

# **Steelhead® Appliance Management Console User's Guide**

Steelhead® DX Appliance

Version 8.5.2  
January 2014



© 2014 Riverbed Technology. All rights reserved.

Riverbed®, Cloud Steelhead®, Granite™, Interceptor®, RiOS®, Steelhead®, Think Fast®, Virtual Steelhead®, Whitewater®, Mazu®, Cascade®, Shark®, AirPcap®, BlockStream™, SkipWare®, TurboCap®, WinPcap®, Wireshark®, TrafficScript®, FlyScript™, WWOS™, and Stingray™ are trademarks or registered trademarks of Riverbed Technology, Inc. in the United States and other countries. Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed Technology. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein cannot be used without the prior written consent of Riverbed Technology or their respective owners.

Akamai® and the Akamai wave logo are registered trademarks of Akamai Technologies, Inc. SureRoute is a service mark of Akamai. Apple and Mac are registered trademarks of Apple, Incorporated in the United States and in other countries. Cisco is a registered trademark of Cisco Systems, Inc. and its affiliates in the United States and in other countries. EMC, Symmetrix, and SRDF are registered trademarks of EMC Corporation and its affiliates in the United States and in other countries. IBM, iSeries, and AS/400 are registered trademarks of IBM Corporation and its affiliates in the United States and in other countries. Juniper Networks and Junos are registered trademarks of Juniper Networks, Incorporated in the United States and other countries. Linux is a trademark of Linus Torvalds in the United States and in other countries. Microsoft, Windows, Vista, Outlook, and Internet Explorer are trademarks or registered trademarks of Microsoft Corporation in the United States and in other countries. Oracle and JInitiator are trademarks or registered trademarks of Oracle Corporation in the United States and in other countries. UNIX is a registered trademark in the United States and in other countries, exclusively licensed through X/Open Company, Ltd. VMware, ESX, ESXi are trademarks or registered trademarks of VMware, Incorporated in the United States and in other countries.

This product includes software developed by the University of California, Berkeley (and its contributors), EMC, and Comtech AHA Corporation. This product is derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm.

NetApp Manageability Software Development Kit (NM SDK), including any third-party software available for review with such SDK which can be found at <http://communities.netapp.com/docs/DOC-1152>, and are included in a NOTICES file included within the downloaded files.

For a list of open source software (including libraries) used in the development of this software along with associated copyright and license agreements, see the Riverbed Support site at <https://support.riverbed.com>.

This documentation is furnished "AS IS" and is subject to change without notice and should not be construed as a commitment by Riverbed Technology. This documentation may not be copied, modified or distributed without the express authorization of Riverbed Technology and may be used only in connection with Riverbed products and services. Use, duplication, reproduction, release, modification, disclosure or transfer of this documentation is restricted in accordance with the Federal Acquisition Regulations as applied to civilian agencies and the Defense Federal Acquisition Regulation Supplement as applied to military agencies. This documentation qualifies as "commercial computer software documentation" and any use by the government shall be governed solely by these terms. All other use is prohibited. Riverbed Technology assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.



Riverbed Technology  
199 Fremont Street  
San Francisco, CA 94105

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

Part Number  
712-00208-01

# Contents

<b>Preface.....</b>	<b>1</b>
About This Guide .....	1
Audience .....	2
Document Conventions.....	2
Product Dependencies and Compatibility.....	3
Hardware and Software Dependencies.....	3
Ethernet Network Compatibility .....	3
SNMP-Based Management Compatibility.....	4
Additional Resources .....	4
Release Notes .....	4
Riverbed Documentation and Support Knowledge Base.....	5
Contacting Riverbed.....	6
Internet .....	6
Technical Support .....	6
Professional Services .....	6
Documentation.....	6
 <b>Chapter 1 - Overview of the Management Console .....</b>	 <b>7</b>
Using the Management Console.....	7
Connecting to the Management Console .....	8
The Home Page.....	9
Navigating in the Management Console .....	10
Getting Help .....	11
Next Steps .....	12
 <b>Chapter 2 - Modifying Host and Network Interface Settings .....</b>	 <b>15</b>
Modifying General Host Settings .....	15
Modifying Base Interfaces .....	18
IPv6 Support.....	18
Modifying In-Path Interfaces .....	25
Configuring a Management In-Path Interface .....	28

<b>Chapter 3 - Configuring In-Path Rules.....</b>	<b>33</b>
In-Path Rules Overview .....	33
Creating In-Path Rules for Packet-Mode Optimization.....	34
Default In-Path Rules .....	35
Configuring In-Path Rules.....	36
<b>Chapter 4 - Configuring Optimization Features.....</b>	<b>47</b>
Configuring General Service Settings.....	47
Enabling Basic Deployment Options .....	47
Enabling Failover.....	48
Configuring General Service Settings.....	49
Enabling Peering and Configuring Peering Rules .....	53
About Regular and Enhanced Automatic Discovery .....	53
Configuring Peering.....	54
Improving Performance.....	58
Configuring CPU and Data Streamlining Settings .....	58
Configuring TCP, Satellite Optimization, and High-Speed TCP .....	59
Optimizing TCP and Satellite WANs .....	59
High-Speed TCP Optimization.....	71
Configuring Service Ports.....	72
Configuring Host Labels.....	73
Creating a Host Label.....	73
Resolving Hostnames .....	74
Viewing the Hostname Resolution Summary .....	74
Modifying Hostnames or Subnets in a Host Label.....	75
Configuring Port Labels.....	75
Modifying Ports in a Port Label .....	76
Configuring FCIP Optimization.....	77
Viewing FCIP Connections .....	78
FCIP Rules (VMAX-to-VMAX Traffic Only).....	78
Configuring SRDF Optimization.....	80
Viewing SRDF Connections .....	82
Setting a Custom Data Reduction Level for an RDF Group .....	82
Creating SRDF Rules (VMAX-to-VMAX Traffic Only) .....	84
Configuring SnapMirror Optimization .....	85
How a Steelhead Appliance Optimizes SnapMirror Traffic.....	86
<b>Chapter 5 - Configuring Network Integration Features.....</b>	<b>91</b>
Configuring Asymmetric Routing Features .....	91
Troubleshooting Asymmetric Routes .....	93
Configuring Connection Forwarding Features .....	94
Configuring IPSec Encryption .....	97

Configuring Subnet Side Rules.....	99
Configuring Flow Statistics .....	101
Enabling Flow Export .....	101
Applying QoS Policies .....	107
QoS Overview .....	107
QoS DX8000 Series Recommendations.....	110
Basic or Advanced Outbound QoS.....	111
QoS Classes.....	111
Configuring Outbound QoS (Basic).....	115
Overview .....	115
Adding a Remote Site .....	117
Adding an Application .....	121
Adding a Service Policy.....	125
Configuring Outbound QoS (Advanced) .....	128
Creating QoS Classes .....	130
Modifying QoS Classes or Rules .....	141
Enabling MX-TCP Queue Policies (Advanced Outbound QoS only).....	143
Configuring Inbound QoS.....	144
How a Steelhead Appliance Identifies and Shapes Inbound Traffic.....	146
Creating Inbound QoS Classes .....	148
Selecting WAN Paths Dynamically.....	153
Configuring Path Selection .....	153
Path Selection Use Cases .....	156
Configuring Simplified Routing Features.....	159
Configuring WCCP .....	160
Verifying a Multiple In-Path Interface Configuration.....	165
Modifying WCCP Group Settings.....	166
Configuring Hardware-Assist Rules.....	166
<b>Chapter 6 - Configuring SSL and a Secure Inner Channel .....</b>	<b>169</b>
Configuring SSL Server Certificates and Certificate Authorities .....	169
How Does SSL Work? .....	169
Prerequisite Tasks .....	170
Configuring CRL Management .....	174
Managing CRL Distribution Points (CDPs).....	175
Configuring Secure Peers .....	176
Secure Inner Channel Overview .....	176
Enabling Secure Peers .....	177
Configuring Peer Trust .....	177
Configuring Advanced and SSL Cipher Settings.....	183
Setting Advanced SSL Options.....	183
Configuring SSL Cipher Settings .....	184

<b>Chapter 7 - Managing Steelhead Appliances .....</b>	<b>187</b>
Starting and Stopping the Optimization Service .....	187
Configuring Scheduled Jobs .....	188
Upgrading Your Software.....	189
Rebooting and Shutting Down the Steelhead Appliance .....	191
Managing Licenses and Model Upgrades.....	191
Flexible Licensing Overview .....	192
Installing a License .....	192
Viewing Permissions .....	194
Managing Configuration Files .....	194
Configuring General Security Settings .....	196
Managing User Permissions .....	197
Capability-Based Accounts .....	197
Managing Password Policy .....	200
Selecting a Password Policy .....	200
Setting RADIUS Servers .....	203
Configuring TACACS+ Access .....	205
Unlocking the Secure Vault .....	206
Configuring a Management ACL.....	207
ACL Management Rules .....	208
Configuring Web Settings .....	210
Managing Web SSL Certificates.....	211
Enabling REST API Access .....	213
 <b>Chapter 8 - Configuring System Administrator Settings .....</b>	 <b>217</b>
Configuring Alarm Settings .....	217
Setting Announcements.....	224
Configuring Email Settings .....	224
Configuring Log Settings.....	226
Filtering Logs by Application or Process .....	228
Configuring the Date and Time .....	230
Current NTP Server Status.....	230
NTP Authentication .....	231
NTP Servers.....	231
Configuring Monitored Ports .....	234
Configuring SNMP Settings.....	235
Configuring SNMP v3 .....	237
SNMP Authentication and Access Control.....	238

<b>Chapter 9 - Viewing Reports and Logs.....</b>	<b>243</b>
Overview .....	244
Navigating the Report Layout .....	245
Viewing Current Connection Reports .....	248
What This Report Tells You .....	248
Viewing Connection History Reports .....	268
What This Report Tells You .....	269
About Report Graphs .....	269
About Report Data .....	269
Viewing Connection Forwarding Reports .....	271
What This Report Tells You .....	271
About Report Graphs .....	271
About Report Data .....	271
Viewing Outbound QoS Reports .....	273
What This Report Tells You .....	273
About Report Graphs .....	273
About Report Data .....	273
Viewing Inbound QoS Reports .....	275
What This Report Tells You .....	275
About Report Graphs .....	275
About Report Data .....	275
Viewing Path Selection Reports .....	277
Viewing Interface Information .....	278
Viewing Path Information .....	279
What This Report Tells You .....	281
Viewing Top Talkers Reports .....	281
What This Report Tells You .....	282
About Report Data .....	282
Viewing Traffic Summary Reports .....	284
What This Report Tells You .....	285
About Report Data .....	285
Viewing WAN Throughput Reports .....	287
What This Report Tells You .....	287
About Report Graphs .....	287
About Report Data .....	288
Viewing Application Statistics Reports .....	289
What This Report Tells You .....	290
About Report Graphs .....	290
About Report Data .....	290
Viewing Application Visibility Reports .....	291
What This Report Tells You .....	292
About Report Graphs .....	292
About Report Data .....	292

Viewing Interface Counter Reports.....	294
What This Report Tells You.....	294
Viewing TCP Statistics Reports .....	295
What This Report Tells You.....	295
Viewing Optimized Throughput Reports .....	296
What This Report Tells You.....	297
About Report Graphs.....	297
About Report Data .....	297
Viewing Bandwidth Optimization Reports .....	299
What This Report Tells You.....	300
About Report Graphs.....	300
About Report Data .....	300
Viewing Peer Reports.....	301
What This Report Tells You.....	302
Viewing SRDF Reports.....	302
What This Report Tells You.....	302
About Report Graphs.....	303
About Report Data .....	303
Viewing SnapMirror Reports .....	304
What This Report Tells You.....	305
About Report Graphs.....	305
About Report Data .....	305
Viewing Alarm Status Reports.....	307
What This Report Tells You.....	317
Viewing CPU Utilization Reports .....	317
What This Report Tells You.....	317
About Report Graphs.....	317
Viewing Memory Paging Reports .....	319
What This Report Tells You.....	319
About Report Graphs.....	319
Viewing TCP Memory Reports.....	320
What This Report Tells You.....	321
About Report Graphs.....	321
About Report Data .....	322
Viewing System Details Reports.....	323
What This Report Tells You.....	323
Checking Network Health Status.....	323
Viewing Logs .....	326
Viewing User Logs .....	326
Viewing System Logs.....	327
Downloading Log Files.....	328
Downloading User Log Files .....	328
Downloading System Log Files.....	329
Generating System Dumps.....	329



Viewing Process Dumps .....	330
Capturing and Uploading TCP Dump Files .....	331
Troubleshooting .....	335
Custom Flag Use Examples.....	335
IPv6 Custom Flag Use Examples.....	336
Stopping a TCP Dump After an Event Occurs.....	336
Viewing a TCP Dump .....	337
Uploading a TCP Dump .....	338
Exporting Performance Statistics.....	339
<b>Appendix A - Steelhead Appliance MIB.....</b>	<b>341</b>
Accessing the Steelhead Appliance Enterprise MIB.....	341
Retrieving Optimized Traffic Statistics by Port.....	342
SNMP Traps.....	342
<b>Appendix B - Steelhead Appliance Ports .....</b>	<b>359</b>
Granite Ports.....	359
Default Ports.....	360
Commonly Excluded Ports .....	360
Interactive Ports Forwarded by the Steelhead Appliance .....	360
Secure Ports Forwarded by the Steelhead Appliance .....	361
<b>Appendix C - Application Signatures for AFE .....</b>	<b>365</b>
List of Recognized Applications.....	365
<b>Index .....</b>	<b>403</b>



# Preface

Welcome to the *Steelhead Appliance Management Console User's Guide*. 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. This preface includes these sections:

- [“About This Guide” on page 1](#)
- [“Product Dependencies and Compatibility” on page 3](#)
- [“Additional Resources” on page 4](#)
- [“Contacting Riverbed” on page 6](#)

---

## About This Guide

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

This guide includes information relevant to these products:

- Riverbed Optimization System (RiOS)
- Riverbed Steelhead appliance (Steelhead appliance)
- Riverbed Steelhead DX appliance (Steelhead DX)
- Riverbed Steelhead CX appliance (Steelhead CX)
- Riverbed Steelhead EX appliance (Steelhead EX)
- Riverbed Virtual Steelhead appliance (VSH)
- Riverbed Whitewater cloud storage appliance Virtual Edition (CSH)
- Riverbed Command-Line Interface (CLI)
- Riverbed Granite Core appliance
- Riverbed Steelhead EX + Granite appliance
- Riverbed Granite Edge appliance
- Riverbed Central Management Console (CMC)
- Riverbed Steelhead Mobile software (Steelhead Mobile)

- Riverbed Steelhead Mobile Client (Mobile Client)
- Riverbed Interceptor appliance (Interceptor appliance)
- Riverbed Virtual Services Platform (VSP)
- Riverbed Services Platform (RSP)
- Riverbed Cascade Product Suite
- Riverbed Cascade Profiler appliance
- Riverbed SkipWare software (SkipWare software)

## Audience

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

This guide is also for users who are using the Riverbed Command-Line Interface as described in the *Riverbed Command-Line Interface Reference Manual*.

## Document Conventions

This guide uses this standard set of typographical conventions.

Convention	Meaning
<i>italics</i>	Within text, new terms and emphasized words appear in italic typeface.
<b>boldface</b>	Within text, CLI commands and GUI controls appear in bold typeface.
Courier	Code examples appear in Courier font: <pre>amnesiac &gt; enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: <b>interface &lt;ipaddress&gt;</b>
[ ]	Optional keywords or variables appear in brackets: <b>ntp peer &lt;addr&gt; [version &lt;number&gt;]</b>
{ }	Required keywords or variables appear in braces: <b>{delete &lt;filename&gt;   upload &lt;filename&gt;}</b>
	The pipe symbol represents a choice between the keyword or variable to the left or right of the symbol (the keyword or variable can be either optional or required): <b>{delete &lt;filename&gt;   upload &lt;filename&gt;}</b>

---

## Product Dependencies and Compatibility

This section provides information about product dependencies and compatibility. It includes this information:

- [“Hardware and Software Dependencies” on page 3](#)
- [“Ethernet Network Compatibility” on page 3](#)
- [“SNMP-Based Management Compatibility” on page 4](#)

### Hardware and Software Dependencies

This 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	<p>Any computer that supports a Web browser with a color image display.</p> <p>The Management Console has been tested with Mozilla Firefox Extended Support Release version 17.0 and Microsoft Internet Explorer v7.0 and v8.0.</p> <p>JavaScript and cookies must be enabled in your Web browser.</p> <p>Internet Explorer v7.0 and v8.0 must refresh reports every 4 minutes due to performance issues. Consider using a different browser to view reports.</p>

### Ethernet Network Compatibility

The Steelhead appliance supports these Ethernet networking standards:

- Ethernet Logical Link Control (LLC) (IEEE 802.2 - 1998)
- Fast Ethernet 100 BASE-TX (IEEE 802.3 - 2008)
- Gigabit Ethernet over Copper 1000 BASE-T and Fiber 1000 BASE-SX (LC connector) and Fiber 1000 BASE LX (IEEE 802.3 - 2008)
- Gigabit Ethernet over Fiber 10GBASE-LR Single Mode and 10GBASE-SR Multimode (IEEE 802.3 - 2008)

The Steelhead appliance ports support these 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 or 10GBASE-LR or 10GBASE-SR, depending on configuration
- **WAN** - 10/100/1000 BASE-TX or 1000 BASE-SX or 1000 BASE-LX or 10GBASE-LR or 10GBASE-SR, depending on configuration

The Steelhead appliance supports VLAN Tagging (IEEE 802.3 - 2008). It does not support the ISL protocol.

All copper interfaces are auto-sensing for speed and duplex (IEEE 802.3 - 2008).

The Steelhead appliance auto-negotiates speed and duplex mode for all data rates and supports full duplex mode and flow control (IEEE 802.3 – 2008).

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.

---

## Additional Resources

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

- [“Release Notes” on page 4](#)
- [“Riverbed Documentation and Support Knowledge Base” on page 5](#)

## Release Notes

An online file containing version release notes supplements the information in this user guide. 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 work-arounds. This file also provides documentation information not covered in the guides or that has been modified since publication.

Examine this file before you begin installation and configuration. 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, go to the Riverbed Support site 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 keywords and strings. To access the Riverbed Knowledge Base, log in to the Riverbed Support site at <https://support.riverbed.com>.

---

## Contacting Riverbed

This section describes how to contact departments within Riverbed.

### Internet

You can learn about Riverbed products at <http://www.riverbed.com>.

### Technical 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, open a trouble ticket by calling 1-888-RVBD-TAC (1-888-782-3822) in the United States and Canada or +1 415 247 7381 outside the United States. You can also go to <https://support.riverbed.com>.

### Professional Services

Riverbed has a staff of professionals who can help you with installation, provisioning, network redesign, project management, custom designs, consolidation project design, and custom coded solutions. To contact Riverbed Professional Services, email [proserve@riverbed.com](mailto:proserve@riverbed.com) or go to [http://www.riverbed.com/services-training/#Consulting\\_Services](http://www.riverbed.com/services-training/#Consulting_Services).

### Documentation

The Riverbed Technical Publications team continually strives to improve the quality and usability of Riverbed documentation. Riverbed appreciates any suggestions you might have about its online documentation or printed materials. Send documentation comments to [techpubs@riverbed.com](mailto:techpubs@riverbed.com).



## CHAPTER 1 Overview of the Management Console

This chapter provides an overview of the Management Console. The Management Console makes managing your Steelhead appliance simpler through a Web browser interface. It includes these sections:

- [“Using the Management Console” on page 7](#)
- [“Next Steps” on page 12](#)

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 *Steelhead Appliance Deployment Guide*.

---

### Using the Management Console

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

- [“Connecting to the Management Console” on page 8](#)
- [“The Home Page” on page 9](#)
- [“Navigating in the Management Console” on page 10](#)
- [“Getting Help” on page 11](#)

---

**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, the system prompts you to inspect and verify the SSL certificate. This is a self-signed certificate that provides encrypted Web connections to the Management Console. The system recreates the certificate when you change the appliance hostname or when the certificate expires.

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.

2. Enter a username and password that you assigned in the configuration wizard of the Steelhead appliance.

Default username: admin

Default password: password

You may also enter monitor, shark, a login from a RADIUS or TACACS+ database, or a login from any local account created using the Role-Based Accounts feature. For details about role-based accounts, see [“Role-Based Accounts” on page 197](#).

Users with administrator (admin) privileges may configure and administer the Steelhead appliance. Users with monitor (monitor) privileges may view the Steelhead appliance reports, user logs, and change their own password. A monitor user cannot make configuration changes. Users with Cascade Shark (shark) privileges may use the Embedded Cascade Shark function for detailed packet analysis through Cascade Pilot.

3. 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 these reports:

- **Optimized LAN Throughput Over Last Week** - Summarizes the throughput or total optimized data transmitted for all applications in the last week. Includes the LAN peak performance in Mbps, the 95th Percentile WAN throughput increase calculated as: 95p LAN / 95p WAN. The home page also includes the LAN average.
- **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 P.M. to 3 P.M. 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.

The system saves settings on a per-user basis. A message appears at the top of each page when multiple users are logged in explaining that user preferences might be overwritten.

## 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, 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 select Bandwidth Optimization to display the page.

This table summarizes the cascading menus.

Menu	Submenus
Home	Displays the Home page.
Configure	<b>Networking</b> - Configure host settings, base interfaces, asymmetric routing, connection forwarding, flow statistics, QoS, path selection, simplified routing, port labels, WCCP, and subnet side rules from this menu. <b>Optimization</b> - Configure optimization features such as in-path rules, protocols, high-speed TCP, peering rules, service ports, SnapMirror, FCIP, and SRDF from this menu. <b>System Settings</b> - Configure alarm settings, announcements, email settings, log settings, monitored ports, and SNMP settings from this menu. <b>Security</b> - Configure general security parameters, RADIUS, TACACS+, the secure vault, Management ACL, Web settings, and user permissions from this menu. <b>Maintenance</b> - Start and stop optimization services, schedule jobs, manage and update licenses, upgrade software, and reboot or shut down the appliance from this menu. <b>My Account</b> - Change your login password and view user permissions from this menu. <b>Configurations</b> - Manage, import, delete, and change your configuration files for the appliance from this menu.
Reports	<b>Networking</b> - 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. <b>Optimization</b> - View optimization reports such as optimized throughput, bandwidth optimization, peers, SRDF reports, and SnapMirror reports from this menu. <b>Diagnostics</b> - Display and download diagnostic reports such as CPU utilization, user and system logs, alarms status, system snapshots, system dumps, TCP dumps, and network health from this menu. <b>Export</b> - Export raw statistics from this menu.
Support	Display online help and appliance documentation; contact information for Riverbed Support; appliance details such as model number, revision type, serial number, software version; and appliance MIB files from this menu.

## Saving Your Configuration

After you change settings, you can apply those changes to the running configuration and you can save your configuration 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.

### To apply changes to the running configuration

- Click **Apply**. The Management Console updates the running configuration; however, your changes are not written to disk.

### To save your changes to disk

- Click **Save**.

## Restarting the Optimization Service

The optimization service is a daemon that executes in the background, performing required operations. Some configuration settings apply to the optimization service. When you change settings for features that depend on the optimization service, you must restart the service for the changes to take effect.

For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

## Logging Out

Ensure that all desired configuration changes are either applied or saved prior to logging out.

### To log out of the current session

- Click **Logout** in the upper-right corner of the screen

## 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 these options:

- **Online Help** - View browser-based online help.
- **Technical Support** - View links and contact information for Riverbed Support.
- **Appliance Details** - View appliance information such as model number, hardware revision type, and serial number currently installed on the appliance.
- **Software Version(s)** - View the RiOS version and build number.
- **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 separate browser window.

## Downloading Documentation

The Riverbed Support site contains PDF versions of the *Steelhead Appliance Management Console User's Guides* 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. Click the link next to Documentation: <https://support.riverbed.com/docs/index.htm>  
The Support site appears.
3. Select the product name.
4. Select the product version from the Documentation For Version drop-down list.
5. Select PDF or HTML next to the document name to download the document.

---

## Next Steps

This table describes a basic approach to configuring the Steelhead appliance.

Task	Reference
1. Become familiar with basic and advanced deployment types.	<i>Steelhead Appliance Deployment Guide</i>
2. Make decisions about where to deploy Steelhead appliances, and what features to use.	Riverbed Professional Services
3. Install appliances and optional interface cards.	<i>Steelhead Appliance Installation and Configuration Guide</i> <i>Network Interface Card Installation Guide</i>
4. Configure optimization traffic with in-path rules.	<a href="#">"Configuring In-Path Rules" on page 36</a>
5. Distribute administrative responsibility by configuring secure access for other administrators, monitor users, or other types of users you choose to create.	<a href="#">"Configuring General Security Settings" on page 196</a> (if applicable)
6. Modify default system administration settings.	<a href="#">"Configuring Alarm Settings" on page 217</a> (if desired)
7. Modify host and network settings you initially set with the installation wizard.	<a href="#">"Modifying General Host Settings" on page 15</a> (if desired)

Task	Reference
8. Save your configuration changes and restart services as necessary.	<a href="#">“Starting and Stopping the Optimization Service” on page 187</a> (as necessary) <a href="#">“Managing Configuration Files” on page 194</a> (as necessary)
9. View logs and reports to verify your deployment.	<a href="#">“Viewing Current Connection Reports” on page 248</a>
10. Troubleshoot (if necessary).	<i>Steelhead Appliance Deployment Guide</i> Riverbed Support





## CHAPTER 2    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 these sections:

- [“Modifying General Host Settings” on page 15](#)
- [“Modifying Base Interfaces” on page 18](#)
- [“Modifying In-Path Interfaces” on page 25](#)

---

### Modifying General Host Settings

You view and modify general host settings in the Configure > Networking > Host Settings page. Use the controls on this page only if you require modifications, additional configuration, or want to verify the DNS configuration:

- **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.
- **Web/FTP Proxy** - Configure proxy addresses for Web or FTP proxy access to the Steelhead appliance.

#### 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 this 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 this 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 <b>Remove Selected</b> .

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 **Web/FTP Proxy**, complete the configuration as described in this table.

Control	Description
Enable Web Proxy	<p>Provides Web proxy access to the Steelhead appliance. Enables the Steelhead appliance to use a Web proxy to contact the Riverbed licensing portal and fetch licenses in a secure environment. You can optionally require user credentials to communicate with the proxy, and you can specify the method used to authenticate and negotiate user credentials.</p> <p>Web proxy access is disabled by default.</p> <p>RiOS supports these proxies: Squid, Blue Coat Proxy SG, Microsoft WebSense, and McAfee Web Gateway.</p>
Web/FTP Proxy	Specify the IP address for the Web or FTP proxy.
Port	Optionally, specify the port for the Web or FTP proxy. The default port is 1080.
Enable Authentication	<p>Optionally, select to require user credentials for use with Web or FTP proxy traffic. Specify the following to authenticate the users:</p> <ul style="list-style-type: none"> <li>• <b>User name</b> - Specify a user name.</li> <li>• <b>Password</b> - Specify a password.</li> <li>• <b>Authentication Type</b> - Select an authentication method from the drop-down list: <ul style="list-style-type: none"> <li>– <b>Basic</b> - Authenticates user credentials by requesting a valid user name and password. This is the default setting.</li> <li>– <b>NTLM</b> - Authenticates user credentials based on an authentication challenge and response.</li> <li>– <b>Digest</b> - Provides the same functionality as basic authentication; however, digest authentication improves security because the system sends the user credentials across the network as a Message Digest 5 (MD5) hash.</li> </ul> </li> </ul>

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

An abbreviated form of the timestamp appears in the left margin of each line. To see the original, full timestamp in the form of a tooltip, hover the mouse over a timestamp. Not all log lines have timestamps, because third-party applications generate some of the logging data.

The log lines highlight errors in red and warnings in yellow.

---

## Modifying Base Interfaces

You 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 controls on this page if you require modifications or additional configuration:

- **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.

## IPv6 Support

RiOS supports auto-discovery and fixed-target rules. By using auto-discovery or fixed-target in-path rules, RiOS can apply transport and application streamlining techniques (similarly as it does for TCP connections over IPv4) to improve the user experience as the transition to IPv6 continues.

IPv6 is enabled by default in RiOS v8.5. The Steelhead appliance support for IPv6 is twofold:

- **Managing Steelhead appliances** - Support for management access using IPv6 IP addresses on primary and auxiliary interfaces.
- **Optimizing IPv6 traffic using Steelhead appliances** - Steelhead appliances can optimize IPv6 traffic.

For details on IPv6 deployments, see the *Steelhead Appliance Deployment Guide*. For details on in-path rules, see [“Configuring In-Path Rules” on page 36](#).

This table lists IPv6 support by feature, and notes any limits and special considerations.

RiOS IPv6 Support Includes	Notes
Conformance with Request for Comments (RFCs) 1981, 2460, 2464, 2710, 3590, 4007, 4291, 4443, 4861, 4862, 4943, 5095, and 5156	
TCP IPv6 traffic interception between source and destination, bandwidth optimization	
Auto-discovery of Steelhead appliances	TCP inner connections between the peer Steelhead appliances is strictly IPv4.
Ability to automatically discover fixed-target and pass-through in-path rules, along with ability to deny and reject IPv6 TCP traffic as configured in the in-path rules	<p>RiOS does not support the Outlook Anywhere and Citrix latency optimization policies for auto-discovery and fixed-target rules.</p> <p>RiOS does not support the neural framing modes Always, TCP Hints, and Dynamic.</p> <p>RiOS does not support the Oracle forms and Oracle forms over SSL preoptimization policies.</p>
HTTP and HTTPS latency optimization for IPv6 TCP traffic	
Ability to configure serial clusters	
Interception of IPv6 traffic for in-path, virtual in-path, and server-side out-of-path configurations	<p>WCCPv6 support is not available. Virtual in-path support is PBR only.</p> <p>Interceptor is not supported.</p>
Intercepting and passing through IPv4 and/or IPv6 traffic depending on the in-path rules	
Ability to detect asymmetric routes for IPv6 TCP traffic; enables connection forwarding of IPv6 TCP traffic in asymmetric conditions	The connection forwarding control channel between the neighbors is strictly IPv4. This requires configuring IPv4 addresses on the Steelhead appliances.
Ability to configure IPv4 and IPv6 addresses on every in-path interface and intercepting and optimizing IPv4 and IPv6 traffic.	
<p>Ability to configure one IPv6 address configuration for every in-path interface.</p> <p>RiOS intercepts and optimizes traffic matching the scope of the IPv6 address configured on the in-path interface. Not applicable for a link-local address configured on the in-path interface.</p>	RiOS passes through IPv6 TCP traffic not matching the scope of the IPv6 address configured on the in-path interface.
<p>Ability to configure IPv6 addresses on any in-path interface.</p> <p>This IPv6-only mode requires configuring only fixed-target in-path rules.</p> <p>IPv6 TCP inner connections only in fixed-target cases.</p>	
Enhanced auto-discovery of Steelhead appliances for IPv6 TCP traffic.	TCP inner connections between the peer Steelhead appliances is IPv4 only.
Simplified routing for IPv6 TCP traffic.	

RiOS IPv6 Support Includes	Notes
Connection forwarding for IPv6 traffic in multi-interface mode.	<p>The control connection between neighbors is still IPv4 only.</p> <p>When multiple interface support on the Configure &gt; Networking &gt; Connection Forwarding page is not enabled, IPv6 traffic is passed through.</p>
Ability to configure peering rules for IPv6 traffic.	The peer client-side Steelhead appliance IP address is IPv4 only.
Ability to configure IPv6 addresses in Single Ended Interception (SEI) rules under Configure > Networking > Transport Settings.	
Global and automatic kickoff for pass-through TCP IPv6 traffic.	
Ability to configure asymmetric VLANs for IPv6 TCP traffic.	

## Features Not Supported with IPv6

The following features are not IPv6 compatible:

- Management In-Path (MIP) Interface
- Transparency
- NetFlow
- RSP
- Path Selection
- QoS
- Host labels
- IPSec
- Automatic address assignment through DHCPv6
- Multicast listener discovery
- IPv6 stateless address auto-configuration
- WCCP using anything other than IPv4 outer connections
- Connection-forwarding neighbor connection using anything other than IPv4
- ICMPv6 redirect messages

### To display and modify the configuration for base interfaces

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

2. Under Primary Interface, complete the configuration as described in this 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	<p>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.</p> <p><b>Important:</b> The primary and in-path interfaces can share the same network subnet. The primary and auxiliary interfaces cannot share the same network subnet.</p>
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	<p>Select this option if you do not use a DHCP server to set the IPv4 address. Specify these settings:</p> <ul style="list-style-type: none"> <li>• <b>IPv4 Address</b> - Specify an IP address.</li> <li>• <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> <li>• <b>Default IPv4 Gateway</b> - 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.</li> </ul>
Specify IPv6 Address Manually	<p>Select this option and specify these settings to set an IPv6 address.</p> <ul style="list-style-type: none"> <li>• <b>IPv6 Auto-Assigned</b> - Displays the link-local address that is automatically generated when IPv6 is enabled on the base interfaces.</li> <li>• <b>IPv6 Address</b> - Specify an IP address using this 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</li> <li>• <b>IPv6 Prefix</b> - 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</li> <li>• <b>IPv6 Gateway</b> - Specify the default gateway IP address. The default gateway must be in the same network as the primary interface.</li> </ul> <p><b>Note:</b> You cannot set an IPv6 address dynamically using a DHCP server.</p>

Control	Description
Speed and Duplex	<p><b>Speed</b> - Select a speed from the drop-down list. The default value is Auto.</p> <p><b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto.</p> <p>If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them manually.</p> <p>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.</p>
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.

3. Under Auxiliary Interface, complete the configuration as described in this table.

Control	Description
Enable Aux Interface	Enables an auxiliary interface, which can be used only 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	<p>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.</p> <p><b>Important:</b> The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.</p>
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	<p>Select this option if you do not use a DHCP server to set the IPv4 address. Specify these settings:</p> <ul style="list-style-type: none"> <li>• <b>IPv4 Address</b> - Specify an IP address.</li> <li>• <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> </ul>
Specify IPv6 Address Manually	<p>Select this option and specify these settings to set an IPv6 address.</p> <ul style="list-style-type: none"> <li>• <b>IPv6 Auto-Assigned</b> - Displays the link-local address that is automatically generated when IPv6 is enabled on the base interfaces.</li> <li>• <b>IPv6 Address</b> - Specify an IP address, using this 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</li> <li>• <b>IPv6 Prefix</b> - 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</li> </ul> <p><b>Note:</b> You cannot set an IPv6 address dynamically using a DHCP server.</p>



Control	Description
Speed and Duplex	<p><b>Speed</b> - Select the speed from the drop-down list. The default value is Auto.</p> <p><b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto.</p> <p>If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them on the device manually.</p> <p>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.</p>
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.

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

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

### To configure routes for IPv4

- Under 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 this 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.
Interface	Select an interface for the IPv4 route from the drop-down menu.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

The Management Console writes your configuration changes to memory.

### To configure routes for IPv6

- Under 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 this 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 from 0 to 128 bits, separated from the address by a forward slash (/).

Control	Description
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.
Interface	Select an interface for the IPv6 route from the drop-down menu.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

The Management Console writes your configuration changes to memory.

---

## Modifying In-Path Interfaces

You 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.
2. To enable link state propagation, under In-Path Settings, complete the configuration as described in this table.

Control	Description
Enable Link State Propagation	<p>Enable to shorten the recovery time of a link failure in physical in-path deployments. Link state propagation (LSP) communicates link status between the devices connected to the Steelhead appliance. When you enable this feature, RiOS monitors the link state of each Steelhead appliance LAN-WAN pair.</p> <p>If either physical port loses link status, the corresponding interface disconnects, blocking the link. This allows a link failure to quickly propagate through a chain of devices. If the link recovers, the Steelhead appliance restores the corresponding interface automatically.</p> <p>LSP is enabled by default.</p> <p><b>Note:</b> You cannot reach an MIP interface when LSP is also enabled and the corresponding in-path interface fails.</p> <p>These Virtual Steelhead appliance configurations do not support LSP:</p> <ul style="list-style-type: none"><li>• VSH models running ESX/ESXi 4.0 or 4.1</li><li>• VSH models running Microsoft Hyper-V</li></ul>

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

Control	Description
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	<p>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.</p> <p><b>Important:</b> 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.</p>
Enable IPv6	<p>Select this check box to assign an IPv6 address. IPv6 addresses are disabled by default. You can only assign one IPv6 address per in-path interface.</p> <p><b>Important:</b> The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.</p>
IPv6 Address	Specify a global or site-local IPv6 address. This IP address is the in-path main interface. You cannot use a DHCP server to assign an IPv6 address automatically.
IPv6 Prefix	<p>Specify the prefix. The prefix length is 0 to 128 bits, separated from the address by a forward slash (/). In the following example, 60 is the prefix:</p> <p><code>2001:38dc:52::e9a4:c5:6282/60</code></p>
IPv6 Gateway	<p>Specify the IPv6 address for the in-path gateway. You can use a link local address. 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.</p> <p><b>Important:</b> 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.</p>

Control	Description
LAN Speed and Duplex WAN Speed and Duplex	<p><b>Speed</b> - Select Auto, 1000, 100, or 10 from the drop-down list. The default value is Auto.</p> <p><b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto.</p> <p>If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them on the device manually.</p> <p>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.</p> <p><b>Note:</b> 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.</p> <p>These guidelines can help you avoid speed and duplex mismatches when configuring the Steelhead appliance:</p> <ul style="list-style-type: none"> <li>• 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 that your switch has the correct setting.</li> <li>• 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.</li> <li>• 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.</li> </ul>
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	<p>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 nontagged (or native VLAN) and is the correct setting if there are no VLANs present.</p> <p>As an example, if the in-path interface is 192.168.1.1 in VLAN 200, you would specify tag 200.</p> <p><b>Note:</b> 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.</p> <p>You must also define in-path rules to apply to your VLANs.</p>

- Under IPv4 Routing Table, 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 <b>Remove Selected</b> .

- Under IPv6 Routing Table, 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.
IPv6 Prefix	Specify the prefix. The prefix length is 0 to 128 bits, separated from the address by a forward slash (/). In the following example, 60 is the prefix: <code>2001:38dc:52::e9a4:c5:6282/60</code>
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 <b>Remove Selected</b> .

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

## Configuring a Management In-Path Interface

You 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 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.

An 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.

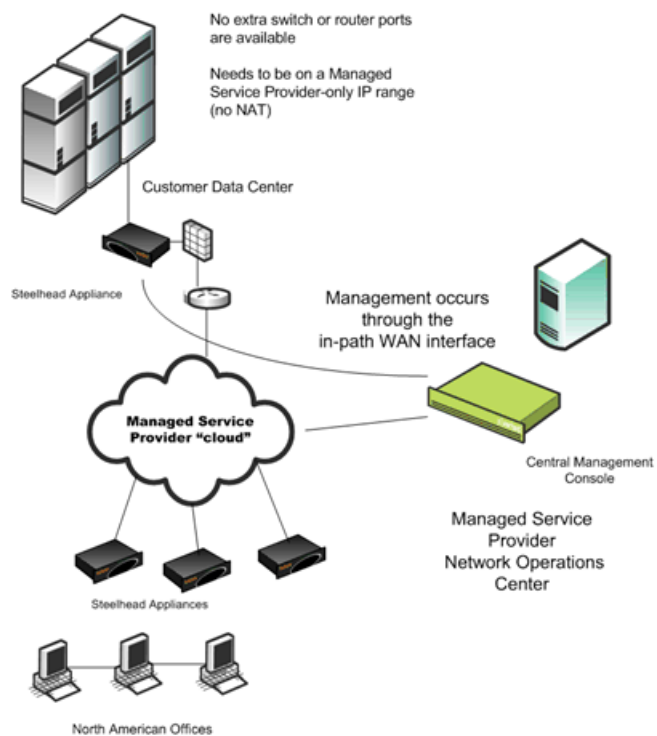
An MIP interface is not accessible if the (physical) LAN and WAN interfaces fail.

---

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

---

**Figure 2-1. Management In-Path Interface Deployment**



## MIP Interface Dependencies

An MIP interface has these dependencies:

- Any connections destined to an MIP interface are not optimized by that Steelhead appliance and do not appear in the Current Connections report.
- An MIP interface cannot reside in the same subnet as the primary or auxiliary interfaces. It cannot share the same subnet with any other interfaces on the Steelhead appliance.
- An MIP interface must be in its own subnet.
- You cannot enable an 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 rerouted onto other paths (where the remaining Steelhead appliances are deployed). For details on fail-to-block, see the *Steelhead Appliance Deployment Guide*.
- You cannot reach an 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. To disable LSP, enter the **no in-path lsp enable** CLI command at the system prompt.
- This feature supports 802.1Q VLAN.
- An MIP interface uses the main routing table.

## Enabling an MIP Interface

Use the controls on this page when you need to enable an 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 this 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.  <b>Note:</b> 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.



Control	Description
IPv4 Subnet Mask	Specify the subnet mask.
VLAN Tag ID	<p>Specifies a numeric VLAN Tag ID.</p> <p>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 nontagged (or native VLAN) and is the correct setting if there are no VLANs present.</p> <p>For example, if the MIP interface is 192.168.1.1 in VLAN 200, you would specify tag 200.</p>

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**

- [“Modifying Base Interfaces” on page 18](#)
- [“Configuring In-Path Rules” on page 36](#)
- [“Configuring General Service Settings” on page 47](#)



## CHAPTER 3      **Configuring In-Path Rules**

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

- [“In-Path Rules Overview” on page 33](#)
- [“Default In-Path Rules” on page 35](#)
- [“Configuring In-Path Rules” on page 36](#)

---

### **In-Path Rules Overview**

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 these 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.

RiOS includes fixed-target, packet-mode optimization in-path rules. RiOS packet-mode optimization includes TCPv4 and UDPv6 traffic. In addition, RiOS enhances connection or flow reporting for packet-mode optimization. To optimize TCPv4 or UDPv6, the client-side and server-side Steelhead appliances must run RiOS v8.5 and later.

---

**Note:** The Steelhead appliance treats the packets for packet-mode optimization rules differently from the in-path rules described in this overview. For details, see [“Creating In-Path Rules for Packet-Mode Optimization” on page 34](#).

---

For details on IPv6 deployment options, see the *Steelhead Appliance Deployment Guide*.

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.
- **Neural Framing Requirements** - Specify never, always, TCP Hints, or Dynamic.
- **WAN Visibility** - Preserve TCP/IP address or port information.


## Creating In-Path Rules for Packet-Mode Optimization

RiOS performs packet-by-packet SDR bandwidth optimization on TCP and UDP flows over both IPv4 and IPv6, using fixed-target, packet-mode optimization in-path rules. This type of in-path rule optimizes bandwidth for applications over any transport protocol.

Sometimes you might want to use the Steelhead appliance optimization to reduce the amount of traffic traversing the WAN. Packet-mode optimization provides a simple approach where the Steelhead appliance looks at a packet, or small group of packets, and performs SDR and LZ on the data payload for data reduction. The host and Steelhead appliance do not create an individual TCP handshake, and the Steelhead appliance reduces payload for packets as the traffic flows through.

The advantage of packet-mode optimization is that it is a universal method that applies data streamlining to diverse protocols. The disadvantage is the lack of performance benefits from transport streamlining or application streamlining, because the Steelhead appliance does not proxy or perform intelligent application prediction.

**Figure 3-1. A Fixed-Target Packet-Mode Optimization Rule Creates an Inner TCPv4 or TCP v6 Channel Between the Steelhead Appliances**



Consider using the typical terminated TCP optimization to improve application latency instead of packet-mode for TCPv4 or TCPv6 traffic. RiOS includes TCP proxy-mode optimization for IPv6 traffic.

## Packet-Mode Optimization Rule Characteristics

When you create a fixed-target packet-mode optimization rule, you define the inner channel characteristics using these controls: source and destination subnet, source destination port or port labels, and DSCP marking.

Packet-mode optimization supports these topologies:

- Physical in-path
- Virtual in-path
  - WCCP/PBR or TCPv4, UDPv4
  - PBR for TCPv6, UDPv6
- Master and backup

Packet-mode optimization does not support these topologies:

- Out-of-path
- Serial cluster
- Interceptor integration

For details, see [“Configuring In-Path Rules” on page 36](#). For design considerations and best practices, see the *Steelhead Appliance Deployment Guide*.

---

## 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, 179, 513, 514, 1494, 2598, 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 (RiOS data store 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 about changing port labels, see [“Configuring Port Labels” on page 75](#).

---

# Configuring In-Path Rules

You 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 [“In-Path Rules Overview” on page 33](#).

For details on IPv6 deployment options, see the *Steelhead Appliance Deployment Guide*.

## To configure in-path rules

1. Choose Configure > Optimization > In-Path Rules to display the In-Path Rules page.
2. Configure the rules as described in this table.

Control	Description
Add a New In-Path Rule	Displays the controls for adding a new rule.

(1 of 10)

Control	Description
Type	<p>Select one of these rule types from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Auto-Discover</b> - 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.</li> <li>• <b>Fixed-Target</b> - Skips the auto-discovery process and uses a specified remote Steelhead appliance as an optimization peer.</li> </ul> <p>You must specify at least one remote target Steelhead appliance (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.</p> <p>In RiOS, a fixed-target rule enables you to optimize traffic end-to-end using IPv6 addresses. You must change the use of all-IP to all-IPv6.</p> <p>If you do not change to all-IPv6, use specific source and destination IPv6 addresses. The inner channel between Steelhead appliances forms a TCP connection using the manually assigned IPv6 address. This method is similar to an IPv4 fixed-target rule and you configure it the same way.</p> <ul style="list-style-type: none"> <li>• <b>Fixed-Target (Packet Mode Optimization)</b> - Skips the auto-discovery process and uses a specified remote Steelhead appliance as an optimization peer to perform bandwidth optimization on TCPv4, TCPv6, UDPv4, or UDPv6 connections.</li> </ul> <p>Packet-mode optimization rules support both physical in-path and master/backup Steelhead configurations.</p> <p>You must specify which TCP or UDP traffic flows need optimization, at least one remote target Steelhead appliance, and, optionally, which ports and backup Steelhead appliances to use.</p> <p>In addition to adding fixed-target packet-mode optimization rules, you must go to Configure &gt; Optimization &gt; General Service Settings, enable packet-mode optimization, and restart the optimization service.</p> <p>Packet-mode optimization rules are unidirectional; a rule on the client-side Steelhead optimizes traffic to the server only. To optimize bidirectional traffic, define two rules:</p> <ul style="list-style-type: none"> <li>• A fixed-target packet-mode optimization rule on the client-side Steelhead to the server.</li> <li>• A fixed-target packet-mode optimization rule on the server-side Steelhead to the client.</li> </ul> <p>Packet-mode optimization rules perform packet-by-packet optimization, as opposed to traffic-flow optimization. After you create the in-path rule to intercept the connection, the traffic flows enter the Steelhead. The Steelhead does not terminate the connection, but instead rearranges the packet headers and payload for SDR optimization. Next, it provides SDR optimization and sends the packets through a TCPv4 or TCPv6 channel to the peer Steelhead appliance. The peer Steelhead appliance decodes the packet and routes it to the destined server. The optimized packets are sent through a dedicated channel to the peer, depending on which in-path rule the packet's flow was matched against.</p> <p>To view packet-mode optimized traffic, choose Reports &gt; Networking &gt; Current Connections or Connection History. You can also enter the <b>show flows</b> CLI command at the system prompt.</p>

(2 of 10)

Control	Description
	<p><b>Requirements:</b></p> <ul style="list-style-type: none"> <li>• Both the client-side Steelhead appliance and the server-side Steelhead appliance must be running RiOS.</li> <li>• To view the packet-mode flows in the Current Connections and Connection History reports, the Steelhead appliance must be running RiOS v8.5.</li> </ul> <p><b>Packet-mode optimization rules do not support:</b></p> <ul style="list-style-type: none"> <li>• Automatic reflection of DSCP markings.</li> <li>• Latency optimization and preoptimization policies. Selecting this rule type automatically sets the preoptimization policy and latency optimization policies to none.</li> <li>• Auto-discovery of the peer Steelhead. Because this is a fixed-target rule, the Steelhead determines the IP address of its peer from the rule configuration.</li> <li>• Connection forwarding, simplified routing, or asymmetric routing.</li> <li>• QoS, MIP interfaces, NetFlow, transparency, or the automatic kickoff feature.</li> <li>• Automatically assigned IPv6 addresses.</li> </ul> <hr/> <ul style="list-style-type: none"> <li>• <b>Pass-Through</b> - 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 optimization service was enabled.)</li> <li>• <b>Discard</b> - 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.</li> <li>• <b>Deny</b> - 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.</li> </ul>

(3 of 10)



Control	Description
Source Subnet	<p>Specify the subnet IP address and netmask for the source network:</p> <ul style="list-style-type: none"> <li>• <b>All IPv4</b> - Maps to 0.0.0.0/0.</li> <li>• <b>All IPv6</b> - Maps to ::/0.</li> <li>• <b>All IP</b> - Maps to all IPv4 and IPv6 networks.</li> </ul> <p><b>Important:</b> In a virtual in-path configuration using packet-mode optimization, do not use the wildcard All IP option for both the source and destination IP addresses on the server-side and client-side Steelhead appliances. Doing so can create a loop between the Steelhead appliances if the server-side Steelhead appliance forms an inner connection with the client-side Steelhead appliance before the client-side Steelhead appliance forms an inner connection with the server-side Steelhead appliance. Instead, configure the rule using the local subnet on the LAN side of the Steelhead appliance.</p> <p>Use this format for an individual subnet IP address and netmask:            XXX.XXX.XXX.XXX/XX (IPv4)            X:X:X::X/XXX (IPv6)</p> <p>When creating a fixed-target packet-mode rule, you must configure an IPv6 address and route for each interface, unless you are optimizing UDP traffic.</p> <p><b>Port or Port Label</b> - Specify the destination port number, port label, or All. Click Port Label to go to the Configure &gt; Networking &gt; Port Labels page for reference.</p>
Destination Subnet	<p>Specify the subnet IP address and netmask for the destination network:</p> <ul style="list-style-type: none"> <li>• <b>All IPv4</b> - Maps to 0.0.0.0/0.</li> <li>• <b>All IPv6</b> - Maps to ::/0.</li> <li>• <b>All IP</b> - Maps to all IPv4 and IPv6 networks.</li> </ul> <p><b>Important:</b> In a virtual in-path configuration using packet-mode optimization, do not use the wildcard All IP option for both the source and destination IP addresses on the server-side and client-side Steelhead appliances. Doing so can create a loop between the Steelhead appliances if the server-side Steelhead appliance forms an inner connection with the client-side Steelhead appliance before the client-side Steelhead appliance forms an inner connection with the server-side Steelhead appliance. Instead, configure the rule using the local subnet on the LAN side of the Steelhead appliance.</p> <p>Use this format for an individual subnet IP address and netmask:            XXX.XXX.XXX.XXX/XX (IPv4)            X:X:X::X/XXX (IPv6)</p> <p>When creating a fixed-target packet mode optimization rule, you must configure an IPv6 address and route for each interface.</p> <p><b>Port or Port Label</b> - Specify the destination port number, port label, or All. Click Port Label to go to the Configure &gt; Networking &gt; Port Labels page for reference.</p>
Target Appliance IP Address	<p>Specify the target appliance address for a fixed-target rule. When the protocol is TCP and you do not specify an IP address, the rule defaults to all IPv6 addresses.</p> <p><b>Port</b> - Specify the target port number for a fixed-target rule.</p>
Backup Appliance IP Address	<p>Specify the backup appliance address for a fixed-target rule.</p> <p><b>Port</b> - Specify the backup destination port number for a fixed-target rule.</p>

(4 of 10)

Control	Description
VLAN Tag ID	<p>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 nontagged connections.</p> <p>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 preexisting VLAN tagging between the LAN and WAN interfaces.</p>
Protocol	<p>(Appears only for fixed-target packet-mode optimization rules.) Select a traffic protocol from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>TCP</b> - Specifies the TCP protocol. Supports TCP-over-IPv6 only.</li> <li>• <b>UDP</b> - Specifies the UDP protocol. Supports UDP-over-IPv4 only.</li> <li>• <b>Any</b> - Specifies all TCP- and UDP-based protocols. This is the default setting.</li> </ul>
Data Reduction Policy	<p>Optionally, if the rule type is Auto-Discover or Fixed Target, you can configure these types of data reduction policies:</p> <ul style="list-style-type: none"> <li>• <b>Normal</b> - Perform LZ compression and SDR.</li> <li>• <b>SDR-Only</b> - Perform SDR; do not perform LZ compression.</li> <li>• <b>Compression-Only</b> - Perform LZ compression; do not perform SDR.</li> <li>• <b>None</b> - Do not perform SDR or LZ compression.</li> </ul> <p>To configure data reduction policies for the FTP data channel, define an in-path rule with the destination port 20 and set its data reduction policy. Setting QoS for port 20 on the client-side Steelhead appliance affects passive FTP, while setting the QoS for port 20 on the server-side Steelhead appliance affects active FTP.</p>

(5 of 10)

Control	Description
Auto Kickoff	<p>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 reestablished and optimized.</p> <p>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.</p> <p>RiOS provides three ways to enable kickoff:</p> <ul style="list-style-type: none"> <li>• Globally for all existing connections in the Configure &gt; Optimization &gt; General Service Settings page.</li> <li>• For a single pass-through or optimized connection in the Current Connections report, one connection at a time.</li> <li>• For all existing connections that match an in-path rule and the rule has kickoff enabled.</li> </ul> <p>In most deployments, you do not want to set automatic kickoff globally because it disrupts <i>all</i> 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.</p> <p><b>Note:</b> 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 might take a while for the connection to reset.</p> <p>By default, automatic kickoff per in-path rule is disabled.</p> <p>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.</p> <p>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.</p> <p>As an 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.</p> <p>Src 10.11.10.10/24 Dst 10.12.10.10/24 Auto Kickoff enabled</p> <p>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.</p> <p><b>Important:</b> 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 &gt; Optimization &gt; General Service Settings page.</p> <p><b>Note:</b> This feature pertains only to auto-discover and fixed-target rule types and is dimmed for the other rule types.</p>

(6 of 10)

Control	Description
Neural Framing Mode	<p>Optionally, if the rule type is 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.</p> <p>Select a neural framing setting:</p> <ul style="list-style-type: none"> <li>• <b>Never</b> - Do not use the Nagle algorithm. The Nagle algorithm is a means of improving the efficiency of TCP/IP networks by reducing the number of packets that need to be sent over the network. It works by combining a number of small outgoing messages and sending them all at once. 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.</li> <li>• <b>Always</b> - 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. This mode is not compatible with IPv6.</li> <li>• <b>TCP Hints</b> - 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. This mode is not compatible with IPv6.</li> <li>• <b>Dynamic</b> - 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. This mode is not compatible with IPv6.</li> </ul> <p>For different types of traffic, one algorithm might be better than others. The considerations include: latency added to the connection, compression, and SDR performance.</p> <p>To configure neural framing for an FTP data channel, define an in-path rule with the destination port 20 and set its data reduction policy. To configure neural framing for a MAPI data channel, define an in-path rule with the destination port 7830 and set its data reduction policy.</p>

(7 of 10)

Control	Description
WAN Visibility Mode	<p>Enables WAN visibility, which pertains to how packets traversing the WAN are addressed. RiOS provides three types of WAN visibility: correct addressing, port transparency, and full address transparency.</p> <p>You configure WAN visibility on the client-side Steelhead appliance (where the connection is initiated).</p> <p>Port, full transparency, and full transparency modes are not compatible with IPv6.</p> <p>Select one of these modes from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Correct Addressing</b> - 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.</li> <li>• <b>Port Transparency</b> - 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.</li> </ul> <p>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.</p> <p>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.</p> <p>Port transparency does not require dedicated port configurations on your Steelhead appliances.</p> <p><b>Note:</b> Port transparency only provides server port visibility. It does not provide client and server IP address visibility, nor does it provide client port visibility.</p> <ul style="list-style-type: none"> <li>• <b>Full Transparency</b> - 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.</li> </ul> <p>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 <i>Steelhead Appliance Deployment Guide</i>.</p> <p>However, if you must see your client or server IP addresses across the WAN, full transparency is your only configuration option.</p> <p><b>Important:</b> 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 <i>Steelhead Appliance Deployment Guide</i>.</p>

(8 of 10)

Control	Description
WAN Visibility Mode ( <i>continued</i> )	<p>RiOS 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:</p> <ul style="list-style-type: none"> <li>• <b>Full Transparency with Reset</b> - 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.</li> </ul> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• For details on configuring WAN visibility and its implications, see the <i>Steelhead Appliance Deployment Guide</i>.</li> <li>• 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.</li> <li>• 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.</li> <li>• 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.</li> <li>• 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 <i>Steelhead Appliance Deployment Guide</i>.</li> <li>• 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.</li> </ul>
Position	<p>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.</p> <p>In general, list rules in this order:</p> <p>1. Deny 2. Discard 3. Pass-through 4. Fixed-target 5. Auto-Discover</p> <p>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.</p>
Description	Describe the rule to facilitate administration.
Enable Rule	Select to enable the in-path rule.
Add	Adds the rule to the list. The Management Console redisplay 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 <b>Remove Selected Rules</b> .

(9 of 10)

Control	Description
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

(10 of 10)

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.

In RiOS v8.5, the default rule maps to all IPv4 and IPv6 addresses (All-IP:\*).

The default rule for TCP traffic, either IPv4 or IPv6, attempts auto-discovery with correct addressing as the WAN visibility mode.

For details on IPv6 deployment options, see the *Steelhead Appliance Deployment Guide*.

### 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 194](#).

### Related Topics

- [“In-Path Rules Overview” on page 33](#)
- [“Default In-Path Rules” on page 35](#)
- [“Configuring General Service Settings” on page 47](#)
- [“Enabling Peering and Configuring Peering Rules” on page 53](#)
- [“Configuring Port Labels” on page 75](#)
- [“Secure Inner Channel Overview” on page 176](#)
- [“Viewing Current Connection Reports” on page 248](#)
- [“Viewing Connection History Reports” on page 268](#)





## CHAPTER 4      **Configuring Optimization Features**

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

- [“Configuring General Service Settings” on page 47](#)
- [“Enabling Peering and Configuring Peering Rules” on page 53](#)
- [“Improving Performance” on page 58](#)
- [“Configuring TCP, Satellite Optimization, and High-Speed TCP” on page 59](#)
- [“Configuring Service Ports” on page 72](#)
- [“Configuring Host Labels” on page 73](#)
- [“Configuring Port Labels” on page 75](#)
- [“Configuring FCIP Optimization” on page 77](#)
- [“Configuring SRDF Optimization” on page 80](#)
- [“Configuring SnapMirror Optimization” on page 85](#)

---

### **Configuring General Service Settings**

You 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, these 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 *Steelhead Appliance 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 about failover redundancy, see the *Steelhead Appliance 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. After 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 RiOS data stores 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.

## Configuring General Service Settings

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.
2. Under In-Path Settings, complete the configuration as described in this 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	<p>Enables <i>kickoff</i> globally. If you enable kickoff, connections that exist when the optimization service is started and restarted are disconnected. When the connections are retried they are optimized.</p> <p>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.</p> <p>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.</p> <p>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.</p>

Control	Description
Enable L4/PBR/WCCP Interceptor Support	<p>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. These redirection methods are available:</p> <ul style="list-style-type: none"> <li>• <b>Layer-4 Switch</b> - You enable Layer-4 switch support when you have multiple Steelhead appliances in your network, so that you can manage large bandwidth requirements.</li> <li>• <b>Policy-Based Routing (PBR)</b> - 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.</li> <li>• <b>Web Cache Communication Protocol (WCCP)</b> - 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.</li> </ul> <p>For details about configuring Layer-4 switch, PBR, and WCCP deployments, see the <i>Steelhead Appliance Deployment Guide</i>.</p>
Enable Optimizations on Interface <interface_name>	<p>Enables in-path support for additional bypass cards.</p> <p>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.</p> <p>The interface names for the bypass cards are a combination of the slot number and the port pairs (inpath&lt;slot&gt;_&lt;pair&gt;, inpath&lt;slot&gt;_&lt;pair&gt;): 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 <i>Network Interface Card Installation Guide</i>.</p>

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

Control	Description
Enable Out-of-Path Support	<p>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.</p> <p>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.</p>

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

Control	Description
Half-Open Connection Limit per Source IP	<p>Restricts half-opened connections on a source IP address initiating connections (that is, the client machine).</p> <p>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).</p> <p>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.</p> <p>The default value is 4096.</p> <p>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.</p> <p><b>Note:</b> 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.</p>
Maximum Connection Pool Size	<p>Specify the maximum number of TCP connections in a connection pool.</p> <p>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.</p> <p>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.</p> <p>The default value is 20. A value of 0 specifies no connection pool.</p> <p><b>Important:</b> You must restart the Steelhead appliance after changing this setting.</p> <p><b>Tip:</b> 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.</p>

5. Under Failover Settings, complete the configuration as described in this 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 warm RiOS data store, and can begin delivering fully optimized performance immediately.  The master and backup Steelhead appliances must be the same hardware model.
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.  <b>Important:</b> You must specify the inpath0_0 interface as the other appliance's in-path IP address.

6. Optionally, under Packet Mode Optimization Settings, complete the configuration as described in this table. For details about packet-mode optimization, see [“Creating In-Path Rules for Packet-Mode Optimization” on page 34](#).

Control	Description
Enable Packet Mode Optimization	Performs packet-by-packet SDR bandwidth optimization on TCP or UDP (over IPv4 or IPv6) flows. This feature uses fixed-target packet mode optimization in-path rules to optimize bandwidth for applications over these transport protocols.  By default, packet-mode optimization is disabled.  Enabling this feature requires an optimization service restart.

7. Click **Apply** to apply your settings.
8. 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 about saving configurations, see [“Managing Configuration Files” on page 194](#).

### Related Topics

- [“Configuring In-Path Rules” on page 36](#)
- [“Enabling Peering and Configuring Peering Rules” on page 53](#)
- [“Configuring Service Ports” on page 72](#)
- [“Modifying In-Path Interfaces” on page 25](#)
- [“Configuring Connection Forwarding Features” on page 94](#)
- [“Configuring Subnet Side Rules” on page 99](#)

## Enabling Peering and Configuring Peering Rules

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

- [“About Regular and Enhanced Automatic Discovery” on page 53](#)
- [“Configuring Peering” on page 54](#)

### About Regular and Enhanced Automatic Discovery

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

Riverbed recommends enhanced auto-discovery for the deployments described in this table.

Deployment Type	Description
Serial Cascade Deployments	<p>Cascade configurations enable optimal multisite deployments where connections between the client and the server might pass through intermediate Steelhead appliances to reach their final destination.</p> <p>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.</p>
Serial Cluster Deployments	<p>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.</p> <p>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.</p> <p>You configure peering rules that define what to do when a Steelhead appliance receives an auto-discovery probe from another Steelhead appliance.</p> <p>You can deploy serial clusters on the client or server-side of the network.</p> <p><b>Note:</b> 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 multiple Steelhead appliance. For example, a serial cluster comprised of two Steelhead appliance models that each have a bandwidth specification of 20-Mbps has a bandwidth specification of 20 Mbps.</p> <p><b>Note:</b> 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.</p>

For details about these deployment types, see the *Steelhead Appliance Deployment Guide*.

## Configuring Peering

You display, add, and modify auto-discovery peering settings in the Configure > Optimization > Peering Rules page.

### To enable enhanced auto-discovery

1. Choose Configure > Optimization > Peering Rules to display the Peering Rules page.
2. Under Settings, complete the configuration as described in this table.

Control	Description
Enable Enhanced Auto-Discovery	<p>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.</p> <p>By default, enhanced auto-discovery peering is enabled. Without 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.</p> <p>IPv6 connections using enhanced auto-discovery use an inner IPv4 channel to the peer Steelhead over a TCP connection. Your network configuration must support IPv4 for use with the inner channels between Steelhead appliances.</p> <p>For detailed information about deployments that require enhanced auto-discovery peering, see the <i>Steelhead Appliance Deployment Guide</i>.</p>

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



## 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 about deployment cases requiring peering rules, see the *Steelhead Appliance 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 36](#).

---

## To configure a peering rule

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

Control	Description
Add a New Peering Rule	Displays the controls for adding a new peering rule.
Rule Type	<p>Determines which action the Steelhead appliance takes on the connection. Select one of these rule types from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b> - 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 <b>Accept</b> 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.</li> <li>• <b>Accept</b> - 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.</li> <li>• <b>Passthrough</b> - 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.</li> </ul>
Insert Rule At	<p>Determines the order in which the system evaluates the rule. Select Start, End, or a rule number from the drop-down list.</p> <p>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.</p> <p>The Rule Type of a matching rule determines which action the Steelhead appliance takes on the connection.</p>
Source Subnet	<p>Specify an IP address and mask for the traffic source, or you can specify All-IP as the wildcard for all IPv4 and IPv6 traffic.</p> <p>Use these formats:</p> <p>XXX.XXX.XXX.XXX/XX (IPv4)</p> <p>X:X:X::X/XXX (IPv6)</p>

Control	Description
Destination Subnet	<p>Specify an IP address and mask pattern for the traffic destination, or you can specify All-IP as the wildcard for all IPv4 and IPv6 traffic.</p> <p>Use these formats:</p> <p>XXX.XXX.XXX.XXX/XX (IPv4)</p> <p>X:X:X::X/XXX (IPv6)</p> <hr/> <p><b>Port</b> - Specify the destination port number, port label, or all.</p>
Peer IP Address	<p>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.</p> <p>The peer client-side Steelhead appliance IP address is IPv4 only.</p>
Description	Specify a description to help you identify the peering relationship.
Add	<p>Adds a peering rule to the list.</p> <p>The Management Console redisplay the Peering Rules table and applies your modifications to the running configuration, which is stored in memory.</p>
Remove Selected Rules	Select the check box next to the name and click <b>Remove Selected Rules</b> .
Move Selected Rules	Select the check box next to the rule and click <b>Move Selected Rules</b> . 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 peers. 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.

When you use this method and add a new remote location in the future, you must 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 36](#)
- [“Configuring General Service Settings” on page 47](#)
- [“Configuring Port Labels” on page 75](#)
- [“Secure Inner Channel Overview” on page 176](#)
- [“Viewing Current Connection Reports” on page 248](#)

---

## Improving Performance

You enable settings to improve network performance in the Configure > Optimization > Performance page. This section describes the default settings and the cases in which you might consider changing the default values.

