] - {- -} Lookup for bravo-
sh81.GEN-VCS78DOM.COM Failed
Oct 13 14:47:06 bravo-sh81 rcud[10014]: [rcud/main/.ERR] - {- -} Failed to join
domain: failed to find DC for domain GEN-VCS78DOM.COM
```

When you encounter this error, choose Configure > Networking > Host Settings and verify that the DNS settings are correct.

Related Topics

- "Configuring SMB Signing" on page 103
- "Configuring MAPI Optimization" on page 123
- "Modifying General Host Settings" on page 39
- "Configuring PFS" on page 161

Configuring Simplified Routing Features

You can enable simplified routing in the Configure > Networking > Simplified Routing page.

Simplified routing collects the IP address for the next hop MAC address from each packet it receives to address traffic. With simplified routing, you can use either the WAN or LAN-side device as a default gateway. The Steelhead appliance learns the right gateway to use by watching where the switch or router sends the traffic, and associating the next-hop Ethernet addresses with IP addresses. Enabling simplified routing eliminates the need to add static routes when the Steelhead appliance is in a different subnet from the client and the server.

Without simplified routing, if a Steelhead appliance is installed in a different subnet from the client or server, you must define one router as the default gateway and static routes for the other routers so that traffic is not redirected back through the Steelhead appliance. In some cases, even with the static routes defined, the ACL on the default gateway can still drop traffic that should have gone through the other router. Enabling simplified routing eliminates this issue.

Simplified routing has the following constraints:

- WCCP cannot be enabled.
- The default route must exist on each Steelhead appliance in your network.

Tip: For detailed configuration information, see the *Riverbed Deployment Guide*.

CSH The AWS Cloud Steelhead does not support simplified routing. The ESX Cloud Steelhead deployed in-path with the Discovery Agent (not with WCCP or PBR) supports simplified routing.

To enable simplified routing

1. Choose Configure > Networking > Simplified Routing to display the Simplified Routing page.

Figure 7-23. Simplified Routing Page

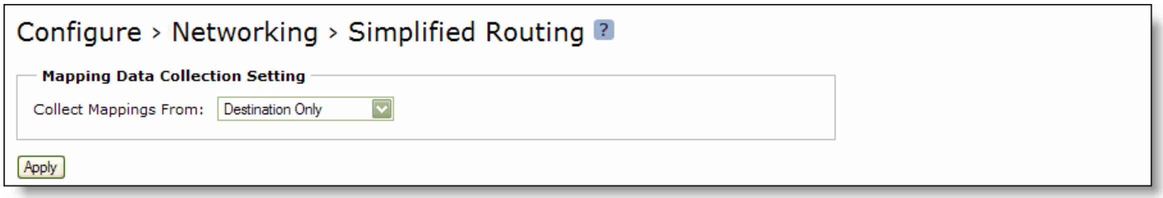

Configure > Networking > Simplified Routing ?

Mapping Data Collection Setting

Collect Mappings From: Destination Only

Apply

2. Under Mapping Data Collection Setting, complete the configuration as described in the following table.

Control	Description
Collect Mappings From	Select one of the following options from the drop-down list: • **None** - Do not collect mappings. • **Destination Only** - Collects destination MAC data. Use this option in connection forwarding deployments. This is the default setting. • **Destination and Source** - Collect mappings from destination and source MAC data. Use this option in connection forwarding deployments. • **All** - Collect mappings for destination, source, and inner MAC data. Also collect data for connections that are *un-NATted* (that is, connections that are not translated using NAT).

3. Click **Apply** to save your settings to the running configuration.

4. Click **Save** to save your settings permanently.

Related Topics

- "About In-Path Rules" on page 27
- "Configuring Connection Forwarding Features" on page 255

Configuring WCCP

You can enable WCCP service groups in the Configure > Networking > WCCP page.

WCCP enables you to redirect traffic that is not in the direct physical path between the client and the server. To enable WCCP, the Steelhead appliance must join a service group at the router. A service group is a group of routers and Steelhead appliances which define the traffic to redirect, and the routers and Steelhead appliances the traffic goes through. You might use one or more service groups to redirect traffic to the Steelheads for optimization.

RiOS v6.1 and later provides additional WCCP configuration, allowing each individual Steelhead appliance in-path interface to be configured as a WCCP client. Each configured in-path interface participates in WCCP service groups as an individual WCCP client, providing flexibility to determine load balancing proportions and redundancy.

In RiOS v6.5 and later, you must enable connection forwarding in a WCCP cluster. A WCCP cluster refers to two or more Steelhead appliances participating in the same service group. By default, RiOS provides load balancing across all participating Steelhead appliances in a WCCP cluster. With connection forwarding enabled, the WCCP load balancing algorithm considers the total number of in-path interfaces of all neighbors in the service group when balancing the traffic load across the interfaces. If you do not enable connection forwarding, the Steelhead with the lowest IP address assigns all traffic flows to itself.

Enabling WCCP is *optional*.

For details on balancing traffic loads in WCCP, see the *Riverbed Deployment Guide*.

Tip: You can also use the CLI to configure WCCP service groups. For detailed configuration information (including configuring the WCCP router), see the *Riverbed Deployment Guide*.

CSH The AWS Cloud Steelhead does not support L4/PBR/WCCP configuration. The ESX Cloud Steelhead supports it.

To enable a WCCP service group

Important: Before configuring your WCCP service group, you must enable L4/PBR/WCCP support in the General Service Settings page. For details, see "Configuring General Service Settings" on page 60.

1. Choose Configure > Networking > WCCP to display the WCCP page.

Figure 7-24. WCCP Page

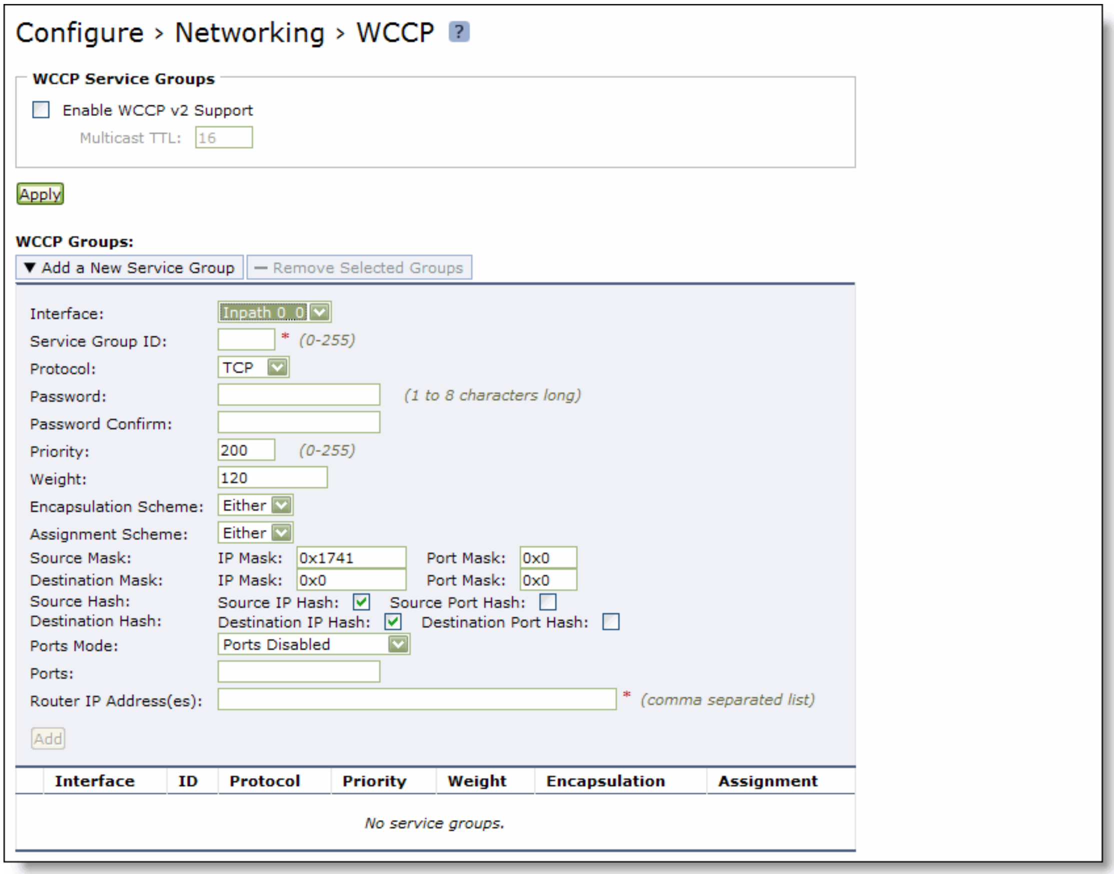

Configure > Networking > WCCP ?

WCCP Service Groups

□ Enable WCCP v2 Support

Multicast TTL: 16

Apply

WCCP Groups:

▼ Add a New Service Group — Remove Selected Groups

Interface:	Inpath 0_0	
Service Group ID:		* (0-255)
Protocol:	TCP	
Password:		(1 to 8 characters long)
Password Confirm:		
Priority:	200	(0-255)
Weight:	120	
Encapsulation Scheme:	Either	
Assignment Scheme:	Either	
Source Mask:	IP Mask: 0x1741	Port Mask: 0x0
Destination Mask:	IP Mask: 0x0	Port Mask: 0x0
Source Hash:	Source IP Hash: ☑	Source Port Hash: □
Destination Hash:	Destination IP Hash: ☑	Destination Port Hash: □
Ports Mode:	Ports Disabled	
Ports:		
Router IP Address(es):		* (comma separated list)

Add

Interface	ID	Protocol	Priority	Weight	Encapsulation	Assignment
No service groups.						

2. Under WCCP Service Groups, complete the configuration as described in the following table.

Control	Description
Enable WCCP v2 Support	Enables WCCP v2 support on all groups added to the Service Group list.
Multicast TTL	Specify the TTL boundary for the WCCP protocol packets. The default value is 16.

3. Click **Apply** to save your settings to the running configuration.

To add, modify, or remove a service group

1. Under WCCP groups, complete the configuration as described in the following table.

Control	Description
Add a New Service Group	Displays the controls for adding a new service group.
Interface	Select a Steelhead appliance interface to participate in a WCCP service group. RiOS v6.1 and later allows multiple Steelhead interfaces to participate in WCCP on one or more routers for redundancy (RiOS v6.0 and earlier allows a single Steelhead interface). If one of the links goes down, the router can still send traffic to the other active links for optimization. You must include an interface with the service group ID. More than one Steelhead appliance in-path interface can participate in the same service group. For WCCP configuration examples, see the *Riverbed Deployment Guide*. If multiple Steelhead appliances are used in the topology, they must be configured as neighbors. RiOS v6.5 and later requires connection forwarding in a WCCP cluster.
Service Group ID	Enables WCCP v2 support on all groups added to the Service Group list. Specify a number from 0 to 255 to identify the service group on the router. A value of 0 specifies the standard HTTP service group. Riverbed recommends that you use WCCP service groups 61 and 62. **Note:** The service group ID is local to the site where WCCP is used. **Note:** The service group number is not sent across the WAN.
Protocol	Select a traffic protocol from the drop-down list: TCP, UDP, or ICMP. The default value is TCP.
Password/Confirm Password	Optionally, assign a password to the Steelhead appliance interface. This password must be the same password that is on the router. WCCP requires that all routers in a service group have the same password. Passwords are limited to 8 characters.
Priority	Specify the WCCP priority for traffic redirection. If a connection matches multiple service groups on a router, the router chooses the service group with the highest priority. The range is 0-255. The default value is 200. The priority value must be consistent across all Steelhead appliances within a particular service group.

Control	Description
Weight	Specify the percentage of connections that are redirected to a particular Steelhead appliance interface, which is useful for traffic load balancing and failover support. The number of TCP, UDP, or ICMP connections a Steelhead appliance supports determines its weight. The more connections a Steelhead appliance model supports, the heavier the weight of that model. In RiOS v6.1 and later you can modify the weight for each in-path interface to manually tune the proportion of traffic a Steelhead interface receives. A higher weight redirects more traffic to that Steelhead interface. The ratio of traffic redirected to a Steelhead interface is equal to its weight divided by the sum of the weights of all the Steelhead interfaces in the same service group. For example, if there are two Steelhead appliances in a service group and one has a weight of 100 and the other has a weight of 200, the one with the weight 100 receives 1/3 of the traffic and the other receives 2/3 of the traffic. However, since it is generally undesirable for a Steelhead with two WCCP in-path interfaces to receive twice the proportion of traffic, for Steelhead appliances with multiple in-paths connected, each of the in-path weights is divided by the number of that Steelhead's interfaces participating in the service group. For example, if there are two Steelhead appliances in a service group and one has a single interface with weight 100 and the other has two interfaces each with weight 200, the total weight will still equal 300 (100 + 200/2 + 200/2). The one with the weight 100 receives 1/3 of the traffic and each of the other's in-path interfaces receives 1/3 of the traffic. The range is 0-65535. The default value corresponds to the number of TCP connections your Steelhead appliance supports. **Failover Support** To enable single in-path failover support with WCCP groups, define the service group weight to be 0 on the backup Steelhead appliance. If one Steelhead appliance has a weight 0, but another one has a non-zero weight, the Steelhead appliance with weight 0 does not receive any redirected traffic. If all the Steelhead appliances have a weight 0, the traffic is redirected equally among them. The best way to achieve multiple in-path failover support with WCCP groups in RiOS v6.1 and later is to use the same weight on all interfaces from a given Steelhead appliance for a given service group. For example, suppose you have Steelhead A and Steelhead B with two in-path interfaces each. When you configure Steelhead A with weight 100 from both inpath0_0 and inpath0_1 and Steelhead B with weight 200 from both inpath0_0 and inpath0_1, RiOS distributes traffic to Steelhead A and Steelhead B in the ratio of 1:2 as long as at least one interface is up on both Steelhead appliances. In a service group, if an interface with a non-zero weight fails, its weight transfers over to the weight 0 interface of the same service group. For details on using the weight parameter to balance traffic loads and provide failover support in WCCP, see the *Riverbed Deployment Guide*.

Control	Description
Encapsulation Scheme	Specifies the method for transmitting packets between a router or a switch and a Steelhead appliance interface. Select one of the following encapsulation schemes from the drop-down list: • **Either** - Use Layer-2 first; if Layer-2 is not supported, GRE is used. This is the default value. • **GRE** - Generic Routing Encapsulation. The GRE encapsulation method appends a GRE header to a packet before it is forwarded. This can cause fragmentation and imposes a performance penalty on the router and switch, especially during the GRE packet de-encapsulation process. This performance penalty can be too great for production deployments. • **L2** - Layer-2 redirection. The L2 method is generally preferred from a performance standpoint because it requires fewer resources from the router or switch than the GRE does. The L2 method modifies only the destination Ethernet address. However, not all combinations of Cisco hardware and IOS revisions support the L2 method. Also, the L2 method requires the absence of L3 hops between the router or switch and the Steelhead appliance.

Control	Description
Assignment Scheme	Determines which Steelhead interface in a WCCP service group the router or switch selects to redirect traffic to for each connection. The assignment scheme also determines whether the Steelhead interface or the router processes the first traffic packet. The optimal assignment scheme achieves both load balancing and failover support. Select one of the following schemes from the drop-down list: • **Either** - Uses Hash assignment unless the router does not support it. When the router does not support Hash, it uses Mask. This is the default setting. • **Hash** - Redirects traffic based on a hashing scheme and the Weight of the Steelhead interface, providing load balancing and failover support. This scheme uses the CPU to process the first packet of each connection, resulting in slightly lower performance. However, this method generally achieves better load distribution. Riverbed recommends Hash assignment for most Steelhead appliances if the router supports it. The Cisco switches that do not support Hash assignment are the 3750, 4000, and 4500-series, among others. Your hashing scheme can be a combination of the source IP address, destination IP address, source port, or destination port. • **Mask** - Redirects traffic operations to the Steelhead appliances, significantly reducing the load on the redirecting router. Mask assignment processes the first packet in the router hardware, using less CPU cycles and resulting in better performance. Mask assignment in RiOS v5.0.1 and earlier is limited to one Steelhead appliance per service group. The Steelhead appliance with the lowest in-path IP address receives all the traffic. This scheme provides high availability. You can have multiple Steelhead appliances in a service group but only the Steelhead appliance with the lowest in-path IP address receives all the traffic. If the Steelhead appliance with the lowest in-path IP address fails, the Steelhead appliance with the next lowest in-path IP address receives all of the traffic. When the Steelhead appliance with the lowest in-path IP address recovers, it again receives all of the traffic. Mask assignment in RiOS v5.0.2 and later supports load-balancing across multiple active Steelhead appliances. This scheme bases load-balancing decisions (for example, which Steelhead appliance in a service group optimizes a given new connection) on bits pulled out, or *masked,* from the IP address and the TCP port packet header fields. Mask assignment in RiOS v6.1 and later supports load-balancing across multiple active Steelhead appliance interfaces in the same service group. The default mask scheme uses an IP address mask of 0x1741, which is applicable in most situations. However, you can change the IP mask by clicking the service group ID and changing the service group settings and flags. In multiple Steelhead environments, it is often desirable to send all users in subnet range to the same Steelhead. Using mask provides a basic ability to leverage a branch subnet and Steelhead to the same Steelhead in a WCCP cluster. For details and best practices for using assignment schemes, see the *Riverbed Deployment Guide*. **Important:** If you use mask assignment you must ensure that packets on every connection and in both directions (client-to-server and server-to-client), are redirected to the same Steelhead appliance. For details, see the *Riverbed Deployment Guide*.

Control	Description
Source	• **IP Mask** - Specify the service group source IP mask. The default value is 0x1741. • **Port Mask** - Specify the service group source port mask. • **IP Hash** - Specify that the router hash the source IP address to determine traffic to redirect. • **Port Hash** - Specify that the router hash the source port to determine traffic to redirect.
Destination	• **IP Mask** - Specify the service group destination IP mask. • **Port Mask** - Specify the service group destination port mask. • **IP Hash** - Specify that the router hash the destination IP address to determine traffic to redirect. • **Port Hash** - Specify that the router hash the destination port to determine traffic to redirect.
Ports Mode	Select one of the following modes from the drop-down list: • **Ports Disabled** - Select to disable the ports. • **Use Source Ports** - The router determines traffic to redirect based on source ports. • **Use Destination Ports** - The router determines traffic to redirect based on destination ports.
Ports	Specify a comma-separated list of up to seven ports that the router will redirect. Use this option only after selecting either the Use Source Ports or the Use Destination Ports mode.
Router IP Address(es)	Specify a multicast group IP address or a unicast router IP address. You can specify up to 32 routers.
Add	Adds the service group.
Remove Selected Groups	Select the check box next to the name and click **Remove Selected Groups**.

2. Click **Apply** to save your settings to the running configuration.

3. Click **Save** to save your settings permanently.

Verifying a Multiple In-Path Interface Configuration

This section describes how to verify that multiple Steelhead appliances are participating in WCCP with one or more routers using a multiple in-path interface configuration.

1. Because the Steelhead appliances are configured as neighbors, messages appear in the log at INFO level when the neighbors connect to each other, and the log displays a list of in-path IP addresses.

2. When the weight computation is about to begin, a message appears in the log at INFO level that the Steelhead interface with the lowest IP address is taking over as the lead cache.

3. When the weight computation is complete, a REDIRECT_ASSIGN WCCP message appears from the Steelhead interface with the lowest IP address. This message includes the load balancing information from the hash or mask value table.

Note: For more WCCP troubleshooting, see the *Riverbed Deployment Guide*.

Modifying WCCP Group Settings

You modify WCCP service group settings, add additional routers to a service group, and set flags for source and destination ports to redirect traffic (that is, the hash table settings) in the Configure > Networking > WCCP Service Group: <group ID> page.

Before you can modify WCCP service group settings, you must create a WCCP service group. For details on creating a WCCP service group, see "Configuring WCCP" on page 312.

When you are modifying service group settings in RiOS v6.1 or later, the service group description includes the interface.

For details on hash table settings for WCCP, see:
http://www.cisco.com/univercd/home/home.htm.

To modify WCCP service group settings

1. Choose Configure > Networking > WCCP to display the WCCP page.

2. Select the service group ID in the Groups list to expand the page.

Figure 7-25. WCCP Service Group: <Group ID> Page

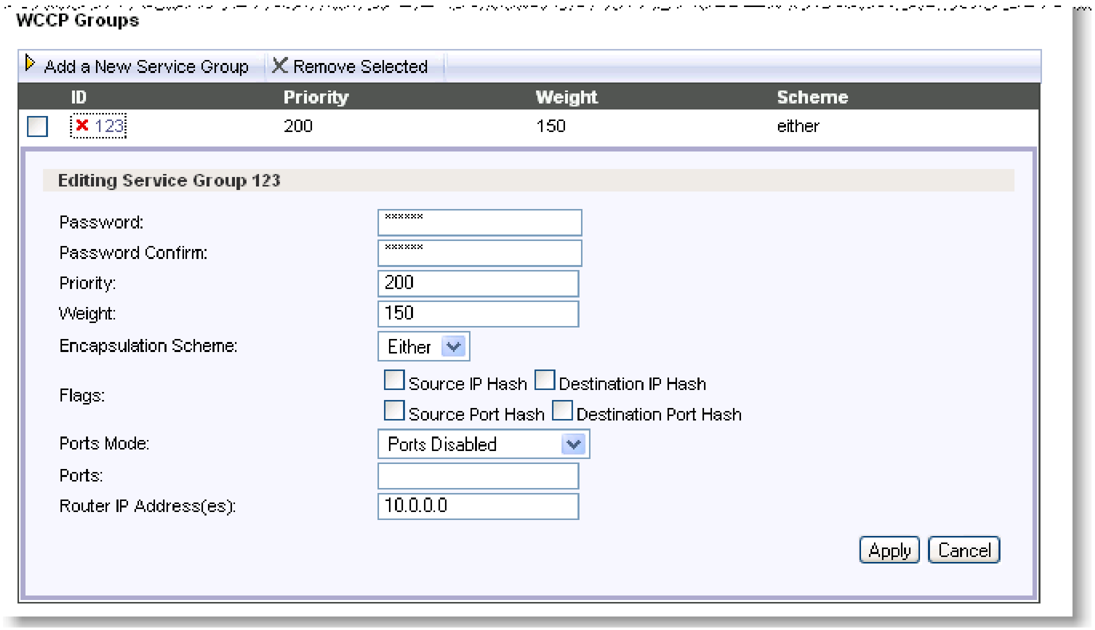

	ID	Priority	Weight	Scheme
□	123	200	150	either

Editing Service Group 123

Password:	******
Password Confirm:	******
Priority:	200
Weight:	150
Encapsulation Scheme:	Either
Flags:	□ Source IP Hash □ Destination IP Hash □ Source Port Hash □ Destination Port Hash
Ports Mode:	Ports Disabled
Ports:	
Router IP Address(es):	10.0.0.0

3. Under Editing Service Group <name><interface>, modify the settings.

4. Click **Apply** to save your settings to the running configuration.

5. Click **Save** to save your settings permanently.

Related Topics

- "Configuring General Service Settings" on page 60
- "Verifying a Multiple In-Path Interface Configuration" on page 318

Configuring Hardware Assist Rules

You configure hardware assist rules in the Configure > Networking > Hardware Assist Rules page. This feature only appears on a Steelhead appliance equipped with one or more Two-Port SR Multimode Fiber 10 Gigabit-Ethernet PCI-E cards.

Hardware Assist rules can automatically bypass all UDP (User Datagram Protocol) connections. You can also configure rules for bypassing specific TCP (Transmission Control Protocol) connections. Automatically bypassing these connections decreases the work load on the local Steelhead appliances because the traffic is immediately sent to the kernel of the host machine or out of the other interface before the Steelhead receives it.

Note: For a hardware assist rule to be applied to a specific 10G bypass card, the corresponding in-path interface must be enabled and have an IP address.

To configure hardware assist rules

1. Choose Configure > Networking > Hardware Assist Rules to display the Hardware Assist Rules page.

Figure 7-26. Hardware Assist Rules Page

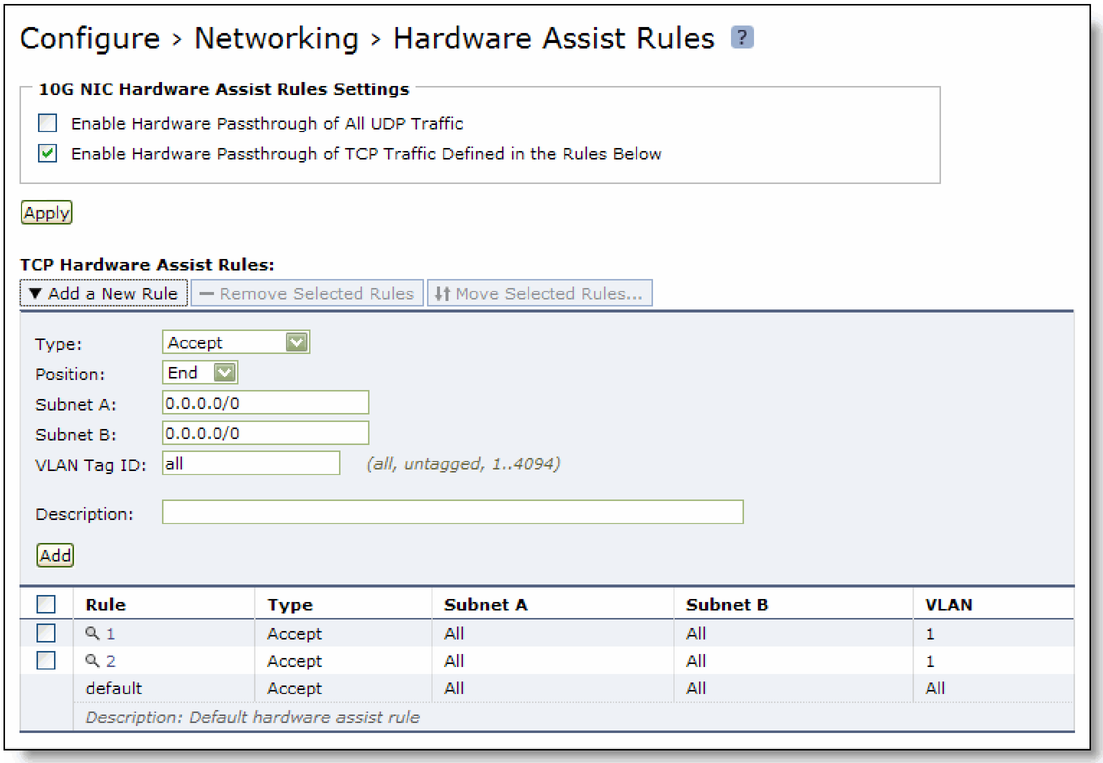

Configure > Networking > Hardware Assist Rules ?

10G NIC Hardware Assist Rules Settings

□ Enable Hardware Passthrough of All UDP Traffic

☑ Enable Hardware Passthrough of TCP Traffic Defined in the Rules Below

TCP Hardware Assist Rules:

Type: Accept

Position: End

Subnet A: 0.0.0.0/0

Subnet B: 0.0.0.0/0

VLAN Tag ID: all (all, untagged, 1..4094)

Description:

□	Rule	Type	Subnet A	Subnet B	VLAN
□	1	Accept	All	All	1
□	2	Accept	All	All	1
	default	Accept	All	All	All
	Description: Default hardware assist rule				

2. Under 10G NIC Hardware Assist Rules Settings, enable pass-through as follows:

   - To automatically pass through all UDP traffic, select the Enable Hardware Passthrough of All UDP Traffic check box.

   - To pass through TCP traffic based on the configured rules, select the Enable Hardware Passthrough of TCP Traffic Defined in the Rules Below check box. TCP pass-through is controlled by rules. The next step describes how to step up hardware assist rules.

Note: All hardware assist rules are ignored unless this check box is selected. No TCP traffic will be passed through.

3. Under TCP Hardware Assist Rules, complete the configuration as described in the following table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Type	Select one of the following rule types: • **Accept** - Accepts rules matching the Subnet A or Subnet B IP address and mask pattern for the optimized connection. • **Pass-Through** - Identifies traffic to be passed through the network unoptimized.

Control	Description
Insert Rule At	Determines the order in which the system evaluates the rule. Select start, end, or a rule number from the drop-down list. The system evaluates rules in numerical order starting with rule 1. If the conditions set in the rule match, then the rule is applied and the system moves on to the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted. In general, filter traffic that is to be unoptimized, discarded, or denied before processing rules for traffic that is to be optimized.
Subnet A	Specify an IP address and mask for the subnet that can be both source and destination together with Subnet B. Use the following format: XXX.XXX.XXX.XXX/XX **Note:** You can specify all or 0.0.0.0/0 as the wildcard for all traffic.
Subnet B	Specify an IP address and mask for the subnet that can be both source and destination together with Subnet A. Use the following format: XXX.XXX.XXX.XXX/XX **Note:** You can specify all or 0.0.0.0/0 as the wildcard for all traffic.
VLAN Tag ID	Optionally, specify a numeric VLAN tag identification number. Select all to specify the rule applies to all VLANs. Select untagged to specify the rule applies to non-tagged connections. **Note:** Pass-through traffic maintains any pre-existing VLAN tagging between the LAN and WAN interfaces. **Note:** To complete the implementation of VLAN tagging, you must set the VLAN tag IDs for the in-path interfaces that the Steelhead appliance uses to communicate with other Steelhead appliances. For details on configuring the in-path interface for the Steelhead appliance, see "Configuring In-Path Rules" on page 28.
Description	Optionally, include a description of the rule.
Add	Adds the new hardware assist rule to the list. You can add up to a maximum number of 50 rules. • RiOS applies the same rule to both LAN and WAN interfaces. • Every 10G card has the same rule set. The Steelhead appliance refreshes the Hardware Assist Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected Rules	Select the check box next to the name and click **Remove Selected Rules**.
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

CHAPTER 8

Configuring System Administrator Settings

This chapter describes how to configure features to assist you in system administration. It includes the following sections:

Configuring Alarm Settings

You can set alarms in the Configure > System Settings > Alarms page.

Enabling alarms is *optional*.

To set alarm parameters

1. Choose Configure > System Settings > Alarms to display the Alarms page.

Figure 8-1. Alarms Page

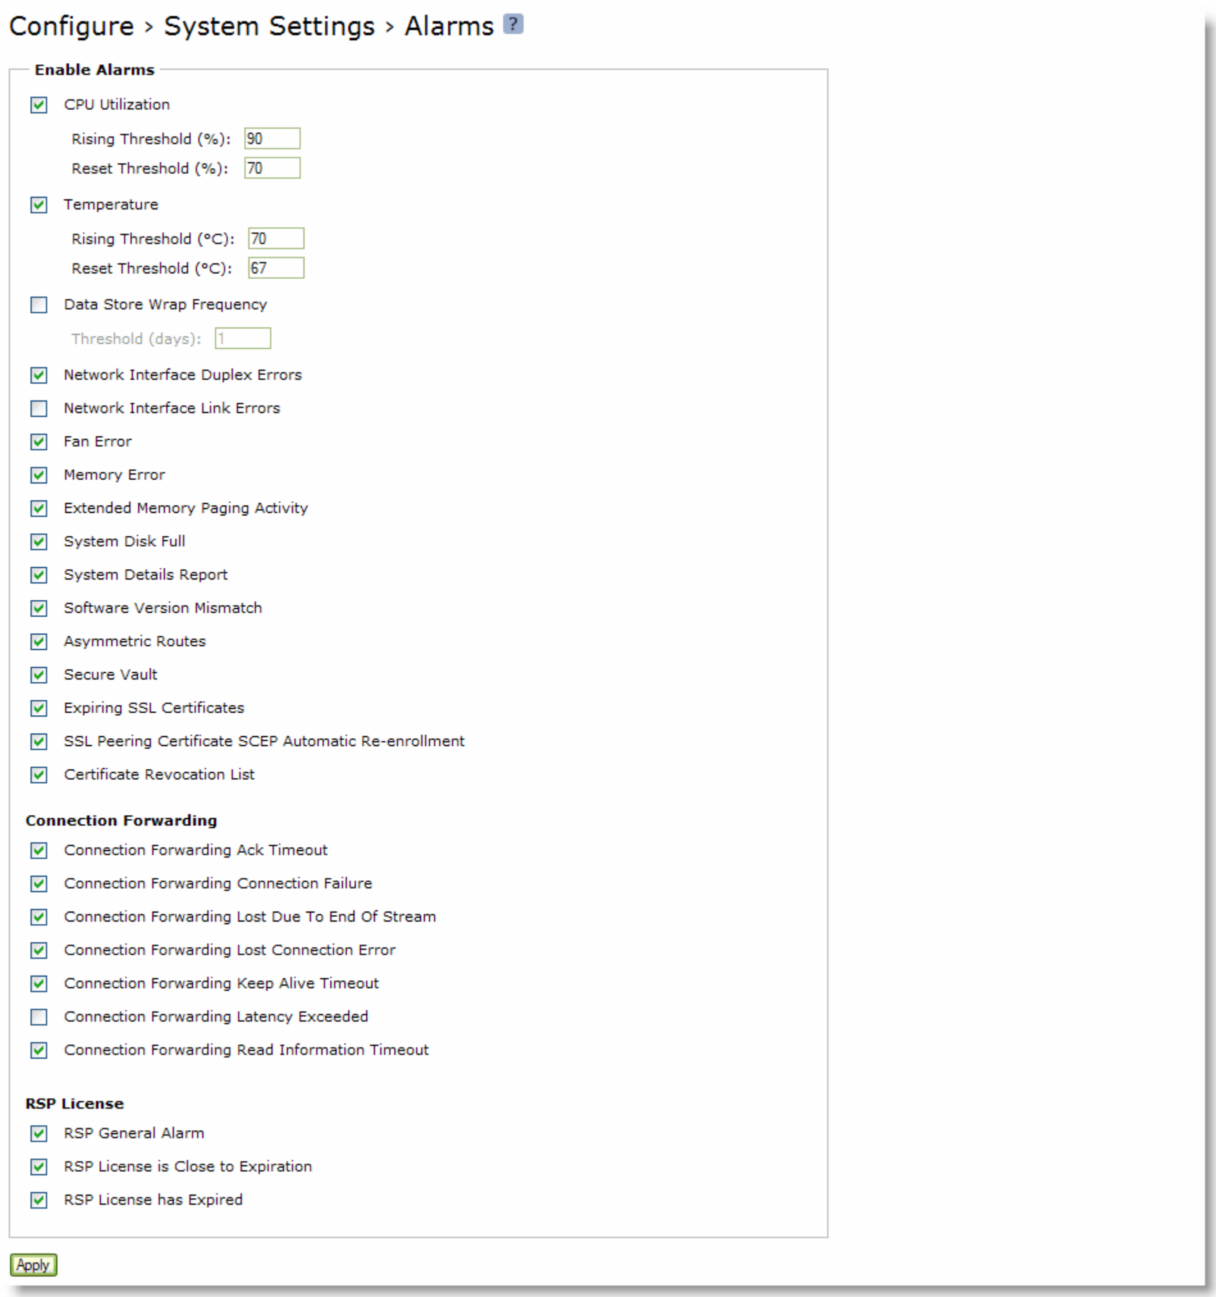

Configure > System Settings > Alarms

Enable Alarms

☑ CPU Utilization

Rising Threshold (%): 90

Reset Threshold (%): 70

☑ Temperature

Rising Threshold (°C): 70

Reset Threshold (°C): 67

□ Data Store Wrap Frequency

Threshold (days): 1

☑ Network Interface Duplex Errors

□ Network Interface Link Errors

☑ Fan Error

☑ Memory Error

☑ Extended Memory Paging Activity

☑ System Disk Full

☑ System Details Report

☑ Software Version Mismatch

☑ Asymmetric Routes

☑ Secure Vault

☑ Expiring SSL Certificates

☑ SSL Peering Certificate SCEP Automatic Re-enrollment

☑ Certificate Revocation List

Connection Forwarding

☑ Connection Forwarding Ack Timeout

☑ Connection Forwarding Connection Failure

☑ Connection Forwarding Lost Due To End Of Stream

☑ Connection Forwarding Lost Connection Error

☑ Connection Forwarding Keep Alive Timeout

□ Connection Forwarding Latency Exceeded

☑ Connection Forwarding Read Information Timeout

RSP License

☑ RSP General Alarm

☑ RSP License is Close to Expiration

☑ RSP License has Expired

2. Under Enable Alarms, complete the configuration as described in the following table.

Control	Description
CPU Utilization	Enables an alarm if the average and peak threshold for the CPU utilization is exceeded. When an alarm reaches the rising threshold, it is activated; when it reaches the lowest or reset threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold. By default, this alarm is enabled, with a rising threshold of 90% and a reset threshold of 70%.
	Rising Threshold - Specify the rising threshold. When an alarm reaches the rising threshold, it is activated. The default value is 90%.
	Reset Threshold - Specify the reset threshold. When an alarm reaches the lowest or reset threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold. The default value is 70%.
Temperature	Enables an alarm when the CPU temperature exceeds the rising threshold. When the CPU returns to the reset threshold, the rising alarm is cleared. The default value for the rising threshold temperature is 80º C; the default reset threshold temperature is 67º C.
	Rising Threshold - Specify the rising threshold (º C). When an alarm reaches the rising threshold, it is activated. The default value is 80º.
	Reset Threshold - Specify the reset threshold (º C). When an alarm reaches the lowest or reset threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold. The default value is 67º.
Data Store Wrap Frequency	Enables an alarm if data in the datastore is replaced with new data before the time period specified. **Threshold** - Specify the number of days before the datastore is replaced. The default value is 1 day.
Network Interface Duplex Errors	Enables an alarm if the system has encountered a large number of packet errors in your network. Make sure the speed and duplex settings on your system match the settings on your switch and router. By default, this alarm is enabled.
Network Interface Link Errors	Enables an alarm and sends an email notification when a link goes down. By default, this alarm is disabled. For WAN/LAN interfaces, an alarm is only triggered if in-path support is enabled for that WAN/LAN pair.
Fan Error	Enables an alarm when an appliance fan error is detected.
Memory Error	Enables an alarm when an appliance memory error is detected.
Extended Memory Paging Activity	Enables the memory paging alarm. If 100 pages are swapped every couple of hours, the system is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support at https://support.riverbed.com. By default, this alarm is enabled.
System Disk Full	Enables an alarm when a system disk full condition is detected. By default, this alarm is enabled.
System Details Report	Enables an alarm if a system component has encountered a problem. By default, this alarm is enabled.

Control	Description
Software Version Mismatch	Enables an alarm if there is a mismatch between software versions in the Riverbed system. By default, this alarm is enabled.
Asymmetric Routes	Enables an alarm if asymmetric routing is detected, an entry is placed in the asymmetric routing table and any subsequent connections from that IP-address pair are passed through unoptimized. Further connections between these hosts are not optimized until that particular asymmetric routing cache entry times out. By default, this alarm is enabled.
Secure Vault	Enables an alarm when an error is detected while initializing the secure vault. This alarm provides links to the Secure Vault page and also appears in the Reports > Diagnostics > Alarm Status page. By default, this alarm is enabled. When the vault is locked, SSL traffic is not optimized and you cannot encrypt the datastore. You can unlock the vault with a password. To unlock the vault, click the link to display the Configure > Security > Secure Vault page and click **Unlock Secure Vault**. When the alarm indicates the password needs to be rekeyed, you can use the default password or reset the password as follows: To clear the alarm using the default password, click **Change Password**. To clear the alarm using a non-default password, specify a new password and click **Unlock**.
Expiring SSL Certificates	Enables an alarm if an SSL certificate is due to expire within 60 days or an expired SSL certificate is detected. By default, this alarm is enabled.
SSL Peering Certificate SCEP Automatic Re-enrollment	Enables an alarm when the Steelhead appliance requests a Simple Certificate Enrollment Protocol (SCEP) server to dynamically re-enroll an SSL peering certificate and the request fails. The Steelhead appliance uses SCEP to dynamically re-enroll a peering certificate to be signed by a certificate authority. The alarm clears automatically when the next automatic re-enrollment succeeds. You can clear the alarm without waiting for the next automatic re-enrollment to succeed with the following CLI command: `protocol ssl peering auto-reenroll last-result clear-alarm` For details, see the *Riverbed Command-Line Interface Reference Manual*. By default, this alarm is enabled.
Certificate Revocation List	Enables an alarm when a Certificate Revocation List (CRL) verification on the server certificate fails. A CRL includes any digital certificates that have been invalidated before their expiration date, including the reasons for their revocation and the names of the issuing certificate signing authorities. A CRL prevents the use of digital certificates and signatures that have been compromised. The certificate authorities that issue the original certificates create and maintain the CRLs. You can clear and disable the alarm with the following CLI command: `no stats alarm crl_error enable` For details, see the *Riverbed Command-Line Interface Reference Manual*. By default, this alarm is enabled.

Control	Description
Connection Forwarding Ack Timeout	Enables an alarm when the connection has been lost because requests have not been acknowledged by a connection forwarding neighbor within the set time-out threshold. This alarm clears automatically the next time all neighbors receive an ACK from this neighbor and the latency of that acknowledgment is less than the set threshold. By default, this alarm is enabled and the time-out period is 1,000 milliseconds (1 second). You can change the time-out period with the following CLI command: `in-path neighbor ack-timer-intvl <milliseconds>` For details, see the *Riverbed Command-Line Interface Reference Manual*. This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.
Connection Forwarding Connection Failure	Enables an alarm when the connection cannot be established with a connection forwarding neighbor. By default, this alarm is enabled. This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.
Connection Forwarding Lost Due To End of Stream	Enables an alarm when the connection is lost since the end of stream was received from the connection forwarding neighbor. By default, this alarm is enabled. This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.
Connection Forwarding Lost Connection Error	Enables an alarm when the connection has been lost with the connection forwarding neighbor due to a communication error. By default, this alarm is enabled. This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.

Connection Forwarding Keep Alive Timeout	Enables an alarm when the connection forwarding neighbor has not responded to a keep-alive message within the specified time-out interval, indicating that the connection has been lost. The alarm clears automatically when all neighbors of the Steelhead appliance are responding to keep-alive messages within the time-out interval.
	By default, this alarm is enabled. The alarm triggers after the number of keep-alive packets that are lost exceeds the keep-alive count. The default keep-alive count is **3** packets and the default keep-alive interval is **1** second.
	You can change the number of packets that must be lost before the alarm triggers and the interval between keep-alive packets with the following CLI commands:
	`show in-path neighbor`
	`in-path neighbor keepalive count <count>`
	`in-path neighbor keepalive interval <seconds>`
	For details, see the *Riverbed Command-Line Interface Reference Manual*.
	This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.

Control	Description

Connection Forwarding Latency Exceeded	Enables an alarm when the amount of latency between connection forwarding neighbors has exceeded the specified threshold. The neighbor latency is the time difference between when the request was sent and the ACK was received.
	By default, this alarm is enabled and the latency threshold is 100 milliseconds.
	The alarm clears automatically when the latency falls below the specified threshold, set with the following CLI commands:
	`stats alarm cf_latency_exceeded rising clear-threshold <threshold>`
	`stats alarm cf_latency exceeded rising error-threshold <threshold>`
	For details, see the *Riverbed Command-Line Interface Reference Manual*.
	This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.

Connection Forwarding Read Information Timeout	Enables an alarm when the Steelhead appliance has timed-out while waiting for an initialization message from the connection forwarding neighbor. By default, this alarm is enabled and the default time-out period is 10,000 milliseconds (10 seconds). You can change the time-out interval with the following CLI command: `in-path neighbor read-timeout <milliseconds>` For details, see the *Riverbed Command-Line Interface Reference Manual*. This alarm includes all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers even if any *one* of the neighbors are in error. Similarly, the alarm clears only when all three neighbors are no longer in error.
RSP General Alarm	(Appears when RSP is installed.) Enables an alarm for general RSP problems including: • No available memory for RSP • An incompatible RSP image is installed. • Virtual Machines are enabled but not currently powered on. • A watchdog activates for any slot that has a watchdog configured. By default, this alarm is enabled.
RSP License is Close to Expiration	(Appears when RSP is installed.) Enables an alarm if an RSP license is due to expire within seven days. By default, this alarm is enabled.
RSP License is Expired	(Appears when RSP is installed.) Enables an alarm when an RSP license has expired. By default, this alarm is enabled.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Related Topics

- "Configuring Email Settings" on page 329
- "Configuring SNMP Settings" on page 336
- "Viewing Process Dumps" on page 481

Setting Announcements

You can create or modify a login message or a message of the day. The login message appears in the Management Console Login page. The message of the day appears in the Home page and when you first log in to the CLI.

To set an announcement

1. Choose Configure > System Settings > Announcements to display the Announcements page.

Figure 8-2. **Announcements Page**

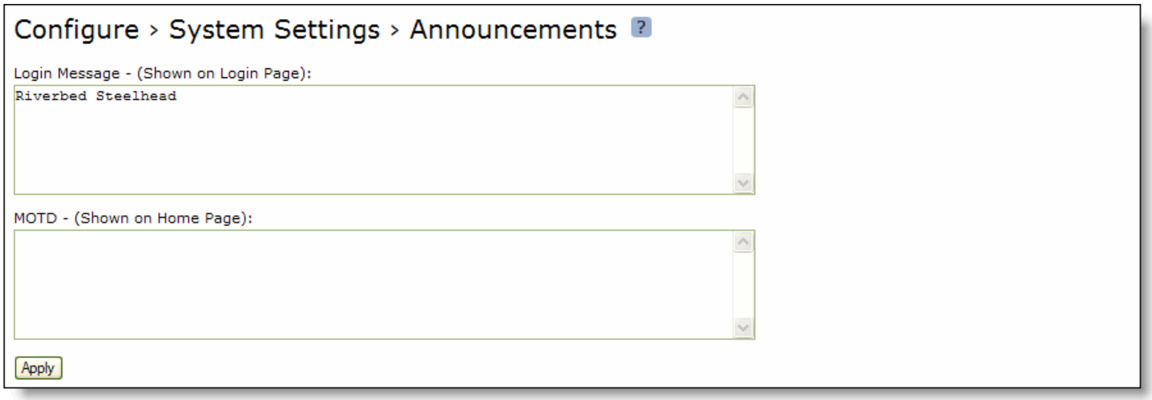

Configure > System Settings > Announcements

Login Message - (Shown on Login Page):

Riverbed Steelhead

MOTD - (Shown on Home Page):

2. Use the controls to complete the configuration as described in the following table.

Control	Description
Login Message	Specify a message in the text box to appear in the Login page.
MOTD	Specify a message in the text box to appear in the Home page.

3. Click **Apply** to view the message before saving.

4. Click **Save** to save your settings permanently.

Configuring Email Settings

You can set email notification parameters for events and failures in the Email Settings page.

By default, email addresses are not specified for event and failure notification.

To set event and failure email notification

1. Choose Configure > System Settings > Email to display the Email Settings page.

Figure 8-3. Email Settings Page

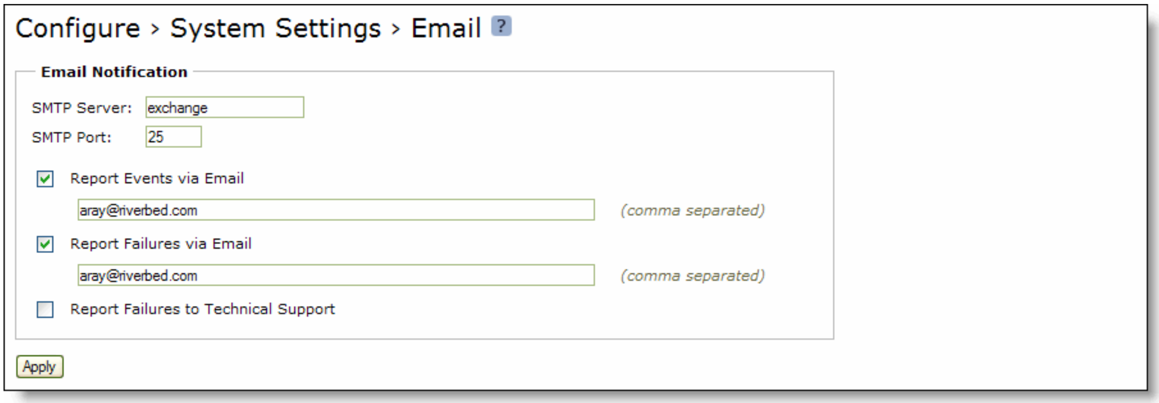

Configure > System Settings > Email ?

Email Notification

SMTP Server: exchange

SMTP Port: 25

☑ Report Events via Email

aray@riverbed.com (comma separated)

☑ Report Failures via Email

aray@riverbed.com (comma separated)

☐ Report Failures to Technical Support

Apply

2. Under Email Notifications, complete the configuration as described in the following table.

Control	Description
SMTP Server	Specify the SMTP server. You must have external DNS and external access for SMTP traffic for this feature to function. **Important:** Make sure you provide a valid SMTP server to ensure that the users you specify receive email notifications for events and failures.
SMTP Port	Specify the port number for the SMTP server.
Report Events via Email	Specify this option to report events through email. Specify a list of email addresses to receive the notification messages.
Report Failures via Email	Specify this option to report failures through email. Specify a list of email addresses to receive the notification messages. Separate addresses by commas.
Report Failures to Technical Support	Specify this option to report serious failures such as system crashes to Riverbed Support. Specify the email addresses to which to send notification messages.. Riverbed recommends that you activate this feature so that problems are promptly corrected. **Important:** This option does not automatically report a disk drive failure. In the event of a disk drive failure, please contact Riverbed Support at support@riverbed.com.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Related Topic

- "Configuring Alarm Settings" on page 323

Configuring Log Settings

You set up local and remote logging in the Configure > System Settings > Logging page.

To set up logging

1. Choose Configure > System Settings > Logging to display the Logging page.

Figure 8-4. **Log Settings Page**

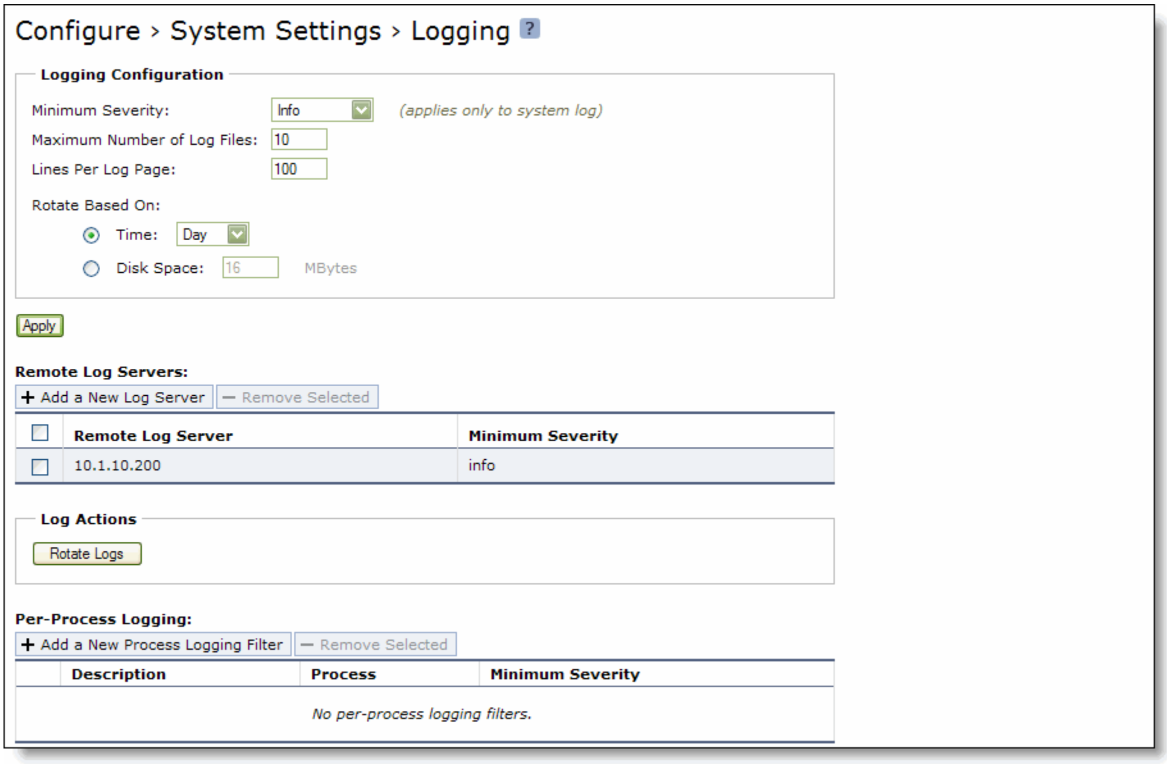

Configure > System Settings > Logging ?

Logging Configuration

Minimum Severity: Info (applies only to system log)
Maximum Number of Log Files: 10
Lines Per Log Page: 100

Rotate Based On:
● Time: Day
○ Disk Space: 16 MBytes

Apply

Remote Log Servers:

+ Add a New Log Server — Remove Selected

□	Remote Log Server	Minimum Severity
□	10.1.10.200	info

Log Actions

Rotate Logs

Per-Process Logging:

+ Add a New Process Logging Filter — Remove Selected

Description	Process	Minimum Severity
	No per-process logging filters.	

2. To rotate the logs immediately, under Log Actions at the bottom of the page, click **Rotate Logs**. After the logs are rotated, the following message appears:

```
logs have been successfully rotated
```

You can also schedule a log rotation based on time or the amount of disk space the log uses, described next.

3. Under Logging Configuration, complete the configuration as described in the following table.

Control	Description
Minimum Severity	Select the minimum severity level for the system log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list: • **Emergency** - Emergency, the system is unusable. • **Alert** - Action must be taken immediately. • **Critical** - Conditions that affect the functionality of the Steelhead appliance. • **Error** - Conditions that probably affect the functionality of the Steelhead appliance. • **Warning** - Conditions that could affect the functionality of the Steelhead appliance, such as authentication failures. • **Notice** - Normal but significant conditions, such as a configuration change. • **Info** - Informational messages that provide general information about system operations. **Note:** This control applies to the system log only. It does not apply to the user log.
Maximum No. of Log Files	Specify the maximum number of logs to store. The default value is 10.
Lines Per Log Page	Specify the number of lines per log page. The default value is 100.
Rotate Based On	Specifies the rotation option: • **Time** - Select Day, Week, or Month from the drop-down list. • **Disk Space** - Specify how much disk space, in megabytes, the log uses before it rotates. The default value is 16 MB. **Note:** The log file size is checked at 10 minute intervals. If there is an unusually large amount of logging activity, it is possible for a log file to grow larger than the set disk space limit in that period of time.

4. Click **Apply** to apply your changes to the running configuration.

5. Click **Save** to save your settings permanently.

To add or remove a log server

1. To add or remove a log server, complete the configuration as described in the following table.

Control	Description
Add a New Log Server	Displays the controls for configuring new log servers.
Server IP	Specify the server IP address.

Control	Description
Minimum Severity	Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list: • **Emergency** - Emergency, the system is unusable. • **Alert** - Action must be taken immediately. • **Critical** - Conditions that affect the functionality of the Steelhead appliance. • **Error** - Conditions that probably affect the functionality of the Steelhead appliance. • **Warning** - Conditions that could affect the functionality of the Steelhead appliance, such as authentication failures. • **Notice** - Normal but significant conditions, such as a configuration change. • **Info** - Informational messages that provide general information about system operations.
Add	Adds the server to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Apply** to apply your changes to the running configuration.

3. Click **Save** to save your settings permanently.

Filtering Logs by Application or Process

RiOS v6.0 and later lets you filter a log by one or more applications or one or more processes. This is particularly useful when capturing data at a lower severity level where a Steelhead appliance might not be able to sustain the flow of logging data the service is committing to disk.

To filter a log

1. Choose Configure > System Settings > Logging to display the Logging page.

Figure 8-5. Filtering a Log

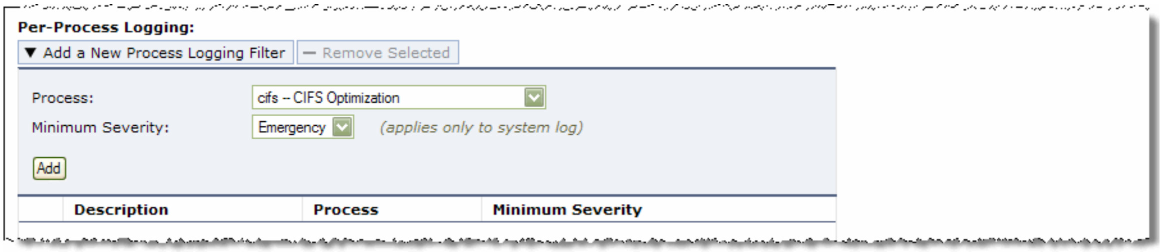

Per-Process Logging:

Process: cifs -- CIFS Optimization

Minimum Severity: Emergency (applies only to system log)

	Description	Process	Minimum Severity

2. Under Per-Process Logging, complete the configuration as described in the following table.

Control	Description
Add a New Process Logging Filter	Displays the controls for adding a process level logging filter.
Process	Select a process to include in the log from the drop-down list: • **cifs** - CIFS Optimization. • **cmcfc** - CMC automatic registration utility. • **rgp** - CMC connector, which handles CMC appliance communication. • **rgpd** - CMC client daemon, the connection manager. • **cli** - Command-Line Interface. • **mgmtd** - Device control and management, which directs the entire device management system. It handles message passing between various management daemons, managing system configuration and general application of system configuration on the hardware underneath through the **hald**. • **http** - HTTP optimization. • **hald** - Hardware Abstraction Daemon, which handles access to the hardware. • **notes** - Lotus Notes optimization. • **mapi** - MAPI optimization. • **nfs** - NFS optimization. • **pm** - Process Manager, which handles launching of internal system daemons and keeps them up and running. • **sched** - Process Scheduler, which handles one-time scheduled events. • **virtwrapperd** - RSP VMWare interface. • **rspd** - RSP Watchdog. • **statsd** - Statistics Collector, which handles queries and storage of system statistics. • **wdt** - Watchdog Timer, the motherboard watchdog daemon. • **webasd** - Web Application Process, which handles the Web user interface. • **domain auth** - Windows Domain Authentication.