### Configuring CPU and Data Streamlining 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 Data Streamlining settings

1. Choose Configure > Optimization > Performance to display the Performance page.
2. Under Data Streamlining, complete the configuration as described in this table.

Setting	Description
Classic	Maximizes data reduction.
Turbo	Maximizes LAN throughput.

3. Click **Apply** to apply your settings.
4. Click **Save** to save your settings permanently.

#### 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 this table.

Setting	Description
Multi-Core Balancing	Enables multicore 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 36](#)

---

## Configuring TCP, Satellite Optimization, and High-Speed TCP

This section describes how to configure TCP, satellite optimization, and high-speed TCP settings. It includes these sections:

- [“Optimizing TCP and Satellite WANs” on page 59](#)
- [“High-Speed TCP Optimization” on page 71](#)

You configure TCP, high-speed TCP, and satellite optimization settings in the **Configure > Optimization > Transport Settings** page.

### Optimizing TCP and Satellite WANs

Riverbed provides satellite WAN optimization to overcome the common sources of performance loss associated with space networking. Satellite optimization allows for more effective use of satellite channels, while providing improved user experiences and increased productivity.

SkipWare, an exclusive technology in the Riverbed product family, senses increases and decreases in bandwidth allocation and automatically adjusts its transmission window in response, without requiring user intervention.

### Optimizing SCPS with SkipWare

RiOS includes compatibility settings for the Space Communications Protocol Standards (SCPS) protocol suite. SCPS is designed to allow communication over challenging environments. Originally, it was developed jointly by NASA and DOD’s USSPACECOM to meet their various needs and requirements. Through a collaborative, multiyear R&D effort, the partnership created the Space Communications Protocol Standards-Transport Protocol (SCPS-TP, commonly referred to as “skips”). This protocol now meets the needs of the satellite and wireless communities.

Unlike TCP, the SCPS protocol was designed to operate in an environment of high latency and limited bandwidth. The first commercial implementation of the SCPS protocol was released under the brand name SkipWare.

To use the SkipWare discovery mechanisms included in RiOS, you must install a SkipWare license. SkipWare is enabled automatically when the license is installed, regardless of which transport optimization method is selected (for example, standard TCP, high-speed TCP, or bandwidth estimation). After installing the SkipWare license, you must restart the optimization service.

The basic RiOS license includes non-SkipWare options such as bandwidth estimation and standard TCP.

To change SkipWare settings, you must have role-based permission to use the Optimization Service role. For details, see [“Managing User Permissions” on page 197](#).

---

**Important:** Each Steelhead appliance supports and can interoperate with another Steelhead appliance running the SCPS RSP package. For details, see the *SCPS Installation Guide*. SCPS RSP package licenses are not valid for use as native RiOS SCPS licenses. Contact support or your sales team for assistance in converting SCPS RSP package licenses to native RiOS SCPS licenses.

---

For details and example satellite deployments, see the *Steelhead Appliance Deployment Guide*.


## SCPS Connection Types

You configure satellite optimization settings depending on the connection type. This section describes the connection types. For details about the SCPS discovery process used in various device scenarios, see the *Steelhead Appliance Deployment Guide*.

### RiOS and SCPS Connection

A RiOS and SCPS connection is established between two Steelhead appliances running RiOS v7.0 or later. Because both Steelhead appliances are SCPS-compatible, this is a double-ended connection that benefits from traditional RiOS optimization (SDR and LZ). A RiOS and SCPS connection works with all RiOS features.


Figure 4-1. RiOS and SCPS Connection



### Single-Ended Interception (SEI) Connection

An SEI connection is established between a single Steelhead appliance running RiOS paired with a third-party device running TCP-PEP (Performance Enhancing Proxy). Both the Steelhead appliance and the TCP-PEP device are using the SCPS protocol to speed up the data transfer on a satellite link or other high-latency links. In the following figure, the Steelhead appliance replaces a third-party device running TCP-PEP in the data center, but the Steelhead appliance can also reside in the branch office. Because there is only one Steelhead appliance that intercepts the connection, this is called a single-ended interception (SEI).

Figure 4-2. Single-Ended Interception Connection



Because a single-ended interception connection communicates with only one Steelhead appliance, it

- performs only sender-side TCP optimization.
- supports virtual in-path deployments such as WCCP and PBR.
- cannot initiate a SCPS connection on a server-side out-of-path Steelhead appliance.
- supports kickoff.
- supports auto-discovery failover (failover is compatible with IPv6).
- co-exists with high-speed TCP.
- does not work with connection forwarding.

To configure satellite optimization for an SEI, you define SEI connection rules. The Steelhead appliance uses SEI connection rules to determine whether to enable or pass-through SCPS connections.

Riverbed recommends that for SEI configurations in which the Steelhead appliance initiates the SCPS connection on the WAN, you add an in-path pass-through rule from the client to the server. While the pass-through rule is optional, without it the Steelhead appliance probes for another Steelhead appliance, and when it does not locate one, will failover. Adding the in-path pass-through rule speeds up setup by eliminating the auto-discovery probe and subsequent failover.

The in-path pass-through rule is not necessary on SEI configurations in which the Steelhead appliance terminates the SCPS connection on the WAN, because in this configuration the Steelhead appliance evaluates only the SEI connection rules table and ignores the in-path rules table.

SEI connections count toward the connection count limit on the Steelhead appliance.

---

**Important:** When server-side network asymmetry occurs in a SEI configuration, the server-side Steelhead appliance creates a bad RST log entry in the asymmetric routing table. This differs from other configurations (non-SCPS) in that the client-side Steelhead appliance typically detects asymmetry because of the bad RST and creates an entry in the asymmetric routing table. In SEI configurations, the Steelhead appliance detects asymmetry and creates asymmetric routing table entries independent of other Steelhead appliances. This results in a TCP proxy only connection between the client-side Steelhead appliance and the server when auto-discovery is disabled. For details about the asymmetric routing table, see [“Configuring Asymmetric Routing Features” on page 91](#).

---

## To configure TCP and SkipWare SCPS Optimization

To properly configure transport settings for your environment, you must understand its characteristics. For information on gathering performance characteristics for your environment, see the *Steelhead Appliance Deployment Guide*.

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

2. Under TCP Optimization, complete the configuration as described in this table.

Control	Description
Auto-Detect	<p>Automatically detects the optimal TCP configuration by using the same mode as the peer Steelhead appliance for inner connections, SkipWare when negotiated, or standard TCP for all other cases. This is the default setting.</p> <p>If you have a mixed environment where several different types of networks terminate into a hub or server-side Steelhead appliance, enable this setting on your hub Steelhead appliance so it can reflect the various transport optimization mechanisms of your remote site Steelhead appliances. Otherwise, you can hard code your hub Steelhead appliance to the desired setting.</p> <p>RiOS advertises automatic detection of TCP optimization to a peer Steelhead appliance through the OOB connection between the appliances.</p> <p>Both the client-side and the server-side Steelhead appliances must be running RiOS v7.0 or later.</p> <p>For single-ended interception connections, use SkipWare per-connection TCP optimization when possible; use standard TCP otherwise.</p>
Standard (RFC-Compliant)	<p>Optimizes non-SCPS TCP connections by applying data and transport streamlining for TCP traffic over the WAN. This control forces peers to use standard TCP as well. For details on data and transport streamlining, see the <i>Steelhead Appliance Deployment Guide</i>. This option clears any advanced bandwidth congestion control that was previously set.</p>
HighSpeed	<p>Enables high-speed TCP optimization for more complete use of long fat pipes (high-bandwidth, high-delay networks). Do not enable for satellite networks.</p> <p>Riverbed recommends that you enable high-speed TCP optimization only after you have carefully evaluated whether it will benefit your network environment. For details about the trade-offs of enabling high-speed TCP, see <b>tcp highspeed enable</b> in the <i>Riverbed Command-Line Interface Reference Manual</i>.</p>
Bandwidth Estimation	<p>Uses an intelligent bandwidth estimation algorithm along with a modified slow-start algorithm to optimize performance in long lossy networks. These networks typically include satellite and other wireless environments, such as cellular networks, longer microwave, or Wi-Max networks.</p> <p>Bandwidth estimation is a sender-side modification of TCP and is compatible with the other TCP stacks in the RiOS system. The intelligent bandwidth estimation is based on analysis of both ACKs and latency measurements. The modified slow-start mechanism enables a flow to ramp up faster in high-latency environments than traditional TCP. The intelligent bandwidth estimation algorithm allows it to learn effective rates for use during modified slow start, and also to differentiate BER loss from congestion-derived loss and manage them accordingly. Bandwidth estimation has good fairness and friendliness qualities toward other traffic along the path.</p> <p>The default setting is off.</p>



Control	Description
SkipWare Per-Connection	<p>Applies TCP congestion control to each SCPS-capable connection. The congestion control uses</p> <ul style="list-style-type: none"> <li>• a pipe algorithm that gates when a packet should be sent after receipt of an ACK.</li> <li>• the NewReno algorithm, which includes the sender's congestion window, slow start, and congestion avoidance.</li> <li>• timestamps, window scaling, appropriate byte counting, and loss detection.</li> </ul> <p>This transport setting uses a modified slow-start algorithm and a modified congestion-avoidance approach. This enables SCPS per-connection to ramp up flows faster in high-latency environments, and handle lossy scenarios, while remaining reasonably fair and friendly to other traffic. SCPS per-connection does a very good job of efficiently filling up satellite links of all sizes. SCPS per-connection is a high-performance option for satellite networks.</p> <p>Riverbed recommends enabling per-connection if the error rate in the link is less than approximately 1%.</p> <p>The Management Console dims this setting until you install a SkipWare license.</p>
SkipWare Error-Tolerant	<p>Enables SkipWare optimization with the error-rate detection and recovery mechanism on the Steelhead appliance.</p> <p>This setting allows the per-connection congestion control to tolerate some loss due to corrupted packets (bit errors), without reducing the throughput, using a modified slow-start algorithm and a modified congestion-avoidance approach. It requires significantly more retransmitted packets to trigger this congestion-avoidance algorithm than the SkipWare per-connection setting.</p> <p>Error-tolerant TCP optimization assumes that the environment has a high BER and that most retransmissions are due to poor signal quality instead of congestion. This maximizes performance in high-loss environments, without incurring the additional per-packet overhead of a FEC algorithm at the transport layer.</p> <p>SCPS error tolerance is a high-performance option for lossy satellite networks.</p> <p>Use caution when enabling error-tolerant TCP optimization, particularly in channels with coexisting TCP traffic, because it can be quite aggressive and adversely affect channel congestion with competing TCP flows.</p> <p>Riverbed recommends enabling error tolerance if the error rate in the link is more than approximately 1%.</p> <p>The Management Console dims this setting until you install a SkipWare license.</p>

Control	Description
Enable Rate Pacing	<p>Imposes a global data-transmit limit on the link rate for all SCPS connections between peer Steelhead appliances, or on the link rate for a Steelhead appliance paired with a third-party device running TCP-PEP (Performance Enhancing Proxy).</p> <p>Rate pacing combines MX-TCP and a congestion-control method of your choice for connections between peer Steelhead appliances and SEI connections (on a per-rule basis). The congestion-control method runs as an overlay on top of MX-TCP and probes for the actual link rate. It then communicates the available bandwidth to MX-TCP.</p> <p>Enable rate pacing to prevent these problems:</p> <ul style="list-style-type: none"> <li>• Congestion loss while exiting the slow-start phase. The slow-start phase is an important part of the TCP congestion-control mechanisms that starts slowly increasing its window size as it gains confidence about the network throughput.</li> <li>• Congestion collapse</li> <li>• Packet bursts</li> </ul> <p>Rate pacing is disabled by default.</p> <p>With no congestion, the slow-start phase ramps up to the MX-TCP rate and settles there. When RiOS detects congestion (either due to other sources of traffic, a bottleneck other than the satellite modem, or because of a variable modem rate), the congestion-control method kicks in to avoid congestion loss and exit the slow-start phase faster.</p> <p>Enable rate pacing on the client-side Steelhead appliance along with a congestion-control method. The client-side Steelhead appliance communicates to the server-side Steelhead appliance that rate pacing is in effect. You must also:</p> <ul style="list-style-type: none"> <li>• Enable Auto-Detect TCP Optimization on the server-side Steelhead appliance to negotiate the configuration with the client-side Steelhead appliance.</li> <li>• Configure an MX-TCP QoS rule to set the appropriate rate cap. If an MX-TCP QoS rule is not in place, the system does not apply rate pacing and the congestion-control method takes effect. You cannot delete the MX-TCP QoS rule when rate pacing is enabled.</li> </ul> <p>The Management Console dims this feature until you install a SkipWare license.</p> <p>Rate pacing does not support IPv6.</p> <p>You can also enable rate pacing for SEI connections by defining a SEI rule for each connection.</p>

Control	Description
Enable Single-Ended Connection Rules Table	<p>Enables transport optimization for single-ended interception connections with no Steelhead appliance peer. These connections appear in the rules table.</p> <p>In RiOS v8.5, you can impose rate pacing for single-ended interception connections with no peer Steelhead appliance. By defining an SEI connection rule, you can enforce rate pacing even when the Steelhead appliance is not peered with a SCPS device and SCPS is not negotiated.</p> <p>To enforce rate pacing for a single-ended interception connection, create an SEI connection rule for use as a transport-optimization proxy, select a congestion method for the rule, and then configure a QoS rule (with the same client/server subnet) to use MX-TCP. RiOS v8.5 accelerates the WAN or LAN-originated proxied connection using MX-TCP.</p> <p>By default, the SEI connection rules table is disabled. When enabled, two default rules appear in the rules table. The first default rule matches all traffic with the destination port set to the interactive port label and bypasses the connection for SCPS optimization.</p> <p>The second default rule matches all traffic with the destination port set to the RBT-Proto port label and bypasses the connection for SCPS optimization.</p> <p>This option does not affect the optimization of SCPS connections between Steelhead appliances.</p> <p>When you disable the table, you can still add, move, or remove rules, but the changes do not take effect until you reenable the table.</p> <p>The Management Console dims the SEI rules table until you install a SkipWare license.</p> <p><b>Enable SkipWare Legacy Compression</b> - Enables negotiation of SCPS-TP TCP header and data compression with a remote SCPS-TP device. This feature enables interoperability with RSP SkipWare packages and TurboIP devices that have also been configured to negotiate TCP header and data compression.</p> <p>Legacy compression is disabled by default.</p> <p>After enabling or disabling legacy compression, you must restart the optimization service.</p> <p>The Management Console dims legacy compression until you install a SkipWare license and enable the SEI rules table.</p> <p>Legacy compression also works with non-SCPS TCP algorithms.</p> <p>These limits apply to legacy compression:</p> <ul style="list-style-type: none"> <li>• This feature is not compatible with IPv6.</li> <li>• Packets with a compressed TCP header use IP protocol 105 in the encapsulating IP header; this might require changes to intervening firewalls to permit protocol 105 packets to pass.</li> <li>• This feature supports a maximum of 255 connections between any pair of end-host IP addresses. The connection limit for legacy SkipWare connections is the same as the appliance-connection limit.</li> <li>• QoS limits for the Steelhead appliance apply to the legacy SkipWare connections.</li> </ul> <p>To view SCPS connections, see <a href="#">“Viewing Current Connection Reports” on page 248</a>.</p>

3. Click **Apply** to save your settings to the running configuration.
4. Click **Save** to save your settings permanently.
5. Click **Restart** to restart the optimization service.

## Configuring Buffer Settings

The buffer settings in the Transport Settings page support high-speed TCP and are also used in data protection scenarios to improve performance. For details about data protection deployments, see the *Steelhead Appliance Deployment Guide*.

To properly configure buffer settings for a satellite environment, you must understand its characteristics. For information on gathering performance characteristics for your environment, see the *Steelhead Appliance 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 this 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.

## Adding Single-Ended Connection Rules

You can optionally add rules to control single-ended SCPS connections. The Steelhead appliance uses these rules to determine whether to enable or pass through SCPS connections.

A Steelhead appliance receiving a SCPS connection on the WAN evaluates only the single-ended connection rules table.

To pass through a SCPS connection, Riverbed recommends setting both an in-path rule and a single-ended connection rule.

### To add a single-ended connection rule

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

2. Under Single-Ended Connection Rules, complete the configuration as described in this table.

Control	Description
Add New Rule	Displays the controls for adding a new rule.
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. As an 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 IPv4 or IPv6 address and mask for the traffic source; otherwise, specify All-IP for all IPv4 and IPv6 traffic.  Use these formats: XXX.XXX.XXX.XXX/XX (IPv4) X:X:X::X/XXX (IPv6)
Destination Subnet	Specify an IPv4 or IPv6 address and mask pattern for the traffic destination; otherwise, specify All-IP for all traffic.  Use these formats: XXX.XXX.XXX.XXX/XX (IPv4) X:X:X::X/XXX (IPv6)
Port or Port Label	Specify the destination port number, port label, or all.  Click Port Label to go to the Configure > Networking > Port Labels page for reference.
VLAN Tag ID	Specify one of the following: a VLAN identification number from 1 to 4094; all to specify that the rule applies to all VLANs; or untagged to specify the rule applies to untagged connections.  RiOS supports VLAN v802.1Q. To configure VLAN tagging, configure SCPS 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 preexisting VLAN tagging between the LAN and WAN interfaces.
Traffic	Specifies the action that the rule takes on a SCPS connection. To allow single-ended interception SCPS connections to pass through the Steelhead appliance unoptimized, disable SCPS Discover and TCP Proxy.  Select one of these options: <ul style="list-style-type: none"> <li>• <b>SCPS Discover</b> - Turns SCPS on and TCP proxy off.</li> <li>• <b>TCP Proxy</b> - Turns SCPS off and TCP proxy on.</li> </ul>

Control	Description
Congestion Control Algorithm	<p>Select a method for congestion control from the drop-down list.</p> <ul style="list-style-type: none"> <li>• <b>Standard (RFC-Compliant)</b> - Optimizes non-SCPS TCP connections by applying data and transport streamlining for TCP traffic over the WAN. This control forces peers to use standard TCP as well. For details on data and transport streamlining, see the <i>Steelhead Appliance Deployment Guide</i>. This option clears any advanced bandwidth congestion control that was previously set.</li> <li>• <b>HighSpeed</b> - Enables high-speed TCP optimization for more complete use of long fat pipes (high-bandwidth, high-delay networks). Do not enable for satellite networks.</li> </ul> <p>Riverbed recommends that you enable high-speed TCP optimization only after you have carefully evaluated whether it will benefit your network environment. For details about the trade-offs of enabling high-speed TCP, see <b>tcp highspped enable</b> in the <i>Riverbed Command-Line Interface Reference Manual</i>.</p> <ul style="list-style-type: none"> <li>• <b>Bandwidth Estimation</b> - Uses an intelligent bandwidth estimation algorithm along with a modified slow-start algorithm to optimize performance in long lossy networks. These networks typically include satellite and other wireless environments, such as cellular networks, longer microwave, or Wi-Max networks.</li> </ul> <p>Bandwidth estimation is a sender-side modification of TCP and is compatible with the other TCP stacks in the RiOS system. The intelligent bandwidth estimation is based on analysis of both ACKs and latency measurements. The modified slow-start mechanism enables a flow to ramp up faster in high latency environments than traditional TCP. The intelligent bandwidth estimation algorithm allows it to learn effective rates for use during modified slow start, and also to differentiate BER loss from congestion-derived loss and deal with them accordingly. Bandwidth estimation has good fairness and friendliness qualities toward other traffic along the path.</p> <ul style="list-style-type: none"> <li>• <b>SkipWare Per-Connection</b> - Applies TCP congestion control to each SCPS-capable connection. This method is compatible with IPv6. The congestion control uses <ul style="list-style-type: none"> <li>• a pipe algorithm that gates when a packet should be sent after receipt of an ACK.</li> <li>• the NewReno algorithm, which includes the sender's congestion window, slow start, and congestion avoidance.</li> <li>• timestamps, window scaling, appropriate byte counting, and loss detection.</li> </ul> </li> </ul> <p>This transport setting uses a modified slow-start algorithm and a modified congestion-avoidance approach. This enables SCPS per connection to ramp up flows faster in high-latency environments, and handle lossy scenarios, while remaining reasonably fair and friendly to other traffic. SCPS per connection does a very good job of efficiently filling up satellite links of all sizes. SkipWare per connection is a high-performance option for satellite networks.</p> <p>The Management Console dims this setting until you install a SkipWare license.</p>

Control	Description
	<ul style="list-style-type: none"><li>• <b>SkipWare Error-Tolerant</b> - Enables SkipWare optimization with the error-rate detection and recovery mechanism on the Steelhead appliance. This method is compatible with IPv6.</li></ul> <p>This method tolerates some loss due to corrupted packets (bit errors), without reducing the throughput, using a modified slow-start algorithm and a modified congestion avoidance approach. It requires significantly more retransmitted packets to trigger this congestion-avoidance algorithm than the SkipWare per-connection setting. Error-tolerant TCP optimization assumes that the environment has a high BER and most retransmissions are due to poor signal quality instead of congestion. This maximizes performance in high-loss environments, without incurring the additional per-packet overhead of a FEC algorithm at the transport layer.</p> <p>Use caution when enabling error-tolerant TCP optimization, particularly in channels with coexisting TCP traffic, because it can be quite aggressive and adversely affect channel congestion with competing TCP flows.</p> <p>The Management Console dims this setting until you install a SkipWare license.</p>

Control	Description
Enable Rate Pacing	<p>Imposes a global data transmit limit on the link rate for all SCPS connections between peer Steelhead appliances or on the link rate for a Steelhead appliance paired with a third-party device running TCP-PEP (Performance Enhancing Proxy).</p> <p>Rate pacing combines MX-TCP and a congestion-control method of your choice for connections between peer Steelhead appliances and SEI connections (on a per-rule basis). The congestion-control method runs as an overlay on top of MX-TCP and probes for the actual link rate. It then communicates the available bandwidth to MX-TCP.</p> <p>Enable rate pacing to prevent these problems:</p> <ul style="list-style-type: none"> <li>• Congestion loss while exiting the slow start phase. The slow-start phase is an important part of the TCP congestion-control mechanisms that starts slowly increasing its window size as it gains confidence about the network throughput.</li> <li>• Congestion collapse</li> <li>• Packet bursts</li> </ul> <p>Rate pacing is disabled by default.</p> <p>With no congestion, the slow start ramps up to the MX-TCP rate and settles there. When RiOS detects congestion (either due to other sources of traffic, a bottleneck other than the satellite modem, or because of a variable modem rate), the congestion-control method kicks in to avoid congestion loss and exit the slow start phase faster.</p> <p>Enable rate pacing on the client-side Steelhead appliance along with a congestion-control method. The client-side Steelhead appliance communicates to the server-side Steelhead appliance that rate pacing is in effect. You must also:</p> <ul style="list-style-type: none"> <li>• Enable Auto-Detect TCP Optimization on the server-side Steelhead appliance to negotiate the configuration with the client-side Steelhead appliance.</li> <li>• Configure an MX-TCP QoS rule to set the appropriate rate cap. If an MX-TCP QoS rule is not in place, rate pacing is not applied and the congestion-control method takes effect. You cannot delete the MX-TCP QoS rule when rate pacing is enabled.</li> </ul> <p>The Management Console dims this setting until you install a SkipWare license.</p> <p>Rate pacing does not support IPv6.</p> <p>You can also enable rate pacing for SEI connections by defining a SEI rule for each connection.</p>
Add	Adds the rule to the list. The Management Console redisplay the SCPS 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 <b>Remove Selected</b> .
Move Selected Rules	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.



---

**Tip:** After you apply your settings, you can verify whether changes have had the desired effect by viewing the Current Connections report. The report summarizes the optimized established connections for SCPS. SCPS connections appear as typical established, optimized or established, or single-ended optimized connections. Click the connection to view details. SCPS connection detail reports display SCPS Initiate or SCPS Terminate under Connection Information. Under Congestion Control, the report displays the congestion control method that the connection is using.

---

## 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) for which 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 about using HS-TCP in data protection scenarios, see the *Steelhead Appliance Deployment Guide*.

---

### HS-TCP Basic Steps

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

Task	Reference
1. Enable high-speed TCP support.	<a href="#">“Optimizing TCP and Satellite WANs” on page 59.</a>
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.	<a href="#">“To configure buffer settings” on page 66.</a>
3. Increase the LAN buffers to 1 MB.	<a href="#">“To configure buffer settings” on page 66.</a>
4. Enable in-path support.	<a href="#">“Configuring General Service Settings” on page 47.</a>

## Configuring Service Ports

You 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.
2. Under Service Port Settings, complete the configuration as described in this 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 this 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 <b>Remove Selected</b> .

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

### Related Topic

- [“Configuring General Service Settings” on page 47](#)

## Configuring Host Labels

You create host labels in the Configure > Networking > Host Labels page.

Host labels are names given to sets of hostnames and subnets to streamline configuration. Host labels provide flexibility because you can create a logical set of hostnames and then apply a rule, such as a QoS rule, to the entire set instead of creating individual rules for each hostname. For example, you can define a set of file servers in a host label, use that host label in a single QoS rule, and apply a policy limiting all IP traffic to and from the servers (independent of what protocol or application is in use).

Other ways to use host labels:

- List multiple dedicated application servers by hostname in a single rule and apply a policy
- List multiple business websites and servers to protect
- List recreational websites to restrict

Host labels are *optional*.

Host labels are not compatible with IPv6.

## Creating a Host Label

### To create a host label

1. Choose Configure > Networking > Host Labels to display the Host Labels page.
2. To add a host label, complete the configuration as described in this table.

Control	Description
Add a New Host Label	Displays the controls to add a new host label.
Name	<p>Specify the label name. These rules apply:</p> <ul style="list-style-type: none"> <li>• Host label names are case-sensitive and can be any string consisting of letters, numbers, the underscore ( _ ), or the hyphen ( - ). There cannot be spaces in host labels.</li> <li>• Riverbed suggests starting the name with a letter or underscore.</li> <li>• To avoid confusion, do not use a number for a host label.</li> <li>• You cannot delete host labels that a QoS rule is using.</li> </ul>
Hostnames/Subnets	<p>Specify a comma-separated list of hostnames and subnets. Hostnames are not case-sensitive. You can also separate hostname and subnet names with spaces or new lines.</p> <p>Use this format:</p> <p>XXX.XXX.XXX.XXX/XX where /XX is a subnet mask value between 0 and 32.</p> <p>A hostname can appear in multiple host labels. You can use up to 100 unique hostnames.</p> <p>A host label can contain up to 64 subnets and hostnames.</p>

Control	Description
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> . You cannot delete host labels that a QoS rule is using.
Add New Host Label	Adds the host label. The page updates the host label table with the new host label. Because the system resolves new hostnames through the DNS, wait a few seconds and then refresh your browser.

## Resolving Hostnames

RiOS resolves hostnames through a DNS server immediately after you add a new host label or after you edit an existing host label. RiOS also automatically re-resolves hostnames once daily. If any problems arise during the automatic or manual hostname resolution, the summary section of the host labels page alerts you quickly that there is a problem.

RiOS relays any changes in IP addresses to QoS after resolving them; you do not need to update the host label in QoS.

When you know that the IP addresses associated with a hostname have been updated in the DNS server, and you do not want to wait until the next scheduled resolution, you can resolve the hostnames manually. After you resolve the hostname cache manually, RiOS schedules the next resolve time to be 24 hours in the future.

### To resolve hostnames through the DNS immediately

- Click the **Resolve Hostnames** button.

### To show or hide the resolved IP addresses of the hostnames

- Select or clear the Show resolved IPs for the hostnames in the table below check box.

When the system resolves a hostname, the elapsed time appears next to the Resolved label.

## Viewing the Hostname Resolution Summary

The summary section displays this information:

- **Unique Hostnames** - The total number of unique hostnames, because a hostname can appear in multiple host labels. You can configure a maximum of 100 unique hostnames.
- **Checking DNS** - The number of unique hostnames that are actively being resolved.
- **Unresolvable** - The number of unique hostnames that cannot be resolved through the DNS because the DNS server is not configured, the DNS server is not reachable due to network connectivity issues, there is a typo in the hostname, and so on.

On rare occasions, if the DNS server goes down after resolving a hostname once, the system keeps the information, even though it might be stale. When this occurs, the following message appears:

“Note: This hostname was resolved successfully at least once in the past but the last attempt failed.”

## Modifying Hostnames or Subnets in a Host Label

You add or delete hostnames or subnets associated with a host label in the Host Labels page.

### To modify hostnames or subnets in a host label

1. Choose **Configure > Networking > Host Labels** to display the Host Labels page.
2. Select the host label name in the Host Label table.
3. Add or delete hostnames or subnets in the Hostnames/Subnets text box.
4. Click **Apply** to save your settings to the running configuration. RiOS immediately applies host label changes to QoS, changing the traffic processing for all rules using the label.
5. Verify that any new hostnames resolve successfully to the expected IP addresses.

### Related Topics

- [“Modifying General Host Settings” on page 15](#)
- [“Creating QoS Classes” on page 130](#)

---

## 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.

This table summarizes the port labels that are provided by default.

Port Type	Description and Ports
Granite	Use this port label to automatically pass-through traffic on Riverbed Granite ports 7950 - 7954 (data transfers), and 7970 (management). Granite delivers block-storage optimization that accelerates access to storage area networks (SANs) across the WAN, decoupling storage from servers and allowing data to reside in one location.
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 (RiOS data store 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, RBT-Proto, secure ports or FTP, you must delete the Interactive, RBT-Proto, Secure, and FTP in-path rules. For details, see [“In-Path Rules Overview” on page 33](#).

For information, on common port assignments see [“Steelhead Appliance Ports” on page 359](#). This feature is *optional*.

### To create a port label

1. Choose **Configure > Networking > Port Labels** to display the Port Labels page.
2. To add a port label, complete the configuration as described in this table.

Control	Description
Add a New Port Label	Displays the controls to add a new port label.
Name	Specify the label name. These rules apply: <ul style="list-style-type: none"> <li>• 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.</li> <li>• The fields in the various rule pages of the Management Console that take a physical port number also take a port label.</li> <li>• To avoid confusion, do not use a number for a port label.</li> <li>• Port labels that are used in in-path and other rules, such as QoS and peering rules, cannot be deleted.</li> <li>• 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.</li> </ul>
Ports	Specify a comma-separated list of ports.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .
Add	Adds the port label.

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

## Modifying Ports in a Port Label

You 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.
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 36](#)
- [“Enabling Peering and Configuring Peering Rules” on page 53](#)
- [“Creating QoS Classes” on page 130](#)

---

## 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.

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) requires configuration of the RiOS SRDF storage optimization module. For details on storage technologies that originate traffic through FC, see the *Steelhead Appliance 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 about data replication deployments, see the *Steelhead Appliance Deployment Guide*.

### To configure FCIP optimization

1. Choose Configure > Optimization > FCIP to display the FCIP page.
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 might 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 78](#).

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 must restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

## 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 must 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 Connection Reports” on page 248](#).

---

**Note:** You can view combined throughput and reduction statistics for two or more FCIP tunnel ports by entering this 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 GigE-based (RF port) originated SRDF traffic between VMAX arrays must isolate DIF headers within the data stream. These DIF headers further interrupt the data stream. When the R1 Symmetrix array is running Enginuity microcode version 5875 or newer, manual FCIP rules are not necessary. In 5875+ environments, RiOS automatically detects the presence of DIF headers and DIF blocksize for GigE-based (RF port) SRDF traffic. To manually isolate the DIF headers when the R1 Symmetrix array is running Enginuity microcode version 5874 or older, you add FCIP rules by defining a match for source or destination IP traffic.

Automatically-detected FCIP settings in Enginuity 5875 and later environments override any manually configured FCIP rules.

### 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 this table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule. Displays the controls for adding a manual rule. Use this control when the R1 Symmetrix array is running Enginuity microcode version 5874 or earlier.
Source IP	Specify the connection source IP address of the FCIP gateway tunnel endpoints. <b>Note:</b> The source IP address cannot be the same as the destination IP address.
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 manual rule to the list. The Management Console redisplay 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 <b>Remove Selected</b> .

3. Click **Apply** to save your settings to the running configuration.
4. Click **Save** to save your settings permanently.
5. You must restart the optimization service after adding or removing a FCIP rule. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

### 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 preexisting 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 this rule.

1. Choose **Configure > Optimization > FCIP** to display the FCIP page.
2. Click **Add a New Rule**.
3. Specify these 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 80](#)

---

## 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 performs the data replication over GigE (instead of the Fibre Channel), using gateways that implement the SRDF protocol.

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 *Steelhead Appliance 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 scalable data replication (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 as originally presented to the Steelhead network.

---

**Note:** Traffic originated through Symmetrix GigE ports (RE ports) requires 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.

---

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 about data replication deployments, see the *Steelhead Appliance Deployment Guide*.

### To configure SRDF optimization

1. Choose Configure > Optimization > SRDF to display the SRDF page.
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 might need additional configuration to isolate and optimize the DIFs embedded within the data payload. For details, see [“Creating SRDF Rules \(VMAX-to-VMAX Traffic Only\)” on page 84](#).

3. Optionally, specify nonstandard individual SRDF port numbers separated by commas. Do not specify a port range.

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 must restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

## 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, RiOS is not optimizing the connection. You must 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. A SRDF protocol error can occur when attempting to optimize traffic originating from the LAN side of the Steelhead appliance. Check the LAN-side Symmetrix array for compatibility.
- If a protocol error does not appear next to the SRDF connection on the client-side Steelhead appliance, RiOS is optimizing the connection normally.

For details, see [“Viewing Current Connection Reports” on page 248](#).

## Setting a Custom Data Reduction Level for an RDF Group

This section describes how to apply custom data reduction levels to remote data facility (RDF) groups.

You can base the data reduction level on the compression characteristics of the data associated with an RDF group to provide SRDF selective optimization. Selective optimization enables you to find the best optimization setting for each RDF group, maximizing the Steelhead appliance use. Selective optimization depends on an R1 Symmetrix array running VMAX Engenuity microcode levels newer than 5874.

For example, you can customize the data reduction level for applications associated with an RDF group when excess WAN bandwidth is available and the application data associated with the group is not reducible. For applications with reducible data, getting maximum reduction might be more important, requiring a more aggressive data reduction level.

You can configure the optimization level from no compression to full scalable data replication (SDR). SDR optimization is the default, and includes LZ compression on the cold, first-pass of the data. You can also configure LZ-compression alone with no SDR.

Consider an example with these types of data:

- Oracle logs (RDF group 1)
- Encrypted check images (RDF group 2)
- Virtual machine images (RDF group 3)

In this example, you can assign LZ-only compression to the Oracle logs, no optimization to the encrypted check images, and default SDR to the virtual machine images. To assign these levels of optimization, you configure the Steelhead appliance to associate specific RE port IP addresses with specific Symmetrix arrays, and then assign a group policy to specific RDF groups to apply different optimization policies.

The data reduction level within a group policy overrides the current default data reduction setting for the storage resources an RDF group represents. This override is distinct per Symmetrix ID.

**To configure a custom data reduction group policy for a Symmetrix ID:**

1. Choose Configure > Optimization > SRDF to display the SRDF page.
2. Under Symmetrix IDs and Group Override Policies, complete the configuration as described in this table.

Control	Description
Add a Symm ID or Group Policy	Displays the tabs for adding a Symmetrix ID or group policy.
Add a Symmetrix ID	Select to display the controls for adding a Symmetrix ID.
Symm ID	Specify the Symmetrix ID. The Symmetrix ID is an alphanumeric string that can contain hyphens and underscores (for example, a standard Symmetrix serial number is 000194900363). Do not use spaces or special characters.  Each Symmetrix ID can have 0-254 group override policies.
Source IPs	Specify the connection source IP address of the Symmetrix DMX or VMAX GigE ports (RE ports) originating the replication.
Add a Group Policy	Select to display the controls for adding a group policy.
RDF Group	Specify the RDF group number. Symmetrix arrays that are serving Open Systems hosts and are using EMC Solutions Enabler report RDF group numbers in decimal, ranging from 1 - 255 (this is the RIOS default).  Mainframe-attached Symmetrix arrays report RDF group numbers in hexadecimal, ranging from 0 - 254.  You cannot add an RDF group until a Symmetrix ID exists.
Symmetrix ID	Specify an IP address of the Symmetrix DMX or VMAX GigE ports (RE ports) originating the replication.
Data Reduction Policy	By default, SDR uses the in-path rule data reduction policy. Select one of these data reduction policies from the drop-down list to override the in-path rule data reduction policy: <ul style="list-style-type: none"> <li>• <b>Default</b> - Performs LZ compression and SDR.</li> <li>• <b>LZ</b> - Performs LZ compression; does not perform SDR.</li> <li>• <b>None</b> - Disables SDR and LZ compression.</li> </ul>
Description	Describe the policy to facilitate administration: for example, Oracle 1 DB.
Add	Adds the ID or policy to the list. The Management Console redisplay the list 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 <b>Remove Selected</b> .

## Creating SRDF Rules (VMAX-to-VMAX Traffic Only)

Environments with GigE-based (RE port) originated SRDF traffic between VMAX arrays must isolate DIF headers within the data stream. These DIF headers further interrupt the data stream. When the R1 Symmetrix array is running Enginuity microcode version 5875 or newer, manual SRDF rules are not necessary. In 5875+ environments, RiOS automatically detects the presence of DIF headers and DIF blocksize for GigE-based (RE port) SRDF traffic. To manually isolate the DIF headers when the R1 Symmetrix array is running Enginuity microcode version 5874 or older, you add SRDF rules by defining a match for source or destination IP traffic.

Automatically-detected SRDF settings in Enginuity 5875 and later environments override any manually configured SRDF rules.

### 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 an SRDF rule

1. Choose Configure > Optimization > SRDF to display the SRDF page.
2. Under Rules, complete the configuration as described in this table.

Control	Description
Add a New Rule	Displays the controls for adding a manual rule. Use this control when the R1 Symmetrix array is running Enginuity microcode version 5874 or earlier.
Source IP	Specify the connection source IP address of the Symmetrix DMX or VMAX GigE ports (RE ports) originating the replication. <b>Note:</b> The source IP address cannot be the same as the destination IP address.
Destination IP	Specify the connection destination IP address of the Symmetrix DMX or VMAX GigE ports (RE ports) receiving the replication.
Enable DIF	Isolates and optimizes the Data Integrity Fields embedded within the SRDF 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 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 manual 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.

Control	Description
Add	Adds the manual rule to the list. The Management Console redisplay 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 <b>Remove Selected</b> .
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 must restart the optimization service after adding or removing a SRDF rule. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

#### To edit an 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 Topics

- [“Configuring FCIP Optimization” on page 77](#)
- [“Viewing SRDF Reports” on page 302](#)

## Configuring SnapMirror Optimization

You enable and modify SnapMirror storage optimization module settings in the Configure > Optimization > SnapMirror page. SnapMirror optimization support is for environments using NetApp ONTAP v7 or Data ONTAP v8 configured for 7-mode.

SnapMirror is used mainly for disaster recovery and replication. To provide maximum protection and ease of management, many enterprises choose to perform SnapMirror operations across the wide-area network. However, WAN links are often costly. Furthermore, the limited bandwidth and high-network latency they provide often severely degrade SnapMirror operations.

The two types of SnapMirror are volume-based and qtree-based. SnapMirror replicates data from one volume or qtree (the source) to another volume or qtree (the mirror). It then periodically updates the mirror to reflect incremental changes to the source. The result of this process is an online, read-only volume (the mirror) that contains the same data as the source volume at the time of the most recent update.

You can use the information on the mirror to:

- Provide quick access to data in the event of a disaster that makes the source volume or qtree unavailable. The secondary copy is nearly identical to the primary copy; every Snapshot copy on the primary copy also exists on the backup copy. You can schedule updates as frequently as every minute.
- Update the source to recover from disaster, data corruption (mirror qtrees only), or user error.
- Archive the data to tape.
- Balance resource loads.
- Back up or distribute the data to remote sites.

Due to the large amount of data transferred, a task such as mirror initialization can take days to complete over a WAN. Some applications use a NetApp storage device called a filer. Filers touch large numbers of storage blocks as they add, delete, and modify files during a typical workday. This results in the marking of many blocks on the filer for mirroring, which means that incremental updates to remote mirror copies might take hours to complete.

## How a Steelhead Appliance Optimizes SnapMirror Traffic

The Steelhead appliance improves the performance of the WAN for NetApp SnapMirror traffic by overcoming limited bandwidth restrictions, high latency, and poor network quality commonly associated with wide-area networks.

In addition, RiOS v8.5 introduces the following features, which further improve WAN performance, visibility, and control of NetApp SnapMirror:

- Present performance statistics and apply optimization policies based on source and destination volume and host pairs.
- Fine-tune network QoS policies for individual volumes, filers, or for SnapMirror as a whole.
- Assign mappings by filer and volume name to one of five volume priorities. Using advanced QoS, you can assign a service class and DSCP value to each volume priority when creating a rule for SnapMirror traffic.
- Collect SnapMirror statistics, such as the total LAN/WAN bytes in and out and the active cycle time.

By default, SnapMirror optimization is disabled. To benefit from SnapMirror optimization, both Steelhead appliances must be running RiOS v8.5.

For details about data replication deployments, see the *Steelhead Appliance Deployment Guide*.

### To configure SnapMirror optimization

1. On the source filer-side Steelhead appliance, choose Configure > Optimization > SnapMirror to display the SnapMirror page.
2. Under SnapMirror Settings, select Enable SnapMirror.
3. By default, RiOS directs all traffic on the standard port 10566 through the SnapMirror module for optimization. Optionally, specify nonstandard individual SnapMirror port numbers, separated by commas.

Do not specify a port range.

The SnapMirror ports field must always contain at least one port.

SnapMirror optimization does not support port 10565 for multipath traffic.



4. Click **Add a New Filer or Volume/QTree**.
5. Select the **Add a Filer** tab.
6. Complete the configuration as described in this table.

Control	Description
Filer Name	Specify the name of the filer. RiOS automatically detects the volumes associated with the filer, or you can optionally add volumes to it later.
IP Addresses	Specify source IPv4 addresses to associate with the filer, separated by a comma. You cannot specify IPv6 addresses.
Filer Default Optimization Policy	<p>You can configure the optimization level from no compression (none) to full Scalable Data Replication (SDR-Default).</p> <p>SDR optimization includes LZ compression on the cold, first-pass of the data. You can also configure LZ-compression alone (LZ-only) with no SDR. For some applications, it might be more important to get maximum throughput with minimal latency, and without compression; for others, getting maximum reduction is more important.</p> <p>Select an optimization policy for the default volumes and qtrees on this filer:</p> <ul style="list-style-type: none"> <li>• <b>SDR-Default</b> - Performs SDR and LZ compression. This is the default policy.</li> <li>• <b>LZ-only</b> - Performs LZ compression only. There is no SDR optimization with this policy.</li> <li>• <b>None</b> - Disables SDR and LZ compression.</li> </ul>
Filer Default SnapMirror Priority	Select a priority for use later in a QoS service class: Highest, High, Medium, Low, Lowest, No Setting. The default priority is Medium. No setting means that there is no priority and the QoS default rules apply.
Description	Optionally, specify a volume description or provide additional comments.
Add	Adds the filer to the list. The Management Console redisplay the Filer 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 <b>Remove Selected</b> .

7. Click **Apply** to save your settings to the running configuration.
8. Click **Save** to save your settings permanently.
9. If you have enabled or disabled SnapMirror optimization or changed a port, you must restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).
10. On the destination filer-side Steelhead appliance, choose **Configure > Optimization > SnapMirror**, select **Enable SnapMirror**, and restart the optimization service.

## Viewing SnapMirror Connections

After completing the SnapMirror configuration on both Steelhead appliances and restarting the optimization service, you can view the SnapMirror connections by choosing **Reports > Optimization > SnapMirror**. For details, see [“Viewing SnapMirror Reports” on page 304](#).

## Adding or Modifying a Filer

This section describes how to create a new filer or make changes to an existing filer. You must add a filer before you can add a volume. SnapMirror needs both a source and a destination IP address for each filer.

### To add a SnapMirror filer

1. Choose Configure > Optimization > SnapMirror to display the SnapMirror page.
2. Click **Add a New Filer or Volume/QTree**.
3. Select the Add a Filer tab.
4. Complete the configuration as described in this table.

Control	Description
Filer Name	Specify the name of the filer. RiOS automatically detects the volumes associated with the filer, or you can optionally add volumes to it later.
IP Addresses	Specify source IPv4 addresses to associate with the filer, separated by a comma. You cannot specify IPv6 addresses.
Filer Default Optimization Policy	<p>You can configure the optimization level from no compression (none) to full Scalable Data Replication (SDR-Default).</p> <p>SDR optimization includes LZ compression on the cold, first-pass of the data. You can also configure LZ-compression alone (LZ-only) with no SDR. For some applications, it might be more important to get maximum throughput with minimal latency, and without compression; for others, getting maximum reduction is more important.</p> <p>Select an optimization policy for the default volumes and qtrees on this filer:</p> <ul style="list-style-type: none"> <li>• <b>SDR-Default</b> - Performs SDR and LZ compression. This is the default policy.</li> <li>• <b>LZ-only</b> - Performs LZ compression only. There is no SDR optimization with this policy.</li> <li>• <b>None</b> - Disables SDR and LZ compression.</li> </ul>
Filer Default SnapMirror Priority	Select a priority for use later in a QoS service class: Highest, High, Medium, Low, Lowest, No Setting. The default priority is Medium. No setting means that there is no priority and the QoS default rules apply.
Description	Optionally, specify a volume description or provide additional comments.
Add	Adds the filer to the list. The Management Console redisplay the Filer 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 <b>Remove Selected</b> .

5. Click **Apply** to save your settings to the running configuration.
6. Click **Save** to save your settings permanently.
7. Restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

### To add a SnapMirror volume or qtree

1. Choose Configure > Optimization > SnapMirror to display the SnapMirror page.
2. Click **Add a New Filer or Volume/QTree**.
3. Select the Add a Volume/QTree tab.
4. Complete the configuration as described in this table.

Control	Description
Volume Name	Specify the name of the volume.
Filer	Select a predefined filer from the drop-down list.
Optimization Policy	<p>By default, the volumes use the same optimization policy as the filer. With this setting, when you change the policy on the filer, the policy setting on the volumes updates automatically.</p> <p>Select an optimization policy for the volume:</p> <ul style="list-style-type: none"> <li>• <b>SDR-Default</b> - Performs SDR and LZ compression. This is the default policy.</li> <li>• <b>Filer-Default</b> - Sets the volume optimization policy to be the same as the filer values. This is the default policy.</li> <li>• <b>LZ-only</b> - Enables LZ compression only. There is no SDR optimization with this policy.</li> <li>• <b>None</b> - Disables SDR and LZ compression.</li> </ul>
SnapMirror Priority	Select a priority for use later in a QoS service class: Highest, High, Filer-Default, Low, Lowest, No Setting. The default priority is Filer-Default, which uses the same priority as the filer. With this setting, when you change the priority on the filer, the priority for the volume updates automatically.
Add	Adds the rule to the list. The Management Console redisplay 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 <b>Remove Selected</b> .

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



## CHAPTER 5      **Configuring Network Integration Features**

This chapter describes how to configure advanced features such as asymmetric routing, connection forwarding, encryption, flow export, QoS, simplified routing, and WCCP.

This chapter includes these topics:

- [“Configuring Asymmetric Routing Features” on page 91](#)
- [“Configuring Connection Forwarding Features” on page 94](#)
- [“Configuring IPSec Encryption” on page 97](#)
- [“Configuring Subnet Side Rules” on page 99](#)
- [“Configuring Flow Statistics” on page 101](#)
- [“Applying QoS Policies” on page 107](#)
- [“Configuring Outbound QoS \(Basic\)” on page 115](#)
- [“Configuring Outbound QoS \(Advanced\)” on page 128](#)
- [“Configuring Inbound QoS” on page 144](#)
- [“Selecting WAN Paths Dynamically” on page 153](#)
- [“Configuring Simplified Routing Features” on page 159](#)
- [“Configuring WCCP” on page 160](#)
- [“Configuring Hardware-Assist Rules” on page 166](#)

For details about basic and advanced deployment types, see the *Steelhead Appliance 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.

The asymmetric routing feature in RiOS v8.5 is compatible with IPv6.

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 automatic 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. This 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.	<ul style="list-style-type: none"> <li>Asymmetric Routing Table: bad RST</li> <li>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)</li> </ul>
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.	<ul style="list-style-type: none"> <li>Asymmetric Routing Table: bad SYN/ACK</li> <li>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)</li> </ul>
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.	<ul style="list-style-type: none"> <li>Asymmetric Routing Table: no SYN/ACK</li> <li>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)</li> </ul>
multiSYN Retransmit	The types of multiSYN retransmits are: <ul style="list-style-type: none"> <li>Probe-filtered occurs when the client-side Steelhead appliance sends out multiple SYN+ frames and does not get a response.</li> <li>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.</li> </ul>	<ul style="list-style-type: none"> <li>Asymmetric Routing Table: probe-filtered(not-AR)</li> <li>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.</li> </ul>

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 multiport 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 94](#) or the *Steelhead Appliance Deployment Guide*.

## Troubleshooting Asymmetric Routes

You can use these 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.

As an 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 Dump Files” on page 331](#).

- **Trace Route** - From the CLI, run the **tracert** 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 **tracert** command with the IP address of the server, then run the **tracert** 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# tracert 10.0.0.4 Type escape sequence to abort.
Tracing the route to 10.0.0.4
 0 10.1.0.1 4 msec 0 msec 4 msec
 1 10.0.0.2 4 msec 4 msec 0 msec
 2 10.0.0.3 4 msec 4 msec 0 msec
 3 10.0.0.4 4 msec 4 msec 0 msec
server# tracert 10.1.0.2 Type escape sequence to abort.
Tracing the route to 10.1.0.2
 0 10.0.0.6 4 msec 0 msec 4 msec
 1 10.0.0.5 4 msec 4 msec 0 msec
 2 10.1.0.1 4 msec 4 msec 0 msec
 3 10.1.0.2 4 msec 4 msec 0 msec
```

For details, see the *Riverbed Command-Line Interface Reference Manual* or the *Steelhead Appliance Deployment Guide*.

### To automatically detect asymmetric routing

1. Choose Configure > Networking > Asymmetric Routing to display the Asymmetric Routing page.

2. Under Asymmetric Routing Settings, complete the configuration as described in this table.

Control	Description
Enable Asymmetric Routing Detection	Detects asymmetric routes in your network.
Enable Asymmetric Routing Pass-Through	<p>Enables pass-through traffic if asymmetric routing is detected.</p> <p>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.</p> <p>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 multiport Steelhead appliance, connection forwarding, or using external ways to redirect packets, such as WCCP or PBR.</p> <p>For details, see the <i>Steelhead Appliance Deployment Guide</i>.</p>
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

### Related Topics

- [“Configuring Connection Forwarding Features” on page 94](#)
- [“Generating System Dumps” on page 329](#)
- [“Viewing Process Dumps” on page 330](#)

## Configuring Connection Forwarding Features

You configure connection forwarding for a network with multiple paths from the server in the **Configure > Networking > Connection Forwarding** page.

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 always 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 must 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.


When RiOS determines an IPv6 incompatibility between connection-forwarding neighbors, it triggers an alarm indicating that a peer Steelhead appliance is incompatible. For details, see [“Configuring Alarm Settings” on page 217](#) and [“Viewing Alarm Status Reports” on page 307](#).



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 appliance with the lowest IP address assigns all traffic flows to itself. For details, see the *Steelhead Appliance Deployment Guide*.

When using connection forwarding in a WCCP cluster with IPv6, Riverbed recommends upgrading all Steelhead appliances in the cluster to RiOS v8.5. You must also enable multiple interface support.

**Figure 5-1. Asymmetric Network**



You can place neighbors in the same physical site or in different sites, but the latency between them must be small because the packets traveling 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 acknowledgment from the furthest neighbor.

**Important:** Connection-forwarding neighbors must use the same WAN visibility mode. For details, see [“Configuring In-Path Rules”](#) on page 36.

For details about connection forwarding, see the *Steelhead Appliance Deployment Guide*.

### To enable connection forwarding

1. Choose **Configure > Networking > Connection Forwarding** to display the Connection Forwarding page.
2. Under Connection Forwarding Settings, complete the configuration as described in this 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. The default value is 1 second.

Control	Description
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.
In-Path Neighbor Failure	<p>Uses the neighbor appliance to optimize new connections if the appliance fails.</p> <p>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.</p> <p>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.</p>
Multiple Interface Support	<p>Enables high availability on Steelhead appliances configured with multiple in-path interfaces and using connection forwarding with another multiport 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.</p> <p>RiOS v6.5 and later requires connection forwarding in a WCCP cluster.</p> <p>You must enable multiple interface support for a connection-forwarding neighbor to work with IPv6.</p>

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 this table.

Control	Description
Add a New Neighbor	Displays the controls to add a new neighbor.
Hostname	Specify a hostname.
In-Path IP Address	<p>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.</p> <p>To use connection forwarding with IPv6, both Steelhead appliances must be running RiOS v8.5 and you must enable multiple interface support.</p>
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 <b>Remove Selected</b> .

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 47](#)
- [“Configuring Asymmetric Routing Features” on page 91](#)

---

## **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 does not support IPsec over IPv6.

RiOS v6.0 and later also provides support for SSL peering beyond traditional HTTPS traffic. For details, see [“Configuring Secure Peers” on page 176](#).

---

**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.
2. Under General Settings, complete the configuration as described in this 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.

Control	Description
Encryption Policy	<p>Select one of these encryption methods from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>DES</b> - Encrypts data using the Data Encryption Standard algorithm. DES is the default value.</li> <li>• <b>NULL</b> - Specifies the null encryption algorithm.</li> <li>• <b>None</b> - Does not apply an encryption policy.</li> <li>• <b>3DES</b> - 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.</li> <li>• <b>AES</b> - 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.</li> <li>• <b>AES256</b> - 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.</li> </ul> <p>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.</p> <p><b>Note:</b> 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.</p>
Authentication Policy	<p>Select one of these authentication methods from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>MD5</b> - Specifies the Message-Digest 5 algorithm, a widely-used cryptographic hash function with a 128-bit hash value. This is the default value.</li> <li>• <b>SHA-1</b> - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA-1 is considered to be the successor to MD5.</li> </ul> <p>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.</p>
Time Between Key Renegotiations	<p>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.</p>
Enter the Shared Secret/Confirm the Shared Secret	<p>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.</p>
Add a New Secure Peer	<p>Displays the controls to add a new secure peer.</p> <ul style="list-style-type: none"> <li>• <b>Peer IP Address</b> - Specify the IP address for the peer Steelhead appliance (in-path interface) for which you want to make a secure connection.</li> </ul>

Control	Description
Add	<p>Adds the peer specified in the Peer IP Address text box.</p> <p>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.</p> <p><b>Note:</b> Adding a peer causes a short service disruption (3-4 seconds) to the peer that is configured to use IPSec security.</p>
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

4. If you have changed an IPSec encryption setting, you must restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

---

**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 and one of these states:

- **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 configure subnet side rules in the Configure > Networking > Subnet Side Rules page.

Subnet side rules are used in a virtual in-path deployment to support Flow Export, to support a client-side Steelhead appliance, or to exempt certain subnets from QoS enforcement.

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

- LAN-bound traffic that traverses the WAN interface of the Steelhead appliance is exempt from QoS enforcement. For details, see [“Bypassing LAN Traffic” on page 113](#).
- Client-side Steelhead appliances configured for virtual in-path deployment can optimize traffic from client-side connections. Otherwise, the appliance does not optimize traffic from client-side connections. In virtual in-path configurations, all traffic flows in and out of one physical interface, and the default subnet side rule causes all traffic to appear to originate from the WAN side of the device.

- Flow export collectors such as NetFlow analyze nonoptimized 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.

FakeIndex is necessary for correct optimized traffic reporting. The fake index feature is enabled by default if you enable the Flow Export option on the Configure > Networking > Flow Statistics page. For details, see the *Steelhead Appliance Deployment Guide*.

### To add subnet side rules

1. Choose Configure > Networking > Subnet Side Rules to display the Subnet Side Rules page.
2. Complete the configuration as described in this 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.
Add	Adds the rule to the subnet map table. The Management Console redisplay 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 <b>Remove Subnet Rules</b> .
Move Subnet Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

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 Topics

- [“Configuring Flow Statistics” on page 101](#)
- [“Applying QoS Policies” on page 107](#)

---

## Configuring Flow Statistics

You enable and configure flow statistic settings in the Configure > Networking > Flow Statistics page. You can also enable flow export to an external collector and to a CascadeFlow collector. CascadeFlow collectors can aggregate information about QoS configuration and other application statistics to send to a Cascade Enterprise Profiler appliance. The Enterprise Profiler summarizes and displays the QoS configuration statistics.

By default, flow export is disabled.

External collectors use information about network data flows to report trends such as the top users, peak usage times, traffic accounting, security, and traffic routing. You can export preoptimization 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.

## Enabling Flow Export

Steelhead appliances support NetFlow v5.0, CascadeFlow, NetFlow v9, and CascadeFlow-compatible. Flow export requires these 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 about configuring flow export in a virtual in-path deployment, see [“Configuring Subnet Side Rules” on page 99](#).

To distinguish between LAN-to-WAN and WAN-to-LAN traffic in virtual in-path deployments, see the *Steelhead Appliance Deployment Guide*.

## To enable flow statistic settings

1. Choose Configure > Networking > Flow Statistics to display the Flow Statistics page.
2. Under Flow Statistics Settings, complete the configuration as described in this table.

Control	Description
Enable Application Visibility	<p>Continuously collects detailed application-level statistics for both passthrough and optimized traffic. The Application Visibility and Application Statistics reports display these statistics. This statistic collection is disabled by default.</p> <p>To view the reports, choose Reports &gt; Networking &gt; Application Statistics or Application Visibility.</p>
Enable WAN Throughput Statistics	<p>Continuously collects WAN throughput statistics, which the WAN Throughput report displays. This statistic collection is enabled by default; however, you can disable the collection to save processing power.</p> <p>To view the WAN throughput statistics, choose Reports &gt; Networking &gt; WAN Throughput.</p>
Enable Top Talkers	<p>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.</p> <p>The most active, heaviest users of WAN bandwidth are called the <i>Top Talkers</i>. 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.</p> <p>You can analyze the Top Talkers for accounting, security, troubleshooting, and capacity planning purposes. You can also export the complete list in CSV format.</p> <p>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).</p> <p>A NetFlow collector is not required for this feature.</p> <p>Optionally, select a time period to adjust the collection interval:</p> <ul style="list-style-type: none"> <li>• <b>24-hour Report Period</b> - For a five-minute granularity (the default setting).</li> <li>• <b>48-hour Report Period</b> - For a ten-minute granularity.</li> </ul> <p>The system also uses the time period to collect SNMP Top Talker statistics. For top talkers displayed in the Top Talker report and SNMP Top Talker statistics, the system updates the Top Talker data ranks either every 300 seconds (for a 24-hour reporting period), or 600 seconds (for a 48-hour reporting period).</p> <p>The system saves a maximum of 300 Top Talker data snapshots, and aggregates these to calculate the top talkers for the 24- or 48-hour reporting period.</p> <p>The system never clears top talker data at the time of polling; however, every 300 or 600 seconds, it replaces the oldest Top Talker data snapshot of the 300 with the new data snapshot.</p> <p>After you change the reporting period, it takes the system one day to update the Top Talker rankings to reflect the new reporting period. In the interim, the data used to calculate the Top Talkers still includes data snapshots from the original reporting period. This delay applies to Top Talker report queries and SNMP Top Talker statistics.</p>



3. Click **Apply** to apply your settings.
4. Click **Save** to save your settings permanently.

## To enable flow export settings

1. Choose Configure > Networking > Flow Statistics to display the Flow Statistics page.
2. Under Flow Export Settings, complete the configuration as described in this 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.
Export QoS and Application Statistics to CascadeFlow Collectors	<p>Sends application-level statistics from all sites to a CascadeFlow collector on a Cascade appliance. Cascade appliances provide central reporting capabilities. The collector aggregates QoS and application statistics to provide visibility using detailed records specific to flows traversing the Steelhead appliance.</p> <p>The Steelhead appliance sends the Cascade appliances an enhanced version of NetFlow called CascadeFlow. CascadeFlow includes:</p> <ul style="list-style-type: none"> <li>• NetFlow v9 extensions for round-trip time measurements that enable you to understand volumes of traffic across your WAN and end-to-end response time.</li> <li>• extensions that enable a Cascade Express appliance to properly measure and report on the benefits of optimization.</li> </ul> <p>After the statistics are aggregated on a Cascade appliance, you can use its central reporting capabilities to:</p> <ul style="list-style-type: none"> <li>• analyze overall WAN use, such as traffic generated by application, most active sites, and so on.</li> <li>• troubleshoot a particular application by viewing how much bandwidth it received, checking for any retransmissions, interference from other applications, and so on.</li> <li>• compare actual application use against your outbound QoS policy configuration to analyze whether your policies are effective. For example, if your QoS policy determines that Citrix should get a minimum of 10 percent of the link, and the application statistics reveal that Citrix performance is unreliable and always stuck at 10 percent, you might want to increase that minimum guarantee.</li> </ul> <p>You must enable outbound QoS on the Steelhead appliance, add a CascadeFlow collector, and enable REST API access before sending QoS configuration statistics to an Enterprise Profiler.</p> <p>To enable QoS, choose Configure &gt; Networking &gt; Outbound Qos (Basic or Advanced). You cannot export statistics for inbound QoS.</p> <p>The collectors appear in the Flow Collector list at the bottom of the Configure &gt; Networking &gt; Flow Statistics page.</p> <p>To enable REST API access, choose Configure &gt; Security &gt; REST API Access.</p> <p>The CascadeFlow collector collects read-only statistics on both pass-through and optimized traffic. When you use CascadeFlow, the Steelhead appliance sends four flow records for each optimized TCP session: ingress and egress for the inner-channel connection, and ingress and egress for the outer-channel connection. A pass-through connection still sends four flow records, even though there are no separate inner- and outer-channel connections. In either case, the Cascade Express appliance merges these flow records together with flow data collected for the same flow from other devices.</p> <p>For details, see the <i>Cascade Deployment Guide</i>.</p>

Control	Description
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.  You can set the time-out period even if the Top Talkers option is enabled.
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 99](#)
- [“Viewing Top Talkers Reports” on page 281](#)
- [“Viewing Application Statistics Reports” on page 289](#)

### To add a Flow collector

1. Under Flow Collectors, complete the configuration as described in this 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	<p>Select one of these versions from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>CascadeFlow</b> - Use with Cascade Profiler v8.4 or later.</li> <li>• <b>CascadeFlow-compatible</b> - Use with Cascade Profiler v8.3.2 or earlier, and select the LAN Address check box.</li> <li>• <b>NetFlow v9</b> - Enables both ingress and egress flow records.</li> <li>• <b>NetFlow v5</b> - Enables ingress flow records.</li> </ul> <p>For details on using NetFlow records with Cascade, see the <i>Cascade Deployment Guide</i>.</p> <p>CascadeFlow and CascadeFlow-compatible are enhanced versions of flow export to the Cascade Product Suite. These versions allow automatic discovery and interface grouping for Steelhead appliances in a Riverbed Cascade Enterprise Profiler appliance or a Cascade Gateway appliance and support WAN and optimization reports in Cascade appliances. For details, see the <i>Cascade Profiler and Cascade Express User's Guide</i> and the <i>Cascade Sensor and Cascade Gateway User's Guide</i>.</p>
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.

Control	Description
LAN Address	<p>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.</p> <p>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.</p>
Capture Interface/Type	<p>Specify the traffic type to export to the flow collector. Select one of these types from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>All</b> - Exports both optimized and nonoptimized traffic.</li> <li>• <b>Optimized</b> - Exports optimized traffic.</li> <li>• <b>Optimized</b> - Exports optimized LAN or WAN traffic when WCCP is enabled.</li> <li>• <b>Passthrough</b> - Exports pass-through traffic.</li> <li>• <b>None</b> - Disables traffic flow export.</li> </ul> <p>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.</p>
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.
Add	Adds the collector to the Collector list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

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

## 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, `-i aux 1.1.1.1` where 1.1.1.1 is the NetFlow collector and aux is the Packet Source Interface).
- Verify that your capture settings are on the correct interface and that traffic is flowing through it.

---

## Applying QoS Policies

This section describes how to set Riverbed Quality of Service (QoS) policies. It includes these sections:

- [“QoS Overview” on page 107](#)
- [“Configuring Outbound QoS \(Basic\)” on page 115](#)
- [“Configuring Outbound QoS \(Advanced\)” on page 128](#)
- [“Configuring Inbound QoS” on page 144](#)

You apply Riverbed QoS policies in the Configure > Networking > Outbound QoS (Basic), Outbound QoS (Advanced), and Inbound QoS pages. This section describes how Steelhead appliances use Riverbed QoS policies to allocate bandwidth and latency priorities.

---

**Note:** For details about QoS, including integrating Steelhead appliances into an existing QoS implementation, see the *Steelhead Appliance Deployment Guide*. The *Steelhead Appliance Deployment Guide* also includes configuration examples and Riverbed QoS best practices.

---

## QoS Overview

QoS is a reservation system for network traffic. In its most basic form, QoS allows organizations to allocate scarce network resources across multiple traffic types of varying importance. Advanced QoS implementations allow organizations to accurately control their applications by the amount of bandwidth they have access to and by their sensitivity to delay.

RiOS and QoS provide the following features:

- **SnapMirror Support** - Use outbound QoS to prioritize SnapMirror replication jobs or shape optimized SnapMirror traffic that is sharing a WAN link with other enterprise protocols. QoS recognizes SnapMirror optimized flows and provisions five different service levels for each packet, based on priorities. You can also distinguish a job priority by filer and volume. Using advanced QoS, you can create a QoS rule for the appropriate site and optionally specify a service class and DSCP marking per priority.
- **Export QoS Configuration Statistics to a CascadeFlow Collector** - CascadeFlow collectors can aggregate information about QoS configuration and other application statistics to send to a Cascade Enterprise Profiler appliance. The Enterprise Profiler summarizes and displays the QoS configuration statistics. For details, see [“Configuring Flow Statistics” on page 101](#).
- **LAN Bypass** - Virtual in-path network topologies in which the LAN-bound traffic traverses the WAN interface might require that you configure the Steelhead appliance to bypass LAN-bound traffic so it is not subject to the maximum root bandwidth limit. RiOS includes LAN bypass for QoS outbound shaping as well as inbound QoS shaping.
- **Host Label Handling** - Specify a range of hostnames and subnets within a single QoS rule.
- **Global DSCP Marking** - By default, the setup of optimized connections and the out-of-band control connections are not marked with a DSCP value. Existing traffic marked with a DSCP value is classified into the default class. If your existing network provides multiple classes of service based on DSCP values, and you are integrating a Steelhead appliance into your environment, you can use the Global DSCP feature to prevent dropped packets and other undesired effects.

- **QoS with IPv6** - RiOS does not support IPv6 traffic for QoS shaping or AFE-based classification. If you enable QoS shaping for a specific interface, all IPv6 packets for that interface are classified to the default class. You can mark IPv6 traffic with an IP TOS value. You can also configure the Steelhead appliance to reflect an existing traffic class from the LAN-side to the WAN-side of the Steelhead appliance.
- **PC-over-IP (PCoIP)** - The Application Flow Engine (AFE) recognizes PCoIP flows and supports provisioning four different service levels for each packet based on PCoIP packet priorities: priorities 0-3 (lowest), priority 4, priority 5, and priorities 6-7 (highest).

PCoIP is a proprietary remote workstation and desktop protocol designed and developed by Teradici. PCoIP compresses and encrypts display updates (pixels only) along with keyboard and mouse events over the wire and enables remote access to workstations and servers from a remote location. Access to these machines is typically performed through either a thin client, a zero client, or a desktop application. Zero clients are ultra-thin clients in which the client mechanism is implemented in hardware or firmware.

- **Substantial Increase in Applications Recognized by the AFE** - The AFE recognizes over 600 application signatures in basic, advanced, and inbound QoS, providing an efficient and accurate way to identify applications for advanced classification and shaping of network traffic. To view the predefined global application list, see [“List of Recognized Applications” on page 365](#)
- **LAN Bypass** - Virtual in-path network topologies in which the LAN-bound traffic traverses the WAN interface might require that you configure the Steelhead appliance to bypass LAN-bound traffic so it is not subject to the maximum root bandwidth limit. Some deployment examples are WCCP or a WAN-side default gateway. The LAN bypass feature allows you to exempt certain subnets from QoS enforcement. For details, see [“Bypassing LAN Traffic” on page 113](#).
- **Bandwidth Enforcement** - Issues a warning when you configure the sum of the bandwidth interfaces as a value greater than the model-specific QoS limits. The warning appears when you save a configured bandwidth limit that exceeds the supported limit. If you receive the warning, adjust the sum of the configured QoS interface values to be lower or equal to the model-specific bandwidth limit and save the configuration.  
  
Riverbed strongly recommends configuring the bandwidth at or below the appliance limit, as problems arise when you exceed it.
- **Inbound QoS** - Allocates bandwidth and prioritizes traffic flowing into the LAN network behind the Steelhead appliance. This provides the benefits of QoS for environments that cannot meet their QoS requirements with outbound QoS. Both the client-side and server-side Steelhead appliances must be running RiOS v7.0 or later. For details, see [“Configuring Inbound QoS” on page 144](#).
- **Steelhead appliance DX series QoS limits** - Steelhead appliance models in the DX series enforce limits on QoS bandwidth and the maximum number of classes, rules, and sites. For details, see [“QoS DX8000 Series Recommendations” on page 110](#).
- More parameters to classify traffic into different marking values. These parameters allow you to classify using header parameters such as VLAN, DSCP, and protocols, and Application Flow Engine (AFE) inspection.
- Instead of using separate rules tables for DSCP marking, marking and shaping are combined in the same table. You can specify a marking value in either a service class or in a rule. When you specify a marking value in a rule, that value takes precedence over the value in a class.

- For optimized traffic, the server-side Steelhead appliance uses its local rules table for QoS marking instead of relying on the rules configured on the client-side Steelhead appliance.

---

**Note:** If you plan to configure QoS marking *without* QoS shaping, set the marking values in the application rules table. When you specify a marking value in a rule, that value takes precedence over the value in a class. If you plan to configure QoS marking *and* QoS shaping, set the marking values in the service policies or classes.

---

- **Simplified Setup** - A basic QoS configuration page streamlines setup for networks that require minimal configuration of network traffic.
- **Application Flow Engine (AFE)** - Allows advanced classification and shaping of network traffic. The Steelhead appliance 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. After an application is identified, you can place it 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 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 to create endless combinations of Layer-7 applications.

To view the predefined global application list, see [“List of Recognized Applications” on page 365](#).

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 the Application Flow Engine.

AFE works with both pass-through and optimized traffic and is available in basic outbound QoS, advanced outbound QoS, and inbound QoS.

- **Port Label Handling** - 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** - Increase the number of rules per site, up to 2000.
- **Site Awareness** - Improves performance and scalability in multisite configurations (outbound QoS only).

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, Voice over IP (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, RDP, and CRM systems) with lower priority, high-bandwidth traffic (for example, backup and replication). RiOS QoS allows you to protect delay-sensitive traffic such as VoIP, as well as other delay-sensitive traffic. 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.

---

**Important:** Flows can be incorrectly classified if there are asymmetric routes in the network when any of the QoS features are enabled.

---

## QoS DX8000 Series Recommendations

Riverbed recommends the maximum bandwidth, class, rules, and sites for the Steelhead appliance DX8000 series shown in this table. The QoS bandwidth recommendations are global across all WAN interfaces and the primary interface.

Traffic that passes through a Steelhead DX but is not destined to the WAN is not subject to these QoS recommendations. Examples of traffic that is not subject to QoS recommendations include routing updates, DHCP requests, and default gateways on the WAN-side of the Steelhead DX that redirect traffic back to other LAN-side subnets.

Steelhead Appliance	Maximum Configurable Root Bandwidth (Mbps)	Recommended Maximum Classes	Recommended Maximum Rules	Recommended Maximum Sites
DX8000	No limit	2,000	2,000	200



## Basic or Advanced Outbound QoS

RiOS provides two types of outbound 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 appliance for the first time.

Advanced outbound QoS supports different bandwidths for different interfaces; basic QoS does not, but you can specify the remote site absolute bandwidth.

Use basic outbound QoS when you:

- currently do not have RiOS QoS configured.
- do not need more granular control and can use the default settings.

Use advanced outbound QoS when you:

- must use the MX-TCP queue. For details, see [“Enabling MX-TCP Queue Policies \(Advanced Outbound QoS only\)” on page 143](#).
- 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 outbound QoS comes with six predefined classes, and you cannot add or delete classes. In advanced outbound QoS and inbound 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 outbound QoS 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 shaping 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 about QoS classes, see the *Steelhead Appliance Deployment Guide*.

---

## Hierarchical Mode (Advanced Outbound QoS)

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, these 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 this 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

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.

Basic outbound and inbound QoS always use flat mode. Advanced outbound QoS can use either flat or hierarchical mode.

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 *Riverbed Central Management Console User's Guide* and the *Steelhead Appliance Deployment Guide*.

---

## Selecting an Outbound QoS Enforcement System

Selecting the appropriate QoS enforcement system depends on the location of WAN bottlenecks for traffic leaving the site.

Use these guidelines when implementing advanced outbound 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.

As an 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 outbound 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 about configuring QoS for a branch office and data center, see the *Steelhead Appliance Deployment Guide*.

---


## Bypassing LAN Traffic

RiOS imposes a maximum limit on the configurable root bandwidth for the WAN interface. The hardware platform determines the maximum limit. The bandwidth limit prevents you from configuring WAN interfaces with a bandwidth rate that exceeds the hardware model limits.


Certain virtual in-path network topologies where the LAN-bound traffic traverses the WAN interface might require that the Steelhead appliance bypass LAN-bound traffic so that it is not included in the rate limit determined by the maximum root bandwidth. Some deployment examples are WCCP or a WAN-side default gateway.

Figure 5-2 and Figure 5-3 illustrate topologies where the default LAN gateway or router is accessible over the WAN interface of the Steelhead appliance. If there are two clients in the local subnet, traffic between the two clients is routable after reaching the LAN gateway. As a result, this traffic traverses the WAN interface of the Steelhead appliance.

**Figure 5-2. In-path Configuration Where Default LAN Gateway is Accessible Over the Steelhead Appliance WAN Interface**



**Figure 5-3. WCCP Configuration Where Default LAN Gateway is Accessible Over the Steelhead Appliance WAN Interface**



In a QoS configuration for these topologies, suppose you have several QoS classes created and the root class is configured with the WAN interface rate. The remainder of the classes use a percentage of the root class. In this scenario, the LAN traffic is rate limited because RiOS classifies it into one of the classes under the root class.

You can use the LAN bypass feature to exempt certain subnets from QoS enforcement, bypassing the rate limit. In RiOS, the LAN bypass feature works when enabled on both inbound and outbound QoS.

### To filter the LAN traffic from the WAN traffic

1. If QoS is not running, choose **Configure > Networking > Inbound or Outbound QoS** and enable the feature.
2. Choose **Configure > Networking > Subnet Side Rules**.
3. Click **Add a Subnet Side Rule**.
4. Select **Start**, **End**, or a rule number from the drop-down list.
5. Specify the client-side Steelhead appliance subnet using the format **<IP address>/<subnet mask>**.
6. Select **Subnet address is on the LAN side of the appliance**.
7. Click **Add**.

To verify the traffic classification, choose Reports > Networking > Inbound QoS or Outbound QoS.

---

**Note:** The Steelhead appliance processes the subnet side LAN rules before the QoS outbound rules.

---

---

**Note:** In virtual-in-path deployment, using subnet side rules is the same for QoS, RSP, and NetFlow. In an in-path deployment RSP and NetFlow discard the subnet side rules.

---

---

## Configuring Outbound QoS (Basic)

This section describes how to configure basic outbound QoS. It contains these sections:

- [“Overview” on page 115](#)
- [“Adding a Remote Site” on page 117](#)
- [“Adding an Application” on page 121](#)
- [“Adding a Service Policy” on page 125](#)

For a QoS overview, see [“Applying QoS Policies” on page 107](#). For information on whether to deploy basic outbound or advanced outbound QoS, see [“Basic or Advanced Outbound QoS” on page 111](#).

### Overview

Basic outbound QoS simplifies QoS configuration by accurately identifying business applications and classifying traffic according to priorities. The Steelhead appliance 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 outbound QoS can prevent recreational applications from interfering with business applications.

Basic outbound 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, see [“List of Recognized Applications” on page 365](#).


Basic outbound 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 outbound QoS.

## Enabling Local WAN Oversubscription

Basic outbound 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 outbound QoS.

Figure 5-4. Bandwidth Oversubscription Feature in Outbound QoS (Basic)



### To enable basic outbound QoS

1. Choose Configure > Networking > Outbound QoS (Basic) to display the Outbound QoS (Basic) page.
2. Under QoS Settings, complete the configuration as described in this table.

Control	Description
Enable QoS Shaping	Enables QoS classification to control the prioritization of different types of network traffic and to ensure that the Steelhead appliance 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.
WAN Bandwidth (kbps)	Specify the interface bandwidth link rate in kbps. The Steelhead appliance automatically sets the bandwidth for the default site to this value.  The link rate is the <i>bottleneck</i> WAN bandwidth, not the interface speed out of the WAN interface into the router or switch. As an 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).  <b>Important:</b> Different WAN interfaces can have different WAN bandwidths; you must enter the bandwidth link rate correctly for QoS to function properly.
Enable QoS on <interface>	Specify a WAN interface <X-Y> to enable.

Control	Description
Enable Local WAN Oversubscription	<p>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.</p> <p>As an example, your data center uplink might be 45 Mbit/s with three remote office sites each with 20 Mbit/s uplinks.</p> <p>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.</p> <p><b>Note:</b> Enabling this option can degrade latency guarantees when the remote sites are fully active.</p>
Enable QoS Marking	<p>Identify traffic using marking values. You can mark traffic using header parameters such as VLAN, DSCP, and protocols. In RiOS 7.0, you can also use Layer-7 protocol information through Application Flow Engine (AFE) inspection to apply DSCP marking values to traffic flows.</p> <p>In RiOS v7.0 and later, the DSCP or IP TOS marking only has local significance. This means you can set the DSCP or IP TOS values on the server-side Steelhead appliance to values different to those set on the client-side Steelhead appliance.</p>
Global DSCP	<p>By default, RiOS does not mark the setup of optimized connections and the out-of-band control connections with a DSCP value. The system classifies existing traffic marked with a DSCP value into the default class. If your existing network provides multiple classes of service based on DSCP values, and you are integrating a Steelhead appliance into your environment, you can select a global DSCP setting from the drop-down list to prevent dropped packets and other undesired effects.</p> <p>You can enable global DSCP marking without enabling QoS shaping.</p>

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 adding a custom application or adjusting the class bandwidth allocations as described in the following sections. When you finish configuring basic outbound QoS, select the Applications tab to make sure the applications belong to the desired class.

## 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 overall maximum number of basic outbound QoS rules is 2000. For details about the maximum number of rules and sites for a Steelhead appliance DX model, see [“QoS DX8000 Series Recommendations” on page 110](#).

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 Steelhead appliances, choose Configure > Networking > Outbound QoS (Basic) to display the Outbound QoS (Basic) page and select the Sites tab.
2. Complete the configuration as described in this table.

Control	Description
Add Site	Displays the controls to define a remote site.
Name	Specify the site name: for example, data center.
Position	Select Start, End, or a number from the drop-down list to rank the site in the QoS index.  <b>Note:</b> The default site, which is tied to the Medium Office policy, cannot be removed and is always listed last.
Remote Subnet	Specify a site subnet. You can define a maximum of 50 subnets per site.  <b>Note:</b> You cannot edit the subnet for the default site.
Remote Link Bandwidth	Specify the maximum WAN bandwidth in kbps.
Service Policy	Optionally, select a service policy from the drop-down list. The default policy is Large Office.
<b>Apply These QoS Settings:</b>	
Service Class	<p>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):</p> <ul style="list-style-type: none"> <li>• <b>Real-Time</b> - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, or video conferencing.</li> <li>• <b>Interactive</b> - Specifies an interactive traffic class: for example, Citrix, RDP, telnet and ssh.</li> <li>• <b>Business Critical</b> - Specifies the high priority traffic class: for example, Thick Client Applications, ERPs, and CRMs.</li> <li>• <b>Normal Priority</b> - Specifies a normal priority traffic class: for example, Internet browsing, file sharing, and email.</li> <li>• <b>Low Priority</b> - Specifies a low priority traffic class: for example, FTP, backup, replication, other high-throughput data transfers, and recreational applications such as audio file sharing.</li> <li>• <b>Best Effort</b> - Specifies the lowest priority.</li> </ul> <p>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.</p> <p><b>Important:</b> The service class describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how <i>important</i> 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.</p>



Control	Description																				
DSCP	<p>Select Inherit from Service Class, a DSCP value from 0 to 63, or all from the drop-down list. This value is required when you enable QoS marking. The default setting is Inherit from Service Class.</p> <p>When you specify a DSCP marking value in a rule, it takes precedence over the value in a service class.</p> <p>The DSCP marking values fall into these classes:</p> <ul style="list-style-type: none"><li>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</li><li>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</li></ul> <table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP 22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table> <ul style="list-style-type: none"><li>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</li></ul>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38																	
Apply These Path Selection Settings	<p>RiOS applies these settings to a flow after it has started but before the AFE has enough information to assign it to a global application (at which point, RiOS uses the global application settings). The system also selects these paths for any flows that do not match an explicit rule.</p>																				
Path 1	<p>Select the predefined path to use as the primary path for the traffic flow or when more AFE is necessary in a secure deployment. The system selects this path for every connection setup, unless it goes down.</p> <p>Selecting a path enables the corresponding DSCP setting.</p> <p>You must specify a distinct end-to-end primary path for each site.</p>																				
Path 2	<p>Select the second predefined path to use as the secondary path if the first path goes down or when more AFE is necessary.</p>																				
Path 3	<p>Select the third predefined path to use as the tertiary path if the first and second paths go down or when more AFE is necessary.</p>																				

Control	Description																				
DSCP	<p>Select Inherit from Application, a DSCP value from 0 to 63, or all from the drop-down list. This value is required when you enable QoS marking. The default setting is Inherit from Service Class.</p> <p>When you specify a DSCP marking value in a rule, it takes precedence over the value in a service class.</p> <p>The DSCP marking values fall into these classes:</p> <ul style="list-style-type: none"><li>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</li><li>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</li></ul> <table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP 22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table> <ul style="list-style-type: none"><li>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</li></ul>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38																	
If paths are configured and all down	Optionally, select how the system handles packets if the default paths go down. These settings are available even when default paths are not set.																				
Relay Traffic	Select to send the packets on the original primary path. This is the default setting.																				
Drop Traffic	<p>Select to drop the packets. Dropping traffic is useful in the situation in which you prefer not to use bandwidth on the secondary (or tertiary) paths in case of failure on the primary path.</p> <p>You do not have to define default paths to drop specific traffic flows; however, you must enable Path Selection on the Configure &gt; Networking &gt; Path Selection page.</p>																				
Add	<p>Adds the site to the list. The Management Console redisplay the Sites table and applies your modifications to the running configuration, which is stored in memory.</p> <p>This button is dimmed until you enter the WAN bandwidth.</p>																				
Remove Site	Select the check box next to the name and click <b>Remove Site</b> .																				
Move Site	Moves the selected sites. Click the arrow next to the desired rule position; the site moves to the new position.																				

3. 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 the basic QoS applications table.

To view the predefined global application list, see [“List of Recognized Applications” on page 365](#).

### To define custom applications or edit existing application definitions

1. Choose Configure > Networking > Outbound QoS (Basic) to display the Outbound QoS (Basic) page.
2. Select the Applications tab.

3. To define a custom application and add it to the application table list on the Applications page, complete the configuration as described in this table.

Control	Description
Add Application	Displays the controls to define an application.
Name	Specify the application name, for example, Outlook Anywhere.
Description	Optionally, describe the application.
Position	Select Start, End, or a number from the drop-down list to rank the application in the QoS index.
<b>For Traffic with the Following Characteristics:</b>	
Local Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Remote Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Protocol	<p>Select All, TCP, UDP GRE, ICMP, IPSec AH, or IPSec ESP from the drop-down list.</p> <p>The default setting is All.</p>
VLAN Tag ID	<p>Optionally, specify a VLAN tag as follows:</p> <ul style="list-style-type: none"> <li>Specify a numeric VLAN tag identification number from 0 to 4094.</li> <li>Specify all to specify the rule applies to all VLANs.</li> <li>Specify none to specify the rule applies to untagged connections.</li> </ul> <p>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.</p>

Control	Description
DSCP	Optionally, specify a DSCP value from 0 to 63, or all to use all DSCP values.
Traffic Type	Select Optimized, Passthrough, or All from the drop-down list. The default setting is All.
Application	<p>Select an application from the drop-down list. To narrow the search start typing the first characters in the application name. When the application name and definition appears, select it from the list.</p> <p>The control is case-sensitive; if you type the entire application name you must match the case as it appears in the list.</p> <p>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.</p>
<b>Apply These QoS Settings:</b>	
Service Class	<p>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):</p> <ul style="list-style-type: none"> <li>• <b>Real-Time</b> - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, or video conferencing.</li> <li>• <b>Interactive</b> - Specifies an interactive traffic class: for example, Citrix, RDP, telnet and ssh.</li> <li>• <b>Business Critical</b> - Specifies the high priority traffic class: for example, Thick Client Applications, ERPs, and CRMs.</li> <li>• <b>Normal Priority</b> - Specifies a normal priority traffic class: for example, Internet browsing, file sharing, and email.</li> <li>• <b>Low Priority</b> - Specifies a low priority traffic class: for example, FTP, backup, replication, other high-throughput data transfers, and recreational applications such as audio file sharing.</li> <li>• <b>Best Effort</b> - Specifies the lowest priority.</li> </ul> <p>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.</p> <p><b>Important:</b> The service class describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how <i>important</i> 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.</p>

Control	Description																				
DSCP	<p>Select Inherit from Service Class, a DSCP value from 0 to 63, or all from the drop-down list. This value is required when you enable QoS marking. The default setting is Inherit from Service Class.</p> <p>When you specify a DSCP marking value in a rule, it takes precedence over the value in a service class.</p> <p>The DSCP marking values fall into these classes:</p> <ul style="list-style-type: none"><li>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</li><li>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</li></ul> <table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP 22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table> <ul style="list-style-type: none"><li>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</li></ul>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38																	
Apply These Path Selection Settings	RiOS applies these paths to a flow after it has started but before the AFE has enough information to assign it to a global application (at which point, RiOS uses the global application settings). It also uses the paths for any flows not matching an explicit rule.																				
Path 1	Select the predefined path to use when more AFE is necessary. The system uses this path for every connection setup unless it goes down. Selecting a path enables the corresponding DSCP setting.																				
Path 2	Select the second predefined path to use as the secondary path if the first path goes down, or when more AFE is necessary.																				
Path 3	Select the third predefined path to use as the tertiary path if the first and second paths go down or when more AFE is necessary.																				
If paths are configured and all down	Optionally, select how the system handles packets if the default paths go down. These settings are available even when default paths are not set.																				
Relay Traffic	Select to send the packets on the original primary path. This is the default setting.																				
Drop Traffic	<p>Select to drop the packets. Dropping traffic is useful in the situation in which you prefer not to use bandwidth on the secondary (or tertiary) paths in case of failure on the primary path.</p> <p>You do not have to define default paths to drop specific traffic flows; however, you must enable Path Selection on the Configure &gt; Networking &gt; Path Selection page.</p>																				
Add	Adds the rule to the list. The Management Console redisplay the Applications table and applies your modifications to the running configuration, which is stored in memory.																				

Control	Description
Remove Application	Select the check box next to the name and click <b>Remove Application</b> .
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 > Outbound QoS (Basic) to display the Outbound QoS (Basic) page.
2. Select the Service Policies tab.

Figure 5-5. Outbound QoS (Basic) Page - Service Policies

The screenshot shows the 'Service Policies' tab in the Outbound QoS (Basic) configuration page. At the top, there are tabs for 'Sites', 'Applications', and 'Service Policies'. Below the tabs, there are buttons for 'Add Service Policy' and 'Remove Service Policy'. The 'Add Service Policy' form includes a 'Name' field, 'Shaping Parameters' (Minimum and Maximum Bandwidth % for Realtime, Interactive, Business-Critical, Normal, Low-Priority, and Best-Effort), and 'Marking Parameters' (DSCP for the same categories). Below the form is an 'Add' button. At the bottom, there is a table of existing service policies.

<input type="checkbox"/>	Name ↑↓	Realtime	Interactive	Business-Critical	Normal	Low-Priority	Best-Effort	Associated Sites
<input type="checkbox"/>	Large_Office	10-100% Reflect	10-100% Reflect	20-100% Reflect	50-100% Reflect	9-100% Reflect	1-100% Reflect	
<input type="checkbox"/>	Larger_Office	20-100% Reflect	20-100% Reflect	20-100% Reflect	20-100% Reflect	19-100% Reflect	1-100% Reflect	
<input type="checkbox"/>	Medium_Office	10-100% Reflect	20-100% Reflect	20-100% Reflect	40-100% Reflect	9-100% Reflect	1-100% Reflect	Default-Site
<input type="checkbox"/>	Small_Office	20-100% Reflect	20-100% Reflect	30-100% Reflect	20-100% Reflect	9-100% Reflect	1-100% Reflect	
<input type="checkbox"/>	Smaller_Office	1-100% Reflect	1-100% Reflect	40-100% Reflect	40-100% Reflect	17-100% Reflect	1-100% Reflect	

3. Complete the configuration as described in this table.

Control	Description
Add Service Policy	Displays the controls to add a service policy.
Name	Specify the policy name: for example, New York Office.
Realtime	Specify the percentage to allocate for the minimum and maximum bandwidth. The minimum 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 minimum and maximum bandwidth.
Business-Critical	Specify the percentage to allocate for the minimum and maximum bandwidth.
Normal	Specify the percentage to allocate for the minimum and maximum bandwidth.
Low-Priority	This is the default service policy; specify the percentage to allocate for the minimum and maximum bandwidth.
Best Effort	Specify the percentage to allocate for the minimum and maximum bandwidth.
Add	Adds the service policy to the list. The Management Console redisplay 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 <b>Remove Service Policy</b> .



4. Under Marking Parameters, complete the configuration as described in this table:

Control	Description																				
DSCP	<div>Displays the controls to select a DSCP marking for each service policy.</div> <div>The DSCP marking values fall into these classes:</div> <div><div><div><div>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</div><div>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</div></div><table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table><div><div>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</div></div></div></div>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP22	AF33 DSCP 30	AF43 DSCP 38																	
Realtime	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Interactive	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Business-Critical	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Normal	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Low-Priority	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Best Effort	Select a DSCP value from 0 to 63, or Reflect (the default setting). Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.																				
Add	Adds DSCP marking to the service policy. The Management Console redisplay 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 <b>Remove Service Policy</b> .																				

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

### To modify the maximum bandwidth and bandwidth guarantees for a service policy

1. Choose Configure > Networking > Outbound QoS (Basic) to display the Outbound QoS (Basic) page.
2. Click the magnifying glass next to a policy name in the policy list and specify the guaranteed and maximum bandwidth percentage.

Figure 5-6. Basic Outbound QoS Page - Modifying a Service Policy

Name	Realtime	Interactive	Business-Critical	Normal	Low-Priority	Best-Effort	Associated Sites
Large_Office	10-100%, R	10-100%, R	20-100%, R	50-100%, R	9-100%, R	1-100%, R	
Larger_Office	20-100%, R	20-100%, R	20-100%, R	20-100%, R	19-100%, R	1-100%, R	
Medium_Office	10-100%, R	20-100%, R	20-100%, R	40-100%, R	9-100%, R	1-100%, R	Default-Site
Small_Office	20-100%, R	20-100%, R	30-100%, R	20-100%, R	9-100%, R	1-100%, R	
Smaller_Office	1-100%, R	1-100%, R	40-100%, R	40-100%, R	17-100%, R	1-100%, R	

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

### Related Topic

- [“Viewing Outbound QoS Reports” on page 273](#)

## Configuring Outbound QoS (Advanced)

You configure advanced outbound QoS in the Configure > Networking > Outbound QoS (Advanced) page. Advanced outbound QoS provides a greater degree of configurability than basic outbound QoS; for example, you can separate rules by sites and you can perform Application Flow Engine matching.

- If you are configuring QoS for the first time, you need to migrate from basic outbound QoS to advanced outbound QoS. For details, see [“Advanced Outbound QoS Steps” on page 129](#).
- If you are upgrading a Steelhead appliance with an existing QoS configuration running RiOS v6.1.x or earlier, the system automatically upgrades to advanced outbound QoS. For details, see [“Advanced Outbound QoS Steps” on page 129](#).

## Advanced Outbound QoS Steps

This table describes the steps for configuring advanced outbound QoS, followed by detailed procedures.

Task	Reference
1. Select each WAN interface and define the bandwidth link rate for each interface.	<a href="#">“To enable advanced outbound QoS” on page 129</a>
2. Select the Enable QoS Shaping check box.	<a href="#">“To enable advanced outbound QoS” on page 129</a>
3. Select either Flat or Hierarchical QoS.	<a href="#">“To enable advanced outbound QoS” on page 129</a>
4. Define the QoS classes for each traffic flow.	<a href="#">“To add an outbound QoS class” on page 130</a>
5. Add sites and define rules for each class or subclass.	<a href="#">“Adding a QoS Site or Rule for Outbound QoS (Advanced)” on page 136</a>

**Important:** If you delete or add new rules, the existing connections are not affected; the changes only effect new connections.

### To enable advanced outbound QoS

1. Choose Configure > Networking > Outbound QoS (Advanced) to display the Outbound QoS (Advanced) page.
2. Under QoS Settings, complete the configuration as described in this table.

Control	Description
Enable QoS Shaping	<p>Enables QoS 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.</p> <p>To disable QoS, clear this check box.</p>
Mode	<p>Specify a QoS structure:</p> <ul style="list-style-type: none"> <li>• Flat mode creates all classes at the same level.</li> <li>• 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.</li> </ul> <p><b>Note:</b> 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.</p> <p><b>Important:</b> 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 enable it only after you are ready for the changes to take effect.</p>

Control	Description
Enable QoS on <interface> with WAN Bandwidth (kbps)	<p>Enables a WAN interface &lt;X-Y&gt;. Specify its bandwidth link rate in kbps. The bandwidth for the default site is automatically set to this value.</p> <p>The link rate is the <i>bottleneck</i> 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).</p> <p><b>Important:</b> Different WAN interfaces can have different WAN bandwidths; you must enter the bandwidth link rate correctly for QoS to function properly.</p>
Enable QoS Marking	Identify traffic using marking values. You can mark traffic using header parameters such as VLAN, DSCP, and protocols. In RiOS 7.0, you can also use Layer-7 protocol information through Application Flow Engine (AFE) inspection to apply DSCP marking values to traffic flows.
Global DSCP	<p>By default, RiOS does not mark the set up of optimized connections and the out-of-band control connections with a DSCP value. The system classifies existing traffic marked with a DSCP value into the default class. If your existing network provides multiple classes of service based on DSCP values, and you are integrating a Steelhead appliance into your environment, you can select a global DSCP setting from the drop-down list to prevent dropped packets and other undesired effects.</p> <p>You can enable global DSCP marking without enabling QoS shaping.</p>

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.

## Creating QoS Classes

Priorities and bandwidths are set by QoS class. You can create multiple classes.

---

**Note:** For details about QoS, see the *Steelhead Appliance Deployment Guide*.

---

### To add an outbound QoS class

1. Choose Configure > Networking > Outbound QoS (Advanced) to display the Outbound QoS (Advanced) page.
2. Under QoS Classes, complete the configuration as described in this table.

Control	Description
Add a New Class	Displays the controls for adding a class.
Name	Specify a name for the QoS class.
<b>Shaping Parameters</b>	

Control	Description
Class Parent	<p>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.</p> <p>Select a class parent from the drop-down list.</p>

Control	Description
Queue	<p>Optionally, select one of these queue methods for the leaf class from the drop-down list (the queue does not apply to the inner class):</p> <ul style="list-style-type: none"> <li>• <b>SFQ</b> - 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.</li> <li>• <b>FIFO</b> - 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.</li> <li>• <b>MX-TCP</b> - Has very different use cases than the other queue parameters. MX-TCP also has secondary effects that you must understand before configuring: <ul style="list-style-type: none"> <li>– 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 guaranteed bandwidth configured on the QoS class.</li> <li>– 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. RiOS v8.5 introduces rate pacing for satellite deployments, which combines MX-TCP with a congestion-control method.</li> <li>– Another use 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.</li> </ul> <p>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 reclassification 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:</p> <ul style="list-style-type: none"> <li>• The QoS rule for MX-TCP is at the top of QoS rules list.</li> <li>• The rule does not use AFE identification.</li> <li>• You only use MX-TCP for optimized traffic. MX-TCP does not work for unoptimized traffic.</li> </ul> <p>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.</p> <p>The link share weight parameter does not apply to MX-TCP queues. When you select the MX-TCP queue, the link share weight parameter does not appear. In RiOS, there is a maximum bandwidth setting for MX-TCP that allows traffic in the MX class to burst to the maximum level if the bandwidth is available.</p> </li> </ul>

Control	Description
	<ul style="list-style-type: none"> <li>• <b>Packet-order</b> - 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.</li> </ul>
Latency Priority	<p>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):</p> <ul style="list-style-type: none"> <li>• <b>Real-Time</b> - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, video conferencing.</li> <li>• <b>Interactive</b> - Specifies an interactive traffic class: for example, Citrix, RDP, telnet and ssh.</li> <li>• <b>Business Critical</b> - Specifies the high priority traffic class: for example, Thick Client Applications, ERPs, and CRMs.</li> <li>• <b>Normal Priority</b> - Specifies a normal priority traffic class: for example, Internet browsing, file sharing, and email.</li> <li>• <b>Low Priority</b> - 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.</li> <li>• <b>Best Effort</b> - Specifies the lowest priority.</li> </ul> <p>These are minimum priority guarantees; if better service is available, it is provided. 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.</p> <p><b>Important:</b> The priority describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how <i>important</i> the traffic is compared to other classes. Typically, you configure low priority for high-throughput, nonpacket delay sensitive applications like FTP, backup, and replication.</p>
Minimum Bandwidth	<p>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.</p> <p>The minimum bandwidth must fall within the bandwidth limit for the Steelhead appliance.</p> <p>In hierarchical mode, excess bandwidth is allocated based on the relative ratios of minimum bandwidth. The total minimum guaranteed bandwidth of all QoS classes must be less than or equal to 100% of the parent class.</p> <p>A default class is automatically created with minimum bandwidth of 10%. Traffic that does not match any of the rules is put into the default class. Riverbed recommends that you change the minimum bandwidth of the default class to the appropriate value.</p> <p>You can adjust the value as low as 0%.</p> <p>The system rounds decimal numbers to the nearest hundredth.</p>

Control	Description
Link Share Weight	<p>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 bandwidth. By default, all the link shares are equal.</p> <p>Classes with a larger weight are allocated more of the excess bandwidth than classes with a lower link share weight.</p> <p>You cannot specify a Link Share Weight in Hierarchical QoS. In Hierarchical QoS, the link share weight is the same proportion as the minimum bandwidth of the class.</p> <p>The system rounds decimal numbers to the nearest hundredth.</p> <p>The link share weight does not apply to MX-TCP queues, so this control does not appear when you select the MX-TCP queue. In RiOS, there is a maximum bandwidth setting for MX-TCP that allows traffic in the MX class to burst to the maximum level if the bandwidth is available.</p>
Maximum Bandwidth	<p>Specify the maximum allowed bandwidth (as a percentage) a class receives as a percentage of the parent class minimum bandwidth. The limit is applied even if there is excess bandwidth available.</p> <p>The system rounds decimal numbers to the nearest hundredth.</p>
Connection Limit	<p>Optionally, specify the maximum number of optimized connections for the class. When the limit is reached, all new connections are passed through unoptimized.</p> <p>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.</p> <p>Connection Limit is supported only with in-path configurations. It is not supported in out-of-path or virtual in-path configurations.</p> <p>Connection Limit does not apply to the packet-order queue or Citrix ICA traffic.</p> <p>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 might subsequently reclassify the traffic flow after applying a more precise match using AFE identification.</p>
<b>Marking Parameters</b>	



Control	Description																				
DSCP	<p>Displays the controls to select a DSCP marking for each service policy.</p> <p>The DSCP marking values fall into these classes:</p> <ul style="list-style-type: none"><li>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</li><li>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</li></ul> <table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP 22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table> <ul style="list-style-type: none"><li>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</li></ul>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38																	
Add	Adds the QoS class.																				
Remove Selected	<p>Select the check box next to the name and click <b>Remove Selected</b>.</p> <p>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.</p>																				

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 outbound QoS configuration. If necessary, migrate to advanced outbound QoS mode, return to basic outbound QoS mode and press **Clear**.
2. Migrate to advanced outbound QoS.

3. Choose **Configure > Networking > Outbound QoS (Advanced)** to display the Outbound QoS (Advanced) 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 for Outbound QoS (Advanced)

Each rule maps a type of network traffic to a QoS class. You can create multiple QoS rules for a class. When multiple QoS rules are created for a class, the rules are followed in the order in which they are shown in the Outbound QoS (Advanced) 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, 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 outbound advanced QoS

1. Choose **Configure > Networking > Outbound QoS (Advanced)** to display the Outbound QoS (Advanced) page.
2. Under QoS Sites and Rules, complete the configuration as described in this table.

Control	Description
Add Site or Rule	Displays the controls to add a QoS site or rule.
Add a	Select either Site or Rule. The default is rule.
Name	Specify a rule or site name.
Subnets	Specify a site subnet. You can define a maximum of 50 subnets per site.

Control	Description
Description	Describe the 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	<p>Inserts a QoS rule for a QoS class. Select Start, End, or a rule number from the drop-down list.</p> <p>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.</p>
Local Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Remote Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Protocol	Select All, TCP, GRE, UDP, ICMP, IPSecAH (Authentication Header), IPSecESP (Encapsulating Security Payload), or a number from the drop-down list. All specifies all TCP and UDP-based protocols.
VLAN Tag ID	Optionally, specify the VLAN tag for the rule. The default value is all.
DSCP	<p>Optionally, select a DSCP level. The default value is all.</p> <p>In RiOS, the DSCP field in a QoS classification rule matches the DSCP value before the packet enters the Steelhead appliance, and the DSCP field in a QoS classification rule for pass-through traffic matches the DSCP value <i>before</i> DSCP marking rules are applied. The DSCP field in a QoS classification rule for optimized traffic matches the DSCP value <i>after</i> DSCP marking rules are applied; that is, it matches the post-marking DSCP value.</p>

Control	Description
Traffic Type	<p>Select All, Optimized, or Passthrough from the drop-down list. The system applies the QoS rules to optimized and pass-through (egress only) traffic.</p> <p><b>Note:</b> Session reliability (port 2598) is not supported with pass-through Citrix traffic.</p>
Application	<p>Select an application from the drop-down list of global applications. To narrow the search, type the first letter of the application.</p> <p>The control is case-sensitive; if you type the entire application name you must match the case as it appears in the list.</p> <p>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.</p> <p>Selecting SnapMirror expands the control to include priorities highest to lowest. Select a priority for the SnapMirror application to separate low-priority traffic from high-priority traffic. SnapMirror classification using a priority supports optimized traffic only.</p> <p>The SnapMirror control also includes a table for assigning service classes and DSCP markings to SnapMirror priorities. Select a service class and a DSCP mark to classify SnapMirror traffic differently based on priority. For SnapMirror traffic without a priority, RiOS uses the service class and DSCP mark from the first two fields below the section titled Apply these QoS Settings. The default settings are Inherit from Service Class.</p> <p>You must enable and configure SnapMirror optimization and assign priorities to filers, volumes, and qtrees for priority-based SnapMirror QoS rules to work.</p> <p>Selecting ICA expands the control to include priorities 0 - 3 and a DSCP marking value. 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. Under DSCP, select Inherit from Service Class, a DSCP value from 0 to 63, or all from the drop-down list. The default setting is Inherit from Service Class.</p> <p>Selecting PCoIP expands the control to include service class by packet priorities 0 - 3 (lowest), priority 4, priority 5, priorities 6 - 7 (highest), and a DSCP marking value based on the priority. Select a priority for the PCoIP application to separate low-priority traffic from high-priority traffic. PCoIP classification using a priority supports optimized and pass-through traffic. Choose either FIFO, SFQ, or packet-order for the queue (packet-order is not required for PCoIP). Under DSCP, select Inherit from Service Class, a DSCP value from 0 to 63, or all from the drop-down list. The default setting is Inherit from Service Class.</p>

Control	Description
Service Class	<p>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):</p> <ul style="list-style-type: none"> <li>• <b>Real-Time</b> - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, or video conferencing.</li> <li>• <b>Interactive</b> - Specifies an interactive traffic class: for example, Citrix, RDP, telnet and ssh.</li> <li>• <b>Business Critical</b> - Specifies the high priority traffic class: for example, Thick Client Applications, ERPs, and CRMs.</li> <li>• <b>Normal Priority</b> - Specifies a normal priority traffic class: for example, Internet browsing, file sharing, and email.</li> <li>• <b>Low Priority</b> - Specifies a low priority traffic class: for example, FTP, backup, replication, other high-throughput data transfers, and recreational applications such as audio file sharing.</li> <li>• <b>Best Effort</b> - Specifies the lowest priority.</li> </ul> <p>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.</p> <p><b>Important:</b> The service class describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how <i>important</i> 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.</p>
DSCP	<p>Select Inherit from Service Class, a DSCP value from 0 to 63, or all from the drop-down list. This value is required when you enable QoS marking. The default setting is Inherit from Service Class.</p> <p>When you specify a DSCP marking value in a rule, it takes precedence over the value in a service class.</p>
<b>Apply These Path Selections</b>	<p>RiOS applies these paths to a flow after it has started but before the AFE has enough information to assign it to a global application (at which point, RiOS uses the global application settings). It also uses the paths for any flows not matching an explicit rule.</p>
Path 1	<p>Select the predefined path to use when more AFE is necessary. Path selection uses this path for every connection setup unless it goes down. Selecting a path enables the corresponding DSCP setting.</p>
Path 2	<p>Select the second predefined path to use when more AFE is necessary, and the first path goes down.</p>
Path 3	<p>Select the third predefined path to use when more AFE is necessary, and the first and second paths go down.</p>

Control	Description																				
DSCP	<p>Select Inherit from Application, a DSCP value from 0 to 63, or all from the drop-down list. This value is required when you enable QoS marking. The default setting is Inherit from Service Class.</p> <p>When you specify a DSCP marking value in a rule, it takes precedence over the value in a service class.</p> <p>The DSCP marking values fall into these classes:</p> <ul style="list-style-type: none"><li>• <b>Expedited forwarding (EF) class</b> - In this class, packets are forwarded regardless of link share of other traffic. The class is suitable for preferential services requiring low delay, low packet loss, low jitter, and high bandwidth.</li><li>• <b>Assured forwarding (AF) class</b> - This class is divided into four subclasses, each containing three drop priorities for more granular classification. The QoS level of the AF class is lower than that of the EF class.</li></ul> <table><tr><th>Drop Priority</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr><tr><td>Low</td><td>AF11 DSCP 10</td><td>AF21 DSCP 18</td><td>AF31 DSCP 26</td><td>AF41 DSCP 34</td></tr><tr><td>Medium</td><td>AF12 DSCP 12</td><td>AF22 DSCP 20</td><td>AF32 DSCP 28</td><td>AF42 DSCP 36</td></tr><tr><td>High</td><td>AF13 DSCP 14</td><td>AF23 DSCP 22</td><td>AF33 DSCP 30</td><td>AF43 DSCP 38</td></tr></table> <ul style="list-style-type: none"><li>• <b>Class selector (CS) class</b> - This class is derived from the IP ToS field.</li></ul>	Drop Priority	Class 1	Class 2	Class 3	Class 4	Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34	Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36	High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38
Drop Priority	Class 1	Class 2	Class 3	Class 4																	
Low	AF11 DSCP 10	AF21 DSCP 18	AF31 DSCP 26	AF41 DSCP 34																	
Medium	AF12 DSCP 12	AF22 DSCP 20	AF32 DSCP 28	AF42 DSCP 36																	
High	AF13 DSCP 14	AF23 DSCP 22	AF33 DSCP 30	AF43 DSCP 38																	
If no configured path is up	Optionally, select how path selection handles packets if the default paths go down. These settings are available even when default paths are not defined.																				
Relay Traffic	Select to relay traffic from the interface. This is the default setting when the default paths are not specified.																				
Drop Traffic	<p>Select to drop the packets. Dropping traffic is useful in the situation in which you prefer not to use bandwidth on the secondary (or tertiary) paths in case of failure on the primary path.</p> <p>You do not have to define default paths to drop specific traffic flows; however, you must enable Path Selection on the Configure &gt; Networking &gt; Path Selection page.</p>																				
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 <b>Move QoS Rules</b> . 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, click the rule name. Enter the changes and click **Apply**.

## Verifying and Saving an Outbound QoS Configuration

After you apply your settings, you can verify whether the traffic is categorized in the correct class by choosing Reports > Networking > Outbound QoS 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 Outbound QoS and Outbound QoS 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 about saving configurations, see [“Managing Configuration Files” on page 194](#).

### Related Topics

- [“Configuring Port Labels” on page 75](#)
- [“Managing Configuration Files” on page 194](#)
- [“Viewing Outbound QoS Reports” on page 273](#)

## Modifying QoS Classes or Rules

You can modify QoS classes and rules in the Outbound QoS (Advanced) page.

### To modify a QoS class or rule

1. Choose Configure > Networking > Outbound QoS (Advanced) to display the Outbound QoS (Advanced) page.
2. Select the class or rule name in the appropriate table.

Figure 5-7. Editing a Class or Rule

**QoS Classes:**

+ Add a New Class    - Remove Selected

<input type="checkbox"/>	Name	Latency Priority	Min BW %	Max BW %	Conn Limit	Queue	DSCP
<input type="checkbox"/>	▼ Default-Site\$\$parent_class	Normal	100.00	100.00		sfq	Reflect
<input type="checkbox"/>	☑ Default-Site\$\$Best-Effort	Best-Effort	1.00	100.00		sfq	10 (AF11)
<p><b>Shaping Parameters:</b></p> <p>Queue: <input type="text" value="sfq"/></p> <p>Minimum Bandwidth: <input type="text" value="1.0"/> %</p> <p>Maximum Bandwidth: <input type="text" value="100.0"/> %</p> <p>Latency Priority: <input type="text" value="Best-Effort"/></p> <p>Connection Limit: <input type="text"/></p> <p><b>Marking Parameters:</b></p> <p>DSCP: <input type="text" value="10 (AF11)"/></p> <p><input type="button" value="Apply"/></p>							
<input type="checkbox"/>	▼ Default-Site\$\$Business-Critical	Business-Critical	20.00	100.00		sfq	11

3. 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 Outbound QoS Configuration to Return to Basic Outbound QoS

In certain situations, it might be necessary to revert from outbound 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 Outbound QoS (Advanced). Reverting to a previously saved configuration restores your entire configuration.

Reverting to basic outbound QoS without using a previous configuration deletes all your current QoS settings.

### To clear an advanced outbound QoS configuration and return to a blank basic QoS configuration

1. Choose **Configure > Networking > Outbound QoS (Basic)** to display the Outbound QoS (Basic) page.
2. Click **Clear**.
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 Outbound QoS (Basic) page appears.

You can now configure basic outbound QoS. For details, see [“Configuring Outbound QoS \(Basic\)” on page 115](#).

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