Minimum Severity	Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list: • **Emergency** - Emergency, the system is unusable. • **Alert** - Action must be taken immediately. • **Critical** - Conditions that affect the functionality of the Steelhead appliance. • **Error** - Conditions that probably affect the functionality of the Steelhead appliance. • **Warning** - Conditions that could affect the functionality of the Steelhead appliance, such authentication failures. • **Notice** - Normal but significant conditions, such as a configuration change. • **Info** - Informational messages that provide general information about system operations.
Add	Adds the filter to the list. The process now logs at the selected severity and higher level.
Remove Selected	Select the check box next to the name and click **Remove Selected** to remove the filter.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Configuring Monitored Ports

You set TCP ports you want to monitor in the Configure > System Settings > Monitored Ports page. The ports you specify appear in the Traffic Summary report. Make sure the description you specify helps you identify the type of traffic on the port.

The Steelhead appliance automatically discovers all the ports in the system that have traffic. Discovered ports, with a label (if one exists), are added to the Traffic Summary report. If a label does not exist then an **unknown** label is added to the discovered port. To change the **unknown** label to a name representing the port, you must add the port with a new label. All statistics for this new port label are preserved from the time the port was discovered.

For details, see “Viewing Traffic Summary Reports” on page 412.

By default, traffic is monitored on ports 21 (FTP), 80 (HTTP), 139 (CIFS:NetBIOS), 443 (SSL), 445 (CIFS:TCP), 1352 (Lotus Notes), 1433 (SQL:TDS), 1748 (SRDF), 3225 (FCIP), 3226 (FCIP), 3227 (FCIP), 3228 (FCIP), 7830 (MAPI), 8777 (RCU), and 10566 (SnapMirror).

To set monitored ports

1. Choose Configure > System Settings > Monitored Ports to display the Monitored Ports page.

Figure 8-6. Monitored Ports Page

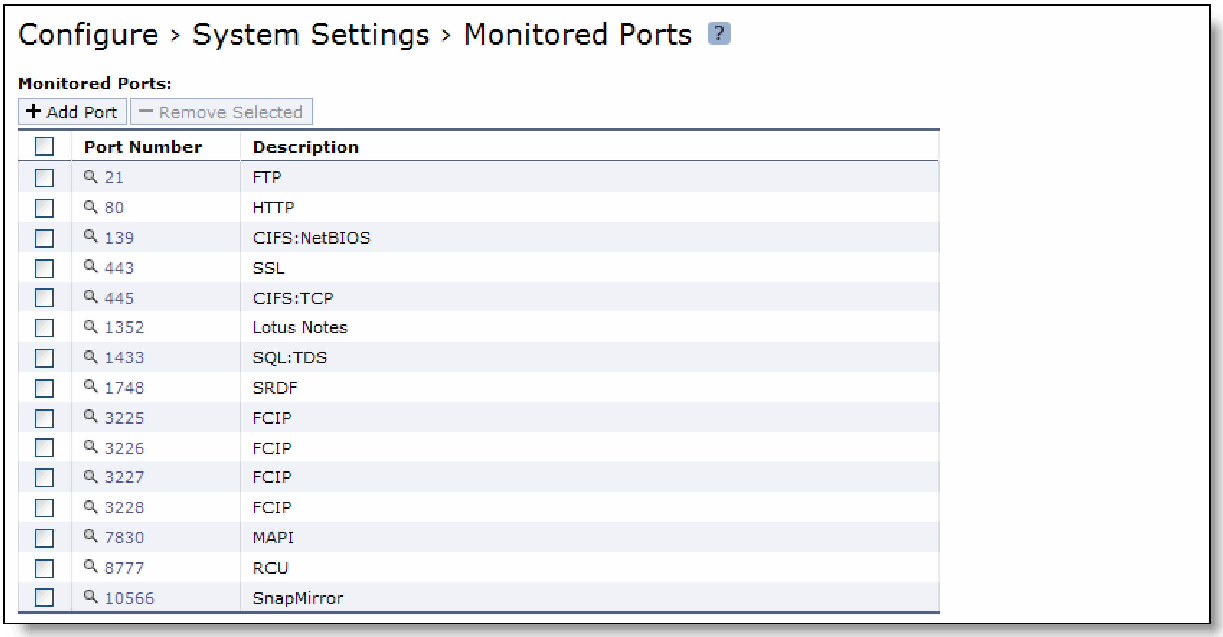

□	Port Number	Description
□	21	FTP
□	80	HTTP
□	139	CIFS:NetBIOS
□	443	SSL
□	445	CIFS:TCP
□	1352	Lotus Notes
□	1433	SQL:TDS
□	1748	SRDF
□	3225	FCIP
□	3226	FCIP
□	3227	FCIP
□	3228	FCIP
□	7830	MAPI
□	8777	RCU
□	10566	SnapMirror

2. Complete the configuration as described in the following table.

Control	Description
Add Port	Displays the controls to add a new port.
Port Number	Specify the port to be monitored.
Port Description	Specify a description of the type of traffic on the port.
Add	Displays the controls for adding a port.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. To modify a monitored port, click the magnifying glass next to the port and complete the configuration as described in the following table.

Control	Description
Port Description	Specify a description of the type of traffic on the port.
Apply Changes	Applies your settings to the running configuration.
Cancel	Cancels your actions.

4. Click **Save** to save your settings permanently.

Configuring SNMP Settings

You configure SNMP contact and trap receiver settings to allow events to be reported to an SNMP entity in the Configure > System Settings > SNMP Basic page.

Traps are messages sent by an SNMP entity that indicate the occurrence of an event. The default system configuration does not include SNMP traps.

RiOS v5.0 provides support for the following:

- SNMP Version 1
- SNMP Version 2c

RiOS v6.0 and later provides support for the following:

- SNMP Version 3, which provides authentication through the User-based Security Model (USM).
- View-Based Access Control Mechanism (VACM), which provides richer access control.

For details on SNMP traps sent to configured servers, see “SNMP Traps” on page 498.

For details on MIBs, see “Accessing the Steelhead Enterprise MIB” on page 497.

To set general SNMP parameters

1. Choose Configure > System Settings > SNMP Basic to display the SNMP Basic page.

Figure 8-7. SNMP Basic Page

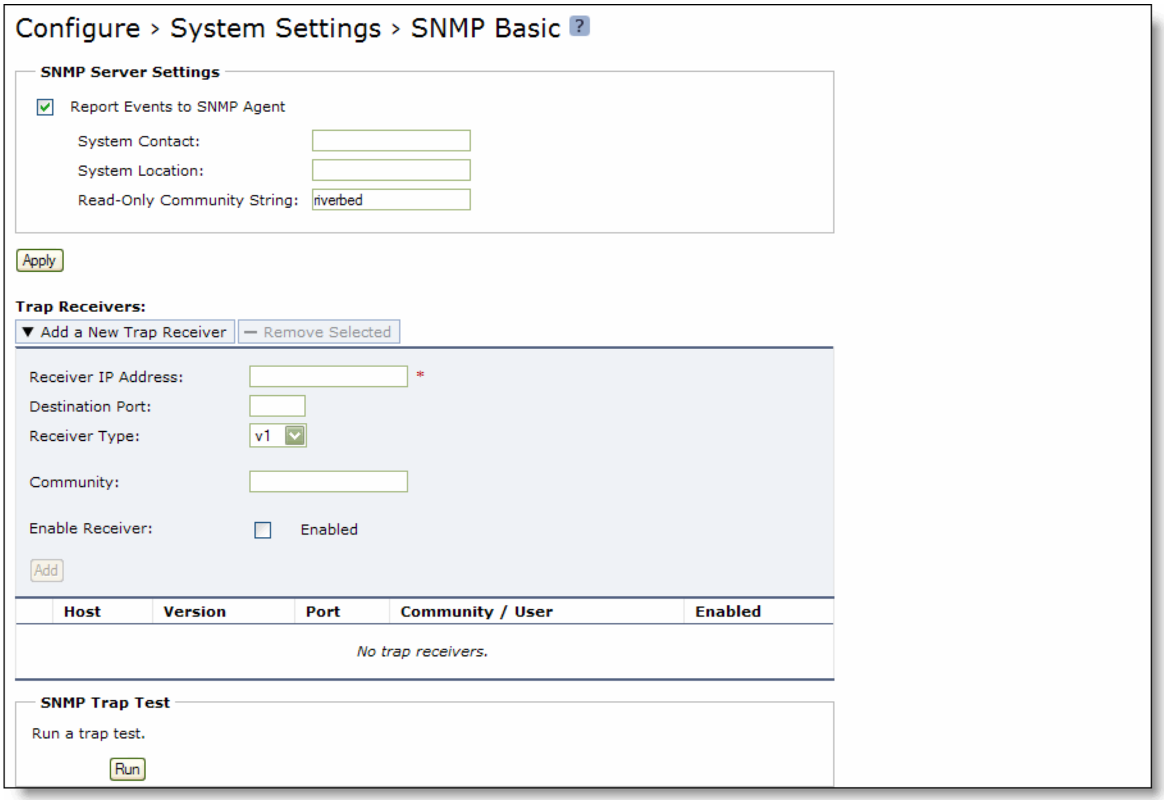

Configure > System Settings > SNMP Basic

SNMP Server Settings

☑ Report Events to SNMP Agent

System Contact:

System Location:

Read-Only Community String: riverbed

Trap Receivers:

Receiver IP Address: *

Destination Port:

Receiver Type: v1

Community:

Enable Receiver: □ Enabled

	Host	Version	Port	Community / User	Enabled

No trap receivers.

SNMP Trap Test

Run a trap test.

2. Under SNMP Server Settings, complete the configuration as described in the following table.

Control	Description
Report Events to SNMP Agent	Enables event reporting to an SNMP entity.
System Contact	Specify the user name for the SNMP contact.
System Location	Specify the physical location of the SNMP system.
Read-Only Community String	Specify a password-like string to identify the read-only community. For example: public. This community string overrides any VACM settings.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To add or remove a trap receiver

1. Under trap receivers, complete the configuration as described in the following table.

Control	Description
Add a New Trap Receiver	Displays the controls to add a new trap receiver.
Receiver IP Address	Specify the destination IP address for the SNMP trap.
Destination Port	Specify the destination port.
Receiver Type	Select SNMP version v1, v2c, or v3 (User-based Security Model).
Community	For v1 or v2 trap receivers, specify the SNMP community name; for example, public or private v3 trap receivers need a remote user with an authentication protocol, and a password and security level.
Enable Receiver	Enables the trap receiver.
Add	Adds a new trap receiver to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

2. Click **Save** to save your settings permanently.

To test an SNMP trap

1. Choose Configure > System Settings > SNMP Basic to display the SNMP Basic page.

2. Under SNMP Trap Test, click **Run**.

Configuring SNMP v3

SNMP v3 provides additional authentication and access control for message security. For example, you can verify the identity of the SNMP entity (manager or agent) sending the message.

Using SNMP v3 is more secure than SNMP v1 or v2; however, it requires more configuration steps to provide the additional security features.

Basic Steps

1. Create the SNMP-server users. Users can be authenticated using either a password or a key.

2. Configure SNMP-server views to define which part of the SNMP MIB tree will be visible.

3. Configure SNMP-server groups, which map users to views, allowing you to control who can view what SNMP information.

4. Configure the SNMP-server access policies that contain a set of rules defining access rights. Based on these rules, the entity decides how to process a given request.

To create users for SNMP v3

1. Choose Configure > System Settings > SNMP v3 to display the SNMP v3 page.

Figure 8-8. SNMP v3 Page

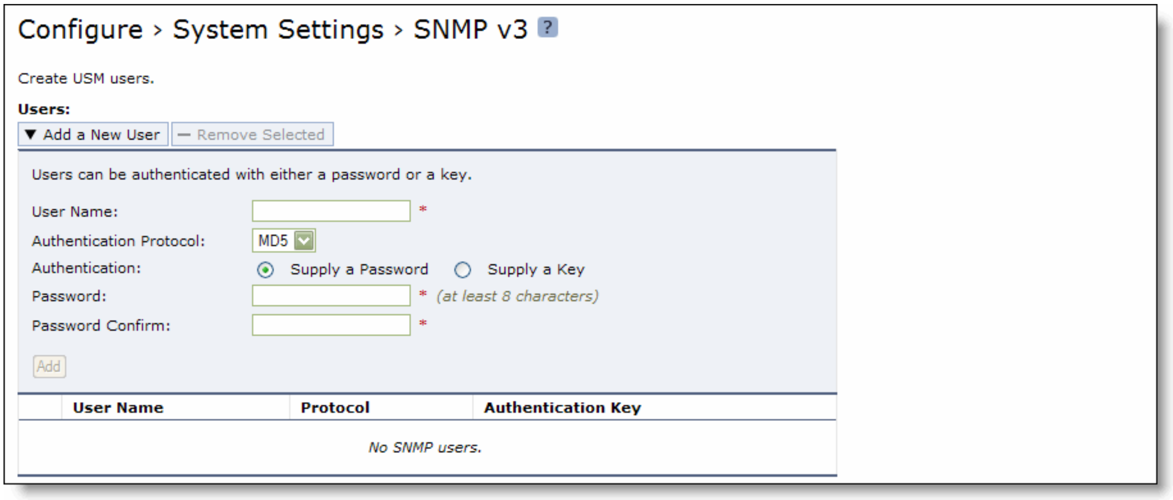

Configure > System Settings > SNMP v3

Create USM users.

Users:

Users can be authenticated with either a password or a key.

User Name: *

Authentication Protocol: MD5

Authentication: ● Supply a Password ○ Supply a Key

Password: * *(at least 8 characters)*

Password Confirm: *

	User Name	Protocol	Authentication Key
		No SNMP users.	

2. Under Users, complete the configuration as described in the following table.

Control	Description
Add a New User	Displays the controls to add a new user.
User Name	Specify the user name.
Authentication Protocol	Select an authentication method from the drop-down list: • **MD5** - Specifies the Message-Digest 5 algorithm, a widely-used cryptographic hash function with a 128-bit hash value. This is the default value. • **SHA** - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.
Authentication	Optionally, select either Supply a Password or Supply a Key to use while authenticating users.
Password/Password Confirm	Specify a password. The password must have a minimum of eight characters. Confirm the password in the Password Confirm text box.
Key	(Appears only when you select Supply A Key.) Specify a unique authentication key. The key is a MD5 or SHA-1 digest created using md5sum or sha1sum.
Add	Adds the user.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

SNMP Authentication and Access Control

The features on this page apply to SNMP v1, v2c, and v3 unless noted otherwise:

- **Security Names** - Identify an individual user (v1 or v2c only).
- **Secure Groups** - Identify a security-name, security model by a group, and referred to by a group-name.
- **Secure Views** - Create a custom view using the VACM that controls who can access which MIB objects under agent management by including or excluding specific OIDs. For example, some users have access to critical read-write control data, while some users have access only to read-only data. For a list of OIDs, see "SNMP Traps" on page 498.
- **Security Models** - A security model identifies the SNMP version associated with a user for the group in which the user resides.
- **Secure Access Policies** - Defines who gets access to which type of information. An access-policy is a comprised of <group-name, security-model, security-level, read-view-name>
  - read-view-name is a preconfigured view that applies to read requests by this security-name.
  - write-view-name is a preconfigured view that applies to write requests by this security-name.
  - notify-view-name is a preconfigured view that applies to write requests to this security-name.

An access-policy is the configurable set of rules, based on which, the entity decides how to process a given request.

To set secure user names

1. Choose Configure > System Settings > SNMP ACLs to display the SNMP ACLs page.

Figure 8-9. **SNMP ACLs Page - Security Names**

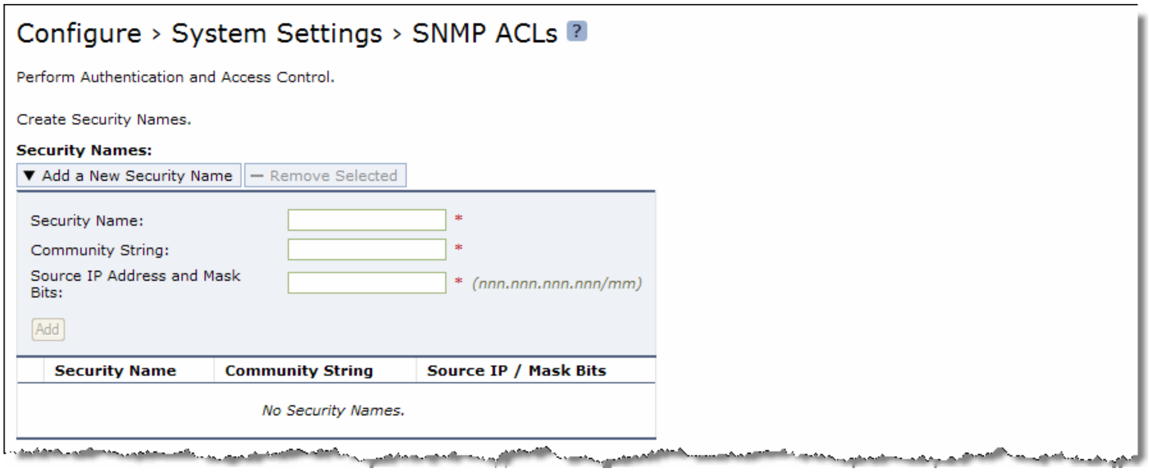

Configure > System Settings > SNMP ACLs ?

Perform Authentication and Access Control.

Create Security Names.

Security Names:

▼ Add a New Security Name — Remove Selected

Security Name: *
Community String: *
Source IP Address and Mask Bits: * *(nnn.nnn.nnn.nnn/mm)*

Add

	Security Name	Community String	Source IP / Mask Bits
		No Security Names.	

2. Under Security Names, complete the configuration as described in the following table.

Control	Description
Add a New Security Name	Displays the controls to add a security name.
Security Name	Specify a name to identify a requestor allowed to issue gets and sets (v1 and v2c only). The security name may make changes to the View Based Access Control Model (VACM) security name configuration. **Note:** This control does not apply to SNMPv3 queries. To restrict v3 USM users to poll from a particular subnet, use the RiOS Management ACL feature, located in the Configure > Security > Management ACL page. **Note:** Traps for v1 and v2c are independent of the security name.
Community String	Specify the password-like community string to control access. Use a combination of uppercase, lowercase, and numerical characters to reduce the chance of unauthorized access to the Steelhead appliance. **Note:** If you specify a read-only community string (located in the SNMP Basic page under SNMP Server Settings), it takes precedence over this community name and allows users to access the entire MIB tree from any source host. If this is not desired, delete the read-only community string. **Note:** To create multiple SNMP community strings on a Steelhead, leave the default public community string and then create a second read-only community string with a different security name. Or, you can delete the default public string and create two new SNMP ACLs with unique names.
Source IP Address and Mask Bits	Specify the host IP address and mask bits to which you permit access using the security name and community string.
Add	Adds the security name.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To set secure groups

1. Choose Configure > System Settings > SNMP ACLs to display the SNMP ACLs page.

Figure 8-10. SNMP ACLs Page - Groups

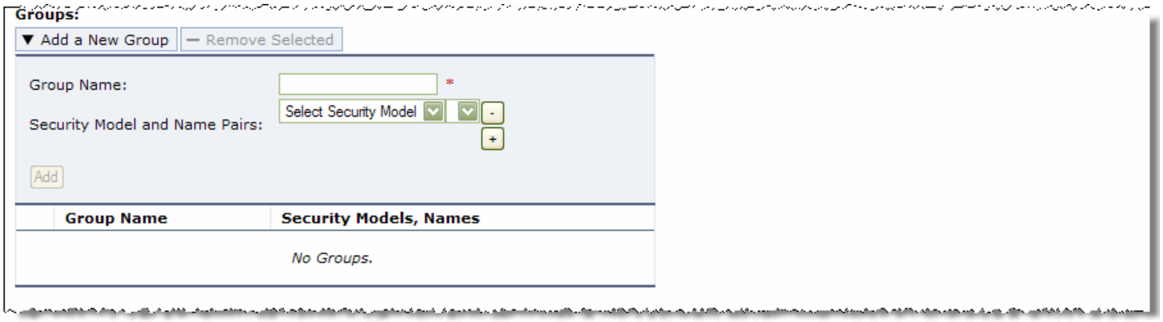

Groups:

▼ Add a New Group — Remove Selected

Group Name: *

Security Model and Name Pairs: Select Security Model

Add

Group Name	Security Models, Names
	No Groups.

2. Under Groups, complete the configuration as described in the following table.

Control	Description
Add a New Group	Displays the controls to add a new group
Group Name	Specify a group name.
Security Models and Name Pairs	Click the **+** button and select a security model from the drop-down list: • v1 or v2c - displays another drop-down menu; select a security name. • v3 (usm) - displays another drop-down menu, select a user. To add another Security Model and Name pair, click the plus sign (**+**).
Add	Adds the group name and security model and name pairs.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To set secure views

1. Choose Configure > System Settings > SNMP ACLs to display the SNMP ACLs page.

Figure 8-11. SNMP ACLs Page - Views

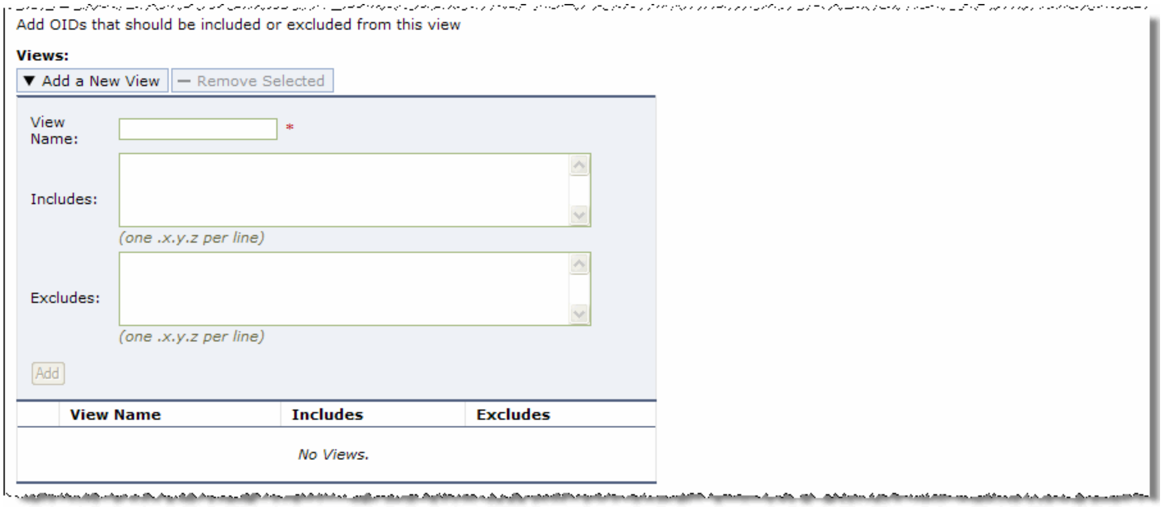

Add OIDs that should be included or excluded from this view

View Name	Includes	Excludes
No Views.		

2. Under Views, complete the configuration as described in the following table.

Control	Description
Add a New View	Displays the controls to add a new view.
View Name	Specify a descriptive view name to facilitate administration.

Control	Description
Includes	Specify the Object Identifiers (OIDs) to include in the view, separated by commas; for example, .1.3.6.1.4.1. By default, the view excludes all OIDs. You can specify .iso or any subtree or subtree branch. You can specify an OID number or use its string form; for example, .iso.org.dod.internet.private.enterprises.rbt.products.steelhead.system.model
Excludes	Specify the OIDs to exclude in the view, separated by commas. By default, the view excludes all OIDs.
Add	Adds the view.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

To add an access policy

1. Choose Configure > System Settings > SNMP ACLs to display the SNMP ACLs page.

Figure 8-12. SNMP ACLs Page

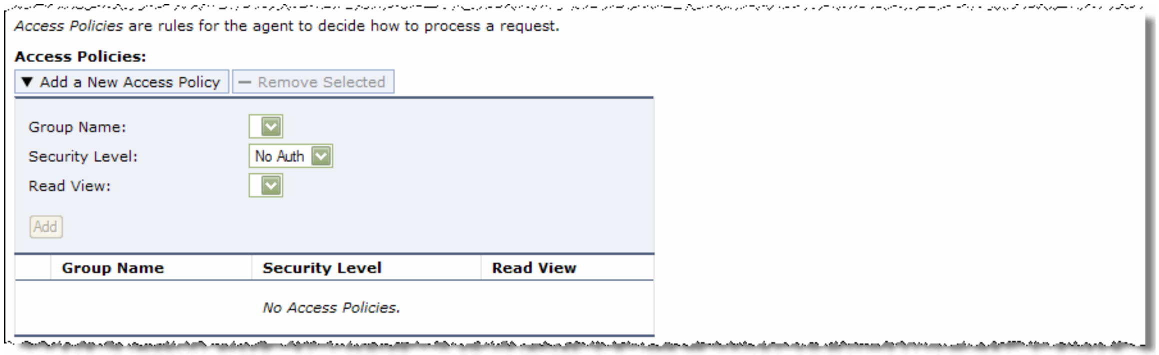

Access Policies are rules for the agent to decide how to process a request.

Access Policies:

Group Name:

Security Level: No Auth

Read View:

	Group Name	Security Level	Read View
		No Access Policies.	

2. Under Access Policies, complete the configuration as described in the following table.

Control	Description
Add a New Access Policy	Displays the controls to add a new access policy.
Group Name	Select a group name from the drop-down list.
Security Level	Determines whether a single atomic message exchange is authenticated. Select one of the following from the drop-down list: • **No Auth** - Does not authenticate packets and does not use privacy. This is the default setting. • **Auth** - Authenticates packets but does not use privacy. **Note:** A security level applies to a group, not to an individual user.
Read View	Select a view from the drop-down list.

Control	Description
Add	Adds the policy to the policy list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

CHAPTER 9 Managing Steelhead Appliances

This chapter describes tasks you perform for routine management of the Steelhead appliance. It includes the following sections:

Starting and Stopping the Optimization Service

You can start, stop, and restart the Steelhead service in the Configure > Maintenance > Services page. You can also use this page to reset the service alarm after it has been triggered.

The Steelhead service is a daemon that executes in the background, performing operations when required.

Many of the Steelhead service commands are initiated at startup. It is important to restart the Steelhead service when you have made changes to your configuration.

Important: Restarting the Steelhead service disrupts existing network connections that are proxied through the Steelhead appliance.

To start, stop, or restart services

1. Choose Configure > Maintenance > Services to display the Services page.

Figure 9-1. Services Page

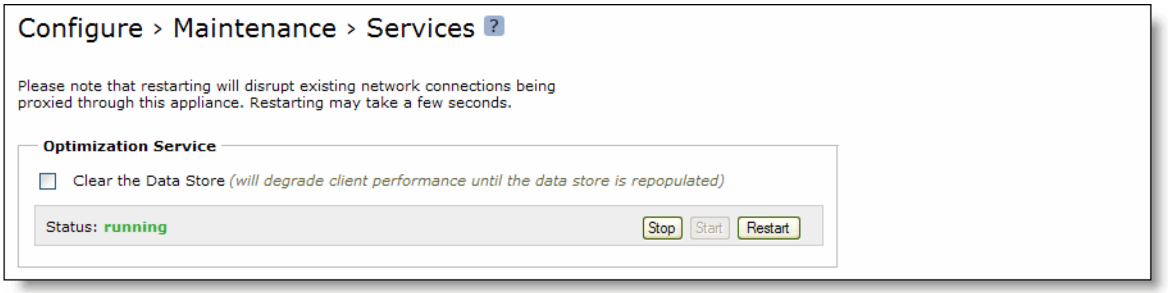

Configure > Maintenance > Services

Please note that restarting will disrupt existing network connections being proxied through this appliance. Restarting may take a few seconds.

Optimization Service

☐ Clear the Data Store *(will degrade client performance until the data store is repopulated)*

Status: running

2. Under Optimization Service click **Stop**, **Start**, or **Restart**.

3. Click **Save** to save your settings permanently.

Tip: To remove data from the datastore, click **Clear the Data Store**. For details, see "Clearing the Datastore" on page 80.

To reset the service alarm

1. Choose Configure > Maintenance > Services to display the Services page. The option to reset the service alarm appears only after the service triggers the Reset Service Alarm.

2. Under Reset Service Alarm, click **Reset Service Alarm**.

3. Click **Save** to save your settings permanently.

Configuring Scheduled Jobs

You can view completed, pending, inactive jobs, as well as jobs that were not completed because of an error in the Configure > Maintenance > Scheduled Jobs page. You can also delete a job, change its status, or modify its properties.

Jobs are commands that are scheduled to execute at a time you specify.

You can use the Management Console to:

- schedule an RSP High Availability transfer.
- schedule a software upgrade.
- generate multiple TCP trace dumps on a specific date and time.

To schedule all other jobs, you must use the Riverbed CLI.

For details on scheduling jobs using the CLI, see the *Riverbed Command-Line Interface Reference Manual*.

To configure scheduled jobs

1. Choose Configure > Maintenance > Scheduled Jobs to display the Scheduled Jobs page.

Figure 9-2. Scheduled Jobs Page

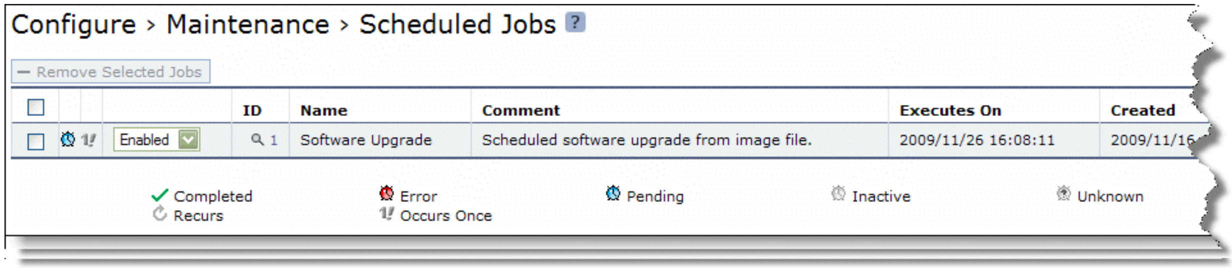

Configure > Maintenance > Scheduled Jobs

□			ID	Name	Comment	Executes On	Created
□		Enabled	1	Software Upgrade	Scheduled software upgrade from image file.	2009/11/26 16:08:11	2009/11/16

2. Select the Job ID number to display details about the job.

3. Under Details for Job <#>, complete the configuration as described in the following table.

Control	Description
Name	Specify a name for the job.
Comment	Specify a comment.
Interval (seconds)	Specify how often the job runs. The default value is 0, which runs the job once.
Executes On	Specify the date on which the job runs.
Enable/Disable Job	Enables the job.
Apply Changes	Applies the changes to the current configuration.
Cancel This Job	Cancels the job.
Execute Now	Runs the job.
Remove Selected Jobs	Select the check box next to the name and click **Remove Selected Jobs**.

4. Click **Save** to save your settings permanently.

Upgrading Your Software

You can upgrade or revert to a backup version of the software in the Configure > Maintenance > Software Upgrade page.

The bottom of the page displays the software version history of the Steelhead appliance, which includes the version number and the software installation date.

To find allowed upgrades between RiOS versions and recommended upgrade paths, use the Software Upgrade tool on the Riverbed Support Site at https://support.riverbed.com. The tool includes all of the recommended intermediate RiOS versions.

CSH VSH You can upgrade software on the ESX Cloud Steelhead and the Virtual Steelhead using the following procedure. To upgrade the software on the AWS Cloud Steelhead, use the Riverbed Cloud Portal. For details, see the *Riverbed Cloud Services Deployment Guide*.

To upgrade or revert software versions

1. Choose Configure > Maintenance > Software Upgrade to display the Software Upgrade page.

Figure 9-3. Software Upgrade Page

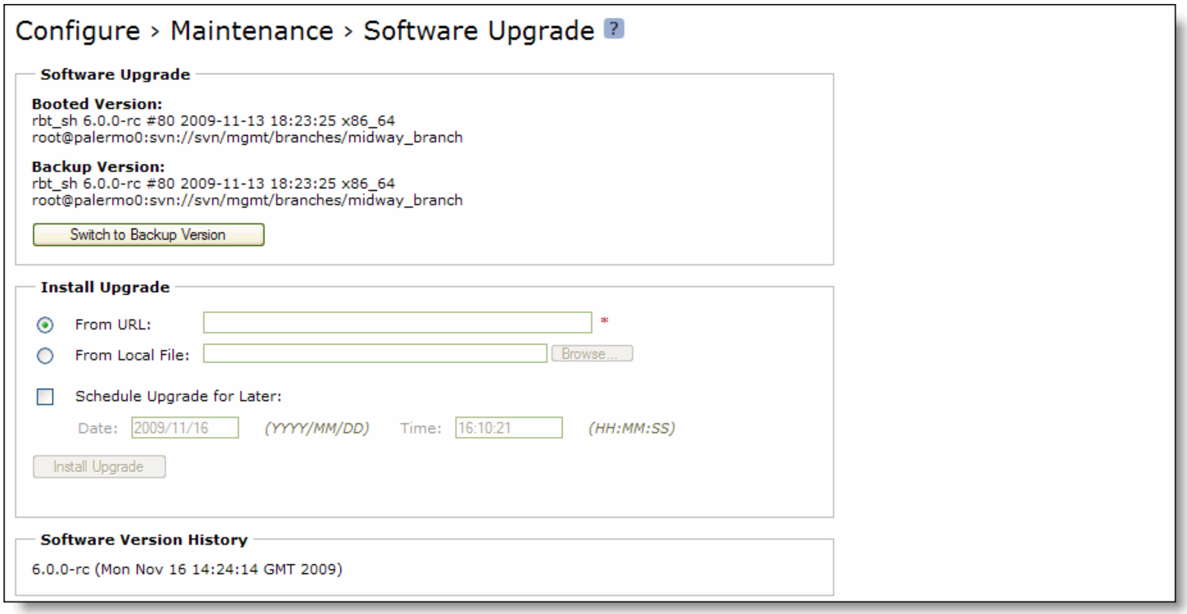

Configure > Maintenance > Software Upgrade ?

Software Upgrade

Booted Version:
rbt_sh 6.0.0-rc #80 2009-11-13 18:23:25 x86_64
root@palermo0:svn://svn/mgmt/branches/midway_branch

Backup Version:
rbt_sh 6.0.0-rc #80 2009-11-13 18:23:25 x86_64
root@palermo0:svn://svn/mgmt/branches/midway_branch

Switch to Backup Version

Install Upgrade

● From URL: ____ *

○ From Local File: ____ Browse...

□ Schedule Upgrade for Later:

Date: 2009/11/16 (YYYY/MM/DD) Time: 16:10:21 (HH:MM:SS)

Install Upgrade

Software Version History

6.0.0-rc (Mon Nov 16 14:24:14 GMT 2009)

2. Under Software Upgrade, complete the configuration as described in the following table.

Control	Description
Switch to Backup Version	Switches to the backup version on the next reboot.
Cancel	Cancels the software version switch on the next reboot.

3. Under Install Upgrade, complete the configuration as described in the following table.

Control	Description
From URL	Click this option and specify the URL. If you specify a URL in the URL text box, the image is uploaded, installed, and the system is rebooted at the time you specify.
From Local File	Click this option and specify the path or click **Browse** to go to the local file directory. If you specify a file to upload in the Local File text box, the image is uploaded immediately, however the image is installed and the system is rebooted at the time you specify.
Schedule Upgrade for Later	Schedules the upgrade process. Specify the date and time to run the upgrade: • **Date and Time** - Use the following formats: YYYY/MM/DD, HH:MM:SS
Install Upgrade	Installs the software upgrade on your system.
Cancel	Cancels your changes.

4. Reboot the Steelhead appliance. For details, see "Rebooting and Shutting Down the Steelhead Appliance" on page 349.

Related Topic

- "Configuring Scheduled Jobs" on page 346

Rebooting and Shutting Down the Steelhead Appliance

You can reboot or shut down the system in the Configure > Maintenance > Reboot/Shutdown page.

Rebooting the system disrupts existing network connections that are currently proxied through it. Rebooting can take a few minutes.

When you shut down the system, connections are broken and optimization ceases. Shutdown can take a few minutes.

To restart the system, you must manually turn on the Steelhead appliance.

To reboot or shut down the system

1. Choose Configure > Maintenance > Reboot/Shutdown to display the Reboot/Shutdown page.

Figure 9-4. Reboot/Shutdown Page

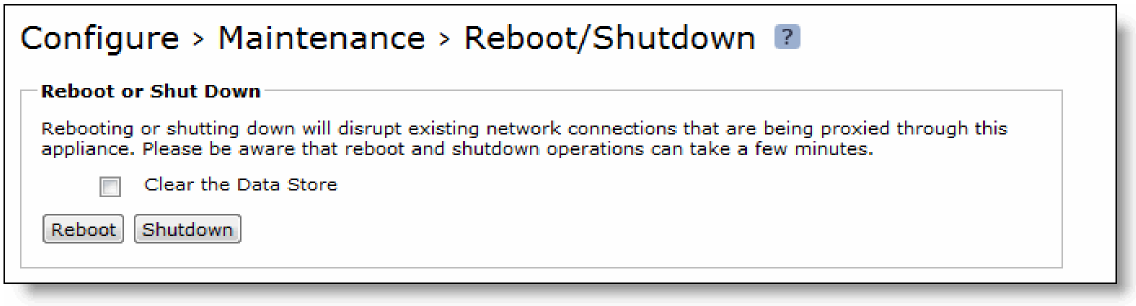

Configure > Maintenance > Reboot/Shutdown ?

Reboot or Shut Down

Rebooting or shutting down will disrupt existing network connections that are being proxied through this appliance. Please be aware that reboot and shutdown operations can take a few minutes.

□ Clear the Data Store

2. Click **Reboot**. After you click **Reboot**, you are logged out of the system and it is rebooted.

3. Click **Shutdown** to shut down the system. After you click **Shutdown**, the system is turned off. To restart the system, you must manually turn on the Steelhead appliance.

Tip: To remove data from the datastore, click **Clear the Data Store**.

Managing Licenses and Model Upgrades

This section describes how to install, update, and remove a license. It also describes how to use flexible licensing to manage model configurations and upgrades. It includes the following sections:

- "Flexible Licensing Overview" on page 350
- "Installing a License" on page 353

- "Upgrading a Model that Requires No Additional Hardware" on page 356
- "Upgrading a Model that Requires Additional Hardware" on page 357
- "Removing a License" on page 357

You perform all license management and Steelhead appliance model upgrades in the Configure > Maintenance > Licenses page.

Flexible Licensing Overview

RiOS provides a flexible way to manage Steelhead appliance licenses, model configurations, and upgrades. Rather than a performing an incremental model upgrade or replacing an appliance, RiOS provides *specification licenses* that configure specific performance characteristics of an appliance. A specification license points to a specific, validated model and includes the required license and the hardware specification. If a model upgrade requires additional hardware, the specification license determines which hardware is necessary to complete the upgrade.

By activating a specification license on an appliance you can transform the capabilities of the appliance to meet performance characteristics for any model within a platform family.

Figure 9-5. Use Specification Licenses to Scale a Steelhead Appliance within a Platform Family

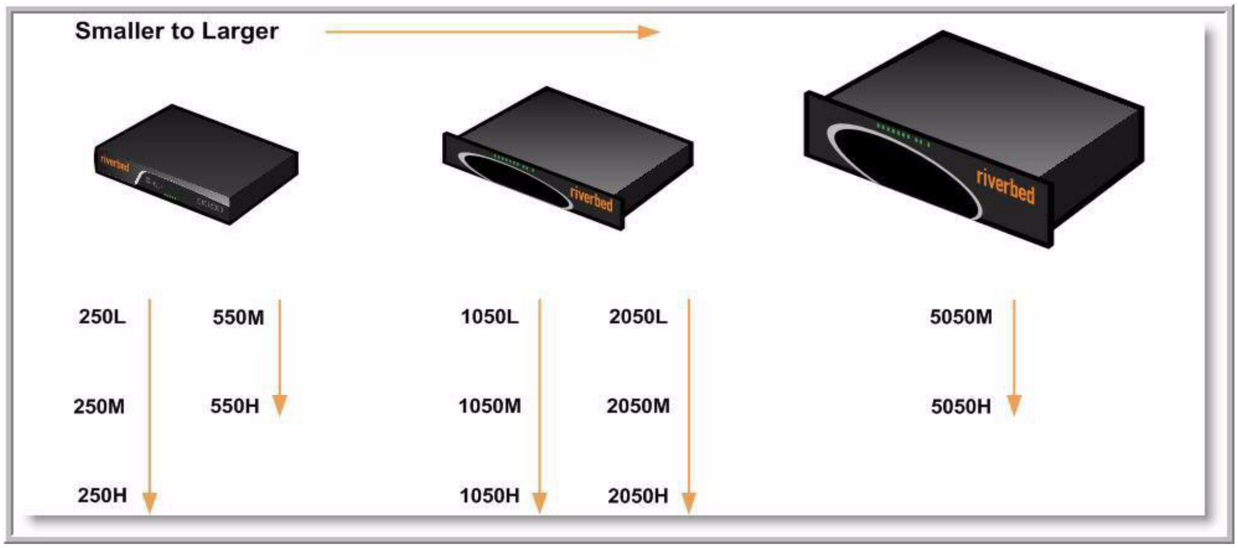

Some model upgrades require new hardware components, as listed in Figure 9-6. For example, to upgrade a model 1050L to a 1050H, you need to install an additional 250 GB of disk space and an additional 2 GB of memory. To accomplish this, order a hardware kit that contains the additional hardware from Riverbed Support or Sales.

After adding the required hardware and license to the Steelhead appliance, activate the hardware specification instead of replacing the appliance.

Figure 9-6. Flexible Licensing and Upgrade Possibilities by Appliance Model

Source Appliance Model	Destination Appliance Model	Upgrade Requires	Minimum RiOS Version	Impact on Datastore	Impact on Appliance Configuration	Reboot Required
250L	250M	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
250L	250H	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
250M	250H	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
550M	550H	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
1050U	1050L	License only	v5.5.9 v6.0.4 v6.1.2	None	None	No
1050L	1050LR	License and Hardware Kit	v5.0.9 v5.5.4	Clears: • datastore • RSP data • PFS • log files	None	Yes
1050L	1050M	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
1050L	1050H	1050H License and Hardware Kit	4.1.9d 5.0.7e 5.5.2d	Clears all data	None	Yes
1050M	1050H	1050H License and Hardware Kit	4.1.9d 5.0.7e 5.5.2d	Clears all data	None	Yes
1050M	1050MR	License and Hardware Kit	v5.0.9 v5.5.4	Clears: • datastore • RSP data • PFS • log files	None	Yes

Source Appliance Model	Destination Appliance Model	Upgrade Requires	Minimum RiOS Version	Impact on Datastore	Impact on Appliance Configuration	Reboot Required
1050H	1050HR	License and Hardware Kit	v5.0.9 v5.5.4	Clears: • datastore • RSP data • PFS • log files	None	Yes
1050LR	1050MR	License only	v5.0.9 v5.5.4	None	None	No
1050LR	1050HR	License and Hardware Kit	v5.0.9 v5.5.4	Clears: • datastore	None	No
1050MR	1050HR	License and Hardware Kit	v5.0.9 v5.5.4	Clears: • datastore	None	No
2050M	2050H	License only	4.1.9d 5.0.7e 5.5.2d	None	None	No
5050L	5050M	License only	v5.5.7 v6.0.2	None	None	No
5050L	5050H	License and Hardware Kit	v5.5.7 v6.0.2	Clears: • datastore • RSP data • PFS • log files	None	Yes
5050M	5050H	5050H License and Hardware Kit	4.1.9d 5.0.7e 5.5.2d	Clears all data in the datastore, and also clears: • all data in the RSP partition • all statistics • all log files	None	Yes

For details on hardware specifications that require hardware upgrades, see the *Upgrade and Maintenance Guide*.

For More Information

The following table describes where to find more information on flexible licensing tasks.

Task	See
Get a license and hardware kit.	Riverbed Support or Sales
Install a license.	"Installing a License" on page 353

Task	See
Update an expired license.	"Installing a License" on page 353
Remove a license.	"Removing a License" on page 357
Upgrade an appliance model without additional hardware.	"Upgrading a Model that Requires No Additional Hardware" on page 356
Upgrade an appliance model with additional hardware.	"Upgrading a Model that Requires Additional Hardware" on page 357 and the *Upgrade and Maintenance Guide*.

Installing a License

This section describes how to install a license on a Steelhead appliance after receiving it from Riverbed Support or Sales.

To install a license

1. Choose Configure > Maintenance > Licenses to display the Licenses page.

Figure 9-7. Licenses Page

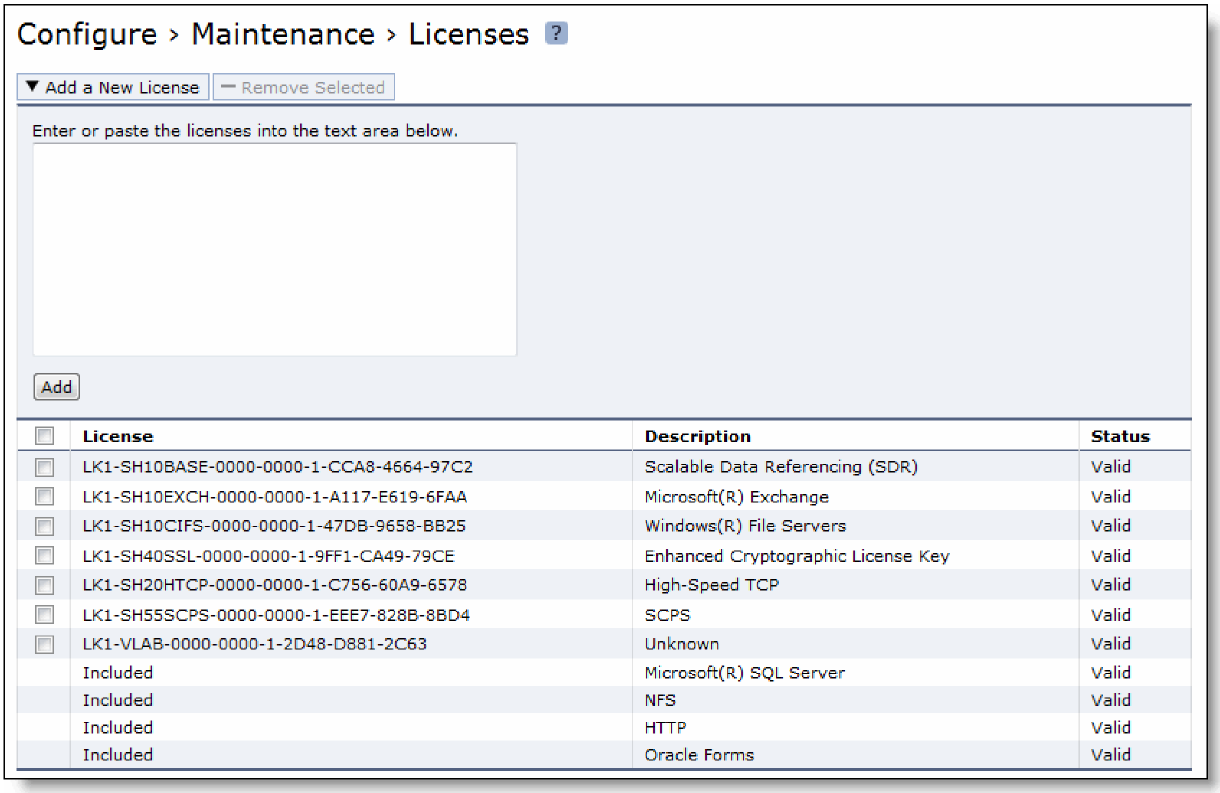

Enter or paste the licenses into the text area below.

□	License	Description	Status
□	LK1-SH10BASE-0000-0000-1-CCA8-4664-97C2	Scalable Data Referencing (SDR)	Valid
□	LK1-SH10EXCH-0000-0000-1-A117-E619-6FAA	Microsoft(R) Exchange	Valid
□	LK1-SH10CIFS-0000-0000-1-47DB-9658-BB25	Windows(R) File Servers	Valid
□	LK1-SH40SSL-0000-0000-1-9FF1-CA49-79CE	Enhanced Cryptographic License Key	Valid
□	LK1-SH20HTCP-0000-0000-1-C756-60A9-6578	High-Speed TCP	Valid
□	LK1-SH55SCPS-0000-0000-1-EEE7-828B-8BD4	SCPS	Valid
□	LK1-VLAB-0000-0000-1-2D48-D881-2C63	Unknown	Valid
	Included	Microsoft(R) SQL Server	Valid
	Included	NFS	Valid
	Included	HTTP	Valid
	Included	Oracle Forms	Valid

2. Complete the configuration as described in the following table.

Control	Description
Add a New License	Displays the controls to add a new license.
Licenses Text Box	Copy and paste the license key provided by Riverbed Support or Sales into the text box. **Tip:** Separate multiple license keys with a space, Tab, or Enter.
Add	Adds the license.

3. Click **Save** to save your settings permanently.

To install a license on a Cloud Steelhead

CSH This feature is only supported by the Cloud Steelhead.

1. Choose Configure > Maintenance > Licenses to display the Cloud Licensing page.

Figure 9-8. Cloud Licensing Page

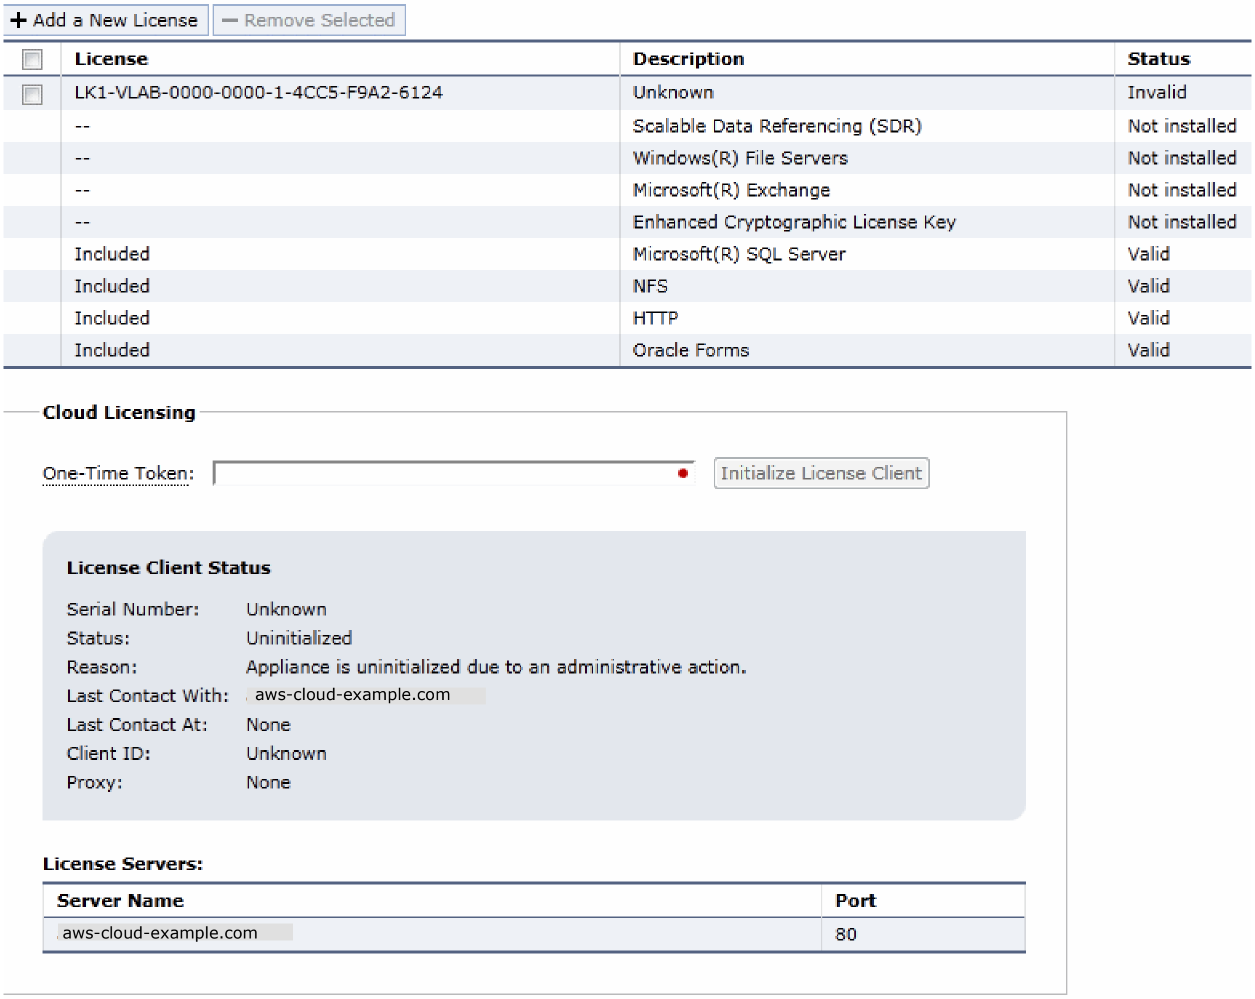

□	License	Description	Status
□	LK1-VLAB-0000-0000-1-4CC5-F9A2-6124	Unknown	Invalid
	--	Scalable Data Referencing (SDR)	Not installed
	--	Windows(R) File Servers	Not installed
	--	Microsoft(R) Exchange	Not installed
	--	Enhanced Cryptographic License Key	Not installed
	Included	Microsoft(R) SQL Server	Valid
	Included	NFS	Valid
	Included	HTTP	Valid
	Included	Oracle Forms	Valid

Cloud Licensing

One-Time Token:

License Client Status

Serial Number:	Unknown
Status:	Uninitialized
Reason:	Appliance is uninitialized due to an administrative action.
Last Contact With:	aws-cloud-example.com
Last Contact At:	None
Client ID:	Unknown
Proxy:	None

License Servers:

Server Name	Port
aws-cloud-example.com	80

2. Complete the configuration as described in the following table.

Control	Description
Add a New License	Displays the controls to add a new license.
Licenses Text Box	Copy and paste the license key provided by Riverbed Support or Sales into the text box. **Tip:** Separate multiple license keys with a space, Tab, or Enter.
Add	Adds the license.

3. Login to the Riverbed Cloud Portal and click the license to display the License Details page.

4. Copy the one-time token on the page into a text editor.

5. Return to the Cloud Licensing page in the Cloud Steelhead and under Cloud Licensing, paste the one-time token in the text field and click **Initialize License Client.**

   If the token is valid, the appliance receives a license. If not, the appliance displays the text field to specify the one-time token. Type the one-time token again and click **Re-Initialize License Client.**

   After the appliance receives a license, the Cloud Licensing section displays the **Refresh License** and **Remove License** buttons. Complete the configuration as described in the following table.

Control	Description
Refresh License	Forces the Cloud Steelhead license to retrieve license updates from the Riverbed Cloud Portal.
Remove License	Deletes the Cloud Steelhead license. As a result, the Cloud Steelhead reverts to the uninitialized state and the optimization service stops.

6. Click **Save** to save your settings permanently.

Model Upgrade Overview

You can use a hardware specification to upgrade a model without requiring remanufacturing from a CD-ROM or compact flash. Some model upgrades require additional hardware. When the appliance has the required hardware, activating the hardware specification upgrades the appliance to the new model number. When the existing hardware is not adequate, a hardware required message appears after the hardware specification description.

Next Steps

After installing a license, the next steps to complete a model upgrade depend on whether the upgrade requires additional hardware.

- If you do not need to add hardware to the Steelhead appliance, see “Upgrading a Model that Requires No Additional Hardware” on page 356.
- If you are upgrading a Steelhead appliance to a model that requires new hardware components, see “Upgrading a Model that Requires Additional Hardware” on page 357.

Upgrading a Model that Requires No Additional Hardware

This section describes how to activate a hardware specification that does not require additional hardware on a Steelhead appliance. Perform the following steps after installing the license.

To activate a hardware specification

1. Stop the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

2. Choose Configure > Maintenance > Licenses to display the Licenses page.

3. The hardware specifications appear at the bottom of the page. The hardware specification description includes the potential bandwidth and connection counts. The current specification appears in bold. Hardware specifications that will be available for the appliance model after the license or required hardware has been installed are included in the list but are dimmed.