### To revert from advanced outbound QoS to a previously saved basic outbound QoS configuration

1. Choose **Configure > Networking > Outbound QoS (Basic)** to display the Outbound QoS (Basic) 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 outbound 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 Outbound 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. Some use case examples are:

- **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.

These exceptions to QoS classes apply to MX-TCP queues:

- There is a maximum bandwidth setting for MX-TCP that allows traffic to burst to the maximum level if the bandwidth is available.
- 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 Application Flow Engine. A traffic flow cannot be classified as MX-TCP and then later classified in a different queue. This reclassification 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 Application Flow Engine identification.

### Basic Steps for MX-TCP

This table describes the basic steps to configure MX-TCP. Enabling this feature is *optional*.

Task	Reference
1. Select either Flat or Hierarchical mode. <b>Note:</b> Selecting a mode does <i>not</i> enable QoS traffic classification. The Enable QoS Shaping and Enforcement check box must be selected and a bandwidth link rate must be set for each WAN interface before traffic optimization begins.	<a href="#">“Selecting an Outbound QoS Enforcement System” on page 113</a>  <a href="#">“To enable advanced outbound QoS” on page 129</a>
2. Select each WAN interface and define the bandwidth link rate for each interface.	<a href="#">“To enable advanced outbound QoS” on page 129</a>
3. Add an MX-TCP class for the traffic flow. Make sure you specify MX-TCP as your queue.	<a href="#">“To add an outbound QoS class” on page 130</a>
4. Define QoS rules to point to the MX-TCP class.	<a href="#">“Adding a QoS Site or Rule for Outbound QoS (Advanced)” on page 136</a>

Task	Reference
5. Select the Enable QoS Classification and Enforcement check box. Your changes take effect immediately.	<a href="#">“To enable advanced outbound QoS” on page 129</a>
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.	<a href="#">“Optimizing TCP and Satellite WANs” on page 59</a>
7. Check and locate the inner connection.	<a href="#">“Viewing Alarm Status Reports” on page 307</a>
8. Check the throughput.	<a href="#">“Viewing Current Connection Reports” on page 248</a>

## Configuring Inbound QoS

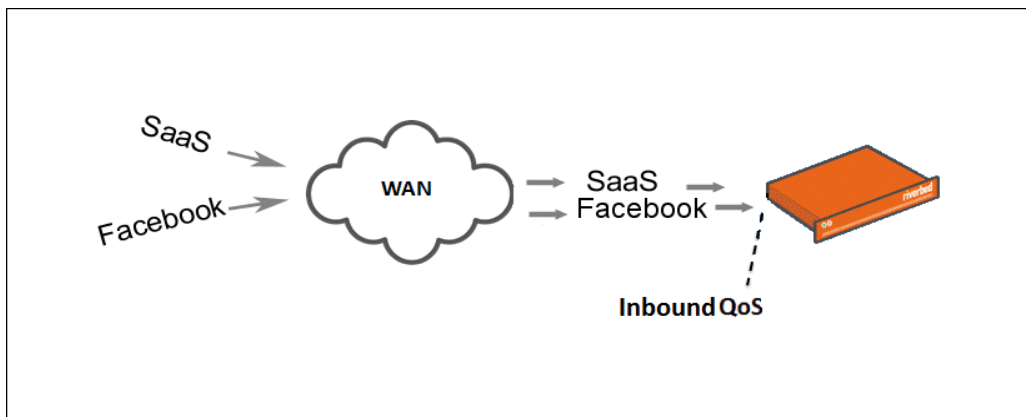
RiOS supports inbound QoS. You configure inbound QoS in the Configure > Networking > Inbound QoS page.

Inbound QoS allocates bandwidth and prioritizes traffic flowing into the LAN network behind the Steelhead appliance. This provides the benefits of QoS for environments that cannot meet their QoS requirements with outbound QoS.

Some examples of environments that can benefit from inbound QoS are:


- A deployment that does not have a Steelhead appliance located at the traffic source (for example, the traffic comes from the Internet, or from servers at a site without a Steelhead appliance).

**Figure 5-8. Guarantee Bandwidth for Incoming Traffic**




- A deployment that has multiple Steelhead appliances located at the traffic source (for example, behind an Interceptor cluster). The Steelhead appliances do not share bandwidth information with each other. As a result, they can overwhelm the branch office site at the remote location.

**Figure 5-9. Data Center with Multiple Steelhead Appliances in a Cluster**




- A branch office receiving data from multiple data centers (either with or without Steelhead appliances). Because the two data centers do not coordinate the amount of bandwidth sent to the branch office, they can overwhelm the link at the branch office, causing degraded performance.

**Figure 5-10. Branch Office that Receives Data From Multiple Data Centers**



Configuring inbound QoS focuses on prioritizing types of traffic using rules and classes just like outbound QoS. The inbound configuration is separate from the outbound configuration. You define the applications on the local network and then create their corresponding shaping policies.

**Figure 5-11. Inbound QoS Overview**



Inbound QoS applies the HFSC shaping policies to the ingress traffic. This addresses environments in which bandwidth constraints exist at the downstream location. When this occurs, the downstream Steelhead appliance (where inbound QoS is enabled) dynamically communicates the bandwidth constraints to the client transmitting the traffic. The client slows down the throughput and the traffic adheres to the configured inbound QoS rule. Inbound QoS, just like outbound QoS, is not a dual-ended Steelhead appliance solution. A single Steelhead appliance performing traffic shaping as needed to avoid network congestion controls inbound WAN traffic on its own.

For details about the HFSC queuing technology, see [“Traffic Classification” on page 110](#) and the *Steelhead Appliance Deployment Guide*.

## How a Steelhead Appliance Identifies and Shapes Inbound Traffic

Inbound rules define the types of traffic flowing into the branch office. As with outbound QoS, the rule can match the traffic based on VLAN, IP header values, TCP/UDP ports, and AFE information. As an example, you can ensure that the voice traffic on the WAN is reserved a fixed bandwidth and this traffic has a higher priority over the recreational Internet traffic.

Inbound classes shape the inbound traffic. The class configuration resembles a flat outbound QoS class configuration. An outbound QoS configuration describes remote sites and services/applications. Inbound QoS describes the local services/applications and how to shape the inbound traffic.

The inbound traffic shaping configuration includes a default shaping class. The QoS scheduler applies the built-in inbound default class constraints and parameters on traffic not placed in any other class by the configured QoS rules. The default shaping class has a 10 percent minimum bandwidth allocation and a 100 percent maximum bandwidth allocation. You cannot delete the default class; however, you can change its bandwidth allocations.

## Inbound QoS Limitations

These limitations apply to inbound QoS traffic shaping.

- Unlike Outbound QoS (Advanced), inbound QoS does not support hierarchical mode; it uses flat mode only. A deployment using inbound QoS will typically not require a hierarchical configuration to support multiple sites across the WAN. Deployment scenarios include configuring hierarchical mode using outbound QoS in the data center and configuring inbound QoS in the branch offices.
- Inbound QoS does not classify and shape traffic received from a peer Steelhead appliance when using the connection forwarding feature. In these configurations, the Steelhead appliance that first intercepts the traffic classifies and shapes it.
- Inbound QoS is not fully compatible with RSP when one or more virtual machines are connected to the in-path data flow. In such scenarios all traffic that enters any Virtual Machine connected to the in-path data flow falls into the default class.
- When packet-mode optimization is enabled, the QoS scheduler places UDP4 traffic into the MX-TCP class. All other traffic goes into the proper class.
- You cannot configure inbound QoS in an out-of-path deployment over a primary or auxiliary interface.
- Inbound QoS does not throttle certain flows such as MX-TCP and UDP bulk traffic flows; however it does provide bandwidth and latency reservation for them.

## Inbound QoS Limits

These limits apply to inbound QoS traffic shaping.

- The maximum number of inbound QoS rules is 500. The maximum number of inbound QoS classes is 200.

## Inbound QoS Steps

This table describes the steps for configuring inbound QoS, followed by detailed procedures.

Task	Reference
1. Select each WAN interface and define the bandwidth link rate for each interface.	<a href="#">“To enable inbound QoS” on page 148</a>
2. Select the Enable Inbound QoS Shaping and Enforcement check box.	<a href="#">“To enable inbound QoS” on page 148</a>
3. Define the QoS classes for each traffic flow.	<a href="#">“To add an inbound QoS class” on page 148</a>
4. Define rules for each class or subclass.	<a href="#">“Adding a QoS Rule (Inbound QoS)” on page 150</a>

**Important:** If you delete or add new rules, the existing connections are not affected; the changes only affect new connections.

### To enable inbound QoS

1. Choose Configure > Networking > Inbound QoS to display the Inbound QoS page.
2. Under WAN Link, complete the configuration as described in this table.

Control	Description
Enable Inbound QoS Shaping and Enforcement	<p>Enables QoS to control the prioritization of different types of inbound network traffic and to ensure that the Steelhead appliance 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.</p> <p>By default, inbound QoS classification is disabled.</p> <p>To disable inbound QoS, clear this check box.</p>
Enable QoS on <interface> with WAN bandwidth (kbps)	<p>Enables a WAN interface &lt;X-Y&gt;. Specify its bandwidth link rate in kbps. The bandwidth for the default site is automatically set to this value.</p> <p>Inbound QoS supports in-path interfaces only; it does not support primary or auxiliary interfaces.</p> <p>The link rate is the <i>bottleneck</i> 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).</p> <p><b>Important:</b> Different WAN interfaces can have different WAN bandwidths; you must enter the bandwidth link rate correctly for QoS to function properly.</p>

## Creating Inbound QoS Classes

Priorities and bandwidths are set by QoS class. You can create up to 200 inbound QoS classes.

### To add an inbound QoS class

1. Choose Configure > Networking > Inbound QoS to display the Inbound QoS page.

2. Under Inbound QoS Classes, complete the configuration as described in this table.

Control	Description
Add a Class	Displays the controls for adding a class.
Class Name	Specify a name for the QoS class.
Priority	<p>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):</p> <ul style="list-style-type: none"> <li>• <b>Real-Time</b> - Specifies real-time traffic class. Give this value to your highest priority traffic; for example, VoIP, video conferencing.</li> <li>• <b>Interactive</b> - Specifies an interactive traffic class: for example, Citrix, RDP, telnet and ssh.</li> <li>• <b>Business Critical</b> - Specifies the high priority traffic class: for example, Thick Client Applications, ERPs, and CRMs.</li> <li>• <b>Normal Priority</b> - Specifies a normal priority traffic class: for example, Internet browsing, file sharing, and email.</li> <li>• <b>Low Priority</b> - 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.</li> <li>• <b>Best Effort</b> - Specifies the lowest priority.</li> </ul> <p>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.</p> <p><b>Important:</b> The latency priority describes only the delay sensitivity of a class, not how much bandwidth it is allocated, nor how <i>important</i> the traffic is compared to other classes. Typically, you configure low latency priority for high-throughput, nonpacket delay sensitive applications like FTP, backup, and replication.</p>
Minimum Bandwidth	<p>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 percent. During contention for bandwidth, the class is guaranteed the amount of bandwidth specified. The class receives more bandwidth if there is unused bandwidth remaining.</p> <p>The minimum bandwidth must fall within the bandwidth limit for the Steelhead appliance.</p> <p>A default class is automatically created with minimum bandwidth of 10 percent. 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.</p> <p>You can adjust the value as low as 0 percent.</p> <p>The system rounds decimal numbers to the nearest hundredth.</p>
Maximum Bandwidth	<p>Specify the maximum allowed bandwidth (as a percentage) a class receives as a percentage of the minimum bandwidth. The limit is applied even if there is excess bandwidth available.</p> <p>The system rounds decimal numbers to the nearest hundredth.</p>



Control	Description
Link Share Weight	Specify the weight for the class. The link share weight determines how the excess bandwidth is allocated among sibling classes. Link share does not depend on the minimum 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.  The system rounds decimal numbers to the nearest hundredth.
Add	Adds the QoS class.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

3. Click **Apply** to apply your settings.
4. Click **Save** to save your settings permanently.

---

**Tip:** The inbound QoS classes appear in the inbound 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.

---

## Adding a QoS Rule (Inbound QoS)

Each rule maps a type of network traffic to a QoS class. You can create multiple QoS rules for a class. When multiple QoS rules are created for a class, the rules are followed in the order in which they appear in the Inbound QoS page and only the first matching rule is applied to the class.

The maximum number of inbound QoS rules is 500.

### To add an inbound QoS rule

1. Choose **Configure > Networking > Inbound QoS** to display the Inbound QoS page.



2. Under Inbound QoS Rules, complete the configuration as described in this table.

Control	Description
Add a Rule	Displays the controls to add a QoS rule.
Name	Specify a rule name.
Insert Rule At	<p>Inserts a QoS rule for a QoS class. Select Start, End, or a rule number from the drop-down list.</p> <p>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.</p>
Description	Describe the rule to facilitate administration.
Remote Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Local Subnet or Host Label	<p>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.</p> <p>Use this format: XXX.XXX.XXX.XXX/XX.</p> <p>—or—</p> <p>Specify a host label. You predefine host labels on the Configure &gt; Optimization &gt; Host Labels page.</p>
Port or Port Label	<p>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.</p> <p>—or—</p> <p>Specify a port label. You predefine port labels on the Configure &gt; Optimization &gt; Port Labels page.</p>
Protocol	Select All, TCP, GRE, UDP, ICMP, IPSecAH (Authentication Header), IPSecESP (Encapsulating Security Payload), or a number from the drop-down list. All specifies all TCP and UDP-based protocols.
Traffic Type	<p>Select All, Optimized, or Passthrough from the drop-down list. The system applies the QoS rules to optimized and pass-through (ingress only) traffic.</p> <p><b>Note:</b> Session reliability (port 2598) is not supported with pass-through Citrix traffic.</p>
DSCP	Optionally, specify a DSCP level.
VLAN Tag ID	Optionally, specify the VLAN tag for the rule.

Control	Description
Application	<p>Select an application from the drop-down list of global applications. To narrow the search, type the first characters in the application name.</p> <p>You can define and add any applications that do not appear in the list.</p> <p>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.</p>
Service Class Name	Select a service class from the drop-down list.
Add	Adds a rule to the inbound QoS rule list.
Remove Rules	Removes the selected rules.
Move	Select the box next to the name and click <b>Move QoS Rules</b> . Click the arrow next to the desired rule position. The rule moves to the new position.

3. Click **Apply** to apply your settings.
4. Click **Save** to save your settings permanently.

#### To modify an inbound QoS class or rule

1. Choose Configure > Networking > Inbound QoS to display the Inbound QoS page.
2. Select the class or rule name in the class or rule table.
3. Modify the settings.
4. Click **Apply** to save your settings to the running configuration.
5. Click **Save** to save your settings permanently.

### Verifying and Saving an Inbound QoS Configuration

After you apply your settings, you can verify whether the traffic is categorized in the correct class by choosing Reports > Networking > Inbound QoS 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 also verify whether the configuration is honoring the bandwidth allocations by reviewing the Inbound QoS report.

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 about saving configurations, see [“Managing Configuration Files” on page 194](#).

#### Related Topics

- [“Configuring Port Labels” on page 75](#)
- [“Managing Configuration Files” on page 194](#)
- [“Viewing Inbound QoS Reports” on page 275](#)

---

## Selecting WAN Paths Dynamically

You configure path selection in the Configure > Networking > Path Selection page.

Path selection is a RiOS v8.5 feature that chooses a predefined WAN gateway for certain traffic flows in real-time, based on availability. You define a path by specifying an egress point and providing a direction for the egressing packets to take. This granular path manipulation enables you to better use and more accurately control traffic flow across multiple WAN circuits.

---

**Important:** Due to the effect of the path selection feature on packet path manipulation, Riverbed recommends that prior to deploying path selection, you contact the Riverbed Path Selection team at [pathselection@riverbed.com](mailto:pathselection@riverbed.com) for network design guidance and the latest updates on supported deployment scenarios.

---

A common use of path selection is to route voice and video over an expensive, high-quality MPLS link, while offloading less time-sensitive business traffic over a less expensive Internet VPN link. This solution provides the right performance levels for your applications and saves on bandwidth costs by optimizing the use of available bandwidth.

Path selection is a transparent operation to the client, server, and any networking devices such as routers or switches.

Path selection works independently of the Steelhead appliance optimization service and functions even if you pause the optimization service or if the optimization service becomes unavailable.

Path selection identifies and processes both optimized and passthrough traffic.

You can configure up to 1024 paths.

## Configuring Path Selection

Configuring path selection involves specifying paths and path preferences for certain traffic. At a high level, you can configure multiple paths for each connection by specifying rules based on various parameters. The Steelhead appliance monitors the state of the path and, based on this, selects the appropriate path for a packet. Selecting appropriate paths for packets provides more control over network link use.

You can define a path based on egress interface and, optionally, next hop (gateway) IP address. You can specify different DSCP marks per-path for a given flow, allowing an upstream router to steer packets based on the observed marking.

For each application, you also specify parameters to monitor path availability; you configure the latency of the path (timeout) and the loss observed (threshold). Path selection uses ICMP pings to monitor the path state dynamically, on a regular schedule (the default is 2 seconds). If the ping responses do not make it back within the probe timeout period, the probe is considered lost. If the system loses the number of packets defined by the probe threshold, it considers the path to be down and triggers an alarm, indicating that the path is unavailable.

Path selection is a dual-ended Steelhead appliance solution. You must deploy two Steelhead appliances using path selection to enforce the return path. To define the return path for traffic and override the original traffic path, you must deploy a Steelhead appliance near the return traffic WAN junction point.

If one path fails, the Steelhead appliance directs traffic through the next available path. When the original path comes back up, the Steelhead appliance redirects the traffic back to it.

Path selection is compatible with all Steelhead appliance transport modes, including fixed-target configuration.

By default, path selection is disabled.

For path selection limits, see [“Path Selection Limits” on page 158](#).

For path selection use case examples, see [“Path Selection Use Cases” on page 156](#).

For more details on path selection, see the *Steelhead Appliance Deployment Guide*.

### To configure path selection

1. Choose Configure > Networking > Path Selection to display the Path Selection page.
2. Under Path Selection Settings, select Enable Path Selection.
3. Click **Apply** to save your settings to the running configuration.
4. Under Path Selections, define a path and the settings the Steelhead appliance uses while probing the path, as described in this table. You prioritize the path in basic or advanced outbound QoS, as described in step 7, after defining it. You can define up to three paths per traffic flow: a primary, a secondary, and a tertiary path.

Control	Description
Add a New Path	Displays the controls to define a path.
Name	Specify a name for the path.
Gateway IP Address	Optionally, specify the IP address of the next-hop gateway for the path. For example, you can specify the primary path for a connection by the next-hop IP address of router 2. The secondary path is next-hop router 3.
Interface	Select a relay interface over which the Steelhead appliance reaches the path.
Remote IP Address	Specify the IP address of the remote host to poll when monitoring the path status.
DSCP	<p>Select the DSCP marking for the ping packet. You must select this option if the service providers are applying QoS metrics based on DSCP marking and each provider is using a different type of metrics.</p> <p>The default marking is reflect. Reflect specifies that the DSCP level or IP ToS value found on pass-through and optimized traffic is unchanged when it passes through the Steelhead appliance.</p>
Timeout	<p>Specify how much time, in seconds, elapses before the system considers the path to be unavailable. The default value is 2 seconds.</p> <p>Path selection uses ICMP pings to probe the paths. If the ping responses do not make it back within this timeout setting and the system loses the number of packets defined by the threshold value, it considers the path to be down and triggers the Path Selection Path Down alarm.</p>
Threshold	<p>Specify how many timed-out probes to count before the system considers the path to be unavailable and triggers the Path Down alarm. The default is 3 failed successive packets.</p> <p>This value also determines the how many probes the system must receive to consider the path to be available.</p> <p>Path selection uses ICMP pings to monitor path availability. If the ping responses do not make it back within the probe timeout and the system loses the number of packets defined by this threshold, it considers the path to be down and triggers the Path Selection Path Down alarm.</p>

Control	Description
Add	Adds the path to the path table. The Management Console redisplay the path table and applies your changes to the running configuration, which is stored in memory.
Remove Selected Paths	Select the check box next to the path name and click <b>Remove Selected Paths</b> .

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

The path and its status appears in the path table.

Repeat steps 4 and 5 to add secondary and tertiary paths for a traffic flow. In case the configured primary path is unavailable, the system uses these alternate paths.

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

The next step is to identify the traffic and assign it to path names. Riverbed recommends that you use the Riverbed Application Flow Engine (AFE) to identify the traffic and steer it along a configured path. To use the AFE in QoS, you select the traffic type from the application list rather than specifying IP header information. Or, you can combine header-based (L3/L4) criteria with an AFE-identified L7 protocol. As an alternative to using AFE, you can specify any of the following combinations of IP header information:

- Source IP
- Destination IP
- Source port address
- DSCP mark
- VLAN tag
- Optimized or unoptimized traffic
- Protocol (TCP, UDP, GRE, and so on)

- Choose **Configure > Networking > Outbound QoS (Basic)** or **Configure > Networking > Outbound QoS (Advanced)**.

Path selection is independent of QoS settings; you can assign path names to traffic without enabling QoS marking or shaping.

- Click **Add Application** for basic QoS, or **Add a Site or Rule** for advanced QoS, depending on your QoS configuration, and identify the traffic flow by selecting an application for AFE or specifying IP header information.

For details, see [“Adding an Application” on page 121](#) and [“Adding a QoS Site or Rule for Outbound QoS \(Advanced\)” on page 136](#).

- Select the preferred primary path name for the application and the DSCP value from the drop-down list. You can define up to three paths in order of priority and three DSCP values in basic QoS.

In advanced QoS, you can define up to three paths for a rule and three DSCP values for a site. The DSCP values can steer traffic based on PBR in an upstream router.

The paths you select cascade from one to the next, based on availability. You must select DSCP values if the service providers are applying QoS metrics based on DSCP marking and each provider is using a different type of metrics.

- Click **Add** to save the settings.

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

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

Path selection processes new flows after you enable it, and does not process preexisting flows.

If the primary path assigned to a connection becomes unavailable, the Steelhead appliance directs traffic through another available path and triggers the Path Selection Path Down alarm. When the original path comes back up, the Steelhead appliance redirects the traffic back to it.

For details on the Path Selection Path Down alarm, see [“Configuring Alarm Settings” on page 217](#) and [“SNMP Traps” on page 342](#).

To verify that the traffic in a multi-interface deployment is exiting the correct interface, choose Reports > Networking > Interface Counters.

To troubleshoot, Riverbed recommends taking TCP dump traces on all WAN and LAN interfaces.

## Path Selection Use Cases

This section describes several different ways to configure path selection. For more use cases, see the *Steelhead Appliance Deployment Guide*.

### Using PBR Routing on a Downstream Router to Select a Path


In this configuration, the Steelhead appliance uses one path going out because there is just one in-path interface. However, the downstream router uses Policy-Based Routing (PBR) based on the DSCP mark on the packet. PBR is a packet redirection mechanism that allows you to define policies to route packets instead of relying on routing protocols. PBR redirects packets to Steelhead appliances that are in a virtual in-path deployment.

To monitor the path, you can configure a remote IP that is reachable only from the MPLS link and another remote IP that is reachable only from the Internet link. You configure DSCP marks for each application in case the Steelhead appliance selects the primary path. Similarly, you configure a DSCP mark for the secondary path for each application.

For details about PBR, see the *Steelhead Appliance Deployment Guide*.

The Steelhead appliance monitors both remote IP addresses and selects the primary path if the remote IP address associated with the primary path is reachable, and the secondary path if it is not reachable. The probes monitoring the remote IP address do not need to have a DSCP marking; however, you can configure a DSCP mark on the probe. This use case assumes that you configure the network so that the probes take the appropriate path.


**Figure 5-12. Path Selection by Downstream Router**



## Using an Interface and Next Hop IP Address to Select a Path

In this configuration, you have multiple paths from the Steelhead appliance to the remote data center. The Steelhead appliance selects which path to use. For each application, you define the primary path as a combination of the outgoing interface (wan0\_0 or wan0\_1) and the next hop IP address. The system must send the probe packets over the exact path that the data packets take.

**Figure 5-13. Path Selection by Relay Interface and Next Hop IP Address**



You can define this type of configuration in one of these ways:

### ***Define the Paths Using Only the Relay Interface***

In this configuration, you define a primary path by specifying the wan0\_0 interface and the secondary path as the wan0\_1 interface. Suppose that the Steelhead appliance selects the primary path for the application. In this case, it does not matter whether it sends the packet to path 1 or path 2. In both cases, the Steelhead appliance selects the MPLS path.

While probing for the remote IP address from wan0\_0, the probe packets use either R1 or R2. The Steelhead appliance cannot monitor both paths because it does not know about them. Since it monitors only one path, it ensures that all data packets are also sent over that path. Assuming that the probe packets are being sent to the remote IP through R1, it cannot use path 2 to send data packets toward the server, because this path might be down. The Steelhead appliance does not route data packets, but simply uses the next hop learnt by probing.

In the case of the secondary path, all packets are sent through path 2 so there is no confusion.

### ***Define the Paths Using an Interface and the Next Hop IP Address***

In this configuration, you specify the next hop as well as the relay interface to use for a given path. This is the simplest case, because the Steelhead appliance does not need to learn anything during probing. The Steelhead appliance does not need to route data packets, because they use the next hop specified in the configuration. The Steelhead appliance sends the packets out of the configured relay.

## Path Selection Limits

These limits apply to path selection:

- You must disable RSP or VSP on a Steelhead EX appliance running v1.0 before enabling path selection.
- You cannot base a path selection on VLANs.
- You cannot use a wildcard for the relay interface in the path definition. For example, you have to specify a relay interface for a path if you are not using PBR.
- You cannot use VLAN transparency for connections that have been rerouted by path selection.
- You cannot configure LAN-side path selection.
- Path selection does not support multiple firewalls. For details on firewalled path selection deployments, see the *Steelhead Appliance Deployment Guide*.
- Path selection does not handle ricochet of probe packets across relay interfaces.
- Path selection does not support L2 WANs.
- Path selection does not support inbound QoS, because the traffic has already traversed the initial path.
- Fully transparent inner connections might require connection forwarding.
- Path selection does not support IPv6 connections or packet-mode flows.
- You must not install any downstream appliance that does source MAC learning a hop away from the WAN side of the Steelhead appliance. Path selection updates a source MAC address of a packet to that of the relay being used to transmit it (IP addresses are unchanged). If source MAC learning is enabled on a downstream Steelhead appliance that is present at next hop, the packets destined to the original source are updated with the MAC address of the Steelhead appliance. When processing the packet, the Steelhead appliance detects that the destination MAC address is that of itself and sends the packet up its stack instead of relaying it forward.
- Path selection does not support WCCP unless it is in DSCP-only mode.
- The Steelhead appliance never takes on the router role or the role of a default gateway. Because path selection is transparent, you do not have to make network design changes to accommodate path selection design.
- Path selection does not update QoS rule changes for long-lived, locally-originated connections such as OOB or connection forwarding cluster and neighbor connections until you restart the optimization service.
- You cannot use path selection with single-ended SCPS connections.

## Related Topics

- [“Applying QoS Policies” on page 107](#)
- [“Configuring Outbound QoS \(Basic\)” on page 115](#)
- [“Configuring Outbound QoS \(Advanced\)” on page 128](#)



---

## 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 these constraints:

- WCCP cannot be enabled.
- The default route must exist on each Steelhead appliance in your network.

---

**Tip:** For detailed configuration information, see the *Steelhead Appliance Deployment Guide*.

---

The simplified routing feature is compatible with IPv6.

### To enable simplified routing

1. Choose **Configure > Networking > Simplified Routing** to display the Simplified Routing page.
2. Under **Mapping Data Collection Setting**, complete the configuration as described in this table.

Control	Description
Collect Mappings From	Select one of these options from the drop-down list: <ul style="list-style-type: none"><li>• <b>None</b> - Do not collect mappings.</li><li>• <b>Destination Only</b> - Collects destination MAC data. Use this option in connection-forwarding deployments. This is the default setting.</li><li>• <b>Destination and Source</b> - Collect mappings from destination and source MAC data. Use this option in connection-forwarding deployments.</li><li>• <b>All</b> - Collect mappings for destination, source, and inner MAC data. Also collect data for connections that are <i>un-NATted</i> (that is, connections that are not translated using NAT).</li></ul>

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

### Related Topics

- [“In-Path Rules Overview” on page 33](#)
- [“Configuring Connection Forwarding Features” on page 94](#)

---

## 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 that 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 Steelhead appliances for optimization.

RiOS allows 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.

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 appliance with the lowest IP address assigns all traffic flows to itself.

In virtual in-path configurations, all traffic flows in and out of one physical interface, and the default subnet side rule causes all traffic to appear to originate from the WAN side of the device. For more information, see [“Configuring Subnet Side Rules” on page 99](#).

Enabling WCCP is *optional*.

WCCP does not support IPv6.

---

**Tip:** You can also use the CLI to configure WCCP service groups. For detailed configuration information (including configuring the WCCP router), see the *Steelhead Appliance Deployment Guide*.

---

### 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 47](#).

---

1. Choose **Configure > Networking > WCCP** to display the WCCP page.
2. Under WCCP Service Groups, complete the configuration as described in this 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 this table.

Control	Description
Add a New Service Group	Displays the controls for adding a new service group.
Interface	<p>Select a Steelhead appliance interface to participate in a WCCP service group.</p> <p>RiOS 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.</p> <p>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 <i>Steelhead Appliance Deployment Guide</i>.</p> <p>If multiple Steelhead appliances are used in the topology, they must be configured as neighbors.</p> <p>RiOS v6.5 and later requires connection forwarding in a WCCP cluster.</p>
Service Group ID	<p>Enables WCCP v2 support on all groups added to the Service Group list.</p> <p>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.</p> <p><b>Note:</b> The service group ID is local to the site where WCCP is used.</p> <p><b>Note:</b> The service group number is not sent across the WAN.</p>
Protocol	Select a traffic protocol from the drop-down list: TCP, UDP, or ICMP. The default value is TCP.
Password/Password Confirm	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	<p>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.</p> <p>The priority value must be consistent across all Steelhead appliances within a particular service group.</p>

Control	Description
Weight	<p>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.</p> <p>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.</p> <p>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.</p> <p>As an 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 (<math>100 + 200/2 + 200/2</math>). 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.</p> <p>The range is 0-65535. The default value corresponds to the number of TCP connections your Steelhead appliance supports.</p> <p><b>Failover Support</b></p> <p>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 nonzero 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.</p> <p>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.</p> <p>In a service group, if an interface with a nonzero weight fails, its weight transfers over to the weight 0 interface of the same service group.</p> <p>For details on using the weight parameter to balance traffic loads and provide failover support in WCCP, see the <i>Steelhead Appliance Deployment Guide</i>.</p>

Control	Description
Encapsulation Scheme	<p>Specifies the method for transmitting packets between a router or a switch and a Steelhead appliance interface. Select one of these encapsulation schemes from the drop-down list:</p> <ul style="list-style-type: none"><li>• <b>Either</b> - Use Layer-2 first; if Layer-2 is not supported, GRE is used. This is the default value.</li><li>• <b>GRE</b> - 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.</li><li>• <b>L2</b> - 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.</li></ul>

Control	Description
Assignment Scheme	<p>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 these schemes from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Either</b> - 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.</li> <li>• <b>Hash</b> - 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.</li> </ul> <p>Your hashing scheme can be a combination of the source IP address, destination IP address, source port, or destination port.</p> <ul style="list-style-type: none"> <li>• <b>Mask</b> - 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.</li> </ul> <p>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.</p> <p>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 <i>masked</i>, from the IP address and the TCP port packet header fields.</p> <p>Mask assignment in RiOS v6.1 and later supports load-balancing across multiple active Steelhead appliance interfaces in the same service group.</p> <p>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.</p> <p>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.</p> <p>For details and best practices for using assignment schemes, see the <i>Steelhead Appliance Deployment Guide</i>.</p> <p><b>Important:</b> 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 <i>Steelhead Appliance Deployment Guide</i>.</p>
Source Mask/Destination Mask	<ul style="list-style-type: none"> <li>• <b>IP Mask</b> – When using the mask assignment scheme, specifies the service group source IP mask. The default value is 0x1741.</li> <li>• <b>Port Mask</b> – When using the mask assignment scheme, specifies the service group source port mask.</li> </ul>

Control	Description
Source Hash/Destination Hash	<ul style="list-style-type: none"> <li>• <b>IP Hash</b> – When using the hash assignment scheme, specifies that the router hashes the source IP address to calculate part of the load distribution and traffic redirection.</li> <li>• <b>Port Hash</b> – When using the hash assignment scheme, specifies that the router hashes the source port to calculate part of the load distribution and traffic redirection.</li> </ul>
Ports Mode	Select one of these modes from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Ports Disabled</b> - Select to disable the ports.</li> <li>• <b>Use Source Ports</b> - The router determines traffic to redirect based on source ports.</li> <li>• <b>Use Destination Ports</b> - The router determines traffic to redirect based on destination ports.</li> </ul>
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 <b>Remove Selected Groups</b> .

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 appliance 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 appliance 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 *Steelhead Appliance 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 about creating a WCCP service group, see [“Configuring WCCP” on page 160](#).

When you are modifying service group settings in RiOS v6.1 or later, the service group description includes the interface.

### 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.
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 47](#)
- [“Verifying a Multiple In-Path Interface Configuration” on page 165](#)

---

## 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 or Two-Port LR Single Mode 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 appliance 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.
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.

RiOS ignores all hardware-assist rules unless you select this check box. No TCP traffic is passed through.
3. Under TCP Hardware Assist Rules, complete the configuration as described in this table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Type	<p>Select a rule type:</p> <ul style="list-style-type: none"> <li>• <b>Accept</b> - Accepts rules matching the Subnet A or Subnet B IP address and mask pattern for the optimized connection.</li> <li>• <b>Pass-Through</b> - Identifies traffic to be passed through the network unoptimized.</li> </ul>
Insert Rule At	<p>Determines the order in which the system evaluates the rule. Select Start, End, or a rule number from the drop-down list.</p> <p>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.</p> <p>In general, filter traffic that is to be unoptimized, discarded, or denied before processing rules for traffic that is to be optimized.</p>
Subnet A	<p>Specify an IP address and mask for the subnet that can be both source and destination together with Subnet B.</p> <p>Use the format XXX.XXX.XXX.XXX/XX.</p> <p><b>Note:</b> You can specify all or 0.0.0.0/0 as the wildcard for all traffic.</p>
Subnet B	<p>Specify an IP address and mask for the subnet that can be both source and destination together with Subnet A.</p> <p>Use the format XXX.XXX.XXX.XXX/XX.</p> <p><b>Note:</b> You can specify all or 0.0.0.0/0 as the wildcard for all traffic.</p>

Control	Description
VLAN Tag ID	<p>Optionally, specify a numeric VLAN tag identification number.</p> <p>Select all to specify the rule applies to all VLANs.</p> <p>Select untagged to specify the rule applies to nontagged connections.</p> <p><b>Note:</b> Pass-through traffic maintains any preexisting VLAN tagging between the LAN and WAN interfaces.</p> <p><b>Note:</b> 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 about configuring the in-path interface for the Steelhead appliance, see <a href="#">“Configuring In-Path Rules” on page 36</a>.</p>
Description	Optionally, include a description of the rule.
Add	<p>Adds the new hardware-assist rule to the list. You can add up to a maximum number of 50 rules.</p> <ul style="list-style-type: none"> <li>• RiOS applies the same rule to both LAN and WAN interfaces.</li> <li>• Every 10G card has the same rule set.</li> </ul> <p>The Steelhead appliance refreshes the hardware-assist rules table and applies your modifications to the running configuration, which is stored in memory.</p>
Remove Selected Rules	Select the check box next to the name and click <b>Remove Selected Rules</b> .
Move Selected Rules	Moves the selected rules. Click the arrow next to the desired rule position; the rule moves to the new position.

## CHAPTER 6      **Configuring SSL and a Secure Inner Channel**

This chapter describes how to configure SSL support. It includes these sections:

- [“Configuring SSL Server Certificates and Certificate Authorities” on page 169](#)
- [“Configuring CRL Management” on page 174](#)
- [“Configuring Secure Peers” on page 176](#)
- [“Configuring Advanced and SSL Cipher Settings” on page 183](#)

---

### **Configuring SSL Server Certificates and Certificate Authorities**

This section provides an overview of SSL support and describes how to configure SSL server certificates and certificate authorities.

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**



RiOS 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 183](#).

RiOS 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. 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 206](#).

## Prerequisite Tasks

Complete these 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 191](#). If you do not have a valid Enhanced Cryptography License Key file, go to <https://sslcrt.riverbed.com> and follow the procedures documented there.
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.	<a href="#">“Configuring Secure Peers” on page 176</a>
2. Set the SSL secure vault password on the client and server-side Steelhead appliance.	<a href="#">“Unlocking the Secure Vault” on page 206</a>
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. Client-side session reuse is enabled by default in RiOS v7.0 and later.	<a href="#">“Setting Advanced SSL Options” on page 183</a>
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.	<a href="#">“Configuring SSL Server Certificates and Certificate Authorities” on page 169</a>
5. Create an in-path rule for the client-side Steelhead appliance.  <b>Out-of-path configurations</b> - 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 these property values: <ul style="list-style-type: none"> <li>• Type - Fixed target</li> <li>• 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.</li> <li>• VLAN Tag - All</li> <li>• Data Reduction Policy - Normal</li> <li>• Neural Framing Mode - Always</li> </ul>	<a href="#">“Configuring In-Path Rules” on page 36</a>

Task	Reference
<p>6. Configure mutual peering trusts so the server-side Steelhead appliance trusts the client-side Steelhead appliance and vice versa. Use one of these approaches:</p> <p><b>Use the secure inner channel and peering lists:</b></p> <ul style="list-style-type: none"> <li>• Configure the inner channel SSL settings as described in <a href="#">“Configuring Secure Peers” on page 176</a>. Both the client-side and server-side Steelhead appliances must be running RiOS v5.0 or later.</li> <li>• To automatically discover Steelhead appliances using self-signed certificates, open your secure application to send some traffic through the Steelhead appliances. The connection is passed through to the server without optimization, but the Steelheads will automatically discover the peers and place them in the self-signed peer <i>gray</i> list.</li> <li>• 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.</li> <li>• Reopen your secure application.</li> </ul> <p>—or—</p> <p><b>Add CA-signed peer certificates:</b></p> <ul style="list-style-type: none"> <li>• Add the PEM certificate of the designated CA as a new trusted entity to the peering trust list for each Steelhead appliance.</li> </ul> <p><b>Tip:</b> 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.</p>	<p><a href="#">“Configuring Secure Peers” on page 176</a></p>
<p>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.</p> <p>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 &gt; Tools &gt; Internet Options &gt; Certificates.</p>	<p><a href="#">“Configuring SSL Server Certificates and Certificate Authorities” on page 169</a></p>
<p>8. On the client and server-side Steelhead appliance, restart the optimization service.</p>	<p><a href="#">“Starting and Stopping the Optimization Service” on page 187</a></p>

## Verifying SSL and Secure Inner Channel Optimization

Use these 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. These 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.
  - In the Management Console, the Traffic Summary report displays encrypted traffic (typically, HTTPS).

---

**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 these 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 these tools to verify that secure inner channels are in use for the selected application traffic types:
  - In the Current Connections report, 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 about the secure inner channel, see [“Secure Inner Channel Overview” on page 176](#).

- **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 block cipher using 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 whether 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 that 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 CRL Management

RiOS 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 appliance in the handshake between the two 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 under Certificate Authorities in the CRL Management 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.
2. Under CRL Settings, complete the configuration as described in this table.

Control	Description
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.

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.

2. 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.

To encrypt and optimize the connection for *non-SSL* traffic, 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 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.

After 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 bidirectional; 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 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 177](#).

---

### To enable secure peering

1. Riverbed recommends using NTP time synchronization or manually synchronizing the clocks on both the server-side and client-side Steelhead appliances. It is critical that the peer Steelhead appliance time is the same for the trust relationship to work.
2. On both the server-side and client-side Steelhead appliances, choose **Configure > Optimization > Secure Peering (SSL)** to display the Secure Peering (SSL) page.
3. Check the **Enable SSL Secure Peering** check box.
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 optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

---

**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 these 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 these 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 these options for configuring peer certificates and Mobile Controller trust:

- **Peering Trust** - Add and view these 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.  
The Steelhead identity certificate details appear, as described in this table.

Control	Description
Issued To/Issued By	<b>Common Name</b> - Specifies the common name of the certificate authority.
	<b>Organization</b> - Specifies the organization name (for example, the company).
	<b>Organization Unit</b> - Specifies the organization unit name (for example, section or department).
	<b>Locality</b> - Specifies the city.
	<b>State</b> - Specifies the state.
	<b>Country</b> - Specifies the country.
	<b>Serial Number</b> - Specifies the serial number (Issued To, only).
Validity	<b>Issued On</b> - Specifies the date the certificate was issued.
	<b>Expires On</b> - Specifies the date the certificate expires.
Fingerprint	Specifies the SSL fingerprint.
Key	<b>Type</b> - Specifies the key type.
	<b>Size</b> - Specifies the size in bytes.

2. To replace an existing certificate, under Certificate, click **Replace** and complete the configuration as described in this 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. <b>Note:</b> The private key is required.
	<b>Local File</b> - Browse to the local file.
	<b>Text</b> - Paste the text content of the file into the text box.
	<b>Decryption Password</b> - Specify the decryption password, if necessary.
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. <b>Note:</b> 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.
	<b>Cipher Bits</b> - Select the key length from the drop-down list. The default value is 1024.
	<b>Common Name (required)</b> - Specify the hostname of the peer.
	<b>Organization Name</b> - Specify the organization name (for example, the company).
	<b>Organization Unit Name</b> - Specify the organization unit name (for example, the section or department).
	<b>Locality</b> - Specify the city.
	<b>State (no abbreviations)</b> - Specify the state.
	<b>Country (2-letter code)</b> - Specify the country (2-letter code only).
	<b>Email Address</b> - Specify the email address of the contact person.
	<b>Validity Period (Days)</b> - Specify how many days the certificate is valid. The default value is 730.
Update Certificate Through SCEP Enrollment	Select to generate a private key and CSR using a Simple Certificate Enrollment Protocol (SCEP) responder. Click the SCEP Management tab to configure the SCEP responder.

3. To export an existing certificate, under Certificate, click **Export** and complete the configuration as described in this 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 this 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 this table.

Control	Description
URL	Specify the URL of the SCEP responder. Use the format http://host[:port]/path[/to/service] Example: http:examplehost:1212/pathtoservice  RiOS v8.5 supports single, two, and three-tier hierarchies to validate the chain certificates it receives.
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 fail then the Steelhead appliance does not poll.
Poll Period	Specify the poll frequency in minutes. The default value is 5.
Change Challenge Passphrase	Specify the challenge password phrase.
Enable Auto Enrollment	Enables automatic re-enrollment of a certificate to be signed by a CA using SCEP. <ul style="list-style-type: none"> <li>• <b>Expiration Threshold</b> - 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.</li> </ul>
Update SCEP Settings	Updates the SCEP settings.

6. To add or remove a Trusted entity, under Peering Trust, complete the configuration as described in this 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 <b>Remove Selected</b> .

7. To add or remove a SCEP Trusted entity, under SCEP Peering Trust, complete the configuration as described in this 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 <b>Remove Selected</b> .

8. 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 this table.

The white, gray, and black peering lists sort the peers by IP address.

---

**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 <b>Remove Selected</b> .

---

**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.

---

9. Click **Apply** to save your settings to the running configuration.
10. Click **Save** to save your settings permanently.
11. Restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

## Verifying the Secure Inner Channel Connections

This section describes what happens when a secure inner channel cannot be established for traffic between Steelhead appliances and how to verify whether connections are using a secure inner channel.

When the Steelhead appliances are configured to use secure inner channels for SSL 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 about 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 that runs into a failure is passed through without optimization if Fallback to No Encryption is disabled. See [“Enabling Secure Peers” on page 177](#).
- 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 Connection Reports” on page 248](#) and [“Verifying SSL and Secure Inner Channel Optimization” on page 173](#).

### **Related Topics**

- [“Configuring In-Path Rules” on page 36](#)
- [“Enabling Peering and Configuring Peering Rules” on page 53](#)
- [“Configuring FCIP Optimization” on page 77](#)
- [“Unlocking the Secure Vault” on page 206](#)
- [“Generating System Dumps” on page 329](#)

---

## **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 these sections:

- [“Setting Advanced SSL Options” on page 183](#)
- [“Configuring SSL Cipher Settings” on page 184](#)

### **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.

2. Complete the configuration as described in this table.

Control	Description
Enable SNI	<p>Enable on the server-side Steelhead appliance while using name-based virtual hosts with SSL. Server name indication (SNI) is a transport layer security extension to the SSL protocol. With SNI, the first SSL client hello handshake message sent to the HTTPS server includes the requested virtual hostname to which the client is connecting. Because the server is aware of the hostname, it returns a host-specific security certificate.</p> <p>Without SNI, an HTTPS server returns a default certificate that satisfies hostnames for all virtual hosts. The SSL connection setup uses the default virtual host configuration for the address where the connection was received. Browser messages warn that certificates have the wrong hostname.</p> <p>With SNI enabled, RiOS provides the hostname. This enables the server to determine the correct named virtual host for the request and set up the connection accordingly from the start.</p> <p>The browser validates the certificate names against the requested URL, and the server-side Steelhead appliance verifies that the selected proxy certificate is compatible with the client hostname. This ensures that the browser does not reject the proxy certificate for the server-side Steelhead appliance.</p> <p>If SNI provides a hostname that does not exactly match the common name or any of the subject alternate names for the certificate on the server-side Steelhead appliance, the system determines that a valid certificate is not present and bypasses that hostname.</p> <p>No configuration is necessary on the client-side Steelhead appliance.</p> <p>The client browser must also support SNI.</p>

3. Click **Apply** to apply your settings.
4. Click **Save** to save your settings permanently.
5. If you have enabled SNI, you must restart the optimization service. For details, see [“Starting and Stopping the Optimization Service” on page 187](#).

## 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.
2. Under Peer Ciphers, complete the configuration on both the server-side and client-side Steelhead appliances, as described in this 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 <b>Remove Selected</b> .



## CHAPTER 7    **Managing Steelhead Appliances**

This chapter describes tasks you perform for routine management of the Steelhead appliance. It includes these sections:

- [“Starting and Stopping the Optimization Service” on page 187](#)
- [“Configuring Scheduled Jobs” on page 188](#)
- [“Upgrading Your Software” on page 189](#)
- [“Rebooting and Shutting Down the Steelhead Appliance” on page 191](#)
- [“Managing Licenses and Model Upgrades” on page 191](#)
- [“Viewing Permissions” on page 194](#)
- [“Managing Configuration Files” on page 194](#)
- [“Configuring General Security Settings” on page 196](#)
- [“Managing User Permissions” on page 197](#)
- [“Managing Password Policy” on page 200](#)
- [“Setting RADIUS Servers” on page 203](#)
- [“Configuring TACACS+ Access” on page 205](#)
- [“Unlocking the Secure Vault” on page 206](#)
- [“Configuring a Management ACL” on page 207](#)
- [“Configuring Web Settings” on page 210](#)
- [“Enabling REST API Access” on page 213](#)

---

### **Starting and Stopping the Optimization Service**

You can start, stop, and restart the optimization service in the Configure > Maintenance > Services page. You can also use this page to reset the optimization service alarm after it has been triggered.

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

Many of the optimization service commands are initiated at startup. It is important to restart the optimization service when you have made changes to your configuration.

---

**Important:** Restarting the optimization 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.
2. Under Optimization Service click **Stop**, **Start**, or **Restart**.
3. Click **Save** to save your settings permanently.

### To reset the optimization service alarm

1. Choose Configure > Maintenance > Services to display the Services page. The option to reset the optimization service alarm appears only after RiOS 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 appliance reboot or shut down.
- generate multiple TCP trace dumps on a specific date and time.

To schedule all other jobs, you must use the Riverbed CLI.

For details about 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.
2. Select Enabled or Disabled from the drop-down list to enable or disable the job.
3. Select the Job ID number to display details about the job.

- Under Details for Job <#>, complete the configuration as described in this table.

Control	Description
Name	Specify a name for the job.
Comment	Specify a comment.
Interval (seconds)	Specify the number of seconds between job recurrences. Specify 0 to run the job one-time only.
Executes on	Specify the start time and end time using the format YYYY/MM/DD HH:MM:SS.
Enable/Disable Job	Select the check box to enable the job, clear the check box to disable the job.
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	Select to enable the job to run at the specified date or clear to prevent the job from running.
Apply Changes	Applies the changes to the current configuration.
Cancel/Remove This Job	Cancels and removes the job.
Execute Now	Runs the job.
Remove Selected Jobs	Select the check box next to the name and click <b>Remove Selected Jobs</b> .

- 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.

### To revert the RiOS software version

- Choose Configure > Maintenance > Software Upgrade to display the Software Upgrade page.
- Under Software Upgrade, complete the configuration as described in this 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.

## To upgrade the RiOS software version

1. Download the software image from the Riverbed Support site to a location such as your desktop. Optionally, in RiOS v8.5, you can download a delta image directly from the Riverbed Support site to the Steelhead appliance. The downloaded image includes only the incremental changes. The smaller file size means a faster download and less load on the network. To download a delta image, skip to step 2.
2. Choose Configure > Maintenance > Software Upgrade to display the Software Upgrade page.
3. Under Install Upgrade, complete the configuration as described in this table.

Control	Description
From URL	<p>Select this option and specify the URL.</p> <p>Use one of the following formats:</p> <p>http://host/path/to/file</p> <p>https://host/path/to/file</p> <p>ftp://user:password@host/path/to/file</p> <p>scp://user:password@host/path/to/file</p>
From Riverbed Support Site	<p>Click this option and select the target release number from the drop-down list. The system uploads and installs the new image immediately after you click <b>Install</b>. To upload and install the image later, schedule another date or time before you click <b>Install</b>.</p>
From Local File	<p>Select this option and specify the path, or click <b>Browse</b> to go to the local file directory.</p> <p>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.</p>
Schedule Upgrade for Later	<p>Schedules the upgrade process. Specify the date and time to run the upgrade: YYYY/MM/DD, HH:MM:SS.</p>
Install	<p>Click to install the software upgrade on your system, unless you schedule it for later.</p> <p>The software image can be quite large; uploading the image to the appliance and installing it can take a few minutes. Downloading a delta image directly from the Riverbed Support site is faster because the downloaded image includes only the incremental changes.</p> <p>As the upgrade progresses, status messages appear.</p> <p>After the installation is complete, the system reminds you to reboot the appliance to switch to the new version of the software.</p>
Cancel	<p>Cancels your changes.</p>

4. Choose Configure > Maintenance > Reboot/Shut Down and click **Reboot**.

The appliance can take a few minutes to reboot. This is normal behavior as the software is configuring the recovery flash device. Do not press Ctrl-C, unplug, or otherwise shut down the system during this first boot. There is no indication displayed during the system boot that the recovery flash device is being configured. After the reboot, the Home page, Software Upgrade, and Support pages of the Management Console display the RiOS version upgrade.

## Related Topic

- [“Configuring Scheduled Jobs” on page 188](#)



---

## Rebooting and Shutting Down the Steelhead Appliance

You can reboot or shut down the system in the Configure > Maintenance > Reboot/Shut Down 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. Shutting down the appliance 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/Shut Down to display the Reboot/Shut Down page.
2. Click **Reboot**. After you click **Reboot**, you are logged out of the system and RiOS reboots.
3. Click **Shut Down** to shut down the system. After you click **Shut Down**, the system is turned off. To restart the system, you must manually turn on the Steelhead appliance.

### To schedule a reboot

1. Choose Configure > Maintenance > Reboot/Shut Down to display the Reboot/Shut Down page.
2. Select Schedule for Later and enter the date and time you would like the reboot to occur.  
The reboot executes at the scheduled time.

---

## 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 these sections:

- [“Flexible Licensing Overview” on page 192](#)
- [“Installing a License” on page 192](#)
- [“Removing a License” on page 193](#)

You perform all license management and Steelhead appliance model upgrades in the Configure > Maintenance > Licenses page.

Steelhead licenses can be permanent or temporary. Permanent licenses do not display an expiration date in their Status column on the Licenses page; temporary licenses display an expiration date in their Status column. For example, evaluation licenses typically expire in 60 days and display a date within that range.

The system warns you two weeks before a license expires with the Expiring License alarm. After a license expires, the system warns with an Expired License alarm. You can add a license to extend the functionality of an expiring licenses. If multiple licenses exists for a feature, the system uses the license with the latest expiration date.

## Flexible Licensing Overview

RiOS provides a flexible way to manage Steelhead appliance licenses, model configurations, and upgrades. Rather than 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.

After adding the required hardware and license to the Steelhead appliance, activate the hardware specification instead of replacing the appliance.

## For More Information

This 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.	<a href="#">“Installing a License” on page 192</a>
Update an expired license.	<a href="#">“Installing a License” on page 192</a>
Remove a license.	<a href="#">“Removing a License” on page 193</a>

## Installing a License

This section describes how to request and fetch a license manually from the Riverbed license portal or install a license manually after receiving it from Riverbed Support or Sales.

RiOS simplifies license management by providing an automated way to fetch and activate licenses for Riverbed products. You no longer have to manually activate individual appliances and install the licenses.

Fetching a license is restricted for read-only users such as monitor and RBM users with read-only access for General Settings (permissions are granted on the Configure > Security > User Permissions page).

### To install a license on a new Steelhead appliance

- Connect a new Steelhead appliance to the network.

The Steelhead appliance automatically contacts the Riverbed license portal and downloads the downloadable licenses. The Licensing page displays a success message or the Alarm Status page reports an actionable error message.

### To replace expired licenses

- Purchase new downloadable licenses to replace the expired license.

At the time of the next scheduled automatic license fetch, the Steelhead appliance automatically contacts the Riverbed license portal and downloads the new licenses. The Licensing page displays a success message or the alarm Status page reports an actionable error message.

### To fetch a license on demand

1. Choose Configure > Maintenance > Licenses to display the Licenses page.
2. Click **Fetch Updates Now**.

The Licensing page displays a success message or the alarm Status page reports an actionable error message.

### To install a license

1. Choose Configure > Maintenance > Licenses to display the Licenses page.

The Licenses page includes a table of licenses with a column showing the date and time the license was installed and the approximate relative time it was installed. The next column shows whether the installation was done manually or automatically.

Below the license table, next to the Fetch Updates Now button, a note displays the date and time of the last update. Normal update results appear in black and any errors appear in red.

2. Complete the configuration as described in this 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. <b>Tip:</b> Separate multiple license keys with a space, Tab, or Enter.
Add	Adds the license.
Fetch Licenses Now	Contacts the Riverbed license portal and downloads all applicable licenses for the Steelhead appliance.

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

### 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 want 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.
2. Under Password, complete the configuration as described in this table.

Control	Description
Change Password	Allows you to add or change your log in password.
New Password/Confirm New Password	Specify a password in the text box. Retype the password in the Confirm New Password text box.
Old Password	(Appears when password policy is enabled and the Minimum Character Difference Between Passwords value is greater than 0). Non-administrators must specify the old password.  Administrators are never required to enter an old password when changing an account password.

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 about setting user permissions, see [“Managing User Permissions” on page 197](#).

---

The My Account page includes a way to clear user preferences if any user settings result in an unsafe state and the Management Console cannot display the page.

User preferences are set for individual users and do not affect the appliance configuration.

### To restore the user preferences for the current user

1. Choose Configure > My Account to display the My Account page.
2. Under User Preferences, click **Restore Defaults**.

---

## 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 optimization 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 saves 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 optimization service for the settings to take effect. For details about restarting the optimization service, see [“Starting and Stopping the Optimization Service” on page 187](#).

---

## To manage configurations

1. Choose Configure > Configurations to display the Configurations page.
2. Under Current Configuration: <filename>, complete the configuration as described in this table.

Control	Description
Current Configuration: <configuration name>	<b>View Running Config</b> - Displays the running configuration settings in a new browser window.
	<b>Save</b> - Saves settings that have been applied to the running configuration.
	<b>Revert</b> - 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 <b>Save</b> .

3. To import a configuration from another appliance, complete the configuration as described in this 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.

Control	Description
Import Shared Data Only	Takes a subset of the configuration settings from the imported configuration and combines them with the current configuration to create a new configuration. Import shared data is enabled by default.
Add	When the Import Shared Data Only check box is selected, activates the imported configuration and makes it the current configuration. This is the default. When the Import Shared Data Only check box is not selected, adds the imported configuration to the Configuration list. It does not become the active configuration until you select it from the list and click <b>Activate</b> .
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .
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 187](#).

**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 this 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 about setting up RADIUS and TACACS+ servers, see the *Steelhead Appliance Deployment Guide*.

### To set general security settings

1. Choose Configure > Security > General Security Settings to display the General Security Settings page.

2. Under Authentication Methods, complete the configuration as described in this 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 these policies from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Remote First</b> - 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.</li> <li>• <b>Remote Only</b> - Only checks the remote server.</li> <li>• <b>Local Only</b> - Only checks the local server. All remote users are mapped to the user specified. Any vendor attributes received by an authentication server are ignored.</li> </ul> <b>Default User</b> - Optionally, select admin, monitor, or shark 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 optimization 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.

As an example, you might have user Jane who can make configuration changes to QoS 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.

### To set the administrator or monitor password

1. Choose Configure > Security > User Permissions to display the User Permissions page.
2. Under Capability-Based Accounts, complete the configuration as described in this table.

Control	Description
admin/monitor	Click the magnifying glass to change the administrator or monitor password.
	<b>Enable Account</b> - Select to enable or clear to disable the administrator or monitor account.
	<b>Change Password</b> - Enables password protection.
	RiOS includes an account control feature that allows you to select a password policy for more security. When you enable account control on the Configure > Security > Password Policy page, a user must use a password.
	When a user has a null password to start with, the administrator can still set the user password with account control enabled. However, once the user or administrator changes the password, it cannot be reset to null as long as account control is enabled.
	<b>Password</b> - Specify a password in the text box.
	<b>Password Confirm</b> - Retype 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 this 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.
Password/New Password Confirm	Specify a password in the text box, and then retype the password for confirmation.
Enable Account	Select the check box to enable the new role-based account.
General Settings	Configures per-source IP connection limit and the maximum connection pooling size.



Control	Description
Network Settings	Configures host and network interface settings, including DNS cache settings and hardware assist rules.
QoS	Enforces QoS policies.
Path Selection	Configures path selection.
Optimization Service	Configures alarms, performance features, SkipWare, 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 <i>Steelhead Appliance Deployment Guide</i> .
SSL Optimization	Configures SSL optimization.
Replication Optimization	Configures the SRDF/A, FCIP, and SnapMirror storage optimization modules.
Security Settings	Configures security settings, including RADIUS and TACACS authentication settings and the secure vault password.
Basic Diagnostics	Customizes system diagnostic logs, including system and user log settings, but does not include TCP dumps.
TCP Dumps	Customizes TCP dump settings and allows use of the Shark function for detailed packet analysis through Cascade Pilot.
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 <b>Remove Selected</b> .

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

**Note:** RiOS ignores the RBM user roles for SCA features. RiOS allows RBM users with DENY permissions in all roles access to SCA Management Console pages and GUI commands.

---

## Managing Password Policy

You can change the password policy and strength in the Configure > Security > Password Policy page.

### Selecting a Password Policy

You can choose one of these password policy templates, depending on your security requirements:

- **Strong** - Sets the password policy to more stringent enforcement settings. Selecting this template automatically prepopulates the password policy with stricter settings commonly required by higher security standards such as for the Department of Defense.
- **Basic** - Reverts the password policy to its predefined settings so you can customize your policy.

#### To set a password policy

1. Choose Configure > Security > Password Policy to display the Password Policy page.
2. Select the Enable Account Control check box to set a password policy. Enabling account control makes password use mandatory.

Passwords for all users expire as soon as account control is enabled. This forces all users to create new passwords that follow the password requirements defined in the password policy. All new passwords are then controlled by the password policy.

The passwords also expire after the number of days specified by the administrator in the Password Policy page. As a consequence of this change, when a user tries to log in to the Management Console and their password has expired, the Expired Password page asks them to change their password. After they change their password, the system automatically logs them in to the Management Console.

RiOS does not allow empty passwords when account control is enabled.

3. Optionally, select either the Basic or Strong template. When you select the basic template, the system prepopulates the page with the secure settings. Also, the system prompts a user logging into the Steelhead appliance after 60 days to change their password. By default, RiOS locks out a user logging into the Steelhead appliance after 300 days without a password change. After the system locks them out, an administrator must unlock the system.

4. Under Password Management, complete the configuration as described in this table.

Control	Description
Login Attempts Before Lockout	Specify the maximum number of unsuccessful login attempts before temporarily blocking user access to the Steelhead appliance. The user is prevented from further login attempts when the number is exceeded. The default for the strong security template is 3.  The lockout expires after the amount of time specified in Timeout for User Login After Lockout elapses.
Timeout for User Login After Lockout	Specify the amount of time, in seconds, that must elapse before a user can attempt to log in after an account lockout due to unsuccessful login attempts. The default for the strong security template is 300.
Days Before Password Expires	Specify the number of days the current password remains in effect. The default for the strong security template is 60. To set the password expiration to 24 hours, specify 0. To set the password expiration to 48 hours, specify 1. Leave blank to turn off password expiration.
Days to Warn User of an Expiring Password	Specify the number of days the user is warned before the password expires. The default for the strong security template is 7.
Days to Keep Account Active After Password Expires	Specify the number of days the account remains active after the password expires. The default for the strong security template is 305. When the time elapses, RiOS locks the account permanently, preventing any further logins.
Days Between Password Changes	Specify the minimum number of days before which passwords cannot be changed.
Minimum Interval for Password Reuse	Specify the number of password changes allowed before a password can be reused. The default for the strong security template is 5.

5. Under Password Characteristics, complete the configuration as described in this table.

Control	Description
Minimum Password Length	Specify the minimum password length. The default for the strong security template is 14 alphanumeric characters.
Minimum Uppercase Characters	Specify the minimum number of uppercase characters required in a password. The default for the strong security template is 1.
Minimum Lowercase Characters	Specify the minimum number of lowercase characters required in a password. The default for the strong security template is 1.
Minimum Numerical Characters	Specify the minimum number of numerical characters required in a password. The default for the strong security template is 1.
Minimum Special Characters	Specify the minimum number of special characters required in a password. The default for the strong security template is 1.
Minimum Character Differences Between Passwords	Specify the minimum number of characters that must be changed between the old and new password. The default for the strong security template is 4.
Maximum Consecutively Repeating Characters	Specify the maximum number of times a character can occur consecutively.
Prevent Dictionary Words	Select to prevent the use of any word that is found in a dictionary as a password. By default, this control is enabled.

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

## Unlocking an Account

RiOS temporarily locks out an account after a user exceeds the configured number of login attempts. Account lockout information appears on the Configure > Security > User Permissions page.

When an account is locked out, the lockout ends after:

- The configured lockout time elapses.
- or—
- The administrator unlocks the account. RiOS never locks out administrator accounts.

### To unlock an account

1. Log in as an administrator (admin).
2. Choose Configure > Security > User Permissions page and click **Clear Login Failure Details**.

When the user logs into their account successfully, RiOS resets the login failure count.

## Resetting an Expired Password

RiOS temporarily locks out an account when its password expires. Passwords expire for one of these reasons:

- An administrator enables account control.
- The expiration time for a password elapses.
- An administrator disables a user account and then enables it.
- An administrator uses a CLI command to encrypt a password.

After a user password expires, the user must update their password within the number of days specified in Days to Keep Account Active After Password Expires. The default value is 305 days. After the time elapses, RiOS locks the account permanently, preventing any further logins.

### To reset the password and unlock the account

1. Log in as an administrator (admin).
2. Choose Configure > Security > User Permissions page and click **Clear Login Failure Details**.
3. Type and confirm the new password and click **Change Password**.

---

**Note:** The password reset feature is separate from the account lockout feature.

---

---

## 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.

Enabling this feature is *optional*.

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.

For details about setting up RADIUS and TACACS+ servers, see the *Steelhead Appliance Deployment Guide*.

### To set RADIUS server authentication

1. Choose Configure > Security > RADIUS to display the RADIUS page.
2. Under Default RADIUS Settings, complete the configuration as described in this 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 this table.

Control	Description
Add a RADIUS Server	Displays the controls for defining a new RADIUS server.
Hostname or IP Address	Specify the hostname or server IP address. RiOS does not support IPv6 server IP addresses.
Authentication Port	Specify the port for the server.
Authentication Type	Select one of these authentication types: <ul style="list-style-type: none"> <li>• <b>PAP</b> - Password authentication protocol (PAP), which validates users before allowing them access to the RADIUS server resources. PAP is the most flexible protocol but is less secure than CHAP.</li> <li>• <b>CHAP</b> - Challenge handshake authentication protocol (CHAP), which provides better security than PAP. CHAP validates the identity of remote clients by periodically verifying the identity of the client using a three-way handshake. This happens at the time of establishing the initial link and might happen again at any time. CHAP bases verification on a user password and transmits an MD5 sum of the password from the client to the server.</li> </ul>
Override the Global Default Key	Overrides the global server key for the server.
	<b>Server Key</b> - Specify the override server key.
	<b>Confirm Server Key</b> - 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 <b>Remove Selected</b> .

**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.

**Note:** To modify RADIUS server settings, click the server IP address in the list of Radius Servers. Use the Status drop-down list to enable or disable a server in the list.

### Related Topic

- [“Configuring General Security Settings” on page 196](#)

---

## Configuring TACACS+ Access

You set up TACACS+ server authentication in the Configure > Security > TACACS+ page.

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.

Enabling this feature is *optional*.

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.

For details about configuring RADIUS and TACACS+ servers to accept login requests from the Steelhead appliance, see the *Steelhead Appliance Deployment Guide*.

### To set a TACACS+ server

1. Choose Configure > Security > TACACS+ to display the TACACS+ page.
2. Under Default TACACS+ Settings, complete the configuration as described in this table.

Control	Description
Set a Global Default Key	Enables 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 this table.

Control	Description
Add a TACACS+ Server	Displays the controls for defining a new TACACS+ server.
Hostname or IP Address	Specify the hostname or 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. The default value is PAP.
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 <b>Remove Selected</b> .

**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 196](#)

---

## 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, the RiOS data store encryption key, and replication or delegate user configuration details. RiOS encrypts and secures these configuration settings 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 unlock or change the password of the secure vault**

1. Choose Configure > Security > Secure Vault to display the Secure Vault page.



2. Under **Unlock Secure Vault**, complete the configuration as described in this table.

Control	Description
Password	Specify a password and click <b>Unlock Secure Vault</b> .  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.
Unlock Secure Vault	Unlocks the vault.

3. Under **Change Password**, complete the configuration as described in this 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 196](#)

## **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 these safeguards to prevent accidental disconnection from the Steelhead appliance, the CMC, and the embedded Shark feature:

- 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.
2. Under Management ACL Settings, complete the configuration as described in this 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 branch service RiOS features, such as DNS caching. These default rules allow access to certain ports required by these features. The list also includes default rules that allow access to the CMC and the embedded Shark feature. If you delete a default ACL rule and need to restore it, see [“Configuring Web Settings” on page 210](#).

## To add an ACL management rule

1. Under Management ACL Settings, complete the configuration as described in this table.

Control	Description
Add a New Rule	Displays the controls for adding a new rule.
Action	<p>Select one of these rule types from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Allow</b> - Allows a matching packet access to the Steelhead appliance. This is the default action.</li> <li>• <b>Deny</b> - Denies access to any matching packets.</li> </ul>
Service	Optionally, select Specify Protocol, or HTTP, HTTPS, SOAP, SNMP, SSH, Telnet. When specified, the Destination Port is dimmed.
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.
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	<p>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).</p> <p>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.</p> <p><b>Note:</b> 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.</p>
Log Packets	Tracks denied packets in the log. By default, packet logging is enabled.
Add	Adds the rule to the list. The Management Console redisplay 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 <b>Remove Selected</b> .
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 about active and passive FTP, see [“QoS Overview” on page 107](#).
- 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.

---

## 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.
2. Under Web Settings, complete the configuration as described in this 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. <b>Important:</b> 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 provides 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 > Security > Web Settings to display the Web Settings page.
2. Under Web Certificate, select the Details tab.

The Steelhead appliance identity certificate details appear, as described in this table.

Control	Description
Issued To/Issued By	<b>Common Name</b> - Specifies the common name of the certificate authority.
	<b>Email</b> - Specifies the organization email.
	<b>Organization</b> - Specifies the organization name (for example, the company).
	<b>Organization Unit</b> - Specifies the organization unit name (for example, section or department).
	<b>Locality</b> - Specifies the city.
	<b>State</b> - Specifies the state.
	<b>Country</b> - Specifies the country.
	<b>Serial Number</b> - Specifies the serial number (Issued To, only).
Validity	<b>Issued On</b> - Specifies the date the certificate was issued.
	<b>Expires On</b> - Specifies the date the certificate expires.
Fingerprint	Specifies the SSL fingerprint.
Key	<b>Type</b> - Specifies the key type.
	<b>Size</b> - 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 this table.

Control	Description
Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)	<p>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.</p> <p><b>Note:</b> The private key is required.</p> <hr/> <p><b>Local File</b> - Browse to the local file.</p> <hr/> <p><b>Text</b> - Paste the text content of the file into the text box.</p> <hr/> <p><b>Decryption Password</b> - Specify the decryption password, if necessary.</p> <hr/> <p><b>Set</b> - Sets the peer.</p>
Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)	<p>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.</p> <p><b>Note:</b> Importing the private key is optional.</p>
Generate New Private Key and Self-Signed Public Certificate	<p>Select this option to generate a new private key and self-signed public certificate.</p> <hr/> <p><b>Cipher Bits</b> - Select the key length from the drop-down list. The default value is 1024.</p> <hr/> <p><b>Common Name (required)</b> - Specify the hostname of the peer.</p> <hr/> <p><b>Organization Name</b> - Specify the organization name (for example, the company).</p> <hr/> <p><b>Organization Unit Name</b> - Specify the organization unit name (for example, the section or department).</p> <hr/> <p><b>Locality</b> - Specify the city.</p> <hr/> <p><b>State (no abbreviations)</b> - Specify the state.</p> <hr/> <p><b>Country (2-letter code)</b> - Specify the country (2-letter code only).</p> <hr/> <p><b>Email Address</b> - Specify the email address of the contact person.</p> <hr/> <p><b>Validity Period (Days)</b> - Specify how many days the certificate is valid. The default value is 730.</p>

4. To generate a CSR, under Web Certificate, select the Generate CSR tab and complete the configuration as described in this table.

Control	Description
Common Name	Specify the common name (hostname).
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.

Control	Description
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.
7. Click **Add**.

---

## Enabling REST API Access

You enable access to the Riverbed REST API in the Configure > Security > REST API Access page.

REST (REpresentational State Transfer) is a framework for API design. REST builds a simple API on top of the HTTP protocol. It is based on generic facilities of the standard HTTP protocol, including the six basic HTTP methods (GET, POST, PUT, DELETE, HEAD, INFO) and the full range of HTTP return codes. You can discover REST APIs by navigating links embedded in the resources provided by the REST API, which follow common encoding and formatting practices.

You can invoke the REST API to enable communication from one Riverbed appliance to another through REST API calls, for example:

- A Cascade Profiler appliance communicating with a Cascade Shark appliance.
- A Cascade Profiler appliance retrieving a QoS configuration from a Steelhead appliance.

For all uses you must preconfigure an access code to authenticate communication between parties and to authorize access to protected resources.

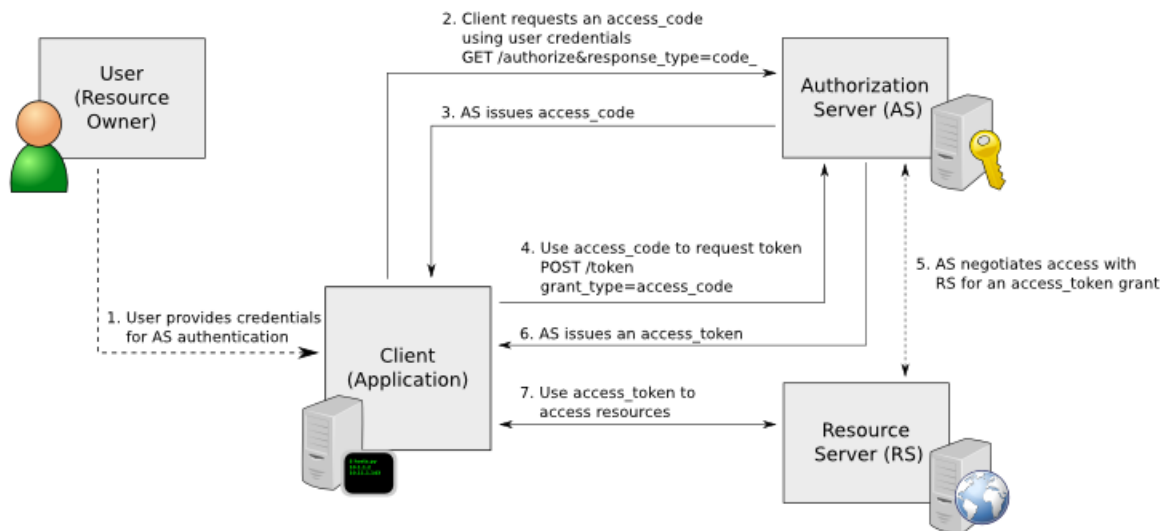
For details, see the *Steelhead Appliance REST API Guide*.

The REST API calls are based on the trusted application flow, a scenario where you download and install an application on some host, such as your own laptop. You trust both the application and the security of the host onto which the application is installed.

For example, suppose you install a Python script on a Linux box that queries QoS policies on a Steelhead appliance and prints a summary as text output. You install the script under your home directory and configure the script with credentials to access the Steelhead appliance. Once set up, you can simply log in to the Linux box and run the script. Because you already preconfigured credentials with the Steelhead appliance, you can run the script without any user interaction after logging in. This enables you to schedule execution through cron, or chain it with other scripts that process the text data and combine it with other functionality.

This basic authentication sequence assumes you have already downloaded the Python script and installed it on a Linux box:

**Figure 7-1. REST API Access Authentication Sequence**



### To enable REST API access

1. Choose **Configure > Security > REST API Access** to display the REST API Access page.
2. Under REST API Access Settings, select the **Enable REST API Access** check box.
3. Click **Apply** to apply your changes to the running configuration.
4. Click **Save** to save your settings permanently.

Before an appliance can access the REST API, you must preconfigure an access code for the system to use to authenticate access.

### To preconfigure the access code

1. Choose **Configure > Security > REST API Access** to display the REST API Access page.
2. Click **Add Access Code**.
3. Under Access Codes, type a description such as the hostname or IP address of the appliance you are using.
4. Create a new code by selecting **Generate New Access Code**, or use an existing code by selecting **Import Existing Access Code**.
5. Click **Add**.  
The access code description appears in the access code table along with the name of the user who created it.
6. Click the access code description to display the access code.



7. Copy the access code from the text field into a text editor such as Notepad.

**To use the access code in your external script**

- Copy the access code copied from the Management Console REST API Access page into the configuration file of your external script. The script uses the access code to make a call to the appliance/system to request an access token. The appliance/system validates the access code and returns an access token for use by the script. Generally the access token is kept by the script for a session only (defined within your script), but note that the script can make many requests using the same access token. These access tokens have some lifetime—usually around an hour—in which they are valid. When they expire, the access code must fetch a new access token. The script uses the access token to make REST API calls with the appliance/system.



## CHAPTER 8      **Configuring System Administrator Settings**

This chapter describes how to configure features to assist you in system administration. It includes these sections:

- [“Configuring Alarm Settings” on page 217](#)
- [“Setting Announcements” on page 224](#)
- [“Configuring Email Settings” on page 224](#)
- [“Configuring Log Settings” on page 226](#)
- [“Configuring the Date and Time” on page 230](#)
- [“Configuring Monitored Ports” on page 234](#)
- [“Configuring SNMP Settings” on page 235](#)

---

### **Configuring Alarm Settings**

You can set alarms in the [Configure > System Settings > Alarms](#) page.

Enabling alarms is *optional*.

RiOS uses hierarchical alarms. The system groups certain alarms into top-level categories, such as the SSL Settings alarm. When an alarm triggers, its parent expands to provide more information. As an example, the System Disk Full top-level parent alarm aggregates over multiple partitions. If a specific partition is full, the System Disk Full parent alarm triggers and the Alarm Status report displays more information regarding which partition caused the alarm to trigger.

Disabling a parent alarm disables its children. You can enable a parent alarm and disable any of its child alarms. You cannot enable a child alarm without first enabling its parent.

The children alarms of a disabled parent appear on the Alarm Status report with a suppressed status. Disabled children alarms of an enabled parent appear on the Alarm Status report with a disabled status. For more details about alarm status, see [“Viewing Alarm Status Reports” on page 307](#).

#### **To set alarm parameters**

1. Choose [Configure > System Settings > Alarms](#) to display the Alarms page.

2. Under Enable Alarms, complete the configuration as described in this table.

Control	Description
Admission Control	<p>Enables an alarm and sends an email notification if the Steelhead appliance enters admission control. When this occurs, the Steelhead appliance optimizes traffic beyond its rated capability and is unable to handle the amount of traffic passing through the WAN link. During this event, the Steelhead appliance continues to optimize existing connections, but new connections are passed through without optimization.</p> <ul style="list-style-type: none"> <li>• <b>Connection Limit</b> - 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.</li> <li>• <b>CPU</b> - 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. The alarm clears automatically when the CPU usage has decreased.</li> <li>• <b>Memory</b> - 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; the alarm clears automatically when the traffic has decreased.</li> <li>• <b>TCP</b> - 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. The alarm clears automatically when the TCP memory pressure has decreased.</li> </ul> <p>By default, this alarm is enabled.</p>
Asymmetric Routing	<p>Enables an alarm if asymmetric routing is detected on the network. This is usually due to a failover event of an inner router or VPN.</p> <p>By default, this alarm is enabled.</p>
Connection Forwarding	<p>Enables an alarm if the system detects 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.</p> <ul style="list-style-type: none"> <li>• <b>Cluster IPv6 Incompatible</b> - Enables an alarm and sends an email notification if a connection-forwarding neighbor is running a RiOS version that is incompatible with IPv6, or if the IP address configuration between neighbors does not match. Neighbors must be running RiOS v8.5.</li> <li>• <b>Multiple Interface</b> - Enables an alarm and sends an email notification if the connection to a Steelhead appliance in a connection forwarding cluster is lost.</li> <li>• <b>Single Interface</b> - Enables an alarm and sends an email notification if the connection to a Steelhead appliance connection-forwarding neighbor is lost.</li> </ul> <p>By default, this alarm is enabled.</p>
CPU Utilization	<p>Enables an alarm and sends an email notification 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.</p> <p>By default, this alarm is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Rising Threshold</b> - Specify the rising threshold. When an alarm reaches the rising threshold, it is activated. The default value is 90 percent.</li> <li>• <b>Reset Threshold</b> - 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 percent.</li> </ul>

Control	Description
Disk Full	<p>Enables an alarm if the system partitions (not the RiOS data store) are full or almost full: for example, RiOS monitors the available space on <b>/var</b> which is used to hold logs, statistics, system dumps, TCP dumps, and so on.</p> <p>By default, this alarm is enabled.</p>
Flash Protection Failure	<p>Enables an alarm if the USB flash drive has not been backed up because there is not enough available space in the <b>/var</b> filesystem directory.</p>
Hardware	<ul style="list-style-type: none"> <li>• <b>Disk Error</b> - Enables an alarm when one or more disks is offline. To see which disk is offline, enter this CLI command from the system prompt:  <pre>show raid diagram</pre> <p>By default, this alarm is enabled.</p> <p>This alarm applies only to the Steelhead appliance RAID Series 3000, 5000, and 6000.</p></li> <li>• <b>Fan Error</b> - Enables an alarm and sends an email notification if a fan is failing or has failed and needs to be replaced. By default, this alarm is enabled.</li> <li>• <b>Flash Error</b> - Enables an alarm when the system detects an error with the flash drive hardware. By default, this alarm is enabled.</li> <li>• <b>IPMI</b> - Enables an alarm and sends an email notification if an Intelligent Platform Management Interface (IPMI) event is detected. (Not supported on all appliance models.)</li> </ul> <p>This alarm triggers when there has been a physical security intrusion. These events trigger this alarm:</p> <ul style="list-style-type: none"> <li>• Chassis intrusion (physical opening and closing of the appliance case)</li> <li>• Memory errors (correctable or uncorrectable ECC memory errors)</li> <li>• Hard drive faults or predictive failures</li> <li>• Power cycle, such as turning the power switch on or off, physically unplugging and replugging the cable, or issuing a power cycle from the power switch controller.</li> </ul> <p>By default, this alarm is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Memory Error</b> - Enables an alarm and sends an email notification if a memory error is detected, for example, when a system memory stick fails.</li> <li>• <b>Other Hardware Error</b> - Enables an alarm if a hardware error is detected. These issues trigger the hardware error alarm: <ul style="list-style-type: none"> <li>• The Steelhead appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration.</li> <li>• 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.</li> <li>• DIMMs are plugged into the Steelhead appliance but RiOS cannot recognize them because: <ul style="list-style-type: none"> <li>– a DIMM is in the wrong slot. You must plug DIMMs into the black slots first and then use the blue slots when all of the black slots are in use.</li> <li>—or—</li> <li>– a DIMM is broken and you must replace it.</li> </ul> </li> <li>• Other hardware issues</li> </ul> <p>By default, this alarm is enabled.</p> <li>• <b>Power Supply</b> - Enables an alarm and sends an email notification if an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted. By default, this alarm is enabled.</li> </li></ul>

Control	Description
Inbound QoS WAN Bandwidth Configuration	<p>Enables an alarm and sends an email notification if the inbound QoS WAN bandwidth for one or more of the interfaces is set incorrectly. You must configure the WAN bandwidth to be less than or equal to the interface bandwidth link rate.</p> <p>This alarm triggers when the system encounters one of these conditions:</p> <ul style="list-style-type: none"> <li>• An interface is connected and the WAN bandwidth is set higher than its bandwidth link rate: for example, if the bandwidth link rate is 1536 kbps, and the WAN bandwidth is set to 2000 kbps.</li> <li>• A non-zero WAN bandwidth is set and QoS is enabled on an interface that is disconnected; that is, the bandwidth link rate is 0.</li> <li>• A previously disconnected interface is reconnected, and its previously configured WAN bandwidth was set higher than the bandwidth link rate. The Management Console refreshes the alarm message to inform you that the configured WAN bandwidth is set higher than the interface bandwidth link rate.</li> </ul> <p>While this alarm appears, the Steelhead appliance puts existing connections into the default class.</p> <p>The alarm clears when you configure the WAN bandwidth to be less than or equal to the bandwidth link rate or reconnect an interface configured with the correct WAN bandwidth.</p> <p>By default, this alarm is enabled.</p>
Licensing	<p>Enables an alarm and sends an email notification if a license on the Steelhead appliance is removed, is about to expire, has expired, or is invalid. This alarm triggers if the Steelhead appliance has no MSPEC license installed for its currently configured model.</p> <ul style="list-style-type: none"> <li>• <b>Appliance Unlicensed</b> - This alarm triggers if the Steelhead appliance has no BASE or MSPEC license installed for its currently configured model. For details about updating licenses, see <a href="#">“Managing Licenses and Model Upgrades” on page 191</a>.</li> <li>• <b>Autolicense Critical Event</b> - This alarm triggers on a Virtual Steelhead appliance when the Riverbed Licensing Portal cannot respond to a license request with valid licenses. The Licensing Portal cannot issue a valid license for one of these reasons: <ul style="list-style-type: none"> <li>– A newer Virtual Steelhead appliance is already using the token, so you cannot use it on the Virtual Steelhead appliance displaying the critical alarm. Every time the Virtual Steelhead appliance attempts to refetch a license token, the alarm retriggers.</li> <li>– The token has been redeemed too many times. Every time the Virtual Steelhead appliance attempts to refetch a license token, the alarm retriggers.</li> </ul> </li> <li>• <b>Autolicense Informational Event</b> - This alarm triggers if the Riverbed Licensing Portal has information regarding the licenses for a Virtual Steelhead appliance. For example, the Virtual Steelhead appliance displays this alarm when the portal returns licenses that are associated with a token that has been used on a different Virtual Steelhead appliance.</li> <li>• <b>Licenses Expired</b> - This alarm triggers if one or more features has at least one license installed, but all of them are expired.</li> <li>• <b>Licenses Expiring</b> - This alarm triggers if the license for one or more features is going to expire within two weeks.</li> </ul> <p><b>Note:</b> The licenses expiring and licenses expired alarms are triggered per feature; for example: if you install two license keys for a feature, LK1-FOO-xxx (expired) and LK1-FOO-yyy (not expired), the alarms do not trigger, because the feature has one valid license.</p> <p>By default, this alarm is enabled.</p>

Control	Description
Link Duplex	<p>Enables an alarm and sends an email notification when an interface was not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex alarm.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable an alarm, choose Configure &gt; System Settings&gt; alarms and select or clear the check box next to the link name.</p>
Link I/O Errors	<p>Enables an alarm and sends an email notification when the link error rate exceeds 0.1 percent while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured LAN connection experiences very few errors.</p> <p>The alarm clears when the rate drops below 0.05 percent.</p> <p>You can change the default alarm thresholds by entering the <b>alarm link_errors err-threshold xxxxx</b> CLI command at the system prompt. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific interface; for example, you can disable the alarm for a link after deciding to tolerate the errors. To enable or disable an alarm, choose Configure &gt; System Settings &gt; alarms and select or clear the check box next to the link name.</p>
Link State	<p>Enables an alarm and sends an email notification if an Ethernet link is lost due to an unplugged cable or dead switch port. Depending on which link is down, the system might no longer be optimizing and a network outage could occur.</p> <p>This condition is often caused by surrounding devices, like routers or switches, interface transitioning. This alarm also accompanies service or system restarts on the Steelhead appliance.</p> <p>For WAN/LAN interfaces, the alarm triggers if in-path support is enabled for that WAN/LAN pair.</p> <p>By default, this alarm is disabled.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable an alarm, choose Configure &gt; System Settings&gt; alarms and select or clear the check box next to the link name.</p>
Memory Paging	<p>Enables an alarm and sends an email notification if memory paging is detected. 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 <a href="https://support.riverbed.com">https://support.riverbed.com</a>.</p> <p>By default, this alarm is enabled.</p>
Neighbor Incompatibility	<p>Enables an alarm if the system has encountered an error in reaching a Steelhead appliance configured for connection forwarding.</p> <p>By default, this alarm is enabled.</p>
Network Bypass	<p>Enables an alarm and sends an email notification if the system is in bypass failover mode.</p> <p>By default, this alarm is enabled.</p>

Control	Description
Optimization Service	<ul style="list-style-type: none"> <li>• <b>Internal Error</b> - Enables an alarm and sends an email notification if the RiOS optimization service encounters a condition that might degrade optimization performance. By default, this alarm is enabled. Go to the Configure &gt; Maintenance &gt; Services page and restart the optimization service.</li> <li>• <b>Service Status</b> - Enables an alarm and sends an email notification if the RiOS optimization service encounters a service condition. By default, this alarm is enabled. The message indicates the reason for the condition. These conditions trigger this alarm: <ul style="list-style-type: none"> <li>• Configuration errors.</li> <li>• A Steelhead appliance reboot.</li> <li>• A system crash.</li> <li>• An optimization service restart.</li> <li>• A user enters the CLI command <b>no service enable</b> or shuts down the optimization service from the Management Console.</li> <li>• A user restarts the optimization service from either the Management Console or CLI.</li> </ul> </li> <li>• <b>Unexpected Halt</b> - Enables an alarm and sends an email notification if the RiOS optimization service halts due to a serious software error. By default, this alarm is enabled.</li> </ul>
Outbound QoS WAN Bandwidth Configuration	<p>Enables an alarm and sends an email notification if the outbound QoS WAN bandwidth for one or more of the interfaces is set incorrectly. You must configure the WAN bandwidth to be less than or equal to the interface bandwidth link rate.</p> <p>This alarm triggers when the system encounters one of these conditions:</p> <ul style="list-style-type: none"> <li>• an interface is connected and the WAN bandwidth is set to higher than its bandwidth link rate: for example, if the bandwidth link rate is 100 Mbps, and the WAN bandwidth is set to 200 Mbps.</li> <li>• a non-zero WAN bandwidth is set and QoS is enabled on an interface that is disconnected; that is, the bandwidth link rate is 0.</li> <li>• a previously disconnected interface is reconnected, and its previously configured WAN bandwidth was set higher than the bandwidth link rate. The Management Console refreshes the alarm message to inform you that the configured WAN bandwidth is set greater than the interface bandwidth link rate.</li> </ul> <p>While this alarm appears, the system puts existing connections into the default class.</p> <p>The alarm clears when you configure the WAN bandwidth to be less than or equal to the bandwidth link rate or reconnect an interface configured with the correct WAN bandwidth.</p> <p>By default, this alarm is enabled.</p>
Path Selection Path Down	<p>Enables an alarm and sends an email notification if the system detects that one of the predefined paths for a connection is unavailable. The path has exceeded either the timeout value for path latency or the threshold for observed packet loss.</p> <p>When a path fails, the Steelhead appliance directs traffic through another available path. When the original path comes back up, the Steelhead appliance redirects the traffic back to it.</p> <p>By default, this alarm is enabled.</p>
Process Dump Creation Error	<p>Enables an alarm and sends an email notification if the system detects an error while trying to create a process dump. This alarm indicates an abnormal condition where RiOS cannot collect the core file after three retries. It can be caused when the /var directory is reaching capacity or other conditions. When the alarm is raised, the directory is blacklisted.</p> <p>By default, this alarm is enabled.</p>



Control	Description
Secure Vault	<p>Enables an alarm and sends an email notification if the system encounters a problem with the secure vault:</p> <ul style="list-style-type: none"> <li>• <b>Secure Vault Locked</b> - Indicates that the secure vault is locked. Go to Configure &gt; Security &gt; Secure Vault and unlock the secure vault.</li> <li>• <b>Secure Vault New Password Recommended</b> - Indicates that the secure vault requires a new, nondefault password. Re-enter the password.</li> <li>• <b>Secure Vault Not Initialized</b> - Indicates that an error has occurred while initializing the secure vault. When the vault is locked, SSL traffic is not optimized and you cannot encrypt the RiOS data store. For details, see <a href="#">“Unlocking the Secure Vault” on page 206</a>.</li> </ul>
Software Compatibility	<p>Enables an alarm and sends an email notification if the system encounters a problem with software compatibility:</p> <ul style="list-style-type: none"> <li>• <b>Peer Mismatch - Needs Attention</b> - Indicates that 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.</li> <li>• <b>Software Version Mismatch - Degraded</b> - Indicates that the appliance is running an incompatible version of system software. To resolve the problem, upgrade your system software. No other action is required as the alarm clears automatically.</li> </ul> <p>By default, this alarm is enabled.</p>
SSL	<p>Enables an alarm if an error is detected in your SSL configuration.</p> <ul style="list-style-type: none"> <li>• <b>SSL Peering Certificates Expiring</b> - Indicates that an SSL certificate is about to expire.</li> <li>• <b>SSL Certificates SCEP</b> - Indicates that an SSL certificate has failed to re-enroll automatically within the SCEP polling interval.</li> </ul> <p>By default, this alarm is enabled.</p>
System Detail Report	<p>Enables an alarm if a system component has encountered a problem.</p> <p>By default, this alarm is disabled (RiOS v7.0.3 and later).</p>
Temperature	<ul style="list-style-type: none"> <li>• <b>Critical Temperature</b> - Enables an alarm and sends an email notification if the CPU temperature exceeds the rising threshold. When the CPU returns to the reset threshold, the critical alarm is cleared. The default value for the rising threshold temperature is 70° C; the default reset threshold temperature is 67° C.</li> <li>• <b>Warning Temperature</b> - Enables an alarm and sends an email notification if the CPU temperature approaches the rising threshold. When the CPU returns to the reset threshold, the warning alarm is cleared. <ul style="list-style-type: none"> <li>• <b>Rising Threshold</b> - Specifies the rising threshold. The alarm activates when the temperature exceeds the rising threshold. The default value is 70 percent.</li> <li>• <b>Reset Threshold</b> - Specifies the reset threshold. The alarm clears when the temperature falls below the reset threshold. The default value is 67 percent.</li> </ul> </li> </ul> <p>After the alarm triggers, it cannot trigger again until after the temperature falls below the reset threshold and then exceeds the rising threshold again.</p>

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 224](#)
- [“Configuring SNMP Settings” on page 235](#)
- [“Viewing Process Dumps” on page 330](#)

---

## 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.
2. Use the controls to complete the configuration as described in this 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 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 page.
2. Under Email Notification, complete the configuration as described in this 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.  <b>Important:</b> 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. Typically you do not need to change the default port 25.

Control	Description
Report Events via Email	<p>Select this option to report alarm events through email. Specify a list of email addresses to receive the notification messages. Separate addresses by spaces, semicolons, commas, or vertical bars.</p> <p>These alarms are events:</p> <ul style="list-style-type: none"> <li>• Admission control</li> <li>• CPU utilization (rising threshold, reset threshold)</li> <li>• Temperature (rising threshold, reset threshold)</li> <li>• Data store wrap frequency</li> <li>• Network interface duplex errors</li> <li>• Network interface link errors</li> <li>• Fan error</li> <li>• Flash error</li> <li>• Hardware error</li> <li>• IPMI</li> <li>• Licensing</li> <li>• Memory error</li> <li>• Neighbor incompatibility</li> <li>• Network bypass</li> <li>• Non-SSL servers detected on upgrade</li> <li>• Optimization service (general service status, optimization service)</li> <li>• Extended memory paging activity</li> <li>• Secure vault</li> <li>• System disk full</li> <li>• Software version mismatch</li> <li>• TCP Stop Trigger scan has started</li> <li>• Asymmetric routes</li> <li>• Expiring SSL peering certificates</li> <li>• SSL peering certificate SCEP automatic re-enrollment</li> <li>• Connection forwarding (ACK timeout, failure, lost EOS, lost ERR, keepalive timeout, latency exceeded, read info timeout)</li> </ul>
Report Failures via Email	<p>Select this option to report alarm failures through email. Specify a list of email addresses to receive the notification messages. Separate addresses by spaces, semicolons, commas, or vertical bars.</p> <p>These alarms are failures:</p> <ul style="list-style-type: none"> <li>• Data store corruption</li> <li>• System details report</li> <li>• Domain join error</li> <li>• RAID</li> <li>• Optimization service - unexpected halt</li> <li>• Critical temperature</li> <li>• Disk error</li> </ul>

Control	Description
Override Default Sender's Address	<p>Select this option to configure the SMTP protocol for outgoing server messages for errors or events. Specify a list of email addresses to receive the notification messages. Separate addresses by commas.</p> <p>You can also configure the outgoing email address sent to the client recipients. The default outgoing address is do-not-reply@hostname.domain. If you do not specify a domain the default outgoing email is do-not-reply@hostname.</p> <p>You can configure the host and domain settings in the Configure &gt; Networking &gt; Host Settings page.</p>
Report Failures to Technical Support	<p>Select this option to report serious failures such as system crashes to Riverbed Support. Riverbed recommends that you activate this feature so that problems are promptly corrected.</p> <p><b>Important:</b> 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.</p>

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 217](#)

## **Configuring Log Settings**

You set up local and remote logging in the Configure > System Settings > Logging page.

By default, the system rotates each log file every 24 hours or if the file size reaches one Gigabyte uncompressed. You can change this to rotate every week or month and you can rotate the files based on file size.

The automatic rotation of system logs deletes your oldest log file, labeled as Archived log #10, pushes the current log to Archived log # 1, and starts a new current-day log file.

### **To set up logging**

1. Choose Configure > System Settings > Logging to display the Logging page.
2. To rotate the logs manually, under Log Actions, click **Rotate Logs**. After the logs are rotated, this message appears:

“logs have been successfully rotated”

When you click **Rotate Logs**, your archived file #1 contains data for a partial day because you are writing a new log before the current 24-hour period is complete.

3. Under Logging Configuration, complete the configuration as described in this table.

Control	Description
Minimum Severity	<p>Select the minimum severity level for the system log messages. The log contains all messages with this severity level or higher. Select one of these levels from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Emergency</b> - Emergency, the system is unusable.</li> <li>• <b>Alert</b> - Action must be taken immediately.</li> <li>• <b>Critical</b> - Conditions that affect the functionality of the Steelhead appliance.</li> <li>• <b>Error</b> - Conditions that probably affect the functionality of the Steelhead appliance.</li> <li>• <b>Warning</b> - Conditions that could affect the functionality of the Steelhead appliance, such as authentication failures.</li> <li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change. This is the default setting.</li> <li>• <b>Info</b> - Informational messages that provide general information about system operations.</li> </ul> <p><b>Note:</b> This control applies to the system log only. It does not apply to the user log.</p>
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	<p>Specifies the rotation option:</p> <ul style="list-style-type: none"> <li>• <b>Time</b> - Select Day, Week, or Month from the drop-down list. The default setting is Day.</li> <li>• <b>Disk Space</b> - Specify how much disk space, in megabytes, the log uses before it rotates. The default value is 16 MB.</li> </ul> <p><b>Note:</b> 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.</p>

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 this 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 these levels from the drop-down list: <ul style="list-style-type: none"><li>• <b>Emergency</b> - Emergency, the system is unusable.</li><li>• <b>Alert</b> - Action must be taken immediately.</li><li>• <b>Critical</b> - Conditions that affect the functionality of the Steelhead appliance.</li><li>• <b>Error</b> - Conditions that probably affect the functionality of the Steelhead appliance.</li><li>• <b>Warning</b> - Conditions that could affect the functionality of the Steelhead appliance, such as authentication failures.</li><li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change. This is the default setting.</li><li>• <b>Info</b> - Informational messages that provide general information about system operations.</li></ul>
Add	Adds the server to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

You can 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.

2. Under Per-Process Logging, complete the configuration as described in this table.

Control	Description
Add a New Process Logging Filter	Displays the controls for adding a process level logging filter.
Process	<p>Select a process to include in the log from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>alarmd</b> - alarm manager, which processes all alarms, including their thresholds and severity.</li> <li>• <b>collectord</b> - Application visibility NetFlow collector.</li> <li>• <b>cli</b> - Command-line interface.</li> <li>• <b>mgmtd</b> - 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 <b>hald</b>.</li> <li>• <b>hald</b> - Hardware abstraction daemon, which handles access to the hardware.</li> <li>• <b>pm</b> - Process manager, which handles launching of internal system daemons and keeps them running.</li> <li>• <b>qosd</b> - QoS scheduler and DPI engine.</li> <li>• <b>sched</b> - Process scheduler, which handles one-time scheduled events.</li> <li>• <b>shark</b> - Embedded Cascade Shark, which enables Cascade Pilot to perform remote packet analysis on trace files captured and stored on this Steelhead appliance without having to transfer large packet capture files across the network.</li> <li>• <b>statsd</b> - Statistics collector, which handles queries, storage, and trending of system statistics.</li> <li>• <b>wdt</b> - Watchdog timer, the motherboard watchdog daemon.</li> <li>• <b>webasd</b> - Web application process, which handles the Web user interface.</li> </ul>
Minimum Severity	<p>Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select a level from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Emergency</b> - Emergency, the system is unusable.</li> <li>• <b>Alert</b> - Action must be taken immediately.</li> <li>• <b>Critical</b> - Conditions that affect the functionality of the Steelhead appliance.</li> <li>• <b>Error</b> - Conditions that probably affect the functionality of the Steelhead appliance.</li> <li>• <b>Warning</b> - Conditions that could affect the functionality of the Steelhead appliance, such authentication failures.</li> <li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change.</li> <li>• <b>Info</b> - Informational messages that provide general information about system operations.</li> </ul>
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 <b>Remove Selected</b> to remove the filter.

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

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

---

## Configuring the Date and Time

You set the system date and time in the Configure > System Settings > Date and Time page.

You can either set the system date and time by entering it manually or assigning an NTP server to the Steelhead appliance. By default, the appliance uses the Riverbed-provided NTP server and these public NTP servers:

- 0.riverbed.pool.ntp.org
- 1.riverbed.pool.ntp.org
- 2.riverbed.pool.ntp.org
- 3.riverbed.pool.ntp.org

### To set the date and time manually

1. Choose Configure > System Settings > Date and Time to display the Date and Time page.
2. Under Date and Time, select **Set Time Manually**.
3. Complete the configuration as described in this table.

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

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

### To use Network Time Protocol (NTP) time synchronization

1. Choose Configure > System Settings > Date and Time to display the Date and Time page.
2. Under Date and Time, select **Use NTP Time Synchronization**.

As a best practice, configure your own internal NTP servers; however, you can use the Riverbed-provided NTP server and public NTP servers. The hard-coded IP address that is preconfigured into every Steelhead appliance is 208.70.196.25. This IP address and the public NTP servers are enabled by default and appear in the requested NTP server list.

## Current NTP Server Status

NTP server state information appears in these server tables:

- **Requested NTP server table** - displays all of the configured NTP server addresses.



- **Connected NTP server table** - displays all of the servers to which the Steelhead appliance is actually connected.

When you request a connection to an NTP server in a public NTP server pool, the server IP address does not map to the actual NTP server to which the Steelhead connects. For example, if you request \*.riverbed.pool.ntp.org, querying the pool address does not return the IP address of the pool hostname, but instead returns the IP address of an NTP server within its pool. For example, when resolving 0.riverbed.pool.ntp.org returns the first NTP server, the connected NTP server table displays the IP address of this first NTP server.

This information appears after an NTP server name:

- Authentication information; unauthenticated appears after the server name when it is not using authentication.
- When RiOS has no NTP information about the current server, nothing appears.

## NTP Authentication

NTP authentication verifies the identity of the NTP server sending timing information to the Steelhead appliance. RiOS v8.5 supports MD5-based Message-Digest Algorithm symmetric keys and Secure Hash Algorithm (SHA1) for NTP authentication. MD5 is a widely used cryptographic hash function that produces a 128-bit (16-byte) hash value. SHA1 is a set of related cryptographic hash functions. SHA1 is considered to be the successor to MD5.

NTP authentication is *optional*.

Configuring NTP authentication involves these steps that you can perform in any order:

- Configure a key ID and a secret pair.
- Configure the key type.
- Configure the NTP server with the key ID.

## NTP Servers

The default NTP configuration points to the Riverbed-provided NTP server IP address 208.70.196.25 and these public NTP servers:

- 0.riverbed.pool.ntp.org
- 1.riverbed.pool.ntp.org
- 2.riverbed.pool.ntp.org
- 3.riverbed.pool.ntp.org

Riverbed recommends synchronizing the Steelhead appliance to an NTP server of your choice.

### To add an NTP server

1. Choose **Configure > System Settings > Date and Time** to display the Date and Time page.
2. Complete the configuration as described in this 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. You can connect to an NTP public server pool; for example, 0.riverbed.pool.ntp.org. When you add an NTP server pool, the server is selected from a pool of time servers.
Version	Select the NTP server version from the drop-down list: 3 or 4.
Enabled/Disabled	Select Enabled from the drop-down list to connect to the NTP server. Select Disabled from the drop-down list to disconnect from the NTP server.
Key ID	Specify the MD5 or SH1 key identifier to use to authenticate the NTP server. The valid range is from 1 - 65534. The key ID must appear on the trusted keys list.
Add	Adds the NTP server to the server list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

### NTP Authentication Keys

NTP authentication uses a key and a shared secret to verify the identity of the NTP server sending timing information to the Steelhead appliance. RiOS encrypts the shared secret text using MD5 or SHA1, and uses the authentication key to access the secret.

#### To add an NTP authentication key

1. Choose **Configure > System Settings > Date and Time** to display the Date and Time page.

2. Complete the configuration as described in this table.

Control	Description
Add a New NTP Authentication Key	Displays the controls to add an authentication key to the key list. Both trusted and untrusted keys appear on the list.
Key ID	Optionally, specify the secret MD5 or SHA1 key identifier for the NTP server. The valid range is from 1 - 65534.
Key Type	Select the authentication key type: MD5 or SHA1.
Secret	<p>Specify the shared secret. You must configure the same shared secret for both the NTP server and the NTP client.</p> <p>The MD5 shared secret</p> <ul style="list-style-type: none"> <li>• is limited to 16 alphanumeric characters or less, or exactly 40-characters hexadecimal.</li> <li>• cannot include white space or #s</li> <li>• cannot be empty</li> <li>• is case sensitive</li> </ul> <p>The SHA1 shared secret</p> <ul style="list-style-type: none"> <li>• is limited to exactly 40-characters hexadecimal</li> <li>• cannot include white space or #s</li> <li>• cannot be empty</li> <li>• is case sensitive</li> </ul> <p>The secret appears in the key list as its MD5 or SHA1 hash value.</p>
Add	Adds the authentication key to the trusted keys list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

### ***NTP Key Information***

NTP keys appear in a list that includes the key ID, type, secret (displays as the MD5 or SHA1 hash value), and whether RiOS trusts the key for authentication.

You can only remove a key from the trust list using the CLI command **ntp authentication trustedkeys**. For details, see the *Riverbed Command-Line Interface Reference Manual*.

## Configuring Monitored Ports

You set the TCP ports 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 284](#).

By default, traffic is monitored on ports 21 (FTP), 80 (HTTP), 135 (EPM), 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), 7919 (IP Blade), 8777 (RCU), 8778 (SMB Signed), 8779 (SMB2), 8780 (SMB2 Signed), and 10566 (SnapMirror).

### To set monitored ports

1. Choose Configure > System Settings > Monitored Ports to display the Monitored Ports page.
2. Complete the configuration as described in this 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 <b>Remove Selected</b> .

3. To modify a monitored port, click the magnifying glass next to the port and complete the configuration as described in this table.

Control	Description
Port Description	Specify a description of the type of traffic on the port.
Apply	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 provides support for these SNMP versions:

- Version 1
- Version 2c
- SNMP Version 3, which provides authentication through the User-based Security Model (USM).
- View-Based Access Control Mechanism (VACM), which provides richer access control.
- SNMP Version 3 authentication using AES 128 and DES encryption privacy.

You set the default community string on the SNMP Basic page. To set more than one SNMP community string, see the Riverbed Knowledge Base article *Can I Have More Than One SNMP Community String?* <https://supportkb.riverbed.com/support/index?page=content&id=S16345>

### To set general SNMP parameters

1. Choose Configure > System Settings > SNMP Basic to display the SNMP Basic page.
2. Under SNMP Server Settings, complete the configuration as described in this table.

Control	Description
Enable SNMP Traps	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. Community strings cannot contain the # (hash) value.

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 this table.

Control	Description
Add a New Trap Receiver	Displays the controls to add a new trap receiver.
Receiver	Specify the destination IPv4 or IPv6 address or hostname for the SNMP trap.
Destination Port	Specify the destination port.

Control	Description
Receiver Type	Select SNMP version v1, v2c, or v3 (user-based security model).
Remote User	(Appears only when you select v3.) Specify a remote user name.
Authentication	(Appears only when you select v3.) Optionally, select either Supply a Password or Supply a Key to use while authenticating users.
Authentication Protocol	<p>(Appears only when you select v3.) Select an authentication method from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>MD5</b> - Specifies the Message-Digest 5 algorithm, a widely used cryptographic hash function with a 128-bit hash value. This is the default value.</li> <li>• <b>SHA</b> - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.</li> </ul>
Password/Password Confirm	(Appears only when you select v3 and Supply a Password.) Specify a password. The password must have a minimum of eight characters. Confirm the password in the Password Confirm text box.
Security Level	<p>(Appears only when you select v3.) Determines whether a single atomic message exchange is authenticated. Select one of these levels from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>No Auth</b> - Does not authenticate packets and does not use privacy. This is the default setting.</li> <li>• <b>Auth</b> - Authenticates packets but does not use privacy.</li> <li>• <b>AuthPriv</b> - Authenticates packets using AES 128 and DES to encrypt messages for privacy.</li> </ul> <p><b>Note:</b> A security level applies to a group, not to an individual user.</p>
Privacy Protocol	(Appears only when you select v3 and AuthPriv.) Select either the AES or DES protocol from the drop-down list. AES uses the AES128 algorithm.
Privacy	(Appears only when you select v3 and AuthPriv.) Select Same as Authentication Key, Supply a Password, or Supply a Key to use while authenticating users. The default setting is Same as Authentication Key.
Privacy Password	(Appears only when you select v3 and Supply a Password.) Specify a password. The password must have a minimum of eight characters. Confirm the password in the Privacy Password Confirm text box.
MD5/SHA Key	(Appears only when you select v3 and Authentication as Supply a Key.) Specify a unique authentication key. The key is either a 32-hexadecimal digit MD5 or a 40-hexadecimal digit SHA digest created using md5sum or sha1sum.
Privacy MD5/SHA Key	(Appears only when you select v3 and Privacy as Supply a Key.) Specify the privacy authentication key. The key is either a 32-hexadecimal digit MD5 or a 40-hexadecimal digit SHA digest created using md5sum or sha1sum.
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, a password, and a security level.
Enable Receiver	Select to enable the new trap receiver. Clear to disable the receiver.
Add	Adds a new trap receiver to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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.

RiOS v7.0 and later supports SNMPv3 message encryption for increased security.

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 is 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.
2. Under Users, complete the configuration as described in this 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: <ul style="list-style-type: none"> <li>• <b>MD5</b> - Specifies the Message-Digest 5 algorithm, a widely used cryptographic hash function with a 128-bit hash value. This is the default value.</li> <li>• <b>SHA</b> - Specifies the Secure Hash Algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.</li> </ul>
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.
Use Privacy Option	Select to use SNMPv3 encryption.
Privacy Protocol	Select either the AES or DES protocol from the drop-down list. AES uses the AES128 algorithm.

Control	Description
Privacy	Select Same as Authentication, Supply a Password, or Supply a Key to use while authenticating users. The default setting is Same as Authentication.
Privacy Password	(Appears only when you select Supply a Password.) Specify a password. The password must have a minimum of eight characters. Confirm the password in the Privacy 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.
MD5/SHA Key	(Appears only when you select Supply a Key.) Specify a unique authentication key. The key is either a 32-hexadecimal digit MD5 or a 40-hexadecimal digit SHA digest created using md5sum or sha1sum.
Add	Adds the user.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

3. 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.
- **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 composed 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.



2. Under Security Names, complete the configuration as described in this table.

Control	Description
Add a New Security Name	Displays the controls to add a security name.
Security Name	<p>Specify a name to identify a requestor allowed to issue gets and sets (v1 and v2c only). The specified requestor can make changes to the view-based access-control model (VACM) security name configuration.</p> <p>This control does not apply to SNMPv3 queries. To restrict v3 USM users from polling a particular subnet, use the RiOS Management ACL feature, located in the Configure &gt; Security &gt; Management ACL page.</p> <p>Traps for v1 and v2c are independent of the security name.</p>
Community String	<p>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.</p> <p>Community strings do not allow printable 7-bit ASCII characters, except for white spaces. Also, the community strings cannot begin with '#' and '-'.</p> <p>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.</p> <p>To create multiple SNMP community strings on a Steelhead appliance, 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.</p>
Source IP Address and Mask Bits	Specify the host IPv4 or IPv6 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 <b>Remove Selected</b> .

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.
2. Under Groups, complete the configuration as described in this 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	<p>Click the + button and select a security model from the drop-down list:</p> <ul style="list-style-type: none"> <li>v1 or v2c - displays another drop-down menu; select a security name.</li> <li>v3 (usm) - displays another drop-down menu, select a user.</li> </ul> <p>To add another Security Model and Name pair, click the plus sign (+).</p>

Control	Description
Add	Adds the group name and security model and name pairs.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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.

2. Under Views, complete the configuration as described in this 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.
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 <b>Remove Selected</b> .

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.

2. Under Access Policies, complete the configuration as described in this 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.

Control	Description
Security Level	<p>Determines whether a single atomic message exchange is authenticated. Select one of these from the drop-down list:</p> <ul style="list-style-type: none"><li>• <b>No Auth</b> - Does not authenticate packets and does not use privacy. This is the default setting.</li><li>• <b>Auth</b> - Authenticates packets but does not use privacy.</li><li>• <b>AuthPriv</b> - Authenticates packets using AES or DES to encrypt messages for privacy.</li></ul> <p>A security level applies to a group, not to an individual user.</p>
Read View	Select a view from the drop-down list.
Add	Adds the policy to the policy list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

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



## CHAPTER 9 Viewing Reports and Logs

This chapter describes how to display system reports and user and system logs to evaluate performance or troubleshoot. It includes these sections:

[“Overview” on page 244](#)

### Networking Reports

- [“Viewing Current Connection Reports” on page 248](#)
- [“Viewing Connection History Reports” on page 268](#)
- [“Viewing Connection Forwarding Reports” on page 271](#)
- [“Viewing Outbound QoS Reports” on page 273](#)
- [“Viewing Inbound QoS Reports” on page 275](#)
- [“Viewing Path Selection Reports” on page 277](#)
- [“Viewing Top Talkers Reports” on page 281](#)
- [“Viewing Traffic Summary Reports” on page 284](#)
- [“Viewing WAN Throughput Reports” on page 287](#)
- [“Viewing Application Statistics Reports” on page 289](#)
- [“Viewing Application Visibility Reports” on page 291](#)
- [“Viewing Interface Counter Reports” on page 294](#)
- [“Viewing TCP Statistics Reports” on page 295](#)

### Optimization Reports

- [“Viewing Optimized Throughput Reports” on page 296](#)
- [“Viewing Bandwidth Optimization Reports” on page 299](#)
- [“Viewing Peer Reports” on page 301](#)
- [“Viewing SRDF Reports” on page 302](#)
- [“Viewing SRDF Reports” on page 302](#)
- [“Viewing SRDF Reports” on page 302](#)
- [“Viewing SnapMirror Reports” on page 304](#)

### Diagnostic Reports

- [“Viewing Alarm Status Reports” on page 307](#)
- [“Viewing CPU Utilization Reports” on page 317](#)
- [“Viewing Memory Paging Reports” on page 319](#)
- [“Viewing TCP Memory Reports” on page 320](#)
- [“Viewing System Details Reports” on page 323](#)
- [“Checking Network Health Status” on page 323](#)
- [“Viewing Logs” on page 326](#)
- [“Downloading Log Files” on page 328](#)
- [“Generating System Dumps” on page 329](#)
- [“Viewing Process Dumps” on page 330](#)
- [“Capturing and Uploading TCP Dump Files” on page 331](#)
- [“Exporting Performance Statistics” on page 339](#)

---

## Overview

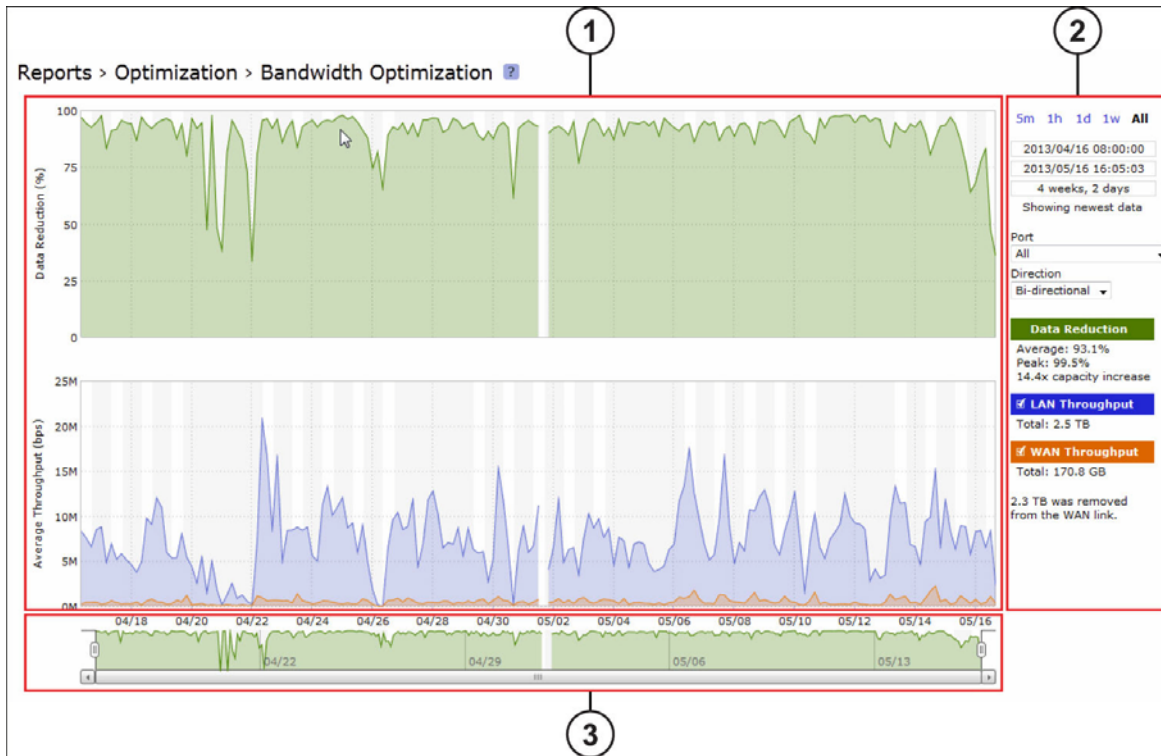
This section describes the report format basics, before describing individual reports.

All of the time-series reports are clear, interactive, and easy to navigate. The statistics presented in this report format are readily accessible, and all updates to the report window appear in real time. This section describes the report format in detail.

## Navigating the Report Layout

The time-series report format not only makes data easily accessible, but also enhances your ability to explore data in context. An example of a typical report appears in [Figure 9-1](#), with the key areas labeled. For details about individual reports, see the report description.

**Figure 9-1. A Time-Series Report**



### 1 Plot Area

The plot area is where the data visualization occurs. Reports can display either a single-pane or dual-pane layout. In a dual-pane layout, both panes remain synchronized with respect to the x-axis. Each pane is capable of having two y-axes (a primary one on the left and a secondary one on the right).

The reports present the majority of data series as simple line series graphs, but some reports display area series graphs where appropriate. The types of area series graphs are:

- **Layered series**, which appear on top of each other in the z direction. These are identified by transparent colors.
- **Stacked area series**, which appear on top of each other in the y direction. RiOS uses stacked area graphs to depict an aggregate broken down into its constituent parts. In this type of graph, each series is a mutually exclusive partition of some aggregate data set, identified by opaque colors. A stacked series is appropriate when the sum of all the series is meaningful.

In the original Steelhead appliance reports, one of the most frequently depicted summary statistic is the peak. The original reports do not let you hover over a specific data point to see what the y values and exact timestamp are in relation to peaks. Consequently, the original reports depict peaks with lines and arrows on the graph itself. The new report format lets you hover over individual data points to view this information, eliminating the need to display a peak line or arrows over peaks.

### To view the timestamp and value of each data series at that time

- Place the mouse pointer over the plot area.

A tool tip displays the timestamp and the value of each data series at that time. The plot area colors the series names appropriately, and the data values have their associated units.

The plot area also displays subtle shading to denote work hours (white background) and non-work hours (gray background). RiOS defines work hours as 8:00 AM to 5:00 PM on weekdays. You cannot configure the work hours.

### To zoom the plot area

1. Place the mouse pointer over the plot area, and then click and hold the left mouse button.
2. Move the mouse left or right and release the left mouse button to zoom in.

## ② Control Panel

Use the control panel to control how much data the chart displays, chart properties, and whether to view or hide the summary statistics.

### To change the chart interval

- Click a link: 5m (five minutes), 1h (one hour), 1d (one day), 1w (one week), or All (all data). All data includes statistics for the last 30 days.

If the current size of the chart window matches any of the links, that link appears in bold black text; the system ignores any clicks on that link. If the time duration represented by any of the links is greater than the total data range of the chart, those links are dimmed.

- **Chart window controls** - More window-related controls appear below the chart window interval links. These controls offer more precise control of the window and also display various window properties. From top to bottom:

- Text field containing the left edge (starting time) of the chart window.
- Text field containing the right edge (ending time) of the chart window.
- Text field containing the chart window interval. The chart window interval in this text field is not always exactly correct, but it is correct to two units (with the units being days, hours, minutes, and seconds). For example, if the chart window interval is exactly two days, three hours, four minutes, and five seconds, this text field displays 2 days, 3 hours.
- Link or static text that represents the chart window state of *attachment* to the end of the chart. When the chart window is attached, the report replaces the link with the static text **Showing newest data**. When the chart is showing newest data, you can see new data points as the system adds them automatically to the chart every 10 seconds. This can be very powerful when you launch a new configuration and need to analyze its impact quickly. You cannot change the 10-second default.

When the chart window is not attached to the end of the chart, the report replaces the static text with a link that displays **Show newest data**. Click this link to slide the chart window to the end of the chart range of data and attach the window.

All three text fields validate your input; if you enter text in an invalid format, an error message appears. If you enter valid text that is logically invalid (for example, an end time that comes before the current start time), an error message appears. With all three text fields, if the focus leaves the field (either because you click outside the field or press Tab), the chart window updates immediately with the new value. Pressing Enter while in one of these fields has the same effect.



## Custom Controls

Below the chart window controls is an optional section of custom, report-specific controls. The custom controls vary for each report. In [Figure 9-1](#), the Bandwidth Optimization report displays Port and Direction drop-down lists.

When you change the value of a custom control, the system sends a new request for data to the server. During this time, the control panel is unavailable and an updating message appears on the chart. When the report receives a response, the system replaces the chart, populates it with the new data, and makes the control panel available again.

## Chart Legend

The chart legend correlates the data series names with line colors and contains a few other features.

You can hide or show individual data series. When a white check box icon appears next to the data series name, you can hide the series from the plot area.

### To hide individual series from the plot area

- Clear the check box next to the data series name.

### To display individual series in the plot area

- Select the check box next to the data series name.

You cannot toggle the visibility of all series, because it does not always make sense to hide a series (for example, if there is only one data series in the chart). For these series, a white check box does not appear next to the series name. In [Figure 9-1](#), you can hide the LAN Throughput and WAN Throughput series, but you cannot hide the Data Reduction series.

The legend also displays statistics. Each report defines any number of statistics for any of the data series in the chart. The system bases the statistics computation on the subset of each data series that is visible in the current chart window. The statistics display changes immediately if you change the chart window. The plot area reflects the changing chart window, as do the associated controls in the control panel.

The reports also support non-series statistics (for example, composite statistics that incorporate the data from multiple data series); these statistics appear at the bottom of the legend, below all the series.

The three most popular statistics calculations are:

- **Average** - the average of all the data points
- **Peak** - the maximum of all the data points
- **Total** - the integral of the series (area under the curve). It is important to note that the total reported under each Throughput color in the chart legend displays the total amount of data transferred during the displayed time interval.

## ③ Navigator

Directly above the scroll bar is the navigator, which shows a much smaller and simpler display of the data in the plot area. The navigator displays only one data series.

Use the navigator to navigate the entire range of chart data. The scroll bar at the bottom shows you which portion of the total data range is displayed in the plot area.

The navigator display can appear very different from the plot area display when an interesting or eye-catching series in the plot area is not the series in the navigator.

### To resize the current chart window

- Move the handles on either side of the chart window in the navigator.

The charts have a minimum chart window size of five minutes, so if you try to resize the chart window to something smaller, the chart window springs back to the minimum size.

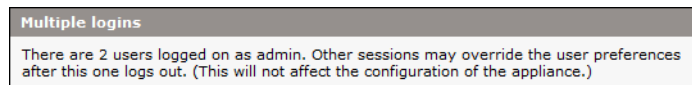
You can also click the data display portion of the navigator (not the scroll bar) and the chart window moves to wherever you clicked.

## Setting User Preferences

You can change report default settings to match your preferred style. When you customize any report-specific settings, the system immediately writes them to disk on the Steelhead appliance. The system saves all of your custom settings, even after you log out, clear your browsing history, or close the browser. When you view the report again, your custom settings are intact.

The system saves report settings on a per-user basis. A message appears at the top of each page when multiple users are logged in, explaining that user preferences might be overwritten.

**Figure 9-2. Multiple Login Alert**



The system also saves the chart window. Whenever you change the chart window, the next time you view any report, the chart window is set to the last chart window used.

## Browser Recommendation

Internet Explorer v7.0 and v8.0 must refresh reports every four minutes due to performance issues. Consider using a different browser to view reports.

---

## Viewing Current Connection Reports

The Current Connections report displays the connections the Steelhead appliance detects, including the connections that are passing through unoptimized.

RiOS v8.5 displays this report using a new layout that shows only relevant data in meaningful visual context. You can search and customize the display using filters to list connections of interest. When you click **Update**, the report retrieves a listing of up to 500 real-time current connections. Navigating to the report or refreshing the page automatically updates the connections display.

## What This Report Tells You

The Current Connections report answers these questions:

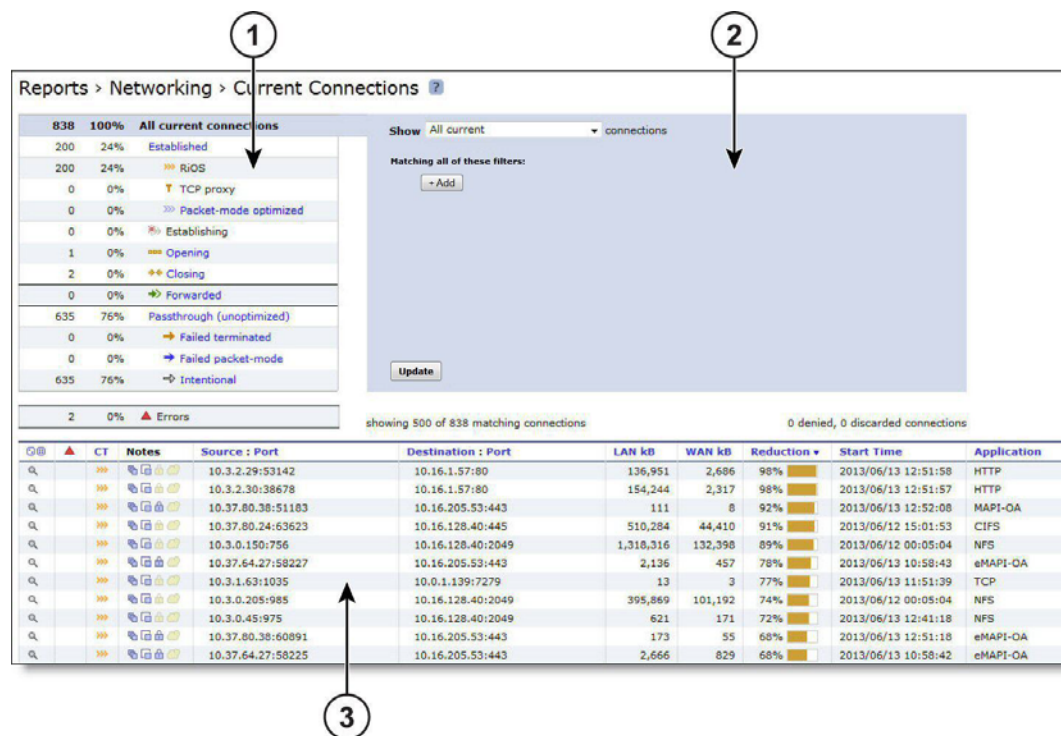
- What traffic is the Steelhead appliance optimizing?
- How many connections are established?
- What is the data reduction on a per-connection basis?
- How many connections are closing?
- How many connections are being passed through either intentionally or unintentionally?

- How many connections are being forwarded by a connection-forwarding neighbor?
- How many connections have been denied or discarded?

### To view the Current Connections report

1. Choose Reports > Networking > Current Connections to display the Current Connections page.

Figure 9-3. Current Connections Report





### 1 Connections Summary





The summary gives you an at-a-glance hierarchical overview of the traffic the Steelhead appliance detects. It displays the total connection numbers for various types of optimization, pass-through, and forwarding. It categorizes the optimized, established connections by type and displays the portion of the total connections each connection type represents.

When you click a connection type such as established, you select it and also drive the show statement in the query area to search for established connections and exclude the other types.

The connections summary displays these connection types:

Connection Type	Icon	Description
All current connections		Displays the total number of connections the Steelhead appliance detects at the time you access the report, refresh the page, or click the <b>Update</b> button. It includes the connections that the Steelhead appliance is passing through unoptimized, and connections that do not appear in the connections table.
Established		Displays the total optimized, active connections.

Connection Type	Icon	Description
		RiOS - Displays the double-ended, non-SCPS connections.
		RiOS + SCPS - Displays the total RiOS and SCPS connections established between two Steelhead appliances running RiOS v7.0 or later. Because both Steelhead appliances are SCPS-compatible, this is a double-ended connection that benefits from traditional RiOS optimization (SDR and LZ).
		SCPS - Displays all current single-ended SCPS-optimized connections as a portion of the total.
		TCP proxy - Displays the total non-SCPS single-ended interception connections. An SEI connection is established between a single Steelhead appliance running RiOS v7.0 or later paired with a third-party device running TCP-PEP (Performance Enhancing Proxy).
		Packet-mode optimized - Displays the total flows that were optimized packet-by-packet with SDR bandwidth optimization. These include TCP and UDP flows over IPv4 or IPv6. Packet-mode flows are considered to be neither single- nor double-ended.
Establishing		Displays the total newly forming, initiating connections. The connection is being established but does not yet have an inner channel.  Establishing connections count toward the connection count limit on the Steelhead appliance because, at any time, they might become a fully opened connection.
Opening		Displays 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.
Closing		Displays 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 disconnected. These connections count 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		Displays the total number of connections that were forwarded when you have configured a connection-forwarding neighbor to manage the connection.  For details about connection forwarding, see <a href="#">“Configuring Connection Forwarding Features” on page 94</a> .

Connection Type	Icon	Description
Passthrough (unoptimized)		Displays 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 connections table that follows this summary.