Figure 9-9. Hardware Specifications Appear in the Licenses Page

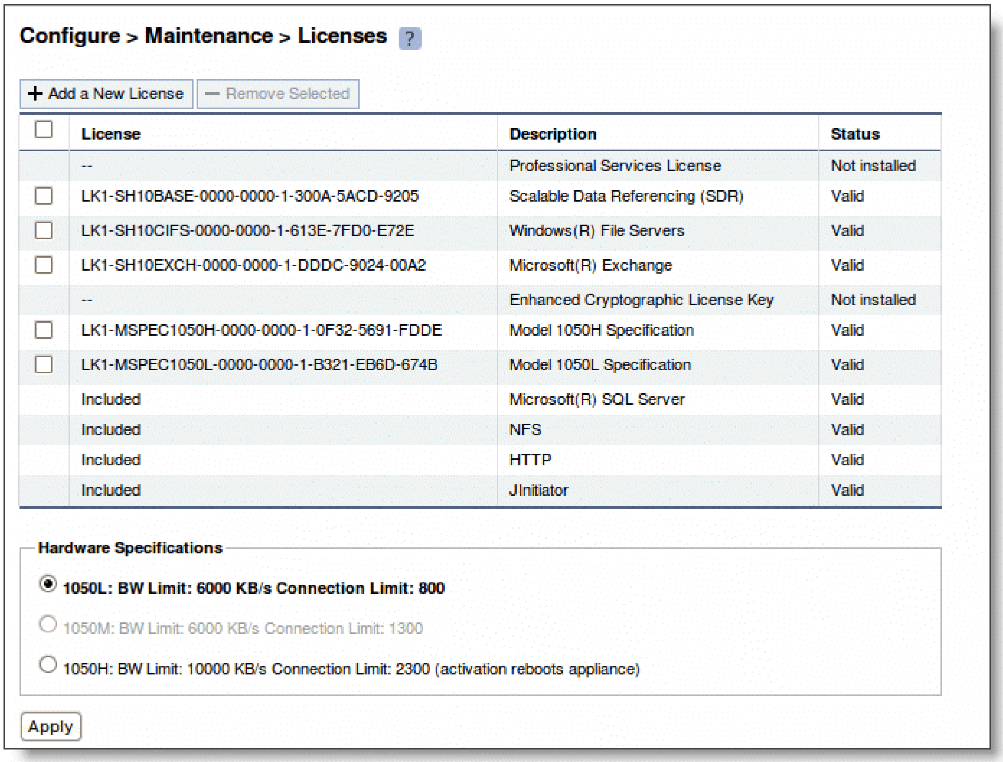

Configure > Maintenance > Licenses

☐	License	Description	Status
	--	Professional Services License	Not installed
☐	LK1-SH10BASE-0000-0000-1-300A-5ACD-9205	Scalable Data Referencing (SDR)	Valid
☐	LK1-SH10CIFS-0000-0000-1-613E-7FD0-E72E	Windows(R) File Servers	Valid
☐	LK1-SH10EXCH-0000-0000-1-DDDC-9024-00A2	Microsoft(R) Exchange	Valid
	--	Enhanced Cryptographic License Key	Not installed
☐	LK1-MSPEC1050H-0000-0000-1-0F32-5691-FDDE	Model 1050H Specification	Valid
☐	LK1-MSPEC1050L-0000-0000-1-B321-EB6D-674B	Model 1050L Specification	Valid
	Included	Microsoft(R) SQL Server	Valid
	Included	NFS	Valid
	Included	HTTP	Valid
	Included	JInitiator	Valid

Hardware Specifications

● **1050L: BW Limit: 6000 KB/s Connection Limit: 800**

○ 1050M: BW Limit: 6000 KB/s Connection Limit: 1300

○ 1050H: BW Limit: 10000 KB/s Connection Limit: 2300 (activation reboots appliance)

4. Select the hardware specification you wish to activate.

5. Click **Apply**.

6. Click **Restart** to restart the optimization service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Upgrading a Model that Requires Additional Hardware

This section describes how to activate a hardware specification that requires additional hardware on a Steelhead appliance. Perform the following steps after installing the license.

To upgrade a model

1. Use the hardware kit to upgrade the disk and memory of the appliance to the new model requirements. For details, see the *Upgrade and Maintenance Guide*.

2. Stop the Steelhead service. For details, see "Starting and Stopping the Optimization Service" on page 345.

3. Choose Configure > Maintenance > Licenses to display the Licenses page. The bottom of the screen lists the available hardware specifications. The current specification appears in bold. The hardware specification description includes the potential bandwidth and connection counts. Hardware specifications that will be available after the license or required hardware has been installed are included in the list but are dimmed.

4. Select the hardware specification you wish to activate. If a hardware specification requires an appliance reboot after activation, the message activation reboots appliance is displayed.

5. Click **Apply**. The appliance reboots and the optimization service restarts.

When the upgrade is complete, the appliance is transformed into the new model. The model number appears on the appliance banner in the upper-right corner of the screen. The appliance retains its original serial number.

Upgrade and Downgrade Limitations

The following limitations apply to Steelhead appliance upgrades and downgrades.

- After using flexible licensing to upgrade a Steelhead appliance, Riverbed does not recommend upgrading a model using a software version prior to RiOS v4.1.6x. Riverbed recommends that you use flexible licensing to perform all subsequent model upgrades on that appliance.
- After adding hardware to a Steelhead appliance and using flexible licensing to upgrade to a 1050H model, you cannot return the appliance to a 1050M or 1050L model.

Removing a License

Riverbed recommends that you keep old licenses in case you ever want to downgrade to an earlier software version; however, in some situations you might want to remove a license.

To remove a license

1. Choose Configure > Maintenance > Licenses to display the Licenses page.

2. Select the license you wish to delete.

3. Click **Remove Selected**.

4. Click **Save** to save your settings permanently.

Viewing Permissions

You can display your system permissions and add or change your login password in the Configure > My Account page.

To display system permissions

1. Choose Configure > My Account to display the My Account page.

Figure 9-10. My Account Page

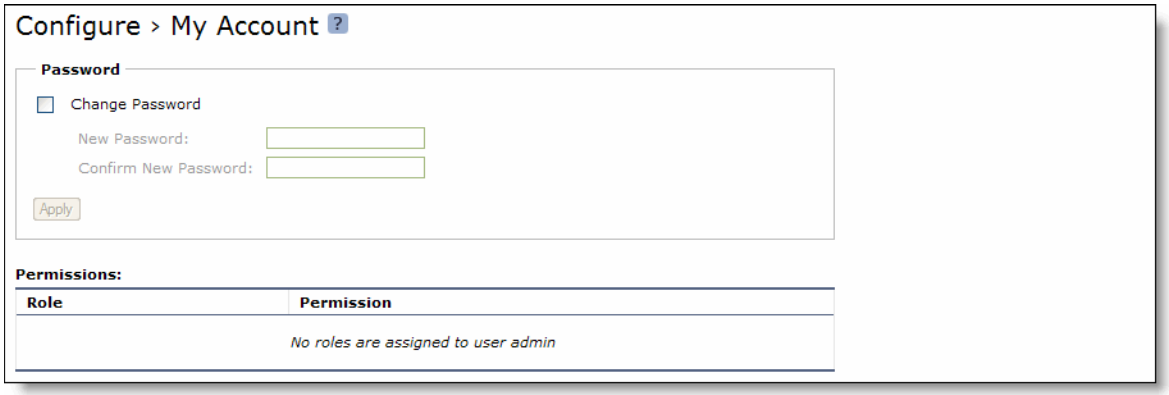

Configure > My Account

Password

□ Change Password

New Password:

Confirm New Password:

Apply

Permissions:

Role	Permission
No roles are assigned to user admin	

2. Under Change Password for <your user name>, complete the configuration as described in the following table.

Control	Description
Change Password	Enables a log in password.
Password/Password Confirm	Specify a password in the text box. The password must have a minimum of six characters. Confirm the password in the Password Confirm text box.

3. Click **Apply** to apply your changes to the running configuration.

   The permissions list displays the roles and permissions assigned to your user name.

Tip: For details on setting user permissions, see "Managing User Permissions" on page 363.

Managing Configuration Files

You can save, activate, import, and revert configurations in the Configure > Configurations page.

Each Steelhead appliance has an active, running configuration and a written, saved configuration.

When you **Apply** your settings in the Management Console, the values are applied to the active running configuration, but the values are not written to disk and saved permanently.

When you **Save** your configuration settings, the values are written to disk and saved permanently. They take effect after you restart the Steelhead service.

Each time you save your configuration settings, they are written to the current running configuration, and a backup is created. For example, if the running configuration is myconfig and you save it, myconfig is backed up to myconfig.bak and myconfig is overwritten with the current configuration settings.

The Configuration Manager is a utility that enables you to save configurations as backups or active configuration backups.

The Configuration Manager also includes an Import Configuration utility to support these common use cases:

- **Replacing a Steelhead appliance** - If you are replacing one Steelhead appliance for another, you can import all of the network information (although not the licenses) and disconnect the old Steelhead appliance before you switch configurations on the new Steelhead appliance.
- **Configuration template for a large deployment** - You can avoid entering the complete Steelhead appliance configuration for every appliance in a large deployment by setting up a template Steelhead appliance and importing template settings to the configuration list.

Important: Some configuration settings require that you restart the Steelhead service for the settings to take effect. For details on restarting the Steelhead service, see "Starting and Stopping the Optimization Service" on page 345.

To manage configurations

1. Choose Configure > Configurations to display the Configurations page.

Figure 9-11. Configurations Page

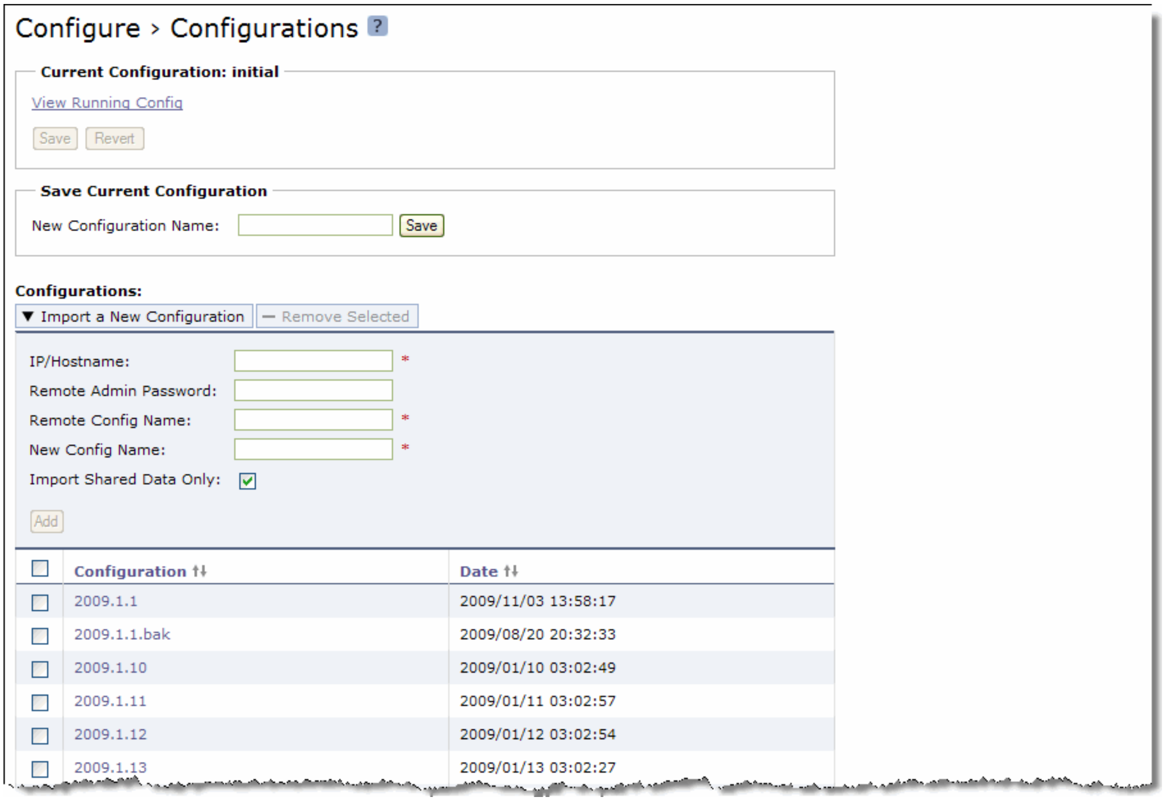

Configure > Configurations ?
View Running Config
Save
Revert
Save Current Configuration
Save
Configurations:
▼ Import a New Configuration
— Remove Selected
Add

Current Configuration: initial

New Configuration Name:

IP/Hostname: *

Remote Admin Password:

Remote Config Name: *

New Config Name: *

Import Shared Data Only: ☑

□	Configuration ⇅	Date ⇅
□	2009.1.1	2009/11/03 13:58:17
□	2009.1.1.bak	2009/08/20 20:32:33
□	2009.1.10	2009/01/10 03:02:49
□	2009.1.11	2009/01/11 03:02:57
□	2009.1.12	2009/01/12 03:02:54
□	2009.1.13	2009/01/13 03:02:27

2. Under Current Configuration: <filename>, complete the configuration as described in the following table.

Control	Description
Current Configuration: <configuration name>	**View Running Config** - Displays the running configuration settings in a new browser window.
	Save - Saves settings that have been applied to the running configuration.
	Revert - Reverts your settings to the running configuration.
Save Current Configuration	Specify a new filename to save settings that have been applied to the running configuration as a new file, and then click **Save**.

3. To import a configuration from another appliance, complete the configuration as described in the following table.

Control	Description
Import a New Configuration	Displays the controls to import a configuration from another appliance.
IP/Hostname	Specify the IP address or hostname of the Steelhead appliance from which you want to import the configuration.
Remote Admin Password	Specify the administrator password for the remote Steelhead appliance.
Remote Config Name	Specify the name of the configuration you want to import from the remote Steelhead appliance.
New Config Name	Specify a new, local configuration name.
Import Shared Data Only	This value is enabled by default. Copies only the following common settings: in-path and out-of-path interface, protocols, CLI and Web, statistics, NTP, SNMP, and alarm settings. The system does not automatically copy the other settings.
Add	Adds the configuration. The imported configuration appears in the Configuration list but does not become the active configuration until you click **Activate**.
Remove Selected	Select the check box next to the name and click **Remove Selected**.
Change Active Configuration	Select the configuration to activate from the drop-down list.

4. Click **Activate**.

5. Restart the Steelhead appliance service. For details, see "Starting and Stopping the Optimization Service" on page 345.

Tip: Select the configuration name to display the configuration settings in a new browser window.

Configuring General Security Settings

You can prioritize local, RADIUS, and TACACS+ authentication methods for the system and set the authorization policy and default user for RADIUS and TACACS+ authorization systems in the Configure > Security > General Settings page.

Important: Make sure to put the authentication methods in the order in which you want authentication to occur. If authorization fails on the first method, the next method is attempted, and so on, until all of the methods have been attempted.

Tip: To set TACACS+ authorization levels (admin or read-only) to allow certain members of a group to log in, add the following attribute to users on the TACACS+ server:

```
        service = rbt-exec {
              local-user-name = “monitor”
        }
```

where you replace monitor with admin for write access.

For details on setting up RADIUS and TACACS+ servers, see the *Riverbed Deployment Guide*.

To set general security settings

1. Choose Configure > Security > General Security Settings to display the General Security Settings page.

Figure 9-12. General Security Settings Page

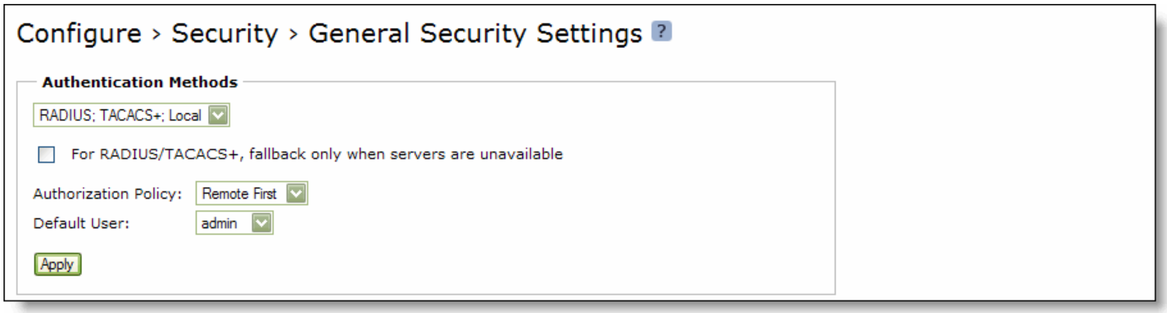

Configure > Security > General Security Settings

Authentication Methods

RADIUS; TACACS+; Local

□ For RADIUS/TACACS+, fallback only when servers are unavailable

Authorization Policy: Remote First

Default User: admin

Apply

2. Under Authentication Methods, complete the configuration as described in the following table.

Control	Description
Authentication Methods	Specifies the authentication method. Select an authentication method from the drop-down list. The methods are listed in the order in which they occur. If authorization fails on the first method, the next method is attempted, and so on, until all of the methods have been attempted.
For RADIUS/TACACS+, fallback only when servers are unavailable.	Specifies that the Steelhead appliance falls back to a RADIUS or TACACS+ server only when all other servers do not respond. This is the default setting. When this feature is disabled, the Steelhead appliance does not fall back to the RADIUS or TACACS+ servers. If it exhausts the other servers and does not get a response, it returns a server failure.
Authorization Policy	Appears only for some Authentication Methods. Optionally, select one of the following policies from the drop-down list: • **Remote First** - Check the remote server first for an authentication policy, and only check locally if the remote server does not have one set. This is the default behavior. • **Remote Only** - Only checks the remote server. • **Local Only** - Only checks the local server. All remote users are mapped to the user specified. Any vendor attributes received by an authentication server are ignored. **Default User** - Optionally, select Admin or Monitor from the drop-down list to define the default authentication policy.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Managing User Permissions

You can change the administrator or monitor passwords and define role-based users in the Configure > Security > User Permissions page.

Capability-Based Accounts

The system has two accounts based on what actions the user can take:

- **Admin** - The administrator user has full privileges. For example, as an administrator you can set and modify configuration settings, add and delete users, restart the Steelhead service, reboot the Steelhead appliance, and create and view performance and system reports.
- **Monitor** - A monitor user can view reports, user logs, and change their password. A monitor user cannot make configuration changes, modify private keys, view logs, or manage cryptographic modules in the system.

Role-Based Accounts

You can also create users, assign passwords to the user, and assign varying configuration roles to the user. A user role determines whether the user has permission to:

- **Read-only** - With read-only privileges you can view current configuration settings but you cannot change them.
- **Read/Write** - With read and write privileges you can view settings and make configuration changes for a feature.
- **Deny** - With deny privileges you cannot view settings or save configuration changes for a feature.

For example, you might have user Jane who can make configuration changes to QoS, PFS, and SSL whereas user John can only view these configuration settings; and finally, user Joe cannot view, change, or save the settings for these features.

Available menu items reflect the privileges of the user. For example, any menu items that a user does not have permission to use are unavailable. When a user selects an unavailable link, the User Permissions page appears.

Important: The Optimization Services role in RiOS v6.1 and later includes permission to access High-Speed TCP (HS-TCP). Consequently, upgrading to RiOS v6.1 or later causes all role-based users with permission for the High-Speed TCP (HS-TCP) role to lose their access. The Administrator must manually re-assign those users who require HS-TCP access with permission for the Optimization Services (GUI) or the Acceleration services (CLI) role. Alternatively, the administrator can create a custom role for the HS-TCP users.

To set the administrator or monitor password

1. Choose Configure > Security > User Permissions to display the User Permissions page.

Figure 9-13. User Permissions Page

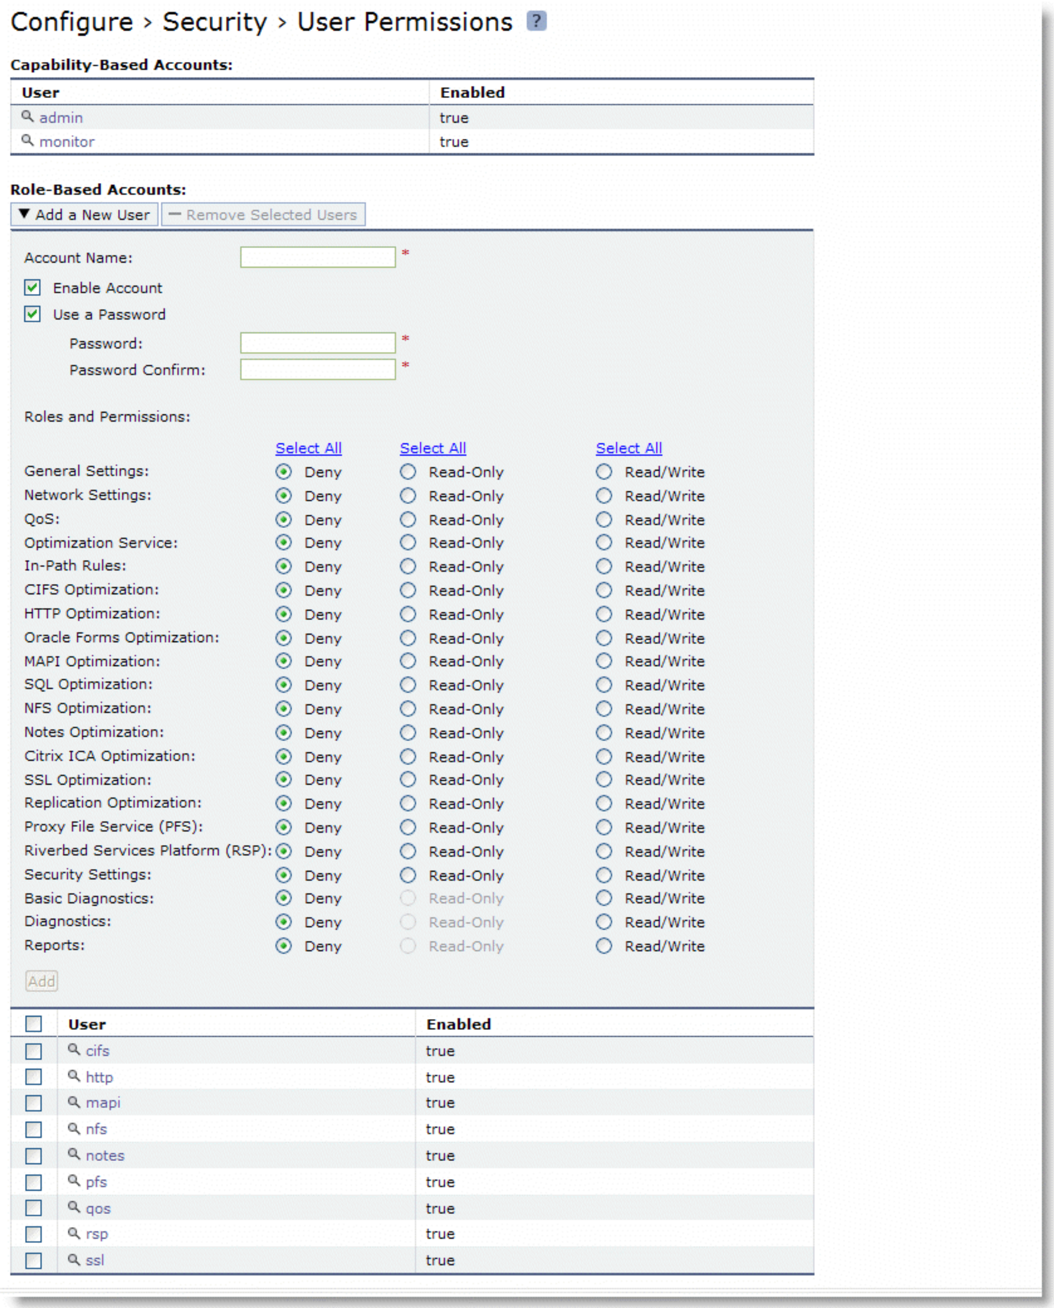

Configure > Security > User Permissions

Capability-Based Accounts:

User	Enabled
admin	true
monitor	true

Role-Based Accounts:

Account Name: *

☑ Enable Account

☑ Use a Password

Password: *

Password Confirm: *

Roles and Permissions:

	Select All	Select All	Select All
General Settings:	● Deny	○ Read-Only	○ Read/Write
Network Settings:	● Deny	○ Read-Only	○ Read/Write
QoS:	● Deny	○ Read-Only	○ Read/Write
Optimization Service:	● Deny	○ Read-Only	○ Read/Write
In-Path Rules:	● Deny	○ Read-Only	○ Read/Write
CIFS Optimization:	● Deny	○ Read-Only	○ Read/Write
HTTP Optimization:	● Deny	○ Read-Only	○ Read/Write
Oracle Forms Optimization:	● Deny	○ Read-Only	○ Read/Write
MAPI Optimization:	● Deny	○ Read-Only	○ Read/Write
SQL Optimization:	● Deny	○ Read-Only	○ Read/Write
NFS Optimization:	● Deny	○ Read-Only	○ Read/Write
Notes Optimization:	● Deny	○ Read-Only	○ Read/Write
Citrix ICA Optimization:	● Deny	○ Read-Only	○ Read/Write
SSL Optimization:	● Deny	○ Read-Only	○ Read/Write
Replication Optimization:	● Deny	○ Read-Only	○ Read/Write
Proxy File Service (PFS):	● Deny	○ Read-Only	○ Read/Write
Riverbed Services Platform (RSP):	● Deny	○ Read-Only	○ Read/Write
Security Settings:	● Deny	○ Read-Only	○ Read/Write
Basic Diagnostics:	● Deny	○ Read-Only	○ Read/Write
Diagnostics:	● Deny	○ Read-Only	○ Read/Write
Reports:	● Deny	○ Read-Only	○ Read/Write

□	User	Enabled
□	cifs	true
□	http	true
□	mapi	true
□	nfs	true
□	notes	true
□	pfs	true
□	qos	true
□	rsp	true
□	ssl	true

2. Under Capability-Based Accounts, complete the configuration as described in the following table.

Control	Description
admin/monitor	Click the magnifying glass to change the administrator or monitor password.
	Enable Account - Select to enable or clear to disable the administrator or monitor account.
	Use a Password - Enables password protection.
	Password - Specify a password in the text box. The password must have a minimum of six characters.
	Password Confirm - Confirm the new administrator password.

Important: A Role-based account cannot modify another role-based or capability account.

3. Under Role-Based Accounts, complete the configuration as described in the following table.

Control	Description
Add a New User	Click to display the controls for creating a new role-based account.
Account Name	Specify a name for the role-based account.
Enable Account	Select the check box to enable the new role-based account.
Use a Password	Select the check box to enable password protection and specify the following: • **Password** - Specify a password in the text box. The password must have a minimum of six characters. • **Password Confirm** - Specify the new password again for confirmation.
General Settings	Configures per-source IP connection limit and the maximum connection pooling size.
Network Settings	Configures host and network interface settings, including DNS cache settings and hardware assist rules.
QoS	Enforces QoS policies.
Optimization Service	Configures alarms, performance features, and HS-TCP and TCP optimization.
In-Path Rules	Configures TCP traffic for optimization and how to optimize traffic by setting in-path rules. This role includes WAN visibility to preserve TCP/IP address or port information. For details about WAN visibility, see the *Riverbed Deployment Guide*.
CIFS Optimization	Configures CIFS optimization settings (including SMB-Signing) and Overlapping Open optimization.
HTTP Optimization	Configures enhanced HTTP optimization settings: URL learning, Parse and Prefetch, Object Prefetch Table, keep-alive, insert cookie, file extensions to prefetch, and the ability to set up HTTP optimization for a specific server subnet.
Oracle Forms Optimization	Optimizes Oracle E-business application content and forms applications.
MAPI Optimization	Optimizes MAPI, and sets Exchange and NSPI ports.
SQL Optimization	Configures MS-SQL optimization.

Control	Description
NFS Optimization	Configures NFS optimization.
Notes Optimization	Configures Lotus Notes optimization.
Citrix ICA Optimization	Configures Citrix ICA optimization.
SSL Optimization	Configures SSL support and the secure inner channel.
Replication Optimization	Configures the SRDF/A and FCIP storage optimization modules.
Proxy File Service	Configures the PFS.
Riverbed Services Platform (RSP)	Adds functionality into a virtualized environment on the client Steelhead appliance. The functionality can include third-party packages such as a firewall security package, a streaming video server, or a package that provides core networking services (for example, DNS and DHCP). This role includes permission to install VMware tools and add subnet side rules. For details, see the *RSP User's Guide*.
Security Settings	Configures security settings, including RADIUS and TACACS authentication settings and the secure vault password.
Basic Diagnostics	Customizes system diagnostic logs, but does not include TCP dumps.
Diagnostics	Customizes system diagnostic logs, including system and user log settings.
Reports	Sets system report parameters.
Add	Adds your settings to the system.
Remove Selected Users	Select the check box next to the name and click **Remove Selected**.

4. Click **Save** to save your settings permanently.

Setting RADIUS Servers

You set up RADIUS server authentication in the Configure > Security > RADIUS page.

RADIUS is an access control protocol that uses a challenge and response method for authenticating users. Setting up RADIUS server authentication is optional.

For details on setting up RADIUS and TACACS+ servers, see the *Riverbed Deployment Guide*.

Enabling this feature is *optional*.

To set RADIUS server authentication

1. Choose Configure > Security > RADIUS to display the RADIUS page.

Figure 9-14. RADIUS Page

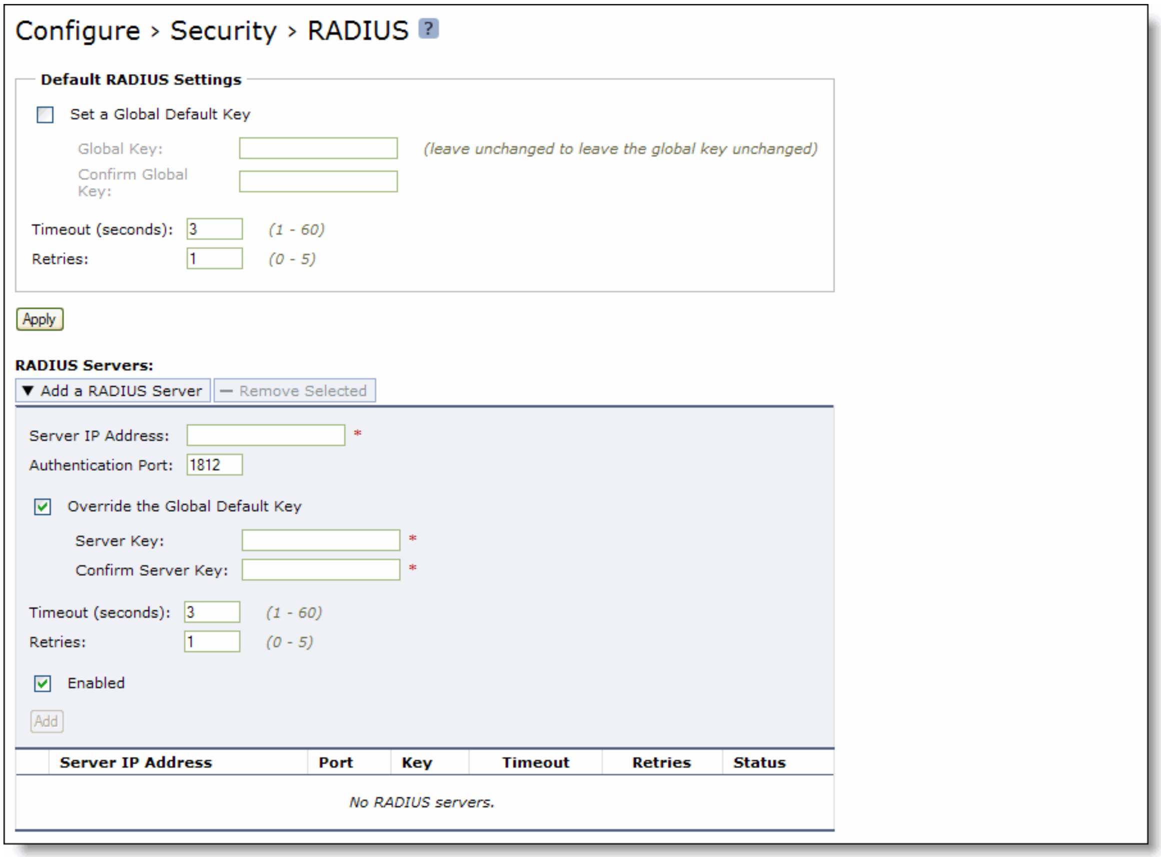

Configure > Security > RADIUS

Default RADIUS Settings

□ Set a Global Default Key

Global Key: (leave unchanged to leave the global key unchanged)

Confirm Global Key:

Timeout (seconds): 3 (1 - 60)

Retries: 1 (0 - 5)

RADIUS Servers:

Server IP Address: *

Authentication Port: 1812

☑ Override the Global Default Key

Server Key: *

Confirm Server Key: *

Timeout (seconds): 3 (1 - 60)

Retries: 1 (0 - 5)

☑ Enabled

	Server IP Address	Port	Key	Timeout	Retries	Status
	No RADIUS servers.					

2. Under Default RADIUS Settings, complete the configuration as described in the following table.

Control	Description
Set a Global Default Key	Enables a global server key for the RADIUS server.
Global Key	Specify the global server key.
Confirm Global Key	Confirm the global server key.
Timeout	Specify the time-out period in seconds (1-60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. The default value is 1.

3. Click **Apply** to apply your changes to the running configuration.

4. To add a new RADIUS server, complete the configuration as described in the following table.

Control	Description
Add a RADIUS Server	Displays the controls for defining a new RADIUS server.
Server IP Address	Specify the server IP address.
Authentication Port	Specify the port for the server.
Override the Global Default Key	Overrides the global server key for the server.
	Server Key - Specify the override server key.
	Confirm Server Key - Confirm the override server key.
Timeout	Specify the time-out period in seconds (1 - 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are 0-5. The default value is 1.
Enabled	Enables the new server.
Add	Adds the RADIUS server to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

Note: If you add a new server to your network and you do not specify these fields at that time, the global settings are applied automatically.

5. Click **Save** to save your settings permanently.

Related Topic

- "Configuring General Security Settings" on page 361

Configuring TACACS+ Access

You set up TACACS+ server authentication in the Configure > Security > TACACS+ page.

Enabling this feature is *optional*.

TACACS+ is an authentication protocol that allows a remote access server to forward a login password for a user to an authentication server to determine whether access is allowed to a given system.

For details on configuring RADIUS and TACACS+ servers to accept login requests from the Steelhead appliance, see the *Riverbed Deployment Guide*.

To set a TACACS+ server

1. Choose Configure > Security > TACACS+ to display the TACACS+ page.

Figure 9-15. TACACS+ Page

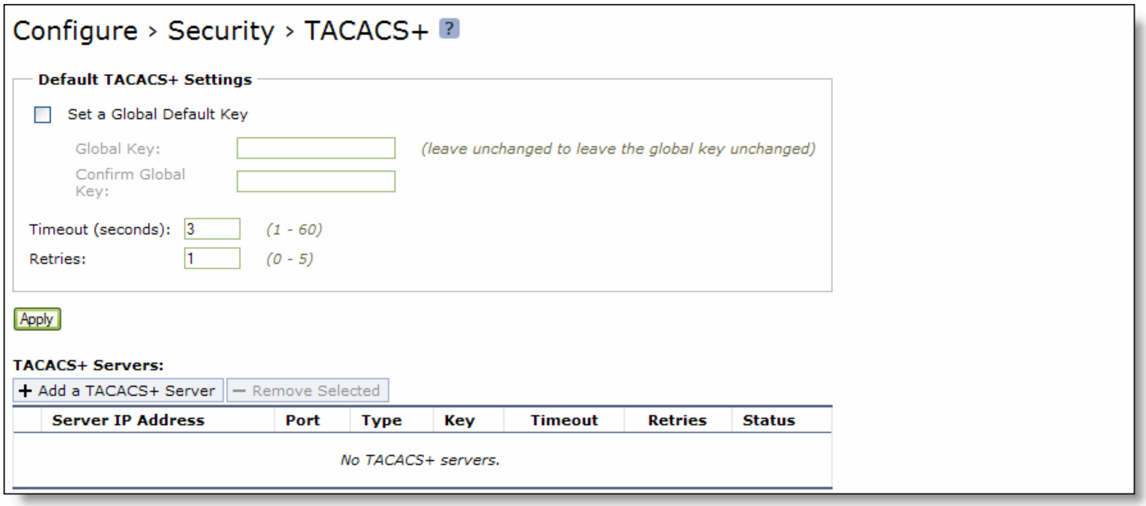

Default TACACS+ Settings

□ Set a Global Default Key

Global Key: (leave unchanged to leave the global key unchanged)

Confirm Global Key:

Timeout (seconds): 3 (1 - 60)

Retries: 1 (0 - 5)

TACACS+ Servers:

	Server IP Address	Port	Type	Key	Timeout	Retries	Status
	No TACACS+ servers.						

2. Under Default TACACS+ Settings, complete the configuration as described in the following table.

Control	Description
Set a Global Default Key	Specify this option to enable a global server key for the server.
Global Key	Specify the global server key.
Confirm Global Key	Confirms the global server key.
Timeout	Specify the time-out period in seconds (1 - 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are 0-5. The default is 1.

3. Click **Apply** to apply your changes to the running configuration.

4. To add or remove a TACACS+ server, complete the configuration as described in the following table.

Control	Description
Add a TACACS+ Server	Displays the controls for defining a new TACACS+ server.
Server IP Address	Specify the server IP address.
Authentication Port	Specify the port for the server. The default value is 49.
Authentication Type	Select either PAP or ASCII as the authentication type.
Override the Global Default Key	Specify this option to override the global server key for the server.
Server Key	Specify the override server key.
Confirm Server Key	Confirm the override server key.
Timeout	Specify the time-out period in seconds (1-60). The default is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are 0-5. The default is 1.
Enabled	Enables the new server.
Add	Adds the TACACS+ server to the list.
Remove Selected	Select the check box next to the name and click **Remove Selected**.

Note: If you add a new server to your network and you do not specify these fields, the system automatically applies the default settings.

5. Click **Save** to save your settings permanently.

Related Topic

- "Configuring General Security Settings" on page 361

Unlocking the Secure Vault

You can unlock and change the password for the secure vault in the Configure > Security > Secure Vault page.

The secure vault contains sensitive information from your Steelhead appliance configuration, including SSL private keys and the datastore encryption key. These configuration settings are encrypted on the disk at all times using AES 256-bit encryption.

Initially the secure vault is keyed with a default password known only to the RiOS software. This allows the Steelhead appliance to automatically unlock the vault during system start up. You can change the password, but the secure vault does not automatically unlock on start up. To optimize SSL connections or to use datastore encryption, the secure vault must be unlocked.

To unlock or change the password of the secure vault

1. Choose Configure > Security > Secure Vault to display the Secure Vault page.

Figure 9-16. Secure Vault Page

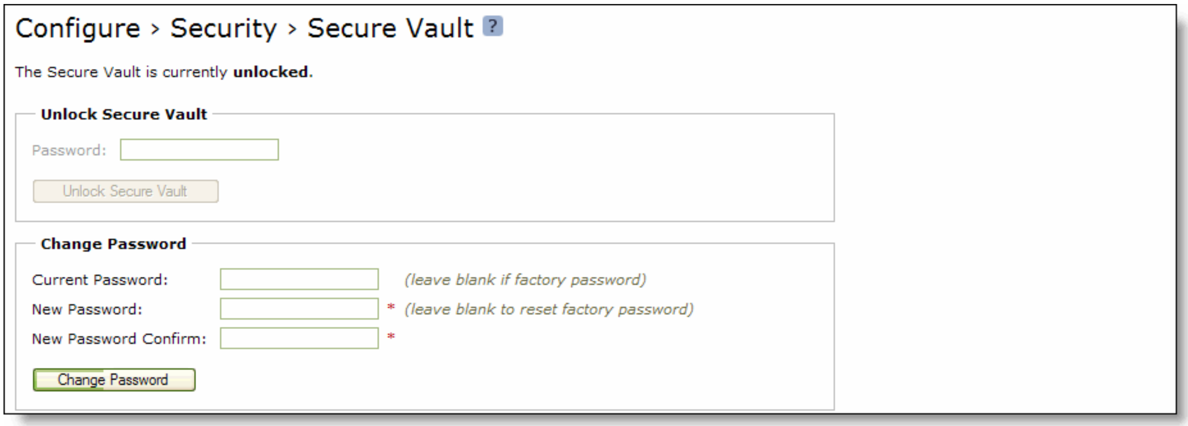

Configure > Security > Secure Vault

The Secure Vault is currently **unlocked**.

Unlock Secure Vault

Password:

Unlock Secure Vault

Change Password

Current Password: *(leave blank if factory password)*

New Password: * *(leave blank to reset factory password)*

New Password Confirm: *

Change Password

2. Under Unlock Secure Vault, complete the configuration as described in the following table.

Control	Description
Password	Specify a password and click **Unlock Secure Vault**. Initially the secure vault is keyed with a default password known only to the RiOS software. This allows the Steelhead appliance to automatically unlock the vault during system start up. You can change the password, but the secure vault does not automatically unlock on start up. To optimize SSL connections or to use datastore encryption, you must unlock the secure vault.
Unlock Secure Vault	Unlocks the vault.

3. Under Change Secure Vault Password, complete the configuration as described in the following table.

Control	Description
Current Password	Specify the current password. If you are changing the default password that ships with the product, leave the text box blank.
New Password	Specify a new password for the secure vault.
New Password Confirm	Confirm the new password for the secure vault.
Change Password	Changes the password for the secure vault.

4. Click **Save** to save your settings permanently.

Related Topic

- "Configuring General Security Settings" on page 361

Configuring a Management ACL

You can secure access to a Steelhead appliance using an internal management Access Control List (ACL) in the Configure > Security > Management ACL page.

Steelhead appliances are subject to the network policies defined by a corporate security policy, particularly in large networks. Using an internal management ACL, you can:

- restrict access to certain interfaces or protocols of a Steelhead appliance.
- restrict inbound IP access to a Steelhead appliance, protecting it from access by hosts that do not have permission without using a separate device (such as a router or firewall).
- specify which hosts or groups of hosts can access and manage a Steelhead appliance by IP address, simplifying the integration of Steelhead appliances into your network.

The Management ACL provides the following safeguards to prevent accidental disconnection from the Steelhead appliance (or the CMC):

- It detects the IP address you are connecting from and displays a warning if you add a rule that denies connections to that address.
- It always allows the default Steelhead appliance ports 7800, 7801, 7810, 7820, and 7850.
- It always allows a previously-connected CMC to connect and tracks any changes to the IP address of the CMC to prevent disconnection.
- It converts well-known port and protocol combinations such as SSH, Telnet, HTTP, HTTPS, SNMP, and SOAP into their default management service and protects these services from disconnection. For example, if you specify protocol 6 (TCP) and port 22, the management ACL converts this port and protocol combination into SSH and protects it from denial.
- It tracks changes to default service ports and automatically updates any references to changed ports in the access rules.

To set up a management ACL

1. Choose Configure > Security > Management ACL to display the Management ACL page.

Figure 9-17. Management ACL Page

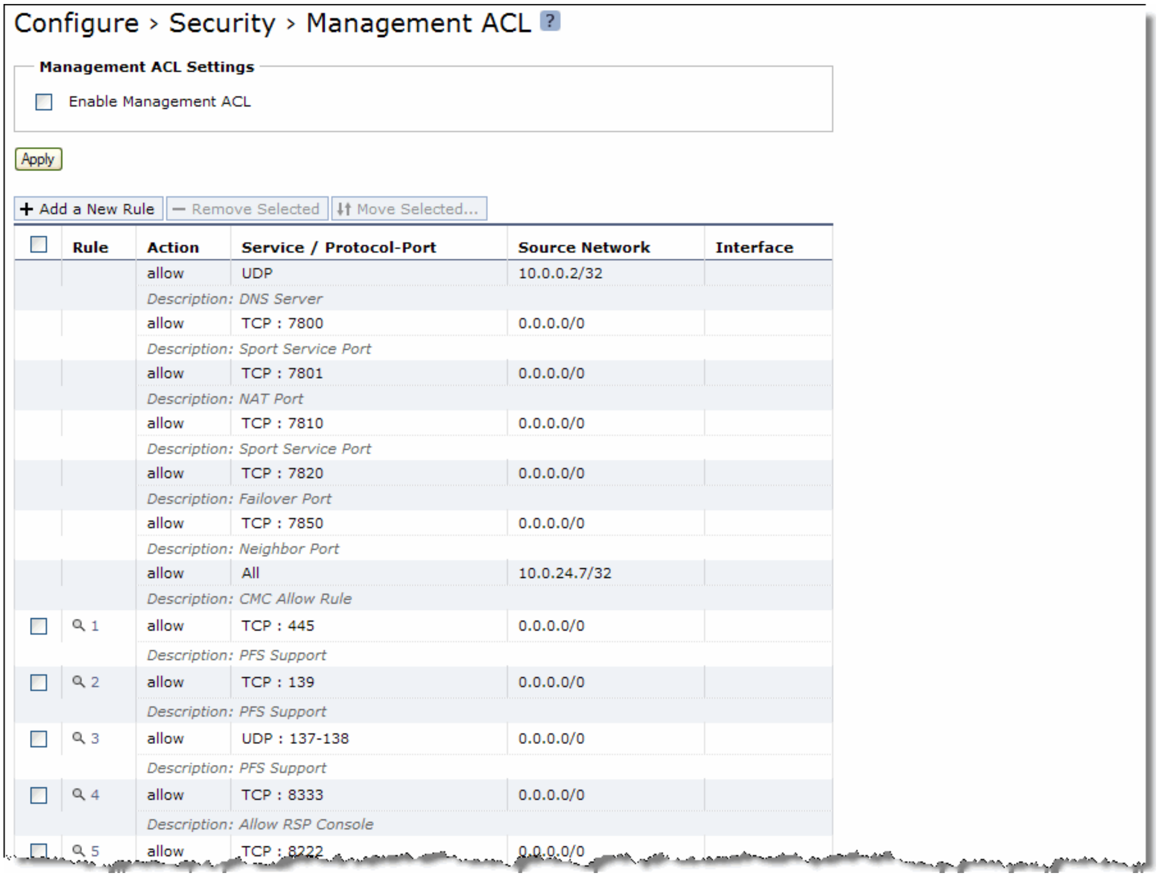

Configure > Security > Management ACL ?

Management ACL Settings

□ Enable Management ACL

Apply

+ Add a New Rule | — Remove Selected | Move Selected...

□	Rule	Action	Service / Protocol-Port	Source Network	Interface
		allow	UDP	10.0.0.2/32	
		Description: DNS Server			
		allow	TCP : 7800	0.0.0.0/0	
		Description: Sport Service Port			
		allow	TCP : 7801	0.0.0.0/0	
		Description: NAT Port			
		allow	TCP : 7810	0.0.0.0/0	
		Description: Sport Service Port			
		allow	TCP : 7820	0.0.0.0/0	
		Description: Failover Port			
		allow	TCP : 7850	0.0.0.0/0	
		Description: Neighbor Port			
		allow	All	10.0.24.7/32	
		Description: CMC Allow Rule			
□	1	allow	TCP : 445	0.0.0.0/0	
		Description: PFS Support			
□	2	allow	TCP : 139	0.0.0.0/0	
		Description: PFS Support			
□	3	allow	UDP : 137-138	0.0.0.0/0	
		Description: PFS Support			
□	4	allow	TCP : 8333	0.0.0.0/0	
		Description: Allow RSP Console			
□	5	allow	TCP : 8222	0.0.0.0/0	

2. Under Management ACL Settings, complete the configuration as described in the following table.

Control	Description
Enable Management ACL	Secures access to a Steelhead appliance using a management ACL.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Important: If you add, delete, edit, or move a rule that could disconnect connections to the Steelhead appliance, a warning message appears. Click **Confirm** to override the warning and allow the rule definition anyway. Use caution when overriding a disconnect warning.

ACL Management Rules

The management ACL contains rules that define a match condition for an inbound IP packet. You set a rule to allow or deny access to a matching inbound IP packet. When you add a rule on a Steelhead appliance, the destination specifies the Steelhead appliance itself, and the source specifies a remote host.

The ACL rules list contains default rules that allow you to use the management ACL with the RiOS features PFS, DNS caching, and RSP. These default rules allow access to certain ports required by these features. The list also includes a default rule that allows access to the CMC. If you delete the default ACL rules for one of these features and need to restore it, see “Restoring Default Access Rules” on page 376.

To add an ACL management rule

1. Under Management ACL Settings, complete the configuration as described in the following table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Action	Select one of the following rule types from the drop-down list: • **Allow** - Allows a matching packet access to the Steelhead appliance. This is the default action. • **Deny** - Denies access to any matching packets.
Service	Optionally, select Specify Protocol, or HTTP, HTTPS, SOAP, SNMP, SSH, Telnet. When specified, the Destination Port is dimmed and unavailable.
Protocol	(Appears only when Service is set to Specify Protocol.) Optionally, select All, TCP, UDP, or ICMP from the drop-down list. The default setting is All. When set to All or ICMP, the Service and Destination Ports are dimmed and unavailable.
Source Network	Optionally, specify the source subnet of the inbound packet; for example, 1.2.3.0/24.
Destination Port	Optionally, specify the destination port of the inbound packet, either a single port value or a port range of port1-port2, where port1 must be less than port2. Leave it blank to specify all ports.
Interface	Optionally, select an interface name from the drop-down list. Select All to specify all interfaces.
Description	Optionally, describe the rule to facilitate administration.
Rule Number	Optionally, select a rule number from the drop-down list. By default, the rule goes to the end of the table (just above the default rule). Steelhead appliances evaluate rules in numerical order starting with rule **1**. If the conditions set in the rule match, then the rule is applied, and the system moves on to the next packet. If the conditions set in the rule do not match, the system consults the next rule. For example, if the conditions of rule 1 do not match, rule 2 is consulted. If rule 2 matches the conditions, it is applied, and no further rules are consulted. **Note:** The default rule, Allow, which allows all remaining traffic from everywhere that has not been selected by another rule, cannot be removed and is always listed last.
Log Packets	Tracks denied packets in the log. By default, packet logging is enabled.
Add	Adds the rule to the list. The Management Console redisplays the Rules table and applies your modifications to the running configuration, which is stored in memory.
Remove Selected	Select the check box next to the name and click **Remove Selected**.
Move Selected	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

2. Click **Save** to save your settings permanently.

Usage Notes

- When you change the default port of services such as SSH, HTTP, HTTPS, on either the client or server-side Steelhead appliance and create a management ACL rule denying that service, the rule will not work as expected. The Steelhead appliance on the other end (either server or client) of an in-path deployment does not know that the default service port has changed, and consequently optimizes the packets to that service port. To work around this problem, add a pass-through rule to the client-side Steelhead appliance for the management interfaces. The pass-through rule prevents the traffic from coming from the local host when optimized.
- A management ACL rule that denies access from port 20 on the server-side Steelhead appliance in an out-of-path deployment prevents data transfer using active FTP. In this deployment, the FTP server and client cannot establish a data connection because the FTP server initiates the SYN packet and the management rule on the server-side Steelhead appliance blocks the SYN packet. To work around this problem:
  - use passive FTP instead of active FTP. With passive FTP, the FTP client initiates both connections to the server. For details on active and passive FTP, see "QoS Classification for the FTP Data Channel" on page 271.

  —or—

  - add a rule to either allow source port 20 on the server-side Steelhead appliance or allow the IP address of the FTP server.

Restoring Default Access Rules

This section describes how to restore the default ACL rules for the RiOS features PFS, RSP, and DNS caching.

To restore the default ACL management rules for PFS

1. Under Management ACL Settings, add a PFS ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	TCP
Destination Port	445
Rule Number	1
Description	PFS Support

2. Click **Add**.

3. Under Management ACL Settings, add a second PFS ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	TCP
Destination Port	139

Property	Value
Rule Number	1
Description	PFS Support

4. Click **Add**.

5. Add a third PFS ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	UDP
Destination Port	137-138
Rule Number	1
Description	PFS Support

6. Click **Add**.

To restore the default ACL management rules for RSP

1. Under Management ACL Settings, add an RSP ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	TCP
Destination Port	8222
Rule Number	1
Description	Allow RSP Console

2. Click **Add**.

3. Under Management ACL Settings, add a second RSP ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	TCP
Destination Port	8333
Rule Number	1
Description	Allow RSP Console

4. Click **Add**.

To restore the default ACL management rules for DNS caching

1. Under Management ACL Settings, add a DNS Caching ACL rule with the following properties.

Property	Value
Type	Allow
Protocol	UDP
Destination Port	53
Rule Number	1
Description	DNS Caching

2. Click **Add**.

Related Topics

- "Configuring PFS" on page 161
- "Enabling DNS Caching" on page 175
- "Installing and Configuring RSP" on page 180

Configuring Web Settings

You can modify Management Console Web user interface and certificate settings in the Configure > Security > Web Settings page.

To modify Web settings

1. Choose Configure > Security > Web Settings to display the Web Settings page.

Figure 9-18. Web Settings Page

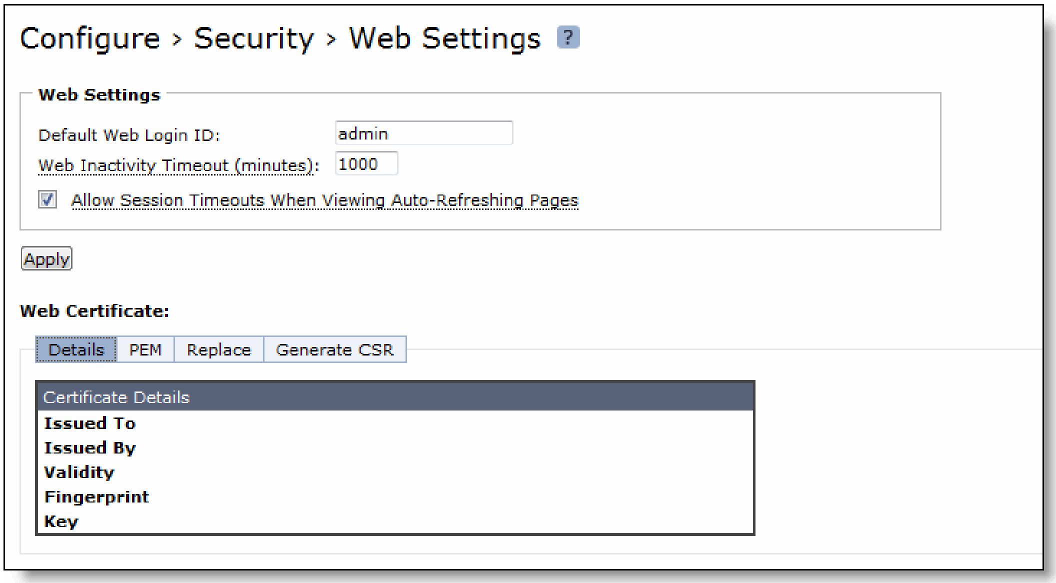

Configure > Security > Web Settings ?

Web Settings

Default Web Login ID: admin

Web Inactivity Timeout (minutes): 1000

☑ Allow Session Timeouts When Viewing Auto-Refreshing Pages

Web Certificate:

Certificate Details
Issued To
Issued By
Validity
Fingerprint
Key

2. Under Web Settings, complete the configuration as described in the following table.

Control	Description
Default Web Login ID	Specify the user name that appears in the authentication page. The default value is admin.
Web Inactivity Timeout	Specify the number of idle minutes before time-out. The default value is 15. A value of 0 disables time-out.
Allow Session Timeouts When Viewing Auto-Refreshing Pages	By default, session time-out is enabled, which stops the automatic updating of the report pages when the session times out. Clear the Allow box to disable the session time-out, remain logged-in indefinitely, and automatically refresh the report pages. **Important:** Disabling this feature poses a security risk.

3. Click **Apply** to apply your changes to the running configuration.

4. Click **Save** to save your settings permanently.

Managing Web SSL Certificates

RiOS v6.5 provides the following additional security features to manage SSL certificates used by the Steelhead appliance Management Console Web user interface using HTTPS.

- Generate the certificate and key pairs on the Steelhead appliance. This overwrites the existing certificate and key pair regardless of whether the previous certificate and key pair was self-signed or user added. The new self-signed certificate lasts for one year (365 days).
- Create certificate signing requests from the certificate and key pairs.
- Replace a signed certificate with one created by an administrator or generated by a 3rd party certificate authority.

To modify Web Certificates

1. Choose Configure > Optimization > Web Settings to display the Web Settings page.

2. Under Web Certificate, select the Details tab.

   The Steelhead identity certificate details appear, as described in the following table.

Control	Description
Issued To/Issued By	**Common Name** - Specifies the common name of the certificate authority.
	Organization - Specifies the organization name (for example, the company).
	Organization Unit - Specifies the organization unit name (for example, section or department).
	Locality - Specifies the city.
	State - Specifies the state.
	Country - Specifies the country.
	Serial Number - Specifies the serial number (Issued To, only).

Control	Description
Validity	**Issued On** - Specifies the date the certificate was issued.
	Expires On - Specifies the date the certificate expires.
Fingerprint	Specifies the SSL fingerprint.
Key	**Type** - Specifies the key type.
	Size - Specifies the size in bytes.

3. To replace an existing certificate, Under Web Certificate, select the Replace tab and complete the configuration as described in the following table.

Control	Description
Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)	Select this option if the existing private key and CA-signed certificate are located in one file. The page displays a Private Key and CA-Signed Public Certificate control for browsing to the key and certificate files or a text box for copying and pasting the key and certificate. **Note**: The private key is required.
	Local File - Browse to the local file.
	Text - Paste the text content of the file into the text box.
	Decryption Password - Specify the decryption password, if necessary.
	Set - Sets the peer.
Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)	Select this option if the existing private key and CA-signed certificate are located in two files. The page displays a Private Key and CA-Signed Public Certificate control for browsing to the key and certificate, or a text box for copying and pasting the key and certificate. **Note:** Importing the private key is optional.
Generate New Private Key and Self-Signed Public Certificate	Select this option to generate a new private key and self-signed public certificate.
	Cipher Bits - Select the key length from the drop-down list. The default value is 1024.
	Common Name (required) - Specify the hostname of the peer.
	Organization Name - Specify the organization name (for example, the company).
	Organization Unit Name - Specify the organization unit name (for example, the section or department).
	Locality - Specify the city.
	State (no abbreviations) - Specify the state.
	Country (2-letter code) - Specify the country (2-letter code only).
	Email Address - Specify the email address of the contact person.
	Validity Period (Days) - Specify how many days the certificate is valid. The default value is 730.

4. To generate a CSR, under Web Certificate, select the Generate CSR tab and complete the configuration as described in the following table.

Control	Description
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).
Locality	Specify the city.
State	Specify the state. Do not abbreviate.
Country	Specify the country (2-letter code only).
Email Address	Specify the email address of the contact person.
Generate CSR	Generates the Certificate Signing Request.

5. Click **Apply** to apply your changes to the running configuration.

6. Click **Save** to save your settings permanently.

CHAPTER 10 Viewing Reports and Logs

This chapter describes how to display system reports and user and system logs. It includes the following sections:

Networking Reports

Optimization Reports

Branch Services Reports

Diagnostic Reports

Viewing Current Connections

The Current Connections report displays the connections that are intercepted by the Steelhead appliance, including the connections that are passed through unoptimized.

The Current Connections report displays the following information:

- A summary of the connection numbers in various stages of optimization, pass-through, and forwarding.
- A summary of the optimized established connections sorted by type.
- An individual connections table, which shows more information about each connection. From this table, you can view more details about each connection and perform operations on it. For example, you can reset connections or send a keep-alive message to the outer remote machine for an optimized connection (the machine that is connected to the Steelhead appliance).

What This Report Tells You

The Current Connections report answers the following questions:

- How many connections are established?
- How many connections are half-open?
- How many connections are half-closed?
- What types of connections are established?
- How many connections are being passed through either intentionally or unintentionally?
- How many connections are being forwarded by a connection forwarding neighbor?
- How many connections are denied or discarded?

Viewing a Current Connections Summary

The summary table displays the total number of optimized and unoptimized connections and the total number of discarded or denied connections, with other statistics described in the following table.

Packet Type	Icon	Description
Established	▶▶▶	Specifies the total established, active connections.
Half-Open	▫▫▫	Specifies the total half-open active connections. A half-open connection is a TCP connection in which the connection has not been fully established. Half-open connections count toward the connection count limit on the Steelhead appliance because, at any time, they might become a fully opened connection. If you are experiencing a large number of half-open connections, consider a more appropriately sized Steelhead appliance.
Half-Closed	⇉⇇	Specifies the total half-closed active connections. Half-closed connections are connections which the Steelhead appliance has intercepted and optimized but are in the process of becoming inactive. These connections are counted toward the connection count limit on the Steelhead appliance. (Half-closed connections might remain if the client or server does not close their connections cleanly.) If you are experiencing a large number of half-closed connections, consider a more appropriately sized Steelhead appliance.
Passthrough	➡	Specifies the total number of connections that were passed through unoptimized. You can view and sort these connections by intentional and unintentional pass-through in the individual connections table that follows this summary.
Forwarded	➡⇨	Specifies the total number of connections that were forwarded when you have configured a connection forwarding neighbor to manage the connection. For details on connection forwarding, see "Configuring Connection Forwarding Features" on page 255.
Total Optimized		Specifies the total number of established, optimized connections plus the half-open and half-closed connections (where the half-open and half-closed connections are TCP connection states).
Total		Specifies the total number of connections intercepted by the Steelhead appliance, including the connections that are passed through unoptimized.
Discarded		Specifies the total number of discarded connections. Discarded packets for the connection that match a **Discard** rule are dropped silently. For details on discard rules, see "Configuring In-Path Rules" on page 28.
Denied		Specifies the total number of denied connections. (When packets for connections match a **Deny** rule, the appliance actively tries to reset the connection.) For details on discard rules, see "About In-Path Rules" on page 27.
Total Rejected		Specifies the total number of connections that were rejected by the system (that is, discarded and denied).

Note: A red x under the Type column indicates that the Steelhead appliance cannot establish the inner optimizing connection with its discovered peer.

Viewing Individual Connections

The individual connections table displays additional information about each connection. Because this report can list thousands of transient connections, the best way to narrow your search for a particular connection is to filter and sort the report.

This report does not allow auto-refresh because doing so could slow down the Steelhead appliance.

Select a column heading to sort the table by that column. For example, to sort the traffic by source IP address, click the **Source** column heading.

Column	Description
🔍	Click the magnifying glass to display the Current Connections Details report. See "Viewing the Current Connection Details" on page 393. Because this report is a snapshot in time, by the time you click the connection, it could be gone or in a different state. If the connection is no longer available, a message tells you that it cannot be found and might be closed.
Type ▶▶▶	**Established** - Indicates that the connection is established and active.
➡	**Intentional Passthrough** - Specifies the connection was intentionally passed through unoptimized.
⇒	**Unintentional Passthrough** - Specifies the connection was unintentionally passed through unoptimized.
▫▫▫	**Half-open** - Specifies the connection is half-open and active. A half-open connection is a TCP connection which has not been fully established.
⇉⇇	**Half-closed** - Specifies the connection is half-closed and active. A half-closed connection has been intercepted and optimized by the Steelhead appliance but is in the process of becoming inactive.
➡⇨	**Forwarded** - Specifies the connection is forwarded by the connection forwarding neighbor managing the connection. For details on connection forwarding, see "Configuring Connection Forwarding Features" on page 255.
	A red x indicates that the Steelhead appliance cannot establish the inner optimizing connection with its discovered peer.
Source:Port	Specifies the connection source IP address and port.
Destination:Port	Specifies the connection destination IP address and port.
Reduction	Specifies the percentage of reduction statistics for the connection.
LAN/KB WAN/KB	Specifies the amount of LAN or WAN throughput in KBs.
Data Start Time	Displays the time the connection was started. This column does not apply to pre-existing connections. Select the column heading to sort data start time in descending order.
Application	Specifies the application associated with the connection. For example: TCP, CIFS, MAPI, (e)MAPI-OA (encrypted MAPI Outlook Anywhere), or HTTP.

Column	Description
Notes	Displays connection icons which indicate the current state of the connection. The connection states can be one of the following:
⧉ / ⧉	**Compression Enabled** - Specifies that LZR compression is enabled**.**
⧉ / ⧉	**SDR Enabled** - Specifies that SDR optimization is enabled.
🔒/🔒	**Encryption Enabled** - Specifies that encryption is enabled on the secure inner channel (WAN). For details, see "Configuring Secure Peers" on page 231.
⚠ / ⚠	**Protocol Error** - Specifies a protocol error for one of the following conditions: • If you have the Optimize Connections with Security Signatures feature enabled (which prevents SMB signing), the Current Connections report displays a protocol error—this is an expected response. For details on preventing SMB signing, see "Configuring CIFS Optimization" on page 97. • If a problem occurs while optimizing encrypted MAPI traffic. For details on enabling optimization of encrypted MAPI traffic, see "Configuring MAPI Optimization" on page 123. • If a problem occurs with SSL optimization or the secure inner channel.

To view the Current Connections report

1. Choose Reports > Networking > Current Connections to display the Current Connections page.

Figure 10-1. Current Connections Page

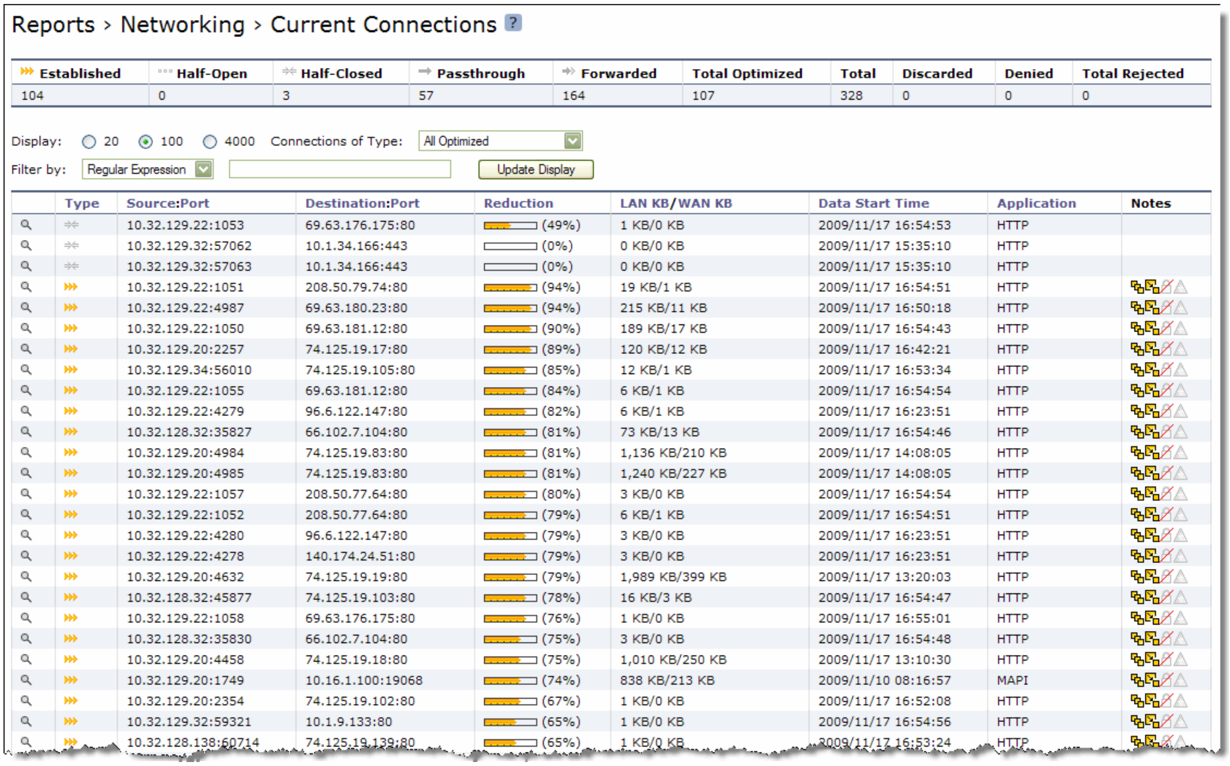

Reports > Networking > Current Connections

⋙ Established	Half-Open	Half-Closed	Passthrough	Forwarded	Total Optimized	Total	Discarded	Denied	Total Rejected
104	0	3	57	164	107	328	0	0	0

Display: ○ 20 ● 100 ○ 4000 Connections of Type: All Optimized

Filter by: Regular Expression

	Type	Source:Port	Destination:Port	Reduction	LAN KB/WAN KB	Data Start Time	Application	Notes
🔍	⇹	10.32.129.22:1053	69.63.176.175:80	(49%)	1 KB/0 KB	2009/11/17 16:54:53	HTTP	
🔍	⇹	10.32.129.32:57062	10.1.34.166:443	(0%)	0 KB/0 KB	2009/11/17 15:35:10	HTTP	
🔍	⇹	10.32.129.32:57063	10.1.34.166:443	(0%)	0 KB/0 KB	2009/11/17 15:35:10	HTTP	
🔍	⋙	10.32.129.22:1051	208.50.79.74:80	(94%)	19 KB/1 KB	2009/11/17 16:54:51	HTTP	△
🔍	⋙	10.32.129.22:4987	69.63.180.23:80	(94%)	215 KB/11 KB	2009/11/17 16:50:18	HTTP	△
🔍	⋙	10.32.129.22:1050	69.63.181.12:80	(90%)	189 KB/17 KB	2009/11/17 16:54:43	HTTP	△
🔍	⋙	10.32.129.20:2257	74.125.19.17:80	(89%)	120 KB/12 KB	2009/11/17 16:42:21	HTTP	△
🔍	⋙	10.32.129.34:56010	74.125.19.105:80	(85%)	12 KB/1 KB	2009/11/17 16:53:34	HTTP	△
🔍	⋙	10.32.129.22:1055	69.63.181.12:80	(84%)	6 KB/1 KB	2009/11/17 16:54:54	HTTP	△
🔍	⋙	10.32.129.22:4279	96.6.122.147:80	(82%)	6 KB/1 KB	2009/11/17 16:23:51	HTTP	△
🔍	⋙	10.32.128.32:35827	66.102.7.104:80	(81%)	73 KB/13 KB	2009/11/17 16:54:46	HTTP	△
🔍	⋙	10.32.129.20:4984	74.125.19.83:80	(81%)	1,136 KB/210 KB	2009/11/17 14:08:05	HTTP	△
🔍	⋙	10.32.129.20:4985	74.125.19.83:80	(81%)	1,240 KB/227 KB	2009/11/17 14:08:05	HTTP	△
🔍	⋙	10.32.129.22:1057	208.50.77.64:80	(80%)	3 KB/0 KB	2009/11/17 16:54:54	HTTP	△
🔍	⋙	10.32.129.22:1052	208.50.77.64:80	(79%)	6 KB/1 KB	2009/11/17 16:54:51	HTTP	△
🔍	⋙	10.32.129.22:4280	96.6.122.147:80	(79%)	3 KB/0 KB	2009/11/17 16:23:51	HTTP	△
🔍	⋙	10.32.129.22:4278	140.174.24.51:80	(79%)	3 KB/0 KB	2009/11/17 16:23:51	HTTP	△
🔍	⋙	10.32.129.20:4632	74.125.19.19:80	(79%)	1,989 KB/399 KB	2009/11/17 13:20:03	HTTP	△
🔍	⋙	10.32.128.32:45877	74.125.19.103:80	(78%)	16 KB/3 KB	2009/11/17 16:54:47	HTTP	△
🔍	⋙	10.32.129.22:1058	69.63.176.175:80	(76%)	1 KB/0 KB	2009/11/17 16:55:01	HTTP	△
🔍	⋙	10.32.128.32:35830	66.102.7.104:80	(75%)	3 KB/0 KB	2009/11/17 16:54:48	HTTP	△
🔍	⋙	10.32.129.20:4458	74.125.19.18:80	(75%)	1,010 KB/250 KB	2009/11/17 13:10:30	HTTP	△
🔍	⋙	10.32.129.20:1749	10.16.1.100:19068	(74%)	838 KB/213 KB	2009/11/10 08:16:57	MAPI	△
🔍	⋙	10.32.129.20:2354	74.125.19.102:80	(67%)	1 KB/0 KB	2009/11/17 16:52:08	HTTP	△
🔍	⋙	10.32.129.32:59321	10.1.9.133:80	(65%)	1 KB/0 KB	2009/11/17 16:54:56	HTTP	△
🔍	⋙	10.32.128.138:60714	74.125.19.139:80	(65%)	1 KB/0 KB	2009/11/17 16:53:24	HTTP	△

2. Use the controls to customize the report as described in the following table.

Control	Description
Display	Click the number of lines (20, 100, or 4000) to display.
Connections of Type	Select one of the following connection types to filter the display: • **All** - Specifies all established active connections, removes any previous filters. • **All Optimized** - Specifies the total optimized active connections. This is the default setting. • **All Established** - Specifies the total established active connections. • **Half-Open** - Specifies the total half-open active connections. A half-open connection is a TCP connection which has not been fully established. Half-open connections count toward the connection count limit on the Steelhead appliance because, at any time, they might become a fully open connection. If you are experiencing a large number of half-open connections, consider a more appropriately sized Steelhead appliance. • **Half-Closed** - Specifies the total half-closed active connections. A half-closed connection has been intercepted and optimized by the Steelhead appliance but is in the process of becoming inactive. These connections are counted toward the connection count limit on the Steelhead appliance. (Half-closed connections might remain if the client or server does not close their connections cleanly.) If you are experiencing a large number of half-closed connections, consider a more appropriately sized Steelhead appliance. • **All Passthrough** - Specifies the total connections passed through intentionally and unintentionally unoptimized. • **Intentional Passthrough** - Specifies the total connections passed through intentionally unoptimized, for any of several reasons such as: an in-path rule for a client-side PBR deployment, an interactive and secure connection, a connection used for troubleshooting, and so on. • **Unintentional Passthrough** - Specifies the total connections passed through unintentionally, unoptimized, for any of several reasons such as misconfigured Steelhead appliances, an unreachable server, a Steelhead appliance reaching the connection limit, and so on. • **Forwarded** - Specifies the total number of connections forwarded by the connection forwarding neighbor managing the connection.

Control	Description

Filter By	Filters the display based on the text specified in the text field and one of the following options:
	• **Regular Expression** - Shows only those connections which match the expression used to filter the display.
	Use the following format in the text field:
	`x.x.x.x[/mask][:port]`
	Examples:
	`10.16.35.1`
	Finds one particular IP
	`10.16.35.1:5001`
	Finds port 5001 on one particular IP
	You can also use the Regular Expression option to show only those connections for which the expression matches the following string:
	<source IP>:<source Port> <destination IP>:<destination Port> <protocol Name>
	where each token in angle brackets is replaced by the connection properties. Use a single space between <source Port> and <destination IP> and between <destination Port> <protocol name>.
	• **Application Protocol** - Filters by the specified application protocol, such as CIFS.
	• **Source IP: Port** - Displays only the connections that originate from the specified subnet and port.
	`0.0.0.0/0:50109`
	Finds any IP with port 50109
	`10.0.0.0/8:443`
	Finds any IP on the 10.0.0.0/8 subnet with port 443
	• **Destination IP: Port** - Displays only the connections received by the specified subnet and port.
	• **Either: IP: Port** - Displays only the connections either originating from or received by the specified subnet and port.
	Notes:
	• The regular expression matching ignores case.
	• The pattern follows Python regular expression syntax. For details, see
	http://www.python.org/doc/howto
	• The Python regex filter used in this report requires extra escaping; to place a backslash (\) into the regular expression, you must specify two (\\).
	Examples:
	^10\\.32\\.
	Finds all connections whose source IP begins with 10.32.
	\\s10\\.1\\.
	Finds all connections whose destination IP begins with 10.1.
	\\b10\\.[13]\\d\\.
	Finds all connections whose source or destination IP begins with 10.1*x* or 10.3*x*, where *x* is a digit from 0 through 9.

Control	Description
	Examples cont'd `:8080 HTTP` Finds all HTTP connections to port 8080.
Update Display	Displays the report.

Note: For information on removing an unknown Steelhead appliance from the current connections list, see "Preventing an Unknown (or Unwanted) Steelhead Appliance from Peering" on page 73.

Viewing the Current Connection Details

The Current Connections report displays details about the connected appliances, such as the source and destination IP address, the peer Steelhead appliance, the inner local port, and so on. You can also perform the following operations:

- For optimized connections, send a keep-alive message to the outer remote machine (the machine that is connected to this appliance).
- Reset any connection, optimized or pass-through.

To view current connection details

1. Choose Reports > Networking > Current Connections to display the Current Connections report.

2. Click the magnifying glass in the first column of the individual connections table to see more details about an individual connection and perform operations on it. Because this report is a snapshot in time, by the time you click, the connection could be gone or in a different state.

Figure 10-2. Current Connections Details for an Optimized Connection

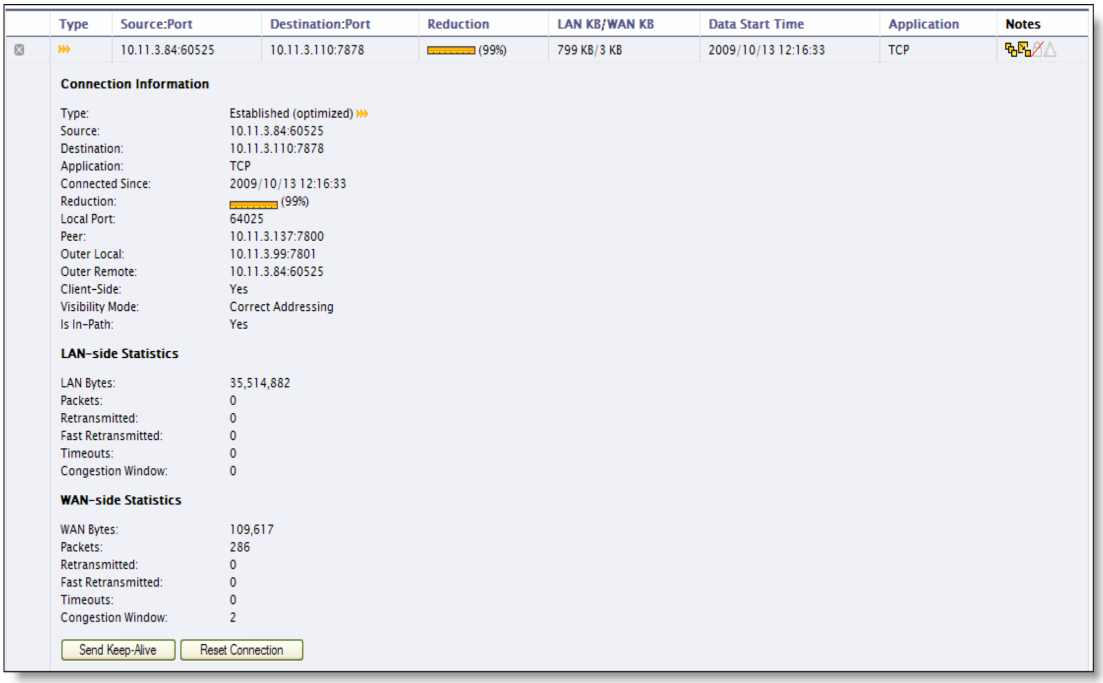

	Type	Source:Port	Destination:Port	Reduction	LAN KB/WAN KB	Data Start Time	Application	Notes
☒	»»	10.11.3.84:60525	10.11.3.110:7878	(99%)	799 KB/3 KB	2009/10/13 12:16:33	TCP	

Connection Information

Type:	Established (optimized) »»
Source:	10.11.3.84:60525
Destination:	10.11.3.110:7878
Application:	TCP
Connected Since:	2009/10/13 12:16:33
Reduction:	(99%)
Local Port:	64025
Peer:	10.11.3.137:7800
Outer Local:	10.11.3.99:7801
Outer Remote:	10.11.3.84:60525
Client-Side:	Yes
Visibility Mode:	Correct Addressing
Is In-Path:	Yes

LAN-side Statistics

LAN Bytes:	35,514,882
Packets:	0
Retransmitted:	0
Fast Retransmitted:	0
Timeouts:	0
Congestion Window:	0

WAN-side Statistics

WAN Bytes:	109,617
Packets:	286
Retransmitted:	0
Fast Retransmitted:	0
Timeouts:	0
Congestion Window:	2

Optimized Connection Details

The following table summarizes details on individual optimized connections.

Field	Description (varies by connection)
Connection Information	**Type** - Specifies whether the connection is established, half-open, or half-closed. **Source** - Specifies the source IP address for the connection. **Destination** - Specifies the destination IP address for the connection. **Application** - Specifies the application corresponding to the connection, for example, NFS. **Connected Since** - Specifies the date and time the connection was established. **Reduction** - Displays the percentage of reduction for the transmitted data. **Local Port** - Specifies the port on which the WAN interface of this appliance receives optimized messages. **Peer** - Specifies the IP address and port for the peer Steelhead appliance. **Outer Local** - Specifies the IP address and port on which this appliance intercepted this connection from the LAN. **Outer Remote** - Specifies the IP address and port for the client or server connected to Outer Local over the LAN. **Client Side** - Specifies whether the connection is a client-side. **Visibility Mode** - Specifies the WAN visibility mode in effect: Correct Addressing, Port Transparency, Full Transparency, or Full Transparency w/Reset. **Is In-path** - Indicates whether the connection is in-path. **In-path rule** - Specifies the name of the in-path rule governing the connection.
WAN and LAN-Side Statistics	**LAN Bytes** - Specifies the total LAN bytes transmitted. **WAN Bytes** - Specifies the total WAN bytes transmitted. **Packets** - Specifies the total number of packets transmitted. **Retransmitted** - Specifies the total packets retransmitted. **Fast Retransmitted** - Specifies the total packets fast retransmitted. Fast retransmit reduces the time a sender waits before retransmitting a lost segment. If an acknowledgement is not received for a particular segment with a specified time (a function of the estimated round-trip delay time), the sender assumes the segment was lost in the network, and retransmits the segment. **Timeouts** - Specifies the number of time-outs. **Congestion Window** - Specifies the WAN or LAN congestion window size.
Send Keep-Alive	For an optimized connection, sends a keep-alive message to the outer remote machine (the machine that is connected to this appliance).
Reset Connection	Sends a RST packet to both the client and server to try to close the connection. You can reset both optimized and pass-through connections. **Note**: If no data is being transferred between the client and server when you click **Reset Connection**, the connection is not reset immediately. It resets the next time the client or server tries to send a message. Therefore, when the application is idle, it may take a while for the connection to disappear.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Tip: To close the connection details report, click the x.

Figure 10-3. Current Connections Details for a Pass-Through Connection

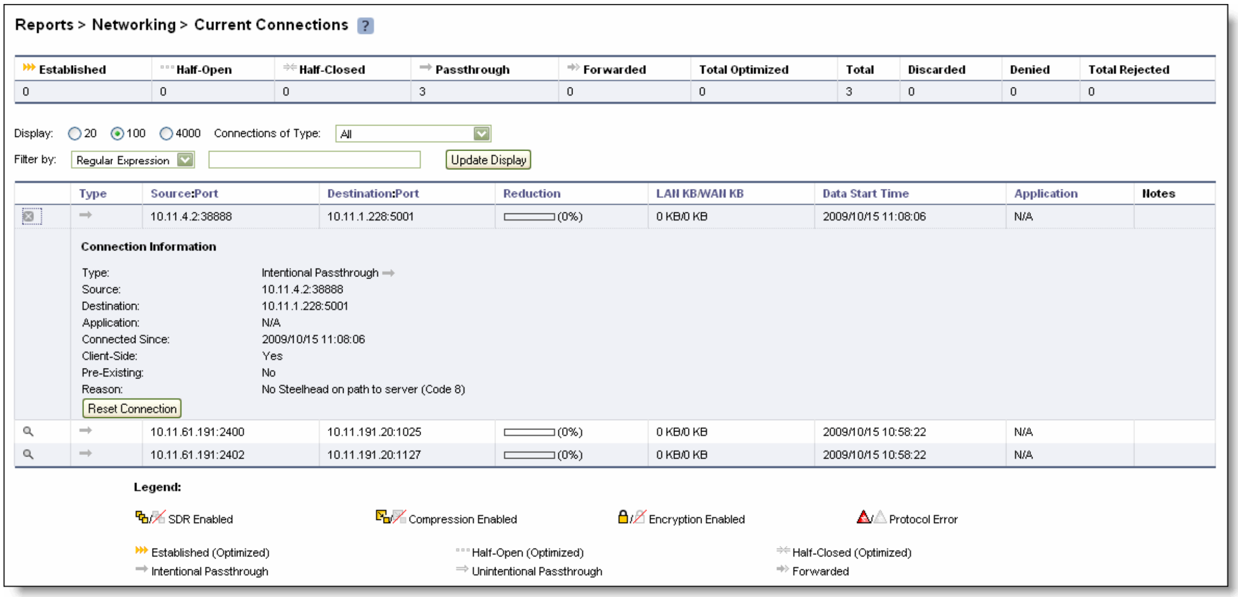

Reports > Networking > Current Connections ?

Established	Half-Open	Half-Closed	Passthrough	Forwarded	Total Optimized	Total	Discarded	Denied	Total Rejected
0	0	0	3	0	0	3	0	0	0

Display: ○ 20 ● 100 ○ 4000 Connections of Type: All

Filter by: Regular Expression

	Type	Source:Port	Destination:Port	Reduction	LAN KB/WAN KB	Data Start Time	Application	Notes
☒	→	10.11.4.2:38888	10.11.1.228:5001	(0%)	0 KB/0 KB	2009/10/15 11:08:06	N/A	
		10.11.61.191:2400	10.11.191.20:1025	(0%)	0 KB/0 KB	2009/10/15 10:58:22	N/A	
		10.11.61.191:2402	10.11.191.20:1127	(0%)	0 KB/0 KB	2009/10/15 10:58:22	N/A	

Connection Information

Type:	Intentional Passthrough
Source:	10.11.4.2:38888
Destination:	10.11.1.228:5001
Application:	N/A
Connected Since:	2009/10/15 11:08:06
Client-Side:	Yes
Pre-Existing:	No
Reason:	No Steelhead on path to server (Code 8)

Individual Pass-Through or Forwarded Connection Details

The following table summarizes details on individual pass-through or forwarded connections.

Field	Description (varies by connection)
Connection Information	**Type** - Specifies whether the pass-through was intentional or unintentional. Displays the Forwarded Reduction Percentage bar for forwarded connections. **Source** - Specifies the source IP address for the connection. **Destination** - Specifies the destination IP address for the connection. **Application** - Specifies the application corresponding to the connection, for example, NFS. **Connected Since** - Specifies the date and time the connection was established. **Client-Side** - Specifies whether the connection is a client-side. **Pre-Existing** - Specifies whether the connection existed before the last restart of the optimization service. **Reason** - Specifies the reason for passing through or forwarding the connection. The pass-through reasons are listed in "Pass-Through Reasons" on page 397.
Reset Connection	Sends a RST packet to both the client and server to try to close the connection. You can reset both optimized and pass-through connections. **Note**: If no data is being transferred between the client and server when you click **Reset Connection**, the connection is not reset immediately. It resets the next time the client or server tries to send a message. Therefore, when the application is idle, it may take a while for the connection to disappear.

Pass-Through Reasons

The following table lists the pass-through reasons.

Value	Pass-through reason (varies by connection)	Description	Action
0	None	None	None
1	Preexisting connection	Connection existed before Steelhead appliance started.	Create a new connection.
2	Connection paused	Steelhead appliance is not intercepting connections.	Check that the service is enabled, in-path is enabled, the neighbor configuration, and whether the Steelhead appliance is in admission control.
3	SYN on WAN side	Client is on the Steelhead appliance WAN side.	Either this is the server-side Steelhead appliance and there is no client-side Steelhead appliance, or the client-side Steelhead appliance did not probe. Check the cabling if it is really the client-side Steelhead appliance.
4	In-path rule	In-path rule matched on the client-side Steelhead appliance is pass-through.	Check the in-path rules.
5	Peering rule	Peering rule matched on the server-side Steelhead appliance is pass-through.	Check the peering rules.
6	Inner failed to establish	Inner connection between Steelhead appliances failed.	Check the connectivity between the client-side Steelhead appliance and the server-side Steelhead appliance.
7	Peer in fixed-target rule down	The target of a fixed-target rule is destined to a failed peer.	Check the connectivity between the client-side Steelhead appliance and the server-side Steelhead appliance.
8	No Steelhead on path to server	No server-side Steelhead appliance.	Check that the server-side Steelhead appliance is up and check that the connection goes through the server-side Steelhead appliance.
9	No route for probe response	No route to send back probe response.	Check in-path gateway on the server-side Steelhead appliance.
10	Out of memory	Memory problem while copying packet.	Check if the Steelhead appliance is out of memory.

Value	Pass-through reason (varies by connection)	Description	Action
11	No room for more TCP options	Not enough space in TCP header to add probe.	This occurs when another device added TCP options before the Steelhead. Take a TCP dump to check which TCP options are in the SYN packet. Search for those options to find out what device uses them.
12	No proxy port for probe response	There is no service port configured on server-side Steelhead appliance.	Configure a service port.
13	RX probe from failover buddy	The connection is intercepted by failover buddy.	No action is necessary.
14	Asymmetric routing	The connection is asymmetric.	Check the asymmetric routing table for reason.
15	Middle Steelhead	The Steelhead appliance is not the first or last Steelhead appliance.	Only happens with the new discovery protocol.
16	Error connecting to server	The server-side Steelhead appliance could not connect to the server.	Only happens with the new discovery protocol.
17	Half open connections above limit	The client has too many half-opened connections.	Check if many connections open quickly from the same client.
18	Connection count above QOS limit	There are too many connections for that QoS class.	Check the QoS class.
19	Reached maximum TTL	The probe has an incorrect TTL.	Take a trace to check the probe.
20	Incompatible probe version	The probe has an incompatible version number.	Check if the new probe format is enabled, it is disabled by default.
21	Too many retransmitted SYNs	The client SYN has been retransmitted too many times.	Check if there is a firewall that does not like the probe TCP option.
22	Connection initiated by neighbor	The connection is intercepted by a neighbor.	No action is necessary.
23	Connection for local host	The connection is to the in-path interface.	No action is necessary.
24	Unknown reason	The pass-through reason does not match any other description.	No action is necessary.
25	Connection from proxy target	Because the connection originates from an IP address which is also the IP address of a fixed target rule, it is not intercepted.	No action is necessary.

Value	Pass-through reason (varies by connection)	Description	Action
26	SYN before SFE outer completes	The client connection was passed-through at the client-side Steelhead appliance and the client's pure SYN was seen at the server-side Steelhead appliance.	Check if there is a firewall that does not like the probe TCP option.
27	Transparent inner on wrong VLAN	The inner connection seen on VLAN is different than the in-path VLAN.	No action is necessary.
28	Transparent inner not for this host		No action is necessary.
29	Error on neighbor side	The neighbor Steelhead appliance returned an error to a connection forwarding request.	Check the health of the configured neighbors.
30	SYN/ACK, but no SYN	There is asymmetric routing - received SYN/ACK but no SYN.	Check your routing.
31	Transparency packet from self	For Riverbed internal use only.	No action is necessary.
32	System is heavily loaded	The Steelhead appliance is experiencing a heavy traffic load.	Contact Riverbed Support. You might require a larger model Steelhead appliance.
33	SYN/ACK at MFE not SFE	There is asymmetric routing around the server-side Steelhead appliance.	Check your routing.
34	Windows branch mode detected	The client-side is a Steelhead Mobile Client. Optimization is occurring between the Steelhead Mobile Client and the server-side Steelhead appliance, so the connection is passed through on the client-side Steelhead appliance.	No action is necessary.
35	Transparent RST to reset firewall state	The optimization service has sent a RST to clear the probe connection created by the Steelhead appliance and to allow for the full transparent inner connection to traverse the firewall.	No action is necessary.
36	Error on SSL inner channel handshake	An inner channel handshake has failed with peer.	Check the SSL configuration on both Steelhead appliances.

Value	Pass-through reason (varies by connection)	Description	Action
37	Ricochet or probe	This pass-through reason is attributed to a flow reported to a v9 NetFlow collector. A probe and packet have been sent by the Steelhead appliance back through itself. For example, in an in-path setup, if a client-side Steelhead appliance gateway is on its WAN side, all packets sent to the client will first go to the gateway and be sent back through the Steelhead appliance on the way to the client.	Packet ricochet can be avoided in many environments by enabling simplified routing.
38	Passthrough due to MAPI admission control	New MAPI connections will be passed through due to high connection count.	New MAPI connections will automatically be optimized when the MAPI traffic has decreased.

Viewing Connection History

The Connection History report summarizes the optimized traffic for the time period specified. It contains the following graphs:

- **Optimized versus Pass Through Connections** - This graph displays the total number of optimized and passed-through connections for the time period specified.
- **Optimized Connections** - This graph displays the total number of optimized, established, half-opened, and half-closed connections for the time period specified.

Note: The graphs in this report plot the *peak* connection history data, not the *average* connection history data. The tables following the graphs display both peaks and averages.

The Connection History report contains the following table of statistics that summarize connection activity.

Packet Type	Description
Optimized	Specifies the total active connections established and optimized, plus the half-open and half-closed connections (where the half-open and half-closed connections are TCP connection states).
Optimized (Active)	Specifies the total number of optimized connections with traffic in the last 60 seconds.
Pass Through	Specifies the total connections passed through, unoptimized.
Optimized (Established)	Specifies the total established connections.

Packet Type	Description
Optimized (Half Opened)	Specifies the total half-opened connections. A half-open connection is a TCP connection which has not been fully established. Half-open connections count toward the connection count limit on the Steelhead appliance because, at any time, they might become a fully-open connection. If you are experiencing a large number of half-open connections, consider a more appropriately sized Steelhead appliance.
Optimized (Half Closed)	Specifies the total half-closed active connections. Half-closed connections are connections which the Steelhead appliance has intercepted and optimized but are in the process of becoming inactive. These connections are counted toward the connection count limit on the Steelhead appliance. (Half-closed connections might remain if the client or server does not close their connections cleanly.) If you are experiencing a large number of half-closed connections, consider a more appropriately sized Steelhead appliance.
Forwarded	Specifies the total forwarded connections.

What This Report Tells You

The Connection History report answers the following questions:

- How many connections were optimized?
- How many connections were passed through, unoptimized?
- How many connections were half-opened?
- How many connections were half-closed?
- How many connections were forwarded to another Steelhead appliance?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Connection History report

1. Choose Reports > Networking > Connection History to display the Connection History page.

Figure 10-4. Connection History Page

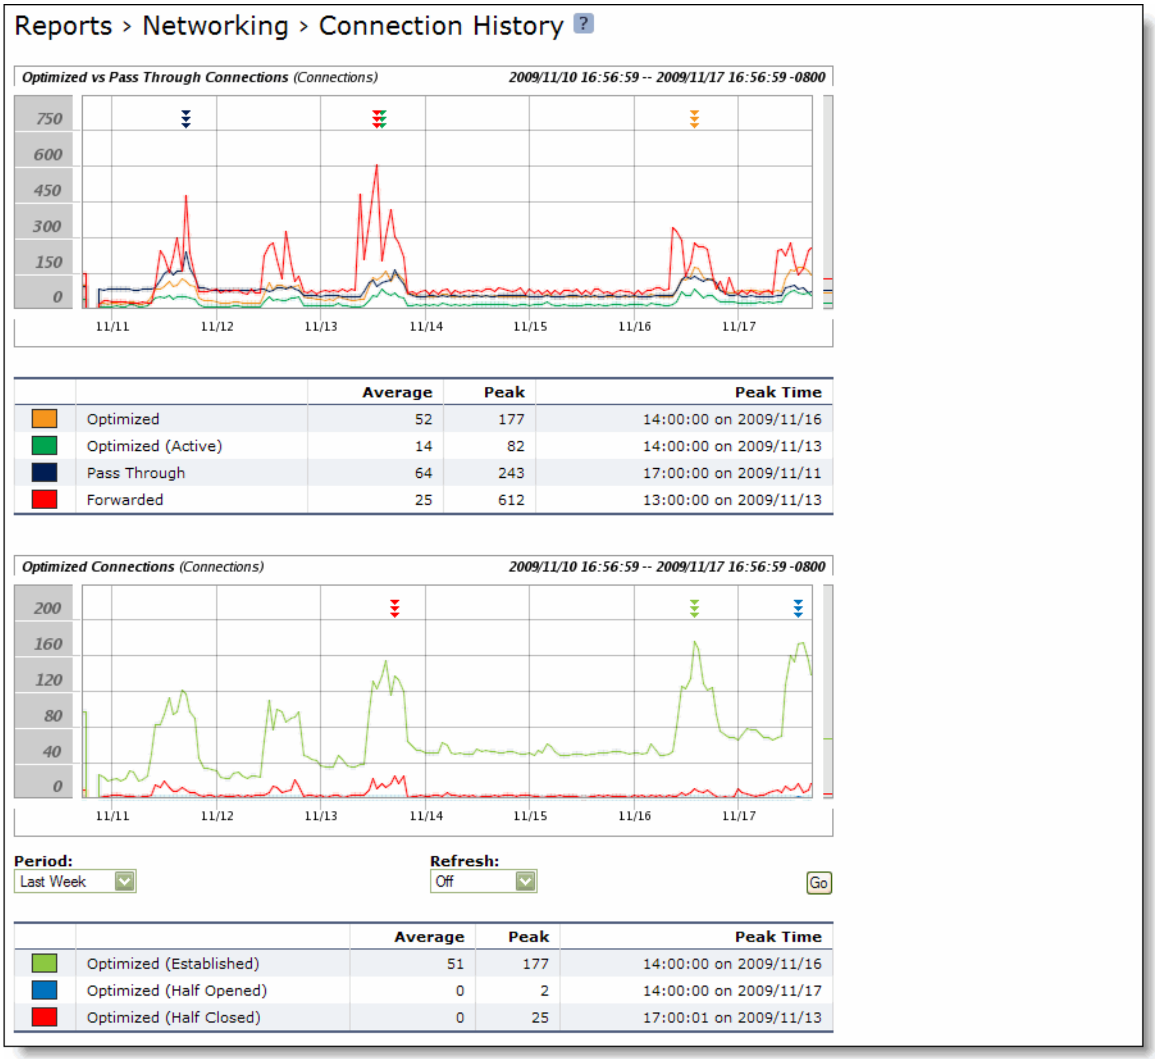

Reports > Networking > Connection History ?

		Average	Peak	Peak Time
🟧	Optimized	52	177	14:00:00 on 2009/11/16
🟩	Optimized (Active)	14	82	14:00:00 on 2009/11/13
⬛	Pass Through	64	243	17:00:00 on 2009/11/11
🟥	Forwarded	25	612	13:00:00 on 2009/11/13

Period: Last Week

Refresh: Off

		Average	Peak	Peak Time
🟩	Optimized (Established)	51	177	14:00:00 on 2009/11/16
🟦	Optimized (Half Opened)	0	2	14:00:00 on 2009/11/17
🟥	Optimized (Half Closed)	0	25	17:00:01 on 2009/11/13

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click Go. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click Off.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Connection Forwarding Reports

The Connection Forwarding report summarizes the number of bytes or packets transferred between the Steelhead appliance and a specified neighbor.

Field	Description
Total Data Sent	Specifies the number of bytes or packets transferred.

You configure neighbors when you enable connection forwarding. For details, see "Configuring Connection Forwarding Features" on page 255.

What This Report Tells You

The Connection Forwarding report answers the following questions:

- How many bytes were transferred between a Steelhead appliance and a specified neighbor?
- How many packets were transferred between a Steelhead appliance and a specified neighbor?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Connection Forwarding report

1. Choose Reports > Networking > Connection Forwarding to display the Connection Forwarding page.

Figure 10-5. Connection Forwarding Page

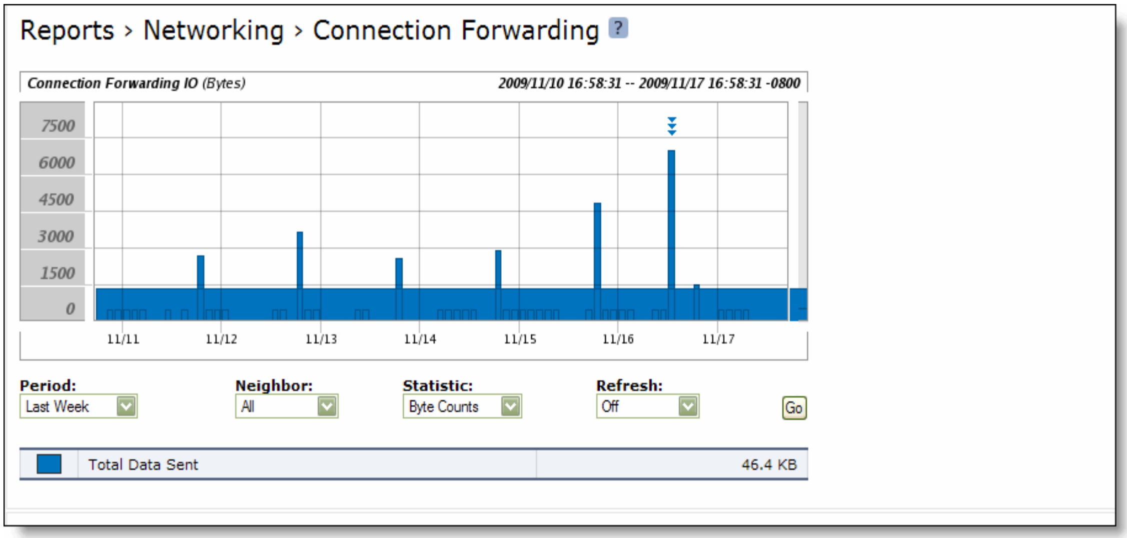

Reports > Networking > Connection Forwarding

Connection Forwarding IO (Bytes) 2009/11/10 16:58:31 -- 2009/11/17 16:58:31 -0800

Period: Last Week
Neighbor: All
Statistic: Byte Counts
Refresh: Off
Go

	Total Data Sent	46.4 KB

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click Go. Use the following format: YYYY/MM/DD HH:MM:SS
Neighbor	Select a specific neighbor or All from the drop-down list.
Statistic	Select either Byte Counts or Packet Counts from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing QoS Stats Dropped Reports

The QoS Stats Dropped report summarizes the number of bits and packets transmitted for the QoS class or an aggregate of all classes for the time period specified.

The QoS Stats Dropped report contains the following graphs:

- **QoS Pre-Enforcement** - Displays the total number of bits sent before enforcement of the QoS parameters for the time period specified.
- **QoS Enforced/Dropped** - Displays the total number of bits dropped after QoS enforcement parameters have been set for the time period specified.

The QoS Stats Dropped report contains the following table of statistics that summarize QoS activity during peak pre-enforcement and peak post-enforcement time periods.

Field	Description
Maximum All Throughput At <time> on <date>	Specifies the date and time of the peak QoS throughput.

What This Report Tells You

The QoS Stats Dropped report answers the following questions:

- How many bits were transmitted over the WAN for the QoS class?
- How many data packets were dropped for the QoS class?
- When did the peak data transmission occur for the QoS class?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the QoS Stats Dropped report

1. Choose Reports > Networking > QoS Stats Dropped to display the QoS Stats Dropped page.

Figure 10-6. QoS Stats Dropped Page

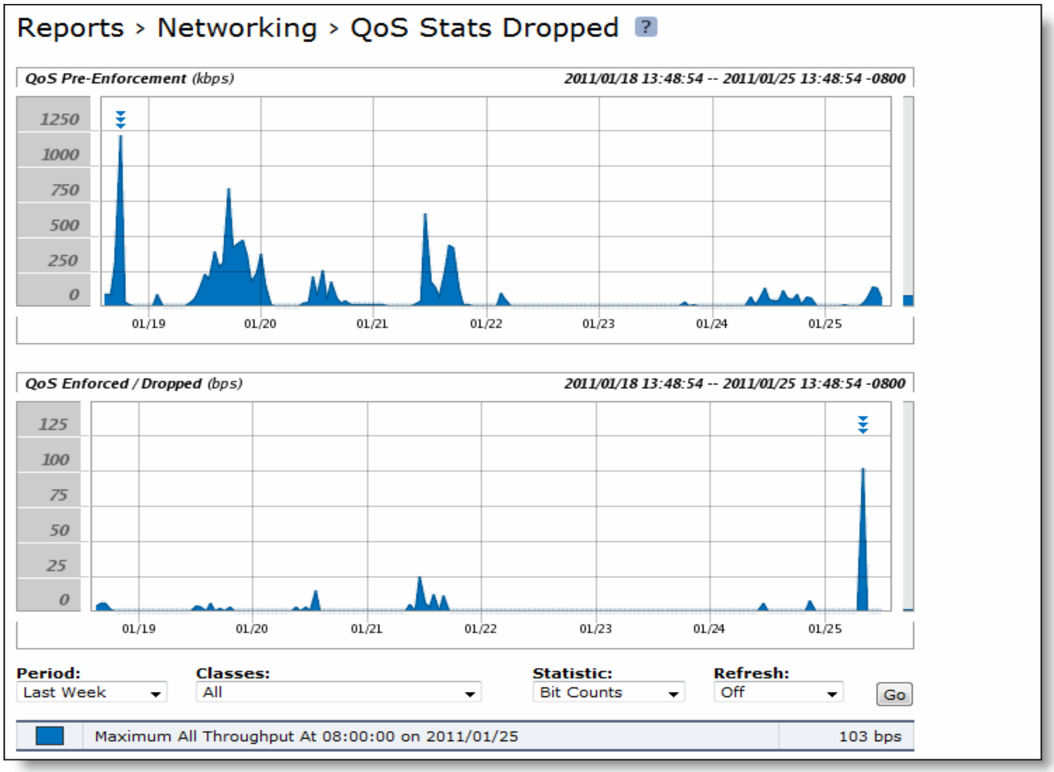

Reports > Networking > QoS Stats Dropped

Period: Last Week

Classes: All

Statistic: Bit Counts

Refresh: Off

🟦	Maximum All Throughput At 08:00:00 on 2011/01/25	103 bps

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, specify the Start Time and End Time and click Go. Use the following format: YYYY/MM/DD HH:MM:SS
Classes	Select a All, a class, or Custom from the drop-down list. Selecting a parent class displays its child classes. For example, the report for an HTTP class with two child classes named WebApp1 and WebApp2 displays statistics for HTTP, WebApp1, and WebApp2. Selecting Custom displays a drop-down list of all the custom classes.
Statistic	Select either Bit Counts or Packet Counts from the drop-down list.

Control	Description
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing QoS Stats Sent Reports

The QoS Stats Sent report includes a graph which summarizes the number of bits and packets transmitted for the QoS class or an aggregate of all classes for the time period specified.

The QoS Stats Dropped report contains the following graphs:

- **QoS Pre-Enforcement** - Displays the total number of bits sent before enforcement of the QoS parameters for the time period specified.
- **QoS Enforced/Sent** - Displays the total number of bits sent after QoS enforcement parameters have been set for the time period specified.

The QoS Stats Sent report contains the following table of statistics that summarize QoS activity during peak pre-enforcement and peak post-enforcement time periods.

Field	Description
Maximum All Throughput At <time> on <date>	Specifies the date and time of the peak QoS throughput.

What This Report Tells You

The QoS Stats Sent report answers the following questions:

- How many bits were transmitted over the WAN for the QoS class?
- How many data packets were sent for the QoS class?
- When did the peak data transmission occur for the QoS class?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the QoS Stats Sent report

1. Choose Reports > Networking > QoS Stats Sent to display the QoS Stats Sent page.

Figure 10-7. QoS Stats Sent Page

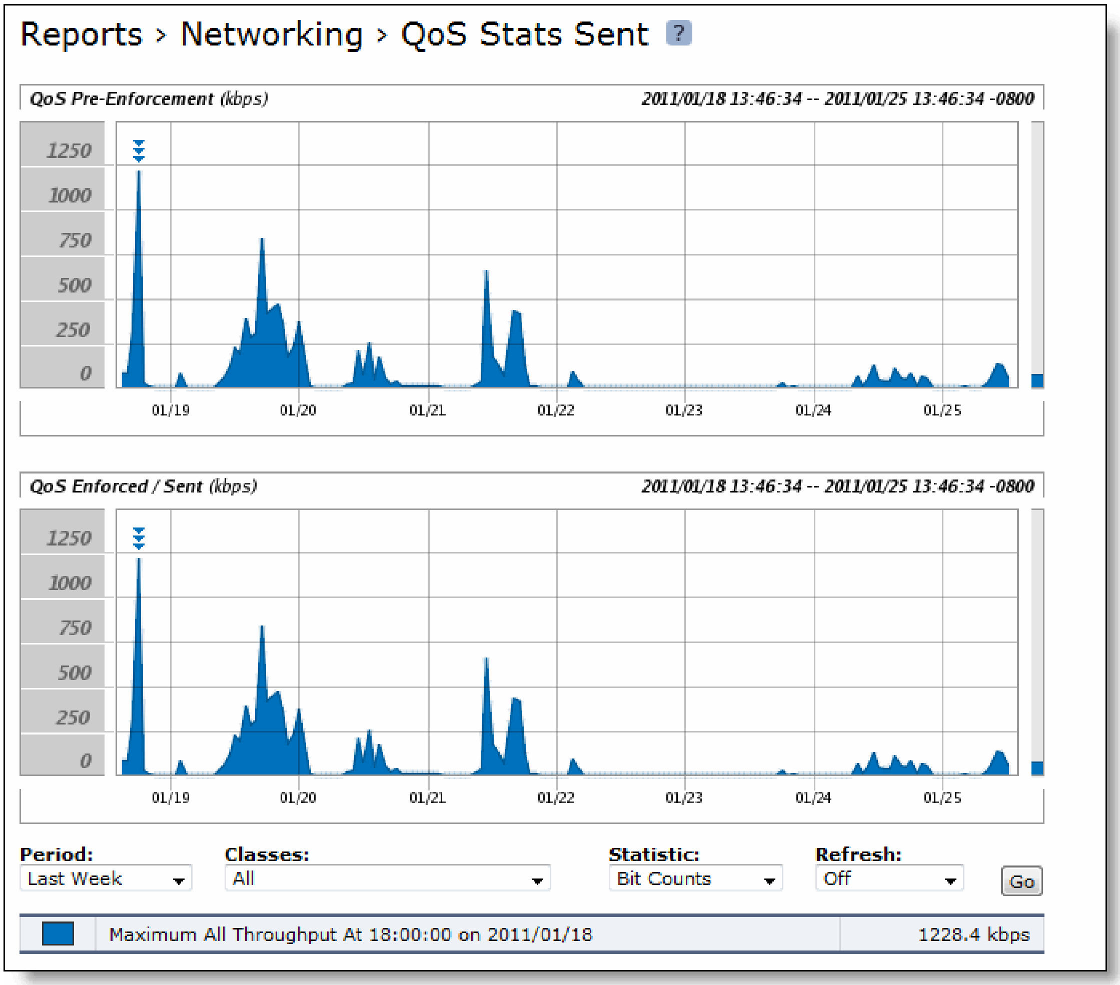

Reports > Networking > QoS Stats Sent

Period: Last Week
Classes: All
Statistic: Bit Counts
Refresh: Off

🟦 Maximum All Throughput At 18:00:00 on 2011/01/18 1228.4 kbps

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, specify the Start Time and End Time. Use the following format: YYYY/MM/DD HH:MM:SS
Classes	Select either All, a class, or Custom from the drop-down list. Selecting a parent class displays its child classes. For example, the report for an HTTP class with two child classes named WebApp1 and WebApp2 displays statistics for HTTP, WebApp1, and WebApp2. Selecting Custom displays a drop-down list of custom classes. You can display a maximum of eight classes.
Statistic	Select either Bit Counts or Packet Counts from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Top Talkers Reports

The Top Talkers report displays the top talking hosts on a per-port basis for the time period specified. The traffic flows that generate the heaviest use of WAN bandwidth are known as the Top Talkers. This report provides WAN visibility for traffic analysis, security monitoring, accounting, load balancing, and capacity planning. It can include both optimized and pass-through traffic.

A traffic flow consists of data sent and received from a single source IP address and port number to a single destination IP address and port number over the same protocol. Only traffic flows that start in the selected time period are shown in the report.

Important: The Top Talkers report includes bytes used for packet headers and is an approximation based on various assumptions.

The Top Talkers report contains the following table of statistics that summarize Top Talker activity.

Field	Description
Rank	Specifies the relative position of the traffic flow WAN bandwidth use.
<Sender> Source IP:Port	Specifies the source IP address and source port for the connection.

Field	Description
<Receiver> Destination IP:Port	Specifies the destination IP address and destination port for the connection.
Byte Count	Specifies the total number of bytes sent and received by the source IP address.

You can export this report in CSV format in the Export report. The CSV format allows you to easily import the statistics into spreadsheets and databases. You can open the CSV file in any text editor. For details, see "Exporting Performance Statistics" on page 490.

Important: Flow Export must be turned on before viewing the Top Talker report. For details, see "Configuring Subnet Side Rules" on page 260.