		Failed terminated - Displays the total number of terminated connections that were passed through unoptimized, because of reasons other than in-path rules.
		Failed packet-mode - Displays the total number of packet-mode flows that were passed through unoptimized, because of reasons other than in-path rules.
		Intentional - Displays the total number of connections that were intentionally passed through unoptimized by in-path rules.
Errors		Displays all connections that have application or transport protocol errors as a portion of the total connections.

## 2 Query Area


The connections summary and the connections table convey a lot of information about connections the Steelhead appliance is detecting. The best way to narrow your search is to filter and sort the report. The query area is where you select a simple or compound connection type for your search and optionally filter the results. The Show search control defines the contents of the connection summary and the connections table.

The simple connection search uses a match against a connection type to display only that type, and excludes the others. If you want to use more advanced criteria, such as including all connections that were started after a certain date, you can add one or more filters to achieve this.

### To display a simple connection type:

1. After Show, select a connection type from the drop-down list:

Figure 9-4. Query Selection



Connection Type	Description
All current	Displays the total number of connections the Steelhead appliance detects, including the connections that are passed through unoptimized. This selection removes any previous selections or filters.
Established	Displays the total optimized, active connections.
Packet-mode optimized	Displays the total connections that were optimized packet-by-packet with SDR bandwidth optimization. These connections include TCP IPv4, TCP IPv6, UDP IPv4, and UDP IPv6 connections.
Opening	<p>Displays 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.</p> <p>If you are experiencing a large number of half-open connections, consider a more appropriately sized Steelhead appliance.</p>
Closing	<p>Displays 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 disconnected. 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.)</p> <p>If you are experiencing a large number of half-closed connections, consider a more appropriately sized Steelhead appliance.</p>
Forwarded	Displays the total number of connections forwarded by the connection-forwarding neighbor managing the connection.
Passthrough (unoptimized)	Displays 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 the connections summary.
Failed terminated	Displays the total number of terminated connections that were passed through unoptimized.
Failed packet-mode	Displays the total number of packet-mode flows that were passed through unoptimized.
Intentional	Displays the total number of connections that were intentionally passed through unoptimized.

## 2. Click **Update**.

### **Filtering the Connections**

Filters provide a powerful way to drill down into large numbers of connections by specifying either simple or complex filter criteria. Each filter further restricts the display.

When you customize filters, the system immediately writes them to disk on the Steelhead appliance. The system saves all of your custom settings even after you log out, clear your browsing history, or close the browser. When you view the report again, your custom settings are intact. The system saves report settings on a per-user basis.

#### **To filter the display (optional):**

1. Click **Add**.
2. Select a filter from the drop-down list. Selecting some filters expands the query with a text input field for additional information. For example, selecting for application from the drop-down list displays a

text input field for the application name. RiOS validates the text input fields as you enter the text (except when you enter a regular expression).

**Figure 9-5. Filtering the Current Connection Display**

The screenshot shows a web interface for filtering connections. At the top, there's a 'Show' dropdown menu currently set to 'All current' with a 'connections' label. Below this, a section titled 'Matching all of these filters:' contains two filter entries. The first entry is 'from source port' with a value of '443'. The second entry is 'from source IP address/mask' with a value of '10.3.4.20/32'. Each entry has a small 'x' icon to its left. Below the filters is a '+ Add' button. At the bottom left of the panel is an 'Update' button.

You can select any combination of these filters:

- **matching regular expression** - Displays a text input field for a regular expression and shows only those connections that match the expression.

Use the following format in the text input field:

`x.x.x.x[/mask][:port]`

**Examples:**

10.16.35.1

Finds one particular IP address

10.16.35.1:5001

Finds port 5001 on one particular IP address

You can also use the regular expression filter 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>`.

**Notes:**

RiOS does not validate the expression. A regular expression can contain special characters and embedded spaces that are unique to the regular expression syntax. For details, see *The Gnu Awk User's Guide*.

The filter matches only against the source, destination, and application name. It does not consider start times, reduction, and byte counts.


The filter separates IP addresses and ports with a colon for matching:

`x.x.x.x:p` for IPv4

`[xxxx:xxxx::xxxx]:p` for IPv6

Upper and lowercase do not matter ("mapi" matches MAPI, MAPI-ENCRYPT, and any other application containing MAPI).

A connection matches if the match string occurs anywhere within it (that is, ":" matches all rows), unless overridden by special regular expression language such as "^" or "\$".

- **from source IP address/mask** - Displays a text input field for the IP address and subnet mask. You can specify an IPv4 or an IPv6 IP address.
  - **from source port** - Displays a text input field for the source port.
  - **to destination IP address/mask** - Displays a text input field for the IP address and subnet mask. You can specify an IPv4 or an IPv6 IP address.
  - **to destination port** - Displays a text input field for the destination port.
  - **that have errors** - Displays connections with either application protocol errors or transport protocol errors.
  - **for application** - Select an application name from the drop-down list. The application filter is only relevant for optimized connections.
  - **that were started before** - Displays a text input field for the date and time. Use this format: YYYY/MM/DD hh:mm:ss.
  - **that were started after** - Displays a text input field for the date and time. Use this format: YYYY/MM/DD hh:mm:ss.
  - **that are single-ended only** - Displays SCPS and TCP proxy connections. Applies only to established connections.
  - **that are double-ended only** - Displays RiOS and RiOS + SCPS connections. Applies only to established connections.
3. To add another filter, click **add filter** again. You can add up to eight filters; they are logically AND-ed together and are all active at any given time. Continue adding filters until your query is complete.
-  To delete a filter, click the delete filter icon.

4. Click **Update**.

### ③ Connections Table

The connections table displays more information about each connection, filtered by the show statement and any filters in the query area. The connections table can show up to 500 connections at a time; it lists the total of all matching connections in the upper-right corner. 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).

For details about the query area, see [“Query Area” on page 251](#).

Connections with IPv6 addresses are split into two rows to accommodate the long address. The report encloses IPv6 addresses in square brackets, and the source address, destination address, and other information appears in different columns.

Icons in the CT and Notes columns indicate the connection type and attributes. Use the mouse to hover over an icon and reveal a tooltip identifying its meaning.



The individual connections table displays additional information about each connection. Because this report can list hundreds of transient connections, you can sort the table by column heading (except for the Notes column). For example, you can sort the connections by source IP address.

**To sort the table by row:**

- Click the table column heading.

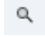








The table contents reload, if necessary. Click the heading again to reverse the order. A small up or down triangle reflects the current bidirectional sort order.





**To reset the connection sample:**

- Click the dice icon on the far left. 

The table contents reappear in the original display. For example, if you sort the display by a particular type, and there are more than 500 connections of that type, click the dice icon to return to the original display.

The connections table displays the following information:

Column	Description
	<p>Click the magnifying glass to display the current connections details. See <a href="#">“Viewing the Current Connection Details” on page 257</a>.</p> <p>Because the details are 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 the connection is closed. To refresh the display, click <b>Update</b>.</p>
	<p><b>Protocol Error</b> - Displays a protocol error for both transport and application conditions. This list contains some of the conditions that trigger errors; it is a small subset of possible error conditions:</p> <ul style="list-style-type: none"> <li>• When the Optimize Connections with Security Signatures feature is enabled (which prevents SMB signing). This is an expected response.</li> <li>• If a SRDF protocol error occurs when attempting to optimize traffic originating from the LAN side of the Steelhead appliance. Check the LAN-side Symmetrix array for compatibility.</li> </ul> <p>Click the magnifying glass for more details about the error.</p>
CT (Connection Type)	<b>Established</b> - Indicates that the connection is established and active.
	
	<b>Intentional Passthrough</b> - Indicates that the connection was intentionally passed through unoptimized because of in-path rules.
	<b>Failed terminated</b> - Indicates that the connection was passed through unoptimized.
	<b>Failed packet-mode</b> - Indicates that the packet-mode flow was passed through unoptimized.
	<b>Establishing</b> - Indicates that the connection is initiating and is not yet fully established. The source and destination ports appear as n/a.
	<b>Opening (Optimized)</b> - Indicates that the connection is half-open and active. A half-open connection is a TCP connection that has not been fully established.
	<b>Closing (Optimized)</b> - Indicates that 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 disconnected.

Column	Description
	<b>Forwarded</b> - Indicates that the connection is forwarded by the connection-forwarding neighbor managing the connection. For details about connection forwarding, see <a href="#">“Configuring Connection Forwarding Features” on page 94</a> .
Notes	Displays connection icons that indicate the current state of the connection. The connection states can be one of these:
	<b>Compression Enabled</b> - Indicates that LZ compression is enabled.
	<b>SDR Enabled</b> - Indicates that SDR optimization is enabled.
	<b>WAN Encryption Enabled</b> - Indicates that encryption is enabled on the secure inner channel (WAN). For details, see <a href="#">“Configuring Secure Peers” on page 176</a> .
Source:Port	Displays the connection source IP address and port.
Destination:Port	Displays the connection destination IP address and port.
LAN/kB WAN/kB	Displays the amount of LAN or WAN throughput, in kilobytes.
Reduction	Displays the degree of WAN traffic optimization as a percentage of LAN traffic sent. Higher percentages mean that fewer bytes were sent over the WAN.
Start Time	Displays the time that the connection was started. This column does not apply to preexisting connections. Select the column heading to sort data start time in ascending or descending order.
Application	Displays the application associated with the connection.

**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 56](#).

## 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 these 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
- Retrieve the most recent data for a connection

The report does not allow the connection details to refresh automatically, because doing so could slow down the Steelhead appliance; however, the connection age updates when you manually refresh the page.

### 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 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. Click **Update** to refresh the display.

Figure 9-6. Current Connections Details for an Optimized Connection

	LAN side	WAN side
Bytes	18,009,907	1,448,489
Packets	233,898	3,824
Retransmitted	0	10
Fast retransmitted	0	0
Timeouts	0	6
Congestion window	5	4

### To close the connection details report

- Click the close icon.

#### 4 Connection Details

The expanded connection details vary, depending on the nature of the connection.

## Optimized Connection Details

This table summarizes details about individual optimized connections.

Data	Description (varies by connection type)
Connection Information	<p><b>Connection Type</b> - Displays the connection type icon and whether the connection is established, opening, or closing.</p> <p><b>Connection Age</b> - Displays the time since the connection was created.</p> <p><b>Transport</b> - Displays the transport protocol name: for example, SSL inner.</p> <p><b>Application</b> - Displays the application corresponding to the connection.</p> <p><b>Client Side</b> - Displays whether this appliance is on the client side.</p> <p><b>In-path</b> - Indicates whether the connection is in-path.</p> <p><b>Protocol</b> - Displays the low-level protocol that RiOS is using inside the packet-mode channel. The protocol can be UDP, TCP, or variants.</p> <p><b>Application Error</b> - Displays the application protocol error, if one exists.</p> <p><b>Transport Error</b> - Displays the transport protocol error, if one exists.</p> <p><b>SaaS application</b> - Displays the SaaS application name, if one exists.</p> <p><b>Cloud application state</b> - Displays the SaaS connection state, if an SaaS application is running.</p> <p><b>SkipWare compression in</b> - Indicates that the single-ended optimized connection is applying Skipware105 compression on incoming data.</p> <p><b>SkipWare compression out</b> - Indicates that the single-ended optimized connection is applying Skipware105 compression on outgoing data.</p> <p><b>Pre-existing asymmetric</b> - Indicates that the connection is travelling an asymmetric route and existed before the last restart of the optimization service.</p> <p><b>Pre-existing</b> - Indicates that the connection existed before the last restart of the optimization service.</p> <p>When relevant, the Notes section displays several details that are binary in nature.</p>
	<p>All optimized connections might show any of the following:</p> <p><b>Client side</b> - Indicates that the Steelhead appliance is on the client side of the connection.</p> <p><b>SDR optimized</b> - Indicates that SDR optimization is enabled.</p> <p><b>LZ compressed</b> - Indicates that LZ compression is enabled.</p>
	<p>Packet-mode optimized connections might show:</p> <p><b>Incomplete parse</b> - Indicates that the inner channel exists but the connection through the channel is not fully formed.</p>
	<p>Optimized, non-packet mode connections might show any of the following:</p> <p><b>In-path</b> - Indicates an in-path connection.</p> <p><b>Single-ended</b> - Indicates that the connection involves only one Steelhead appliance.</p> <p><b>WAN encrypted</b> - Indicates that encryption is enabled on the secure inner channel (WAN).</p>

Data	Description (varies by connection type)
	<p>At least one of these items appear for SCPS connections:</p> <p><b>SCPS initiate WAN</b> - Indicates that the Steelhead appliance has initiated the SCPS connection on the WAN.</p> <p><b>SCPS initiate LAN</b> - Indicates that the Steelhead appliance has initiated the SCPS connection on the LAN.</p> <p><b>SCPS terminate WAN</b> - Indicates that the Steelhead appliance has terminated the SCPS connection on the WAN.</p> <p><b>SCPS terminate LAN</b> - Indicates that the Steelhead appliance has terminated the SCPS connection on the LAN.</p>
WAN and LAN-Side Statistics	<p><b>LAN Bytes</b> - Displays the total LAN bytes transmitted.</p> <p><b>WAN Bytes</b> - Displays the total WAN bytes transmitted.</p> <p><b>Retransmitted</b> - Displays the total packets retransmitted.</p> <p><b>Fast Retransmitted</b> - Displays 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.</p> <p><b>Timeouts</b> - Displays the number of packet transmissions that timed out because no ACK was received.</p> <p><b>Congestion Window</b> - Displays number of unACKed packets permitted, adjusted automatically by the Steelhead appliance, depending on WAN congestion.</p>

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

### Individual Pass-Through or Forwarded Connection Details

This table summarizes details about individual pass-through or forwarded connections.

Data	Description (Varies by Connection)
Connection Information	<p><b>Connection Type</b> - Displays a connection type icon and whether the pass-through was intentional or unintentional. Displays the forwarded reduction percentage bar for forwarded connections.</p> <p><b>Connection Age</b> - Displays the time since the connection was created.</p> <p><b>Transport</b> - Displays the transport protocol name.</p> <p><b>Application</b> - Displays the application corresponding to the connection.</p> <p><b>Client-Side</b> - Displays whether the connection is on the client side.</p> <p><b>Pre-Existing</b> - Displays whether the connection existed before the last restart of the optimization service.</p> <p><b>Passthrough Reason</b> - Displays the reason for passing through or forwarding the connection.</p>

### Pass-Through Reasons

This table shows the connection pass-through reasons.

Value	Pass-through Reason (Varies by Connection)	Description	Action
0	None	None	None

Value	Pass-through Reason (Varies by Connection)	Description	Action
1	Preexisting connection	Connection existed before Steelhead appliance started.	Create a 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.
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 appliance. Take a TCP dump to check which TCP options are in the SYN packet. Search for those options to learn what device uses them.

Value	Pass-through Reason (Varies by Connection)	Description	Action
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 when the enhanced auto-discovery protocol is enabled.
16	Error connecting to server	The server-side Steelhead appliance could not connect to the server.	Only happens when the enhanced auto-discovery protocol is enabled.
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.
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.



Value	Pass-through Reason (Varies by Connection)	Description	Action
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	An inner channel handshake has failed with peer.	Check the SSL configuration on both Steelhead appliances.
37	Netflow only: Ricochet packet of optimized connection	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.
39	A SYN or RST packet contains data		
40	Failed to discover SCPS device	RiOS cannot find a SCPS device.	

Value	Pass-through Reason (Varies by Connection)	Description	Action
41	No matching client/server IPv6 scope	RiOS cannot set up the outer channel connection.	RiOS passes all packets through until it creates the outer channel.
42	Failed to create sport outer channel	RiOS cannot set up the outer channel connection.	RiOS passes all packets through until it creates the outer channel.
43	Flows not matching in-path rule	RiOS cannot match this traffic flow to any packet-mode optimization in-path rule. A packet-mode optimization rule defines the inner channel characteristics.	RiOS passes all packets through while the flow is in this state. Go to Configure > Optimization > In-Path rules to add a fixed-target packet-mode optimization in-path rule.
44	Packet mode channel setup pending	RiOS is attempting to set up the inner IPv4 or IPv6 channel connection.	RiOS passes all packets through until it creates the inner IPv4 or IPv6 channel.
45	Peer does not support packet-mode optimization	The peer Steelhead Appliance to which RiOS needs to establish the inner IPv4 or IPv6 channel connection does not support packet-mode optimization or packet-mode optimization is not enabled.	RiOS stops trying to optimize connections using packet-mode optimization with the peer.
46	Generic Flow error	<p>A packet-mode optimization traffic flow transitions to this state when RiOS encounters one of these unrecoverable errors:</p> <ul style="list-style-type: none"> <li>• There is not enough memory to set up the inner channel.</li> <li>• The system has requested that RiOS kill the traffic flow.</li> </ul> <p>When RiOS receives this error, the Steelhead appliance abandons all attempts to optimize the flow.</p>	RiOS passes the flow through for its lifetime.
47	Failed to cache sock pointer	While configured for packet-mode optimization, RiOS cannot locate the socket pointer used to exchange packets through the inner channel. The system is attempting to write packets to the ring, but the socket is closed. This can occur when the optimization service shuts down unexpectedly.	Go to Configure > Maintenance > Services and restart the optimization service.
48	Packet mode optimization disabled	The connection is being passed through because packet-mode optimization is disabled.	Go to Configure > Optimization > In-path Rules and enable packet-mode optimization.

Value	Pass-through Reason (Varies by Connection)	Description	Action
49	Optimizing local connections only	On a Steelhead EX appliance, the connection is being passed through because it did not originate locally.	
50	Netflow only: probe packet of optimized connection		
51	IPv6 connection forwarding requires multi-interface support	RiOS is passing the connection through because the client-side Steelhead appliance is configured without multi-interface connection forwarding. This configuration does not support IPv6.	Go to Configure > Networking > Connection Forwarding and enable multiple interface support.
52	Neighbor does not support IPv6	RiOS is passing the connection through because a connection-forwarding neighbor does not support IPv6.	Upgrade the connection-forwarding neighbor to RiOS v8.0 or later.

### SaaS Connection Details

This table shows the SaaS connection details.

Value	Reason	Description	Action
0	None	None	None
1	Optimized connection	Connection is redirected through the SCA to a SaaS service.	No action is necessary.

### Pass-Through Reasons for SaaS Connections

This table lists the connection pass-through reasons for SaaS connections.

Value	Pass-through Reason (Varies by Connection)	Description	Action
3	Not a supported SaaS destination	Connection is through a SaaS service that is not supported, subscribed to, or enabled.	No action is necessary; however, if you want to optimize this destination IP address, contact Riverbed Support.
9	Cloud proxy is down	Connection is not redirected through SCA because the redirection service encountered an error.	Contact Riverbed Support.
11	Failed to append CP code	Connection is not redirected through SCA because of a packet processing error.	Contact Riverbed Support.
12	SYN retransmit (backhailed)	Connection is not redirected through SCA because too many SYN retransmits were received from the client.	Check if there is a firewall that does not allow inbound or outbound UDP packets for the Steelhead appliance.

Value	Pass-through Reason (Varies by Connection)	Description	Action
13	SYN retransmit (direct)	Connection is not redirected through SCA because too many SYN retransmits were received from the client.	Check if there is a firewall that does not allow inbound or outbound UDP packets for the Steelhead appliance.
14	Passing to downstream Steelhead	Connection is not redirected through SCA because admission control is reached and there is a Steelhead appliance downstream that might optimize the connection.	No action is necessary.
15	Passthrough SYN retransmit	Connection is not redirected through SCA because too many SYN retransmits were received from the client.	Check if there is a firewall that does not allow inbound or outbound UDP packets for the Steelhead appliance.
16	Rejected by cloud proxy	Connection is not redirected through SCA because the SCA network rejected the connection.	Contact Riverbed Support.
17	Invalid Entitlement code	Connection is not redirected through SCA because of an invalid SCA configuration.	Contact Riverbed Support.
18	Invalid timestamp	Connection is not redirected through SCA because the clock on the Steelhead appliance is not synchronized.	Check the date and time settings on the Steelhead appliance.
19	Invalid customer ID	Connection is not redirected through SCA because of an invalid SCA configuration.	Contact Riverbed Support.
20	Invalid ESH ID	Connection is not redirected through SCA because of an invalid SCA configuration.	Contact Riverbed Support.
21	Invalid SaaS ID	Connection is not redirected through SCA because of an invalid SCA configuration.	Contact Riverbed Support.
22	Connection limit reached	Connection is not redirected through SCA because the subscription limit for the number of connections is reached.	Contact Riverbed Support. You might require a higher SCA license.
23	Bandwidth limit reached	Connection is not redirected through SCA because the subscription limit for bandwidth used is reached.	Contact Riverbed Support. You might require a higher SCA license.

## 5 Tools

This section provides buttons that perform an operation on a single connection. It also provides a link to log information.

**Figure 9-7. Tools**




You can perform these operations:

Control	Description
Send Keep-Alive	For an optimized connection, click to send a keep-alive message to the outer remote machine (the machine that is connected to this appliance). This operation is not available for a pass-through connection.  This button is dimmed for users logged in as a monitor user.
Refresh Data	Click to retrieve the most recent data for the connection.
Reset Connection	Click to send a RST packet to both the client and server to try to close the connection. You can reset both optimized and pass-through connections. You cannot reset a forwarded connection.  <b>Note:</b> If no data is being transferred between the client and server when you click <b>Reset Connection</b> , 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 might take a while for the connection to disappear.  This button is dimmed for users logged in as a monitor user.
Log for this Steelhead	Click to go to the System Logs page.

## 6 Network Topology

This section shows a graphical representation of the connection source-to-destination network topology and information associated with the different elements. This graphic varies depending on the connection type and is only relevant for optimized connections. It does not appear for pass-through connections.

**Figure 9-8. Connection Topology**



The topology shows this information:

- All of the IP addresses and port numbers associated with the connection.

- Transparency mode, which describes the visibility of each actual IP address and port on the Steelhead appliances to each other, for terminated connections only. For details, see [“Configuring In-Path Rules” on page 36](#).
- Channel ID and type for packet-mode flows only. For details, see [“Configuring In-Path Rules” on page 36](#).
- Congestion control, including the method in use to mitigate WAN congestion. For details on congestion-control types, see [“Configuring TCP, Satellite Optimization, and High-Speed TCP” on page 59](#).
- Steelhead appliance models and RiOS versions.
- A YOU ARE HERE label identifies the Steelhead appliance whose page you are viewing.

## 7 LAN/WAN Table

This table shows raw tallies for LAN and WAN connections to summarize data about channel processing for a specific connection. The table varies by type of connection.

**Figure 9-9. LAN/WAN Table**

	LAN side	WAN side
Bytes	392,113,072	14,838,685
Packets	63,021	10,780
Retransmitted	0	12
Fast retransmitted	0	12
Timeouts	0	0
Congestion window	81	13

Use this table to answer questions such as:

- For any given channel, how many bytes (or packets) did the channel receive and subsequently transmit?
- Which channels have processed the most traffic? The least traffic?
- What error types and quantities were encountered for traffic inbound from the WAN?
- What error types and quantities were encountered for traffic inbound from the LAN?

## Viewing Connection History Reports

The Connection History report shows connection counts for a variety of connection types for the time period specified.

This report includes IPv6 and packet-mode optimized traffic in RiOS v8.5.

For details about the report format, see [“Overview” on page 244](#).

The Connection History report contains these statistics that summarize connection activity.

Connection Type	Description
Optimized	Displays the total connections established and optimized, plus the half-open and half-closed connections (where half-open and half-closed are TCP connection states).

Connection Type	Description
Optimized (Active)	Displays the total active connections established, optimized, and flowing.
Passthrough	Displays the total connections passed through unoptimized.
Forwarded	Displays the total number of connections forwarded by the connection-forwarding neighbor managing the connection.
Optimized (Half Open)	<p>Displays the percentage of half-opened connections represented in the optimized connection total. A half-open connection is a TCP connection that 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.</p> <p>If you are experiencing a large number of half-opened connections, consider a more appropriately sized Steelhead appliance.</p>
Optimized (Half Closed)	<p>Displays the percentage of half-closed active connections represented in the optimized connection total. Half-closed connections are connections that the Steelhead appliance has intercepted and optimized but are in the process of being disconnected. 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 its connections cleanly.)</p> <p>If you are experiencing a large number of half-closed connections, consider a more appropriately sized Steelhead appliance.</p>

The navigator shadows the optimized series.

## What This Report Tells You

The Connection History report answers these questions:

- How many connections were optimized?
- How many connections were passed through, unoptimized?
- What is the percentage of half-opened connections represented in the total optimized connections?
- What is the percentage of half-closed connections represented in the total optimized connections?

## About Report Graphs

Use the mouse to hover over a specific data point to see the values and exact timestamp.


## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The data is collected at a 5-minute granularity for the entire month.

To view the Connection History report

- 1. Choose Reports > Networking > Connection History to display the Connection History page.

Figure 9-10. Connection History Page



- 2. Use the controls to change the report display as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>



---

## Viewing Connection Forwarding Reports

The Connection Forwarding report summarizes the data throughput between the Steelhead appliance and a specified neighbor (or all neighbors).

For details about the report format, see [“Overview” on page 244](#).

Data Series	Description
Throughput	Displays the throughput in bits per second.

The navigator shadows the throughput series.

You configure neighbors when you enable connection forwarding. For details, see [“Configuring Connection Forwarding Features” on page 94](#).

### What This Report Tells You

The Connection Forwarding report answers this question:

- How many bytes were transferred between a Steelhead appliance and a specified neighbor?

### About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.


### About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The data is collected at a 5-minute granularity for the entire month.

## To view the Connection Forwarding report

1. Choose Reports > Networking > Connection Forwarding to display the Connection Forwarding page.

Figure 9-11. Connection Forwarding Page



2. Use the controls to customize the report as described in this table.

Control	Description
Time interval	Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.  Time intervals that do not apply to a particular report are dimmed.  For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.  You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.
Neighbor	Select a neighbor from the drop-down list or All to display all neighbors.

---

## Viewing Outbound QoS Reports

The Outbound QoS report summarizes the number of bits per second or packets per second transmitted for either a set of QoS classes (up to seven) or an aggregate total of all classes for the time period specified.

For details about the report format, see [“Overview” on page 244](#).

### What This Report Tells You

The Outbound QoS report answers these questions:

- Is outbound QoS working correctly?
- How many bits or packets per second were transmitted over the WAN for the QoS classes?
- How many bits or packets per second were sent and dropped for the QoS classes?

### About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

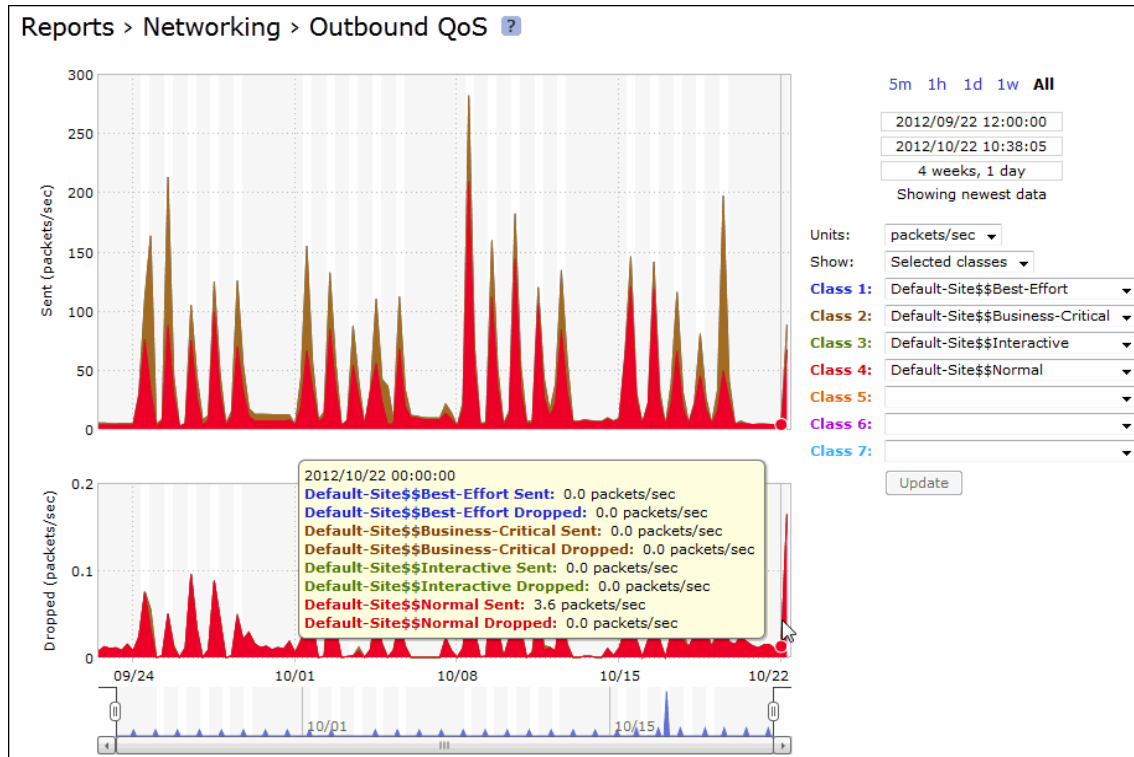
### About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled with a granularity of 5 minutes for the day, 1 hour for the last week, and 2 hours for the rest of the month.

## To view the Outbound QoS report

1. Choose Reports > Networking > Outbound QoS to display the Outbound QoS page.

Figure 9-12. Outbound QoS Page



2. Use the controls to change the report display as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>Because the system aggregates data on the hour, request hourly time intervals. For example, setting a time interval to 08:30:00 to 09:30:00 from 2 days ago does not create a data display, whereas setting a time interval to 08:00:00 to 09:00:00 from 2 days ago will display data.</p> <p>When you request a custom time interval to view data beyond the aggregated granularity, the data is not visible because the system is no longer storing the data. For example, the following custom time intervals do not return data because the system automatically aggregates data older than 7 days into 2-hour data points:</p> <ul style="list-style-type: none"> <li>Setting a 1-hour time period that occurred 2 weeks ago.</li> <li>Setting a 75-minute time period that occurred more than 1 week ago.</li> </ul> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>

Control	Description
Units	Select either packets/sec or bps from the drop-down list.
Classes	<p>Select Total or Selected classes from the drop-down list. Selected classes lets you narrow the report by choosing from drop-down lists of classes (up to seven). You cannot select a class more than once.</p> <p>Click <b>Update</b> to change the QoS class selection without updating the chart.</p> <p>When the report display includes the total classes, the data series appear as translucent; selected classes appear as opaque.</p> <p>When the report display includes the total classes, the navigator shadows the total sent series. When the report display includes selected classes, the navigator shadows the first non-empty sent series. A data series can be empty if you create a QoS class but it has not seen any traffic yet.</p> <p>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.</p> <p>When a selected class has descendent classes, the report aggregates the statistics for the entire tree of classes. It displays the aggregated tree statistics as belonging to the selected class.</p>

## Viewing Inbound QoS Reports

The Inbound QoS report displays received and dropped throughputs for a variety of inbound QoS class configurations (up to seven) or an aggregate total of all classes for the time period specified.

For details about the report format, see [“Overview” on page 244](#).

### What This Report Tells You

The Inbound QoS report answers these questions:

- How many bits or packets per second were transmitted over the WAN for the QoS classes?
- How many bits or packets per second were received and dropped for the QoS classes?

### About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

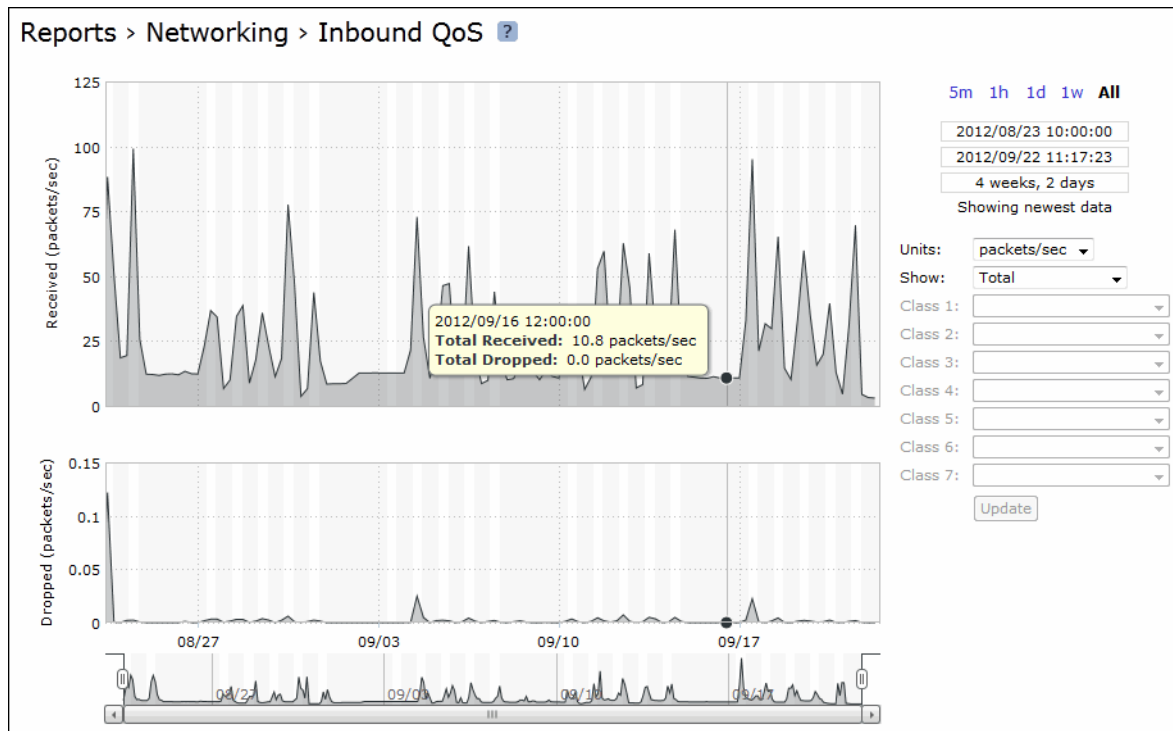
### About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled with a granularity of 5 minutes for the day, 1 hour for the last week, and 2 hours for the rest of the month.

## To view the Inbound QoS report

1. Choose Reports > Networking > Inbound QoS to display the Inbound QoS page.

Figure 9-13. Inbound QoS Page



2. Use the controls to customize the report as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>Because the system aggregates data on the hour, request hourly time intervals. For example, setting a time interval to 08:30:00 to 09:30:00 from 2 days ago does not create a data display, whereas setting a time interval to 08:00:00 to 09:00:00 from 2 days ago will display data.</p> <p>When you request a custom time interval to view data beyond the aggregated granularity, the data is not visible because the system is no longer storing the data. For example, these custom time intervals do not return data because the system automatically aggregates data older than 7 days into 2-hour data points:</p> <ul style="list-style-type: none"> <li>• Setting a 1-hour time period that occurred 2 weeks ago.</li> <li>• Setting a 75-minute time period that occurred more than 1 week ago.</li> </ul> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Units	Select either packets/sec or bps from the drop-down list.
Classes	<p>Select Total or Selected classes from the drop-down list. Selected classes lets you narrow the report by choosing from drop-down lists of classes (up to seven). You cannot select a class more than once.</p> <p>Click <b>Update</b> to change the QoS class selection without updating the chart.</p> <p>When the report display includes the total classes, the data series appear as translucent; selected classes appear as opaque.</p> <p>When the report display includes the total classes, the navigator shadows the total received series. When the report display includes selected classes, the navigator shadows the first non-empty received series. A data series can be empty if you create a QoS class but it has not seen any traffic yet.</p> <p>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.</p> <p>When a selected class has descendent classes, the report aggregates the statistics for the entire tree of classes. It displays the aggregated tree statistics as belonging to the selected class.</p>

## Viewing Path Selection Reports

The Path Selection report provides real-time visibility into path selection activity. View this report to see how your path selection configuration is working for all in-path interfaces and user-defined paths. Seeing path selection in operation can help refine your configuration.

For details on enabling path selection, see [“Configuring Path Selection” on page 153](#).

The top of the report displays whether path selection is currently enabled or disabled. When path selection is disabled, the interface and path tables do not appear.

## Viewing Interface Information

The interfaces table displays statistics about the relay interfaces associated with a path. The table abbreviates the interface name to ...0\_0, ...0\_1, and so on, instead of inpath0\_0, and inpath0\_1.

For each in-path interface the table displays information about the amount of traffic (in Megabytes) going through it::

Data	Description
Megabytes In	Displays the amount of traffic received through this interface.
Megabytes Relayed	Displays the amount of traffic passed through this interface; relayed data is not using path selection.
Megabytes Dropped	Displays the amount of traffic dropped from this interface; the data is lost.
Megabytes Out	Displays the amount of traffic transmitted from this interface.

The report uses this calculation for each interface:

$$\text{Megabytes In} = \text{Megabytes Relayed} + \text{Megabytes Dropped} + \text{Megabytes Out}$$

**Figure 9-14. Path Selection Page - Interfaces**

Reports > Networking > Path Selection ?

Path Selection Enabled: Yes

Paths Interfaces

	Inpath	Megabytes In	Megabytes Relayed	Megabytes Dropped	Megabytes Out
	inpath0_0	44.30	8.68	9.37	🔍 26.25
	inpath0_1	40.89	8.71	9.22	🔍 22.96
	inpath1_0	30.52	8.57	6.52	🔍 15.43
	inpath1_1	31.66	8.66	4.81	🔍 18.18







To view more information about the data transmitted from an interface, click the magnifying glass under the Megabytes Out column.

**Figure 9-15. Path Selection Page - Interfaces**

Path Selection Enabled: Yes

Paths

Interfaces

	Inpath	Megabytes In	Megabytes Relayed	Megabytes Dropped	Megabytes Out										
	inpath0_0	44.30	8.68	9.37	 26.25										
	Sorted Breakdown of "Megabytes Out":														
	<table><tr><th>Megabytes</th><th>Interface</th></tr><tr><td>8.60</td><td>inpath0_0</td></tr><tr><td>7.73</td><td>inpath1_1</td></tr><tr><td>6.96</td><td>inpath0_1</td></tr><tr><td>2.96</td><td>inpath1_0</td></tr></table>					Megabytes	Interface	8.60	inpath0_0	7.73	inpath1_1	6.96	inpath0_1	2.96	inpath1_0
Megabytes	Interface														
8.60	inpath0_0														
7.73	inpath1_1														
6.96	inpath0_1														
2.96	inpath1_0														
	inpath0_1	40.89	8.71	9.22	 22.96										
	inpath1_0	30.52	8.57	6.52	 15.43										
	inpath1_1	31.66	8.66	4.81	 18.18										

A table breaks down the outbound traffic from the interface. The table is sorted by data in descending order (high to low). The total of the values under the megabytes column add up to the amount of megabytes out for the interface. The Interface column names the interface used to transmit the data.

## Viewing Path Information

The paths table displays the path selection configuration, state, and statistics. If path selection is disabled, the status of all paths is disabled.

Use the radio buttons to filter the contents of the table based on the status of the path. You can display all paths or just those paths that are either up or down.

Data	Description
Name	Displays the path name, sorted alphabetically.
Status	Indicates whether the server on the other end is pingable (Up) or not (Down).
Interface	Displays the interface name over which the path is reached.
Gateway IP	Displays the next-hop gateway IP address for the path.
Destination MAC	Displays the destination MAC address used to send a packet over the path.
Remote IP	Displays the remote endpoint IPv4 address of the path.

To see more information about a particular path, click the path name to expand the row.

Figure 9-16. Path Selection Page

**Reports > Networking > Path Selection**

Interfaces Paths

**Paths Table:** Filter by Status: ☒ All ☐ "Up" ☐ "Down"

Name	Status	Interface	Gateway IP	Destination MAC	Remote IP
<a href="#">Path_1</a>	Up	inpath0_0	--	00:50:56:93:44:32	10.11.6.4
<a href="#">Path_2</a>	Down	inpath0_0	--	00:00:00:00:00:00	10.12.1.18
<a href="#">Path_3</a>	Up	inpath0_0	--	00:50:56:93:44:32	10.11.6.4

**Probe**

DSCP: 0x20001

Timeout (sec): 2

Threshold (sec): 3

Requests: 11

Response Relay Mismatch: 0

Ricochet Ignored: 0

Ricochet Relay Mismatch: 0

**Miscellaneous**

Bytes: 0

VLAN: 65535

Source MAC: 00:0e:b6:91:4a:ef

The information that appears is split into two columns: probe and miscellaneous. The probe column displays the probe packet settings described in this table.

Data	Description
DSCP	Displays the DSCP mark defined on this path for a given flow.
Timeout (sec)	Displays the number of seconds defined on this path to monitor path latency availability. If the ICMP ping responses do not make it back within this many seconds, the system considers the probe to be lost. The default is 2 seconds.
Threshold (sec)	Displays how many lost packets the system counts before it considers the path to be unavailable and triggers the Path Down alarm. The default is 3 failed successive packets.
Requests	Displays the number of path monitor probe requests (ICMP probes) that have been sent out thus far by the Steelhead appliance to monitor the status of the path.
Response Relay Mismatch	Displays how many times a path monitor probe request was sent from a different interface than the interface that received the path monitor probe reply. The system considers this a relay mismatch and increments this counter.

The miscellaneous column displays the information described in this table.

Data	Description
Bytes	Displays the total number of bytes of data that have been sent on the given path.
VLAN	Displays the VLAN ID. When there is no VLAN associated with the path, the report displays "None."
Source MAC	Displays the source MAC address of the in-path interface defined for the given path.

## What This Report Tells You

The Path Selection report answers these questions:

- What is the best way to balance traffic flows across all available paths?
- How much traffic is going through in-path interface 0\_0?
- Which paths are currently unavailable?
- Is the Steelhead appliance directing traffic through the configured paths?

### To view the Path Selection report

- Choose Reports > Networking > Path Selection to display the Path Selection page.

Figure 9-17. Path Selection Page

Reports > Networking > Path Selection ?						
Interfaces		Paths				
Paths Table: <span style="float: right;">Filter by Status: <input checked="" type="radio"/> All <input type="radio"/> "Up" <input type="radio"/> "Down"</span>						
	Name	Status	Interface	Gateway IP	Destination MAC	Remote IP
	<a href="#">Path_1</a>	Up	inpath0_0	--	00:50:56:93:44:32	10.11.6.4
	<a href="#">Path_2</a>	Down	inpath0_0	--	00:00:00:00:00:00	10.12.1.18
	<a href="#">Path_3</a>	Up	inpath0_0	--	00:50:56:93:44:32	10.11.6.4

Related Topics: [Configure: Path Selection](#)

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 first single IP address and port number to a second single IP address and port number over the same protocol. Only traffic flows that start in the selected time period are shown in the report.

The Top Talkers report does not include IPv6 traffic.

**Important:** The Top Talkers report includes bytes used for packet headers and is an approximation based on various assumptions.

The Top Talkers report contains this table of statistics that summarize Top Talker activity.

Column	Description
Rank	Displays the relative position of the traffic flow WAN bandwidth use.
<Sender> IP Address 1:Port	Displays the first IP address and port for the connection.
<Receiver> IP Address 2:Port	Displays the second IP address and port for the connection.
Byte Count	Displays the total number of bytes sent and received by the first 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 339](#).

---

**Important:** Flow Export must be turned on before viewing the Top Talker report. For details, see [“Configuring Subnet Side Rules” on page 99](#).

---

## What This Report Tells You

The Top Talkers report answers this question:

- Who were the top talking hosts on a per-port basis?

## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The data is collected at a 5-minute granularity for the entire month.

## To view the Top Talkers report

1. Choose Reports > Networking > Top Talkers to display the Top Talkers page.

Figure 9-18. Top Talkers Page

Reports > Networking > Top Talkers ?

Chart: By Conversation ▼ Period: All ▼ Count: 50 Protocol: Both ▼ Traffic Type: Both ▼ Go

**Top Conversations:**

Rank ↑↓	IP Address 1:Port ↑↓	IP Address 2:Port ↑↓	Byte Count ↑↓
1	10.32.129.24:52013	208.117.254.166:80	85.7 MB
2	10.0.4.104:443	10.32.129.135:63714	54.4 MB
3	10.0.4.104:443	10.32.129.135:63708	41.3 MB
4	10.32.129.24:42396	192.235.0.6:80	35.5 MB
5	10.32.128.20:57157	74.125.0.144:80	31.4 MB
6	10.32.129.20:64487	204.93.181.179:1935	28.9 MB
7	10.32.128.20:55279	208.117.254.160:80	24.1 MB
8	10.0.4.104:443	10.32.129.136:56167	23.0 MB
9	10.32.3.82:7800	10.32.128.20:55259	20.6 MB
10	10.0.4.104:443	10.32.129.135:55580	18.2 MB
11	10.32.129.24:42689	192.235.0.6:80	17.8 MB
12	10.0.4.104:443	10.32.129.136:56165	16.0 MB
13	10.0.4.104:443	10.32.129.135:63715	15.7 MB
14	10.0.4.104:443	10.32.129.135:54068	15.7 MB
15	10.16.205.53:41698	10.32.129.136:51215	13.7 MB
16	10.16.205.53:7830	10.32.129.136:56278	13.5 MB
17	10.32.128.20:50695	74.125.0.138:80	13.2 MB
18	10.32.128.20:52162	169.233.6.159:48573	12.3 MB
19	10.32.128.20:50539	65.78.184.147:31504	11.0 MB
20	10.32.129.136:57264	74.125.224.74:443	9506.8 kB
21	10.32.128.20:55921	208.117.254.21:80	8903.3 kB
22	10.32.128.20:51740	173.194.29.216:80	8851.6 kB
23	10.32.128.20:55106	68.224.132.68:11428	7672.8 kB
24	10.32.128.20:49405	74.125.224.71:80	7609.7 kB
25	10.0.4.104:443	10.32.129.135:64501	7082.5 kB
26	10.32.128.20:49404	74.125.224.71:80	6874.0 kB
27	10.16.5.10:7810	10.32.129.135:51022	6834.7 kB

2. Use the controls to customize the report as described in this 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	<p>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.</p> <p><b>Note:</b> Top Talker statistics are not persistent between service restarts.</p>
Count	<p>Specify how many top pairs of IP 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 IP address 1. The default value is 50.</p> <p><b>Note:</b> You can export the complete list of top talkers to a file in CSV format using the Export report.</p>
Protocol	Select Both, TCP, or UDP from the drop-down list. The default value is Both.

Control	Description
Traffic Type	Select Both, Optimized, or Passthrough from the drop-down list. The default value is Both.
Go	Displays the report.

**Tip:** 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. For details about setting ports to be monitored, see [“Configuring Monitored Ports” on page 234](#).

The Steelhead appliance automatically discovers all the ports in the system that have traffic. The discovered port and its label (if one exists) are added to the report. If a label does not exist, 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 re-add the port with a new label. All statistics for this new port label are preserved from the time the port was discovered.

---

**Note:** The Traffic Summary report displays a maximum of 16 ports and pie slices for the traffic types comprising more than 0.005 percent of the total traffic (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 traffic that comprises less than 0.005 percent of the total is not included in the Traffic Summary report, but is aggregated into the Bandwidth Optimization report.

---

The Traffic Summary report provides this table of statistics that describe data activity for the application and the time period you specify.

Column	Description
Port	Displays the TCP/IP port number and application for each row of statistics.
Reduction	Displays the amount of application data reduction.
LAN Data	Displays the amount of application data on the LAN.
WAN Data	Displays 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 these questions:

- How much data reduction has occurred?
- What was the percentage of the total traffic for each port?

## About Report Data

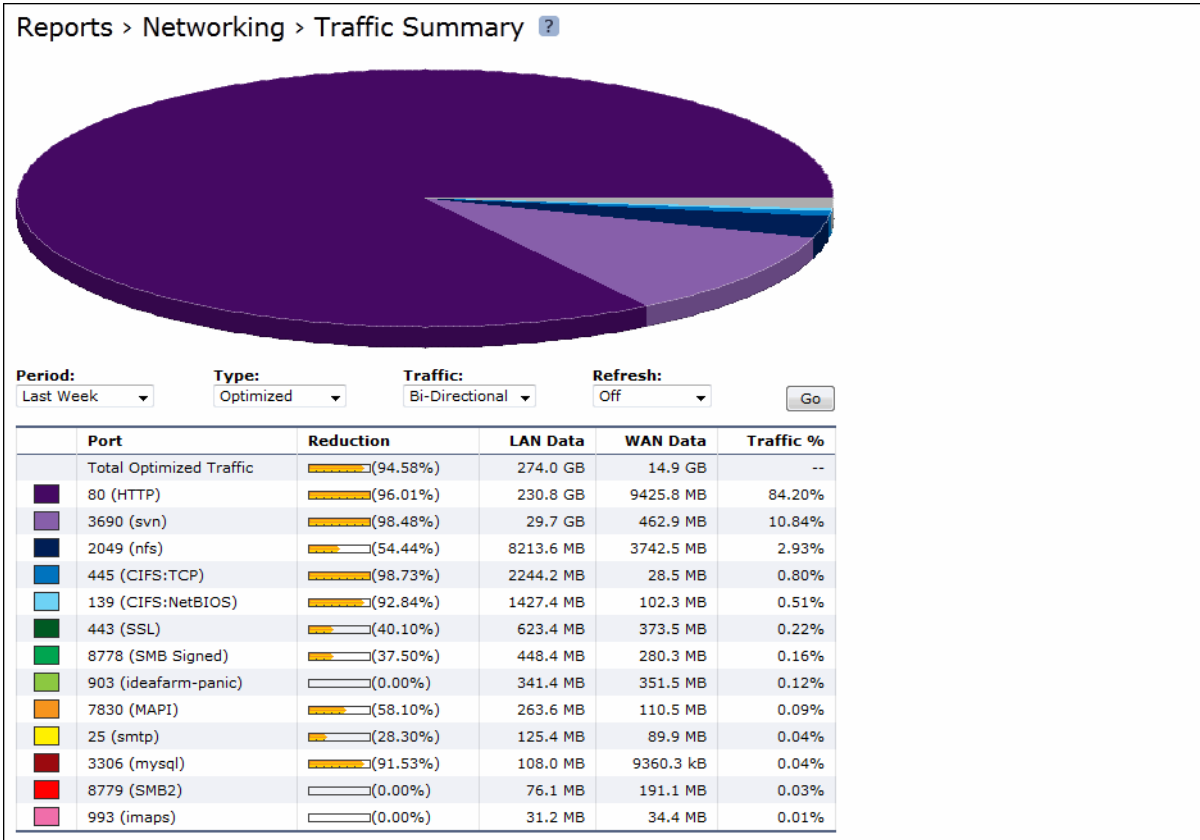
The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The Traffic Summary report displays these data granularities:

- Last 1 hour's worth of data is available at 10-second granularity.
- Last 1 day's worth of data is available at 5-minute granularity.
- Last 1 week's worth of data is available at 1-hour granularity.
- Last 1 month's worth of data is available at 2-hour granularity.

To view the Traffic Summary report

- 1. Choose Reports > Networking > Traffic Summary to display the Traffic Summary page.

Figure 9-19. Traffic Summary Page



- 2. Use the controls to customize the report as described in this 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 <b>Go</b> . Use the following format: YYYY/MM/DD HH:MM:SS
Refresh	Select a refresh rate from the drop-down list: <ul style="list-style-type: none"><li>• To refresh the report every 10 seconds, select 10 seconds.</li><li>• To refresh the report every 30 seconds, select 30 seconds.</li><li>• To refresh the report every 60 seconds, select 60 seconds.</li><li>• To turn refresh off, click <b>Off</b>.</li></ul>
Go	Displays the report.



## Viewing WAN Throughput Reports

The WAN Throughput report summarizes the WAN throughput for the time period specified. In standard in-path and virtual in-path deployments, the throughput is an aggregation of all data the system transmits out of all WAN interfaces. In a server-side out-of-path configuration, the report summarizes all data the system transmits out of the primary interface.

For details about the report format, see [“Overview” on page 244](#).

You must choose Configure > Networking > Flow Statistics and enable WAN Throughput Statistics to view data on this report. WAN throughput statistics are enabled by default.

The WAN Throughput report does not include any traffic that is bypassed, either by an in-path interface in hardware bypass, or the portion of traffic that is bypassed by hardware-assist rules on supported Fiber 10 Gigabit-Ethernet in-path cards.

The WAN Throughput report includes a WAN link throughput graph that provides these statistics describing data activity for the time period you specify.

Data Series	Description
Peak Throughput	Displays the peak data activity.
Average Throughput	<p>Displays the average and total throughput.</p> <p>RiOS calculates the WAN average at each data point by taking the number of bytes transferred, converting that to bits, and then dividing by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p> <p>The total throughput shows the data amount transferred during the displayed time interval.</p> <p>The average that appears below the Average Throughput is an average of all displayed averages.</p>

The navigator shadows the Peak throughput series.

In some configurations, RiOS transmits LAN traffic out of WAN interfaces: for example, virtual in-path deployments and deployments using the default gateway on the WAN side without simplified routing. In such deployments, you can configure subnet side rules to decide which channel traffic is not destined for the WAN. For details, see [“Configuring Subnet Side Rules” on page 99](#).

## What This Report Tells You

The WAN Throughput report answers these questions:

- What was the average WAN throughput?
- What was the peak WAN throughput?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.


## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled with a granularity of 5 minutes for the day, 1 hour for the last week, and 2 hours for the rest of the month.

### To view the WAN Throughput report

1. Choose Reports > Optimization > WAN Throughput to display the WAN Throughput page.

Figure 9-20. WAN Throughput Page



2. Use the controls to change the display as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>Because the system aggregates data on the hour, request hourly time intervals. For example, setting a time interval to 08:30:00 to 09:30:00 from 2 days ago does not create a data display, whereas setting a time interval to 08:00:00 to 09:00:00 from 2 days ago will display data.</p> <p>When you request a custom time interval to view data beyond the aggregated granularity, the data is not visible because the system is no longer storing the data. For example, these custom time intervals do not return data because the system automatically aggregates data older than 7 days into 2-hour data points:</p> <ul style="list-style-type: none"> <li>• Setting a 1-hour time period that occurred 2 weeks ago.</li> <li>• Setting a 75-minute time period that occurred more than 1 week ago.</li> </ul> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>

## Viewing Application Statistics Reports

The Application Statistics report provides a tabular summary or a graph of the traffic flowing through a Steelhead appliance for the time period specified. You can view up to seven applications in a stacked view.

You must enable application visibility on the Configure > Networking > Flow Statistics page before the Application Statistics report can gather and display statistics. For details, see [“Configuring Flow Statistics” on page 101](#).

RiOS collects application statistics for all data transmitted out of the WAN and primary interfaces and commits samples every 5 minutes. Let the system collect statistics for awhile to view the most meaningful data display.

For details about the report format, see [“Overview” on page 244](#).

The Application Statistics report includes these statistics for each listed application, traffic direction, and the time period you specify.

Data Series	Description
Average bps	Displays the average data activity in all flows of an application in bits per second. The minimum sample granularity is 5 minutes.
Per Flow Average bps	<p>Displays the average trended throughput in all traffic flows of an application in bits per second. This indicates how bandwidth intensive an application is per user or flow.</p> <p>RiOS calculates the WAN average at each data point by taking the number of bytes transferred, converting that to bits, and then dividing by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p>

Data Series	Description
Peak bps	Displays the peak data activity in bits per second. For larger granularity data points, this represents the largest 5 minute average within. For 5 minutes, this is the same as the average.
Per flow Peak bps	Displays the peak trended data activity per traffic flow in bits per second. This is the largest per flow 5-minute bps within a larger sample.

This report displays applications within their protocol hierarchy. For example, Facebook appears as TCP > HTTP > Facebook.

This report lists unrecognized applications by their server port; for example, TCP > Unknown (port 5001).

## What This Report Tells You

The Application Statistics report answers this question:

- How much bandwidth is a particular application using?

## About Report Graphs

While viewing the application statistics in a graph, use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

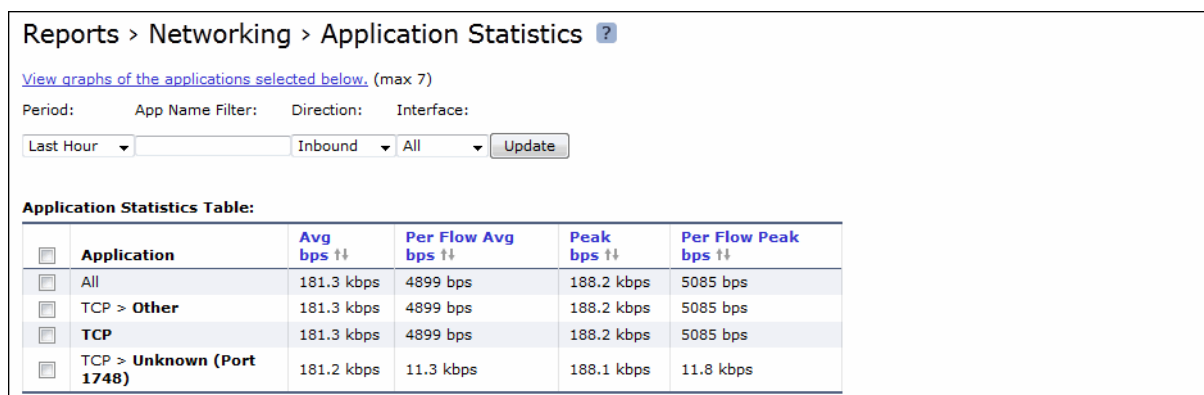
## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The data is collected at a 5-minute granularity for the entire month.

### To view the Application Statistics report

1. Choose Reports > Networking > Application Statistics to display the Application Statistics page.
2. Click View graphs of the applications selected below to switch from a tabular display to a graph.

Figure 9-21. Application Statistics Page



3. Use the controls to change the report display as described in this 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 <b>Go</b> . Use the format YYYY/MM/DD HH:MM:SS.
App Name Filter	Click a protocol or application name (for example, TCP, LDAP, SharePoint) to show only the selection.  You can select only one filter at a time. For example, if the report is filtering on UDP and you click TCP, the report displays all TCP entries and clears the UDP filter.
Direction	Select the traffic direction from the drop-down list. The default is outbound LAN > WAN traffic.
Interface	Select an interface from the drop-down list. The default is all WAN and primary interfaces.
Update	Click to update the chart without updating the application selection.

## Viewing Application Visibility Reports

The Application Visibility report summarizes the traffic flowing through a Steelhead appliance classified by the application for the time period specified. This report provides application level visibility into layer-7 and shows the application dynamics for pass-through and optimized traffic.

You must enable application visibility on the Configure > Networking > Flow Statistics page before the Application Visibility report can gather and display statistics. Application Visibility is enabled by default. For details, see [“Configuring Flow Statistics” on page 101](#).

For details about the report format, see [“Overview” on page 244](#).

This report does not include IPv6 traffic.

The Application Visibility report includes these statistics for each listed application, traffic direction, and the time period you specify.

Data Series	Description
App Throughput	Displays the throughput for all traffic flows in bits per second. The minimum sample granularity is 5 minutes.  <b>Throughput Peak</b> - Hover the mouse over the data series to display the peak data activity in bits per second. For larger granularity data points, this represents the largest 5 minute average within. For 5 minutes, this is the same as the average.  <b>Throughput Average</b> - Hover the mouse over the data series to display the average trended throughput for all traffic flows in kbps.

Data Series	Description
Per-Flow Throughput	<p>Displays the throughput per traffic flow in bits-per-second.</p> <p><b>Per-Flow Peak</b> - Hover the mouse over the data series to display the peak trended data activity per traffic flow in bits per second. This is the largest per flow 5-minute bps within a larger sample.</p> <p><b>Per-Flow Average</b> - Hover the mouse over the data series to display the average trended throughput in all traffic flows of an application in bits per second. This indicates how bandwidth-intensive an application is per user or flow.</p> <p>RiOS calculates the WAN average at each data point by taking the number of bytes transferred, converting that to bits, and then dividing by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p>

---

The navigator shadows the Per-flow throughput series.

## What This Report Tells You

The Application Visibility report answers this question:

- How much bandwidth is a particular application using?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.


## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled. The data is collected at a 5-minute granularity for the entire month.

## To view the Application Visibility report

1. Choose Reports > Networking > Application Visibility to display the Application Visibility page.

Figure 9-22. Application Visibility Page



2. Use the controls to change the report display as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Interface	Select an interface from the drop-down list. The default is all interfaces.
Direction	Select the traffic direction from the drop-down list. The default is outbound LAN > WAN traffic.
Application Name	<p>Select an application from the drop-down list. To narrow the search, start typing the first characters in the application name. When the application name and definition appears, select it from the list.</p> <p>You can select up to 7 applications.</p> <p>Click <b>Update Chart</b> to update the chart without changing the application selection.</p>

---

## Viewing Interface Counter Reports

The Interface Counters report summarizes the statistics for the interfaces. It also displays the IP address, speed, duplex, MAC address, and current status of each interface.

In RiOS v8.5, this report includes interfaces configured with IPv6 addresses.

For automatically negotiated speed and duplex settings, the Interface Counters report displays the speed at which they are negotiated.

Interface statistics display the data accumulated since the last reboot.

The Interface Counters report displays the statistics described in this table.

Column	Description
Interface	<b>LAN</b> - Displays statistics for the LAN interface. <b>WAN</b> - Displays statistics for the WAN interface. <b>Primary</b> - Displays statistics for the primary interface. <b>Aux</b> - Displays statistics for the auxiliary interface. <b>Inpath</b> - Displays statistics for the in-path interface.
IP	Displays the IP address (if applicable) for the interface.
Ethernet	Displays 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	Displays true or false to indicate whether the link is up or down.
Receive Packets	Displays the total number of packets, packets discarded, errors encountered, packets overrun, frames sent, and multicast packets sent.
Transmit Packets	Displays 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 these questions:

- How many packets is the appliance transmitting or receiving?
- 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 9-23. Interface Counters Page

Reports > Networking > Interface Counters ?

The network interface statistics have been collected since the system was booted 20 hours, 47 minutes ago (or since the statistics were last cleared).

**Interface Statistics:**

Interface	IP	Ethernet	Link	Receive Packets	Transmit Packets
primary	10.11.140.52/16	MAC: 00:0E:B6:03:6C:A0 Speed: 1000Mb/s (auto) Duplex: full (auto)	true	4846 packets 0 discards 0 errors 0 overruns 0 frames 0 multicast	1231 packets 0 discards 0 errors 0 overruns 0 carriers 0 collisions
aux	10.3.2.52/21	MAC: 00:0E:B6:4A:FD:71 Speed: 100Mb/s (auto) Duplex: full (auto)	true	390120 packets 0 discards 0 errors 0 overruns 0 frames 1 multicast	31958 packets 0 discards 0 errors 0 overruns 0 carriers 0 collisions
inpath0_0 (main)	10.11.142.52/16	MAC: 00:0E:B6:9B:7D:CC Speed: N/A Duplex: N/A	N/A	64907 packets 0 discards 0 errors 0 overruns 0 frames 0 multicast	18435 packets 0 discards 0 errors 0 overruns 0 carriers 0 collisions
lan0_0 (main)	N/A	MAC: 00:0E:B6:9B:7D:CC Speed: 1000Mb/s (auto) Duplex: full (auto)	true	4289 packets 0 discards 0 errors 0 overruns 0 frames 1449 multicast	45987 packets 0 discards 0 errors 0 overruns 0 carriers 0 collisions
wan0_0 (main)	N/A	MAC: 00:0E:B6:9B:7D:CD Speed: 1000Mb/s (auto) Duplex: full (auto)	true	64199 packets 0 discards 0 errors 0 overruns 0 frames 41137 multicast	22561 packets 0 discards 0 errors 0 overruns 0 carriers 0 collisions

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 this table of statistics that summarize TCP activity.

Packet Type	Description
Packets Received	Displays the total packets received.
Packets Sent	Displays the total TCP packets sent.
Packets Retransmitted	Displays the total TCP packets retransmitted.
Packets Fast Retransmitted	Displays 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	Displays the number of time-outs.
Loss Events	Displays the total number of loss events.

## What This Report Tells You


The TCP Statistics report answers these 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 9-24. TCP Statistics Page**



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 for the port, traffic direction, and time period specified.

For details about the report format, see [“Overview” on page 244](#).

The Optimized Throughput report includes LAN and WAN link throughput graphs that include these statistics describing data activity for the port, traffic direction, and the time period you specify.

Data Series	Description
LAN Peak	Displays the peak data activity.
LAN P95	Displays the 95th percentile for data activity. The 95th percentile is calculated by taking the peak of the lower 95 percent of inbound and outbound throughput samples.
LAN Average	<p>Displays the average throughput.</p> <p>RiOS calculates the LAN average at each data point by taking the number of bytes transferred, converting that to bits, and then dividing by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p> <p>The average that appears below the LAN Average is an average of all displayed averages.</p>
WAN Peak	Displays the peak data activity.
WAN P95	Displays the 95th percentile for data activity. The 95th percentile is calculated by taking the peak of the lower 95 percent of inbound and outbound throughput samples.

Data Series	Description
WAN Average	<p>Displays the average throughput.</p> <p>RiOS calculates the WAN average at each data point by taking the number of bytes transferred, converting that to bits, and then dividing by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p> <p>The average that appears below the WAN Average is an average of all displayed averages.</p>

The navigator shadows the WAN Peak series.

## What This Report Tells You

The Optimized Throughput report answers these questions:

- What was the average WAN and LAN throughput?
- What was the peak WAN and LAN throughput?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

## About Report Data

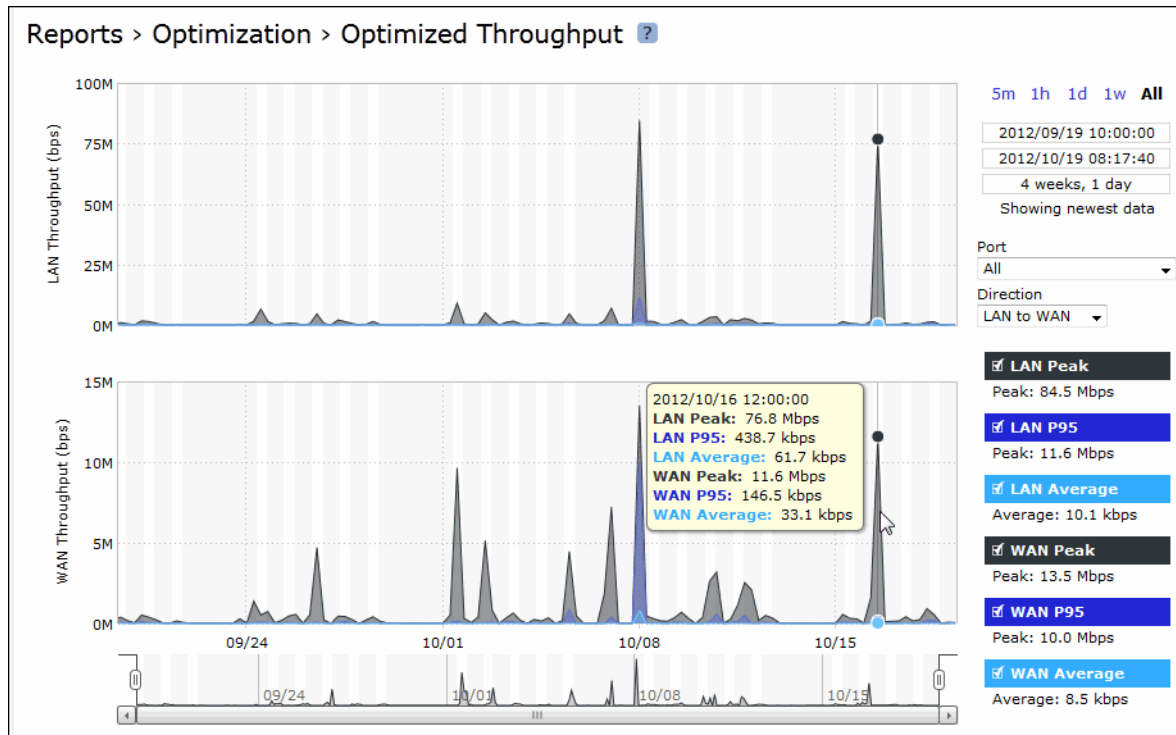
The Riverbed system reports on performance for periods up to one month. However, due to performance and disk space considerations, data representation in reports for periods longer than the last 5 minutes are interpolated from aggregate data points. The Optimized Throughput report displays these data granularities:

- Last 1 hour's worth of data is available at 10-second granularity.
- Last 1 day's worth of data is available at 5-minute granularity.
- Last 1 week's worth of data is available at 1-hour granularity.
- Last 1 month's worth of data is available at 2-hour granularity.

## To view the Optimized Throughput report

1. Choose Reports > Optimization > Optimized Throughput to display the Optimized Throughput page.

Figure 9-25. Optimized Throughput Page



2. Use the controls to change the display as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>Because the system aggregates data on the hour, request hourly time intervals. For example, setting a time interval to 08:30:00 to 09:30:00 from 2 days ago does not create a data display, whereas setting a time interval to 08:00:00 to 09:00:00 from 2 days ago will display data.</p> <p>When you request a custom time interval to view data beyond the aggregated granularity, the data is not visible because the system is no longer storing the data. For example, these custom time intervals do not return data because the system automatically aggregates data older than 7 days into 2-hour data points:</p> <ul style="list-style-type: none"> <li>• Setting a 1-hour time period that occurred 2 weeks ago.</li> <li>• Setting a 75-minute time period that occurred more than 1 week ago.</li> </ul> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Direction	Select a traffic direction (Bi-directional, WAN to LAN, or LAN to WAN) from the drop-down list.
Port	Select a port or All to display all of the TCP ports on which the Steelhead appliance has seen traffic. The list appends the port name to the number where available.

## Viewing Bandwidth Optimization Reports

The Bandwidth Optimization report summarizes the overall inbound and outbound bandwidth improvements on your network. You can create reports according to the time period, port, and traffic direction of your choice.

For details about the report format, see [“Overview” on page 244](#).

The Bandwidth Optimization report includes these statistics describing bandwidth activity for the time period you specify.

Data Series	Description
Data Reduction %	<p>Displays the peak and total decrease of data transmitted over the WAN, according to this calculation:</p> $(\text{Data In} - \text{Data Out}) / (\text{Data In})$ <p>Displays the capacity increase x-factor below the peak and total data reduction percentages.</p>
WAN and LAN Throughput	<p>Depending on which direction you select, specifies one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Bi-Directional</b> - traffic flowing in both directions</li> <li>• <b>WAN-to-LAN</b> - inbound traffic flowing from the WAN to the LAN</li> <li>• <b>LAN-to-WAN</b> - outbound traffic flowing from the LAN to the WAN</li> </ul>

The navigator shadows the data reduction series.

## What This Report Tells You

The Bandwidth Optimization report answers these questions:

- How much data reduction has occurred?
- How much data was removed from the WAN link?
- How much data was sent/received on the LAN/WAN ports?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

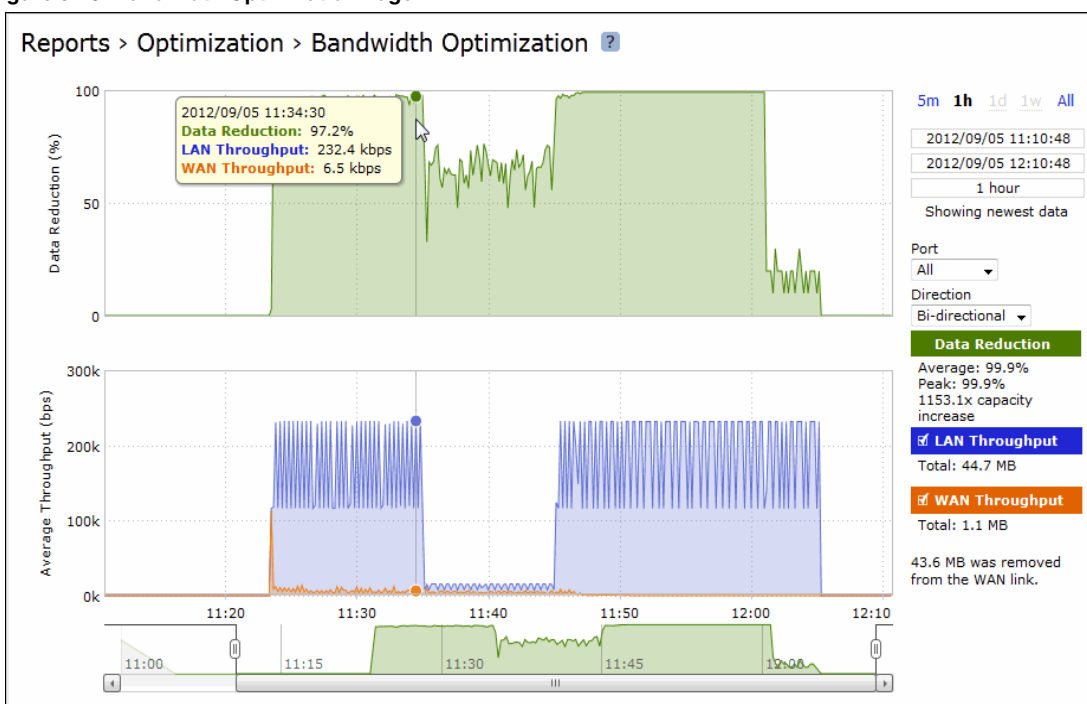
## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled with a granularity of 5 minutes for the day, 1 hour for the last week, and 2 hours for the rest of the month.

### To view a Bandwidth Optimization report

1. Choose Reports > Optimization > Bandwidth Optimization to display the Bandwidth Optimization page.

Figure 9-26. Bandwidth Optimization Page



2. Use the controls to customize the report as described in this table.



Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 week (1w), All, or type a custom date. All includes statistics for the past 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>Because the system aggregates data on the hour, request hourly time intervals. For example, setting a time interval to 08:30:00 to 09:30:00 from 2 days ago does not create a data display, whereas setting a time interval to 08:00:00 to 09:00:00 from 2 days ago will display data.</p> <p>When you request a custom time interval to view data beyond the aggregated granularity, the data is not visible because the system is no longer storing the data. For example, these custom time intervals do not return data because the system automatically aggregates data older than 7 days into 2-hour data points:</p> <ul style="list-style-type: none"> <li>• Setting a 1-hour time period that occurred 2 weeks ago.</li> <li>• Setting a 75-minute time period that occurred more than 1 week ago.</li> </ul> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Port	Select a port or All to select all ports from the drop-down list.
Direction	Select a traffic direction (Bi-Directional, WAN to LAN, or LAN to WAN) from the drop-down list.


## Viewing Peer Reports

The Peers report summarizes the peer Steelhead appliances. The Peers report contains this table of statistics that summarize connection peer activity.

Column	Description
Name	Specifies the name of the peer appliance.
IP Address	Specifies the IP address of the peer appliance.
Model	Specifies the appliance model.
Version	Specifies the appliance version.
Licenses	Specifies the current appliance licenses.

The report includes both connected and unconnected peers. The connected icon appears next to a connected peer. A dimmed icon indicates that the peer is disconnected:

	STAIPUTER2	10.32.128.21	Steelhead Mobile	3.1.3c-vantiv #133_12_1	CIFS, MAPI, ORACLE-FORMS, SCPS, SSL
	DKEY-W7	10.32.129.23	Steelhead Mobile	3.1.3c #133_12	CIFS, MAPI, ORACLE-FORMS, SCPS, SSL

 Indicates that the peer is connected

For details about configuring peering, see [“Enabling Peering and Configuring Peering Rules”](#) on page 53.

## What This Report Tells You

The Peers report answers these questions:

- How many peers are connected to the Steelhead appliance?
- How many peers are disconnected from the Steelhead appliance?

### To view the Peers report

1. Choose Reports > Optimization > Peers to display the Peers page.
2. To view only connected peers, select the Hide Disconnected Peers check box.  
Select a report column heading to sort the column in ascending or descending order.  
To open the Management Console for a peer, click the peer name or IP address.

---

## Viewing SRDF Reports

The SRDF report presents information regarding optimized throughput and data reduction for EMC's Symmetrix Remote Data Facility (SRDF) protocol. You can view a summary of performance statistics for all optimized SRDF traffic, as well as drill into metrics for all remote data facility (RDF) groups for a specific EMC Symmetrix array or an individual RDF group within an array.

For details about the report format, see [“Overview” on page 244](#).

---

**Note:** You can also check the total optimized SRDF traffic throughput by viewing the Reports > Optimization > Optimized Throughput report.

---

SRDF reports contain this information:

Data Series	Description
Data Reduction	Specifies the percentage of total decrease in overall data transmitted (when viewing all Symmetrix RDF groups).
WAN/LAN Throughput	Specifies the total throughput transmitted over the WAN and LAN.

When the report display includes all Symmetrix RDF groups or a single RDF group for a single Symmetrix ID, the navigator shadows the LAN/WAN throughput series. When the report display includes all RDF groups for a single Symmetrix ID, the navigator shadows the group 1 LAN/WAN throughput series.

## What This Report Tells You

The SRDF report answers these questions:

- How much total SRDF traffic is the Steelhead appliance processing over time?
- How much data reduction is being delivered overall?
- How much data reduction is being delivered for individual RDF groups?
- Which Symmetrix array is generating the most SRDF traffic?



- How are SRDF traffic patterns changing over time?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

## About Report Data

The Riverbed system reports for periods up to one month. Due to performance and disk space considerations, data representation in reports for periods longer than the latest five minutes are interpolated between data points obtained by aggregating multiple 10-second samples. The display granularity decreases with time passed since data was sampled. The data is collected at a 5-minute granularity for the entire month.

### To view the SRDF report

1. On the client-side Steelhead appliance, choose Reports > Optimization > SRDF to display the SRDF page.

Figure 9-27. SRDF Page



2. Use the controls to customize the report as described in this table.

Control	Description
Time interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>You can view the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Symmetrix (RDF Group)	<p>Select a Symmetrix server ID from the drop-down list to view detailed statistics per Symmetrix ID.</p> <p>Use the <b>protocol srdf</b> CLI command to map a logical Symmetrix ID to its set of network IP addresses. For example, the following commands create a Symmetrix ID, Sym1, and associate it with traffic originating from IP addresses 10.12.61.42 and 10.12.61.43:</p> <pre>protocol srdf symm id Sym1 address 10.12.61.42 protocol srdf symm id Sym1 address 10.12.61.43</pre> <p>RiOS maps SRDF traffic originating from IP addresses that have not been mapped to a Symmetrix ID to the default Symmetrix ID, represented by DefaultSymm for this field.</p> <p>Select an RDF group number from the drop-down list to view data reduction information for individual RDF groups. You can use data reduction information to fine-tune the optimization settings for those RDF groups. The Steelhead appliance automatically identifies and summarizes information by RDF group based on the SRDF traffic seen by the Steelhead appliance.</p> <p>Peak lines appear after one hour for RDF group detail reports.</p>
Traffic Type	<p>Select either LAN or WAN to display the amount of data transmitted over the LAN/WAN during the selected time period.</p>

### Related Topic

- [“Configuring SRDF Optimization” on page 80](#)