What This Report Tells You

The Top Talkers report answers the following question:

- Who were the top talking hosts on a per-port basis?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Top Talkers report

1. Choose Reports > Networking > Top Talkers to display the Top Talkers page.

Figure 10-8. Top Talkers Page

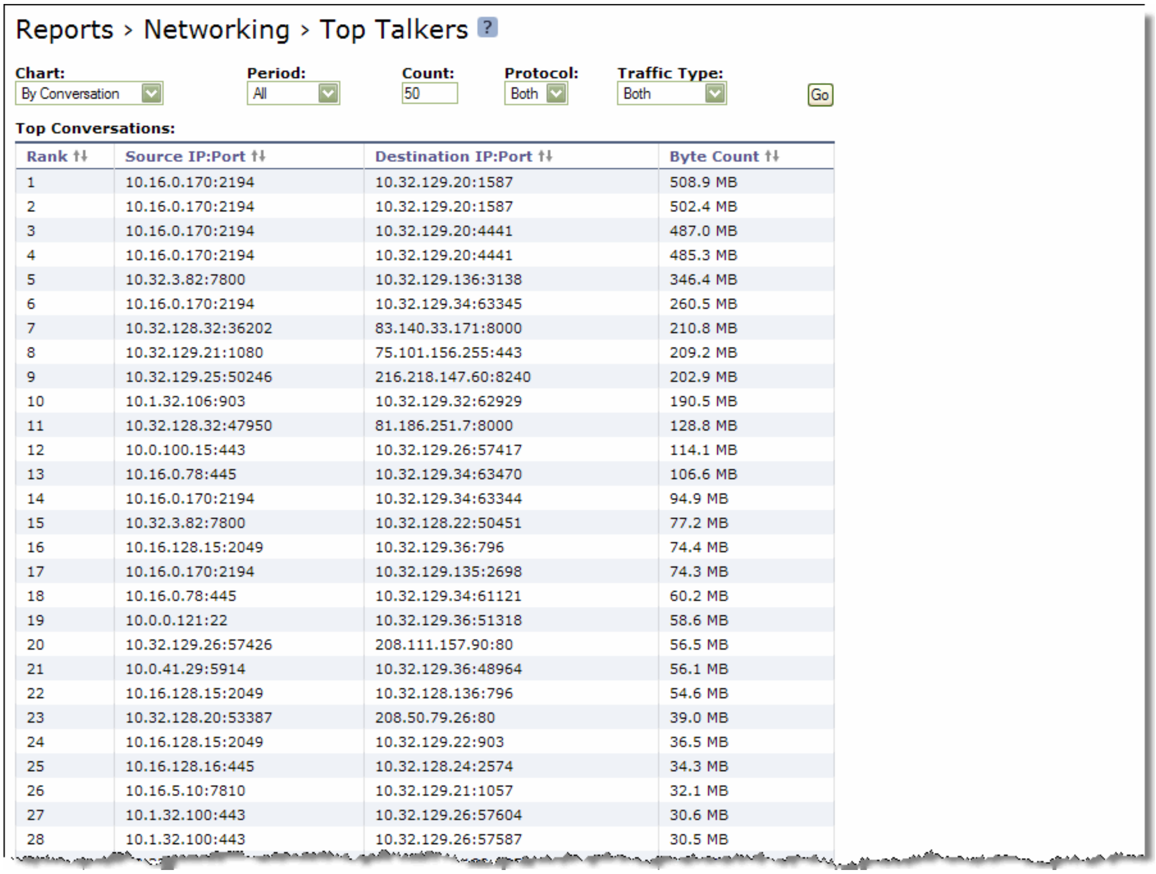

Reports > Networking > Top Talkers

Chart: By Conversation

Period: All

Count: 50

Protocol: Both

Traffic Type: Both

Top Conversations:

Rank	Source IP:Port	Destination IP:Port	Byte Count
1	10.16.0.170:2194	10.32.129.20:1587	508.9 MB
2	10.16.0.170:2194	10.32.129.20:1587	502.4 MB
3	10.16.0.170:2194	10.32.129.20:4441	487.0 MB
4	10.16.0.170:2194	10.32.129.20:4441	485.3 MB
5	10.32.3.82:7800	10.32.129.136:3138	346.4 MB
6	10.16.0.170:2194	10.32.129.34:63345	260.5 MB
7	10.32.128.32:36202	83.140.33.171:8000	210.8 MB
8	10.32.129.21:1080	75.101.156.255:443	209.2 MB
9	10.32.129.25:50246	216.218.147.60:8240	202.9 MB
10	10.1.32.106:903	10.32.129.32:62929	190.5 MB
11	10.32.128.32:47950	81.186.251.7:8000	128.8 MB
12	10.0.100.15:443	10.32.129.26:57417	114.1 MB
13	10.16.0.78:445	10.32.129.34:63470	106.6 MB
14	10.16.0.170:2194	10.32.129.34:63344	94.9 MB
15	10.32.3.82:7800	10.32.128.22:50451	77.2 MB
16	10.16.128.15:2049	10.32.129.36:796	74.4 MB
17	10.16.0.170:2194	10.32.129.135:2698	74.3 MB
18	10.16.0.78:445	10.32.129.34:61121	60.2 MB
19	10.0.0.121:22	10.32.129.36:51318	58.6 MB
20	10.32.129.26:57426	208.111.157.90:80	56.5 MB
21	10.0.41.29:5914	10.32.129.36:48964	56.1 MB
22	10.16.128.15:2049	10.32.128.136:796	54.6 MB
23	10.32.128.20:53387	208.50.79.26:80	39.0 MB
24	10.16.128.15:2049	10.32.129.22:903	36.5 MB
25	10.16.128.16:445	10.32.128.24:2574	34.3 MB
26	10.16.5.10:7810	10.32.129.21:1057	32.1 MB
27	10.1.32.100:443	10.32.129.26:57604	30.6 MB
28	10.1.32.100:443	10.32.129.26:57587	30.5 MB

2. Use the controls to customize the report as described in the following table.

Control	Description
Chart	Select the report display from the drop-down list: By Conversation, By Sender, By Receiver, By Host, or By Application Port. The default setting is By Conversation.
Period	You can view the traffic statistics for the past hour, the past 24 hours, or all available hours. All is the default setting, which displays statistics for the entire duration the Steelhead appliance has gathered statistics. This can be up to 2 days, depending on how long the service has been up and the traffic volume. Select All, Last Hour, or Last Day from the drop-down list. The default setting is All. **Note:** Top Talker statistics are not persistent between service restarts.
Count	Specify how many top pairs of source and destination addresses and ports with the highest total traffic (sent and received) appear in the report. Each pair shows the number of bytes and packets sent and received at the source address. The default value is 50. **Note:** You can export the complete list of top talkers to a file in CSV format using the Export report.
Protocol	Select Both, TCP, or UDP from the drop-down list. The default value is Both.
Traffic Type	Select Both, Optimized, or Passthrough from the drop-down list. The default value is Both.
Go	Displays the report.

Note: The Top Talkers data does not exactly match the Traffic Summary data, the Bandwidth Optimization data, or specific connection data that appears when you select a particular connection in the Current Connections report. This is due to packet headers, packet retransmits, and other TCP/IP effects that flow export collectors see, but RiOS does not. Consequently, the reports are proportional but not equivalent.

Tip: Select a Top Talkers report column heading to sort the column in ascending or descending order.

Viewing Traffic Summary Reports

The Traffic Summary report provides a percentage breakdown of the amount of TCP traffic going through the system by the port and type of traffic. For details on setting ports to be monitored, see "Configuring Monitored Ports" on page 335.

The Steelhead appliance automatically discovers all the ports in the system that have traffic. The discovered port with a label (if one exists) is added to the report. If a label does not exist then an unknown label is added to the discovered port.

If you want to change the unknown label to a name representing the port, you must add the port with a new label. All statistics for this new port label are preserved from the time the port was discovered. For details on adding ports to be monitored, see "Configuring Monitored Ports" on page 335.

Note: The Traffic Summary report displays a maximum of 16 ports and pie slices for the top 100 traffic types (by destination port). When there are more than 16 ports, the report displays 15 individual ports and aggregates the remaining ports into the 16th slice. The 16th slice is always gray. Any ports aggregated into the 16th slice are also gray. Any additional traffic beyond the top 100 ports is not included in the Traffic Summary report, but is aggregated into the Bandwidth Optimization report.

The Traffic Summary report provides the following table of statistics that describe data activity for the application and the time period you specify.

Field	Description
Port	Specifies the TCP/IP port number and application for each row of statistics.
Reduction	Specifies the amount of application data reduction.
LAN Data	Specifies the amount of application data on the LAN.
WAN Data	Specifies the amount of application data on the WAN.
Traffic %	Indicates the percentage of the total traffic each port represents.

What This Report Tells You

The Traffic Summary report answers the following questions:

- How much data reduction has occurred?
- What was the percentage of the total traffic for each port?

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Traffic Summary report

1. Choose Reports > Networking > Traffic Summary to display the Traffic Summary page.

Figure 10-9. Traffic Summary Page

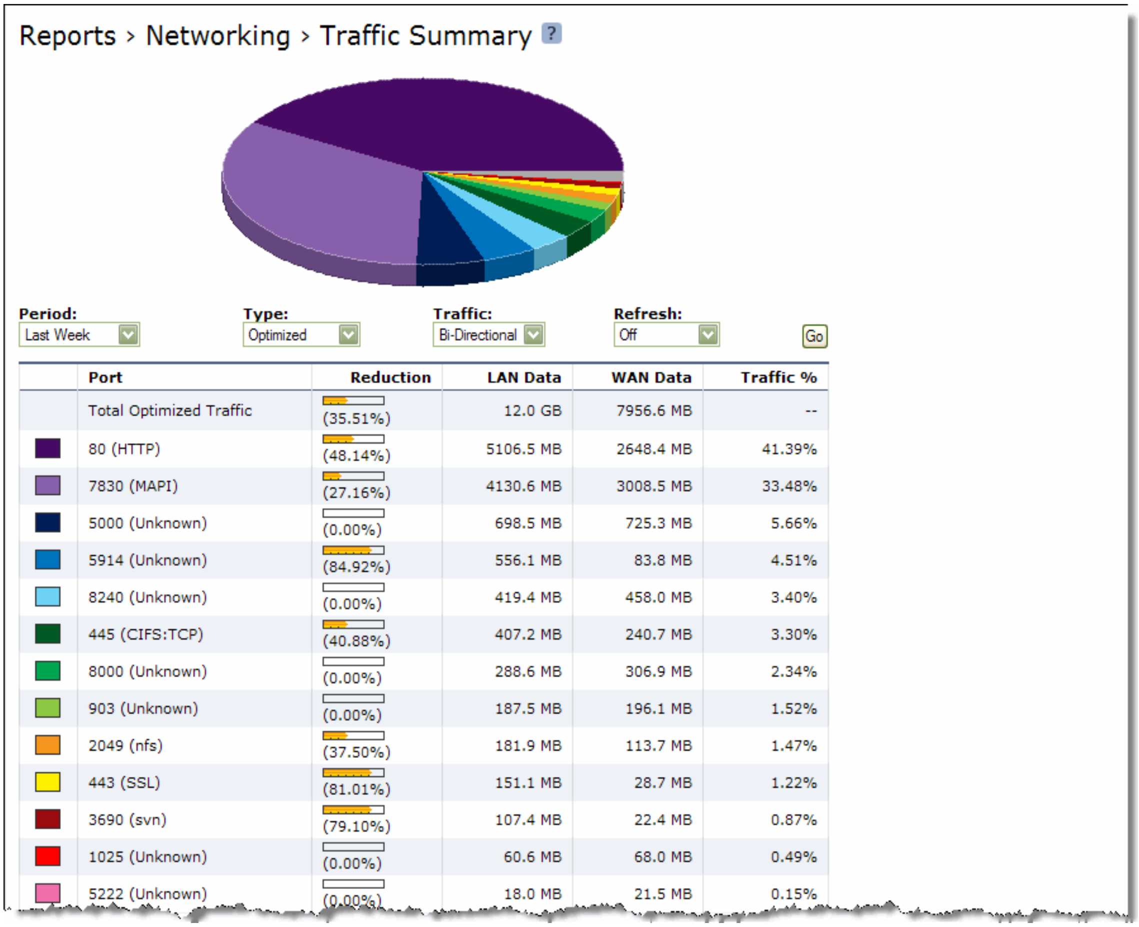

Reports > Networking > Traffic Summary

Period: Last Week

Type: Optimized

Traffic: Bi-Directional

Refresh: Off

	Port	Reduction	LAN Data	WAN Data	Traffic %
	Total Optimized Traffic	(35.51%)	12.0 GB	7956.6 MB	--
🟪	80 (HTTP)	(48.14%)	5106.5 MB	2648.4 MB	41.39%
🟪	7830 (MAPI)	(27.16%)	4130.6 MB	3008.5 MB	33.48%
🟦	5000 (Unknown)	(0.00%)	698.5 MB	725.3 MB	5.66%
🟦	5914 (Unknown)	(84.92%)	556.1 MB	83.8 MB	4.51%
🟦	8240 (Unknown)	(0.00%)	419.4 MB	458.0 MB	3.40%
🟩	445 (CIFS:TCP)	(40.88%)	407.2 MB	240.7 MB	3.30%
🟩	8000 (Unknown)	(0.00%)	288.6 MB	306.9 MB	2.34%
🟩	903 (Unknown)	(0.00%)	187.5 MB	196.1 MB	1.52%
🟧	2049 (nfs)	(37.50%)	181.9 MB	113.7 MB	1.47%
🟨	443 (SSL)	(81.01%)	151.1 MB	28.7 MB	1.22%
🟥	3690 (svn)	(79.10%)	107.4 MB	22.4 MB	0.87%
🟥	1025 (Unknown)	(0.00%)	60.6 MB	68.0 MB	0.49%
🟪	5222 (Unknown)	(0.00%)	18.0 MB	21.5 MB	0.15%

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Application	Select an application from the drop-down list. The default value is All.
Traffic	Select a traffic direction (Bi-Directional, WAN-to-LAN, or LAN-to-WAN) from the drop-down list.

Control	Description
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Interface Counters

The Interface Counters report summarizes the statistics for the primary, in-path LAN and WAN, and auxiliary interfaces. It also displays the IP address, speed, duplex, MAC address, and current status of each interface. For auto-negotiated speed and duplex settings the Interface Counters report displays the speed at which they are negotiated.

The Interface Counters report displays the statistics described in the following table.

Counter	Description
Interface	**LAN** - Specifies the IP address for the LAN interface.
	WAN - Specifies the IP address for the WAN interface.
	Primary - Specifies the IP address for the primary interface.
	Auxiliary Interface - Specifies the IP address for the auxiliary interface.
IP	Specifies the IP address for the interface.
Ethernet	Specifies the MAC address, speed, and duplex setting for the interface. Use this information to troubleshoot speed and duplex problems. Make sure the speed for the Steelhead appliance matches the WAN or LAN interfaces. Riverbed recommends setting the speed to 100 and duplex to full.
Link	Specifies true or false to indicate whether the link is up or down.
Receive Packets	Specifies the total number of packets, packets discarded, errors encountered, packets overrun, frames sent, and multicast packets sent.
Transmit Packets	Specifies the total number packets, packets discarded, errors encountered, packets overrun, carriers used, and collisions encountered.

Note: If you have multiple dual port, four-port, or six-port bypass cards installed, the Configure > Networking > Interface Counters report displays the interface statistics for each LAN and WAN port.

What This Report Tells You

The Interface Counters report answers the following questions:

- How many packets are being transmitted?
- Are there any errors occurring during the packet transmissions?
- What is the current status of the interface?

To view interface counters

- Choose Reports > Networking > Interface Counters to display the Interface Counters page.

Figure 10-10. Interface Counters Page

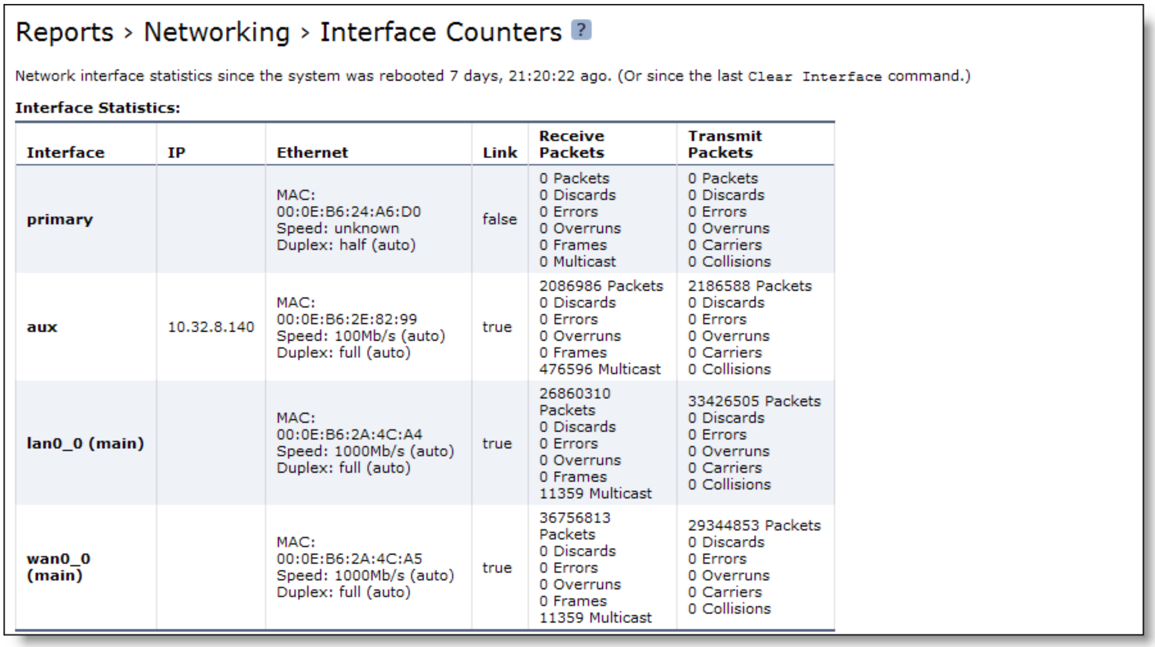

Reports > Networking > Interface Counters ?

Network interface statistics since the system was rebooted 7 days, 21:20:22 ago. (Or since the last `Clear Interface` command.)

Interface Statistics:

Interface	IP	Ethernet	Link	Receive Packets	Transmit Packets
primary		MAC: 00:0E:B6:24:A6:D0 Speed: unknown Duplex: half (auto)	false	0 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 0 Multicast	0 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions
aux	10.32.8.140	MAC: 00:0E:B6:2E:82:99 Speed: 100Mb/s (auto) Duplex: full (auto)	true	2086986 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 476596 Multicast	2186588 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions
lan0_0 (main)		MAC: 00:0E:B6:2A:4C:A4 Speed: 1000Mb/s (auto) Duplex: full (auto)	true	26860310 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 11359 Multicast	33426505 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions
wan0_0 (main)		MAC: 00:0E:B6:2A:4C:A5 Speed: 1000Mb/s (auto) Duplex: full (auto)	true	36756813 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 11359 Multicast	29344853 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing TCP Statistics Reports

The TCP Statistics report summarizes TCP statistics for the appliance.

The TCP Statistics report contains the following table of statistics that summarize TCP activity.

Packet Type	Description
Packets Received	Specifies the total packets received.
Packets Sent	Specifies the total TCP packets sent.
Packets Retransmitted	Specifies the total TCP packets retransmitted.

Packet Type	Description
Packets Fast Retransmitted	Specifies the total TCP packets fast retransmitted. Fast retransmit is an enhancement to TCP which reduces the time a sender waits before retransmitting a lost segment. If an acknowledgement is not received for a particular segment with a specified time (a function of the estimated round-trip delay time), the sender assumes the segment was lost in the network, and retransmits the segment.
Time-outs	Specifies the number of time-outs.
Loss Events	Specifies the total number of loss events.

What This Report Tells You

The TCP Statistics report answers the following questions:

- How many TCP packets have been sent and received?
- How many TCP packets have been retransmitted?
- How many time-outs have occurred?
- How many loss events have occurred?

To view the TCP Statistics report

- Choose Reports > Networking > TCP Statistics to display the TCP Statistics page.

Figure 10-11. TCP Statistics Page

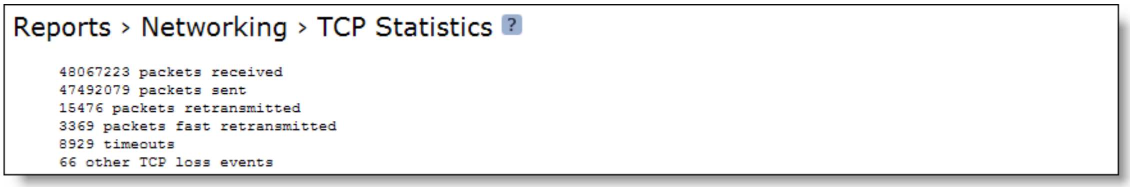

Reports › Networking › TCP Statistics ?

```
48067223 packets received
47492079 packets sent
15476 packets retransmitted
3369 packets fast retransmitted
8929 timeouts
66 other TCP loss events
```

Tip: To print the report, choose **File** > **Print** in your Web browser to open the Print dialog box.

Viewing Optimized Throughput Reports

The Optimized Throughput report summarizes the throughput or total TCP data transmitted for the application and time period specified.

The Optimized Throughput report includes LAN and WAN Link Throughput graphs which include the following statistics that describe data activity for the application and the time period you specify.

Field	Description
Peak WAN/LAN Throughput at <time> on <date>	Specifies the date and time of the peak data activity.

Field	Description
95th Percentile WAN/LAN Throughput	Specifies the 95th percentile for data activity. The 95th percentile is calculated by taking the peak of the lower 95% of inbound and outbound throughput samples.
Average WAN/LAN Throughput	Specifies the average amount of data transmitted.

What This Report Tells You

The Optimized Throughput report answers the following questions:

- What was the average throughput?
- What was the peak throughput?
- At what time did the peak throughput occur?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Optimized Throughput report

1. Choose Reports > Optimization > Throughput to display the Optimized Throughput page.

Figure 10-12. Optimized Throughput Page

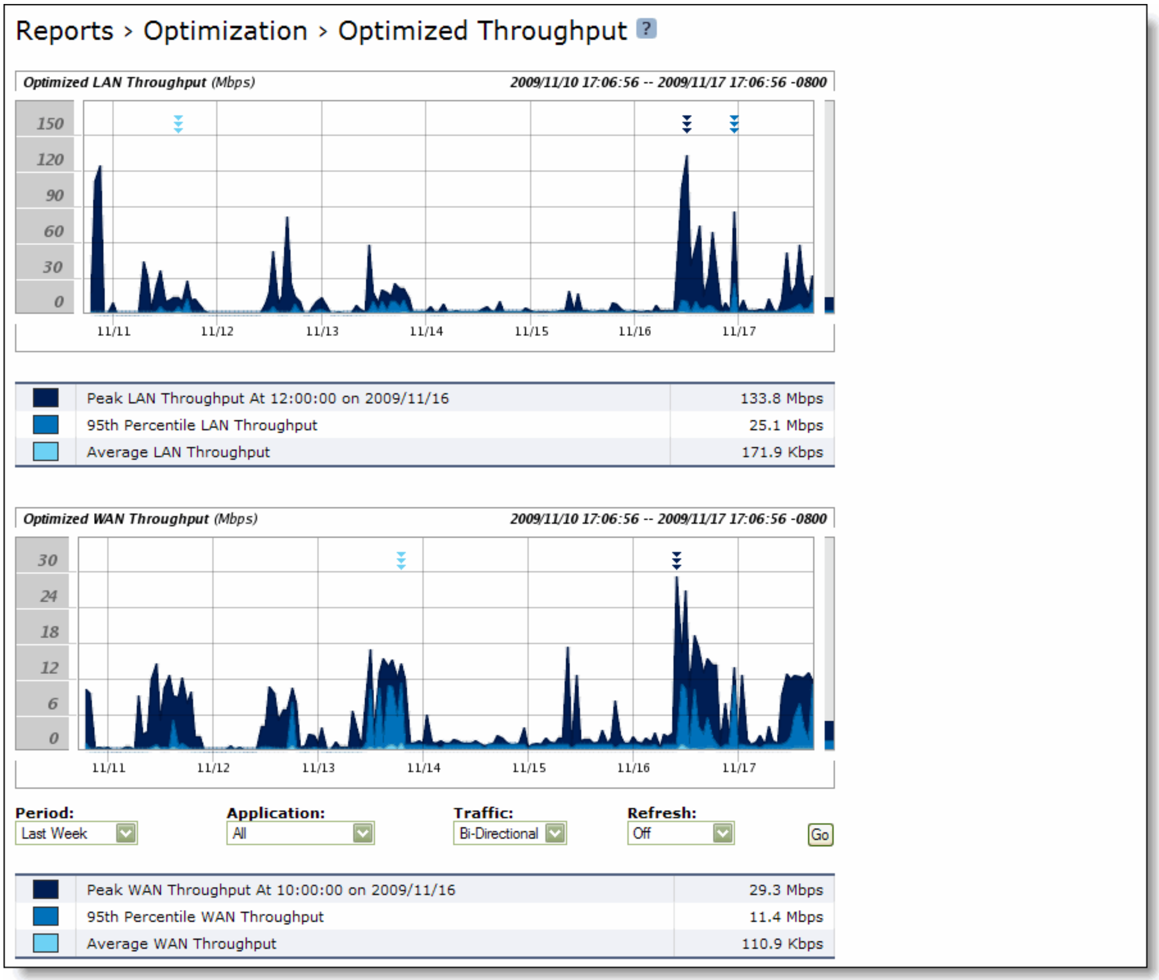

■	Peak LAN Throughput At 12:00:00 on 2009/11/16	133.8 Mbps
■	95th Percentile LAN Throughput	25.1 Mbps
■	Average LAN Throughput	171.9 Kbps

Period: Last Week
Application: All
Traffic: Bi-Directional
Refresh: Off

■	Peak WAN Throughput At 10:00:00 on 2009/11/16	29.3 Mbps
■	95th Percentile WAN Throughput	11.4 Mbps
■	Average WAN Throughput	110.9 Kbps

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Application	Select an application from the drop-down list. The default value is All.
Traffic	Select a traffic direction (Bi-Directional, WAN-to-LAN, or LAN-to-WAN) from the drop-down list.

Control	Description
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Bandwidth Optimization Reports

The Bandwidth Optimization report summarizes the overall inbound and outbound bandwidth improvements for TCP on your network. You can create reports according to the time period of your choice, application, and type of traffic.

The Bandwidth Optimization report includes the following table of statistics that describe bandwidth activity for the time period you specify.

Field	Description
WAN Data	Depending on which Traffic type you select, specifies one the following: • **Bi-Directional** - bytes sent and received on the WAN and LAN ports • **WAN-to-LAN** - bytes received on the WAN and sent out the LAN • **LAN-to-WAN** - bytes received on the LAN and sent out the WAN
LAN Data	Depending on which Traffic type you select, specifies one the following: • **Bi-Directional** - bytes sent and received on the WAN and LAN ports • **WAN-to-LAN** - bytes received on the WAN and sent out the LAN • **LAN-to-WAN** - bytes received on the LAN and sent out the WAN
Total Data Reduction %	Specifies the total decrease of data transmitted over the WAN, according to the following calculation: (Data In – Data Out)/(Data In)
Peak Data Reduction Occurred At <time> on <date>	Specifies the date and time that the peak data reduction occurred.
Optimized Bandwidth Capacity Increase	Specifies the increase in the amount of data transmitted over the WAN, according to the following calculation: 1/(1-Reduction Rate)

What This Report Tells You

The Bandwidth Optimization report answers the following questions:

- How much bandwidth optimization has occurred?

- What was the average and peak amount of data sent?
- What was the overall increase in the amount of data that can be transmitted using the Steelhead appliance?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view a Bandwidth Optimization report

1. Choose Reports > Optimization > Bandwidth Optimization to display the Bandwidth Optimization page.

Figure 10-13. Bandwidth Optimization Page

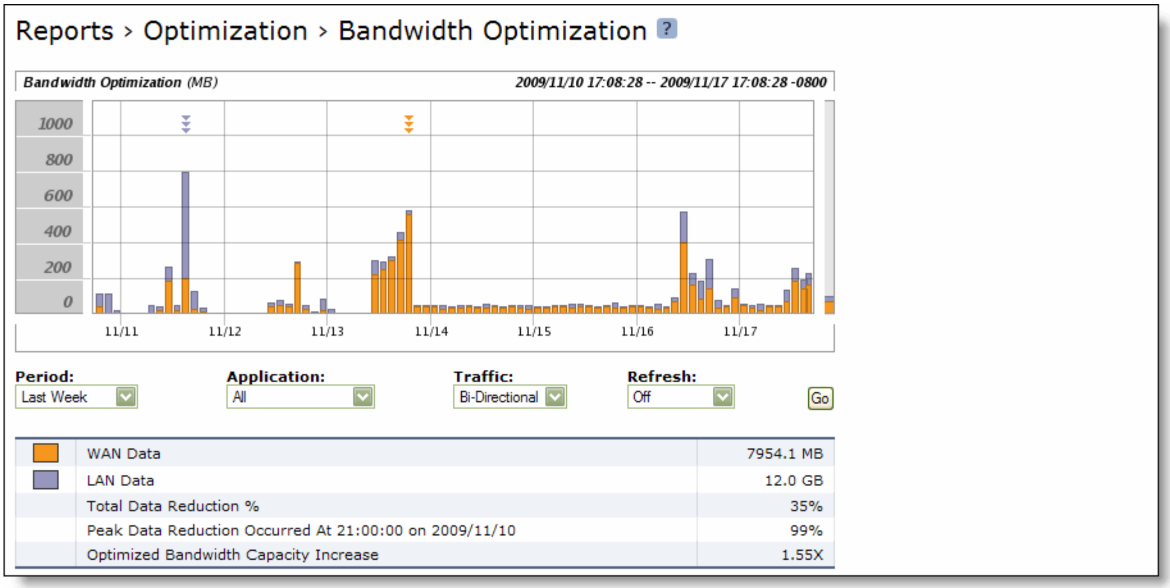

Reports > Optimization > Bandwidth Optimization ?

Period: Last Week

Application: All

Traffic: Bi-Directional

Refresh: Off

Go

🟧	WAN Data	7954.1 MB
🟪	LAN Data	12.0 GB
	Total Data Reduction %	35%
	Peak Data Reduction Occurred At 21:00:00 on 2009/11/10	99%
	Optimized Bandwidth Capacity Increase	1.55X

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Application	Select an application from the drop-down list. The default value is All.
Traffic	Select a traffic direction (Bi-Directional, WAN-to-LAN, or LAN-to-WAN) from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Reduction Reports

The Data Reduction report summarizes the percent reduction of data transmitted by an application such as FTP, HTTP, NetBIOS and TCP, traffic in CIFS, and MAPI.

The Data Reduction report includes the following table of statistics that describe data reduction for the application and the time period you specify.

Field	Description
Total Data Reduction %	Specifies the total decrease of data transmitted over the WAN.
Peak Data Reduction At <time> on <date>	Specifies the date and time that the peak data reduction occurred.
Optimized Bandwidth Capacity Increase	Specifies the increase in the amount of the data that can be transmitted over the WAN.

What This Report Tells You

The Data Reduction report answers the following questions:

- What was the total reduction in the amount of data that can be transmitted for each application?
- What was the peak reduction in the amount of data transmitted for each application?
- What was the total increase of data transmitted for the application and time period specified?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Data Reduction report

1. Choose Reports > Optimization > Data Reduction to display the Data Reduction page.

Figure 10-14. Data Reduction Page

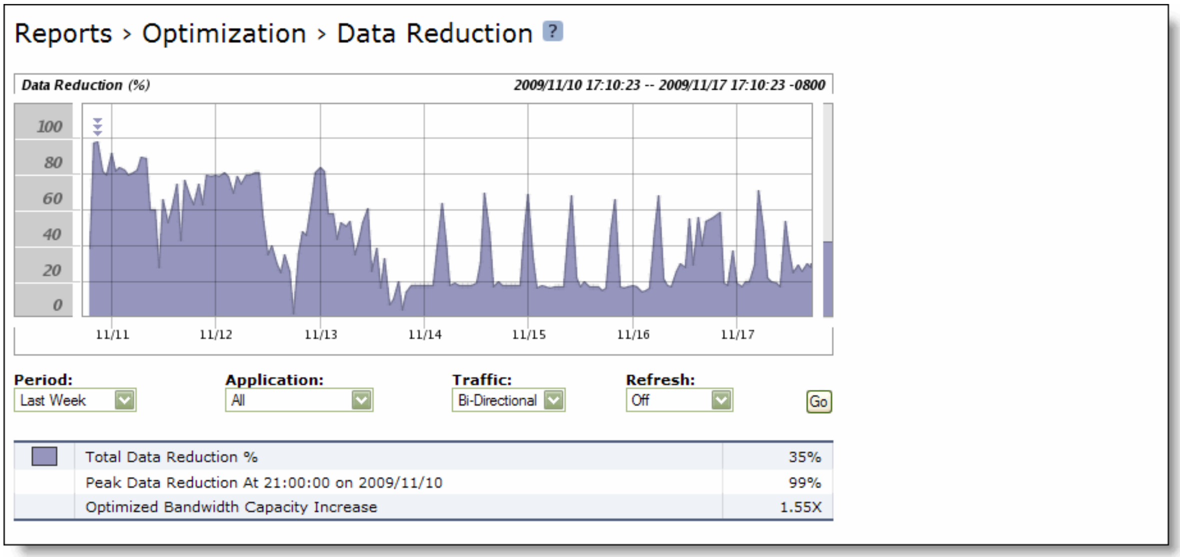

Reports > Optimization > Data Reduction ?

Period: Last Week

Application: All

Traffic: Bi-Directional

Refresh: Off

▨	Total Data Reduction %	35%
	Peak Data Reduction At 21:00:00 on 2009/11/10	99%
	Optimized Bandwidth Capacity Increase	1.55X

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, specify the Start Time and End Time. Use the following format: YYYY/MM/DD HH:MM:SS
Application	Select an application from the drop-down list. The default value is All.
Traffic	Select a traffic direction (Bi-Directional, WAN-to-LAN, or LAN-to-WAN) from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select **1**0 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Connected Appliances Reports

The Connected Appliances report lists the remote Steelhead appliances that are connected to the system.

Note: Steelhead appliances might remain listed in the report for a short time after they have been shut down or renamed. Select Show Only Online Peers to clear the report of any peers that were once connected but are no longer available.

What This Report Tells You

The Connected Appliances report answers the following questions:

- What remote Steelhead appliances are connected to this Steelhead appliance?
- Is there an incompatibility issue between the Steelhead appliance software versions?

To view the Connected Appliances report

1. Choose Reports > Appliance > Connected Appliances to display the Connected Appliances page.

2. By default, all peers are displayed, including Steelhead appliances that have been shut down or renamed. To view only the online peers, select Show Only Online Peers.

Figure 10-15. Connected Appliances Page

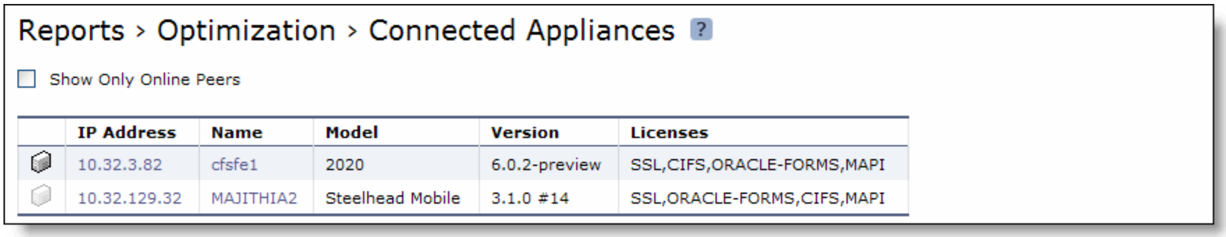

Reports > Optimization > Connected Appliances ?

□ Show Only Online Peers

	IP Address	Name	Model	Version	Licenses
■	10.32.3.82	cfsfe1	2020	6.0.2-preview	SSL,CIFS,ORACLE-FORMS,MAPI
□	10.32.129.32	MAJITHIA2	Steelhead Mobile	3.1.0 #14	SSL,ORACLE-FORMS,CIFS,MAPI

Tip: Select the appliance name or IP address in the Connected Appliance list to open the Management Console for the specified system in a new browser window.

Tip: Place the cursor over the icon to the left of the IP address to display the appliance license status: expired or valid.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Connection Pooling

The Connection Pooling report summarizes the current connection pool of connections to peer appliances. It contains the following table of statistics that summarize connection pooling activity.

Field	Description
Total Requests	Specifies the total number of requests for connections to peer appliances.
Total Hits	Specifies the total number of successful connections and connections that are serviced by already existing inner channel connections.
Peak Hits At <time> on <date>	Specifies the date and time of the peak number of successful connections and connections that are serviced by already existing inner channel connections.

The connection pool holds many idle TCP connections up to the maximum pool size. When a client requests a new connection to a previously visited server, the pool manager checks the pool for unused connections, returns one if available, and then replenishes the pool with another idle connection. For details on setting the maximum pool size, see “Configuring General Service Settings” on page 60.

Note: A slight delay might occur during the time it takes the pool manager to check for an unused connection, pull the connection out of the pool, and then refill it. In a Steelhead appliance with a very active connection count, this report might indicate a high amount of requests before the pool manager has time to establish new connections and refill the pool. On a very busy link, the entire pool could drain before the pool manager refills it. In addition, the pool manager refills the pool one connection at a time, so when the appliance receives bursty connection requests, it might take some time to refill the pool. A couple of bursts in succession can drain the pool. Network congestion can also lengthen the pool refill time.

What This Report Tells You

The Connection Pooling report answers the following questions:

- How large is the pool of connections?
- How many connections occurred?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Connection Pooling report

1. Choose Reports > Optimization > Connection Pooling to display the Connection Pooling page.

Figure 10-16. Connection Pooling Page

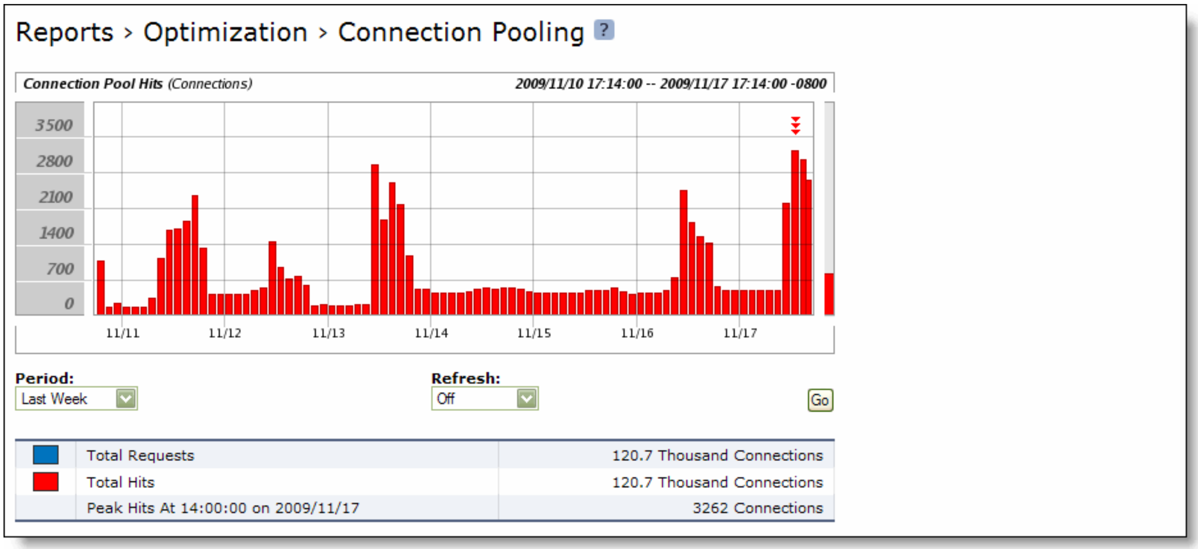

Reports › Optimization › Connection Pooling

Period: Last Week

Refresh: Off

🟦	Total Requests	120.7 Thousand Connections
🟥	Total Hits	120.7 Thousand Connections
	Peak Hits At 14:00:00 on 2009/11/17	3262 Connections

2. Use the controls to customize the report, as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing CIFS Prepopulation Share Log Reports

The prepopulation share logs provide detail regarding the initial copy of the share or the last share synchronization.

To view CIFS prepopulation share logs

1. Choose Configure > Optimization > CIFS Prepopulation to display the CIFS Prepopulation page.

2. Select the remote path for the share.

3. Click **Initial Copy** or **Last Sync**.

The initial copy or last sync log appears.

Figure 10-17. CIFS Prepopulation Share Initial Copy Log

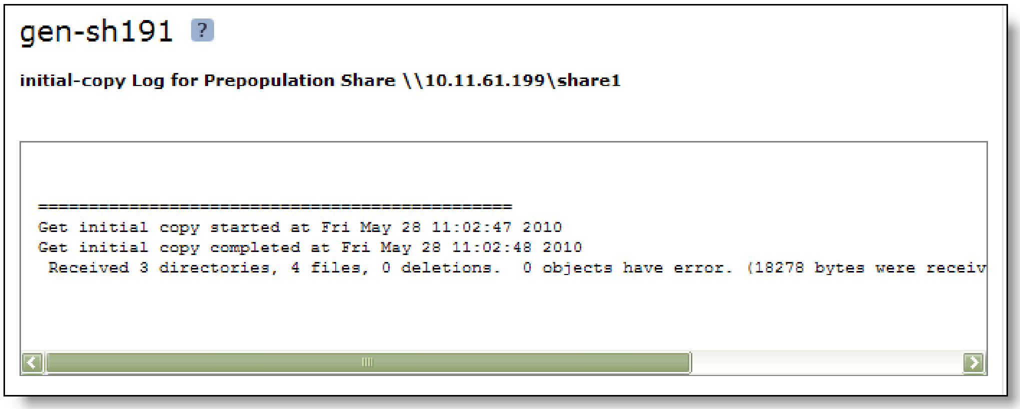

gen-sh191

initial-copy Log for Prepopulation Share \\10.11.61.199\share1

```
==================================================
Get initial copy started at Fri May 28 11:02:47 2010
Get initial copy completed at Fri May 28 11:02:48 2010
 Received 3 directories, 4 files, 0 deletions.  0 objects have error. (18278 bytes were receiv
```

Figure 10-18. CIFS Prepopulation Share Sync Log

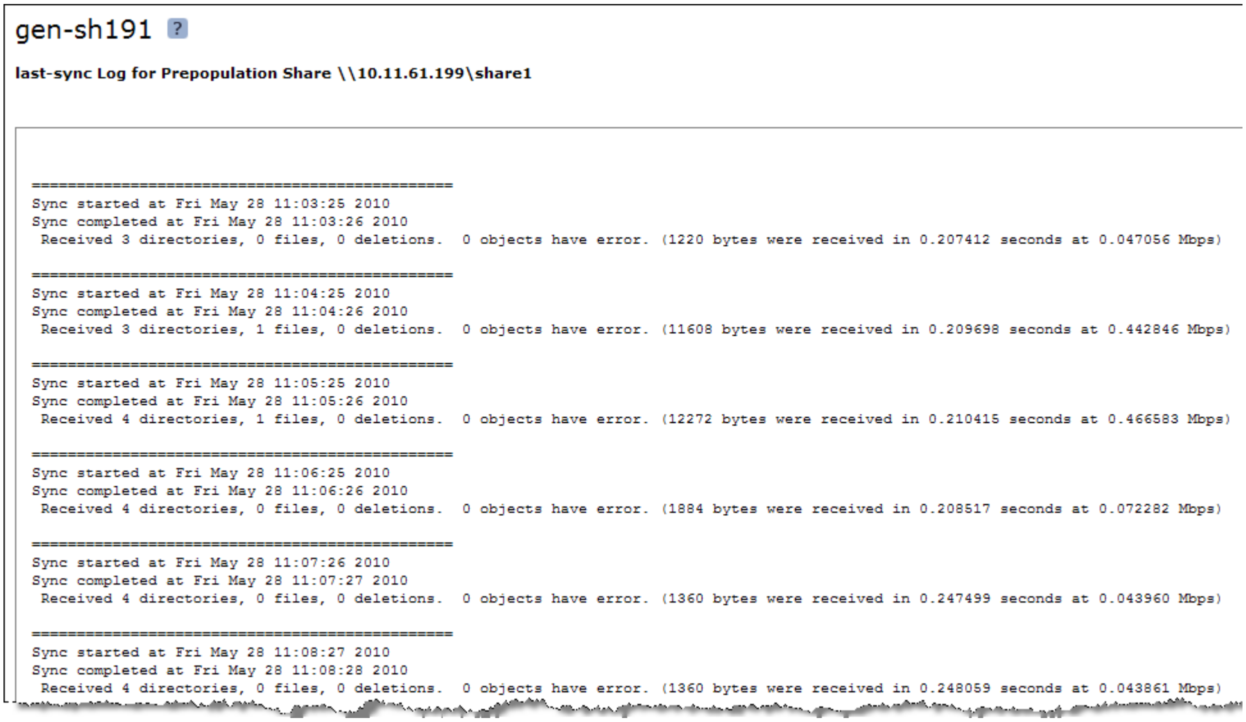

gen-sh191

last-sync Log for Prepopulation Share \\10.11.61.199\share1

```
==================================================
Sync started at Fri May 28 11:03:25 2010
Sync completed at Fri May 28 11:03:26 2010
 Received 3 directories, 0 files, 0 deletions.  0 objects have error. (1220 bytes were received in 0.207412 seconds at 0.047056 Mbps)

==================================================
Sync started at Fri May 28 11:04:25 2010
Sync completed at Fri May 28 11:04:26 2010
 Received 3 directories, 1 files, 0 deletions.  0 objects have error. (11608 bytes were received in 0.209698 seconds at 0.442846 Mbps)

==================================================
Sync started at Fri May 28 11:05:25 2010
Sync completed at Fri May 28 11:05:26 2010
 Received 4 directories, 1 files, 0 deletions.  0 objects have error. (12272 bytes were received in 0.210415 seconds at 0.466583 Mbps)

==================================================
Sync started at Fri May 28 11:06:25 2010
Sync completed at Fri May 28 11:06:26 2010
 Received 4 directories, 0 files, 0 deletions.  0 objects have error. (1884 bytes were received in 0.208517 seconds at 0.072282 Mbps)

==================================================
Sync started at Fri May 28 11:07:26 2010
Sync completed at Fri May 28 11:07:27 2010
 Received 4 directories, 0 files, 0 deletions.  0 objects have error. (1360 bytes were received in 0.247499 seconds at 0.043960 Mbps)

==================================================
Sync started at Fri May 28 11:08:27 2010
Sync completed at Fri May 28 11:08:28 2010
 Received 4 directories, 0 files, 0 deletions.  0 objects have error. (1360 bytes were received in 0.248059 seconds at 0.043861 Mbps)
```

The logs contain the following statistics that summarize prepopulation share activity.

Field	Description
Initial-copy log	Displays the date and time the initial share copy started and completed. The log includes how many directories, files, and bytes were received and how long it took to receive them. The log also lists any errors or deletions.
Last-sync log	Displays the date and time the last share synchronization started and completed. The log includes how many directories, files, and bytes were received and how long it took to receive them. The log also lists any errors or deletions.

Tip: To print the log, choose File > Print in your Web browser to open the Print dialog box.

Related Topic

- "Configuring CIFS Prepopulation" on page 86

Viewing HTTP Reports

The HTTP report summarizes HTTP optimization statistics for the time period specified. The HTTP report contains the HTTP (%) Hits graph, which displays the following statistics that summarize HTTP data activity. For details, see "Configuring HTTP Optimization" on page 111.

Field	Description
Total Hits %	Specifies the total percentage of HTTP objects served by all optimization schemes: • **URL Learning** - The Steelhead appliance learns associations between a base request and a follow-on request. Instead of saving each object transaction, the Steelhead appliance saves only the request URL of object transactions in a Knowledge Base and then generates related transactions from the list. • **Parse and Prefetch** - The Steelhead appliance determines which objects are going to be requested for a given Web page and prefetches them so that they are readily available when the client makes its requests. • **Object Prefetch** - The Steelhead appliance stores object prefetches from HTTP GET requests for cascading style sheets, static images, and Java scripts in the Object Prefetch Table.
Parse and Prefetch Hit %	Specifies the percentage of embedded objects that were successfully found and prefetched.
URL Learning Hit %	Specifies the percentage of base requests and follow-on requests that were found.
Object Prefetch Hit %	Specifies the percentage of object prefetches that were stored.
Total Objects Requested	Specifies the number of HTTP objects requested by clients.
Total Objects Hit	Specifies the total number of HTTP objects served locally by all optimization schemes: URL Learning, Parse and Prefetch, and Object Prefetch.
Parse and Prefetch Hits	Specifies how many embedded objects were served locally due to Parse and Prefetch.
URL Learning Hits	Specifies how many URLs were served due to URL learning.
Object Prefetch Hits	Specifies how many requests were found and served from the Object Prefetch Table.
Misses	Specifies the total number of HTTP objects, URLs, and object prefetches which were requested but not stored or prefetched.

What This Report Tells You

The HTTP report answers the following questions:

- What was the overall percent increase in HTTP data transmitted over the WAN?
- How many HTTP objects were requested?
- How many HTTP objects were successfully obtained and transmitted over the WAN?
- How many object prefetches and prefetch hits occurred per HTTP object?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the HTTP report

1. Choose Reports > Optimization > HTTP to display the HTTP page.

Figure 10-19. HTTP Page

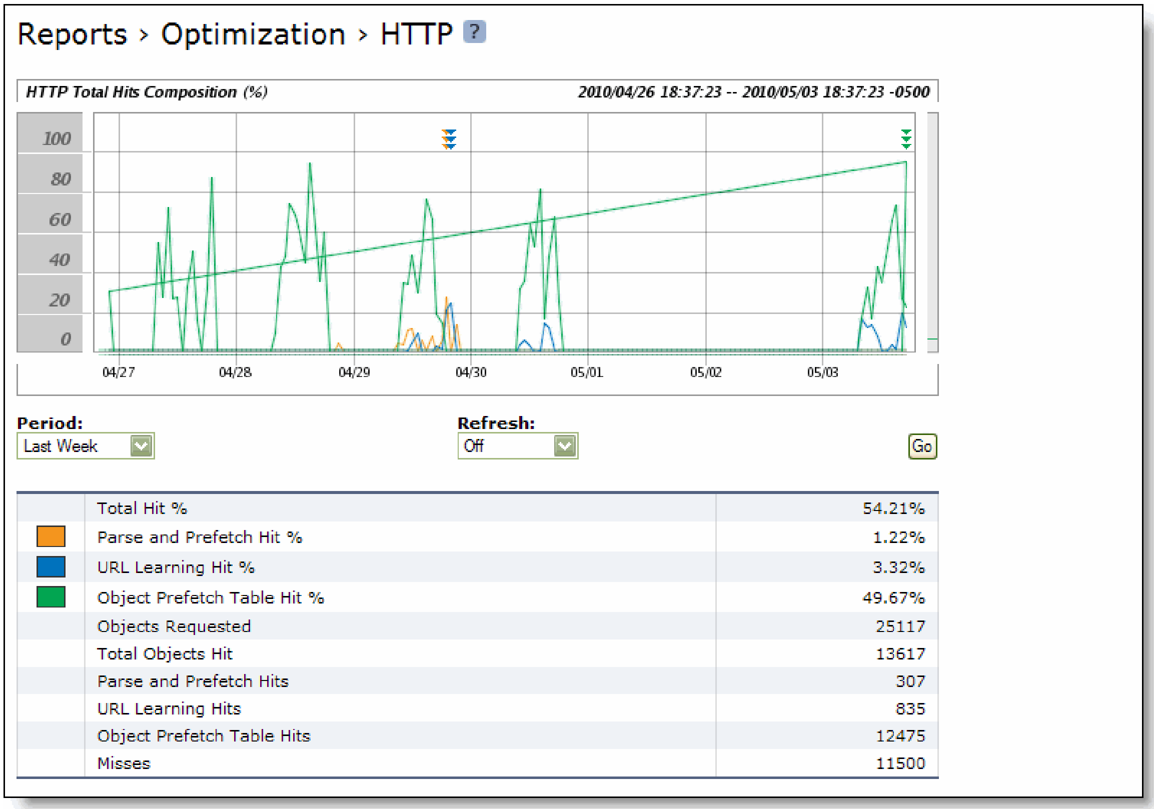

Reports > Optimization > HTTP ?

HTTP Total Hits Composition (%) 2010/04/26 18:37:23 -- 2010/05/03 18:37:23 -0500

Period: Last Week

Refresh: Off

Go

	Total Hit %	54.21%
🟧	Parse and Prefetch Hit %	1.22%
🟦	URL Learning Hit %	3.32%
🟩	Object Prefetch Table Hit %	49.67%
	Objects Requested	25117
	Total Objects Hit	13617
	Parse and Prefetch Hits	307
	URL Learning Hits	835
	Object Prefetch Table Hits	12475
	Misses	11500

2. Use the controls to customize the report, as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing NFS Reports

The NFS report summarizes NFS optimization statistics for the time period specified.

The NFS report contains the following table of statistics that summarize NFS activity.

Field	Description
Local Responses	Specifies the number of NFS calls that were responded to locally.
Remote Responses	Specifies the number of NFS calls that were responded to remotely (that is, calls that traversed the WAN to the NFS server).
Total Delayed	Specifies the delayed calls which were responded to locally but not immediately (for example, reads which were delayed while a read ahead was occurring and were responded to from the data in the read ahead).
Total Reduction %	Specifies the percentage decrease of NFS calls over the WAN. For example, you might see an 85% reduction in NFS data (see the Data Reduction or the Traffic Summary report) and a 55% reduction in the number of NFS calls over the WAN (NFS Statistics report).
Peak Reduction % At <time> on <date>	Specifies the percentage of reduction for the date and time.
Capacity Increase	Specifies the increase in the number of NFS calls that can be transmitted over the WAN.

What This Report Tells You

The NFS report answers the following questions:

- How many NFS calls were answered locally and remotely?
- How many delayed calls occurred for NFS activity?
- What is the reduction in the number of NFS calls that went to the server?
- What was the overall decrease in NFS calls transmitted over the WAN?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the NFS report

1. Choose Reports > Optimization > NFS to display the NFS page.

Figure 10-20. NFS Page

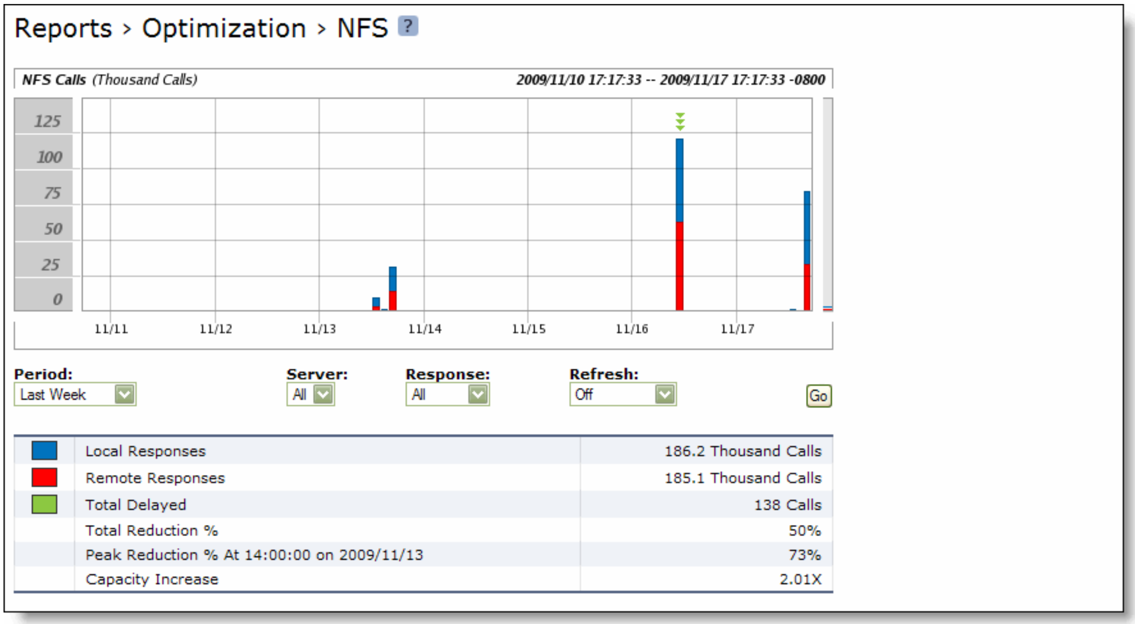

Reports > Optimization > NFS

Period: Last Week

Server: All

Response: All

Refresh: Off

Go

🟦	Local Responses	186.2 Thousand Calls
🟥	Remote Responses	185.1 Thousand Calls
🟩	Total Delayed	138 Calls
	Total Reduction %	50%
	Peak Reduction % At 14:00:00 on 2009/11/13	73%
	Capacity Increase	2.01X

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Server	Select the server for which you want to collect statistics from the drop-down list.
Response	Select a response (All, Local, Remote, or Delayed) from the drop-down list. The default value is All.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing SSL Reports

The SSL report summarizes the SSL connection requests and connection rate for the time period specified.

The SSL report contains the following graphs:

- **SSL Connection Requests (Connections)** - Summarizes the connection requests for the time period specified. The Connection Requests graph includes the following table of statistics that describe data activity for the application and the time period you specify.

Field	Description
Number of Established Sessions	Specifies the number of established SSL connections.
Number of Requests	Specifies the number of SSL requests.
Number of Failed Connections	Specifies the number of failed SSL connections.
Number of Current Established Connections	Specifies the total number of current established SSL connections.

- **SSL Connection Rate (Connections Per Second)** - Summarizes the average number of successfully completed SSL connections in one second. The SSL connection rate is also called SSL TPS (SSL Transactions per Second). The Connection Rate graph includes the following table of statistics that describe data activity for the application and the time period you specify.

Field	Description
Average Connection Rate	Specifies the average connection rate for SSL connections.
Peak Connection Rate At <time> on <date>	Specifies the peak connection rate for SSL connections for the date and time.

What This Report Tells You

The SSL report answers the following questions:

- What is the number of established SSL connections?
- What is the number of SSL requests during a specified period of time?
- What is the number of failed connections during a specified period of time?
- What is the number of concurrent connections open at the current time?

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the SSL report

1. Choose Reports > Optimization > SSL to display the SSL page.

Figure 10-21. SSL Page

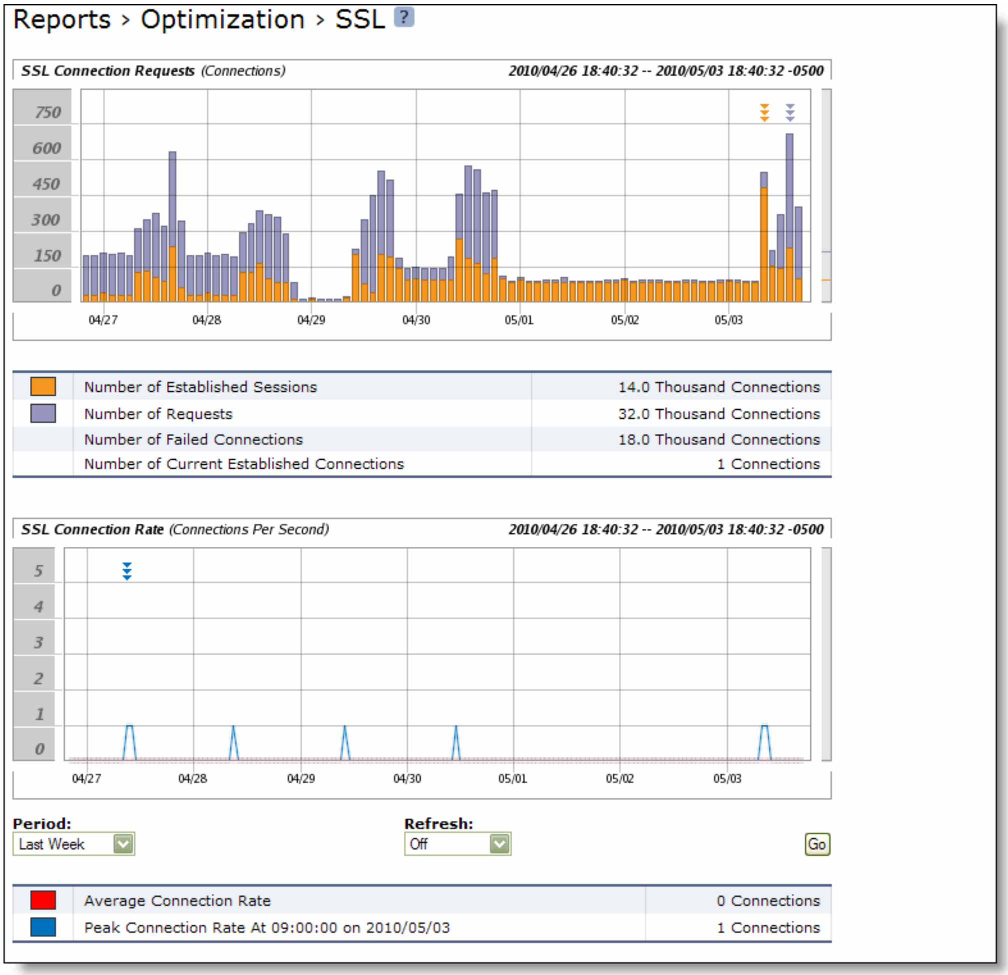

🟧	Number of Established Sessions	14.0 Thousand Connections
🟪	Number of Requests	32.0 Thousand Connections
	Number of Failed Connections	18.0 Thousand Connections
	Number of Current Established Connections	1 Connections

Period: Last Week

Refresh: Off

Go

🟥	Average Connection Rate	0 Connections
🟦	Peak Connection Rate At 09:00:00 on 2010/05/03	1 Connections

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store Status Reports

The Data Store Status report summarizes the current status and state of the data store synchronization process.

If you have enabled data store synchronization, it summarizes the state of the replication process. For details, see "Synchronizing Peer Datastores" on page 78.

The Data Store Status report contains the following table of statistics that summarizes data store activity.

Field	Description
Synchronization Connection	Indicates the status of the connection between the synched Steelheads.
Synchronization Catch-Up	Indicates the status of transferring data between the synched Steelheads. Catch-Up is used for synching data that was not synched during the Keep-Up phase.
Synchronization Keep-Up	Indicates the status of transferring new incoming data between the synched Steelheads.
Data Store Percentage Used (Since Last Clear)	Specifies the percentage of the datastore that is used.

What This Report Tells You

The Data Store Status report answers the following questions:

- Is the synchronization connection active?
- Is the Steelhead appliance in the Catch-up or Keep-up phase of datastore synchronization?
- What percentage of the datastore is unused?

To view the Data Store Status report

- Choose Reports > Optimization > Data Store to display the Date Store Status page.

Figure 10-22. Data Store Status Page

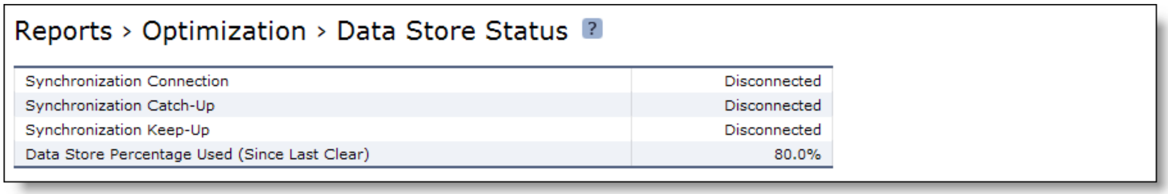

Reports > Optimization > Data Store Status

Synchronization Connection	Disconnected
Synchronization Catch-Up	Disconnected
Synchronization Keep-Up	Disconnected
Data Store Percentage Used (Since Last Clear)	80.0%

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store SDR-Adaptive Reports

The Data Store SDR-Adaptive report summarizes:

- how much adaptive compression is occurring in the datastore using legacy mode. The report combines both the percentage due to local and remote adaptive compression (as signaled by the peers).
- the percentage of the traffic, in bytes, which is adapted to in-memory-only (or transient), compared to the total SDR traffic (SDR-adaptive mode).

The report contains the following table of statistics that summarizes datastore adaptive compression activity, shown as a percent of total SDR data.

Note: You must enable the SDR-Adaptive setting before creating this report. For details, see "Setting an Adaptive Streamlining Mode" on page 83.

Field	Description
Maximum Compression Due To Disk Pressure at <time> on <date>	Specifies the maximum compression due to disk pressure for the date and time.
Minimum Compression Due To Disk Pressure at <time> on <date>	Specifies the minimum compression due to disk pressure for the date and time.
Average Compression Due To Disk Pressure	Specifies the average compression due to disk pressure for the date and time.
Maximum Compression Due To In-Path Rule at <time> on <date>	Specifies the maximum compression due to in-path rule for the date and time.
Minimum Compression Due To In-Path Rule at <time> on <date>	Specifies the minimum compression due to in-path rule for the date and time.
Average Compression Due To In-Path Rule	Specifies the average compression due to in-path rule for the date and time.
Maximum In-Memory SDR Due To Disk Pressure at <time> on <date>	Specifies the maximum in-memory SDR due to disk pressure for the date and time.
Minimum In-Memory SDR Due To Disk Pressure at <time> on <date>	Specifies the minimum in-memory SDR due to disk pressure for the date and time.
Average In-Memory SDR Due To Disk Pressure	Specifies the average in-memory SDR due to disk pressure for the date and time.
Maximum In-Memory SDR Due To In-Path Rule at <time> on <date>	Specifies the maximum in-memory SDR due to in-path rule for the date and time.
Minimum In-Memory SDR Due To In-Path Rule at <time> on <date>	Specifies the minimum in-memory SDR due to in-path rule for the date and time.
Average In-Memory SDR Due To In-Path Rule	Specifies the average in-memory SDR due to in-path rule for the date and time.

What This Report Tells You

The Data Store SDR-Adaptive report answers the following question:

- What is relative adaptive compression when SDR-Adaptive is enabled at various times of the day?

To view the Data Store SDR-Adaptive report

1. Choose Reports > Optimization > Data Store SDR-Adaptive to display the Data Store SDR-Adaptive page.

Figure 10-23. Data Store SDR-Adaptive Page

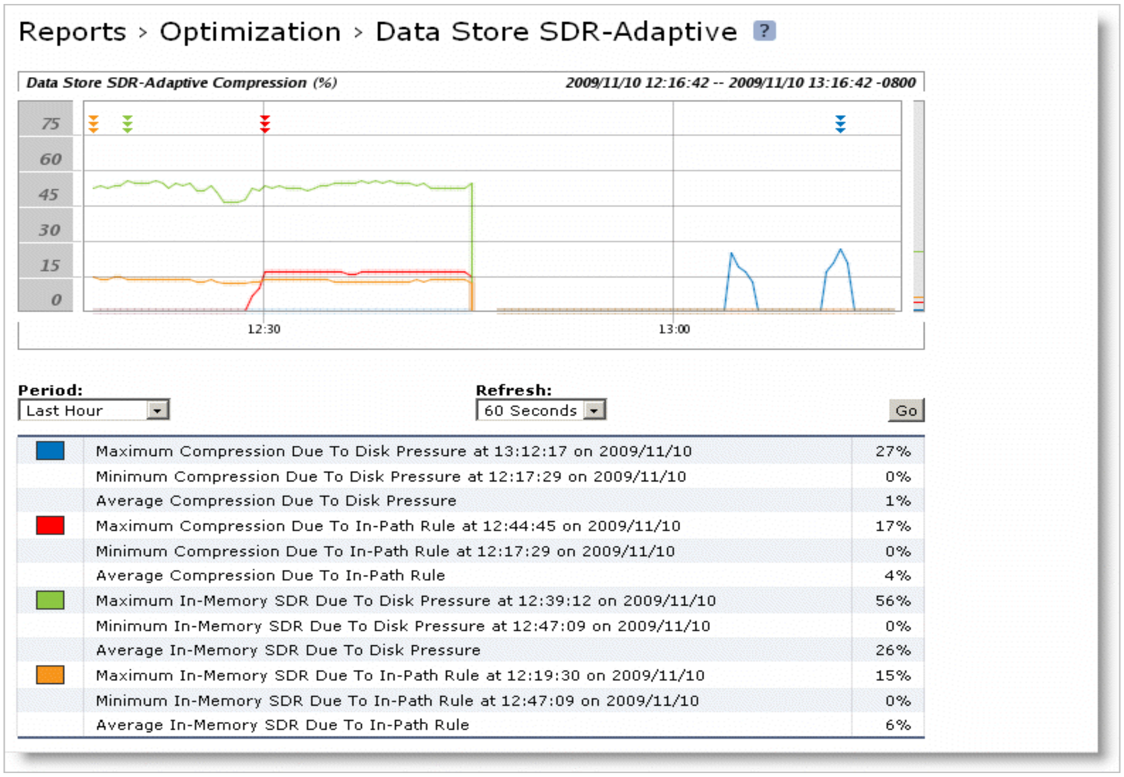
Reports > Optimization > Data Store SDR-Adaptive ?
Data Store SDR-Adaptive Compression (%)
2009/11/10 12:16:42 -- 2009/11/10 13:16:42 -0800
75
60
45
30
15
0
12:30
13:00
Go

Period: Last Hour

Refresh: 60 Seconds

🟦	Maximum Compression Due To Disk Pressure at 13:12:17 on 2009/11/10	27%
	Minimum Compression Due To Disk Pressure at 12:17:29 on 2009/11/10	0%
	Average Compression Due To Disk Pressure	1%
🟥	Maximum Compression Due To In-Path Rule at 12:44:45 on 2009/11/10	17%
	Minimum Compression Due To In-Path Rule at 12:17:29 on 2009/11/10	0%
	Average Compression Due To In-Path Rule	4%
🟩	Maximum In-Memory SDR Due To Disk Pressure at 12:39:12 on 2009/11/10	56%
	Minimum In-Memory SDR Due To Disk Pressure at 12:47:09 on 2009/11/10	0%
	Average In-Memory SDR Due To Disk Pressure	26%
🟧	Maximum In-Memory SDR Due To In-Path Rule at 12:19:30 on 2009/11/10	15%
	Minimum In-Memory SDR Due To In-Path Rule at 12:47:09 on 2009/11/10	0%
	Average In-Memory SDR Due To In-Path Rule	6%

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store Cost Reports

The Data Store Cost report summarizes the relative cost of doing datastore operations.

The Data Store Cost report includes a throughput graph which displays the following statistic that describes datastore segment throughput for the date and the time period you specify.

For details on datastore disk performance, see "Viewing Data Store Disk Load Reports" on page 441.

Field	Description
Maximum Cost	Displays the peak datastore segment throughput. A low value indicates that the total load on the datastore is healthy and not constrained by resources like disk. When the y-axis on this graph shows a high peak value for significant periods of time, the Steelhead appliance might be experiencing disk pressure during those times. Use this report in conjunction with the Data Store Disk Load report to identify disk pressure. For details, see "Optimizing the Datastore for High-Throughput Environments" on page 83.

What This Report Tells You

The Data Store Cost report answers the following questions:

- Is the datastore load healthy?
- Is there possible disk pressure?
- What is the relative load at different times of the day?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

To view the Data Store Cost report

1. Choose Reports > Optimization > Data Store Cost to display the Data Store Cost page.

Figure 10-24. Data Store Cost Page

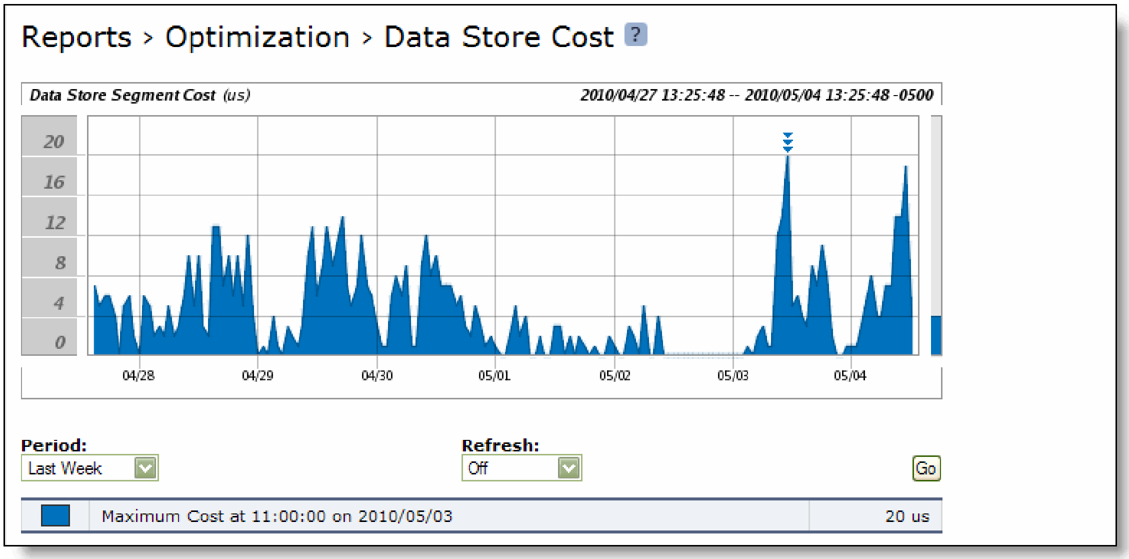

Reports > Optimization > Data Store Cost

Period: Last Week

Refresh: Off

🟦	Maximum Cost at 11:00:00 on 2010/05/03	20 us

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store Disk Load Reports

The Data Store Disk Load report summarizes the datastore disk load due to SDR only as related to the benchmarked capacity of the datastore. Consider any value under 90% as healthy. Any value higher than a sustained load over 90% is considered high and might indicate disk pressure. When a value is consistently higher than 90%, contact Riverbed Support for guidance on reconfiguring the datastore to alleviate disk pressure. The report contains the following table of statistics that summarizes the datastore disk load.

Field	Description
Maximum Disk Load	Specifies the maximum percentage of the datastore disk load.
Average Disk Load	Specifies the average percentage of the datastore disk load.
Minimum Disk Load	Specifies the minimum percentage of the datastore disk load.

For details on datastore disk performance, see “Viewing Data Store Cost Reports” on page 439.

What This Report Tells You

The Data Store Disk Load report answers the following questions:

- Is there any indication of disk pressure?
- What is the disk load at different times of the day?

To view the Data Store Disk Load report

1. Choose Reports > Optimization > Data Store Disk Load to display the Data Store Disk Load page.

Figure 10-25. Data Store Disk Load Page

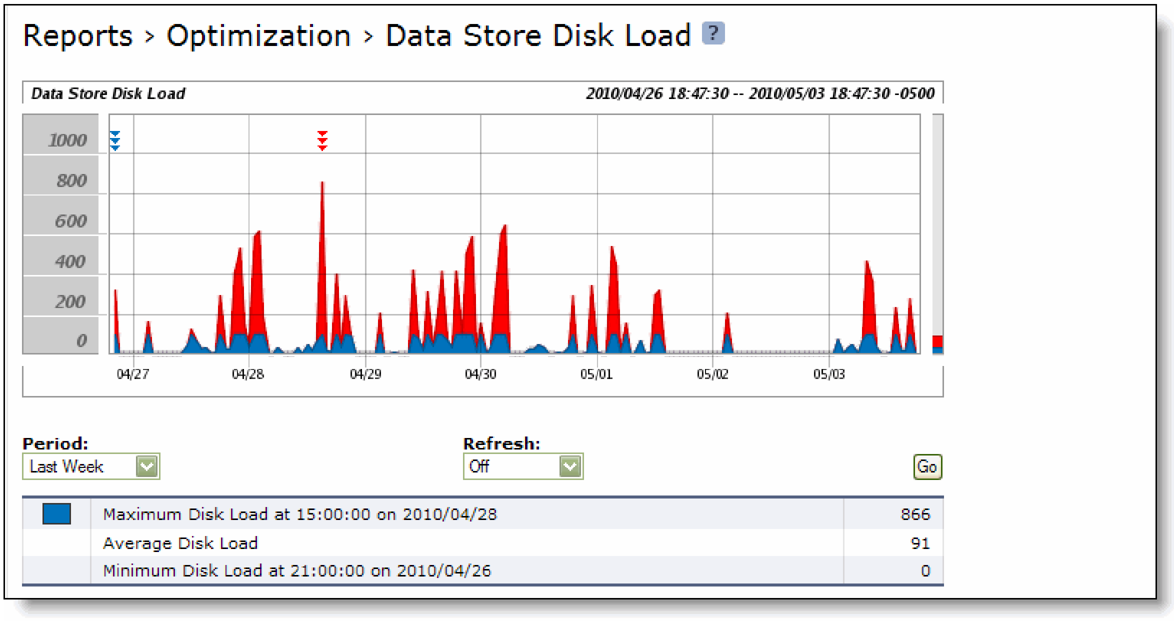

Reports > Optimization > Data Store Disk Load ?

Period: Last Week

Refresh: Off

Go

🟦	Maximum Disk Load at 15:00:00 on 2010/04/28	866
	Average Disk Load	91
	Minimum Disk Load at 21:00:00 on 2010/04/26	0

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store Read Efficiency Reports

The Data Store Read Efficiency report summarizes how efficiently the datastore disk is performing for the time period specified.

The Data Store Read Efficiency report includes a graph which displays a percentage breakdown of how much of each segment page has data in it for the time period you specify. This graph indicates how efficiently the datastore is using a page after a disk read operation.

Field	Description
Maximum Read Efficiency	Specifies the maximum disk segment page utilization range as a percent of bytes used after reading a page.
Minimum Read Efficiency	Specifies the minimum disk segment page utilization range as a percent of bytes used after reading a page.
Average Read Efficiency	Specifies the average disk segment page utilization range as a percent of bytes used after reading a page.

What This Report Tells You

The Data Store Read Efficiency report answers the following question:

- What percent of the disk data that is read from the datastore is actually used for active connections?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

To view the Data Store Read Efficiency report

1. Choose Reports > Optimization > Data Store Read Efficiency to display the Data Store Read Efficiency page.

Figure 10-26. Data Store Read Efficiency Page

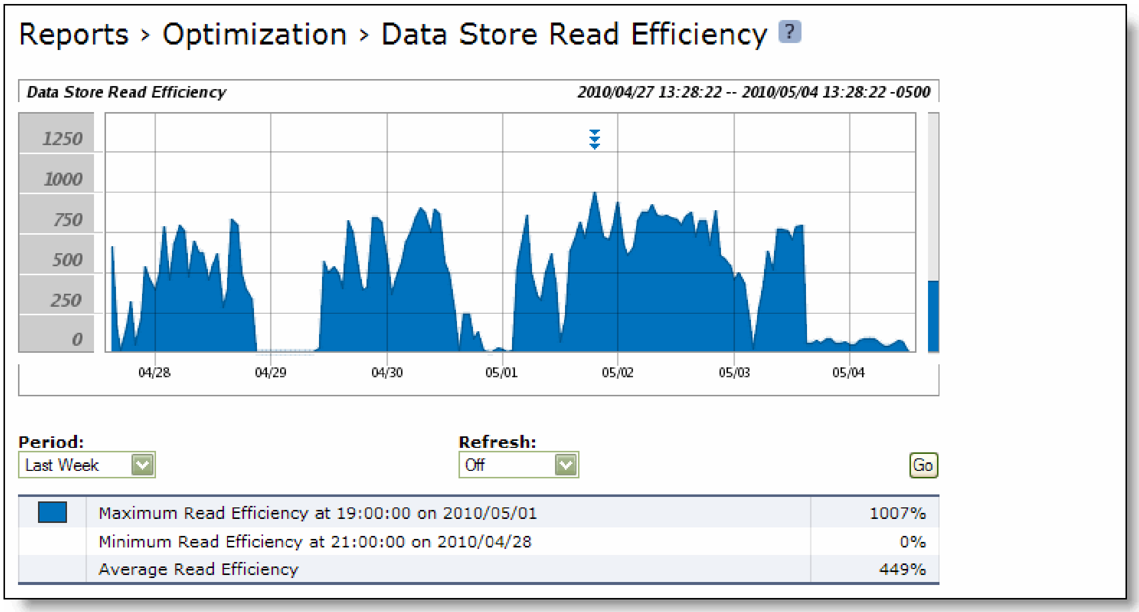

Reports > Optimization > Data Store Read Efficiency ?

Period: Last Week

Refresh: Off

🟦	Maximum Read Efficiency at 19:00:00 on 2010/05/01	1007%
	Minimum Read Efficiency at 21:00:00 on 2010/04/28	0%
	Average Read Efficiency	449%

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store Hit Rate Reports

The Data Store Hit Rate report summarizes how many times the data-store disk and memory have seen a data segment. A hit is a data segment that has been seen before by the datastore in the system. When a hit occurs, the system sends the reference to the data segment rather than the actual data over the WAN.

The Data Store Hit Rate report contains the following table of statistics that summarize datastore activity.

Control	Description
Total Hits	Specifies the total number of hits against the datastore. A hit is a data segment that has been seen before by the datastore in the system. If a hit has occurred, the system sends the reference to the data rather than the actual data over the WAN.
Total Misses	Specifies the number of misses that occurred. A miss is an unmatched data segment—the datastore has not seen the data segment before and must send all the data across the WAN. The data is LZ compressed, if LZ compression is enabled. For details on setting optimization policies, see "Configuring the Datastore" on page 76.
Peak Hits At <time> on <date>	Specifies the number of hits for the date and time.
Peak Misses At <time> on <date>	Specifies the number of misses for the date and time.

What This Report Tells You

The Data Store Hit Rate report answers the following questions:

- How much optimization is occurring?
- How much optimization occurred through SDR hits?
- How much data traversed the WAN without optimization?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the Data Store Hit Rate report

1. Choose Reports > Optimization > Data Store Hit Rate to display the Data Store Hit Rate page.

Figure 10-27. Data Store Hit Rate Page

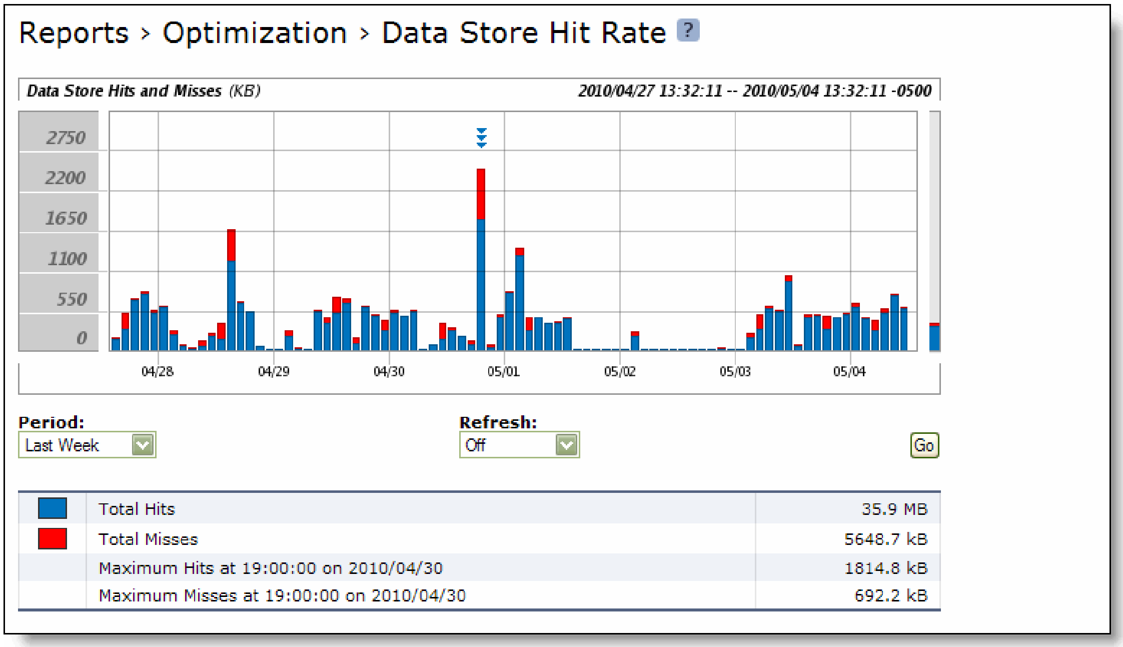

Reports > Optimization > Data Store Hit Rate ?

Data Store Hits and Misses (KB) 2010/04/27 13:32:11 -- 2010/05/04 13:32:11 -0500

Period: Last Week

Refresh: Off

Go

🟦	Total Hits	35.9 MB
🟥	Total Misses	5648.7 kB
	Maximum Hits at 19:00:00 on 2010/04/30	1814.8 kB
	Maximum Misses at 19:00:00 on 2010/04/30	692.2 kB

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Data Store IO Reports

The Data Store IO report summarizes how the datastore disk IO is performing for the time period specified. It measures how many random reads and writes are occurring, where a low value indicates the most random IO and larger values indicate more sequential IO. This report displays the following graphs:

- Data Store Cluster Average Reads - Plots the read cluster sizes for the time period you specify.
- Data Store Cluster Average Writes - Plots the write cluster sizes for the time period you specify.

The Data Store IO report cluster graphs include the following statistics:

Field	Description
Average Cluster Reads	Specifies the average cluster read size.
Average Cluster Writes	Specifies the average cluster write size.
Peak Cluster Reads	Specifies the peak cluster read size.
Peak Cluster Writes	Specifies the peak cluster write size.

The Data Store IO report also displays the following page graphs:

- Data Store Page Reads - Plots the page reads for the time period you specify.
- Data Store Page Writes - Plots the page writes for the time period you specify.

The Data Store IO report page graphs include the following statistics:

Field	Description
Total Page Reads	Specifies the total page read counts.
Total Page Writes	Specifies the total page write counts.

What This Report Tells You

The Data Store IO report answers the following questions:

- Is there any indication of disk pressure?
- What was the average cluster read and write size for the time period?
- What was the peak cluster read and write sizes for the time period?
- What was the average page read and write count for the time period?
- What was the peak page read and write count for the time period?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

To view the Data Store IO report

1. Choose Reports > Optimization > Data Store IO to display the Data Store IO page.

Figure 10-28. Data Store IO Page

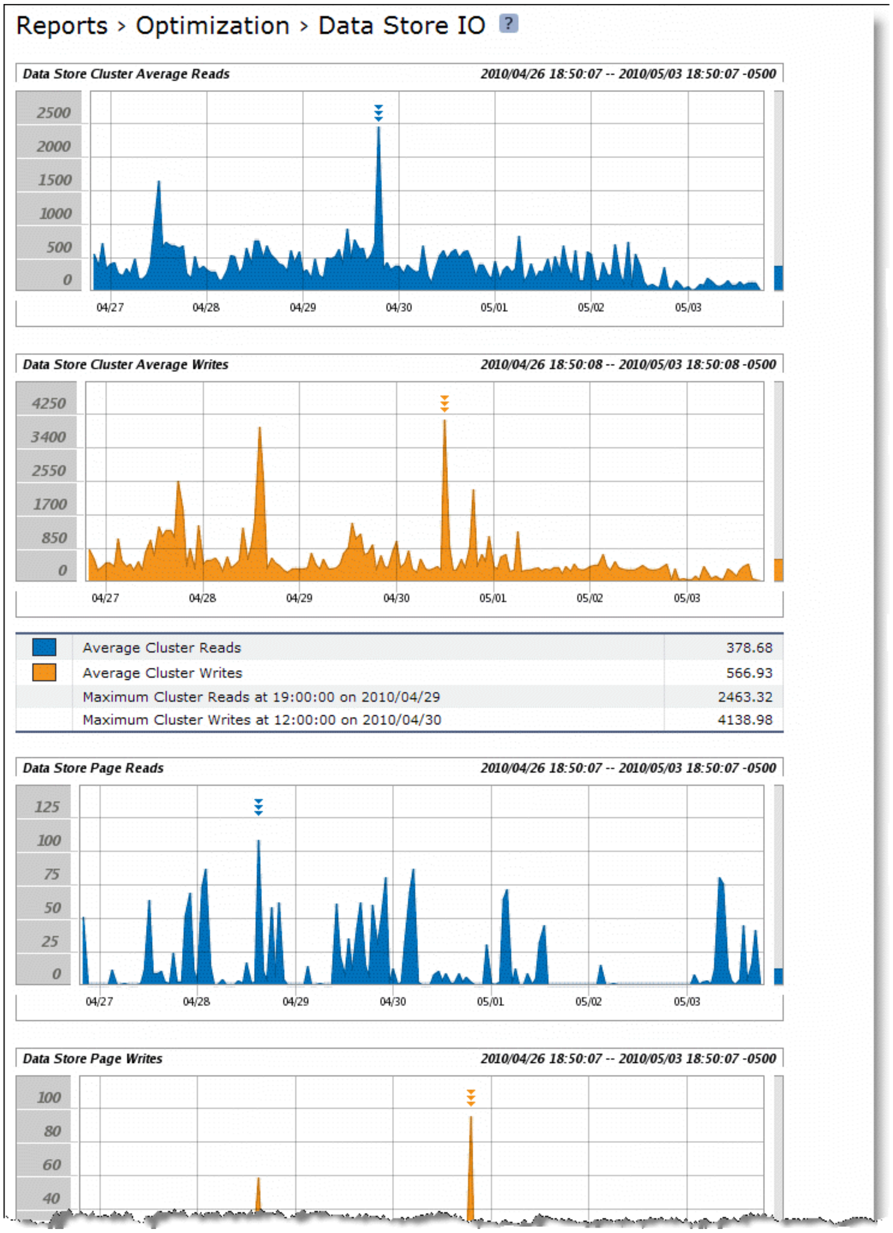

Reports > Optimization > Data Store IO

Data Store Cluster Average Reads 2010/04/26 18:50:07 -- 2010/05/03 18:50:07 -0500

Data Store Cluster Average Writes 2010/04/26 18:50:08 -- 2010/05/03 18:50:08 -0500

🟦	Average Cluster Reads	378.68
🟧	Average Cluster Writes	566.93
	Maximum Cluster Reads at 19:00:00 on 2010/04/29	2463.32
	Maximum Cluster Writes at 12:00:00 on 2010/04/30	4138.98

Data Store Page Reads 2010/04/26 18:50:07 -- 2010/05/03 18:50:07 -0500

Data Store Page Writes 2010/04/26 18:50:07 -- 2010/05/03 18:50:07 -0500

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing PFS Share Reports

The PFS Share Status report provides information about your PFS shares: the local name of the share and the status of the share. It displays the following table of statistics that summarize PFS share activity.

Packet Type	Description
Local Name	Specifies the local share name.
Sharing	Specifies whether sharing is enabled for the share.
Syncing	Specifies whether synchronization is enabled for the share.
Status	Specifies the current share status.
Last Sync Time	Specifies when the last attempt to synchronize the share occurred.
Last Sync Status	Specifies whether the last share synchronization was successful.

For details, see "Adding PFS Shares" on page 166.

What This Report Tells You

The PFS Share Status report answers the following questions:

- What action is occurring on the share?
- Is the share ready for synchronization?
- Is a synchronization currently occurring?
- When was the last time a synchronization occurred?
- What was the last synchronization status?

To view the PFS Shares report

- Choose Reports > Branch Services > PFS Shares to display the PFS Shares report.

Figure 10-29. PFS Shares Report

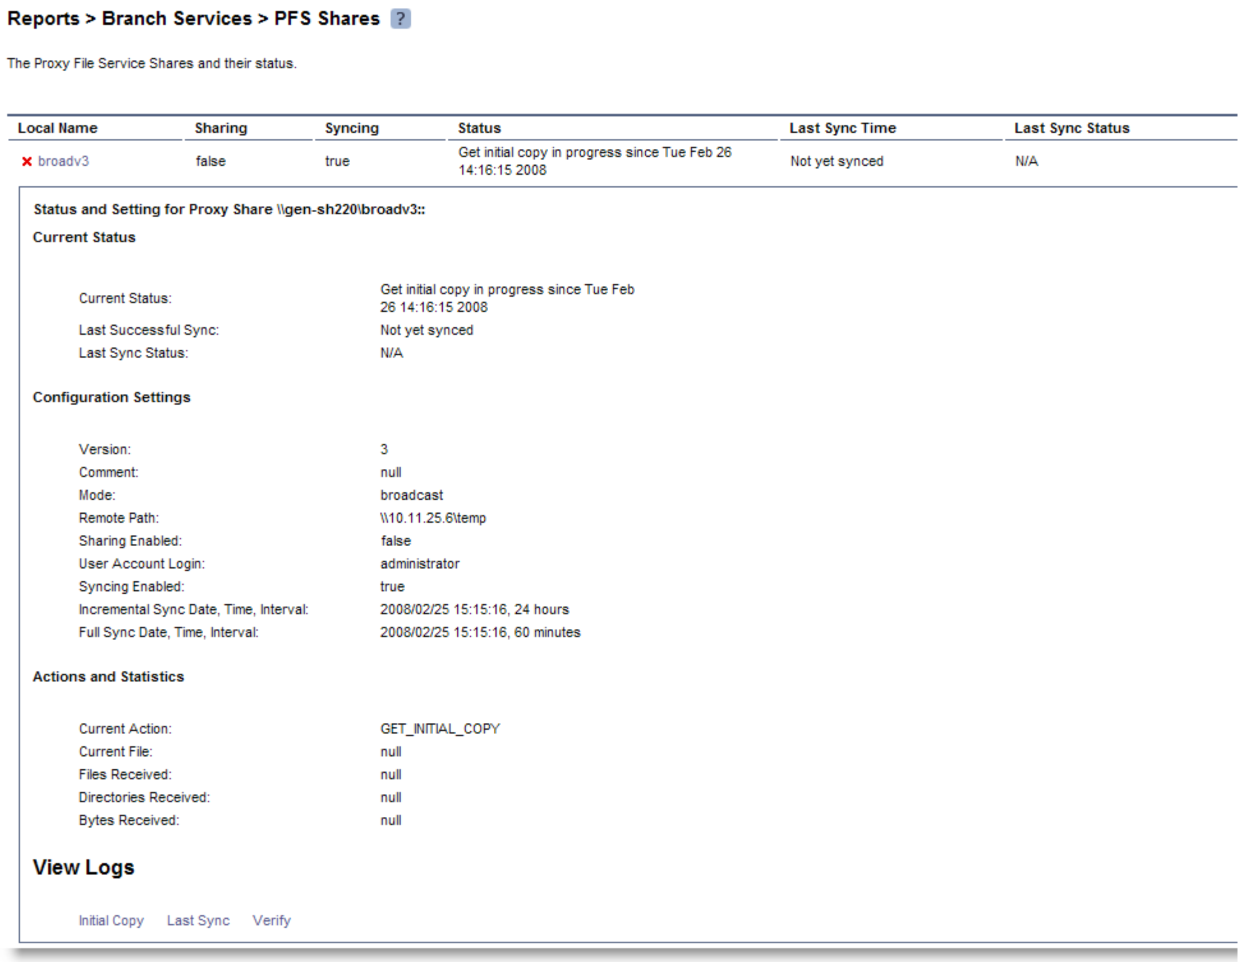

Reports > Branch Services > PFS Shares ?

The Proxy File Service Shares and their status.

Local Name	Sharing	Syncing	Status	Last Sync Time	Last Sync Status
✖ broadv3	false	true	Get initial copy in progress since Tue Feb 26 14:16:15 2008	Not yet synced	N/A

Status and Setting for Proxy Share \\gen-sh220\broadv3::

Current Status

Current Status:	Get initial copy in progress since Tue Feb 26 14:16:15 2008
Last Successful Sync:	Not yet synced
Last Sync Status:	N/A

Configuration Settings

Version:	3
Comment:	null
Mode:	broadcast
Remote Path:	\\10.11.25.6\temp
Sharing Enabled:	false
User Account Login:	administrator
Syncing Enabled:	true
Incremental Sync Date, Time, Interval:	2008/02/25 15:15:16, 24 hours
Full Sync Date, Time, Interval:	2008/02/25 15:15:16, 60 minutes

Actions and Statistics

Current Action:	GET_INITIAL_COPY
Current File:	null
Files Received:	null
Directories Received:	null
Bytes Received:	null

View Logs

Initial Copy Last Sync Verify

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing PFS Share Logs

You can view PFS share logs from the PFS Share Logs page.

The PFS share logs contain more detail regarding the initial copy of the share and the last share synchronization. You can use the share log to verify the share.

To view the PFS share logs

1. Choose Reports > Branch Services > PFS Shares to display the PFS Shares page.

Figure 10-30. PFS Shares Report

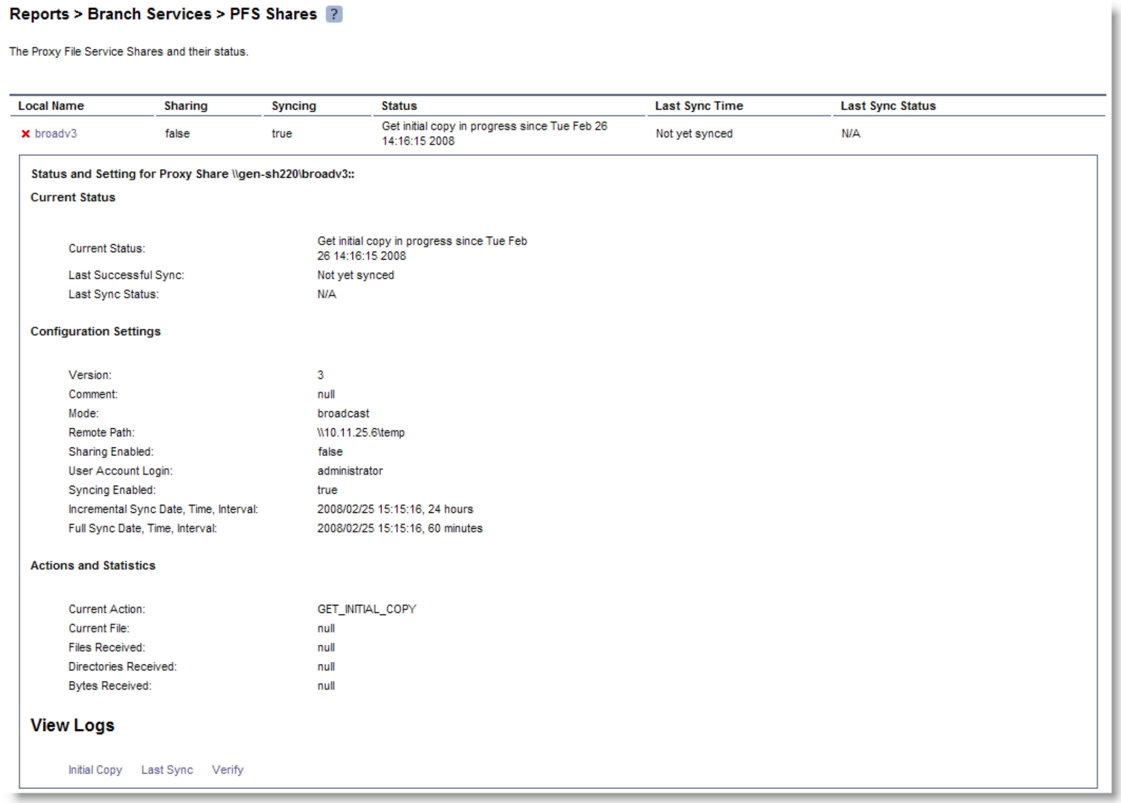

Reports > Branch Services > PFS Shares

The Proxy File Service Shares and their status.

Local Name	Sharing	Syncing	Status	Last Sync Time	Last Sync Status
✕ broadv3	false	true	Get initial copy in progress since Tue Feb 26 14:16:15 2008	Not yet synced	N/A

Status and Setting for Proxy Share \\gen-sh220\broadv3::

Current Status

Current Status:	Get initial copy in progress since Tue Feb 26 14:16:15 2008
Last Successful Sync:	Not yet synced
Last Sync Status:	N/A

Configuration Settings

Version:	3
Comment:	null
Mode:	broadcast
Remote Path:	\\10.11.25.6\temp
Sharing Enabled:	false
User Account Login:	administrator
Syncing Enabled:	true
Incremental Sync Date, Time, Interval:	2008/02/25 15:15:16, 24 hours
Full Sync Date, Time, Interval:	2008/02/25 15:15:16, 60 minutes

Actions and Statistics

Current Action:	GET_INITIAL_COPY
Current File:	null
Files Received:	null
Directories Received:	null
Bytes Received:	null

2. Select a share name.

3. Use the controls to customize the report as described in the following table.

Control	Description
Initial Copy	Displays details on when the first share synchronization occurred.
Last Sync	Displays details on when the last share synchronization occurred.
Verify	Displays a list of differences between the share on the Steelhead appliance and the origin-file server.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing PFS Data Reports

The PFS Data report provides information about how much data was sent and received by PFS. For details, see “Adding PFS Shares” on page 166.

The PFS Data report contains the following table of statistics that summarize PFS activity.

Packet Type	Description
Total Bytes Sent	Specifies the total number of bytes sent over the WAN.
Peak Bytes Sent At <time> on <date>	Specifies the peak number of bytes sent for the date and time.
Total Bytes Received	Specifies the total number of bytes received over the WAN.
Peak Bytes Received At <time> on <date>	Specifies the peak number of bytes received for the date and time.

What This Report Tells You

The PFS Data report answers the following questions:

- How many bytes were sent over the WAN?
- How many bytes were received over the WAN?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the PFS Data report

1. Choose Reports > Branch Services > PFS Data to display the PFS Data page.

Figure 10-31. PFS Data Page

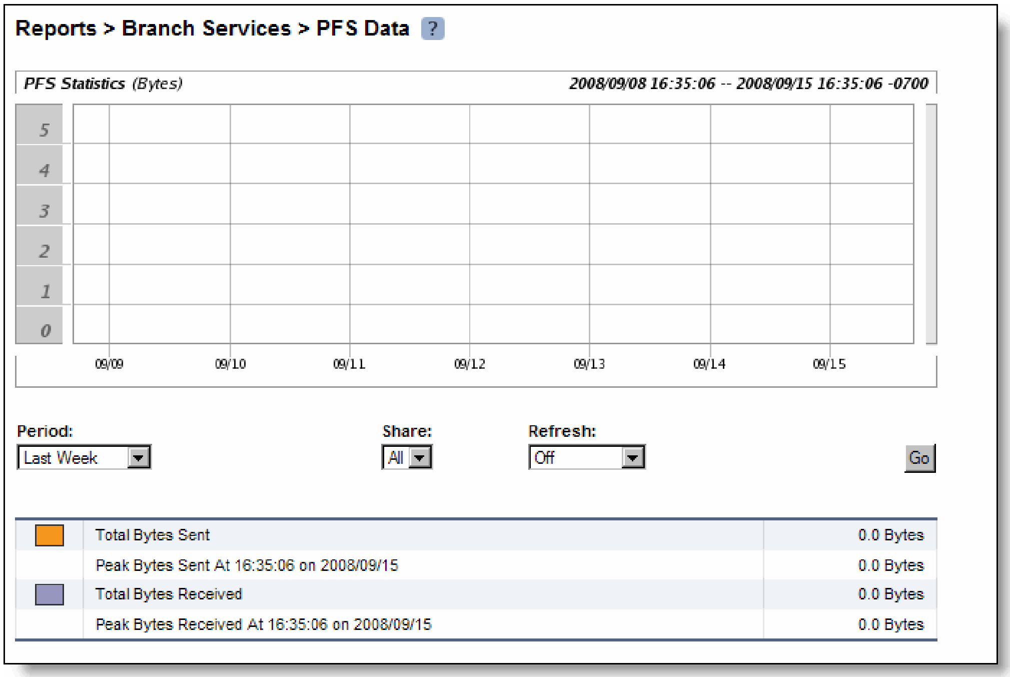

Reports > Branch Services > PFS Data ?

PFS Statistics (Bytes) *2008/09/08 16:35:06 -- 2008/09/15 16:35:06 -0700*

Period: Last Week

Share: All

Refresh: Off

Go

🟧	Total Bytes Sent	0.0 Bytes
	Peak Bytes Sent At 16:35:06 on 2008/09/15	0.0 Bytes
🟪	Total Bytes Received	0.0 Bytes
	Peak Bytes Received At 16:35:06 on 2008/09/15	0.0 Bytes

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Share	Specify a specific share or select All from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing DNS Cache Hits

The DNS Cache Hits report provides a DNS cache hits graph for the time period specified. It contains the following table of statistics that summarize DNS activity.

Control	Description
Total Requests	Specifies the total number of DNS requests.
Total Hits	Specifies the total number of cache hits.
Hit %	Specifies the percentage of cache hits.

What This Report Tells You

The DNS Cache Hits report answers the following questions:

- How many DNS requests occurred?
- How many DNS entries were retrieved from the cache?
- What percentage of DNS requests were cached?
- What is the average number of cached entries?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the DNS Cache Hits report

1. Choose Reports > Branch Services > DNS Cache Hits to display the DNS Cache Hits page.

Figure 10-32. DNS Cache Hits Page

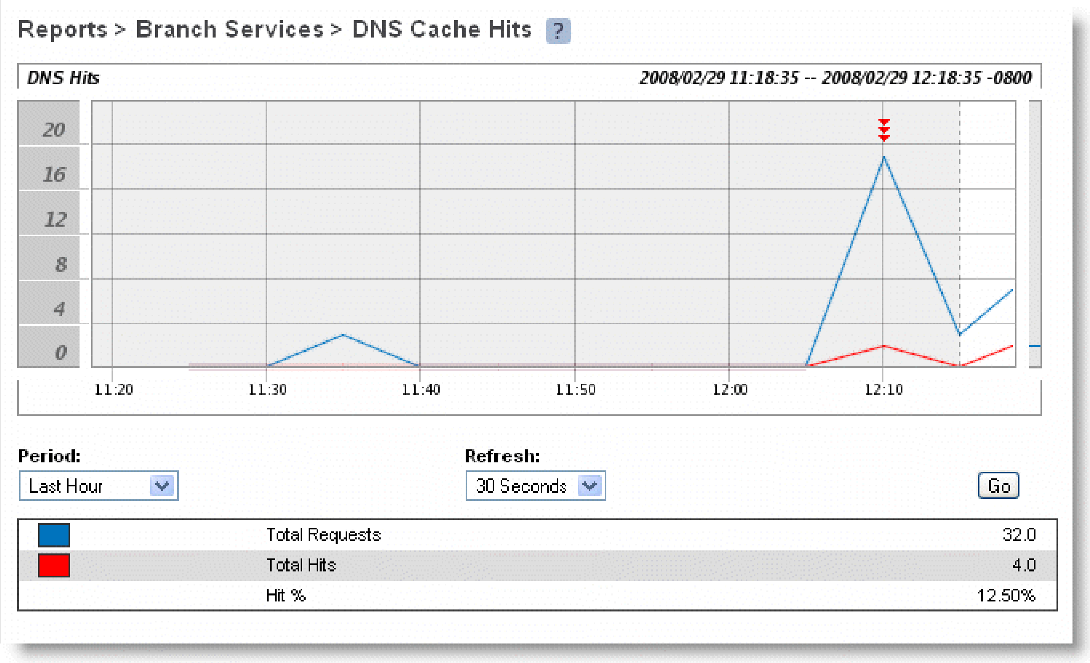

Reports > Branch Services > DNS Cache Hits

DNS Hits 2008/02/29 11:18:35 -- 2008/02/29 12:18:35 -0800

Period: Last Hour

Refresh: 30 Seconds

🟦	Total Requests	32.0
🟥	Total Hits	4.0
	Hit %	12.50%

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing DNS Cache Utilization

The DNS Cache Utilization report provides a DNS cache utilization graph for the time period specified. It contains the following table of statistics that summarize DNS cache activity.

Field	Description
Average Cache Memory Utilization	Specifies the average cache memory used.
Average Cache Entries	Specifies the average number of entries in the cache.

What This Report Tells You

The DNS Cache Utilization report answers the following questions:

- How much cache memory is used?
- What is the average cache memory used?
- How many DNS entries are in the cache?
- What is the average number of DNS entries in the cache?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the DNS Cache Utilization report

1. Choose Reports > Branch Services > DNS Cache Utilization to display the DNS Cache Utilization page.

Figure 10-33. DNS Cache Utilization Page

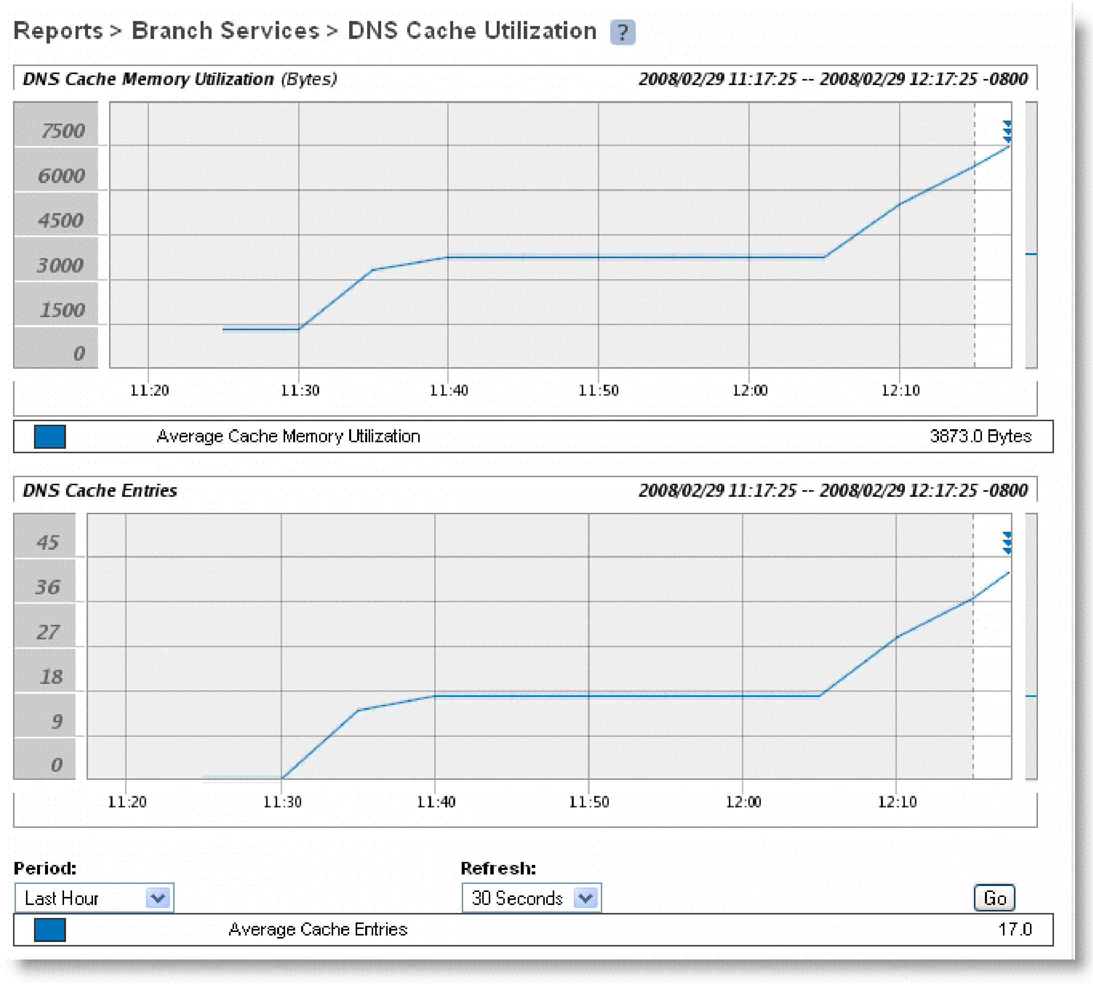

Reports > Branch Services > DNS Cache Utilization ?

DNS Cache Memory Utilization (Bytes) 2008/02/29 11:17:25 -- 2008/02/29 12:17:25 -0800

🟦 Average Cache Memory Utilization 3873.0 Bytes

DNS Cache Entries 2008/02/29 11:17:25 -- 2008/02/29 12:17:25 -0800

Period: Last Hour

Refresh: 30 Seconds

🟦 Average Cache Entries 17.0

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

3. To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing RSP Statistics Reports

The RSP VNI IO report provides a utilization graph for a package and a VNI for the time period specified. It contains the following table of statistics that summarize RSP VNI activity.

Field	Description
LAN In	Specifies the byte count or number of packets coming into the LAN.
LAN Out	Specifies the byte count or number of packets going out of the LAN.
WAN In	Specifies the byte count or number of packets coming into the WAN.
WAN Out	Specifies the byte count or number of packets going out of the WAN.
Package In	Specifies the byte count or number of packets coming into the package.
Package Out	Specifies the byte count or number of packets going out of the package.

What This Report Tells You

The RSP VNI IO report answers the following questions:

- How much traffic is coming in and going out the LAN direction?
- How much traffic is coming in and going out the WAN direction?
- How much traffic is coming in and going out of a package?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

About Report Data

The Riverbed system polls bandwidth and connection metrics every second and reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the Last 5 Minutes are interpolated from aggregate data points.

To view the RSP VNI IO report

1. Choose Reports > Branch Services > RSP VNI IO to display the RSP VNI IO page.

Figure 10-34. RSP VNI IO Page

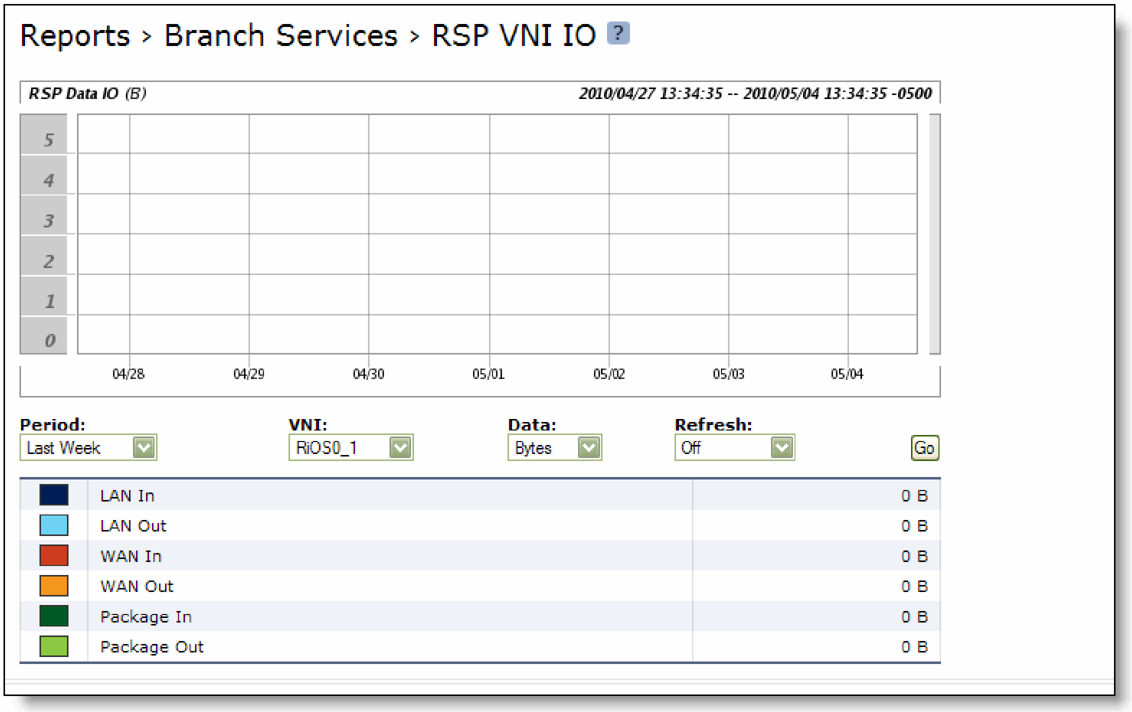

Reports > Branch Services > RSP VNI IO

Period: Last Week

VNI: RiOS0_1

Data: Bytes

Refresh: Off

Go

🟦	LAN In	0 B
🟦	LAN Out	0 B
🟥	WAN In	0 B
🟧	WAN Out	0 B
🟩	Package In	0 B
🟩	Package Out	0 B

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of Last Minute, Last 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Slot	Select a slot from the drop-down list.
VNI	Select a VNI from the drop-down list. You must first select a slot.
Data	Select either Bytes or Packets.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Alarm Status Reports

The Alarm Status report provides status for the Steelhead appliance alarms. It includes the following alarm information.

Alarm	Reason
Admission Control	Indicates the system connection limit has been reached. Additional connections are passed through unoptimized. The alarm clears when the Steelhead appliance moves out of this condition.
Asymmetric Routing	Indicates the system is experiencing asymmetric traffic. Indicates OK if the system is not experiencing asymmetric traffic. In addition, any asymmetric traffic is passed through, and the route appears in the Asymmetric Routing table. For details on the Asymmetric Routing table, see "Configuring Asymmetric Routing Features" on page 251.
Certificate Revocation List Expirations	Indicates a failure with the Certificate Revocation List (CRL) verification on the server certificates. A CRL includes any digital certificates that have been invalidated before their expiration date, including the reasons for their revocation and the names of the issuing certificate signing authorities. A CRL prevents the use of digital certificates and signatures that have been compromised.
Connection Forwarding	Indicates the system has detected a problem with a connection forwarding neighbor. The connection forwarding alarms are inclusive of all connection forwarding neighbors. For example, if a Steelhead appliance has three neighbors, the alarm triggers if any *one* of the neighbors are in error. In the same way, the alarm clears only when all three neighbors are no longer in error. The following issues trigger the connection forwarding alarm: • The connection forwarding neighbor has not sent a keep-alive message within the time-out period to the neighbor Steelhead appliance(s), indicating that the connection has been lost. • The connection cannot be established with a connection forwarding neighbor. • The connection has been closed by the connection forwarding neighbor. • The connection has been lost with the connection forwarding neighbor due to an error. • The connection has been lost because requests have not been acknowledged by a connection forwarding neighbor within the set threshold. • The Steelhead appliance has timed out while waiting for an initialization message from a connection forwarding neighbor. • The amount of latency between connection forwarding neighbors has exceeded the specified threshold.