## Viewing SnapMirror Reports

The SnapMirror report displays how much benefit SnapMirror optimization is providing for a given filer (or all filers) and traffic type in the time period specified. You can use this report to view optimization outcomes for a filer, all volumes for a single filer, or a single filer for a volume or qtree. You can drill down to specific optimization statistics for a volume or a qtree.

SnapMirror captures and reports only traffic flowing in the LAN-to-WAN direction.

For details about the report format, see [“Overview” on page 244](#).

SnapMirror reports contain this information:

Data Series	Description
Peak LAN/WAN Throughput	Displays the peak LAN/WAN data activity. The system stores peak statistics in terms of bytes transferred over the LAN, but calculates the normal throughput using a granularity of 10 seconds.
Average LAN/WAN Throughput	<p>Displays the average LAN/WAN data activity. The system stores non-peak statistics as the number of bytes transferred over the LAN/WAN, and calculates the throughput by converting bytes to bits and then dividing the result by the granularity.</p> <p>For instance, if the system reports 100 bytes for a data point with a 10-second granularity, RiOS calculates:</p> $100 \text{ bytes} * 8 \text{ bits/byte} / 10 \text{ seconds} = 80 \text{ bps}$ <p>This means that 80 bps was the average throughput over that 10-second period.</p> <p>The total throughput shows the data amount transferred during the displayed time interval.</p>
Data Reduction	<p>Specifies the percentage of total decrease in overall data transmitted (when viewing all SnapMirror filers). The system calculates data reduction as (total LAN data - total WAN data) / total LAN data.</p> <p>You can use data reduction information to fine-tune the optimization settings for a filer, a filer and a volume, or a filer, volume, and qtree.</p>

The navigator shadows the Throughput series.

## What This Report Tells You

The SnapMirror report answers this question:

- How much total SnapMirror traffic is the Steelhead appliance processing over time?

## About Report Graphs

Use the mouse to hover over a specific data point to see what the y values and exact timestamp were in relation to peaks.

## About Report Data

The Riverbed system reports for periods up to one month. Due to performance and disk space considerations, data representation in reports for periods longer than the latest five minutes are interpolated between data points obtained by aggregating multiple 10-second samples. The display granularity decreases with time passed since data was sampled. The data is collected at a 5-minute granularity for the entire month.

## To view the SnapMirror report

1. On the client-side Steelhead appliance, choose Reports > Optimization > SnapMirror to display the SnapMirror page.

Figure 9-28. SnapMirror Page Displaying All Filers

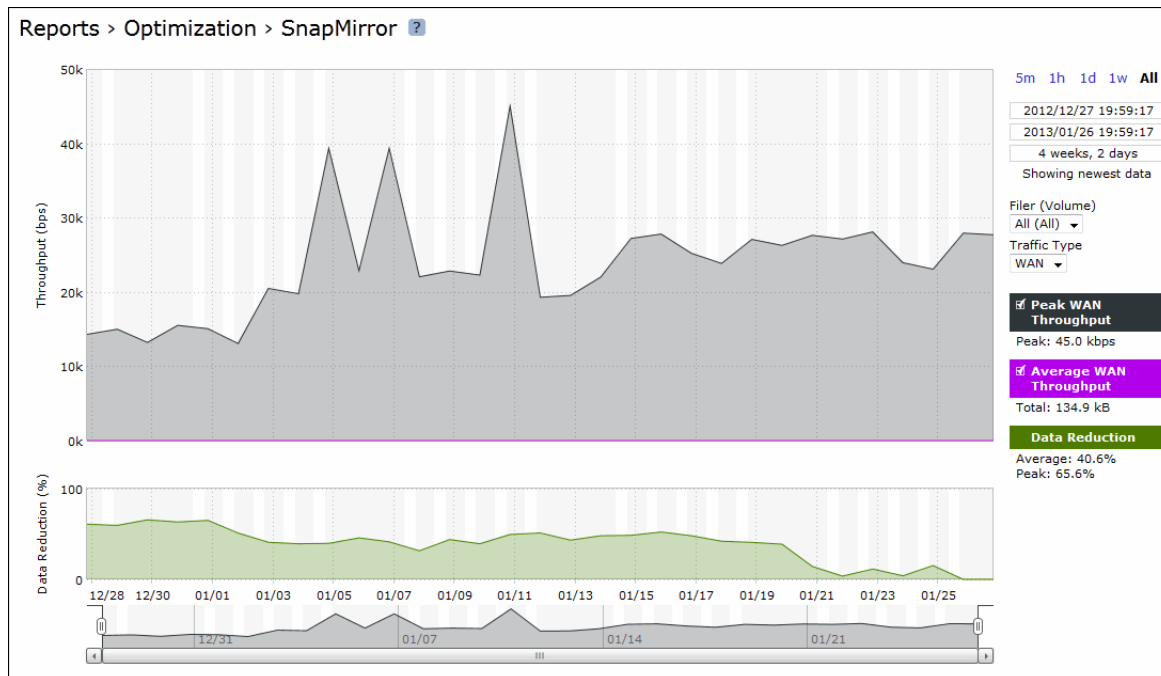
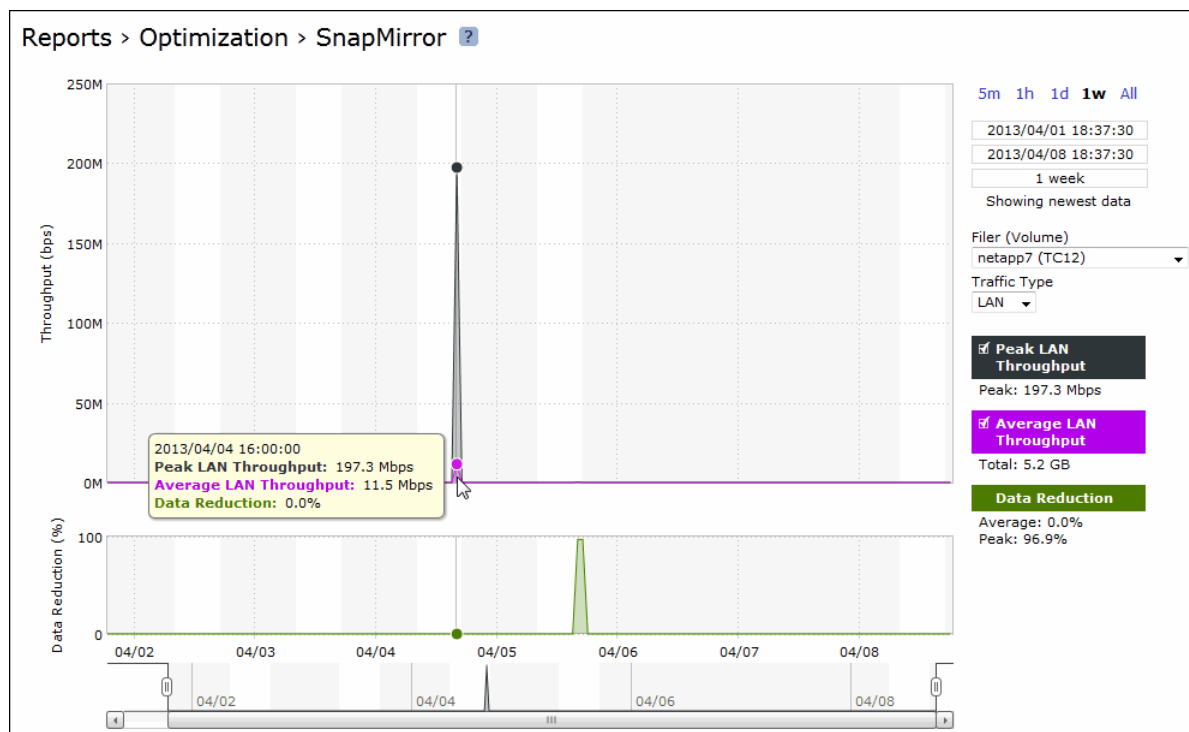


Figure 9-29. SnapMirror Page For a Filer and Volume



---

## Viewing Alarm Status Reports

The Alarm Status report provides status for the Steelhead appliance alarms.

The Steelhead appliance tracks key hardware and software metrics and alerts you of any potential problems so you can quickly discover and diagnose issues.

RiOS groups certain alarms into top-level categories, such as the SSL Settings alarm. When an alarm triggers, its parent expands to provide more information. For example, the System Disk Full top-level alarm aggregates over multiple partitions. If a specific partition is full, the System Disk Full alarm triggers and the Alarm Status report displays more information regarding which partition caused the alarm to trigger.

The health of an appliance falls into one of these states:

- **Needs Attention** - Accompanies a healthy state to indicate management-related issues not affecting the ability of the Steelhead appliance to optimize traffic.
- **Degraded** - The Steelhead appliance is optimizing traffic but the system has detected an issue.
- **Admission Control** - The Steelhead appliance is optimizing traffic but has reached its connection limit.
- **Critical** - The Steelhead appliance might or might not be optimizing traffic; you must address a critical issue.

The Alarm Status report includes this alarm information.

Alarm	Steelhead Appliance State	Reason
Admission Control	Admission Control	<ul style="list-style-type: none"> <li>• <b>Connection Limit</b> - Indicates that 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.</li> <li>• <b>CPU</b> - Indicates that the Steelhead appliance has entered admission control due to high CPU use. During this event, the Steelhead appliance continues to optimize existing connections, but passes through new connections without optimization. The alarm clears automatically when the CPU usage decreases.</li> <li>• <b>Memory</b> - Indicates that 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 decreases.</li> <li>• <b>TCP</b> - Indicates that 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. The alarm clears automatically when the TCP memory pressure decreases.</li> </ul>
Asymmetric Routing	Needs Attention	Indicates that 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 about the Asymmetric Routing table, see <a href="#">“Configuring Asymmetric Routing Features” on page 91</a> .

Alarm	Steelhead Appliance State	Reason
Connection Forwarding	Degraded	<p>Indicates that 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 <i>one</i> of the neighbors is in error. In the same way, the alarm clears only when all three neighbors are no longer in error.</p> <ul style="list-style-type: none"> <li>• <b>Cluster IPv6 Incompatible</b> - Indicates that a connection-forwarding neighbor is running a RiOS version that is incompatible with IPv6. Neighbors must be running RiOS v8.5. The Steelhead appliance neighbors pass through IPv6 connections when this alarm triggers.</li> <li>• <b>Multiple Interface</b> - Indicates that the connection to a Steelhead appliance in a connection forwarding cluster is lost.</li> <li>• <b>Single Interface</b> - Indicates that the connection to a Steelhead appliance connection-forwarding neighbor is lost.</li> </ul> <p>These issues trigger the single connection-forwarding alarm:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• The connection cannot be established with a connection-forwarding neighbor.</li> <li>• The connection has been closed by the connection-forwarding neighbor.</li> <li>• The connection has been lost with the connection-forwarding neighbor due to an error.</li> <li>• The connection has been lost because requests have not been acknowledged by a connection-forwarding neighbor within the set threshold.</li> <li>• The Steelhead appliance has timed out while waiting for an initialization message from a connection-forwarding neighbor.</li> <li>• The amount of latency between connection-forwarding neighbors has exceeded the specified threshold.</li> </ul>
CPU Utilization	Degraded	<p>Indicates that 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 <a href="#">“Configuring Alarm Settings” on page 217</a>.</p> <p>If your alarm thresholds are correct, reboot the Steelhead appliance. For details, see <a href="#">“Rebooting and Shutting Down the Steelhead Appliance” on page 191</a>.</p>
Disk Full		<p>Indicates that the system partitions (not the RiOS data store) are full or almost full. For example, RiOS monitors the available space on <b>/var</b>, which is used to hold logs, statistics, system dumps, TCP dumps, and so on.</p> <p>Examine the directory to see if it is storing an excessive amount of snapshots, system dumps, or TCP dumps that you could delete. You could also delete any RiOS images that you no longer use.</p>
Flash Protection Failure	Critical	<p>Indicates that the USB flash drive has not been backed up because there is not enough available space in the <b>/var</b> filesystem directory.</p> <p>Examine the <b>/var</b> directory to see if it is storing an excessive amount of snapshots, system dumps, or TCP dumps that you could delete. You could also delete any RiOS images that you no longer use.</p>

Alarm	Steelhead Appliance State	Reason
Hardware	Either Critical or Degraded, depending on the state	<ul style="list-style-type: none"> <li>• <b>Disk Error</b> - Indicates that one or more disks is offline. To see which disk is offline, enter this CLI command from the system prompt: <code>show raid diagram</code>  This alarm applies only to the Steelhead appliance RAID Series 3000, 5000, and 6000.</li> <li>• <b>Fan Error</b> - Indicates that a fan is failing or has failed and must be replaced.</li> <li>• <b>Flash Error</b> - Indicates an error with the flash drive hardware. At times, the USB flash drive that holds the system images might become unresponsive; the Steelhead appliance continues to function normally. When this error triggers you cannot perform a software upgrade, as the Steelhead appliance is unable to write a new upgrade image to the flash drive without first power cycling the system.  To reboot the appliance, go to the Configure &gt; Maintenance &gt; Reboot/Shut Down page or enter the CLI <b>reload</b> command to automatically power cycle the Steelhead appliance and restore the flash drive to its proper function.</li> <li>• <b>IPMI</b> - Indicates an Intelligent Platform Management Interface (IPMI) event (not supported on all appliance models).  This alarm triggers when there has been a physical security intrusion. These events trigger this alarm: <ul style="list-style-type: none"> <li>• chassis intrusion (physical opening and closing of the appliance case)</li> <li>• memory errors (correctable or uncorrectable ECC memory errors)</li> <li>• hard drive faults or predictive failures</li> <li>• power supply status or predictive failure</li> </ul> By default, this alarm is enabled.</li> <li>• <b>Memory Error</b> - Indicates a memory error (for example, when a system memory stick fails).</li> <li>• <b>Other Hardware Error</b> - Indicates an issue with one of the following: <ul style="list-style-type: none"> <li>• the Steelhead appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration.</li> <li>• 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.</li> <li>• DIMMs are plugged into the Steelhead appliance but RiOS cannot recognize them because: <ul style="list-style-type: none"> <li>– a DIMM is in the wrong slot. You must plug DIMMs into the black slots first and then use the blue slots when all of the black slots are in use.</li> <li>— or —</li> <li>– a DIMM is broken and you must replace it.</li> </ul> </li> <li>• other hardware issues exist.</li> </ul> By default, this alarm is enabled.</li> <li>• <b>Power Supply</b> - Indicates an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted.</li> </ul>



Alarm	Steelhead Appliance State	Reason
		<ul style="list-style-type: none"> <li>• <b>RAID</b> - Indicates an error with the RAID array (for example, missing drives, pulled drives, drive failures, and drive rebuilds). An audible alarm might also sound. To see if a disk has failed, enter this CLI command from the system prompt:  <pre>show raid diagram</pre> For drive rebuilds, if a drive is removed and then reinserted, the alarm continues to be triggered until the rebuild is complete. Rebuilding a disk drive can take 4-6 hours. This alarm applies only to the Steelhead appliance RAID Series 3000, 5000, and 6000.</li> </ul>
Inbound QoS WAN Bandwidth Configuration	Degraded (Needs Attention)	<p>Indicates that the inbound QoS WAN bandwidth for one or more of the interfaces is set incorrectly. You must configure the WAN bandwidth to be less than or equal to the interface bandwidth link rate.</p> <p>This alarm triggers when the system encounters one of these conditions:</p> <ul style="list-style-type: none"> <li>• An interface is connected and the WAN bandwidth is set higher than its bandwidth link rate: for example, if the bandwidth link rate is 1536 kbps, and the WAN bandwidth is set to 2000 kbps.</li> <li>• A non-zero WAN bandwidth is set and QoS is enabled on an interface that is disconnected; that is, the bandwidth link rate is 0.</li> <li>• A previously disconnected interface is reconnected, and its previously configured WAN bandwidth was set higher than the bandwidth link rate. The Management Console refreshes the alarm message to inform you that the configured WAN bandwidth is set higher than the interface bandwidth link rate.</li> </ul> <p>While this alarm appears, the Steelhead appliance puts existing connections into the default class.</p> <p>The alarm clears when you configure the WAN bandwidth to be less than or equal to the bandwidth link rate or reconnect an interface configured with the correct WAN bandwidth.</p> <p>By default, this alarm is enabled.</p>

Alarm	Steelhead Appliance State	Reason
Licensing	Needs Attention, Degraded, or Critical, depending on the state	<p>Indicates whether your licenses are current.</p> <ul style="list-style-type: none"> <li>• <b>Appliance Unlicensed</b> - This alarm triggers if the Steelhead appliance has no BASE or MSPEC license installed for its currently configured model. For details about updating licenses, see <a href="#">“Managing Licenses and Model Upgrades” on page 191</a>.</li> <li>• <b>Autolicense Critical Event</b> - This alarm triggers on a Virtual Steelhead appliance when the Riverbed Licensing Portal cannot respond to a license request with valid licenses. The Licensing Portal cannot issue a valid license for one of these reasons: <ul style="list-style-type: none"> <li>– A newer Virtual Steelhead appliance is already using the token, so you cannot use it on the Virtual Steelhead appliance displaying the critical alarm. Every time the Virtual Steelhead appliance attempts to refetch a license token, the alarm retriggers.</li> <li>– The token has been redeemed too many times. Every time the Virtual Steelhead appliance attempts to refetch a license token, the alarm retriggers.</li> </ul> <p>Discontinue use of the other Virtual Steelhead appliance or contact Riverbed Support.</p> </li> <li>• <b>Autolicense Informational Event</b> - This alarm triggers if the Riverbed Licensing Portal has information regarding the licenses for a Virtual Steelhead appliance. For example, the Virtual Steelhead appliance displays this alarm when the portal returns licenses that are associated with a token that has been used on a different Virtual Steelhead appliance. <p>Make sure that any previous Virtual Steelhead appliances that were licensed with that token are no longer running. The alarm clears automatically the next time the Virtual Steelhead appliance fetches the licenses from the Licensing Portal.</p> </li> <li>• <b>Licenses Expired</b> - This alarm triggers if one or more features has at least one license installed, but all of them are expired.</li> <li>• <b>Licenses Expiring</b> - This alarm triggers if the license for one or more features is going to expire within two weeks.</li> </ul> <p><b>Note:</b> The licenses expiring and licenses expired alarms are triggered per feature; for example: if you install two license keys for a feature, LK1-FOO-xxx (expired) and LK1-FOO-yyy (not expired), the alarms do not trigger, because the feature has one valid license.</p> <p>If the Licenses Expiring alarm triggers, the system status changes to Needs Attention. The Licenses Expired alarm changes the system status to Degraded. Depending on the expiring license, other alarms might trigger simultaneously. For example, if the MSPEC or SH10BASE license expires, the Appliance Unlicensed alarm triggers and changes the health to Critical.</p>

Alarm	Steelhead Appliance State	Reason
Link Duplex	Degraded	<p>Indicates that an interface was not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex error.</p> <p>Choose Configure &gt; Networking &gt; Base Interfaces and examine the Steelhead appliance link configuration. Next, examine the peer switch user interface to check its link configuration. If the configuration on one side is different from the other, traffic is sent at different rates on each side, causing many collisions.</p> <p>To troubleshoot, change both interfaces to automatic duplex negotiation. If the interfaces do not support automatic duplex, configure both ends for full duplex.</p> <p>You can enable or disable the alarm for a specific interface. To disable an alarm, choose Configure &gt; System Settings &gt; Alarms and select or clear the check box next to the link alarm.</p>
Link I/O Errors	Degraded	<p>Indicates that the error rate on an interface has exceeded 0.1 percent while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured LAN connection experiences very few errors. The alarm clears when the error rate drops below 0.05 percent.</p> <p>You can change the default alarm thresholds by entering the <b>alarm link_errors err-threshold xxxxx</b> CLI command at the system prompt. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p> <p>To troubleshoot, try a new cable and a different switch port. Another possible cause is electromagnetic noise nearby.</p> <p>You can enable or disable the alarm for a specific interface; for example, you can disable the alarm for a link after deciding to tolerate the errors. To enable or disable an alarm, choose Configure &gt; System Settings &gt; Alarms and select or clear the check box next to the link name.</p>
Link State	Degraded	<p>Indicates that the system has lost one of its Ethernet links due to an unplugged cable or dead switch port. 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.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable the alarm, choose Configure &gt; System Settings &gt; Alarms and select or clear the check box next to the link name.</p>
Memory Error	Degraded	<p>Indicates that the system has detected a memory error. A system memory stick might be failing. First, try reseating the memory first. If the problem persists, contact Riverbed Support for an RMA replacement as soon as practically possible.</p>
Memory Paging	Degraded	<p>Indicates that 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, reboot the Steelhead appliance. For details, see <a href="#">“Rebooting and Shutting Down the Steelhead Appliance” on page 191</a>. If rebooting does not solve the problem, contact Riverbed Support at <a href="https://support.riverbed.com">https://support.riverbed.com</a>.</p>
Neighbor Incompatibility	Degraded	<p>Indicates that the system has encountered an error in reaching a Steelhead appliance configured for connection forwarding. For details, see <a href="#">“Configuring Connection Forwarding Features” on page 94</a>.</p>

Alarm	Steelhead Appliance State	Reason
Network Bypass	Critical	<p>Indicates that the system is in bypass failover mode. If the Steelhead appliance is in bypass failover mode, restart the optimization service.</p> <p>If restarting the service does not resolve the problem, reboot the Steelhead appliance.</p> <p>If rebooting does not resolve the problem, shut down and restart the Steelhead appliance. For details, see <a href="#">“Rebooting and Shutting Down the Steelhead Appliance” on page 191</a>, and <a href="#">“Starting and Stopping the Optimization Service” on page 187</a>.</p>
Optimization Service	Critical	<ul style="list-style-type: none"> <li>• <b>Internal Error</b> - The optimization service has encountered a condition which might degrade optimization performance. Go to the Configure &gt; Maintenance &gt; Services page and restart the optimization service.</li> <li>• <b>Unexpected Halt</b> - The optimization service has halted due to a serious software error. See if a system dump was created. If so, retrieve the system dump and contact Riverbed Support immediately. For details, see <a href="#">“Viewing Logs” on page 326</a>.</li> <li>• <b>Service Status</b> - The optimization service has encountered an optimization service condition. The message indicates the reason for the condition: <ul style="list-style-type: none"> <li>• optimization service is not running This message appears after an optimization restart. For more information, review the Steelhead appliance logs.</li> <li>• in-path optimization is not enabled This message appears if an in-path setting is disabled for an in-path Steelhead appliance. For more information, review the Steelhead appliance logs.</li> <li>• optimization service is initializing This message appears after a reboot. The alarm clears. For more information, review the Steelhead appliance logs.</li> <li>• optimization service is not optimizing This message appears after a system crash. For more information, review the Steelhead appliance logs.</li> <li>• optimization service is disabled by user This message appears after entering the CLI command <b>no service enable</b> or shutting down the optimization service from the Management Console. For more information, review the Steelhead appliance logs.</li> <li>• optimization service is restarted by user This message appears after the optimization service is restarted from either the CLI or Management Console. You might want to review the Steelhead appliance logs for more information.</li> </ul> </li> </ul>

Alarm	Steelhead Appliance State	Reason
Outbound QoS WAN Bandwidth Configuration	Degraded (Needs Attention)	<p>Indicates that the outbound QoS WAN bandwidth for one or more of the interfaces is set incorrectly. You must configure the WAN bandwidth to be less than or equal to the interface bandwidth link rate.</p> <p>This alarm triggers when the system encounters one of these conditions:</p> <ul style="list-style-type: none"> <li>• An interface is connected and the WAN bandwidth is set higher than its bandwidth link rate: for example, if the bandwidth link rate is 1536 kbps, and the WAN bandwidth is set to 2000 kbps.</li> <li>• A non-zero WAN bandwidth is set and QoS is enabled on an interface that is disconnected; that is, the bandwidth link rate is 0.</li> <li>• A previously disconnected interface is reconnected, and its previously configured WAN bandwidth was set higher than the bandwidth link rate. The Management Console refreshes the alarm message to inform you that the configured WAN bandwidth is set higher than the interface bandwidth link rate.</li> </ul> <p>While this alarm appears, the Steelhead appliance puts existing connections into the default class.</p> <p>The alarm clears when you configure the WAN bandwidth to be less than or equal to the bandwidth link rate or reconnect an interface configured with the correct WAN bandwidth.</p> <p>By default, this alarm is enabled.</p>
Path Selection Path Down	Degraded	<p>Indicates that one of the predefined paths for a connection is unavailable because it has exceeded either the timeout value for path latency or the threshold for observed packet loss.</p> <p>When a path fails, the Steelhead appliance directs traffic through another available path. When the original path comes back up, the Steelhead appliance redirects the traffic back to it.</p>
Process Dump Creation Error	Degraded	<p>Indicates that the system has detected an error while trying to create a process dump. This alarm indicates an abnormal condition in which RiOS cannot collect the core file after three retries. It can be caused when the <code>/var</code> directory, which is used to hold system dumps, is reaching capacity or other conditions. When this alarm is raised, the directory is blacklisted.</p> <p>Contact Riverbed Support to correct the issue.</p>
Secure Vault	Degraded	<p>Indicates a problem with the secure vault.</p> <ul style="list-style-type: none"> <li>• <b>Secure Vault Locked</b> - Needs Attention - Indicates that the secure vault is locked.</li> <li>• <b>Secure Vault New Password Recommended</b> - Degraded - Indicates that the secure vault requires a new, non-default password. Reenter the password.</li> <li>• <b>Secure Vault Not Initialized</b> - Critical - Indicates that an error has occurred while initializing the secure vault.</li> </ul>

Alarm	Steelhead Appliance State	Reason
Software Compatibility	Needs Attention or Degraded, depending on the state	<p>Indicates that there is a mismatch between software versions in the Riverbed system.</p> <ul style="list-style-type: none"> <li>• <b>Peer Mismatch</b> - Needs Attention - Indicates that 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.</li> <li>• <b>Software Version Mismatch</b> - Degraded - Indicates that the appliance is running an incompatible version of system software. To resolve the problem, upgrade your system software. No other action is required as the alarm clears automatically.</li> </ul> <p>By default, this alarm is enabled.</p>
SSL	Needs Attention	<p>Indicates that an error has been detected in your secure vault or SSL configuration. For details about checking your settings, see <a href="#">“Verifying SSL and Secure Inner Channel Optimization” on page 173</a>.</p> <ul style="list-style-type: none"> <li>• <b>SSL Peering Certificates Expiring</b> - Indicates that an SSL certificate is about to expire.</li> <li>• <b>SSL Certificates SCEP</b> - Indicates that an SSL certificate has failed to re-enroll automatically within the SCEP polling interval</li> </ul>
System Detail Report	Degraded	<p>Indicates that the system has detected a problem with an optimization or system module. For details, see <a href="#">“Viewing System Details Reports” on page 323</a>.</p>
Temperature	Critical or Warning	<ul style="list-style-type: none"> <li>• <b>Critical</b> - Indicates that the CPU temperature has exceeded the critical threshold. The default value for the rising threshold temperature is 80° C; the default reset threshold temperature is 67° C.</li> <li>• <b>Warning</b> - Indicates that the CPU temperature is about to exceed the critical threshold.</li> </ul>

## What This Report Tells You

The Alarm Status report answers this 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.

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 all of the CPU cores used in the system within the time period specified. You can display individual cores or an overall average, or both.

For details about the report format, see [“Overview” on page 244](#).

### *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 these questions:

- How much of the CPU is being used?
- What is the average and peak percentage of the CPU being used?


## About Report Graphs

Use the mouse to hover over a specific data point to see the values and exact timestamp.

## To view the CPU Utilization report

1. Choose Reports > Diagnostics > CPU Utilization to display the CPU Utilization page.

Figure 9-30. CPU Utilization Page



2. Use the controls to customize the report as described in this table.

Control	Description
Time Interval	<p>Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days.</p> <p>Time intervals that do not apply to a particular report are dimmed.</p> <p>For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS.</p> <p>You can quickly see the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.</p>
Display Mode	<p>Select one of these displays from the drop-down menu:</p> <ul style="list-style-type: none"> <li>• <b>Brief</b> - Displays the CPU percentages for each RiOS core individually. The individual cores appear with a number and a color in the data series. To hide or display a core in the plot area, select or clear the check box next to the core name.</li> <li>• <b>Detailed</b> - Displays the CPU utilization percentage of all cores combined as a systemwide average.</li> </ul>



---

## Viewing Memory Paging Reports

The Memory Paging report provides the rate at which memory pages are swapped out to disk.

For details about the report format, see [“Overview” on page 244](#).

The Memory Page report includes this statistic that describes memory paging activity for the time period you specify.

Data Series	Description
Page Swap Out Rate	Specifies the total number of pages swapped per second. 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 <a href="https://support.riverbed.com">https://support.riverbed.com</a> .

---

### What This Report Tells You

The Memory Paging report answers this question:

- How many memory pages are swapping out?


### About Report Graphs

Use the mouse to hover over a specific data point to see the values and exact timestamp.

To view the Memory Paging report

- 1. Choose Reports > Diagnostics > Memory Paging to display the Memory Paging page.

Figure 9-31. Memory Paging Page



- 2. Use the controls to customize the report as described in this table.

Control	Description
Time Interval	Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days. Time intervals that do not apply to a particular report are dimmed. For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS. You can quickly see the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.

## Viewing TCP Memory Reports

The TCP Memory report simplifies the analysis of unexplainable throughput degradations, stalled and timed-out connections, and other network-related problems by providing the history of the TCP memory consumption and any TCP memory pressure events detected during network traffic processing. Use this report to gather preliminary information before calling Riverbed Support to troubleshoot an issue.

For details about the report format, see [“Overview” on page 244](#).

The TCP Memory report includes two graphs. The TCP usage graph provides the absolute number of memory bytes allocated by the TCP subsystem. This graph includes these statistics that describe TCP memory activity for the time period you specify.

Data Series	Description
Max Threshold	Displays the maximum amount of memory bytes that the TCP stack can allocate for its needs.
Cutoff Threshold	Displays the number of memory bytes allocated until the TCP memory allocation subsystem does not apply memory-saving mechanisms and rules. As soon as the TCP memory consumption reaches the cutoff limit, the TCP stack enters a “memory pressure” state. This state applies several important limitations that restrict memory use by incoming and transmitted packets. In practice, this means that part of the incoming packets can be discarded, and user space code is limited in its abilities to send data.
Enable Threshold	Displays the lower boundary of TCP memory consumption, when the memory pressure state is cleared and the TCP stack can use the regular memory allocation approach again.
Memory Usage	Displays the average memory consumption by the TCP/IP stack.
Memory Pressure	Displays the maximum percentage of time that the kernel has spent under TCP memory pressure.

The navigator shadows the memory usage series.

In many cases, even an insignificant increase in network traffic can cause TCP memory pressure, leading to negative consequences. There are many conditions that can cause TCP memory pressure events. However, all of them can be sorted into these two categories to identify the bottleneck in the data transfer chain:

- **Slow client cases** - Occur when the receiver (client) is not able to accept data at the rate the client-side Steelhead appliance or the server-side Steelhead appliance transfers data. This condition usually causes two TCP memory pressure points—one on the sender's side and another one on the receiver's (client's) side. The slow client on the sender's side (usually the client-side Steelhead appliance) is characterized by a large amount of unsent data collected in the send socket buffers. Incorrect Steelhead appliance settings, such as overly large send buffers, can trigger TCP memory pressure, even with relatively normal network traffic.
- **Fast server cases** - Occur when the sender is able to transfer data faster than the receiver can accept it. This condition can be triggered not only because of insufficient CPU resources, but also because of an insufficient disk transfer rate (especially with a cold and warm data pattern). The most common causes of this problem are a lack of processing power on the Steelhead appliance and a large receive buffer setting.

## What This Report Tells You

The TCP Memory report answers these questions:

- How much time is the kernel spending under TCP memory pressure?
- What is the average TCP memory consumption for the Steelhead appliance?

## About Report Graphs

Use the mouse to hover over a specific data point to see the values and exact timestamp.

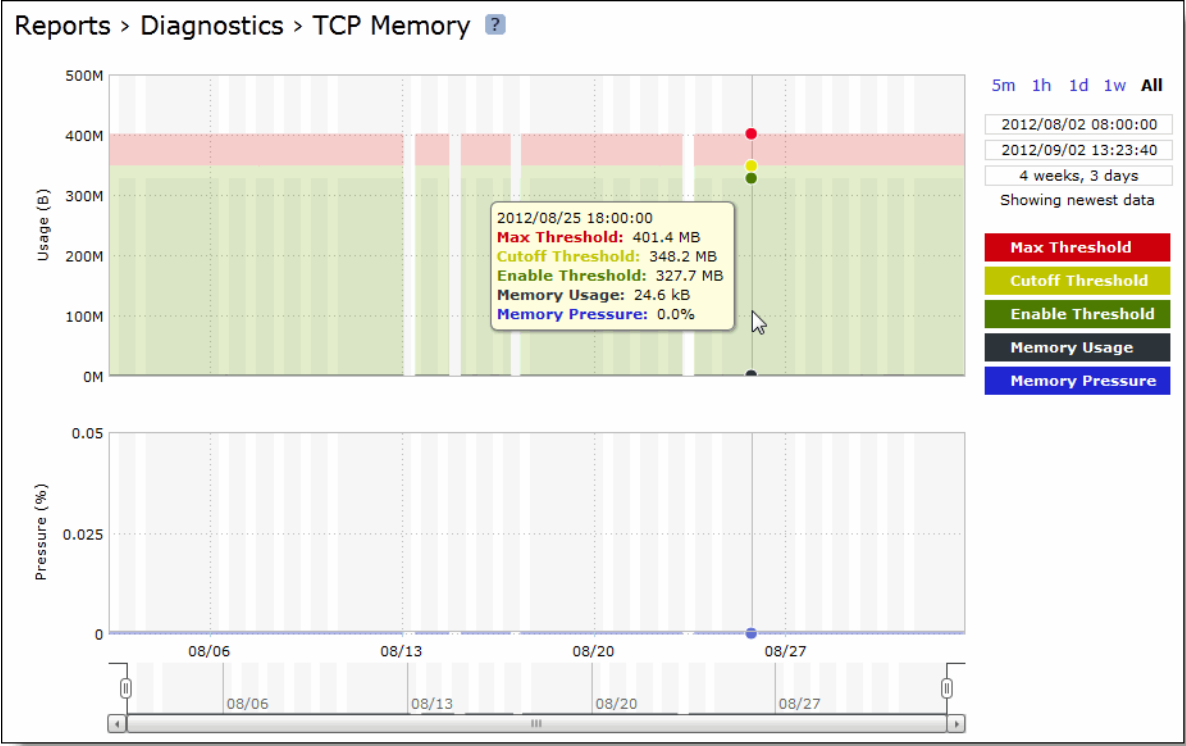
## About Report Data

The Riverbed system reports on performance for periods up to one month. Due to performance and disk space considerations, the display granularity decreases with time passed since the data was sampled with a granularity of 5 minutes for the day, 1 hour for the last week, and 2 hours for the rest of the month.

### To view the TCP Memory report

1. Choose Reports > Diagnostics > TCP Memory to display the TCP Memory page.

Figure 9-32. TCP Memory Page



2. Use the controls to customize the report as described in this table.

Control	Description
Time Interval	Select a report time interval of 5 minutes (5m), 1 hour (1h), 1 day (1d), 1 week (1w), All, or type a custom date. All includes statistics for the last 30 days. Time intervals that do not apply to a particular report are dimmed. For a custom time interval, enter the start time and end time using the format YYYY/MM/DD HH:MM:SS. You can quickly see the newest data and see data points as they are added to the chart dynamically. To display the newest data, click Show newest data.

## 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.

Column	Description
Module	<p>Displays 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.</p> <p>This report examines these modules:</p> <ul style="list-style-type: none"> <li>• <b>CPU</b> - Displays information about idle time, system time, and user time per CPU.</li> <li>• <b>Memory</b> - Displays information about the total, used, and free memory by percentage and in KBs.</li> <li>• <b>Intercept</b> - Click the right arrow to view statistics for message queue, GRE, and WCCP. Also includes table length and watchdog status.</li> <li>• <b>Secure Peering</b> - Click the right arrow and submodule name to view details for secure inner channels, including information about certificate and private key validity, peer Steelhead appliance trust, and blacklisted servers.</li> <li>• <b>Splice Policy</b> - Displays details about the splice policy in use.</li> </ul>
Status	<p>Displays one of these results:</p> <ul style="list-style-type: none"> <li>• OK (Green)</li> <li>• Warning (Yellow)</li> <li>• Error (Red)</li> <li>• Disabled (Gray). Appears when you manually disable the module.</li> </ul>

## What This Report Tells You

The System Details report answers this question:

- Is there a problem with one particular application module or does the issue affect multiple modules?

### To view the System Details report

- Choose Reports > Diagnostics > System Details to display the System Details page.

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

## Checking Network Health Status

You can run diagnostic tests on Steelhead appliance connectivity in the Reports > Diagnostics > Network Health Check page.

The network health check provides a convenient way to troubleshoot connectivity issues by running a set of general diagnostic tests. Viewing the test results can pinpoint any issues with appliance connectivity and significantly speed problem resolution.

## To run diagnostic tests

1. Choose Reports > Diagnostics > Network Health Check to display the Network Health Check page.
2. Complete the configuration as described in this table.

Control	Description
Gateway Test	<p>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.</p> <ul style="list-style-type: none"> <li>• <b>Internet Protocol</b> - Select IPv4 or IPv6 from the drop-down list.</li> <li>• <b>Run</b> - Click to run the test.</li> </ul> <p>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.</p> <p>If the test fails and only some packets are lost, check your duplex settings and other network conditions that might cause dropped packets.</p>
Cable Swap Test	<p>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).</p> <p>By default, this test is disabled.</p> <p><b>Important:</b> Certain network topologies might cause an incorrect result for this test. For the following topologies, Riverbed recommends that you confirm the test result manually:</p> <ul style="list-style-type: none"> <li>• Steelhead appliances deployed in virtual in-path mode.</li> <li>• Server-side Steelhead appliances that receive significant amounts of traffic from nonoptimized sites.</li> <li>• Steelhead appliances that sit in the path between other Steelheads that are optimizing traffic.</li> </ul> <p>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.</p>
Duplex Test	<p>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.</p> <ul style="list-style-type: none"> <li>• <b>Interface</b> - Specify an interface to test.</li> <li>• <b>IP Address</b> - Specify an IPv4 or IPv6 address that is on the testing interface side.</li> <li>• <b>Run</b> - Click to run the test.</li> </ul> <p>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.</p> <p>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.</p> <p>The test output records the percentage of any lost packets and number of collisions.</p> <p><b>Note:</b> For accurate test results, traffic must be running through the Steelhead appliance.</p>

Control	Description
Peer Reachability Test	<p>Select to send a test probe to a specified peer and await the probe response. If a response is not received, the test fails.</p> <p><b>Tip:</b> To view the current peer appliances, choose Reports &gt; Optimization &gt; Peers.</p> <ul style="list-style-type: none"> <li>• <b>IP Address</b> - Specify the IPv4 or IPv6 address of the peer appliance to test.</li> <li>• <b>Run</b> - Click to run the test.</li> </ul> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• This test might not be accurate when the peer Steelhead appliance is configured out-of-path.</li> <li>• Do not specify the primary or auxiliary IP of the same Steelhead appliance displayed in the Peers report (the primary or aux IP to which the Steelhead appliance is connected).</li> </ul> <p>If the test fails, ensure that there are no firewalls, IDS/IPS, VPNs, or other security devices which might be stripping or dropping connection packets between Steelhead appliances.</p>
IP Port Reachability Test	<p>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 <b>telnet</b>s to the port.</p> <ul style="list-style-type: none"> <li>• <b>Interface</b> - Optionally, specify an interface to test.</li> <li>• <b>IP Address</b> - Specify the IP4 or IPv6 address to test.</li> <li>• <b>Port</b> - Optionally, specify a port to test.</li> <li>• <b>Run</b> - Click to run the test.</li> </ul> <p>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.</p>
Run Selected	Runs the selected tests.
View or Hide Test Output	Click to view or hide the test results.

## Viewing the Test Status

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.

## Viewing the Test Results

The Results column displays one of these 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 > Network Health Check to display the Network Health Check page.

2. Under the test name, click **View Test Output**.

To print the test results, click **View Test Output** and 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 326](#)
- [“Viewing System Logs” on page 327](#)
- [“Downloading Log Files” on page 328](#)

## 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.
2. Use the controls to customize the log as described in this table.

Control	Description
Show	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 these options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Page</b> - Specify the number of pages you want to display.</li> <li>• <b>Time</b> - Specify the time for the log you want to display.</li> </ul>
Filter	Select one of these filtering options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Regular expression</b> - Specify a regular expression on which to filter the log.</li> <li>• <b>Error or higher</b> - Displays Error level logs or higher.</li> <li>• <b>Warning or higher</b> - Displays Warning level logs or higher.</li> <li>• <b>Notice or higher</b> - Displays Notice level logs or higher.</li> <li>• <b>Info or higher</b> - Displays Info level logs or higher.</li> </ul>
Go	Displays the report.

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 326](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

---

**Note:** If the continuous log does not appear after clicking Launch Continuous Log, 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 appliance 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.
2. Use the controls to customize the report as described in this table.

Control	Description
Show	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 these options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Page</b> - Specify the number of pages you want to display.</li> <li>• <b>Time</b> - Specify the time for the log you want to display.</li> </ul>
Regular Expression Filter	Select one of these filtering options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Regular expression</b> - Specify a regular expression on which to filter the log.</li> <li>• <b>Error or higher</b> - Displays Error level logs or higher.</li> <li>• <b>Warning or higher</b> - Displays Warning level logs or higher.</li> <li>• <b>Notice or higher</b> - Displays Notice level logs or higher.</li> <li>• <b>Info or higher</b> - Displays Info level logs or higher.</li> </ul>
Go	Displays the report.

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 327](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

---

**Note:** If the continuous log does not appear after clicking Launch Continuous Log, 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 appliance 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 interface

---

---

## 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 328](#)
- [“Downloading System Log Files” on page 329](#)

### 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.

The User Logs Download page displays up to ten archived log files plus the current day log file. By default, the system rotates each file every 24 hours or if the file size reaches one Gigabyte uncompressed. You can change this to rotate every week or month in the Configure > System Settings > Logging page. Additionally, you can rotate the files based on file size.

The automatic rotation of system logs deletes your oldest log file, labeled as Archived log #10, pushes the current log to Archived log # 1, and starts a new current-day log file.

#### To download user logs

1. Choose Reports > Diagnostics > User Logs Download to display the User Logs Download page.
2. Click the log name in the Download Plain Text column or the Download Compressed column.
3. Open or save the log (these procedures vary depending on which browser you are using).
4. Click **Rotate Logs** to manually archive the current log to a numbered archived log file and then clear the log so that it is empty again.

When you click **Rotate Logs**, your archived file #1 contains data for a partial day because you are writing a new log before the current 24-hour period is complete.

## 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.

The System Logs Download page displays up to ten archived log files plus the current day log file. By default, the system rotates each file every 24 hours or if the file size reaches one Gigabyte uncompressed. You can change this to rotate every week or month in the Configure > System Settings > Logging page. Additionally, you can rotate the files based on file size.

The automatic rotation of system logs deletes your oldest log file, labeled as Archived log #10, pushes the current log to Archived log # 1, and starts a new current-day log file.

### To download system logs

1. Choose Reports > Diagnostics > System Logs Download to display the System Logs Download page.
2. Click the log name in the Download Plain Text column or the Download Compressed column.
3. Open or save the log (these procedures vary depending on which browser you are using).
4. Click **Rotate Logs** to manually archive the current log to a numbered archived log file and then clear the log so that it is empty again.

When you click **Rotate Logs**, your archived file #1 contains data for a partial day because you are writing a new log before the current 24-hour period is complete.

---

## Generating System Dumps

You can generate, display, and download system dumps in the System Dumps 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 generate a system dump

1. Choose Reports > Diagnostics > System Dumps to display the System Dumps page.
2. Under Generate System Dump, select the type of information to include in the report:
  - **Include Statistics** - Select to collect and include CPU, memory, and other statistics in the system dump (this option is enabled by default). These statistics are useful while analyzing traffic patterns to correlate to an issue. The system adds the statistics to a file in the system dump called stats.tgz.  
  
In RiOS v8.5, you can collect and include application visibility statistics in a compressed archive file called app\_vis.db. For details, see [“Viewing Application Visibility Reports” on page 291](#).
  - **Include All Logs** - Removes the 50 MB limit for compressed log files, to include all logs in the system dump.
3. Click **Generate System Dump**.

Because generating a system dump can take a while, a spinner appears during the system dump creation. When the system dump is complete, its name appears in the list of links to download.

### To view system dump files

1. Choose Reports > Diagnostics > System Dumps to display the System Dumps page.
2. Click **Download** to view a previously saved system dump.
3. Select the filename to open a file or save the file to disk.
4. To remove a log, check the box next to the name and click **Remove Selected**.

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

### To upload a system dump file to Riverbed support

1. Choose Reports > Diagnostics > System Dumps to display the System Dumps page.
2. Select the filename.
3. Optionally, specify a case number that corresponds to the system dump. Riverbed Support recommends using a case number; for example, 194170.

You can also enter the CLI command **file debug dump upload URL** to specify a URL instead of a case number. When you specify a URL, the dump file goes directly to the URL.

If the URL points to a directory on the upload server, it must have a trailing /

For example:

ftp://ftp.riverbed.com/incoming/

(not ftp://ftp.riverbed.com/incoming)

The file name as it exists on the appliance will then match the file name on the upload server.

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

4. Click **Upload**.

Because uploading a system dump can take a while, the status appears during the upload. When the system dump finishes uploading, the date, time, and a status of either uploaded (appears in green) or failed (appears in red). An explanation appears for uploads that fail.

---

## 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.
2. Select the filename to open a file or save the file to disk.

To remove an entry, check the box next to the name and click **Remove Selected**.

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

### To download a process dump file to Riverbed support

1. Choose Reports > Diagnostics > System Dumps to display the Process Dumps page.
2. Click **Download** to receive a copy of the previously saved process dump.
3. Select the filename to open a file or save the file to disk.
4. To remove a log, check the box next to the name and click **Remove Selected**.

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

### To upload a process dump file to Riverbed support

1. Choose Reports > Diagnostics > Process Dumps to display the Process Dumps page.
2. Optionally, specify a case number that corresponds to the process dump. Riverbed Support recommends using a case number: for example, 194170.

You can also enter the CLI command **file process dump upload URL** to specify a URL instead of a case number. When you specify a URL, the dump file goes directly to the URL.

If the URL points to a directory on the upload server, it must have a trailing /.

For example:

ftp://ftp.riverbed.com/incoming/

(not ftp://ftp.riverbed.com/incoming)

The file name as it exists on the appliance will then match the file name on the upload server.

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

3. Click **Upload**.

Because uploading a process dump can take a while, a progress bar appears during the upload. When the process dump finishes uploading, the date, time, and a status of either uploaded (appears in green) or failed (appears in red). An explanation appears for uploads that fail.

---

## Capturing and Uploading TCP Dump Files

You can create, download, and upload TCP capture files in the Reports > Diagnostics > TCP Dumps page.

Capture files contain summary information for every Internet packet received or transmitted on the interface to help diagnose problems in the system.

RiOS provides an easy way to create and retrieve multiple capture files from the Management Console. You can create capture files from multiple interfaces at the same time, limit the size of the capture file, and schedule a specific date and time to create a capture file. Scheduling and limiting a capture file by time or size allows unattended captures.

RiOS v7.0 and later supports remote capture analysis using the Cascade Pilot software on capture files created and stored on the Steelhead appliance without transferring the entire packet capture across the network. The Steelhead appliance includes this functionality as Embedded Cascade Shark. Embedded Cascade Shark software enables on-demand packet capture on Steelhead appliances at remote sites, and control and analysis of packet captures on remote Steelhead appliances directly from Pilot. You can use Embedded Cascade Shark to drill down to deliver microlevel flow resolution for analysis using Riverbed's XML-based protocol on top of an HTTPS connection for transferring data to Pilot. You do not need to transfer full packets until you need them.

---

**Note:** You cannot upload a capture file to the Steelhead appliance using Pilot.

---

The top of the TCP Dumps page displays a list of existing capture files and the bottom of the page displays controls to create a capture file. The bottom of the page also includes the capture files that are currently running, and controls to create a trigger that stops a capture when a specific event occurs. The Running Capture Name list includes captures running at a particular time. It includes captures started manually and also any captures that were scheduled previously and are now running.

### To capture TCP dumps

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
2. Complete the configuration as described in this table.

Control	Description
Enable Cascade Shark	<p>Permits remote capture analysis using Cascade Pilot software. When enabled, the capture files appear in the list of completed capture files in Pilot for more detailed analysis. You do not need a separate license to use the Embedded Cascade Shark function. It is disabled by default.</p> <p>Embedded Cascade Shark uses port 61898 during a typical interaction with Pilot (for example, status requests, views configuration, and output transfer). It uses port 61899 during packet transfers.</p> <p>When Embedded Cascade Shark is disabled, Shark functionality shuts down and no processes listen on port 61898 or 61899.</p> <p>To use Embedded Cascade Shark, you must create a password for the user name Shark.</p> <p>Capture files do not appear in Pilot until they are complete.</p> <p>You can create a capture file without enabling Shark. The capture file appears in the TCP dump list in Pilot the next time you enable Shark and point Pilot to process TCP dumps from this Steelhead appliance.</p> <p>You must be able to reach the Steelhead appliance with Shark enabled from the computer running Pilot.</p> <p>For details, see the <i>Riverbed Cascade Product Suite Deployment Guide</i>.</p>
Add a New TCP Dump	Displays the controls for creating a capture file.

Control	Description
Capture Name	<p>Specify the name of the capture file. Use a unique filename to prevent overwriting an existing capture file. The default filename uses this format:</p> <p><i>hostname_interface_timestamp.cap</i></p> <p><i>hostname</i> is the hostname of the Steelhead appliance, <i>interface</i> is the name of the interface selected for the trace (for example, lan0_0, wan0_0), and <i>timestamp</i> is in the YYYY-MM-DD-HH-MM-SS format.</p> <p>If this capture file relates to an open Riverbed Support case, specify the capture filename <i>case_number</i> where <i>number</i> is your Riverbed Support case number: for example, <i>case_12345</i>.</p> <p><b>Note:</b> The .cap file extension is not included with the filename when it appears in the capture queue.</p>
Endpoints	<p>Specify IP addresses and port numbers to capture packets between them:</p> <p><b>IPs</b> - Specify IP addresses of endpoints on <i>one side</i>. Separate multiple IP addresses using commas. You can enter IPv6 addresses separated by commas. The default setting is all IP addresses.</p> <p><b>Ports</b> - Specify ports on <i>one side</i>. Separate multiple ports using commas. The default setting is all ports.</p> <p>—and—</p> <p><b>IPs</b> - Specify IP addresses of endpoints on the <i>other side</i>. Separate multiple IP addresses using commas. You can enter IPv6 addresses separated by commas. The default setting is all IP addresses.</p> <p><b>Ports</b> - Specify ports on the <i>other side</i>. Separate multiple ports using commas. The default setting is all ports.</p> <p>To capture traffic flowing in only one direction or to enter a custom command, use the CLI <b>tcpdump</b> command. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p>
Capture Interfaces	<p>Captures packet traces on the selected interfaces. You can select all interfaces or a base, in-path, or RSP interface. The default setting is none. You must specify a capture interface.</p> <p>If you select several interfaces at a time, the data is automatically placed into separate capture files.</p> <p>When path selection is enabled, Riverbed recommends that you collect packet traces on all LAN and WAN interfaces.</p>

Control	Description
Capture Parameters	<p>These parameters let you capture information about dot1q VLAN traffic. You can match traffic based on VLAN-tagged or untagged packets, or both. You can also filter by port number or host IP address and include or exclude ARP packets. Select one of these parameters for capturing VLAN packets:</p> <ul style="list-style-type: none"> <li>• <b>Capture Untagged Traffic Only</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– All untagged VLAN traffic.</li> <li>– Untagged 7850 traffic and ARP packets. You must also specify <b>or arp</b> in the custom flags field on this page.</li> <li>– Only untagged ARP packets. You must also specify <b>and arp</b> in the custom flags field on this page.</li> </ul> </li> <li>• <b>Capture VLAN-Tagged Traffic Only</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– Only VLAN-tagged traffic.</li> <li>– VLAN-tagged packets with host 10.11.0.6 traffic and ARP packets. You must also specify <b>10.11.0.6</b> in the IPs field, and specify <b>or arp</b> in the custom flags field on this page.</li> <li>– VLAN-tagged ARP packets only. You must also specify <b>and arp</b> in the custom flags field on this page.</li> </ul> </li> <li>• <b>Capture both VLAN and Untagged Traffic</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– All VLAN traffic.</li> <li>– Both tagged and untagged 7850 traffic and ARP packets. You must also specify the following in the custom flags field on this page: (port 7850 or arp) or (vlan and (port 7850 or arp))</li> <li>– Both tagged and untagged 7850 traffic only. You must also specify <b>7850</b> in one of the port fields on this page. No custom flags are required.</li> <li>– Both tagged and untagged ARP packets. You must also specify the following in the custom flags field on this page: (arp) or (vlan and arp)</li> </ul> </li> </ul>
Capture Duration (Seconds)	<p>Specify a positive integer to set how long the capture runs, in seconds. The default value is 30. Specify 0 or continuous to initiate a continuous trace.</p> <p>For continuous capture, Riverbed recommends specifying a maximum capture size and a non-zero rotate file number to limit the size of the TCP dump.</p>
Maximum Capture Size	<p>Specify the maximum capture file size, in MB. The default value is 100. After the file reaches the maximum capture size, TCP dump starts writing capture data into the next file, limited by the Number of Files to Rotate field.</p> <p>Riverbed recommends a maximum capture file size of 1024 MB (1 GB).</p>
Buffer Size	<p>Optionally, specify the maximum amount of data, in kB, allowed to queue while awaiting processing by the capture file. The default value is 154 kB.</p>
Snap Length	<p>Optionally, specify the snap length value for the capture file, which equals the number of bytes captured for each packet. Having a snap length smaller than the maximum packet size on the network enables you to store more packets, but you might not be able to inspect the full packet content. Specify 0 for a full packet capture. The default value is 1518 bytes.</p>



Control	Description
Number of Files to Rotate	<p>Specify how many capture files to keep for each interface before overwriting the oldest file. To stop file rotation, you can specify 0; however, Riverbed recommends rotating files, because stopping the rotation can fill the disk partition.</p> <p>This limits the number of files created to the specified number, and begins overwriting files from the beginning, thus creating a rotating buffer.</p> <p>The default value is five files per interface. The maximum value is a 32-bit integer.</p>
Custom Flags	<p>Specify custom flags as additional statements within the filter expression. Custom flags are added to the end of the expression created from the Endpoints fields and the Capture Parameters radio buttons (pertaining to VLANs).</p> <p>If you require an “and” statement between the expression created from other fields and the expression that you are entering in the custom flags field, you must include the “and” statement at the start of the custom flags field.</p> <p>Do not use host, src, or dst statements in the custom flags field. Although it is possible in trivial cases to get these to start without a syntax error, they do not capture GRE-encapsulated packets that some modes of Steelhead appliance communications use, such as WCCP deployments or Interceptor connection-setup traffic. Riverbed recommends using bidirectional filters by specifying endpoints.</p> <p>For complete control of your filter expression, use the CLI <b>tcpdump</b> command. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p> <p>For examples, see <a href="#">“Custom Flag Use Examples” on page 335</a>.</p>
Schedule Dump	Schedules the capture to run at a later date and time.
Start Date	Specify a date to initiate the capture, in this format: YYYY/MM/DD.
Start Time	Specify a time to initiate the capture, in this format: HH:MM:SS.
Add	Adds the capture request to the capture queue.

## Troubleshooting

If your command results in a syntax error with an immediate or scheduled TCP dump, this message appears:

“Error in tcpdump command. See System Log for details.”

Review the system log to see the full tcpdump command attempt. Check the expression for issues such as a missing “and,” as well as contradictory instructions such as looking for VLAN-tagged traffic AND non-tagged traffic.

## Custom Flag Use Examples

The examples in this table focus on the custom flag entry but rely on other fields to create a complete filter.

Filter Purpose	Custom Flag
To capture all traffic on VLAN 10 between two specified endpoints: 1.1.1.1 and 2.2.2.2	and vlan 10
To capture any packet with a SYN or an ACK	tcp[tcpflags] & (tcp-syn   tcp-ack) != 0

Filter Purpose	Custom Flag
To capture any packet with a SYN	tcp[tcpflags] & (tcp-syn) != 0 —or— tcp[13] & 2 == 2
To capture any SYN to or from host 1.1.1.1	and (tcp[tcpflags] & (tcp-syn) != 0) —or— and (tcp[13] & 2 == 2)

## IPv6 Custom Flag Use Examples

The examples in this table focus on the custom flag entry, but rely on other fields to create a complete filter.

To build expressions for TCP dump, IPv6 filtering does not currently support the TCP, UDP, and other upper-layer protocol types that IPv4 does. Also, these IPv6 examples are based on the assumption that only a single IPv6 header is present.

Filter Purpose	Custom Flag
To capture all FIN packets to or from host 2001::2002	and (ip6[53] & 1!=0)
To capture all IPv6 SYN packets	ip6 or proto ipv6 and (ip6[53] & 2 == 2)

## Stopping a TCP Dump After an Event Occurs

Capture files offer visibility into intermittent network issues, but the amount of traffic they capture can be overwhelming. Also, because rotating logs is common, after a capture logs an event, the Steelhead appliance log rotation can overwrite debugging information specific to the event.

RiOS v8.5.x makes troubleshooting easier because it provides a trigger that can stop a continuous capture after a specific log event occurs. The result is a smaller file to help pinpoint what makes the event happen.

The stop trigger continuously scans the system logs for a search pattern. When it finds a match, it stops all running captures.

### To stop a capture after a specific log event

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
2. Schedule a capture.
3. In the Pattern text box, enter a Perl regular expression (regex) to find in a log. RiOS compares the Perl regex against each new line in the system logs and the trigger stops if it finds a match.

The simplest regex is a word or a string of characters. For example, if you set the pattern to "Limit," the trigger matches the line "Connection Limit Reached."

Notes:

- Perl regular expressions are case-sensitive.
- Perl treats the space character " " like any other character in a regex.

- Perl reserves some characters, called metacharacters, for use in regex notation. The metacharacters are:

`{ } [ ] ( ) ^ $ . | * + ? \`

You can match a metacharacter by putting a backslash before it. For example, to search for a backslash in the logs, you must enter two backslashes (`\\`) as the pattern.

- The pattern follows Perl regular expression syntax. For details, go to:  
<http://perldoc.perl.org/perlre.html>
- You cannot change the pattern while a scan is running. You must stop the scan before changing a pattern.
- You do not need to wrap the pattern with the metacharacters to match the beginning or end of a line (`^` `$`) or with the wildcard character (`*`).

4. Specify the amount of time to pause before stopping all running captures when RiOS finds a match. This gives the system some time to log more data without abruptly cutting off the capture. The default is 30 seconds. Specify 0 for no delay; the capture stops immediately.

After a trigger has fired, the capture can stop by itself before the delay expires; for example, the capture duration can expire.

5. Click **Start Scan**.

When the scan stops, RiOS sends an email to all email addresses on the Configure > System Settings > Email page appearing under Report Events via Email. The email notifies users that the trigger has fired.

The page indicates “Last Triggered: Never” if a TCP Dump stop trigger has never triggered on the Steelhead appliance. After the delay duration of the stop trigger, RiOS displays the last triggered time.

Before changing the Perl regular expression or amount of delay, you must first stop the process.

### To stop a running scan

- Click **Stop Scan** to halt the background process that monitors the system logs. RiOS dims this button when the stop trigger is idling.

## Stop Trigger Limitations

These limitations apply to the trigger:

- You cannot create a trigger to stop a specific capture; the trigger affects all running captures.
- If the search pattern contains a typo, the trigger might never find a match.
- Only one instance of a trigger can run at one time.

## Viewing a TCP Dump

The top of the TCP Dumps page displays a list of existing captures.

### To view a capture file

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
2. Under Stored TCP Dumps, select the capture name to open the file.

3. Click **Download** to view a previously saved capture file.
4. To remove a capture file, select the check box next to the name and click **Remove Selected**.

### To print a capture file

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
2. Under Download Link, select the capture filename to open the file.
3. When the file opens, choose File > Print in your Web browser to open the Print dialog box.

### To stop a running capture

1. Choose Reports > Diagnostics > TCP Dumps to display the TCP Dumps page.
2. Select the capture filename in the Running Capture Name list.
3. Click **Stop Selected Captures**.

## Uploading a TCP Dump

Riverbed offers a couple of ways to upload capture files to the support server for sharing with the support team while diagnosing issues.

### To upload the capture file to Riverbed Support

1. In continuous mode, on the TCP Dumps page, select the running capture and click **Stop Selected Captures**.

For timed captures that are complete, skip to step 2.

The capture appears as a download link in the list of Stored TCP Dumps.

2. Select the capture filename.
3. Optionally, specify a case number that corresponds to the capture. Riverbed Support recommends using a case number: for example, 194170.

To specify a URL instead of a case number, you must use the CLI. You can enter the CLI command **file tcpdump upload URL**. When you specify a URL, the capture file goes directly to the URL.

If the URL points to a directory on the upload server, it must have a trailing backslash (/).

For example:

`ftp://ftp.riverbed.com/incoming/`

(not `ftp://ftp.riverbed.com/incoming`)

The filename as it exists on the appliance will then match the filename on the upload server.

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

4. Click **Upload**.

Because uploading a capture file can take a while, a progress bar displays the percentage of the total upload completed, the case number (if applicable), and the date and time the upload began. When the capture file finishes uploading, the date, time, and a status of either uploaded (appears in green) or failed (appears in red).

Successful uploads show the status, the case number (if applicable), and the date and time the upload finished.

For uploads that fail, an explanation, the case number (if applicable), and the upload starting date and time appear.

---

## Exporting Performance Statistics

You export performance statistics in CSV format in the Reports > Export Report Data page. 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 Report Data to display the Export Report Data page.
2. Use the controls to customize the report, as described in this table.

Control	Description
Report	Select the type of report you want to export from the drop-down list.
Period	Select a report time interval of custom, last five minutes, last hour, last day, last week, or last month.
Email Delivery	Sends the report to an email address.
Email Address	Specify the email address of the recipient.
Export	Exports the report data.



## APPENDIX A 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:

- [“Accessing the Steelhead Appliance Enterprise MIB” on page 341](#)
- [“SNMP Traps” on page 342](#)

---

### Accessing the Steelhead Appliance Enterprise MIB

The Steelhead appliance enterprise MIB monitors device status and peers. It 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 product documentation.

The following guidelines describe how to download and access the Steelhead appliance enterprise MIB using common MIB browsing utilities:

- You can download the Steelhead appliance enterprise MIB file (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 extension to the type required by the utility you have chosen.
- 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 appliance MIB object bwPortTable to display a table of statistics for optimized traffic by port, the command retrieves only the monitored ports. The monitored ports include the default TCP ports and any ports you add. To view the monitored ports that 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, as in the following 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 triggers 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 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 includes support for SNMP v3.

You can view Steelhead appliance health at the top of each Management Console page, 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 that you can quickly discover and diagnose issues. The health of an appliance falls into one of the following states:

- **Healthy** - The Steelhead appliance is functioning and optimizing traffic.
- **Needs Attention** - Accompanies a healthy state to indicate management-related issues not affecting the ability of the Steelhead appliance to optimize traffic.
- **Degraded** - The Steelhead appliance is optimizing traffic but the system has detected an issue.
- **Admission Control** - The Steelhead appliance is optimizing traffic but has reached its connection limit.
- **Critical** - The Steelhead appliance might or might not be optimizing traffic; you must address a critical issue.



The following table summarizes the SNMP traps sent from the system to configured trap receivers and their effect on the Steelhead appliance health state.

Trap and OID	Steelhead Appliance State	Text	Description
procCrash (enterprises.17163.1.1.4.0.1)	Healthy	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)	Healthy	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)	N/A	This alarm is deprecated.	N/A

Trap and OID	Steelhead Appliance 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)	Deprecated	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.
storeCorruption (enterprises.17163.1.1.4.0.9)	Critical	The data store is corrupted.	Indicates that the RiOS data store is corrupt or has become incompatible with the current configuration.  If the alarm was triggered by an unintended change to the configuration, change the configuration to match the previous RiOS data store settings. Then restart the optimization service without clearing the data store to reset the alarm.  Typical configuration changes that require an optimization restart with a clear RiOS data store are enabling enhanced peering or changing the data store encryption.

Trap and OID	Steelhead Appliance 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 a system dump 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)	Healthy	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)	Healthy	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)	Healthy	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 Appliance State	Text	Description
linkError (enterprises.17163.1.1.4.0.17)	Degraded	An interface on the appliance has lost its link.	<p>The system has lost one of its Ethernet links, typically due to an unplugged cable or dead switch port. 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.</p> <p>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.</p>
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 as 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).	<p>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:</p> <pre>show hardware error-log all</pre>
configChange (enterprises.17163.1.1.4.0.24)	Healthy	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.

Trap and OID	Steelhead Appliance State	Text	Description
datastoreWrapped (enterprises.17163.1.1.4.0.25)	Healthy	The datastore has wrapped around.	The RiOS data store 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 that the RiOS data store is undersized. If a message is received every seven days or less, investigate traffic patterns and RiOS data store 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 Appliance 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.
sslPeeringSCEPAutoReenrollError (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 RiOS data store synchronization between two Steelhead appliances has been disrupted and the RiOS data stores 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 RiOS data store cannot be encrypted. Check the Alarm Status page for more details. The alarm clears when the secure vault is unlocked.

Trap and OID	Steelhead Appliance State	Text	Description
secureVaultNeedsRekey (enterprises.17163.1.1.4.0.39)	Needs Attention	If you wish to use a nondefault password for the secure vault, the password must be rekeyed. Please see the Knowledge Base solution 5592 for more details.	<p>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.</p> <p>For details, check the Alarm Status page and see the Knowledge Base solution 5592.</p> <p>The alarm clears when you verify the default password or reset the password.</p>
secureVaultInitError (enterprises.17163.1.1.4.0.40)	Critical	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)	Healthy	The current appliance configuration has been saved.	<p>A configuration has been saved either by entering the <code>write mem</code></p> <p>CLI command or by clicking <b>Save</b> in the Management Console. This message is for security notification purposes only; no other action is necessary.</p>
tcpDumpStarted (enterprises.17163.1.1.4.0.42)	Healthy	A TCP dump has been started.	<p>A user has started a TCP dump on the Steelhead appliance by entering a <code>tcpdump</code></p> <p>or</p> <p><code>tcpdump -x</code></p> <p>command from the CLI. This message is for security notification purposes only; no other action is necessary.</p>
tcpDumpScheduled (enterprises.17163.1.1.4.0.43)	Healthy	A TCP dump has been scheduled.	<p>A user has started a TCP dump on the Steelhead appliance by entering a <code>tcpdump</code></p> <p>or</p> <p><code>tcpdump -x</code></p> <p>command with a scheduled start time from the CLI. This message is for security notification purposes only; no other action is necessary.</p>
newUserCreated (enterprises.17163.1.1.4.0.44)	Healthy	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.

Trap and OID	Steelhead Appliance State	Text	Description
diskError (enterprises.17163.1.1.4.0.45)	Degraded	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)	Degraded	Accumulated SSD write cycles passed predefined level.	Triggers on Steelhead appliance models using Solid State Disks (SSDs). An SSD has reached 95 percent of its write cycle limit. Contact Riverbed Support.
cliUserLogin (enterprises.17163.1.1.4.0.47)	Healthy	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.
cliUserLogout (enterprises.17163.1.1.4.0.48)	Healthy	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)	Healthy	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)	Healthy	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)	Healthy	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.



Trap and OID	Steelhead Appliance State	Text	Description
systemDiskFullError (enterprises.17163.1.1.4.0.54)	Degraded	One or more system partitions is full or almost full.	The alarm clears when the system partitions fall below usage thresholds.
domainJoinError (enterprises.17163.1.1.4.0.55)	Degraded	An attempt to join a domain failed.	<p>An attempt to join a Windows domain has failed.</p> <p>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, this error message appears:</p> <pre>lt-kinit: krb5_get_init_creds: Clock skew too great</pre> <p>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 the time 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.</p> <p>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, these error messages appear:</p> <pre>Failed to join domain: failed to find DC for domain &lt;domain name&gt; Failed to join domain : No Logon Servers</pre> <p>Additionally, the domain join alarm triggers and messages similar to the following appear in the logs:</p> <pre>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</pre> <p>When you encounter this error, go to the <b>Configure &gt; Networking &gt; Host Settings</b> page and verify that the DNS settings are correct.</p> <p>To verify the time settings, go to the <b>Configure &gt; System Settings &gt; Date and Time</b> page.</p>

Trap and OID	Steelhead Appliance State	Text	Description
certsExpiringError (enterprises.17163.1.1.4.0.56)	Needs Attention	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.
licenseError (enterprises.17163.1.1.4.0.57)	Critical	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)	Either Critical or Degraded, depending on the state	Hardware error detected.	<p>Indicates that the system has detected a problem with the Steelhead appliance hardware. These issues trigger the hardware error alarm:</p> <ul style="list-style-type: none"> <li>the Steelhead appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration</li> <li>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</li> <li>other hardware issues</li> </ul> <p>The alarm clears when you add the necessary hardware, remove the unqualified hardware, or resolve other hardware issues.</p>
sysdetailError (enterprises.17163.1.1.4.0.59)	Needs Attention	Error is found in System Detail Report.	A top-level module on the system detail report is in error. For details, choose Reports > Diagnostics > System Details.
neighborIncompatibility (enterprises.17163.1.1.4.0.61)	Degraded	Serial cascade misconfiguration has been detected.	Check your automatic peering configuration. Restart the optimization service to clear the alarm.
flashError (enterprises.17163.1.1.4.0.62)	Needs Attention	Flash hardware error detected.	<p>At times, the USB flash drive that holds the system images might become unresponsive; the Steelhead appliance continues to function normally. When this alarm triggers, you cannot perform a software upgrade, as the system is unable to write a new upgrade image to the flash drive without first power cycling the system.</p> <p>To reboot the appliance, go to the Configure &gt; Maintenance &gt; Reboot/ Shut Down page or enter the CLI <b>reload</b> command to automatically power cycle the Steelhead appliance and restore the flash drive to its proper function.</p>

Trap and OID	Steelhead Appliance State	Text	Description
lanWanLoopError (enterprises.17163.1.1.4.0.63)	Critical	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.

Trap and OID	Steelhead Appliance State	Text	Description