CPU Utilization	Indicates the system has reached the CPU threshold for any of the CPUs in the Steelhead appliance. If the system has reached the CPU threshold, check your settings. For details, see "Configuring Alarm Settings" on page 323. If your alarm thresholds are correct, reboot the Steelhead appliance. For details, see "Rebooting and Shutting Down the Steelhead Appliance" on page 349. **Note**: If more than 100 MBs of data is moved through a Steelhead appliance while performing PFS synchronization, the CPU utilization might become high and result in a CPU alarm. This CPU alarm is not cause for concern.

Alarm	Reason
Data Store Error	Indicates the RiOS datastore is corrupt or has become incompatible with the current configuration. To clear the RiOS datastore of data, choose Configure > Maintenance > Services, select **Clear Data Store** and click **Restart** to restart the optimization service. For details, see “Starting and Stopping the Optimization Service” on page 345. If the alarm was caused by an unintended change to the configuration, you can change the configuration to match the previous RiOS datastore settings. Then restart the optimization service without clearing the datastore to reset the alarm. Typical configuration changes that require a restart with a clear RiOS datastore are enabling the extended peer table or changing the datastore encryption. For details, see “Enabling Peering and Configuring Peering Rules” on page 67 and “Configuring the Datastore” on page 76.

Datastore Not Zeroed	(Appears only on Virtual Steelheads) Indicates that the datastore disk has not been set up for optimal performance. A Virtual Steelhead will achieve the best performance if the datastore is zeroed, meaning that the storage blocks allocated to the datastore are sequentially located. This notification indicates that allocated blocks may be fragmented, which may lead to some degradation in performance.
	If possible, you should zero the datastore for best performance after first installation and after each model upgrade that requires a change in the size of the datastore. The performance penalty for not doing so should be small in most cases, but in some circumstances could be larger. The time needed to complete zeroing the datastore depends on the size and throughput of the underlying storage device, and can range from 30 minutes to several hours.
	To zero the datastore disk on a Virtual Steelhead, enter the following CLI commands at the system prompt:
	`enable`
	`config t`
	`datastore zero`
	The amount of time it takes to zero the datastore disk depends on its size as well as the throughput of the underlying storage device. For example, on the V250 series (with a 44 GB datastore disk), it takes approximately 45 minutes. On the V550 series (with an 80 GB disk), it takes approximately 1 1/2 hours. The optimization service restarts automatically when the datastore is set up.
	To cancel the operation, press Ctrl+C within the first five minutes. Do not start the optimization service from another console while zeroing the datastore.
	The alarm clears after the disk is zeroed or the datastore wraps at least once.

Data Store Synchronization	Indicates the system has detected a problem with the synchronized datastores. For details, see “Synchronizing Peer Datastores” on page 78.
Disk Condition	Indicates the system has detected a problem with a disk or a Solid State Drive. • If the Disk Error alarm is triggered, the message “Disk(s) *x*, *y*, and *z* are reporting errors” appears. The disk is either failing, invalid, or missing. • If the SSD Wear alarm is triggered, the message “Disk *x* is approaching its write cycle limit” appears. (Appears only on Steelhead appliance models 7050L or 7050M).
Domain Join Error	Indicates an attempt to join a Windows domain has failed. For details, see “Troubleshooting a Domain Join Failure” on page 309.
Fan Error	Indicates the system has detected a problem with the fans. Fans for many systems can be replaced. Contact Riverbed Support at https://support.riverbed.com and file a trouble ticket to order a replacement fan. For details on replacing fans, see the *Upgrade and Maintenance Guide*.

Alarm	Reason
Flash Error	Indicates the system has detected an error with the flash drive hardware. At times, the USB flash drive that holds the system images may become unresponsive. When this happens, the system is unable to write a new upgrade image to the flash drive without first power cycling the system. Reboot using either Configure > Maintenance > Reboot/Shutdown or the CLI **reload** command to automatically power cycle the Steelhead appliance and restore the flash drive to proper function.
Hardware Error	Indicates the system has detected a problem with the Steelhead appliance hardware. The following issues trigger the hardware error alarm: • the Steelhead appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration • the Steelhead appliance is using a memory Dual In-line Memory Module (DIMM), a hard disk, or a NIC that is not qualified by Riverbed • an RSP upgrade requires additional memory or a memory replacement • other hardware issues The alarm clears when you add the necessary hardware, remove the non-qualified hardware, or resolve other hardware issues.
IPMI	Indicates there has been a physical security intrusion triggering an Intelligent Platform Management Interface (IPMI) error. The following events trigger the IPMI alarm: • chassis intrusion (physical opening and closing of the appliance case) • memory errors (correctable or uncorrectable ECC memory errors) • hard drive faults or predictive failures • power supply status or predictive failure The option to reset the alarm appears only after the service triggers the IPMI alarm. To reset the alarm, click **Clear the IPMI alarm now**.
Licensing	Indicates whether your licenses are current. For details on updating licenses, see "Managing Licenses and Model Upgrades" on page 349.
Link State	Indicates the system has detected a link that is down. You are notified through SNMP traps, email, and alarm status.
Memory Paging	Indicates the system has reached the memory paging threshold. If 100 pages are swapped approximately every two hours the Steelhead appliance is functioning properly. If thousands of pages are swapped every few minutes, then reboot the Steelhead appliance. For details, see "Rebooting and Shutting Down the Steelhead Appliance" on page 349. If rebooting does not solve the problem, contact Riverbed Support at https://support.riverbed.com.
Neighbor Incompatibility	Indicates the system has encountered an error in reaching a Steelhead appliance configured for connection forwarding. For details, see "Configuring Connection Forwarding Features" on page 255.
Network Bypass	Indicates the system is in bypass failover mode. If the Steelhead appliance is in bypass failover mode, restart the Steelhead service. If restarting the service does not resolve the problem, reboot the Steelhead appliance. If rebooting does not resolve the problem, shutdown and restart the Steelhead appliance. For details, see "Rebooting and Shutting Down the Steelhead Appliance" on page 349, and "Starting and Stopping the Optimization Service" on page 345.
NFS V2/V4 Alarm	Indicates the system has triggered a v2 or v4 NFS alarm. For details, see "Configuring NFS Optimization" on page 131.

Alarm	Reason
Non-443 SSL Servers Detected	Indicates that during a RiOS upgrade (for example, from v5.5 to v6.0), the system has detected a pre-existing SSL server certificate configuration on a port other than the default SSL port 443. SSL traffic might not be optimized. To restore SSL optimization, you can add an in-path rule to the client-side Steelhead appliance to intercept the connection and optimize the SSL traffic on the non-default SSL server port. After adding an in-path rule, you must clear this alarm manually by entering the following CLI command: `stats alarm non_443_ssl_servers_detected_on_upgrade clear`
Optimization Service	Indicates the service is not running or there is a software error in the Steelhead service. The Steelhead service continues to function, but an error message appears in the logs that you should investigate. For details, see "Viewing Logs" on page 475.
Power Supply	Indicates an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted.
Prepopulation or Proxy File Service Configuration	Indicates there has been a PFS or prepopulation operation error. If an operation error is detected, restart the Steelhead service and PFS. For details, see "Configuring PFS" on page 161, and "Starting and Stopping the Optimization Service" on page 345.
Prepopulation or Proxy File Service Operation	Indicates a synchronization operation has failed. If an operation failure is detected, attempt the operation again. For details, see "Adding PFS Shares" on page 166.
Process Dump Staging Directory Inaccessible	Indicates that the system has detected an error while trying to create a process dump. Please contact Riverbed Support to correct the issue.
RAID	Indicates the system has encountered RAID errors (for example, missing drives, pulled drives, drive failures, and drive rebuilds). For drive rebuilds, if a drive is removed and then reinserted, the alarm continues to be triggered until the rebuild is complete. **Important**: Rebuilding a disk drive can take 4-6 hours. **Note**: RAID status applies only to the Steelhead appliance Series 3000, 5000, and 6000.
Riverbed Services Platform	Indicates the system has detected a problem with RSP. The following issues trigger the RSP alarm: • The available memory for RSP is negative. • The installed RSP image is incompatible with the current release. • A watchdog activates for any slot that has a watchdog configured. This can indicate that an RSP package or a VM has failed and is blocking or bypassing traffic. For details, see "Specifying Watchdog Settings" on page 191. • Virtual machines are enabled but are not currently powered on.
RSP License	Indicates an RSP license will expire within seven days or has already expired.
Secure Vault Alarm	Indicates the secure vault is locked or an error has occurred while initializing the secure vault. When the vault is locked, SSL traffic is not optimized and you cannot encrypt the datastore. For details, see "Unlocking the Secure Vault" on page 370.
SMB Signing Alert	Indicates the system has detected an SMB signing error. For details, see "Configuring CIFS Optimization" on page 97.

Alarm	Reason

Software Version Mismatch	Indicates there is a mismatch between software versions in your network. If a software mismatch is detected, resolve the mismatch by upgrading or reverting to a previous version of the software. For details, see "Upgrading Your Software" on page 347.
	Note: If a software version mismatch occurs and you are running v.1.2 and client-side v.2.1 Steelhead appliances, you must set the correct version of the Steelhead service protocol on the client-side v.2.1 appliances using the Steelhead CLI:
	`sh> peer <addr> version min 5`
	`sh> peer <addr> version max 5`

SSL Alarms	Indicates an error has been detected in your SSL configuration. For details on checking your settings, see "Verifying SSL and Secure Inner Channel Optimization" on page 216.
SSL Peering Certificate SCEP Automatic Re-enrollment	Indicates an SSL peering certificate has failed to re-enroll automatically within the Simple Certificate Enrollment Protocol (SCEP) polling interval.
System Details Report	Indicates that the system has detected a problem with an optimization or system module. For details, see "Viewing System Details Reports" on page 465.
System Disk Full	Indicates the system partitions (not the datastore) are almost full. For example, **/var** which is used to hold logs, statistics, system dumps, TCP dumps, and so on.
Temperature	Indicates the CPU temperature has exceeded the critical threshold. The default value for the rising threshold temperature is 80º C; the default reset threshold temperature is 67º C.

What This Report Tells You

The Alarm Status report answers the following question:

- What is the current status of the Steelhead appliance?

To view the Alarm Status report

- Choose Reports > Diagnostics > Alarm Status to display the Alarm Status page. Alternately, you can select the current system status that appears in the status box in the upper-right corner of each screen (**Healthy**, **Admission Control**, **Degraded**, or **Critical**) to display the Alarm Status page.

Figure 10-35. Alarm Status Page

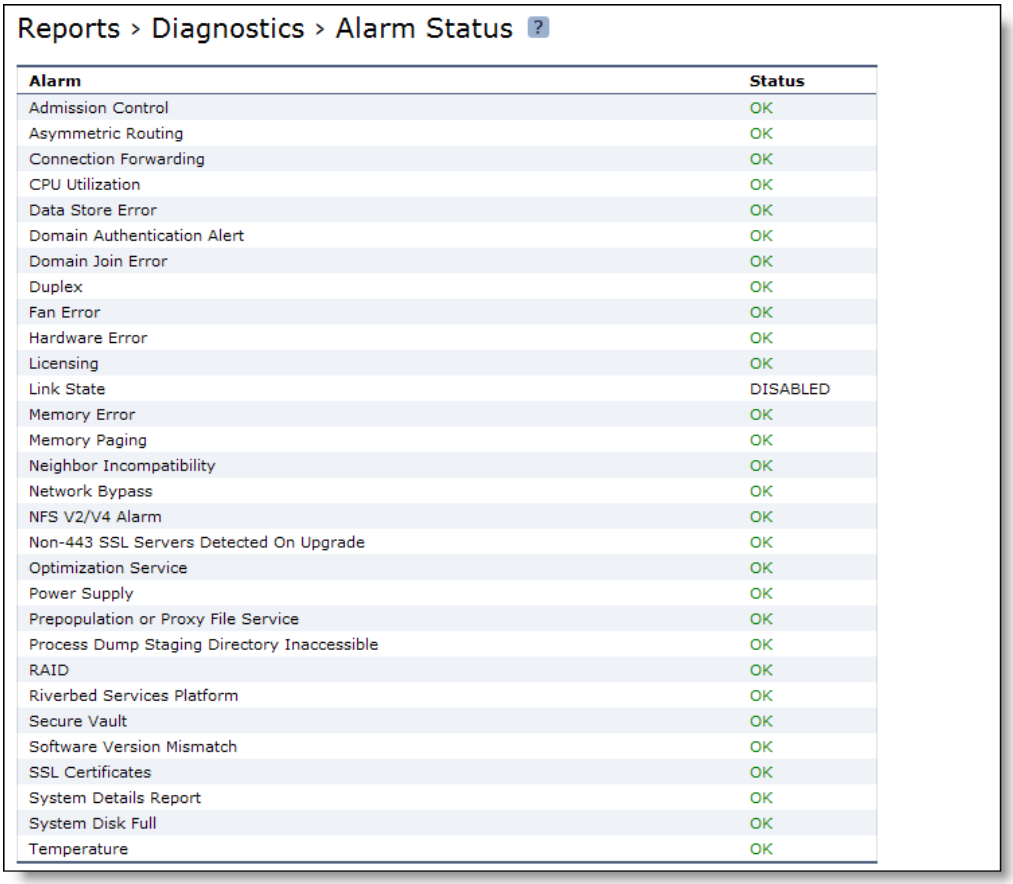

Reports > Diagnostics > Alarm Status ?

Alarm	Status
Admission Control	OK
Asymmetric Routing	OK
Connection Forwarding	OK
CPU Utilization	OK
Data Store Error	OK
Domain Authentication Alert	OK
Domain Join Error	OK
Duplex	OK
Fan Error	OK
Hardware Error	OK
Licensing	OK
Link State	DISABLED
Memory Error	OK
Memory Paging	OK
Neighbor Incompatibility	OK
Network Bypass	OK
NFS V2/V4 Alarm	OK
Non-443 SSL Servers Detected On Upgrade	OK
Optimization Service	OK
Power Supply	OK
Prepopulation or Proxy File Service	OK
Process Dump Staging Directory Inaccessible	OK
RAID	OK
Riverbed Services Platform	OK
Secure Vault	OK
Software Version Mismatch	OK
SSL Certificates	OK
System Details Report	OK
System Disk Full	OK
Temperature	OK

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing System Details Reports

The System Details report takes a current snapshot of the system to provide a one-stop report you can use to check for any issues with the Steelhead appliance. The report examines key system components; for example, the CPU and memory. Use this report to gather preliminary system information before calling Riverbed Support to troubleshoot an issue.

Field	Description
Module	Specifies the Steelhead appliance module. Select a module name to view details. A right arrow to the left of a module indicates that the report includes detailed information about a submodule. Click the arrow to view submodule details. This report examines the following modules: • CPU - Displays information on idle time, system time, and user time per CPU. • Memory - Displays information on total, used, and free memory by percentage and in KBs. • CIFS - Click the right arrow and the submodule name to view details for unexpected shutdowns and round trip statistics. • HTTP - Click the right arrow and the submodule name to view details for the URL Learning, Parse and Prefetch, and Object Prefetch Table optimization schemes. • Intercept - Click the right arrow to view statistics for message queue, GRE, and WCCP. Also includes table length and watchdog status. • Lotus Notes • MAPI - Click the right arrow and the submodule name to view details for: Accelerators - displays how many accelerator objects have been created for readahead, writebehind, and cached-mode folder synchronization. One accelerator object corresponds to the optimization of one particular Outlook action. • Readahead is for downloading an email attachment (in non-cached Outlook mode or for public folders). • Writebehind is for uploading an email attachment. • Cache-sync is for downloading the new contents of a folder (in cached mode). Requests and responses - displays the number of MAPI round trips used and saved. Includes the number of responses and faults along with the fault reason; for example, access denied. MAPI decryption and encryption (RPCCR) - displays whether MAPI decryption and encryption is enabled. Includes the number of client and server-side Steelhead appliance encrypted MAPI sessions, along with details on how many sessions were not encrypted, how many sessions were successfully decrypted and encrypted, how many sessions were passed-through, and how many experienced an authentication failure. Connection sessions - displays the number of client and server-side Steelhead appliance MAPI sessions, counting the number of MAPI 2000, 2003, 2007, and pass-through sessions. • MS-SQL - Displays whether MS-SQL optimization is enabled. • Oracle Forms - Click the right arrow and submodule name to view details for native and HTTP mode key • Secure Peering - Click the right arrow and submodule name to view details for secure inner channels, including information on certificate and private key validity, peer Steelhead appliance trust, and blacklisted servers. • SSL - Displays whether SSL optimization is enabled and details about the SSL configuration such as which advanced settings are in use. Click the right arrow and the submodule name to view details for the SSL outer and inner channels.

Field	Description
Status	Displays one of the following results: OK (Green) Warning (Yellow) Error (Red) Disabled (Gray). Appears when you manually disable the module.

What This Report Tells You

The System Details report answers the following questions:

- Is there a problem with one particular application module or does the issue affect more than one module?

To view the System Details report

- Choose Reports > Diagnostics > System Details to display the System Details page.

Figure 10-36. System Details Page

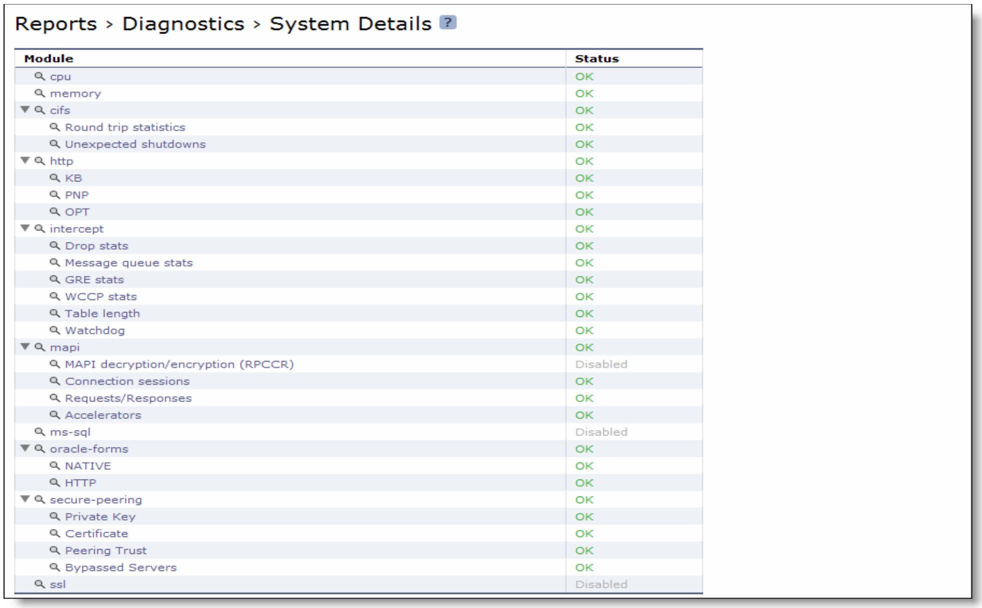

Reports > Diagnostics > System Details

Module	Status
cpu	OK
memory	OK
▼ cifs	OK
Round trip statistics	OK
Unexpected shutdowns	OK
▼ http	OK
KB	OK
PNP	OK
OPT	OK
▼ intercept	OK
Drop stats	OK
Message queue stats	OK
GRE stats	OK
WCCP stats	OK
Table length	OK
Watchdog	OK
▼ mapi	OK
MAPI decryption/encryption (RPCCR)	Disabled
Connection sessions	OK
Requests/Responses	OK
Accelerators	OK
ms-sql	Disabled
▼ oracle-forms	OK
NATIVE	OK
HTTP	OK
▼ secure-peering	OK
Private Key	OK
Certificate	OK
Peering Trust	OK
Bypassed Servers	OK
ssl	Disabled

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing CPU Utilization Reports

The CPU Utilization report summarizes the percentage of the CPU used within the time period specified.

General Usage Guidelines

Typically, a Steelhead appliance operates on approximately 30-40 percent CPU capacity during non-peak hours and approximately 60-70 percent capacity during peak hours. No single Steelhead appliance CPU usage should exceed 90 percent.

What This Report Tells You

The CPU Utilization report answers the following questions:

- How much of the CPU is being used?
- What is the average and peak percentage of the CPU being used?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

To view the CPU Utilization report

1. Choose Reports > Diagnostics > CPU Utilization to display the CPU Utilization page.

Figure 10-37. CPU Utilization Page

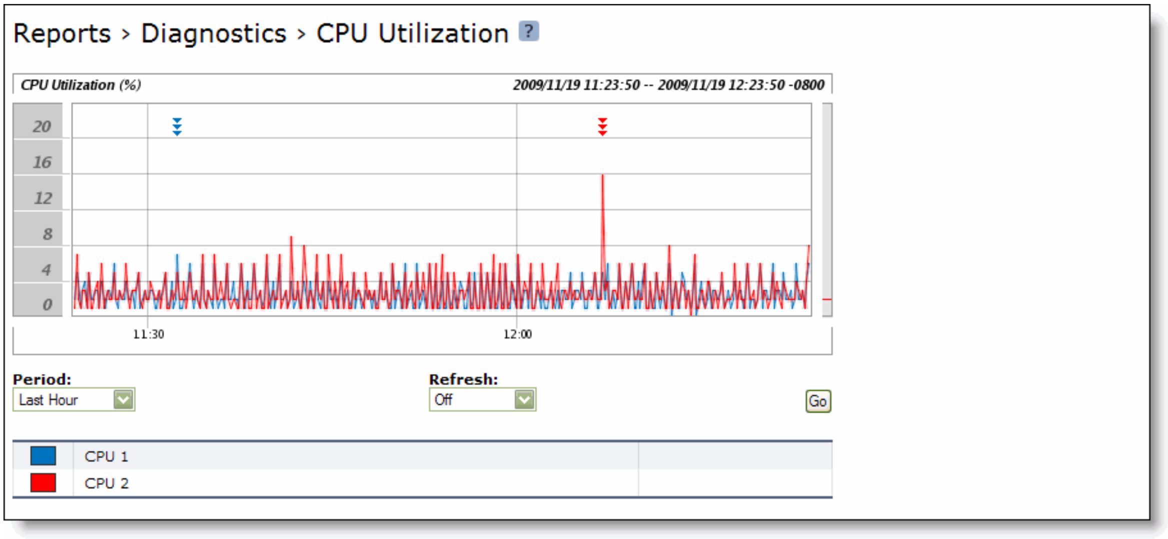

Reports > Diagnostics > CPU Utilization

CPU Utilization (%) 2009/11/19 11:23:50 -- 2009/11/19 12:23:50 -0800

Period: Last Hour

Refresh: Off

Go

🟦 CPU 1

🟥 CPU 2

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select Last Minute, 5 Minutes, Last Hour, Last Day, Last Week, Last Month, or Custom from the drop-down list. For Custom, enter the Start Time and End Time and click **Go**. Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select 10 seconds. • To refresh the report every 30 seconds, select 30 seconds. • To refresh the report every 60 seconds, select 60 seconds. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Disk Status Reports

The Disk Status report appears only on Steelhead appliance models 7050L and 7050M. It alerts you to a disk failure or recovery. A disk failure or recovery can occur when the optimization service is:

- not running.
- running, but idle because there is no traffic.
- handling optimized connections but not using the disk.
- writing to the disk.
- reading from the disk.

The Disk Status report includes the following information.

Field	Description
Disk	Specifies the disk number.
Status	Specifies the disk status: Failed, Missing, Invalid, or Online. If the report displays a failed disk status, go to Riverbed Support at https://support.riverbed.com
Task	Specifies the system component.

What This Report Tells You

The Disk Status report answers the following questions:

- How many disks are on the Steelhead appliance?
- What is the current status of each disk?
- What function is the disk performing?

To view the Disk Status report

- Choose Reports > Diagnostics > Disk Status to display the Disk Status page. This menu item appears only on Steelhead appliance models 7050L and 7050M.

Figure 10-38. Disk Status Page

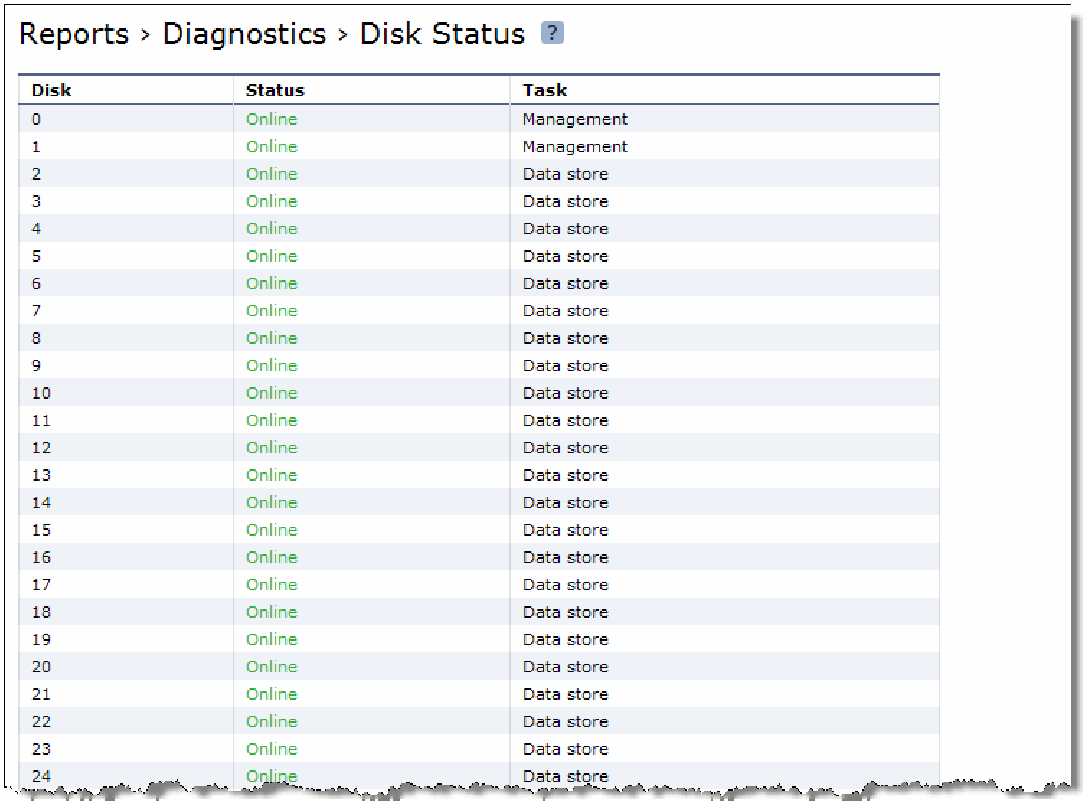

Reports > Diagnostics > Disk Status

Disk	Status	Task
0	Online	Management
1	Online	Management
2	Online	Data store
3	Online	Data store
4	Online	Data store
5	Online	Data store
6	Online	Data store
7	Online	Data store
8	Online	Data store
9	Online	Data store
10	Online	Data store
11	Online	Data store
12	Online	Data store
13	Online	Data store
14	Online	Data store
15	Online	Data store
16	Online	Data store
17	Online	Data store
18	Online	Data store
19	Online	Data store
20	Online	Data store
21	Online	Data store
22	Online	Data store
23	Online	Data store
24	Online	Data store

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Memory Paging Reports

The Memory Paging report provides the total number of memory pages, per second, utilized in the time period specified. It includes the following table of statistics that describe memory paging activity for the time period you specify.

Field	Description
Total Pages Swapped Out	Specifies the total number of pages swapped. If 100 pages are swapped approximately every two hours the Steelhead appliance is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support at https://support.riverbed.com.

Field	Description
Average Pages Swapped Out	Specifies the average number of pages swapped. If 100 pages are swapped every couple of hours the Steelhead appliance is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support at https://support.riverbed.com.
Peak Pages Swapped Out At <time> on <date>	Specifies the date and time that the peak number of pages were swapped.

What This Report Tells You

The Memory Paging report answers the following questions:

- How much memory is being used?
- What is the average and peak amount of memory pages swapped?

About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as GBs of bandwidth, percent (%) of data reduction, connection counts, and the like.

Three triangles near the top margin of the graph point to the value on the x-axis (the time) at which the peak occurred.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

Pie chart graphs do not indicate peaks or averages. Pie chart graphs represent the aggregate for the time period selected.

To view the Memory Paging report

1. Choose Reports > Diagnostics > Memory Paging to display the Memory Paging page.

Figure 10-39. Memory Paging Page

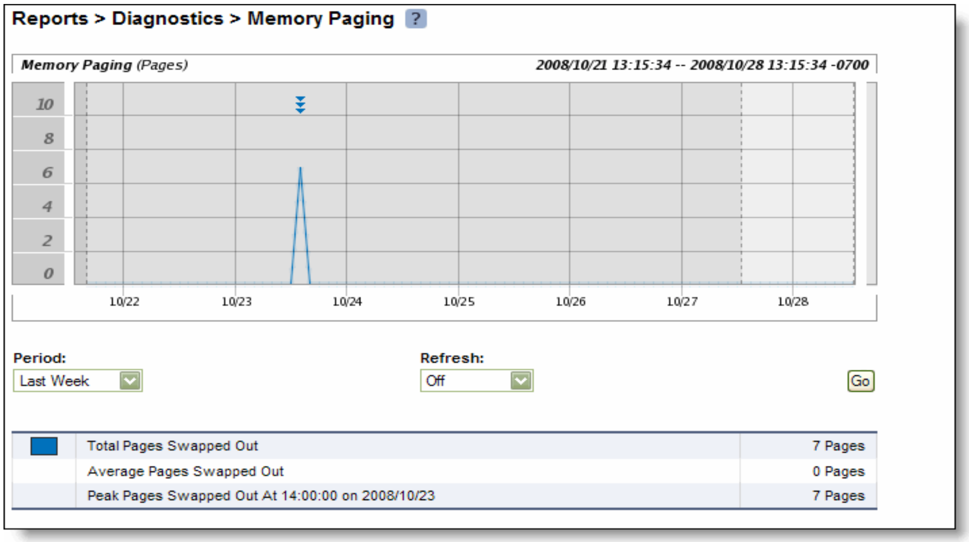

Period: Last Week

Refresh: Off

🟦	Total Pages Swapped Out	7 Pages
	Average Pages Swapped Out	0 Pages
	Peak Pages Swapped Out At 14:00:00 on 2008/10/23	7 Pages

2. Use the controls to customize the report as described in the following table.

Control	Description
Period	Select a period of **Last Hour** or **Last Day** from the drop-down list.
Refresh	Select a refresh rate from the drop-down list: • To refresh the report every 10 seconds, select **10 seconds**. • To refresh the report every 30 seconds, select **30 seconds**. • To refresh the report every 60 seconds, select **60 seconds**. • To turn refresh off, click **Off**.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Logs

Steelhead appliance log reports provide a high-level view of network activity. You can view both user and system logs.

- "Viewing User Logs" on page 475
- "Viewing System Logs" on page 477

Viewing User Logs

You can view user logs in the Reports > Diagnostics > User Logs page. The user log filters messages from the system log to display messages that are of immediate use to the system administrator.

View user logs to monitor system activity and to troubleshoot problems. For example, you can monitor who logged in, who logged out, and who entered particular CLI commands, alarms and errors. The most recent log events are listed first.

To view and customize user logs

1. Choose Reports > Diagnostics > User Logs to display the User Logs page.

Figure 10-40. User Logs Page

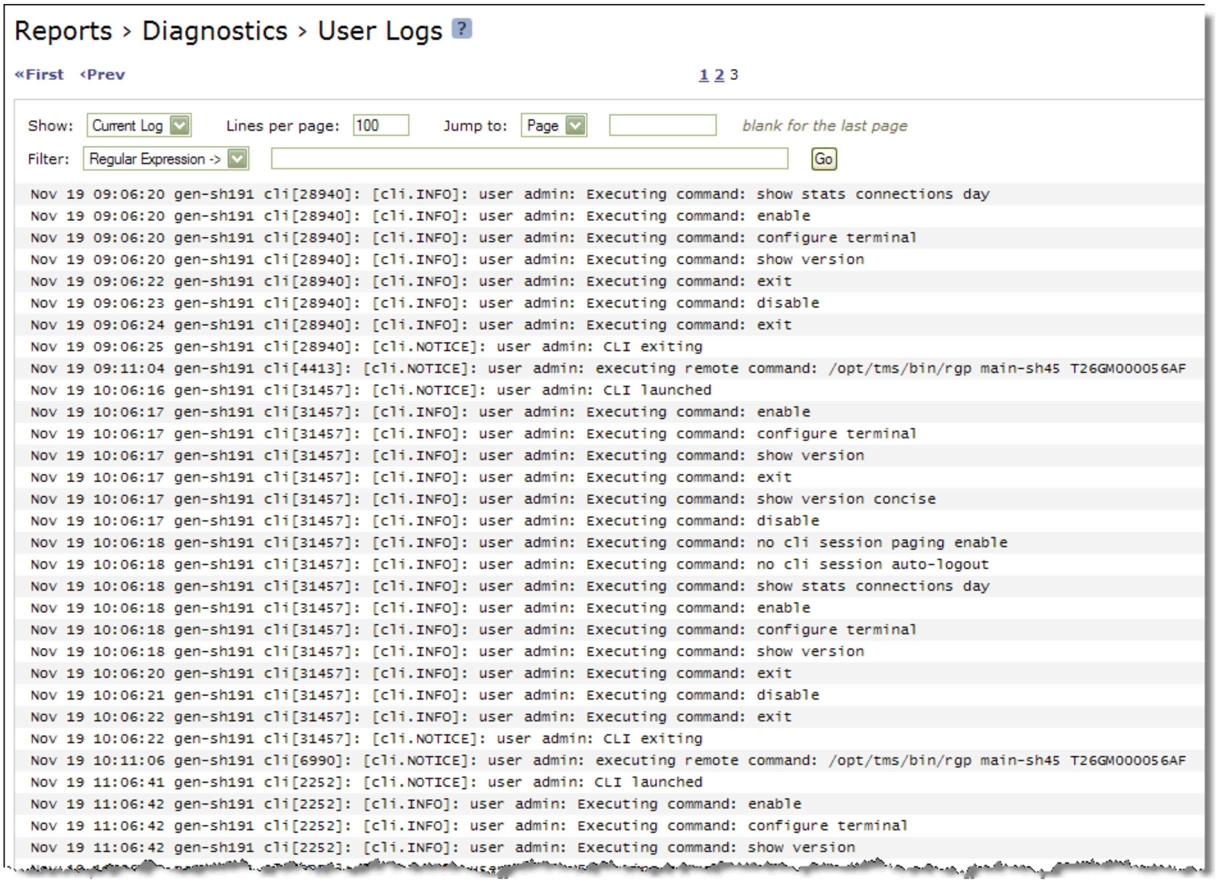

Show: Current Log Lines per page: 100 Jump to: Page *blank for the last page*

Filter: Regular Expression ->

```
Nov 19 09:06:20 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: show stats connections day
Nov 19 09:06:20 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: enable
Nov 19 09:06:20 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: configure terminal
Nov 19 09:06:20 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: show version
Nov 19 09:06:22 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: exit
Nov 19 09:06:23 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: disable
Nov 19 09:06:24 gen-sh191 cli[28940]: [cli.INFO]: user admin: Executing command: exit
Nov 19 09:06:25 gen-sh191 cli[28940]: [cli.NOTICE]: user admin: CLI exiting
Nov 19 09:11:04 gen-sh191 cli[4413]: [cli.NOTICE]: user admin: executing remote command: /opt/tms/bin/rgp main-sh45 T26GM000056AF
Nov 19 10:06:16 gen-sh191 cli[31457]: [cli.NOTICE]: user admin: CLI launched
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: enable
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: configure terminal
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: show version
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: exit
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: show version concise
Nov 19 10:06:17 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: disable
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: no cli session paging enable
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: no cli session auto-logout
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: show stats connections day
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: enable
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: configure terminal
Nov 19 10:06:18 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: show version
Nov 19 10:06:20 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: exit
Nov 19 10:06:21 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: disable
Nov 19 10:06:22 gen-sh191 cli[31457]: [cli.INFO]: user admin: Executing command: exit
Nov 19 10:06:22 gen-sh191 cli[31457]: [cli.NOTICE]: user admin: CLI exiting
Nov 19 10:11:06 gen-sh191 cli[6990]: [cli.NOTICE]: user admin: executing remote command: /opt/tms/bin/rgp main-sh45 T26GM000056AF
Nov 19 11:06:41 gen-sh191 cli[2252]: [cli.NOTICE]: user admin: CLI launched
Nov 19 11:06:42 gen-sh191 cli[2252]: [cli.INFO]: user admin: Executing command: enable
Nov 19 11:06:42 gen-sh191 cli[2252]: [cli.INFO]: user admin: Executing command: configure terminal
Nov 19 11:06:42 gen-sh191 cli[2252]: [cli.INFO]: user admin: Executing command: show version
```

2. Use the controls to customize the log as described in the following table.

Control	Description
Current Log	Select one of the archived logs or **Current Log** from the drop-down list.
Lines per Page	Specify the number of lines you want to display in the page.
Jump to	Select one of the following options from the drop-down list: • **Page** - Specify the number of pages you want to display. • **Time** - Specify the time for the log you want to display.

Control	Description
Filter	Select one of the following filtering options from the drop-down list: • **Regular expression** - Specify a regular expression on which to filter the log. • **Error or higher** - Displays Error level logs or higher. • **Warning or higher** - Displays Warning level logs or higher. • **Notice or higher** - Displays Notice level logs or higher. • **Info or higher** - Displays Info level logs or higher.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

You can continuously display new lines as the log grows and appends new data.

To view a continuous log

1. Choose Reports > Diagnostics > User Logs to display the User Logs page.
2. Customize the log as described in "To view and customize user logs" on page 476.
3. Click the continuous log icon in the upper-right corner of the page.

Note: If the continuous log does not appear after clicking the icon, a pair of Steelhead appliances might be optimizing HTTP traffic between the user's Web browser and the primary or auxiliary interface of the Steelhead for which the user is viewing the log, and they are buffering the HTTP response.

To display the continuous log, you can switch to HTTPS because the Steelhead appliances will not optimize HTTPS traffic. Alternatively, you can configure the other Steelhead appliances to pass-through traffic on the primary or auxiliary interfaces for port 80.

Viewing System Logs

You can view system logs in the Reports > Diagnostics > System Logs page. View System logs to monitor system activity and to troubleshoot problems. The most recent log events are listed first.

To customize system logs

1. Choose Reports > Diagnostics > System Logs to display the System Logs page.

Figure 10-41. System Logs Page

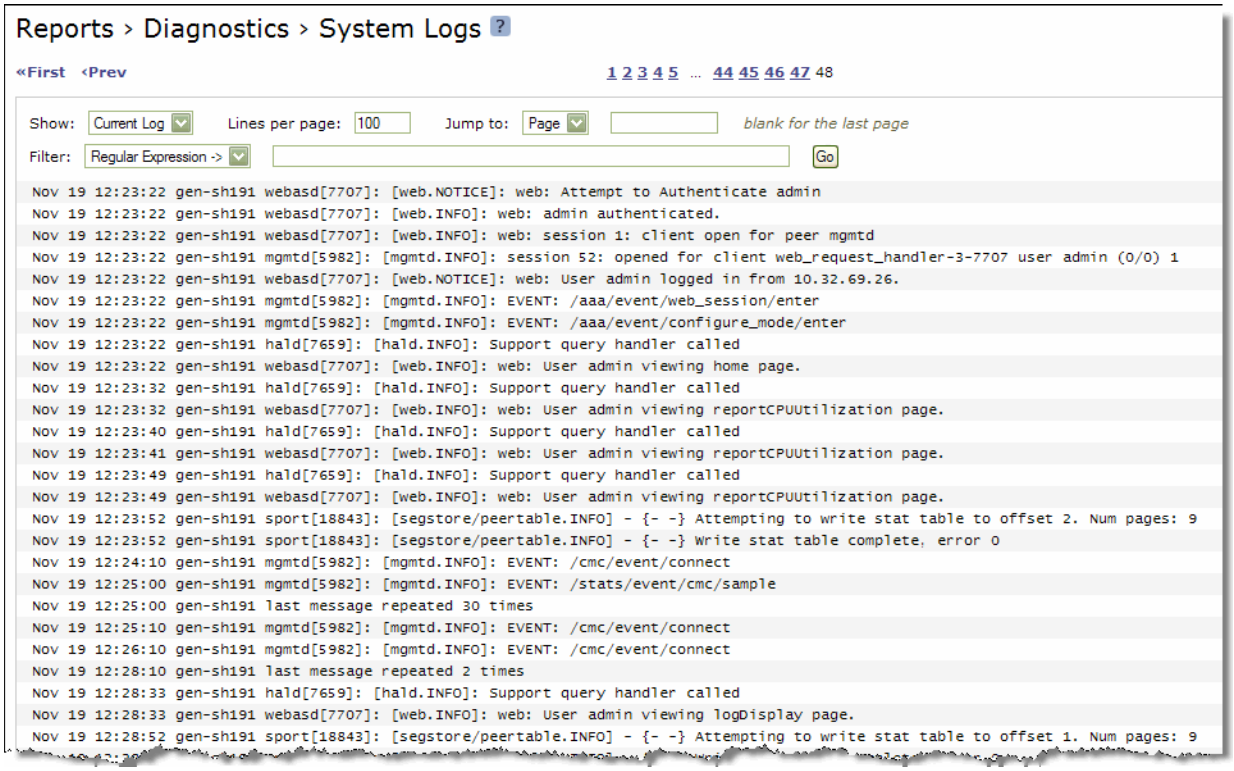

Reports > Diagnostics > System Logs

Show: Current Log Lines per page: 100 Jump to: Page

Filter: Regular Expression ->

```
Nov 19 12:23:22 gen-sh191 webasd[7707]: [web.NOTICE]: web: Attempt to Authenticate admin
Nov 19 12:23:22 gen-sh191 webasd[7707]: [web.INFO]: web: admin authenticated.
Nov 19 12:23:22 gen-sh191 webasd[7707]: [web.INFO]: web: session 1: client open for peer mgmtd
Nov 19 12:23:22 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: session 52: opened for client web_request_handler-3-7707 user admin (0/0) 1
Nov 19 12:23:22 gen-sh191 webasd[7707]: [web.NOTICE]: web: User admin logged in from 10.32.69.26.
Nov 19 12:23:22 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /aaa/event/web_session/enter
Nov 19 12:23:22 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /aaa/event/configure_mode/enter
Nov 19 12:23:22 gen-sh191 hald[7659]: [hald.INFO]: Support query handler called
Nov 19 12:23:22 gen-sh191 webasd[7707]: [web.INFO]: web: User admin viewing home page.
Nov 19 12:23:32 gen-sh191 hald[7659]: [hald.INFO]: Support query handler called
Nov 19 12:23:32 gen-sh191 webasd[7707]: [web.INFO]: web: User admin viewing reportCPUUtilization page.
Nov 19 12:23:40 gen-sh191 hald[7659]: [hald.INFO]: Support query handler called
Nov 19 12:23:41 gen-sh191 webasd[7707]: [web.INFO]: web: User admin viewing reportCPUUtilization page.
Nov 19 12:23:49 gen-sh191 hald[7659]: [hald.INFO]: Support query handler called
Nov 19 12:23:49 gen-sh191 webasd[7707]: [web.INFO]: web: User admin viewing reportCPUUtilization page.
Nov 19 12:23:52 gen-sh191 sport[18843]: [segstore/peertable.INFO] - {- -} Attempting to write stat table to offset 2. Num pages: 9
Nov 19 12:23:52 gen-sh191 sport[18843]: [segstore/peertable.INFO] - {- -} Write stat table complete, error 0
Nov 19 12:24:10 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /cmc/event/connect
Nov 19 12:25:00 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /stats/event/cmc/sample
Nov 19 12:25:00 gen-sh191 last message repeated 30 times
Nov 19 12:25:10 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /cmc/event/connect
Nov 19 12:26:10 gen-sh191 mgmtd[5982]: [mgmtd.INFO]: EVENT: /cmc/event/connect
Nov 19 12:28:10 gen-sh191 last message repeated 2 times
Nov 19 12:28:33 gen-sh191 hald[7659]: [hald.INFO]: Support query handler called
Nov 19 12:28:33 gen-sh191 webasd[7707]: [web.INFO]: web: User admin viewing logDisplay page.
Nov 19 12:28:52 gen-sh191 sport[18843]: [segstore/peertable.INFO] - {- -} Attempting to write stat table to offset 1. Num pages: 9
```

2. Use the controls to customize the report as described in the following table.

Control	Description
Current Log	Select one of the archived logs or Current Log from the drop-down list.
Lines per page	Specify the number of lines you want to display in the page.
Jump to	Select one of the following options from the drop-down list: • **Page** - Specify the number of pages you want to display. • **Time** - Specify the time for the log you want to display.
Filter	Select one of the following filtering options from the drop-down list: • **Regular expression** - Specify a regular expression on which to filter the log. • **Error or higher** - Displays Error level logs or higher. • **Warning or higher** - Displays Warning level logs or higher. • **Notice or higher** - Displays Notice level logs or higher. • **Info or higher** - Displays Info level logs or higher.
Go	Displays the report.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

To view a continuous log

1. Choose Reports > Diagnostics > System Logs to display the System Logs page.
2. Customize the log as described in "To customize system logs" on page 478.
3. Click the continuous log icon in the upper-right corner of the page.

Note: If the continuous log does not appear after clicking the icon, a pair of Steelhead appliances might be optimizing the HTTP traffic between the user's Web browser and the primary or auxiliary interface of the Steelhead for which the user is viewing the log, and they are buffering the HTTP response.

To display the continuous log, you can switch to HTTPS because the Steelhead appliances will not optimize HTTPS traffic. You might want to configure the other Steelhead appliances to pass-through traffic on the primary or auxiliary interfaces for port 80.

Downloading Log Files

This section describes how to download user and system log files.

You can download both user and system logs.

- "Downloading User Log Files" on page 479
- "Downloading System Log Files" on page 480

Downloading User Log Files

You can download user logs in the User Logs Download page. Download user logs to monitor system activity and to troubleshoot problems.

To download user logs

- Choose Reports > Diagnostics > User Logs Download to display the User Logs Download page.

Figure 10-42. User Logs Download Page

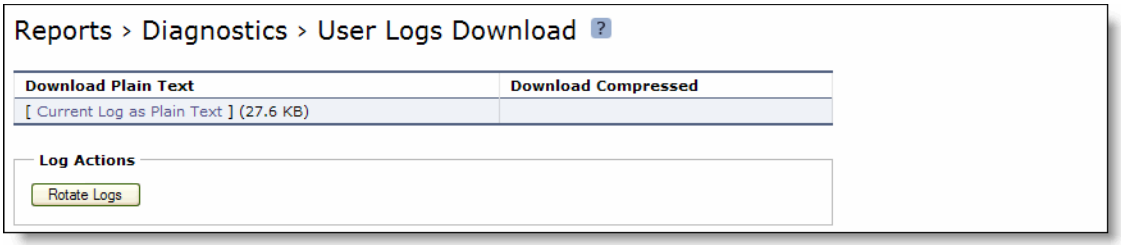

Reports > Diagnostics > User Logs Download ?

Download Plain Text	Download Compressed
[Current Log as Plain Text] (27.6 KB)	

Downloading System Log Files

You can download system logs in the System Logs Download page. Download system logs to monitor system activity and to troubleshoot problems.

To download system logs

1. Choose Reports > Diagnostics > System Logs Download to display the System Logs Download page.

Figure 10-43. System Logs Download Page

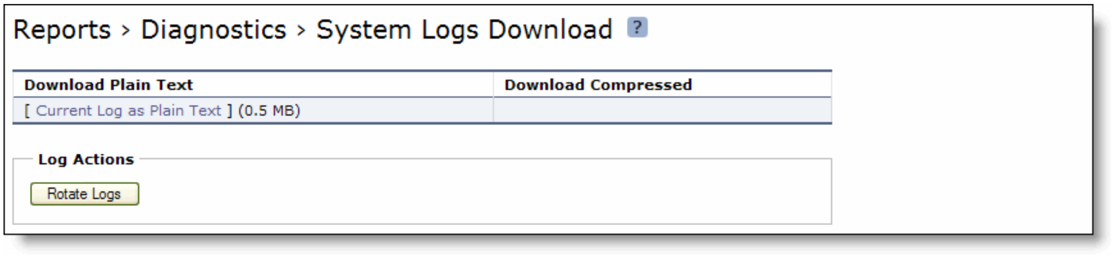

Reports > Diagnostics > System Logs Download

Download Plain Text	Download Compressed
[Current Log as Plain Text] (0.5 MB)	

2. Select the name of the log to display the dialog box to display or save the log to disk.

3. Click **Rotate Logs** to archive the current log to a numbered archived log file and then clear the log so that it is empty again.

Viewing the System Dumps List

You can display and download system dumps in the System Dump page. A system dump contains a copy of the kernel data on the system. System dump files can help you diagnose problems in the system.

To view system dump files

1. Choose Reports > Diagnostics > System Dumps to display the System Dumps page.

Figure 10-44. System Dumps Page

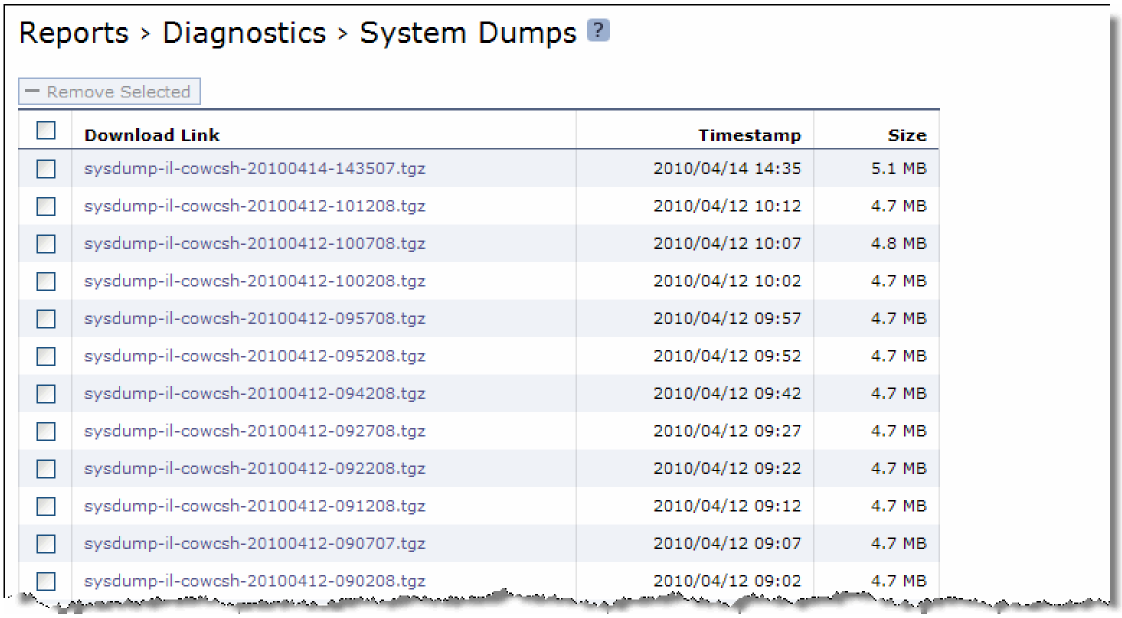

Reports > Diagnostics > System Dumps ?

□	Download Link	Timestamp	Size
□	sysdump-il-cowcsh-20100414-143507.tgz	2010/04/14 14:35	5.1 MB
□	sysdump-il-cowcsh-20100412-101208.tgz	2010/04/12 10:12	4.7 MB
□	sysdump-il-cowcsh-20100412-100708.tgz	2010/04/12 10:07	4.8 MB
□	sysdump-il-cowcsh-20100412-100208.tgz	2010/04/12 10:02	4.7 MB
□	sysdump-il-cowcsh-20100412-095708.tgz	2010/04/12 09:57	4.7 MB
□	sysdump-il-cowcsh-20100412-095208.tgz	2010/04/12 09:52	4.7 MB
□	sysdump-il-cowcsh-20100412-094208.tgz	2010/04/12 09:42	4.7 MB
□	sysdump-il-cowcsh-20100412-092708.tgz	2010/04/12 09:27	4.7 MB
□	sysdump-il-cowcsh-20100412-092208.tgz	2010/04/12 09:22	4.7 MB
□	sysdump-il-cowcsh-20100412-091208.tgz	2010/04/12 09:12	4.7 MB
□	sysdump-il-cowcsh-20100412-090707.tgz	2010/04/12 09:07	4.7 MB
□	sysdump-il-cowcsh-20100412-090208.tgz	2010/04/12 09:02	4.7 MB

2. Click **Download Link** to view a previously saved system dump.

3. Select the filename to open a file or save the file to disk.

4. Click **Include Statistics**.

5. Click **Generate System Dump** to generate a new system dump.

Tip: To remove an entry, check the box next to the name and click **Remove Selected**.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Viewing Process Dumps

You can display and download process dumps in the Process Dumps page. A process dump is a saved copy of memory including the contents of all memory, bytes, hardware registers, and status indicators. It is periodically taken to restore the system in the event of failure. Process dump files can help you diagnose problems in the system.

To view process dump files

1. Choose Reports > Diagnostics > Process Dumps to display the Process Dumps page.

Figure 10-45. Process Dumps Page

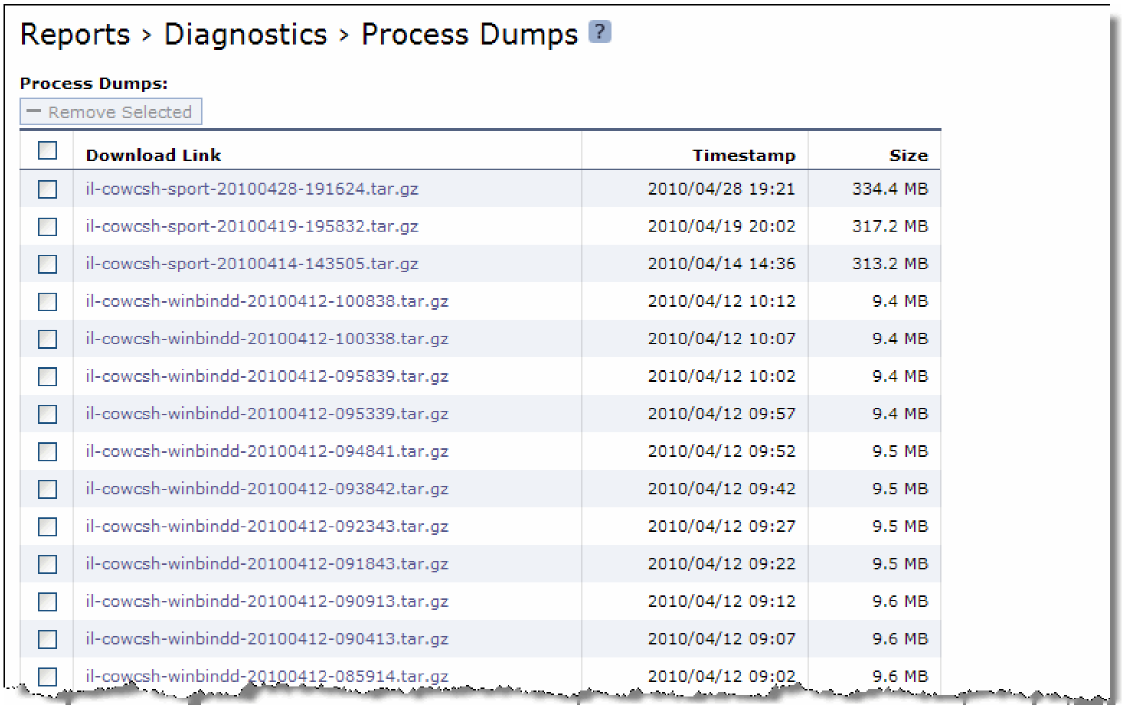

Reports > Diagnostics > Process Dumps

Process Dumps:

— Remove Selected

□	Download Link	Timestamp	Size
□	il-cowcsh-sport-20100428-191624.tar.gz	2010/04/28 19:21	334.4 MB
□	il-cowcsh-sport-20100419-195832.tar.gz	2010/04/19 20:02	317.2 MB
□	il-cowcsh-sport-20100414-143505.tar.gz	2010/04/14 14:36	313.2 MB
□	il-cowcsh-winbindd-20100412-100838.tar.gz	2010/04/12 10:12	9.4 MB
□	il-cowcsh-winbindd-20100412-100338.tar.gz	2010/04/12 10:07	9.4 MB
□	il-cowcsh-winbindd-20100412-095839.tar.gz	2010/04/12 10:02	9.4 MB
□	il-cowcsh-winbindd-20100412-095339.tar.gz	2010/04/12 09:57	9.4 MB
□	il-cowcsh-winbindd-20100412-094841.tar.gz	2010/04/12 09:52	9.5 MB
□	il-cowcsh-winbindd-20100412-093842.tar.gz	2010/04/12 09:42	9.5 MB
□	il-cowcsh-winbindd-20100412-092343.tar.gz	2010/04/12 09:27	9.5 MB
□	il-cowcsh-winbindd-20100412-091843.tar.gz	2010/04/12 09:22	9.5 MB
□	il-cowcsh-winbindd-20100412-090913.tar.gz	2010/04/12 09:12	9.6 MB
□	il-cowcsh-winbindd-20100412-090413.tar.gz	2010/04/12 09:07	9.6 MB
□	il-cowcsh-winbindd-20100412-085914.tar.gz	2010/04/12 09:02	9.6 MB

2. Select the filename to open a file or save the file to disk.

Tip: To remove an entry, check the box next to the name and click **Remove Selected**.

Tip: To print the report, choose File > Print in your Web browser to open the Print dialog box.

Capturing and Uploading TCP Dumps

You can capture, download, and upload TCP dumps in the Reports > Diagnostics > TCP Dumps page. TCP trace dump files contain summary information for every Internet packet received or transmitted on the interface. TCP trace dump files can help diagnose problems in the system.

RiOS provides an easy way to capture and retrieve multiple TCP trace dumps from the Management Console. You can generate trace dumps from multiple interfaces at the same time, limit the size of the trace dump, and schedule a specific date and time to generate a trace dump. Scheduling and limiting a trace dump by time or size allows unattended captures.

The top of the TCP Dumps page displays a list of existing TCP trace dumps and the bottom of the page displays controls to create a new trace dump. It also includes the trace dumps that are currently running. The Running Capture Name list includes TCP trace dumps running at a particular time. It includes TCP trace dumps started manually and also any dumps which were scheduled previously and are now running.

To capture TCP trace dumps

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.

Figure 10-46. TCP Dumps Page

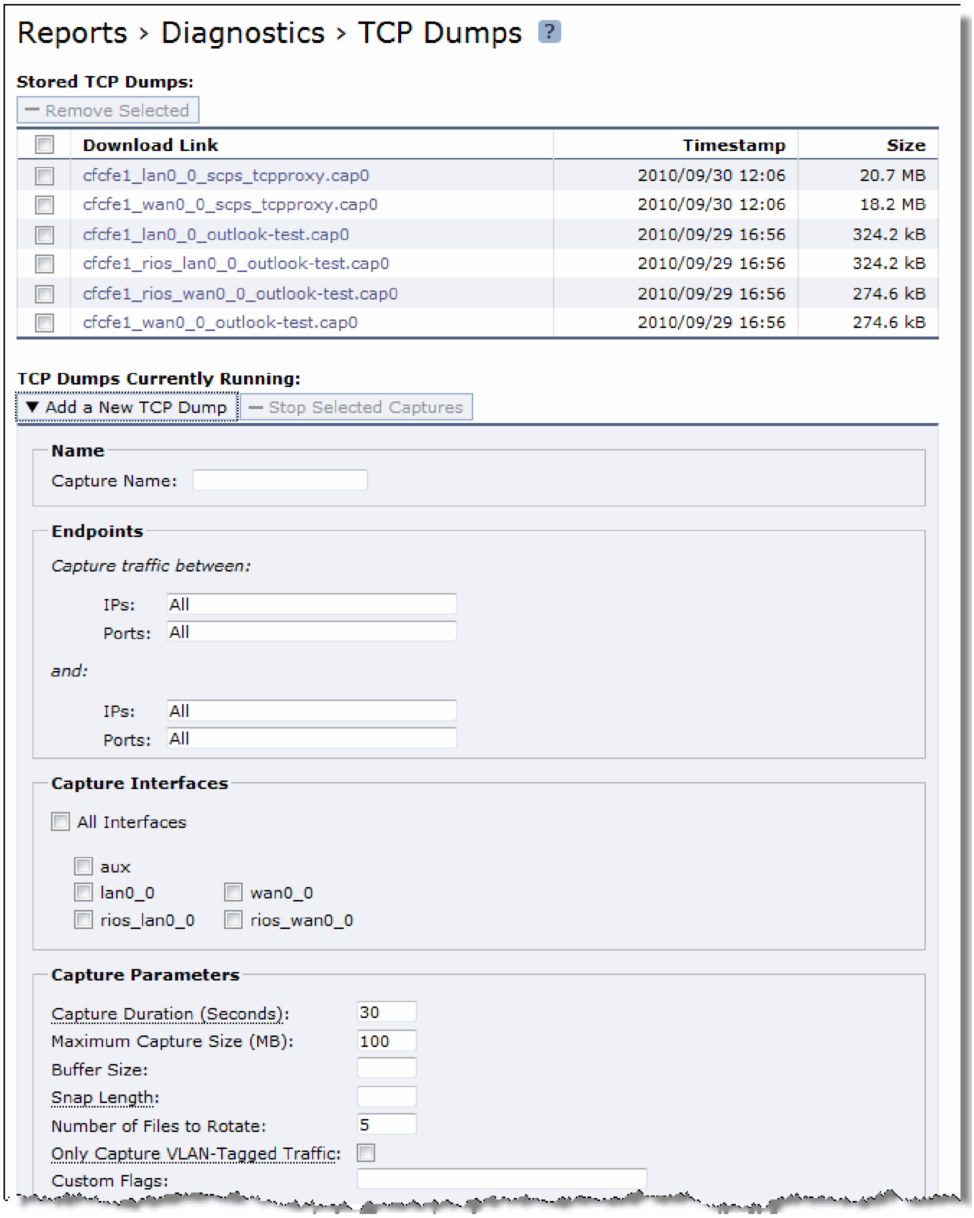

Reports > Diagnostics > TCP Dumps ?

Stored TCP Dumps:

— Remove Selected

□	Download Link	Timestamp	Size
□	cfcfe1_lan0_0_scps_tcpproxy.cap0	2010/09/30 12:06	20.7 MB
□	cfcfe1_wan0_0_scps_tcpproxy.cap0	2010/09/30 12:06	18.2 MB
□	cfcfe1_lan0_0_outlook-test.cap0	2010/09/29 16:56	324.2 kB
□	cfcfe1_rios_lan0_0_outlook-test.cap0	2010/09/29 16:56	324.2 kB
□	cfcfe1_rios_wan0_0_outlook-test.cap0	2010/09/29 16:56	274.6 kB
□	cfcfe1_wan0_0_outlook-test.cap0	2010/09/29 16:56	274.6 kB

TCP Dumps Currently Running:

▼ Add a New TCP Dump — Stop Selected Captures

Name

Capture Name:

Endpoints

Capture traffic between:

IPs: All
Ports: All

and:

IPs: All
Ports: All

Capture Interfaces

□ All Interfaces

□ aux
□ lan0_0 □ wan0_0
□ rios_lan0_0 □ rios_wan0_0

Capture Parameters

Capture Duration (Seconds): 30
Maximum Capture Size (MB): 100
Buffer Size:
Snap Length:
Number of Files to Rotate: 5
Only Capture VLAN-Tagged Traffic: □
Custom Flags:

2. Complete the configuration as described in the following table.

Control	Description
Add a New TCP Dump	Displays the controls for creating a TCP trace dump.
Capture Name	Specify the name of the capture file. The default filename uses the following format: *hostname_interface_timestamp*.cap Where *hostname* is the hostname of the Steelhead appliance, *interface* is the name of the interface selected for the trace (for example, lan0_0, wan0_0), and *timestamp* is in the YYYY-MM-DD-HH-MM-SS format. If this trace dump relates to an open Riverbed Support case, specify the capture filename case_*number* where number is your Riverbed Support case number; for example, case_12345. **Note:** The .cap file extension is not included with the filename when it appears in the capture queue.
Capture Traffic Between	**IPs** - Specify the source IP addresses. Separate multiple IP addresses with a comma to include all addresses bidirectionally. The default setting is all IP addresses. **Ports** - Specify the source ports. Separate multiple ports with a comma. The default setting is all ports. *and:* **IPs** - Specify the destination IP addresses. Separate multiple IP addresses with a comma to include all addresses bidirectionally. The default setting is all IP addresses. **Ports** - Specify the destination ports. Separate multiple ports with a comma. The default setting is all ports.
Capture Interfaces	Captures the TCP trace dump on the selected interface(s). You can select all interfaces or a physical, MIP, or RSP interface. The default setting is none. You must specify a capture interface. If you select several interfaces at a time, the data is automatically placed into separate capture files.
Capture Duration (Seconds)	Specify how long the capture runs, in seconds. The default value is 30. Leave this value blank to initiate a continuous trace. When a continuous trace reaches the maximum space allocation of 100 MB, the oldest file is overwritten.
Maximum Capture Size (MB)	Specify the maximum capture file size in MBs. The default value is 100. The recommended maximum capture file size is 1024 MBs (1 GB).
Buffer Size	Optionally, specify the maximum number of packets allowed to queue up while awaiting processing by the TCP trace dump. The default value is 154.
Snap Length	Optionally, specify the snap length value for the trace dump. Specify 0 for a full packet capture (recommended for CIFS, MAPI, and SSL traces). The default value is 1518.
Number of Files to Rotate	Specify how many TCP trace dump files to rotate. The default value is 5.
Only Capture VLAN Packets	Captures only VLAN-tagged packets within a trace dump for a trunk port (802.1Q). Enabling this setting filters the trace dump by capturing only VLAN-tagged packets. This setting applies to physical interfaces only because logical interfaces (inpath0_0, mgmt0_0) do not recognize VLAN headers.

Control	Description
Custom Flags	Specify custom flags to capture unidirectional traces. Examples: To capture all traffic to or from a single host host x.x.x.x To capture all traffic between a pair of hosts host x.x.x.x and host y.y.y.y To capture traffic between two hosts and two Steelhead inner channels: (host x.x.x.x and host y.y.y.y) or (host a.a.a.a and host b.b.b.b)
Schedule Dump	Schedules the trace dump to run at a later date and time.
Start Date	Specify a date to initiate the trace dump in the following format: YYYY/MM/DD
Start Time	Specify a time to initiate the trace dump in the following format: HH:MM:SS
Add	Adds the TCP trace dump to the capture queue.

Note: If a problem occurs with an immediate or scheduled TCP dump, the following message appears: “Error in tcpdump command. See System Log for details.” Check the trace dump for any syntax errors.

To view TCP trace dump files

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.

2. Under Stored TCP Dumps, select the trace dump name to open the file.

Tip: To print the TCP dump, select the trace dump filename under Stored TCP Dumps. When the file opens, choose File > Print in your Web browser to open the Print dialog box.

Tip: To remove an entry, check the box next to the name in the TCP dump list and click **Remove Selected**.

To stop a running TCP trace dump

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.

2. Select the trace dump filename in the Running Capture Name list.

3. Click **Stop Selected Captures**.

To upload the trace to Riverbed Support

In continuous mode, after you complete the capture, perform the following steps:

(For timed TCP dumps, start with step 2.)

1. On the TCP Dumps page, select the running TCP Dump and click **Stop Selected Captures**.

   The trace appears as a download link in the list of TCP dumps stored on the Steelhead appliance.

2. Select the top file in the Stored TCP Dumps list and save it locally.

   This file should contain the current date.

3. Compress (zip) the file and follow the upload instructions to share it with Riverbed Support:

   Attach the file(s) to your case at

   https://support.riverbed.com/cases/viewcases.htm

   —or—

   Upload the file(s) to ftp://ftp.riverbed.com/incoming

   (for FTP, be sure the file is prefixed with case_*number*).

   ```
   ftp ftp.riverbed.com
   User: anonymous
   Password: your_email@address
   ftp> cd /incoming
   ftp> bi
   ftp> put case_12345-tcpdump.zip
   ```

Checking Steelhead Appliance Health Status

You can run diagnostic tests on a Steelhead appliance in the Reports > Diagnostics > Health Check page.

RiOS v6.0 and later provides a convenient way to troubleshoot any Steelhead appliance problems by running a set of general diagnostic tests. Viewing the test results can pinpoint any issues with the appliance and significantly speed problem resolution.

To run diagnostic tests

1. Choose Reports > Diagnostics > Health Check to display the Health Check page.

Figure 10-47. Health Check Page

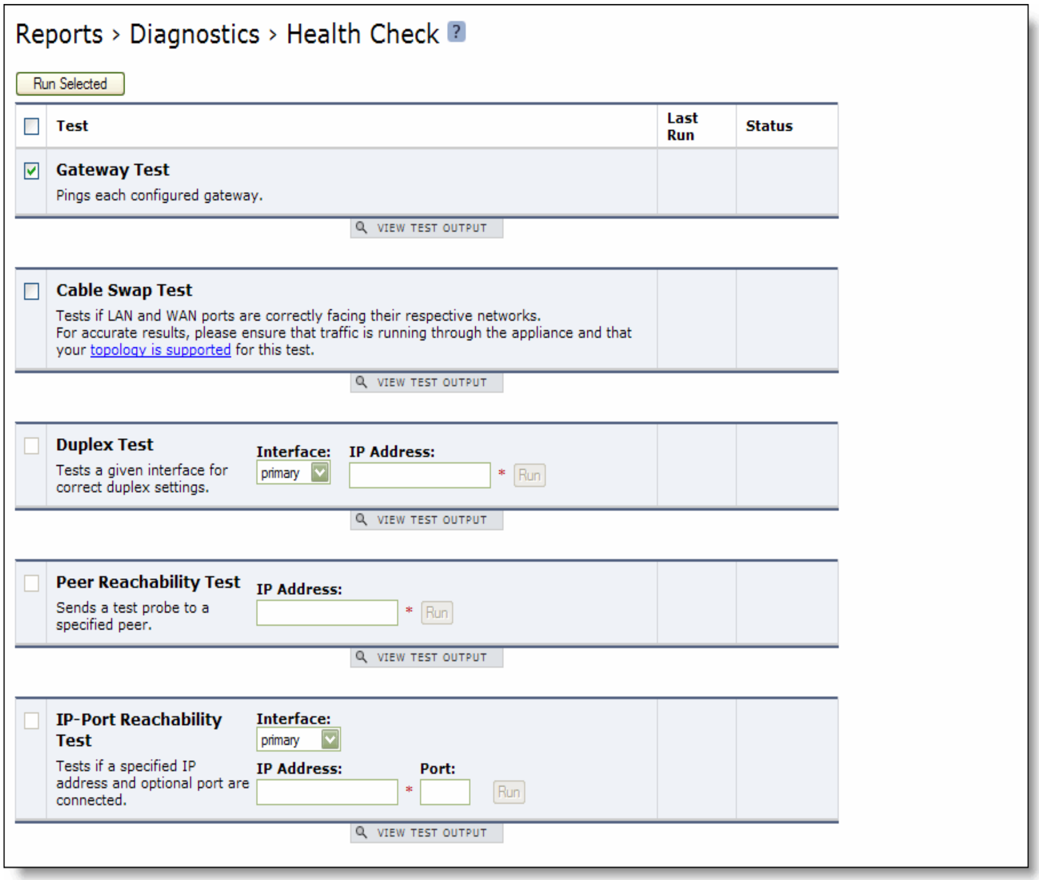

Reports > Diagnostics > Health Check

	Test	Last Run	Status
☑	**Gateway Test** Pings each configured gateway.		
□	**Cable Swap Test** Tests if LAN and WAN ports are correctly facing their respective networks. For accurate results, please ensure that traffic is running through the appliance and that your topology is supported for this test.		
□	**Duplex Test** Tests a given interface for correct duplex settings. Interface: primary IP Address: *		
□	**Peer Reachability Test** Sends a test probe to a specified peer. IP Address: *		
□	**IP-Port Reachability Test** Tests if a specified IP address and optional port are connected. Interface: primary IP Address: * Port:		

2. Complete the configuration as described in the following table.

Control	Description
Gateway Test	Determines if each configured gateway is connected correctly. Run this test to ping each configured gateway address with 4 packets and record the number of failed or successful replies. The test passes if all 4 packets are acknowledged. The default packet size is 64 bytes. If the test fails and all packets are lost, ensure the gateway IP address is correct and the Steelhead appliance is on the correct network segment. If the gateway is reachable from another source, check the connections between the Steelhead appliance and the gateway. If the test fails and only some packets are lost, check your duplex settings and other network conditions that might cause dropped packets.
Cable Swap Test	Ensures that the WAN and LAN cables on the Steelhead appliance are connected to the LAN and WAN of the network. The test enumerates the results by interface (one row entry per pair of bypass interfaces). By default, this test is disabled. **Important:** Certain network topologies might cause an incorrect result for this test. For the following topologies, Riverbed recommends that you confirm the test result manually: • Steelhead appliances deployed in virtual in-path mode. • Server-side Steelhead appliances that receive significant amounts of traffic from non-optimized sites. • Steelhead appliances that sit in the path between other Steelheads that are optimizing traffic. If the test fails, ensure a straight-through cable is not in use between an appliance port and a router, or that a crossover cable is not in use between an appliance port and a switch.
Duplex Test	Determines if the speed and duplex settings match on each side of the selected interface. If one side is different from the other, then traffic is sent at different rates on each side, causing a great deal of collision. This test runs the ping utility for 5 seconds with a packet size of 2500 bytes against the interface. • **Interface** - Specify an interface to test. • **IP Address** - Specify an IP address that is on the testing interface side. • **Run** - Click to run the test. The test passes if the system acknowledges 100% of the packets and receives responses from all packets. If any packets are lost, the test fails. If the test fails, ensure the speed and duplex settings of the appliance's Ethernet interface matches that of the switch ports to which it is connected. The test output records the percentage of any lost packets and number of collisions. **Note:** For accurate test results, traffic must be running through the Steelhead appliance.

Control	Description
Peer Reachability Test	Select to send a test probe to a specified peer and await the probe response. If a response is not received, the test fails. **Tip**: To view the current peer appliances, choose Reports > Optimization > Connected Appliances. • **IP Address** - Specify the IP address of the peer appliance to test. • **Run** - Click to run the test. **Notes**: • This test might not be accurate when the peer Steelhead appliance is configured out-of-path. • Do not specify the primary or auxiliary IP of the same Steelhead appliance displayed in the Connected Appliances report (the primary or aux IP to which the Steelhead appliance is connected). If the test fails, ensure that there are no firewalls, IDS/IPS, VPNs, or other security devices which may be stripping or dropping connection packets between Steelhead appliances.
IP Port Reachability Test	Select to determine whether a specified IP address and optional port is correctly connected. If you specify only an IP address, the test sends an ICMP message to the IP address. If you specify a port number, the test **telnets** to the port. • **Interface** - Optionally, specify an interface to test. • **IP Address** - Specify the IP address to test. • **Port** - Optionally, specify a port to test. • **Run** - Click to run the test. If the test fails, ensure that dynamic or static routing on your network is correctly configured and that the remote network is reachable from hosts on the same local subnet as this appliance.
Run Selected	Runs the selected tests.
View or Hide Test Output	Click to view or hide the test results.

Viewing the Test Status and Results

The Last Run column displays the time and date the last test was run.

The Status column displays **Initializing** temporarily while the page loads. When the test starts, the Status column displays **Running**, and then the test result appears in the Results column.

The Results column displays one of the following test results:

- **Passed**.
- **Failed**.
- **Undetermined** - A test with an undetermined status indicates that the test could not accurately determine a pass or fail test status.

To view diagnostic test results

1. Choose Reports > Diagnostics > Health Check to display the Health Check page.

2. Under the test name, click **View Test Output**.

Tip: To print the test results, click **View Test Output** and choose File > Print in your Web browser to open the Print dialog box.

Exporting Performance Statistics

You can export performance statistics in CSV format in the Export report. The CSV format allows you to easily import the statistics into spreadsheets and databases. You can open the CSV file in any text editor.

The CSV file contains commented lines (comments beginning with the # character) at the beginning of the file. These comments report what host generated the file, the report that was generated, time boundaries, the time the export occurred, and the version of the Steelhead appliance the file was exported from. The statistical values are provided in columns: the first column is the date and time of the statistic sample, the columns that follow contain the data.

To export statistics

1. Choose Reports > Export to display the Export page.

Figure 10-48. Export Page

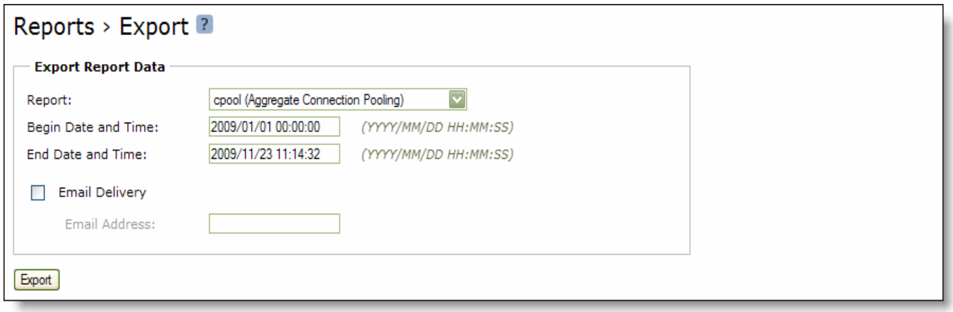

Report: cpool (Aggregate Connection Pooling)

Begin Date and Time: 2009/01/01 00:00:00 (YYYY/MM/DD HH:MM:SS)

End Date and Time: 2009/11/23 11:14:32 (YYYY/MM/DD HH:MM:SS)

□ Email Delivery

Email Address:

2. Use the controls to customize the report, as described in the following table.

Control	Description
Report	Select the report you want to export from the drop-down list.
Start Date (YYYY/MM/DD HH:MM:SS)	Specify a start date and time in the following format: YYYY/MM/DD HH:MM:SS
End Date (YYYY/MM/DD HH:MM:SS)	Specify an end date and time in the following format: YYYY/MM/DD HH:MM:SS
Report Type	Sends the report through email.
Email Recipient	Specify the email address of the recipient.
Export	Exports the report data.

APPENDIX A Steelhead Appliance Ports

This appendix provides a reference to ports used by the system. It includes the following sections:

Default Ports

The following table summarizes Steelhead appliance default ports with the port label: RBT-Proto.

Default Ports	Description
7744	Datastore synchronization port
7800	In-path port for appliance-to-appliance connections
7801	Network Address Translation (NAT) port
7810	Out-of-path server port
7820	Failover port for redundant appliances
7850	Connection forwarding (neighbor) port
7860	Interceptor appliance
7870	Steelhead Mobile

Note: Because optimization between Steelhead appliances typically takes place over a secure WAN, it is not necessary to configure company firewalls to support Steelhead-specific ports. If there are one or more firewalls between two Steelhead appliances, ports 7800 and 7810, must be passed through firewall devices located between the pair of Steelhead appliances. Also, SYN and SYN/ACK packets with the TCP option 76 must be passed through firewalls for auto-discovery to function properly. For the CMC, port 22 must be passed through for the firewall to function properly.

Commonly Excluded Ports

This section summarizes the ports that are commonly excluded from optimization in the Steelhead appliance.

If you have multiple ports that you want to exclude, create a port label and list the ports.

Application	Ports
PolyComm (video conferencing)	1503, 1720-1727, 3230-3253, 5060
Cisco IPTel	2000

Interactive Ports Forwarded by the Steelhead Appliance

A default in-path rule with the port label Interactive is automatically created in your system. This in-path rule automatically passes through traffic on interactive ports (for example, Telnet, TCP ECHO, remote logging, and shell).

Tip: If you do not want to automatically forward these ports, simply delete the Interactive rule in the Management Console.

The following table lists the interactive ports that are automatically forwarded by the Steelhead appliance.

Port	Description
7	TCP ECHO
23	Telnet
37	UDP/Time
107	Remote Telnet Service
179	Border Gateway Protocol
513	Remote Login
514	Shell
1494	Citrix
1718-1720	h323gatedisc
2000-2003	Cisco SCCp
2427	Media Gateway Control Protocol Gateway
2598	Citrix
2727	Media Gateway Control Protocol Call Agent
3389	MS WBT Server, TS/Remote Desktop
5060	SIP
5631	PC Anywhere

Port	Description
5900-5903	VNC
6000	X11

Secure Ports Forwarded by the Steelhead Appliance

A default in-path rule with the port label Secure is automatically created in your system. This in-path rule automatically passes through traffic on commonly secure ports (for example, ssh, https, and smtps).

Tip: If you do not want to automatically forward these ports, simply delete the Secure rule in the Management Console.

The following table lists the common secure ports that are automatically forwarded by the Steelhead appliance.

Type	Port	Description
ssh	22/tcp	SSH Remote Login Protocol
tacacs	49/tcp	TACACS+
https	443/tcp	http protocol over TLS/SSL
smtps	465/tcp	# SMTP over SSL (TLS)
nntps	563/tcp	nntp protocol over TLS/SSL (was snntp)
imap4-ssl	585/tcp	IMAP4+SSL (use 993 instead)
sshell	614/tcp	SSLshell
ldaps	636/tcp	ldap protocol over TLS/SSL (was sldap)
ftps-data	989/tcp	ftp protocol, data, over TLS/SSL
ftps	990/tcp	ftp protocol, control, over TLS/SSL
telnets	992/tcp	telnet protocol over TLS/SSL
imaps	993/tcp	imap4 protocol over TLS/SSL
pop3s	995/tcp	pop3 protocol over TLS/SSL (was spop3)
l2tp	1701/tcp	l2tp
pptp	1723/tcp	pptp
tftps	3713/tcp	TFTP over TLS

The following table contains the uncommon ports automatically forwarded by the Steelhead appliance.

Type	Port	Description
nsiiops	261/tcp	IIOP Name Service over TLS/SSL
ddm-ssl	448/tcp	DDM-Remote DB Access Using Secure Sockets
corba-iiop-ssl	684/tcp	CORBA IIOP SSL

Type	Port	Description
ieee-mms-ssl	695/tcp	IEEE-MMS-SSL
ircs	994/tcp	irc protocol over TLS/SSL
njenet-ssl	2252/tcp	NJENET using SSL
ssm-cssps	2478/tcp	SecurSight Authentication Server (SSL)
ssm-els	2479/tcp	SecurSight Event Logging Server (SSL)
giop-ssl	2482/tcp	Oracle GIOP SSL
ttc-ssl	2484/tcp	Oracle TTC SSL
groove	2492	GROOVE
syncserverssl	2679/tcp	Sync Server SSL
dicom-tls	2762/tcp	DICOM TLS
realsecure	2998/tcp	Real Secure
orbix-loc-ssl	3077/tcp	Orbix 2000 Locator SSL
orbix-cfg-ssl	3078/tcp	Orbix 2000 Locator SSL
cops-tls	3183/tcp	COPS/TLS
csvr-sslproxy	3191/tcp	ConServR SSL Proxy
xnm-ssl	3220/tcp	XML NM over SSL
msft-gc-ssl	3269/tcp	Microsoft Global Catalog with LDAP/SSL
networklenss	3410/tcp	NetworkLens SSL Event
xtrms	3424/tcp	xTrade over TLS/SSL
jt400-ssl	3471/tcp	jt400-ssl
seclayer-tls	3496/tcp	securitylayer over tls
vt-ssl	3509/tcp	Virtual Token SSL Port
jboss-iiop-ssl	3529/tcp	JBoss IIOP/SSL
ibm-diradm-ssl	3539/tcp	IBM Directory Server SSL
can-nds-ssl	3660/tcp	Candle Directory Services using SSL
can-ferret-ssl	3661/tcp	Candle Directory Services using SSL
linktest-s	3747/tcp	LXPRO.COM LinkTest SSL
asap-tcp-tls	3864/tcp	asap/tls tcp port
topflow-ssl	3885/tcp	TopFlow SSL
sdo-tls	3896/tcp	Simple Distributed Objects over TLS
sdo-ssh	3897/tcp	Simple Distributed Objects over SSH
iss-mgmt-ssl	3995/tcp	ISS Management Svcs SSL
suucp	4031/tcp	UUCP over SSL
wsm-server-ssl	5007/tcp	wsm server ssl

Type	Port	Description
sip-tls	5061/tcp	SIP-TLS
imqtunnels	7674/tcp	iMQ SSL tunnel
davsrcs	9802/tcp	WebDAV Source TLS/SSL
intrepid-ssl	11751/tcp	Intrepid SSL
rets-ssl	12109/tcp	RETS over SSL

APPENDIX B

Steelhead Appliance MIB

This appendix provides a reference to the Steelhead Enterprise MIB and SNMP traps. These tools allow for easy management of the Steelhead appliances and straightforward integration into existing network management systems.

This appendix includes the following sections:

Note: RiOS v6.0 and later includes support for integration of a configurable XML/SOAP API. The SOAP API allows a broad set of reporting and management actions to be taken from external Network Management Systems (for example, HP OpenView). Most statistics are exposed and many configuration actions can be taken remotely. For information on the SOAP API, see the *Steelhead Appliance API Guide*.

Accessing the Steelhead Enterprise MIB

The Steelhead Enterprise MIB monitors device status, peers, and provides network statistics for seamless integration into network management systems such as Hewlett Packard OpenView Network Node Manager, PRTG, and other SNMP browser tools.

For details on configuring and using these network monitoring tools, consult their individual Web sites.

The following guidelines describe how to download and access the Steelhead Enterprise MIB using common MIB browsing utilities:

- You can download the Steelhead Enterprise MIB (STEELHEAD-MIB.txt) from the Support page of the Management Console or from the Riverbed Support site at https://support.riverbed.com and load it into any MIB browser utility.
- Some utilities might expect a file type other than a text file. If this occurs, change the file type to the one expected.
- Some utilities assume that the root is mib-2 by default. If the utility sees a new node, such as enterprises, it might look under mib-2.enterprises. If this occurs, use .iso.org.dod.internet.private.enterprises.rbt as the root.

- Some command-line browsers might not load all MIB files by default. If this occurs, find the appropriate command option to load the STEELHEAD-MIB.txt file. For example, for NET-SNMP browsers: **snmpwalk -m all**.

Retrieving Optimized Traffic Statistics By Port

When you perform an snmpwalk on the Steelhead MIB object bwPortTable to display a table of statistics for optimized traffic by port, it retrieves only the monitored ports. The monitored ports include the default TCP ports and any ports you add. To view the monitored ports this object returns, choose Configure > System Settings > Monitored Ports or enter the following CLI command at the system prompt:

```
show stats settings bandwidth ports
```

To retrieve statistics for an individual port, perform an smnpget for that port. For example:

```
.iso.org.dod.internet.private.enterprises.rbt.products.steelhead.statistics.bandwidth.
bandwidthPerPort.bwPort Table.bwPortEntry.bwPortOutLan.port_number
```

SNMP Traps

Every Steelhead appliance supports SNMP traps and email alerts for conditions that require attention or intervention. An alarm fires for most, but not every, event and the related trap is sent. For most events, when the condition clears, the system clears the alarm and also sends out a clear trap. The clear traps are useful in determining when an event has been resolved.

This section describes the SNMP traps. It does not list the corresponding clear traps.

RiOS v6.0 and later includes support for SNMP v3.

You can view Steelhead appliance health at the top of each Management Console page, and by entering the CLI **show info** command, and through SNMP (health, systemHealth).

The Steelhead appliance tracks key hardware and software metrics and alerts you of any potential problems so you can quickly discover and diagnose issues. The health of an appliance falls into one of the following states:

- **Healthy** - The Steelhead is functioning and optimizing traffic.
- **Needs Attention** - Accompanies a healthy state to indicate management-related issues not affecting the ability of the Steelhead to optimize traffic.
- **Degraded** - The Steelhead is optimizing traffic but the system has detected an issue.
- **Admission Control** - The Steelhead is optimizing traffic but has reached its connection limit.
- **Critical** - The Steelhead may or may not be optimizing traffic; you need to address a critical issue.

The following table summarizes the SNMP traps sent out from the system to configured trap receivers and their effect on the Steelhead appliance health state.

Trap and OID	Steelhead State	Text	Description
procCrash (enterprises.17163.1.1.4.0.1)		A procCrash trap signifies that a process managed by PM has crashed and left a core file. The variable sent with the notification indicates which process crashed.	A process has crashed and subsequently been restarted by the system. The trap contains the name of the process that crashed. A system snapshot associated with this crash has been created on the appliance and is accessible via the CLI or the Management Console. Riverbed Support might need this information to determine the cause of the crash. No other action is required on the appliance as the crashed process is automatically restarted.
procExit (enterprises.17163.1.1.4.0.2)		A procExit trap signifies that a process managed by PM has exited unexpectedly, but not left a core file. The variable sent with the notification indicates which process exited.	A process has unexpectedly exited and been restarted by the system. The trap contains the name of the process. The process might have exited automatically or due to other process failures on the appliance. Review the release notes for known issues related to this process exit. If none exist, Contact Riverbed Support to determine the cause of this event. No other action is required on the appliance as the crashed process is automatically restarted.
cpuUtil (enterprises.17163.1.1.4.0.3)	Degraded	The average CPU utilization in the past minute has gone above the acceptable threshold.	Average CPU utilization has exceeded an acceptable threshold. If CPU utilization spikes are frequent, it might be because the system is undersized. Sustained CPU load can be symptomatic of more serious issues. Consult the CPU Utilization report to gauge how long the system has been loaded and also monitor the amount of traffic currently going through the appliance. A one-time spike in CPU is normal but Riverbed recommends reporting extended high CPU utilization to Riverbed Support. No other action is necessary as the alarm clears automatically.
pagingActivity (enterprises.17163.1.1.4.0.4)	Degraded	The system has been paging excessively (thrashing).	The system is running low on memory and has begun swapping memory pages to disk. This event can be triggered during a software upgrade while the optimization service is still running but there can be other causes. If this event triggers at any other time, generate a debug sysdump and send it to Riverbed Support. No other action is required as the alarm clears automatically.
smartError (enterprises.17163.1.1.4.0.5)		SMART has sent an event about a possible disk error.	A disk is about to fail. Contact Riverbed Support immediately. **Note:** Applicable to models 100, 200, 510, 520, 1010, 1020, 2010, 2510, and 2511 only.

Trap and OID	Steelhead State	Text	Description
peerVersionMismatch (enterprises.17163.1.1.4.0.6)	Degraded	Detected a peer with a mismatched software version.	The appliance has encountered another appliance which is running an incompatible version of system software. Refer to the CLI, Management Console, or the SNMP peer table to determine which appliance is causing the conflict. Connections with that peer will not be optimized, connections with other peers running compatible RiOS versions are unaffected. To resolve the problem, upgrade your system software. No other action is required as the alarm clears automatically.
bypassMode (enterprises.17163.1.1.4.0.7)	Critical	The appliance has entered bypass (failthru) mode.	The appliance has entered bypass mode and is now passing through all traffic unoptimized. This error is generated if the optimization service locks up or crashes. It can also be generated when the system is first turned on or turned off. If this trap is generated on a system that was previously optimizing and is still running, contact Riverbed Support.
raidError (enterprises.17163.1.1.4.0.8)	Degraded	An error has been generated by the RAID array.	A drive has failed in a RAID array. Consult the CLI or Management Console to determine the location of the failed drive. Contact Riverbed Support for assistance with installing a new drive, a RAID rebuild, or drive reseating. The appliance continues to optimize during this event. After the error is corrected, the alarm clears automatically. **Note:** Applicable to models 3010, 3510, 3020, 3520, 5010, 5520, 6020, and 6120 only.
storeCorruption (enterprises.17163.1.1.4.0.9)	Critical	The data store is corrupted.	Indicates the RiOS datastore is corrupt or has become incompatible with the current configuration. To clear the RiOS datastore of data, choose Configure > Maintenance > Services, select **Clear Data Store** and click **Restart** to restart the optimization service. If the alarm was caused by an unintended change to the configuration, you can change the configuration to match the previous RiOS datastore settings. Then restart the service without clearing the datastore to reset the alarm. Typical configuration changes that require a restart with a clear RiOS datastore are enabling the extended peer table or changing the datastore encryption.

Trap and OID	Steelhead State	Text	Description
admissionMemError (enterprises.17163.1.1.4.0.10)	Admission Control	Admission control memory alarm has been triggered.	The appliance has entered admission control due to memory consumption. The appliance is optimizing traffic beyond its rated capability and is unable to handle the amount of traffic passing through the WAN link. During this event, the appliance continues to optimize existing connections, but new connections are passed through without optimization. No other action is necessary as the alarm clears automatically when the traffic has decreased.
admissionConnError (enterprises.17163.1.1.4.0.11)	Admission Control	Admission control connections alarm has been triggered.	The appliance has entered admission control due to the number of connections and is unable to handle the amount of connections going over the WAN link. During this event, the appliance continues to optimize existing connections, but new connections are passed through without optimization. No other action is necessary as the alarm clears automatically when the traffic has decreased.
haltError (enterprises.17163.1.1.4.0.12)	Critical	The service is halted due to a software error.	The optimization service has halted due to a serious software error. See if a core dump or sysdump was created. If so, retrieve and contact Riverbed Support immediately.
serviceError (enterprises.17163.1.1.4.0.13)	Degraded	There has been a service error. Please consult the log file.	The optimization service has encountered a condition which might degrade optimization performance. Consult the system log for more information. No other action is necessary.
scheduledJobError (enterprises.17163.1.1.4.0.14)		A scheduled job has failed during execution.	A scheduled job on the system (for example, a software upgrade) has failed. To determine which job failed, use the CLI or the Management Console.
confModeEnter (enterprises.17163.1.1.4.0.15)		A user has entered configuration mode.	A user on the system has entered a configuration mode from either the CLI or the Management Console. A log in to the Management Console by user admin sends this trap as well. This is for notification purposes only; no other action is necessary.
confModeExit (enterprises.17163.1.1.4.0.16)		A user has exited configuration mode.	A user on the system has exited configuration mode from either the CLI or the Management Console. A log out of the Management Console by user admin sends this trap as well. This is for notification purposes only; no other action is necessary.

Trap and OID	Steelhead State	Text	Description
linkError (enterprises.17163.1.1.4.0.17)	Degraded	An interface on the appliance has lost its link.	The system has lost one of its Ethernet links due to a network event. Check the physical connectivity between the Steelhead appliance and its neighbor device. Investigate this alarm as soon as possible. Depending on what link is down, the system might no longer be optimizing and a network outage could occur. This is often caused by surrounding devices, like routers or switches interface transitioning. This alarm also accompanies service or system restarts on the Steelhead appliance.
nfsV2V4 (enterprises.17163.1.1.4.0.18)	Degraded	NFS v2/v4 alarm notification.	The Steelhead appliance has detected that either NFSv2 or NFSv4 is in use. The Steelhead appliance only supports NFSv3 and passes through all other versions. Check that the clients and servers are using NFSv3 and reconfigure if necessary.
powerSupplyError (enterprises.17163.1.1.4.0.19)	Degraded	A power supply on the appliance has failed (not supported on all models).	A redundant power supply on the appliance has failed on the appliance and needs to be replaced. Contact Riverbed Support for an RMA replacement as soon as practically possible.
asymRouteError (enterprises.17163.1.1.4.0.20)	Needs Attention	Asymmetric routes have been detected, certain connections might not have been optimized because of this.	Asymmetric routing has been detected on the network. This is very likely due to a failover event of an inner router or VPN. If so, no action needs to be taken. If not, contact Riverbed Support for further troubleshooting assistance.
fanError (enterprises.17163.1.1.4.0.21)	Degraded	A fan has failed on this appliance (not supported on all models).	A fan is failing or has failed and needs to be replaced. Contact Riverbed Support for an RMA replacement as soon practically possible.
memoryError (enterprises.17163.1.1.4.0.22)	Degraded	A memory error has been detected on the appliance (not supported on all models).	A memory error has been detected. A system memory stick might be failing. Try reseating the memory first. If the problem persists, contact Riverbed Support for an RMA replacement as soon as practically possible.
ipmi (enterprises.17163.1.1.4.0.23)	Degraded	An IPMI event has been detected on the appliance. Please check the details in the alarm report on the Web UI (not supported on all models).	An Intelligent Platform Management Interface (IPMI) event has been detected. Check the Alarm Status page for more detail. You can also view the IPMI events on the Steelhead appliance, by entering the CLI command: `show hardware error-log all`

Trap and OID	Steelhead State	Text	Description
configChange (enterprises.17163.1.1.4.0.24)		A change has been made to the system configuration.	A configuration change has been detected. Check the log files around the time of this trap to determine what changes were made and whether they were authorized.
datastoreWrapped (enterprises.17163.1.1.4.0.25)		The datastore has wrapped around.	The datastore on the Steelhead appliance went through an entire cycle and is removing data to make space for new data. This is normal behavior unless it wraps too quickly, which might indicate the datastore is undersized. If message is received every seven days or less, investigate traffic patterns and datastore sizing.
temperatureWarning (enterprises.17163.1.1.4.0.26)	Degraded	The system temperature has exceeded the threshold.	The appliance temperature is a configurable notification. By default, this notification is set to trigger when the appliance reached 70 degrees Celsius. Raise the alarm trigger temperature if it is normal for the Steelhead appliance to get that hot, or reduce the temperature of the Steelhead appliance.
temperatureCritical (enterprises.17163.1.1.4.0.27)	Critical	The system temperature has reached a critical stage.	This trap/alarm triggers a critical state on the appliance. This alarm occurs when the appliance temperature reaches 90 degrees Celsius. The temperature value is not user-configurable. Reduce the appliance temperature.
cfConnFailure (enterprises.17163.1.1.4.0.28)	Degraded	Unable to establish connection with the specified neighbor.	The connection cannot be established with a connection forwarding neighbor. This alarm clears automatically the next time all neighbors connect successfully.
cfConnLostEos (enterprises.17163.1.1.4.0.29)	Degraded	Connection lost since end of stream was received from the specified neighbor.	The connection has been closed by the connection forwarding neighbor. This alarm clears automatically the next time all neighbors connect successfully.
cfConnLostErr (enterprises.17163.1.1.4.0.30)	Degraded	Connection lost due to an error communicating with the specified neighbor.	The connection has been lost with the connection forwarding neighbor due to an error. This alarm clears automatically the next time all neighbors connect successfully.
cfKeepaliveTimeout (enterprises.17163.1.1.4.0.31)	Degraded	Connection lost due to lack of keep-alives from the specified neighbor.	The connection forwarding neighbor has not responded to a keep-alive message within the time-out period, indicating that the connection has been lost. This alarm clears automatically when all neighbors of the Steelhead appliance are responding to keep-alive messages within the time-out period.

Trap and OID	Steelhead State	Text	Description
cfAckTimeout (enterprises.17163.1.1.4.0.32)	Degraded	Connection lost due to lack of ACKs from the specified neighbor.	The connection has been lost because requests have not been acknowledged by a connection forwarding neighbor within the set time-out threshold. This alarm clears automatically the next time all neighbors receive an ACK from this neighbor and the latency of that acknowledgment is less than the set time-out threshold.
cfReadInfoTimeout (enterprises.17163.1.1.4.0.33)	Degraded	Timeout reading info from the specified neighbor.	The Steelhead appliance has timed out while waiting for an initialization message from the connection forwarding neighbor. This alarm clears automatically when the Steelhead appliance is able to read the initialization message from all of its neighbors.
cfLatencyExceeded (enterprises.17163.1.1.4.0.34)	Degraded	Connection forwarding latency with the specified neighbor has exceeded the threshold.	The amount of latency between connection forwarding neighbors has exceeded the specified threshold. The alarm clears automatically when the latency falls below the specified threshold.
sslPeeringSCEPAutoReenroll Error (enterprises.17163.1.1.4.0.35)	Needs Attention	There is an error in the automatic re-enrollment of the SSL peering certificate.	An SSL peering certificate has failed to re-enroll with the Simple Certificate Enrollment Protocol (SCEP).
crlError (enterprises.17163.1.1.4.0.36)	Needs Attention	CRL polling fails.	The polling for SSL peering CAs has failed to update the Certificate Revocation List (CRL) within the specified polling period. This alarm clears automatically when the CRL is updated.
datastoreSyncFailure (enterprises.17163.1.1.4.0.37)	Degraded	Data store sync has failed.	The datastore synchronization between two Steelhead appliances has been disrupted and the datastores are no longer synchronized.
secureVaultNeedsUnlock (enterprises.17163.1.1.4.0.38)	Needs Attention	SSL acceleration and the secure data store cannot be used until the secure vault has been unlocked.	The secure vault is locked. SSL traffic is not being optimized and the datastore cannot be encrypted. Check the Alarm Status page for more details. The alarm clears when the secure vault is unlocked.
secureVaultNeedsRekey (enterprises.17163.1.1.4.0.39)	Needs Attention	If you wish to use a non-default password for the secure vault, the password must be rekeyed. Please see the Knowledge Base solution 5592 for more details.	The secure vault password needs to be verified or reset. Initially, the secure vault has a default password known only to the RiOS software so the Steelhead appliance can automatically unlock the vault during system startup. For details, check the Alarm Status page and see the Knowledge Base solution 5592. The alarm clears when you verify the default password or reset the password.

Trap and OID	Steelhead State	Text	Description
secureVaultInitError (enterprises.17163.1.1.4.0.40)		An error was detected while initializing the secure vault. Please contact Riverbed Support.	An error occurred while initializing the secure vault after a RiOS software version upgrade. Contact Riverbed Support.
configSave (enterprises.17163.1.1.4.0.41)		The current appliance configuration has been saved.	A configuration has been saved either by entering the `write mem` CLI command or by clicking **Save** in the Management Console. This message is for security notification purposes only; no other action is necessary.

tcpDumpStarted (enterprises.17163.1.1.4.0.42)		A TCP dump has been started.	A user has started a TCP dump on the Steelhead appliance by entering a
			`tcpdump`
			or
			`tcpdump-x`
			command from the CLI. This message is for security notification purposes only; no other action is necessary.

tcpDumpScheduled (enterprises.17163.1.1.4.0.43)		A TCP dump has been scheduled.	A user has started a TCP dump on the Steelhead appliance by entering a
			`tcpdump`
			or
			`tcpdump-x`
			command with a scheduled start time from the CLI. This message is for security notification purposes only; no other action is necessary.

newUserCreated (enterprises.17163.1.1.4.0.44)		A new user has been created.	A new Role-Based Management user has been created using the CLI or the Management Console. This message is for security notification purposes only; no other action is necessary.
diskError (enterprises.17163.1.1.4.0.45)		Disk error has been detected.	A disk error has been detected. A disk might be failing. Try reseating the memory first. If the problem persists, contact Riverbed Support.
wearWarning (enterprises.17163.1.1.4.0.46)		Accumulated SSD write cycles passed predefined level.	Triggers on Steelhead appliance models 7050L and 7050M. A Solid State Disk (SSD) has reached 95% of its write cycle limit. Contact Riverbed Support.
cliUserLogin (enterprises.17163.1.1.4.0.47)		A user has just logged-in via CLI.	A user has logged in to the Steelhead appliance using the Command Line Interface. This message is for security notification purposes only; no other action is necessary.

Trap and OID	Steelhead State	Text	Description
cliUserLogout (enterprises.17163.1.1.4.0.48)		A CLI user has just logged-out.	A user has logged out of the Steelhead appliance using the Command Line Interface using the Quit command or ^D. This message is for security notification purposes only; no other action is necessary.
webUserLogin (enterprises.17163.1.1.4.0.49)		A user has just logged-in via the Web UI.	A user has logged in to the Steelhead appliance using the Management Console. This message is for security notification purposes only; no other action is necessary.
webUserLogout (enterprises.17163.1.1.4.0.50)		A user has just logged-out via the Web UI.	A user has logged out of the Steelhead appliance using the Management Console. This message is for security notification purposes only; no other action is necessary.
trapTest (enterprises.17163.1.1.4.0.51)		Trap Test	An SNMP trap test has occurred on the Steelhead appliance. This message is informational and no action is necessary.
admissionCpuError (enterprises.17163.1.1.4.0.52)	Admission Control	Optimization service is experiencing high CPU utilization.	The appliance has entered admission control due to high CPU use. During this event, the appliance continues to optimize existing connections, but new connections are passed through without optimization. No other action is necessary as the alarm clears automatically when the CPU usage has decreased.
admissionTcpError (enterprises.17163.1.1.4.0.53)	Admission Control	Optimization service is experiencing high TCP memory pressure.	The appliance has entered admission control due to high TCP memory use. During this event, the appliance continues to optimize existing connections, but new connections are passed through without optimization. No other action is necessary as the alarm clears automatically when the TCP memory pressure has decreased.
systemDiskFullError (enterprises.17163.1.1.4.0.54)		One or more system partitions is full or almost full.	The alarm clears when the system partitions fall below usage thresholds.

Trap and OID	Steelhead State	Text	Description

domainJoinError (enterprises.17163.1.1.4.0.55)		An attempt to join a domain failed.	An attempt to join a Windows domain has failed.
			The number one cause of failing to join a domain is a significant difference in the system time on the Windows domain controller and the Steelhead appliance. When the time on the domain controller and the Steelhead appliance do not match, the following error message appears:
			`lt-kinit: krb5_get_init_creds: Clock skew too great`
			Riverbed recommends using NTP time synchronization to synchronize the client and server clocks. It is critical that the Steelhead appliance time is the same as on the Active Directory controller. Sometimes an NTP server is down or inaccessible, in which case there can be a time difference.You can also disable NTP if it is not being used and manually set the time. You must also verify that the time zone is correct.
			A domain join can fail when the DNS server returns an invalid IP address for the Domain Controller. When a DNS misconfiguration occurs during an attempt to join a domain, the following error messages appear:
			`Failed to join domain: failed to find DC for domain <domain name>`
			`Failed to join domain : No Logon Servers`
			Additionally, the Domain Join alarm triggers and messages similar to the following appear in the logs:
			`Oct 13 14:47:06 bravo-sh81 rcud[10014]: [rcud/main/.ERR] - {--} Failed to join domain: failed to find DC for domain GEN-VCS78DOM.COM`
			When you encounter this error, go to the Configure > Networking > Host Settings page and verify that the DNS settings are correct.

certsExpiringError (enterprises.17163.1.1.4.0.56)		Some x509 certificates may be expiring.	The service has detected some x.509 certificates used for Network Administration Access to the Steelhead appliance that are close to their expiration dates. The alarm clears when the x.509 certificates are updated.

Trap and OID	Steelhead State	Text	Description
licenseError (enterprises.17163.1.1.4.0.57)		The main Steelhead license has expired, been removed, or become invalid.	A license on the Steelhead appliance has been removed, has expired, or is invalid. The alarm clears when a valid license is added or updated.
hardwareError (enterprises.17163.1.1.4.0.58)		Hardware error detected.	A hardware error has been detected.
sysdetailError (enterprises.17163.1.1.4.0.59)	Needs Attention	Error is found in System Detail Report.	Triggers an alarm when any top-level module on the system detail report is in error.
admissionMapiError (enterprises.17163.1.1.4.0.60)	Degraded	New MAPI connections will be passed through due to high connection count.	The total number of MAPI optimized connections have exceeded the maximum admission control threshold. By default, the maximum admission control threshold is 85% of the total maximum optimized connection count for the client-side Steelhead appliance. The Steelhead appliance reserves the remaining 15% so the MAPI admission control does not affect the other protocols. The 85% threshold is applied only to MAPI connections. RiOS is now passing through MAPI connections from new clients but continues to intercept and optimize MAPI connections from existing clients (including new MAPI connections from these clients). RiOS continues optimizing non-MAPI connections from all clients. This alarm is disabled by default. The alarm clears automatically when the MAPI traffic has decreased; however, it can take one minute for the alarm to clear. **Important:** MAPI admission control cannot solve a general Steelhead appliance Admission Control Error (enterprises.17163.1.1.4.0.11); however, it can help to prevent it from occurring.
neighborIncompatibility (enterprises.17163.1.1.4.0.61)	Degraded	Serial cascade misconfiguration has been detected.	Check your auto-peering configuration. Restart the optimization service to clear the alarm.

Trap and OID	Steelhead State	Text	Description
flashError (enterprises.17163.1.1.4.0.62)		Flash hardware error detected.	At times, the USB flash drive that holds the system images may become unresponsive. When this happens, the system is unable to write a new upgrade image to the flash drive without first power cycling the system. Reboot using either the Management Console or the CLI **reload** command to automatically power cycle the Steelhead appliance and restore the flash drive to proper function.
lanWanLoopError (enterprises.17163.1.1.4.0.63)		LAN-WAN loop detected. System will not optimize new connections until this error is cleared.	A LAN-WAN network loop has been detected between the LAN and WAN interfaces on a Virtual Steelhead. This can occur when you connect the LAN and WAN virtual NICs to the same vSwitch or physical NIC. This alarm triggers when a Virtual Steelhead starts up, and clears after you connect each LAN and WAN virtual interface to a distinct virtual switch and physical NIC (through the vSphere Networking tab) and then reboot the Virtual Steelhead.

Acronyms and Abbreviations

AAA. Authentication, Authorization, and Accounting.

ACL. Access Control List.

ACK. Acknowledgment Code.

ACS. (Cisco) Access Control Server.

AD. Active Directory.

ADS. Active Directory Services.

AES. Advanced Encryption Standard.

APT. Advanced Packaging Tool.

AR. Asymmetric Routing.

ARP. Address Resolution Protocol.

BDP. Bandwidth-Delay Product.

BW. Bandwidth.

CA. Certificate Authority.

CAD. Computer Aided Design.

CDP. Cisco Discovery Protocol.

CHD. Computed Historical Data.

CIFS. Common Internet File System.

CLI. Command-Line Interface.

CMC. Central Management Console.

CPU. Central Processing Unit.

CRL. Certificate Revocation List.

CRM. Customer Relationship Management.

CSR. Certificate Signing Request.

CSV. Comma-Separated Value.

DC. Domain Controller.

DER. Distinguished Encoding Rules.

DES. Data Encryption Standard.

DHCP. Dynamic Host Configuration Protocol.

DID. Deployment ID.

DMZ. Demilitarized Zone.

DNS. Domain Name Service.

DR. Data Replication.

DSA. Digital Signature Algorithm.

DSCP. Differentiated Services Code Point.

ECC. Error-Correcting Code.

ERP. Enterprise Resource Planning.

ESD. Electrostatic Discharge.

FCIP. Fiber Channel over IP

FDDI. Fiber Distributed Data Interface.

FIFO. First in First Out.

FIPS. Federal Information Processing Standards.

FSID. File System ID.

FTP. File Transfer Protocol.

GB. Gigabytes.

GMT. Greenwich Mean Time.

GPO. Group Policy Object.

GRE. Generic Routing Encapsulation.

GUI. Graphical User Interface.

HFSC. Hierarchical Fair Service Curve.

HSRP. Hot Standby Routing Protocol.

HSTCP. High-Speed Transmission Control Protocol.

HTTP. HyperText Transport Protocol.

HTTPS. HyperText Transport Protocol Secure.

ICA. Independent Computing Architecture.

ICMP. Internet Control Message Protocol.

ID. Identification Number.

IETF. Internet Engineering Task Force.

IGP. Interior Gateway Protocol.

IKE. Internet Key Exchange.

IOS. (Cisco) Internetwork Operating System.

IP. Internet Protocol.

IPMI. Intelligent Platform Management Interface.

IPSec. Internet Protocol Security Protocol.

ISL. InterSwitch Link. Also known as Cisco InterSwitch Link Protocol.

L2. Layer-2.

L4. Layer-4.

LAN. Local Area Network.

LDAP. Lightweight Directory Access Protocol.

LED. Light-Emitting Diode.

LRU. Least Recently Used.

LZ. Lempel-Ziv.

MAC. Media Access Control.

MAPI. Messaging Application Protocol Interface.

MDI, MDI-X. Medium Dependent Interface-Crossover.

MEISI. Microsoft Exchange Information Store Interface.

MIB. Management Information Base.

MOTD. Message of the Day.

MS GPO. Microsoft Group Policy Object.

MS SMS. Microsoft Systems Management Server.

MS-SQL. Microsoft Structured Query Language.

MSFC. Multilayer Switch Feature Card.

MSI Package. Microsoft Installer Package.

MTU. Maximum Transmission Unit.

MX-TCP. Max-Speed TCP.

NAS. Network Attached Storage.

NAT. Network Address Translate.

NFS. Network File System.

NIS. Network Information Services.

NSPI. Name Service Provider Interface.

NTLM. Windows NT LAN Manager.

NTP. Network Time Protocol.

OSI. Open System Interconnection.

OSPF. Open Shortest Path First.

PAP. Password Authentication Protocol.

PBR. Policy-Based Routing.

PCI. Peripheral Component Interconnect.

PEM. Privacy Enhanced Mail.

PFS. Proxy File Service.

PKCS12. Public Key Cryptography Standard #12.

PRTG. Paessler Router Traffic Grapher.

PSU. Power Supply Unit.

QoS. Quality of Service.

RADIUS. Remote Authentication Dial-In User Service.

RAID. Redundant Array of Independent Disks.

RCU. Riverbed Copy Utility.

ROFS. Read-Only File System.

RPC. Remote Procedure Call.

RSA. Rivest-Shamir-Adleman Encryption Method by RSA Security.

RSP. Riverbed Services Platform.

SA. Security Association.

SAP. System Application Program.

SCP. Secure Copy Program.

SCEP. Simple Certificate Enrollment Protocol.

SCPS. Space Communications Protocol Standards.

SDR. Scalable Data Referencing.

SDR-A. Scalable Data Referencing - Adaptive.

SDR-M. Scalable Data Referencing - Memory.

SEL. System Event Log.

SFQ. Stochastic Fairness Queuing.

SMB. Server Message Block.

SMI. Structure of Management Information.

SMTP. Simple Mail Transfer Protocol.

SNMP. Simple Network Management Protocol.

SOAP. Simple Object Access Protocol

SPAN. Switched Port Analyzer.

SQL. Structured Query Language.

SRDF. Symmetric Remote Data Facility

SRDF/A. Symmetric Remote Data Facility/Asynchronous

SSH. Secure Shell.

SSL. Secure Sockets Layer.

SYN. Synchronize.

SYN/ACK. Synchronize/Acknowledgement.

TA. Transaction Acceleration.

TACACS+. Terminal Access Controller Access Control System.

TCP. Transmission Control Protocol.

TCP/IP. Transmission Control Protocol/Internet Protocol.

ToS. Type of Service.

TP. Transaction Prediction.

TTL. Time to Live.

U. Unit.

UDP. User Diagram Protocol.

UNC. Universal Naming Convention.

URL. Uniform Resource Locator.

USM. User-based Security Model.

UTC. Universal Time Code.

VACM. View-Based Access Control Model.

VGA. Video Graphics Array.

VLAN. Virtual Local Area Network.

VoIP. Voice over IP.

VWE. Virtual Window Expansion.

WAN. Wide Area Network.

WCCP. Web Cache Communication Protocol.

WOC. WAN Optimization Controller.

XOR. Exclusive OR logic.

Index