optimizationServiceStatusError (enterprises.17163.1.1.4.0.64)	Critical	Optimization service currently not optimizing any connections.	<p>The optimization service has encountered an optimization service condition. The message indicates the reason for the condition:</p> <ul style="list-style-type: none"> <li>optimization service is not running This message appears after a configuration file error. For more information, review the Steelhead appliance logs.</li> <li>in-path optimization is not enabled This message appears if an in-path setting is disabled for an in-path Steelhead appliance. For more information, review the Steelhead appliance logs.</li> <li>optimization service is initializing This message appears after a reboot. The alarm clears on its own; no other action is necessary. For more information, review the Steelhead appliance logs.</li> <li>optimization service is not optimizing This message appears after a system crash. For more information, review the Steelhead appliance logs.</li> <li>optimization service is disabled by user This message appears after entering the CLI command <b>no service enable</b> or shutting down the optimization service from the Management Console. For more information, review the Steelhead appliance logs.</li> <li>optimization service is restarted by user This message appears after the optimization service is restarted from either the CLI or Management Console. You might want to review the Steelhead appliance logs for more information.</li> </ul>

Trap and OID	Steelhead Appliance State	Text	Description
upgradeFailure (enterprises.17163.1.1.4.0.65)	Needs attention	Upgrade failed and the system is running the previous image.	<p>A RiOS upgrade has failed and the Steelhead appliance is running the previous RiOS version. Check the banner message in the Management Console to view more information. The banner message displays which upgrade failed along with the RiOS version the Steelhead appliance has reverted to and is currently running.</p> <p>Check that the upgrade image is correct for your Steelhead appliance.</p> <p>Verify that the upgrade image is not corrupt. You can use the MD5 checksum tool provided on the Riverbed Support site for the verification.</p> <p>After you have confirmed that the image is not corrupt, try to upgrade the RiOS software again. If the upgrade continues to fail, contact Riverbed Support.</p>
licenseExpiring (enterprises.17163.1.1.4.0.66)	Needs Attention	One or more licensed features will expire within the next two weeks.	<p>Choose Configure &gt; Maintenance &gt; Licenses and look at the Status column to see which licenses are about to expire. One or more feature licenses are scheduled to expire within two weeks.</p> <p>This alarm is triggered per feature. Suppose you installed two license keys for a feature, LK1-FOO-xxx, which is going to expire in two weeks, and LK1-FOO-yyy, which is not expired. Because one license for the feature is valid, the alarm does not trigger.</p>
licenseExpired (enterprises.17163.1.1.4.0.67)	Degraded	One or more licensed features have expired.	<p>Choose Configure &gt; Maintenance &gt; Licenses and look at the Status column to see which licenses have expired. One or more feature licenses have expired.</p> <p>This alarm is triggered per feature. Suppose you installed two license keys for a feature, LK1-FOO-xxx (expired), and LK1-FOO-yyy (not expired). Because one license for the feature is valid, the alarm does not trigger.</p>
clusterDisconnectedSHAlertError (enterprises.17163.1.1.4.0.68)	Degraded	A cluster Steelhead has been reported as disconnected.	<p>Choose Configure &gt; Networking &gt; Connection Forwarding and verify the configuration for both this Steelhead appliance and the neighbor Steelhead appliance. Verify that the neighbor is reachable from this Steelhead appliance.</p> <p>Next, check that the optimization service is running on both Steelhead appliances.</p> <p>This error clears when the configuration is valid.</p>

Trap and OID	Steelhead Appliance State	Text	Description
linkDuplex (enterprises.17163.1.1.4.0.70)	Degraded	An interface on the appliance is in half-duplex mode	<p>Indicates that an interface was not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>Choose Configure &gt; Networking &gt; Base Interfaces and examine the Steelhead appliance link configuration. Next, examine the peer switch user interface to check its link configuration. If the configuration on one side is different from the other, traffic is sent at different rates on each side, causing many collisions.</p> <p>To troubleshoot, change both interfaces to automatic duplex negotiation. If the interfaces do not support automatic duplex, configure both ends for full duplex.</p>
linkIoErrors (enterprises.17163.1.1.4.0.71)	Degraded	An interface on the appliance is suffering I/O errors	<p>Indicates that the error rate on an interface has exceeded 0.1 percent while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured LAN connection should experience very few errors. The alarm clears when the error rate drops below 0.05 percent.</p> <p>To troubleshoot, try a new cable and a different switch port. Another possible cause is electromagnetic noise nearby.</p> <p>You can change the default alarm thresholds by entering the <b>alarm link_errors err-threshold xxxxxx</b> CLI command at the system prompt. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p>

Trap and OID	Steelhead Appliance State	Text	Description
clusterIpv6IncompatiblePeerError (enterprises.17163.1.1.4.0.74)	Degraded	A cluster Steelhead has been reported as IPv6 incompatible.	<p>The optimization service has encountered a peer Steelhead appliance IPv6 incompatibility. The message indicates the reason for the condition:</p> <ul style="list-style-type: none"> <li>• Not all local inpath interfaces configured for IPv6</li> </ul> <p>This message indicates that the peer Steelhead appliance is IPv6 capable and its IP address configuration is correct, but the IP address configuration on the local Steelhead appliance does not match the configuration on the peer Steelhead appliance. The mismatch means that there is at least one relay on the local appliance that is not IPv4 or IPv6 capable. An IPv4 address is necessary for routing between neighbors and an IPv6 address is necessary for v6 optimization.</p> <li>• Not all peer inpath interfaces configured for IPv6</li> <p>This message indicates that the local Steelhead appliance is IPv6 capable and its IP address configuration is correct, but the IP address configuration on the peer Steelhead appliance does not match the configuration on the local Steelhead appliance. The mismatch means that there is at least one relay on the peer that is not IPv4 or IPv6 capable. An IPv4 address is necessary for routing between neighbors and an IPv6 address is necessary for v6 optimization.</p> <li>• Cluster IPv6 Incompatible</li> <p>Indicates that a connection-forwarding neighbor is running a RiOS version that is incompatible with IPv6. The Steelhead appliance neighbors pass through IPv6 connections when this alarm triggers.</p>

Trap and OID	Steelhead Appliance State	Text	Description
flashProtectionFailed (enterprises.17163.1.1.4.0.75)	Critical	Flash disk hasn't been backed up due to not enough free space on /var filesystem.	<p>Indicates that the USB flash drive has not been backed up because there is not enough available space in the /var filesystem directory.</p> <p>Examine the /var directory to see if it is storing an excessive amount of snapshots, system dumps, or TCP dumps that you could delete. You could also delete any RiOS images that you no longer use.</p>
pathSelectionPathDown (enterprises.17163.1.1.4.0.77)	Degraded	Path Selection - A path went down.	<p>Indicates that one of the predefined paths for a connection is unavailable because it has exceeded either the timeout value for path latency or the threshold for observed packet loss.</p> <p>When a path fails, the Steelhead appliance directs traffic through another available path. When the original path comes back up, the Steelhead appliance redirects the traffic back to it.</p>



## APPENDIX B Steelhead Appliance Ports

This appendix provides a reference to ports used by the system. It includes these sections:

- [“Granite Ports” on page 359](#)
- [“Commonly Excluded Ports” on page 360](#)
- [“Interactive Ports Forwarded by the Steelhead Appliance” on page 360](#)
- [“Secure Ports Forwarded by the Steelhead Appliance” on page 361](#)

---

### Granite Ports

This table lists and describes the Steelhead EX Granite default ports with the port label Granite.

Default Ports	Description
7950	Data requests for data blocks absent in Edge appliance from the data center
7951	New data created at the Edge to the data center
7952	Prefetch data for which Granite has highest confidence (for example, file read ahead)
7953	Prefetch data for which Granite has medium confidence (for example, boot)
7954	Prefetch data for which Granite has lowest confidence (for example, prepopulation)
7970	Management information exchange between Edge and Core appliances

---

## Default Ports

This table summarizes Steelhead appliance default ports with the port label: RBT-Proto.

Default Ports	Description
7744	RiOS data store 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 appliance-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 automatic 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, delete the Interactive rule in the Management Console.

---

This 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
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, delete the Secure rule in the Management Console.

This 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+
kerberos	88	Kerberos
rtsp	322	rtsp over TLS/SSL

Type	Port	Description
https	443/tcp	http protocol over TLS/SSL
smtps	465/tcp	# SMTP over SSL (TLS)
nntp	563/tcp	nntp protocol over TLS/SSL (was snntp)
imap4-ssl	585/tcp	IMAP4+SSL (use 993 instead)
ssh	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
telnet	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
tftp	3713/tcp	TFTP over TLS
operations manager	5723	Microsoft Operations Manager

This table contains the uncommon ports automatically forwarded by the Steelhead appliance.

Type	Port	Description
nsiops	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
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

Type	Port	Description
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
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 C Application Signatures for AFE

This appendix provides a reference to the application signatures recognized by the Application Flow Engine (AFE). It includes the following section:

- [“List of Recognized Applications” on page 365](#)

---

### List of Recognized Applications

AFE recognizes over 700 application signatures. These application signatures provide an efficient and accurate way to identify applications for advanced classification of network traffic in QoS.

You can verify the application signatures available in your specific RiOS version from within the Management Console. Type the first few letters of the application in the Application Protocol or Application field for QoS configuration. As you type the name of an application, a menu appears and lists available applications that match your typing.

These tables list and describe application signatures recognized by the Application Flow Engine. The tables are organized by application type.

Collaboration Applications	Description	First Available In
Citrix Jedi	An online streaming connection protocol for streaming real-time data.	v8.0
Citrix Online	An online service that includes GoToMyPC, GoToMeeting, GoToWebinar, and GoToTraining.	v8.0
GoToMeeting	A remote meeting and desktop sharing software that enables the user to meet with other computer users, customers, clients, or colleagues via the Internet in real-time.	v6.5
Groupwise	A messaging and collaborative software platform from Novell that supports email, calendaring, personal information management, instant messaging, and document management.	v8.0
HL7	A medical information exchange standard for exchanging information between medical applications.	v8.0
Livemeeting	A Microsoft commercial web-conferencing service.	v6.5
Microsoft Lync	A Microsoft voice, video, file transfer, and video sharing communications platform.	v8.5.1
Notes	An IBM enterprise collaboration suite, Lotus Notes.	v8.0

Collaboration Applications	Description	First Available In
Meeting Maker	A cross-platform personal calendar and group scheduling software application from PeopleCube.	v8.0
NetMeeting	A VoIP and multi-point video conferencing client included in many versions of Microsoft Windows.	v8.0
SharePoint	A Microsoft collaboration, file sharing, and web publishing system.	v6.5
WebEx	A Cisco online meeting and web-conferencing application.	v6.5

Database Applications	Description	First Available In
BLIDM	A Britton-Lee integrated database manager.	v8.0
dBase	The first widely used database management system (DBMS) for microcomputers. A major upgrade was released as dBASE III, and ported to a wider variety of platforms, adding UNIX and VMS.	v8.0
DEOS	A distributed external object port 76 for TCP and UDP	v8.0
INGRES-NET	An IngresNET service.	v8.0
LDAP	A protocol for reading and editing directories over an IP network.	v6.5
Mini SQL	A lightweight database management system.	v8.0
MS OLAP	An online analytical processing capability that is a component of Microsoft SQL Server.	v8.0
MS SQL	A relational database server produced by Microsoft.	v8.0
MySQL	A relational database management system (RDBMS) that runs as a server, providing multi-user access to a number of databases.	v6.5
Oracle	An object-relational database management system (ORDBMS) produced and marketed by Oracle Corporation.	v6.5
Oracle SQLNET	The networking software that enables remote data access between programs and the Oracle Database, or among multiple Oracle databases.	v8.0
PostgreSQL	An open source object-relational database system.	v8.0
RIS	The relational interface system (RIS) which is Intergraph Corporation's middleware for connecting client software and DBMS.	v8.0
SQL Services	A protocol for service type SQL Services registered with IANA on port 118 TCP/UDP.	v8.0

Email Applications	Description	First Available In
126.com	A free Web mail service of Netease.	v8.5
Exchange	The Microsoft Exchange email, scheduling, and contact services.	v6.5
Facebook-Messages	The Facebook email and instant messaging service.	v8.5
Gmail	The Google hosted Web mail service that allows private access to email, file storage, and instant messaging.	v6.5



Email Applications	Description	First Available In
Hotmail	A free Web-based email service by Microsoft as part of their Windows Live range of online services. It features unlimited storage and can be accessed via a Web browser. Formerly known as MSN Hotmail and now commonly referred to simply as Hotmail.	v6.5
IMAP	An Internet standard protocol for accessing email on a remote server.	v8.0
Infostore	A Microsoft Exchange information store.	v8.0
MAILQ	A protocol for service type MAILQ registered with IANA on port 174 TCP/UDP	v8.0
MAPI	The protocol that Microsoft Outlook uses to communicate with Microsoft Exchange.	v6.5
MTA	The Microsoft Exchange Mail Transfer Agent.	v8.0
NI Mail	A protocol for service type NI MAIL registered with IANA on port 61 TCP/UDP.	v8.0
PCMAIL	A protocol for service type pmail-srv registered with IANA on port 158 TCP/UDP.	v8.0
POP2	A protocol used by local email clients to retrieve email from a remote server.	v8.0
POP3	A protocol used by local email clients to retrieve email from a remote server.	v6.5
QMTP	An email transmission protocol that is designed to have better performance than Simple Mail Transfer Protocol (SMTP), the de facto standard.	v8.0
REMAIL	A protocol for service type re-mail-ck registered with IANA on port 50 TCP/UDP.	v8.0
RFRI	The Microsoft Exchange referral interface.	v8.0
SMTP	An Internet standard for email transmission across Internet Protocol (IP) networks.	v6.5
Store Admin	An Exchange Store server administration tool.	v8.0
Sysatt	The Microsoft Exchange system attendant service.	v8.0
XNS Mail	The Xerox networking services mail.	v8.0

File Transfer Applications	Description	First Available In
4Shared	A file-sharing service that provides search functions and enables users to upload and download files to their accounts and share links with other people.	v8.0
ACR-NEMA	A standard for handling, storing, printing, and transmitting information in medical imaging.	v8.0

File Transfer Applications	Description	First Available In
AFP	A network protocol that offers file services for Mac OS X and original Mac OS. In Mac OS X, AFP is one of several file services supported including Server Message Block (SMB), Network File System (NFS), File Transfer Protocol (FTP), and WebDAV. AFP currently supports Unicode file names, POSIX and access control list permissions, resource forks, named extended attributes, and advanced file locking.	v8.0
Apple Update	A software tool by Apple Computer that installs the latest version of Apple software.	v8.0
AppleJuice	A semi-centralized peer-to-peer file sharing network similar to the original eDonkey network.	v8.0
AppleJuice GUI	An AppleJuice host running a GUI that represents traffic between itself and a host running the AppleJuice Core.	v8.0
Ares	An open-source peer-to-peer file sharing application.	v8.5
Astraweb	A Usenet/newsgroup service provider.	v6.5
Auditd	The userspace daemon to the Linux auditing system, which is responsible for writing audit records to the disk.	v8.0
AVG	A free downloadable antivirus software solution made by AVG Technologies.	v8.0
Avira	A free, downloadable antivirus program that is part of the Avira security product suite.	v8.0
BackBlaze	An online backup tool that allows Windows and Mac OS X users to back up their data to an offsite data center.	v8.0
BFTP	A protocol for service type background file transfer program registered with IANA on port 152 TCP/UDP.	v8.0
BigUpload	A secure uploading, transferring, and file sharing service.	v8.5
BitDefender	An antivirus software solution for varying levels of AV protection.	v8.0
BITS	A file transfer protocol used primarily for Microsoft updates.	v6.5
BitTorrent	A peer-to-peer file sharing protocol used for transferring large amounts of data.	v6.5
BlazeFS	A remote file sharing system designed specifically for the Mac OS. Once running, Blaze is transparent to the user and to the client application. It appears as if users are accessing files on a local hard drive.	v8.0
Boxnet	An online file sharing and storage Web site.	v8.5
CFDPTKT	A protocol for service type CFDPTKT registered with IANA on port 120 TCP/UDP.	v8.0
CIFS	The common internet file system used to provide shared access to directories, files, printers, serial ports, and miscellaneous communication devices between nodes on a network.	v6.5
CNETdownload	An Internet download directory launched in 1996 as part of CNET.	v8.5
Commvault	A software solution for enterprise data backup and storage management.	v8.0
Datei.to	A general browsing and file transfer hosting service.	v8.5

File Transfer Applications	Description	First Available In
Deposit-Files	An online file storage service.	v8.5
DirectConnect	An open-source Windows client for the DirectConnect protocol and Advanced DirectConnect protocol that allows users to connect to a central hub and download files directly from one another.	v8.0
Divshare	A media delivery, sharing, and publishing system.	v8.5
Docstoc	An electronic business document repository and online store.	v8.0
Dropbox	An online file hosting and sharing service.	v8.0
eDonkey	A peer sharing application for storage and distribution of large files.	v8.0
Eset	An Eset Antivirus/Security software solution.	v8.0
Extratorrent	A free download torrent system for movies, music, software, and so on.	v8.5
F-Prot	An antivirus software solution for varying levels of AV protection.	v8.0
FASP	A high-speed, secure file transfer protocol.	v8.0
Filer.cx	A file hosting service that provides free web space for documents, pictures, music, and movies.	v8.0
FileServe	A file hosting service.	v8.5
Filesonic	A general browsing and file transfer service from data storage Web site Filesonic.	v8.5
FilesTube	A file search engine that searches various file sharing and uploading sites like rapidshare, megaupload, mediafire, hotfile, netload, filesonic, and 4shared.	v8.0
FTP	A protocol used to transfer files from a file server to a local machine.	v6.5
FTP Control	The FTP (File Transfer Protocol) control used to manage FTP data transfers from a file server to a local machine.	v8.0
FTP Data	The FTP data flow	v8.0
FTPS	The FTP control used over TLS/SSL.	v8.0
FTPDATA	The FTP data over TLS/SSL	v8.0
FXP	A protocol that provides a method of data transfer that uses the FTP protocol to transfer data from one remote server to another (inter-server) without routing this data through the client's connection.	v8.0
Giganews	A popular Usenet/newsgroup service provider.	v6.5
Gnutella	A large peer-to-peer file-sharing network.	v6.5
GPFS	A high-performance shared-disk clustered file system.	v8.0
GSIFTP	An FTP enhancement that uses GSI security.	v8.0
HiveStor	An open source software program that integrates existing and new commodity hardware to provide a reliable storage network with no single point of failure.	v8.0

File Transfer Applications	Description	First Available In
Hotfile	A free download/upload management tool that increases the speed and stability of your downloads/uploads.	v8.0
iCloud	The cloud data storage and computing service from Apple.	v8.5
ifile.it	An online backup service Web site.	v8.5
iMesh	A media and file-sharing client with online social network features.	v8.0
ImageShack	A general browsing and file transfer service from image hosting Web site ImageShack.	v8.5
Kaspersky	An antivirus software solution for varying levels of AV protection at home and at work.	v8.0
Kat	A torrent download site.	v8.5
Kazaa	A popular file-sharing client that provides unlimited streaming of free music files.	v6.5
KFTP	A file transfer protocol with Kerberos authentication and encryption.	v8.0
KFTPDATA	A protocol for service type Kerberos FTP Data registered with IANA on port 6620 TCP/UDP.	v8.0
Letitbit	A Russian file hosting Web site.	v8.5
Manolito	A free peer-to-peer file sharing network. Users can download music, create play lists, and use instant messaging to chat with friends.	v8.0
MC-FTP	An encrypted multicast file transfer program that transfers files to multiple receivers simultaneously.	v8.0
McAfee	A free, downloadable antivirus software solution, and premium AV software solutions for home and office.	v8.0
McIDAS	A protocol for service type McIDAS Data Transmission registered with IANA on port 112 TCP/UDP.	v8.0
MediaFire	An online file hosting service that enables users to upload, download, manage and share documents, presentations, videos, images and more.	v8.0
Megashares	A file sharing and storage site.	v8.5
MegaUpload	An online storage and file delivery service that includes video browsing through MegaVideo, MegaLive, MegaPix, and Megabox.	v8.0
MSDN	The group within Microsoft responsible for networking with developers and testers.	v8.0
Multiupload	A forwarding site for major upload sites such as Megaupload, Upload King, depositfiles, hotfile, Uploadhere, Zshare, Filesonic, Fileserve, and Wupload.	v8.5
MUTE-net	A peer-to-peer file sharing network that uses a routing algorithm inspired by ant colonies. MUTE-net has not been maintained since April 2007, although software compatibility with the MUTE network has been updated since then.	v8.0
NFA	A network file system that acts as a client for a remote file access protocol, providing access to files on a server.	v8.0

File Transfer Applications	Description	First Available In
NFS	A network file system protocol that enables a user on a client computer to access files over a network.	v6.5
NI FTP	The network independent file transfer program.	v8.0
NNTP	An Internet application protocol used for transporting Usenet news articles (netnews) between news servers and for reading and posting articles by end-user client applications.	v6.5
NovaBACKUP	A data protection and availability software solution that offers support for multi-OS environments and is capable of handling thousands of servers and petabytes of information.	v8.0
OFTP	A protocol used for EDI (electronic data interchange) between two communications business partners.	v8.0
OFTPS	An FTP protocol used over SSL/TLS primarily for electronic data interchange between two communications business partners.	v8.0
Online-File-Folder	A storage service provided by godaddy.com.	v8.5
Paltalk File Transfer	A software file transfer application that allows users to communicate through instant messaging, voice, and video chat.	v8.0
Panda	An antivirus software solution for varying levels of AV protection.	v8.0
Pando	A free file-sharing application that uses both peer-to-peer and client-server architectures. Users are able to send files that might be too large to send through email.	v8.0
PFTP	A file transfer protocol that transfers files, directories, and data to other hosts running pftp.	v8.0
PutLocker	An online file hosting site.	v8.5
QFT	The queued file transport protocol.	v8.0
RapidShare	An online file hosting and sharing service.	v8.0
SBNTBCST	A file transfer protocol.	v8.0
SFTP	A secure file transfer protocol typically used with the SSH protocol.	v8.0
Share P2P	A closed-source p2p application developed in Japan.	v8.0
SkyDrive	A Microsoft online storage service that is part of Windows Live.	v6.5
SuperNews	A Usenet/newsgroup service provider.	v6.5
Swift RVFP	A protocol for service type swift-rvf registered with IANA on port 97 TCP/UDP.	v8.0
TFTP	A lightweight file transfer protocol.	v8.0
TFTPS	A lightweight file transfer protocol over SSL/TLS.	v8.0
Torrentz	A Finland-based metasearch engine for BitTorrent.	v8.5
Uploading	An online paid cloud storage service.	v8.5
Usenet	A worldwide distributed Internet discussion system. Users read and post messages (called articles or posts, and collectively termed news) to one or more categories, known as newsgroups.	v6.5
UUCP	A protocol for service type UUCP Path Service registered with IANA on port 117 TCP/UDP.	v8.0

File Transfer Applications	Description	First Available In
WebDAV	A Web-based distributed authoring and versioning system that enables users to collaboratively edit and manage files on a remote Web server.	v8.0
Windows Update	A service provided by Microsoft that enables users to get software patches and updates for MS Windows and other programs, including Internet Explorer, over the Internet.	v8.0
WinMX	WinMX is a free peer-to-peer file sharing program. It runs on the Windows operating system; however, the official WinMX Web site and WinMX servers have been offline since 2005 as a result of an increased presence of dummy files on the site, which led to a lawsuit. The application now operates through third-party modifications.	v8.0
Winny	A Japanese peer-to-peer (P2P) file-sharing program.	v8.0
Xunlei	A download manager that supports file transfers using HTTP, FTP, eDonkey, and BitTorrent protocols	v8.0
Yahoo Msg File Transfer	A file transfer protocol for traffic within the Yahoo Messenger client.	v8.0
YouSendIt	A Web-based secure digital file delivery company that lets users securely send, receive, and track files on demand.	v8.0
ZanNet	A combination Windows 95 network client and UNIX server that provides Windows 95 network drive access to your server files. Intended to replace both File Transfer Protocol (FTP) and Telnet programs, ZanNet accesses Web pages and remote files over your current Internet connection.	v8.0
Game Applications	Description	First Available In
4399COM	A general browsing and game play on a Chinese casual gaming web site.	v8.5
Battle.net	A premium gaming service provided by Blizzard Entertainment.	v8.0
Bet365	An online gaming and betting Web site.	v8.5
Doof	A free online gaming Web site.	v8.5
Evony	An Adobe Flash-based multiplayer online game, set in medieval times.	v8.5
Farmville	A real-time farm simulation game developed by Zynga, available as an application on Facebook and as an app on the Apple iPhone.	v8.0
IMGames	A protocol for service-type IMGames registered with IANA on port 1077 TCP/UDP.	v8.0
Mafiawars	A multiplayer browser game created by Zynga. It is on several social networking sites and on the iPhone.	v8.0
Steam	An online gaming social networking Web site. Users can purchase, download, and play games, as they connect with friends and groups with similar interests.	v8.0
Steam Client	A client for digital-based distribution used with HTTP traffic for Steam store browsing and news updates.	v8.0

Game Applications	Description	First Available In
Steam DLC	The downloadable content from Steam (games, updates, and so on).	v8.0
Steam Game	The Steam online gaming traffic.	v8.0
Steam Social	The Steam social traffic (friends network, peer-to-peer voice chat).	v8.0
Xbox LIVE	An online multiplayer gaming and digital media delivery service created and operated by Microsoft Corporation.	v8.0
Y8	A general browsing, game play, and streaming media Web site from Y8.com.	v8.5
Zynga	A social network game developer of browser-based games that work both stand-alone and as application widgets on social networking Web sites such as Facebook and MySpace.	v8.0

Messaging Applications	Description	First Available In
AIM	An instant messaging and presence application that enables users to conduct person-to-person instant messaging, chat room messaging, peer-to-peer file sharing, and Facebook support, among other features.	v6.5
APNS	An Apple push notification service that opens a constant IP connection to forward notifications from its servers to Apple devices.	v8.5
C2DM	A service that helps developers send data from servers to their applications on Android devices. The service provides a simple, lightweight mechanism that servers can use to tell mobile applications to contact the server directly, to fetch updated application or user data. The C2DM service handles all aspects of message queueing and delivery to the target application running on the target device.	v8.0
Fring	A mobile messaging application to control messaging flow traffic.	v8.5
Google Talk	The VoIP application used with Google IM/chat.	v6.5
Google Talk Gadget	The Flash-based Google Talk IM client.	v8.0
Hushmail	A Web-based email service offering PGP-encrypted email, file storage, and vanity domain service.	v8.5
ICQ	An instant messaging computer program.	v8.0
IMO.im	An instant messenger service that allows for third-party authentication.	v8.5
IRC	A popular form of real-time Internet text messaging.	v6.5
ISCHAT	An integrated set of live voice, chat, and email response services that enable online businesses to deliver just-in-time, personalized, interactive assistance to each visitor (now known as ATG Live Help).	v8.0
Koolim	A Web messaging site that combines all of the most popular instant messaging services together.	v8.5

Messaging Applications	Description	First Available In
Line2	A mobile VoIP application that allows the user to add a second line to their iPhone or Android device, or to give a phone number to an iPad. Data is transferred over WiFi, cellular data, or cellular voice connections. Text messaging is supported for US-based customers only.	v8.0
Open-Webmail	A simple Web mail service.	
Meebo	An instant messaging (IM) Web platform that can connect with numerous IM networks.	v8.0
MPM	An Internet Message Protocol - RFC 753	v8.0
MSMQ	A messaging protocol that enables applications running on separate servers/processes to communicate in a failsafe manner.	v8.0
MSN2Go	A third-party service for Windows Live Messenger.	v8.5
MSNP	An instant messaging protocol developed by Microsoft for use by the .NET Messenger Service and the instant messaging clients that connect to it, such as Windows Live Messenger.	v8.0
MSP	An application layer protocol used to send a short message between nodes on a network.	v8.0
OSCAR	Open System for Communication in Realtime is AOL's flagship instant messaging and presence information protocol. Currently, OSCAR is in use for two main AOL instant messaging systems: ICQ and AIM.	v8.0
Paltalk Chat	The Paltalk instant messaging text messaging traffic.	v8.0
Paltalk Messenger	An Internet and downloadable chat service that enables users to communicate through instant messaging, voice, and video chat	v8.0
Pinger	A software application that enables you to send and receive free texts (real SMS) with your own free texting number.	v8.0
QOTD	An Internet protocol defined in RFC 865. It is intended for testing and measurement purposes.	v8.0
QQ	A free instant messaging computer program in mainland China.	v8.0
Skype	A proprietary service that enables users to chat, make voice and video calls, and transfer files over the Internet.	v6.5
Skype Auth	The Skype IP authentication and registration.	v8.0
Skype Out	The service that allows Skype users to call phone numbers, including landline and mobile phones, for a fee.	v8.0
Skype p2p	The Skype peer-to-peer traffic, chat, file-transfer, voice, and video.	v8.0
Skype Probe	The Skype discovery probe, used to locate open ports and automatically detect a local Web proxy.	v8.0
Tango	A free mobile video communications service that works on the PC, iPhone, iPod touch, Windows Phone 7, hundreds of Android phones and tablets, and 3G, 4G and Wi-Fi.	v8.5
Vchat	An Internet conferencing protocol.	v8.0
Viber	A mobile application for iPhone and Android that enables users to make free phone calls and send text messages to anyone else using the installed application.	v8.5



Messaging Applications	Description	First Available In
WhatsApp	WhatsApp Messenger is a cross-platform mobile messaging application that enables message exchanges without paying for SMS.	v8.0
Windows Live	A collection of Microsoft online services.	v8.0
XMPP	The extensible messaging and presence protocol, an open technology for real-time communication.	v6.5
Yahoo Messenger	The Yahoo instant messaging client.	v6.5

Networking Applications	Description	First Available In
3COM-TSMUX	A queuing protocol for service type 3COM-tsmux, registered with IANA on port 106 TCP/UDP.	v8.0
8021Q	A protocol that enables nodes on different VLANs to communicate with one another through a network switch with Network Layer (Layer-3) capabilities, or through a router.	v8.0
914CG	A Texas Instruments 914C/G terminal protocol for service type 914c-g (alias: 914c/g) registered with IANA on port 211 TCP/UDP.	v8.0
ACA Services	A DEC Application Control Architecture Services protocol for service type ACAS registered with IANA on Port 62 TCP/UDP.	v8.0
ACI	The application communication interface registered with IANA on Port 187 TCP/UDP.	v8.0
Active Directory	The Microsoft Active Directory protocol.	v6.5
ActiveSync	The Microsoft Exchange ActiveSync notifications, on IANA port 1034/TCP and 1034/UDP.	v6.5
AD Backup	The Microsoft Active Directory backup service.	v8.0
AD DRS	The Microsoft Active directory replication services.	v8.0
AD DSAOP	The Microsoft Active Directory DSAOP services.	v8.0
AD DSROL	The Microsoft Active Directory domain services that help administrators securely manage users, computers, and other devices on the network and facilitates resource sharing and collaboration between users.	v8.0
AD FRS	The Microsoft Active Directory file replication service.	v8.5
AD NSP	The Microsoft Active Directory name service provider.	v8.0
AD Restore	The Microsoft Active Directory restore service.	v8.0
AD XDS	The Microsoft Active Directory Extended Directory Service that enables AD to be extended to store custom data that is of interest to the enterprise.	v8.0
AED512	AED 512 Emulation Service	v8.0
Alias	port 1187/TCP and 1187/UDP	v8.0
ANET	ATEXSSTR	v8.0
ANSA Notify	ANSA REX Notify	v8.0
ANSA REX Trader	ANSA REX Trader	v8.0

Networking Applications	Description	First Available In
Apple ARP	The Apple Computer system that enables AppleTalk protocol to work over networks other than LocalTalk, such as Ethernet or Token Ring.	v8.0
AppleShare	AppleShare IP WebAdmin	v8.0
AppleTalk	A proprietary suite of protocols developed by Apple Inc. for networking computers.	v8.0
ARCISDMS	Protocol for service type Arcisdms registered with IANA on port 262 TCP/UDP.	v8.0
Ariel	Ariel hardware and software scans articles, photos, and other documents and transmits the electronic images to other Ariel workstations anywhere in the world using either FTP or email. Also converts them to PDF for easy delivery.	v8.0
ARNS	port 384/TCP and 384/UDP	v8.0
ARP	A computer networking protocol for determining a network host's link layer or hardware address when only its Internet layer (IP) or network layer address is known.	v8.0
ASA	port 386/TCP and 386/UDP	v8.0
ATM FATE	Frame-based ATM Transport over Ethernet	v8.0
ATM MPOA	Multi-protocol over ATM	v8.0
AUDIT	Unisys Audit SITP	v8.0
Aurora	A link layer communications protocol for use on point-to-point serial links. Developed by Xilinx, it is intended for use in high-speed (tens of gigabits/second or more) connections internally in a computer.	v8.0
BGMP	Border Gateway Multicast Protocol	v8.0
BGP	BGP (Border Gateway Protocol) is the protocol backing the core routing decisions on the Internet.	v6.5
BH611	Protocol for service type bh611 registered with IANA on port 354 TCP/UDP.	v8.0
BHEVENT	Protocol for service type bhevent registered with IANA on port 357 TCP/UDP.	v8.0
BHFHS	Protocol for service type bhfhs registered with IANA on port 248 TCP/UDP.	v8.0
BHMDS	Protocol for service type bhmds registered with IANA on port 310 TCP/UDP.	v8.0
Blackjack	port 1025/TCP and 1025/UDP	v8.0
Bnet	port 415/TCP and 415/UDP	v8.0
Cableport AX	Protocol for service type Cable port A/X registered with IANA on port 282 TCP/UDP	v8.0
CALlic	Computer Associates Int'l License Server	v8.0
CAP	port 1026/TCP and 1026/UDP	v8.0
CDC	Certificate Distribution Center	v8.0

Networking Applications	Description	First Available In
Cisco DRP	(DRP) Director Response Protocol enables the Cisco DistributedDirector product to query routers (DRP Server Agents) in the field for Border Gateway Protocol (BGP) and Interior Gateway Protocol (IGP) routing table metrics between distributed servers and clients.	v8.0
Cisco FNATIVE	Used for Cisco Proprietary Protocols on Cisco Catalyst Network Analysis Modules.	v8.0
Cisco GDP	The Gateway Discovery Protocol (GDP) allows hosts to dynamically detect the arrival of new routers, as well as determine when a router goes down.	v8.0
Cisco SYSMAINT	Cisco SYSMAINT	v8.0
Cisco TNATIVE	Cisco TNATIVE	v8.0
CL1	Network Innovations CL/1	v8.0
CLDAP	The connectionless lightweight directory access protocol is an application protocol for accessing and maintaining distributed directory information services over an Internet protocol network using UDP.	v8.5
Clearcase	A software tool for revision control (for example, configuration management or SCM) of source code and other software development assets. It is developed by the Rational Software division of IBM. ClearCase forms the base of revision control for many large- and medium-sized businesses and can handle projects with hundreds or thousands of developers.	v8.0
CLOANTO	The cloanto.net infrastructure provides redundant hosting, email, and telecommunications services.	v8.0
Coda Auth	Coda Authentication Service, part of Coda file system services, developed by Carnegie Mellon University. Protocol for service type codaauth2 registered with IANA on port 370 TCP/UDP.	v8.0
CompressNET	CompressNET is a commercial WAN compression protocol.	v8.0
COMSCM	SCM Microsystems is a leading provider of solutions for secure access, secure identity, and secure exchange.	v8.0
CORBA	Common Object Request Broker Architecture (CORBA) is a standard defined by the Object Management Group (OMG) that enables software components written in multiple computer languages and running on multiple computers to work together (that is, it supports multiple platforms). Domino Internet Inter-ORB Protocol (DIIOP) is CORBA over IIOP for Lotus Domino. DIIOP allows external programs to attach to and manipulate Domino databases. DIIOP is frequently used to allow Java-based and other non-CORBA programs to connect to Lotus Domino.	v8.0
corerjd	Protocol for service type corerjd registered with IANA on port 284 TCP/UDP	v8.0
Covia CI	Covia Communications Integrator	v8.0
CSISGWP	port 348/TCP and 348/UDP	v8.0
CSNET-NS	CSNET Mailbox Nameserver	v8.0
CVHOSTD	port 442/TCP and 442/UDP	v8.0

Networking Applications	Description	First Available In
DASP	This protocol is designed to provide an unordered, reliable, secure session for full-duplex datagram exchange that can be implemented for low-power wireless networks and low-cost devices.	v8.0
DATEX-ASN	An application profile specification that uses protocols to address the Application Layer (Layer 7 of the OSI Reference Model), the Presentation Layer (Layer 6 of the OSI Reference Model), and that defines the Session Layer (Layer 5 of the OSI Reference Model) as null.	v8.0
DCAP	An application layer protocol used between workstations and routers to transport SNA/NetBIOS traffic over TCP sessions.	v8.0
DCCP	DCCP (Datagram Congestion Control Protocol) is a transport protocol used for congestion control. Applications include Internet telephony and video/audio streaming.	v8.0
DCE/RPC	Distributed Computing Environment / Remote Procedure Calls is the remote procedure call system developed for the Distributed Computing Environment (DCE).	v8.0
DEC Auth	DEC Auth	v8.0
DEC Debug	Decladebug is a source code debugger targeted at debugging software on the local machine or a remote Digital UNIX box.	v8.0
DECVMS	port 441/TCP and 441/UDP	v8.0
DHCP	DHCP (Dynamic Host Configuration Protocol) is an automatic configuration protocol used for assigning IP addresses.	v6.5
DHCPv6	DHCP (Dynamic Host Configuration Protocol) for IPv6	v8.0
Direct	A protocol for service type Direct registered with IANA on port 242 TCP/UDP.	v8.0
DIXIE	A lightweight Directory Assistance protocol.	v8.0
DLS	A directory location service that provides information on the location (addresses) and protocols needed to access white pages name servers.	v8.0
DNA-CML	A protocol for service type DNA-CML registered with IANA on port 436 TCP/UDP.	v8.0
DNS	A domain name system that provides hostname resolution for finding hosts on a network.	v6.5
DNSIX	The Department of Defense Intelligence Information System Network Security for Information Exchange (DNSIX) is a collection of security requirements for networking defined by the U.S. Defense Intelligence Agency.	v8.0
DPSI	A protocol for service type Desktop Paging Software, Inc registered with IANA on port 315 TCP/UDP.	v8.0
DSFGW	A protocol for service type dsfgw registered with IANA on port 438 TCP/UDP.	v8.0
DSP	The display support protocol.	v8.0
DSP3270	The display systems protocol for service type dsp3270 registered with IANA on port 246 TCP/UDP	v8.0

Networking Applications	Description	First Available In
DSSETUP	The Microsoft Active Directory's Directory Services Setup.	v8.0
DTAG	A telecommunications company headquartered in Bonn, Germany. Deutsche Telekom AG is the largest telecommunications company in Europe.	v8.0
DTK	A deception toolkit designed to make it appear to attackers as if the system running it has a large number of widely known vulnerabilities.	v8.0
EGP	The exterior gateway protocol, an obsolete routing protocol for the Internet.	v8.0
EMBLNDT	port 394/TCP and 394/UDP	v8.0
EMFIS	The EMFIS Service	v8.0
EntrustTime	The EntrustTime protocol.	v8.0
Epmap	The Microsoft EPMap (End Point Mapper), also known as DCE/RPC Locator service, is used to remotely manage services.	v6.5
ESRO	The Efficient Short Remote Operations service is a Remote Procedure Call service.	v8.0
ETH	A framing protocol that carries data to and from LANs (Local Area Networks).	v8.0
ETOS	A protocol for service type NEC Corporation registered with IANA on port 377/378 TCP/UDP.	v8.0
Fatmen	A protocol for service type Fatmen Server registered with IANA on port 347 TCP/UDP.	v8.0
FileMaker	A computer software company formed in 1998 from Claris as a wholly owned subsidiary of Apple Inc. FileMaker develops, supports and markets two relational database programs; FileMaker and Bento. Filemaker is available for both Mac OS X and Microsoft Windows operating systems and is aimed toward business use, or home users with high-end needs. Bento, aimed at the home user or basic small business user, is a Mac OS X application with additional versions available for the iPhone and iPad.	v8.0
GACP	A protocol for service type Gateway Access Control Protocol registered with IANA on port 190 TCP/UDP.	v8.0
Genesis PPP	A protocol for service type Genesis Point-to-Point Trans Net registered with IANA on port 103 TCP/UDP.	v8.0
Genie	An old network management and diagnostic protocol.	v8.0
GENRAD	A protocol for service type GENRAD-MUX registered with IANA on port 176 TCP/UDP.	v8.0
GIST	A protocol for service type Q-most encapsulation for general Internet signalling transport messages registered with IANA on port 270 UDP.	v8.0
Gss License	A protocol for service type GSS X License Verification registered with IANA on port 128 TCP/UDP.	v8.0

Networking Applications	Description	First Available In
Hassle	A protocol for service type Hassle registered with IANA on port 375 TCP/UDP. HASSLE is a networking application that enables users to execute remote jobs that have a transfer component built in. HASSLE is flexible as it only transfers the data and the parameters. The execution at the remote site is automatically accomplished through a code generator.	v8.0
HBCI	A bank-independent protocol for online banking developed and used by German banks.	
HDAP	A Microsoft HDA protocol for service type hdap registered with IANA on port 263 TCP/UDP.	v8.0
HEMS	A protocol for service type hems registered with IANA on port 151 TCP/UDP.	v8.0
Hostname	A protocol for service type hostname registered with IANA on port 101 TCP/UDP.	v8.0
HP Perf	The Performance Data Collector for HP OpenVMS (TDC) that gathers performance data for OpenVMS systems. By default, TDC periodically collects and stores data in a file. Subsequently, user applications can retrieve and analyze data from the file.	v8.0
HTTPMGT	A protocol for service type HTTP-mgmt registered with IANA on port 280 TCP/UDP.	v8.0
Hyper-G	A publishing system with hypertext features more advanced than those available with the Hypertext Transfer Protocol and today's Web browser.	v8.0
IASD	A protocol for service type IASD registered with IANA on port 432 TCP/UDP.	v8.0
IBM APP	A protocol for service type IBM application registered with IANA on port 385 TCP/UDP.	v8.0
IBM OPC	A protocol that automatically plans, controls, and monitors your production workload to maximize and optimize throughput, but lets you intervene manually when required. This protocol is for service type IBM Operations Planning and Control Start registered with IANA on port 4.	v8.0
ICAD	A knowledge-based engineering (KBE) system based upon the Lisp programming language. ICAD has an open architecture that can use all the power and flexibility of the underlying language.	v8.0
ICP	An Intelligent Communication Protocol registered with IANA port 1112 TCP/UDP.	v8.0
Ident	A protocol that helps identify the user of a particular TCP connection.	v8.0
IDP	A close descendant of PUP's internetwork protocol, and roughly corresponds to the Internet Protocol (IP) layer in TCP/IP.	v8.0
IGMP	An Internet group management protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships.	v8.0
IMSP	An interactive mail support protocol on port 406 TCP/UDP.	v8.0

Networking Applications	Description	First Available In
InBusiness	A protocol for service type inbusiness on TCP port 244, used to connect to the administrative functions on the Dayna Communications InBusiness line of small office network equipment.	v8.0
IP	The principal communications protocol used for relaying datagrams (packets) across an internetwork using the Internet protocol suite.	v8.0
IPv6	The Internet protocol (IP) version 6.	v8.0
IPX	The IPX/SPX protocol stack, supported by the Novell NetWare network operating system.	v8.0
IS-99	A data services option standard for wideband spread spectrum cellular systems that provides asynchronous data transmission capability on TIA/EIA/IS-95-using ports 379 and 380 TCP/UDP.	v8.0
ISAKMP	A protocol defined by RFC 2408 for establishing Security Associations (SA) and cryptographic keys in an Internet environment. ISAKMP only provides a framework for authentication and key exchange and is designed to be key exchange independent.	v8.0
ISI Graphics	An ISI graphics language protocol for service type isi-gl registered with IANA on port 55 TCP/UDP.	v8.0
ISO-TSAP	An ISO transport service access point class protocol for service type iso-tsap registered with IANA on port 102 TCP/UDP.	v8.0
ISOIP	An ISO internetworking protocol for service type iso-ip registered with IANA on port 147 TCP/UDP.	v8.0
JARGON	A protocol for service type jargon registered with IANA on port 148 TCP/UDP.	v8.0
Java RMI	A Java application programming interface that performs the object-oriented equivalent of remote procedure calls (RPC).	v8.0
Kblock	A protocol for service type K-BLOCK registered with IANA on Port 287 TCP/UDP. K-Block protects unattended logged-in terminals from unauthorized access in OpenVMS environments.	v8.0
Kerberos	A network authentication protocol designed to provide strong authentication for client/server applications by using secret-key cryptography.	v6.5
KIS	A protocol for service type KIS Protocol registered with IANA on port 186 TCP/UDP	v8.0
KNETCMP	KNET/VM Command/Message Protocol	v8.0
Kryptolan	port 398/TCP and 398/UDP	v8.0
LA-Maint	IMP Logical Address maintenance	v8.0
Legent	The protocols for service type legent-1 (Legent Corporation) registered on IANA on port 373 TCP/UDP and service type legent-2 (Legent Corporation) registered on IANA port 374 TCP/UDP.	v8.0
LINK	A protocol for service type LINK registered with IANA on port 245 TCP/UDP	v8.0
LLMNR	A link-local multicast name resolution protocol used by Windows for local name resolution.	v6.5

Networking Applications	Description	First Available In
Locus Conn	A protocol for service type Locus PC-Interface Conn Server registered with IANA on port 127 TCP/UDP.	v8.0
Locus Map	A protocol for service type Locus PC-Interface Net Map Service registered with IANA on port 125 TCP/UDP.	v8.0
LSARPC	The Microsoft Active Directory Local Security Authority Subsystem Service.	v8.0
Magenta Logic	A protocol for service type Magenta Logic registered with IANA on port 313 TCP/UDP.	v8.0
MANET	The mobile Ad-hoc networks protocol.	v8.0
Masqdiabler	A system that allows authorized LAN users to manipulate the network interface, usually a modem, that gives Internet access on a Linux box without having to use Telnet.	v8.0
MATIP	An application protocol for airline reservation, ticketing, and messaging systems to use over a TCP/IP network.	v8.0
MDNS	A multicast DNS protocol that uses familiar DNS programming interfaces to a smaller network without the conventional DNS server.	v6.5
Meta5	A business intelligence tool that allows users to visually create reports that can access multiple corporate data source. Registered with IANA on port 393 TCP/UDP	v8.0
Metagram	A protocol for service type Metagram Relay registered with IANA on port 99 TCP/UDP.	v8.0
MF Cobol	A micro focus Cobol directory service protocol for service type mfcobol, registered with IANA on port 86 TCP/UDP.	v8.0
MFTP	A communication protocol designed for file sharing. This protocol is used by clients such as eMule and eDonkey and, in its extended implementation, by the Overnet network.	v8.0
MIT Spooler	A protocol for service type MIT Dover Spooler, registered with IANA on port 91 TCP/UDP	v8.0
mit-ml-dev	A protocol for the MIT ML device, registered with IANA on port 83 TCP/UDP.	v8.0
MobileIP	An Internet Engineering Task Force (IETF) standard communications protocol that is designed to allow mobile device users to move from one network to another while maintaining a permanent IP address.	v8.0
MortgageWare	A product developed by Interlinq Software Corp that automates all components of the loan originating process.	v8.0
MPLS Multicast	MPLS (Multiprotocol Label Switching) multicast traffic	v8.0
MPLS Unicast	MPLS (Multiprotocol Label Switching) unicast traffic	v8.0
MPP	Netix Message Posting Protocol is a network protocol that is used for posting messages from a computer to a mail service host.	v8.0
MPTN	port 397/TCP and 397/UDP	v8.0
MS CRS	A Microsoft Content Replication System protocol used on ports 507/TCP and 507/UDP.	v8.0
MSG	port 29/TCP and 29/UDP; port 31/TCP and 31/UDP	v8.0



Networking Applications	Description	First Available In
Multiplex	Network Innovations Multiplex	v8.0
MUMPS	Plus Five's MUMPS	v8.0
NAMP	Neighbor Aware Multicast Routing Protocol	v8.0
NCED	port 404/TCP and 404/UDP	v8.0
NCLD	port 405/TCP and 405/UDP	v8.0
NDS Auth	A software module from Symantec Corporation	v8.0
NetBIOS	NetBIOS (Network Basic Input/Output System).	v6.5
Netinfo	port 1033/TCP and 1033/UDP	v8.0
Netlogon	The Microsoft Net Logon service verifies logon requests, and it registers, authenticates, and locates domain controllers.	v8.0
NETSC	Protocols for service type netsc-prod registered with IANA on port 154 TCP/UDP and service type netsc-dev registered with IANA on port 155 TCP/UDP	v8.0
NetScout	port 395/TCP and 395/UDP	v8.0
Netware	A network operating system developed by Novell, Inc. It initially used cooperative multitasking to run various services on a personal computer, with network protocols based on the archetypal Xerox Network Systems stack.	v8.0
NIP	port 376/TCP and 376/UDP	v8.0
NNSP	An Internet application protocol used for transporting Usenet news articles (netnews) between news servers and for reading and posting articles by end-user client applications.	v8.0
NPP	Network Printing enables users in locations geographically separate from each other and from their print devices to produce documents for themselves and others.	v8.0
NSIIOPS	IIOP Name Service	v8.0
NSRMP	port 359/TCP and 359/UDP	v8.0
NSS	NSS Routing	v8.0
NSSTP	Nebula Secure Segment Transfer Protocol, IANA port 1036/TCP and 1036/UDP	v8.0
NTP	NTP (Network Time Protocol) is used for synchronizing the clocks of computer systems over a network.	v6.5
NXEdit	Protocol for service type NXEdit registered with IANA on port 126 TCP/UDP	v8.0
NXTSTEP	NextStep Window Server	v8.0
OCBinder	OCBinder	v8.0
OCS	Microsoft Office Communications Server 2007 R2 delivers streamlined communications to users, so everyone in an organization can communicate with the right person, right away, from the applications they use most.	v8.0
OCServer	OCServer	v8.0

Networking Applications	Description	First Available In
ODMR	An SMTP extension standardized in RFC 2645 that relays email to after the authenticating the sender. It uses the extended SMTP command ATRN. It is similar to the ETRN command but works with dynamically assigned IP addresses.	v8.0
Onmux	Protocol for service type Onmux registered with IANA on port 417 TCP/UDP	v8.0
Openport	Protocol for service type Openport registered with IANA on port 260 TCP/UDP	v8.0
OSUNMS	OSU Network Monitoring System	v8.0
PAWSERV	Allows you to analyze transaction performance and behavioral problems by providing a platform for investigating logs and other historical data	v8.0
PDAP	port 344/TCP and 344/UDP	v8.0
PersonalLink	Personal Link	v8.0
PIM	PIM (Protocol Independent Multicast) is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet.	v8.0
PIP	PIP	v8.0
PKIX Timestamp	The PKIX TS specifies the format of packets, along with some possible transport protocols and some verifications to be done by the server and the client.	v8.0
PPP Discovery	PPPoE (Point-to-point Protocol over Ethernet) discovery messages	v8.0
PPP Session	PPPoE (Point-to-point Protocol over Ethernet) session messages	v8.0
Printer	A standard network protocol for remote printing as well as for managing print jobs, media size, resolution, and so forth. Like all IP-based protocols, IPP can run locally or over the Internet to printers hundreds or thousands of miles away. Unlike other printing protocols, IPP also supports access control, authentication, and encryption, making it a much more capable and secure printing solution than older ones.	v8.0
PRINTSRV	Network PostScript	v8.0
PROFILE	PROFILE Naming System	v8.0
PROSPERO	Prospero Directory Service is a name server based on the virtual system model.	v8.0
PTP	A high-precision time protocol for synchronization used in measurement and control systems residing on a local area network. Accuracy in the sub-microsecond range might be achieved with low-cost implementations.	v8.0
PUP	One of the two earliest internetwork protocol suites. The entire suite provided routing and packet delivery, as well as higher level functions such as a reliable byte stream, along with numerous applications.	v8.0
PWDGEN	Password Generator Protocol, rfc 972.	v8.0

Networking Applications	Description	First Available In
Qbik	Qbik has developed sophisticated & user friendly software specializing in Internet connectivity and security. Our products allow users to manage their Internet connections (WinGate), connect remote offices together (WinGate VPN), and combat network security issues (NetPatrol).	v8.0
Radius	Provides centralized Authentication, Authorization, and Accounting (AAA) management for computers to connect and use a network service.	v6.5
RADIUS-ACCT	A client component of Remote Access, Virtual Private Network, and Network Access servers that authenticates users or devices before granting network access and accounts for service use by communicating with the RADIUS server.	v8.5
RAP	Route Access Protocol, a general protocol for distributing routing information at all levels of the Internet.	v8.0
RARP	An obsolete computer networking protocol used by a host computer to request its Internet Protocol (IPv4) address from an administrative host, when it has available its Link Layer or hardware address, such as a MAC address.	v8.0
ResCap	ResCap Resolution Protocol	v8.0
RIP	RIP (Routing Information Protocol) is a dynamic routing protocol.	v6.5
RLP	RLP (Resource Location Protocol) is used to help find network services.	v8.0
RMT	port 411/TCP and 411/UDP	v8.0
RPC2PMAP	An ONC RPC service that runs on network nodes that provide other ONC RPC services.	v8.0
RRP	port 648/TCP and 648/UDP	v8.0
RSVD	RSVD	v8.0
RSVP	Resource Reservation Protocol, a control protocol designed to reserve resources across a network.	v8.0
Rsync	A software application for UNIX systems that synchronizes files and directories from one location to another while minimizing data transfer using delta encoding when appropriate. An important feature of rsync not found in most similar programs/protocols is that the mirroring takes place with only one transmission in each direction. rsync can copy or display directory contents and copy files, optionally using compression and recursion.	v8.0
SAMR	Microsoft Active Directory Security Account Manager	v8.0
SCCM	System Center Configuration Manager (CM07 or SCCM or ConfigMgr or Configuration Manager), formerly Systems Management Server (SMS), is a systems management software product by Microsoft for managing large groups of Windows-based computer systems. Configuration Manager provides remote control, patch management, software distribution, operating system deployment, network access protection, and hardware and software inventory.	v8.0
SCOI2DLG	port 360/TCP and 360/UDP	v8.0
SCSI-ST	A set of standards for physically connecting and transferring data between computers and peripheral devices.	v8.0

Networking Applications	Description	First Available In
SCTP	A Transport Layer protocol, serving in a similar role to the popular protocols TCP and UDP. It provides some of the same service features of both: it is message-oriented like UDP and ensures reliable, in-sequence transport of messages with congestion control like TCP.	v8.0
SecurSight	An architecture that combines authentication, authorization, and secure communications. The primary goal of this architecture is to secure access to network resources, while providing a migration path from legacy authentication and authorization methods to a public-key infrastructure.	v8.0
Semantix	A protocol for service type Semantix registered with IANA on Port 361 TCP/UDP. The Semantix ASN.1 compiler is an open source ASN.1 compiler.	v8.0
SEND	The Secure Neighbor Discovery protocol.	v8.0
SET	A standard secure electronic transaction protocol for securing credit card transactions over insecure networks.	v8.0
SGCP	A communications protocol used within a voice over Internet protocol system. It has been superseded by MGCP, an implementation of the media gateway control protocol architecture.	v8.0
Shrinkwrap	A protocol for service type Shrinkwrap registered with IANA on Port 358 TCP/UDP.	v8.0
Silverplatter	A protocol used for the delivery of information in a digital library across the network. SilverPlatter Information, Inc. was one of the first companies to produce commercial reference databases on CD-ROMs.	v8.0
SLOW	A slow protocols dissector that implements support of the link aggregation control protocol and OAM.	v8.0
SMAKYNET	SMAKYNET Protocol	v8.0
Smart SDP	port 426/TCP and 426/UDP	v8.0
SMPTE	port 420/TCP and 420/UDP	v8.0
SMSP	port 413/TCP and 413/UDP	v8.0
SNET	Sirius Systems	v8.0
SNPP	A protocol that defines a method by which a pager can receive a message over the Internet. It is supported by most major paging providers, and serves as an alternative to the paging modems used by many telecommunications services.	v8.0
SoftPC	A protocol developed by Insignia Solutions for service type softpc registered with IANA on Port 215 TCP/UDP.	v8.0
SRC	An IBM System Resource Controller that facilitates the management and control of complex subsystems. The SRC is a subsystem controller.	v8.0
SRMP	The Spider Remote Monitoring protocol.	v8.0
SRS Send	port 362/TCP and 362/UDP	v8.0
SSDP	The Simple Service Discovery Protocol, used for discovery of Universal Plug-and-Play services.	v8.0

Networking Applications	Description	First Available In
STUN	The session traversal utilities used in NAT traversal for applications with real-time voice, video, messaging, and other interactive communications.	v8.0
Sun RPC	A protocol for service type Sun Remote Procedure Call registered with IANA on Port 111 TCP/UDP. Sun RPC is a widely deployed remote procedure call system.	v8.0
SURMEAS	A protocol for service type Survey Measurement registered with IANA on Port 243 TCP/UDP.	v8.0
SVRLOC	A service discovery protocol that allows computers and other devices to find services in a local area network without prior configuration. SLP has been designed to scale from small, unmanaged networks to large enterprise networks. It has been defined in RFC 2608 and RFC 3224 as Standards Track document.	v8.0
Sybase SQL	A comprehensive suite of solutions that provides data management, synchronization and data exchange technologies that enable the rapid development and deployment of database-powered applications in remote and mobile environments.	v8.0
SynOptics	A network management protocol that has been changed many times through acquisitions. SynOptics Communications is credited with having invented the concept of the modular Ethernet hub and high-speed Ethernet networking over copper twisted-pair and fiber optic cables.	v8.0
TAC News	A protocol for service type TAC News registered with IANA on port 98 TCP/UDP.	v8.0
TACACS	A remote authentication protocol used to communicate with an authentication server commonly used in UNIX networks. TACACS allows a remote access server to communicate with an authentication server in order to determine if the user has access to the network.	v6.5
TCP	A core transmission control protocol of the Internet protocol suite that enables reliable communication between hosts.	v8.0
TCPMUX	A multiplexing service that might be accessed with a network protocol to contact any one of a number of available TCP services of a host on a single, well-known port number.	v8.0
TCX Flash	TCX Flash redirection and acceleration software from Wyse.	v8.5
TCX Multimedia	TCX multimedia redirection and acceleration software from Wyse.	v8.5
TCX Sound	TCX sound software from Wyse receives and transmits high-quality audio.	v8.5
TCX USB	TCX USB imaging creator software from Wyse.	v8.5
Texar	A policy-based authorization solution to securely control what people can do with highly valuable or critical data across the extended enterprise.	v8.0
Timbuktu	A remote control software product developed by Motorola. Remote control software allows a user to control another computer across the local network or the Internet, viewing its screen and using its keyboard and mouse as if he or she were sitting in front of it. Timbuktu is compatible with computers running both Mac OS X and Windows.	v8.0

Networking Applications	Description	First Available In
Time	A network protocol in the Internet Protocol Suite defined in 1983 in RFC 868. Its purpose is to provide a site-independent, machine readable date and time.	v6.5
Tobit	A client/server software solution for NetWare or Windows NT Server that enables users on a LAN to send and receive faxes directly from their network-connected PC.	v8.0
UAAC	A protocol for service type uaac registered with IANA on Port 145 TCP/UDP.	v8.0
UARPS	A protocol for service type Unisys ARPs registered with IANA on Port 219 TCP/UDP.	v8.0
UDP	A core user datagram protocol of the Internet protocol suite that enables low overhead and unreliable communication between hosts, often for real-time data transfer.	v8.0
UIS	A protocol for service type UIS registered with IANA on Port 390 TCP/UDP.	v8.0
ULSTPROC	A ListProcessor, ListProc for short, used as a powerful mailing list agent to track thousands of people subscribed to any number of mailing lists.	v8.0
Unidata LDM	A collection of cooperating programs that select, capture, manage, and distribute arbitrary data products. The system is designed for event-driven data distribution, and is currently used in the Unidata Internet Data Distribution (IDD) project. The LDM system includes network client and server programs and their shared protocols.	v8.0
UNIFY	A protocol for service type Unify registered with IANA on port 181 TCP/UDP	v8.0
UPS	An electrical apparatus that provides emergency power to a load when the input power source, typically the utility mains, fails. A UPS differs from an auxiliary or emergency power system or standby generator in that it will provide instantaneous or near-instantaneous protection from input power interruptions by means of one or more attached batteries and associated electronic circuitry for low power users, and or by means of diesel generators and flywheels for high power users.	v8.0
UTMP	A file on UNIX-like systems that keeps track of all system log in and log out activity. It was never a part of any official UNIX standard, such as Single UNIX Specification, and was obsoleted with introduction of utmpx and corresponding APIs	v8.0
vetTCP	A protocol for service type vetTCP registered with IANA on port 78 TCP/UDP	v8.0
VMNET	A protocol for service type VMNET registered with IANA on port 175 TCP/UDP	v8.0
VMPWCS	A protocol for service type VM PWSCS registered with IANA on port 214 TCP/UDP	v8.0
VSLMP	A protocol for service type vslmp registered with IANA on port 312 TCP/UDP	v8.0
WCCP	A Cisco-developed content-routing protocol that provides a mechanism to redirect traffic flows in real-time to web-caches.	v8.0

Networking Applications	Description	First Available In
WebFilter	A WebFilter Remote Monitor, IANA port 1046/TCP and 1046/UDP.	v8.0
Whois	A query and response protocol that is widely used for querying databases that store the registered users or assignees of an Internet resource, such as a domain name, an IP address block, or an autonomous system, but is also used for a wider range of other information.	v8.0
WINS	The Microsoft implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	v6.5
Wyse TCX	A suite of TCX collaborative processing virtualization solutions from Wyse.	v8.5
X.224	A protocol component used in establishing RDP connections.	v8.0
X.25	An ITU-T standard protocol suite for packet switched wide area network (WAN) communication.	v8.0
Xbone	A system for the automated deployment, management, coordination, and monitoring of IP overlay networks.	v8.0
XDMCP	The X Display Manager Control protocol.	v8.0
Xfer	A utility used for DNS zone transfers.	v8.0
XNS	Xerox Network Services.	v8.0
XNS Authentication	Xerox networking services authentication	v8.0
XNS Clearinghouse	The Xerox networking services Clearinghouse protocol.	v8.0
XNS Time	XNS Time Protocol	v8.0
Xyplex	Networking products from Xyplex Networks.	v8.0
Z3950	ANSI Z39.50 is a client-server protocol for searching and retrieving information from remote computer databases.	v8.0
Zebra	A high-performance, general-purpose structured text indexing and retrieval engine. It reads structured records in a variety of input formats (email, XML, MARC) and allows access to them through exact boolean search expressions and relevance-ranked free-text queries.	v8.0

Network Monitoring Applications	Description	First Available In
Chargen	A device or software that produces static or animated text (such as crawls and rolls) for keying into a video stream. Modern character generators are computer-based, and can generate graphics as well as text.	v8.0
Cisco SLA	A control protocol that enables delivery of time-based network and services performance data used in monitoring Service Level Agreements (SLAs).	v8.0
CMIP	The common management information protocol for service type CMIP/TCP Manager registered with IANA on port 163 TCP/UDP.	v8.0

Network Monitoring Applications	Description	First Available In
CTF	The DECnet-Plus Common Trace Facility, used to collect and display information about specific protocol exchanges between systems.	v8.0
Daytime	A service in the Internet Protocol Suite, defined in 1983 in RFC 867. It is intended for testing and measurement purposes in computer networks.	v8.0
DCP	An application level protocol optimized for the integration, monitoring and control of devices on a network.	v8.0
Discard	A service in the Internet Protocol Suite defined in RFC 863. It is intended for testing, debugging, and measurement purposes.	v8.0
Echo	A service in the Internet Protocol Suite defined in RFC 862. It was originally proposed for testing and measurement of round-trip times in IP networks.	v8.0
Finger	A simple network protocol for the exchange of human-oriented status and user information.	v8.0
ICMP	A core protocol of the Internet Protocol Suite, chiefly used by the operating systems of networked computers to send error messages indicating, for instance, that a requested service is not available or that a host or router could not be reached. ICMP can also be used to relay query messages.	v6.5
ICMPv6	The implementation of the Internet Control Message Protocol (ICMP) for Internet Protocol version 6 (IPv6).	v8.0
Opalis Robot	A comprehensive system management and automation solution. It provides real-time monitoring, notification, corrective action and event driven job scheduling to proactively manage your Windows NT/W2K environment.	v8.0
SAP HostControl	An SAP Host Control Agent protocol used for viewing logs and traces of a remote host.	v6.5
SGMP	A protocol to manage and monitor gateways from a controlling entity. SGMP preceded SNMP.	v6.5
SMUX	A computer networking protocol used in implementing the Simple Network Management Protocol. It defines communications between the SNMP Agent and other processes.	v8.0
SNMP	An Internet-standard protocol for managing devices on IP networks.	v6.5
STATSRV	A statistics service for collecting STAT data from hosts.	v8.0
Syslog	A standard for logging program messages.	v8.0
Systat	An internet protocol for system diagnostic information in the form of a list of users currently logged into the system.	v8.0
Tivoli	The central driving mechanism for operations in the IBM Tivoli (Integrated Service Management software) environment.	v8.0
Tripwire	A free software security and data integrity tool useful for monitoring and alerting on specific file change(s) on a range of systems.	v8.0
UMA	A protocol for service type universal management architecture registered with IANA on port 144 TCP/UDP.	v8.0



Proxy Applications	Description	First Available In
Avocent	A protocol for service type Avocent Proxy Protocol registered with IANA on port 1078 TCP/UDP.	v8.0
Hopster	An application that tunnels other applications to bypass firewalls.	v8.0
GlypeProxy	A free Web-based proxy script written in PHP.	v8.5
Privax	A Web anonymity network, aimed at providing people the freedom to surf anonymously online. Privax offers free Web proxy services (The Privax Network), which allows anyone in the world to surf anonymously under their IP address.	v8.0
SOCKS	An Internet protocol that facilitates the routing of network packets between client-server applications through a proxy server.	v8.0
SureSome	A Web proxy that encrypts and tunnels Web traffic.	v8.5
Surrogafier	A Web proxy that tunnels Web traffic.	v8.5
Tor	A free software implementation of second-generation onion routing, a system which claims to enable its users to communicate more anonymously on the Internet.  TorBrowser is a software product designed to make it extremely difficult to determine where the user is located and the Web sites visited. TorBrowser also lets you access sites which are lost.	v8.0

Remote Access Applications	Description	First Available In
Citrix CGP	The Citrix common gateway protocol.	v8.0
Citrix ICA	A proprietary protocol for an application server system, designed by Citrix Systems.	v8.0
Citrix IMA	The Citrix IMA (Independent Management Architecture) protocol is used for server-to-server communication in performing functions such as licensing and server load updates, all of which occur behind the scenes.	v8.0
Citrix Licensing	A protocol for service type Citrix Licensing registered with IANA on port 7279 TCP/UDP.	v8.0
Citrix RTMP	A protocol for service time Citrix RTMP registered with IANA on port 2897 TCP/UDP	v8.0
Citrix SLG	A storage link gateway that enables automated discovery and one-click access to native storage services using any of the leading storage architectures and protocols, including DAS, NAS, SAN, iSCSI, and Fibre Channel.	v8.0
Citrix WANScaler	A WAN accelerator that improves application performance for branch office users.	v8.0
ERPC	A protocol for service type ERPC registered with IANA on port 121 TCP/UDP.	v8.0
HP VMM	A protocol for service type HP VMM Control/Agent registered with IANA on port 1124/1125 TCP/UDP.	v8.0

Remote Access Applications	Description	First Available In
Ktelnet	A protocol that provides telnet clients with authentication and encryption, FTP clients with Kerberos authentication, proxy server functionality, the ability to run through NAT firewalls and firewalls supporting HTTP proxy with CONNECT.	v8.0
KVM	A protocol for service type KVM-via-IP Management Service registered with IANA on port 1132 TCP/UDP.	v8.0
KWDB	A protocol for service type KWDB Remote Communication registered with IANA on port 1127 TCP/UDP.	v8.0
LogMein	A host software application that provides remote access and PC desktop control.	v8.0
PCoIP	A remote access protocol that compresses, encrypts, and encodes a computing experiences at a data center and transmits it across any standard IP network.	v8.0
RDP	A remote desktop protocol that provides users with a graphical interface to another computer.	v6.5
Remote Job Service	A protocol for service type netrjs-1 - 4 registered with IANA on ports 71-74 TCP/UDP.	v8.0
Remote Telnet	A protocol for service type Remote Telnet Service registered with IANA on port 107 TCP/UDP.	v8.0
RJE	The remote job entry processes: sending jobs to mainframe computers from remote workstations, and receiving output from mainframe jobs at remote workstations.	v8.0
rlogin	A software utility for UNIX-like computer operating systems that allows users to log in on another host through a network.	v8.0
RSH	A remote shell service that enables a user to access a remote host and execute a single command upon it without requiring the login and logout steps.	v8.0
SCCM Remote Control	A system center configuration manager that controls a client remotely.	v8.0
SNA Gateway	A server that enables users to exchange information and share resources between configured OpenVMS systems in DECnet and/or TCP/IP environments in a bidirectional manner.	v8.0
SSH	A network protocol that allows data to be exchanged using a secure channel between two networked devices.	v6.5
Su-Mit Telnet	A protocol for service type su-mit-tg registered with IANA on port 89 TCP/UDP.	v8.0
SUPDUP	A protocol that provides for login to a remote system over a network with terminal-independent output.	v8.0
TeamViewer	A remote access application that controls any computer on the Internet. TeamViewer needs to be running on both machines in order to connect. While the main focus is remote control, it also includes desktop sharing, file transfer, and presentation features.	v8.0
Telnet	A network protocol used on the Internet or local area networks to provide a bidirectional, interactive, text-oriented communications facility using a virtual terminal connection.	v8.0

Social Networking Applications	Description	First Available In
Bebo	A social networking Web site where users can post blogs, photographs, music, videos and questionnaires to which other users can answer.	v8.0
Chinacom	A Chinese social media site.	v8.5
Classmates	A social networking site based on high school yearbooks.	v8.5
Delicious	A social bookmarking service for sharing, storing, and discovering bookmarks.	v8.5
Facebook	A social networking service.	v6.5
Facebook-Apps	Any add-ons developed for Facebook; generally games, puzzles, gifts, classifieds and so on.	v8.0
Facebook-Event	An application for creating and editing Facebook events.	v8.5
Facebook-Post	An application that provides interactions with Facebook walls.	v8.5
Flixster	A social movie site allowing users to share movie ratings, discover new movies and meet others with similar movie taste. The site has expanded to include a Facebook app, a MySpace app, and an app for Bebo and Orkut.	v8.0
Foursquare	A location-based social networking Web site for mobile devices where users check in their locations and can find out where their friends are located.	v8.5
FriendFeed	A real-time feed aggregator from social media sites.	v8.0
Friendster	A social networking service.	v6.5
Google +	A social networking service provided by Google Inc. that incorporates existing and new Google services	v8.0
Hi5	A social gaming Web site that allow for third-party authentication through Facebook.	v8.5
LinkedIn	A business-oriented social networking site.	v6.5
Match	An online dating Web site.	v8.5
Meetup	An online social networking portal that facilities offline group meetings in various locations around the world.	v8.5
Multiply	A social shopping site that connects merchants with shoppers, offering both products and services.	v8.5
MySpace	A social networking service.	v6.5
Orkut	A social networking application.	v8.0
Pinterest	An online pinboard to collect and share interests.	v8.5
Plaxo	An online address book and social networking service that provides automatic updating of contact information.	v8.0
Reddit	A social news Web site.	v8.0
Sourceforge	A Web-based source code repository.	v8.5

<b>Social Networking Applications</b>	<b>Description</b>	<b>First Available In</b>
Tagged	A social networking site that allows members to browse the profiles of other members, play games, and share tags and virtual gifts.	v8.5
TwitPic	A picture posting and delivery service.	v8.0
Twitter	A social networking and microblogging service.	v6.5
Yelp	A social networking, user review, and local search service	v8.0
<b>Streaming Media Applications</b>	<b>Description</b>	<b>First Available In</b>
56COM	A video sharing Web site in China.	v8.5
adnStream	A Spanish video streaming Web site.	v8.5
AfreecaTV	A South Korean video streaming service.	v8.5
Dailymotion	A video sharing service Web site.	v8.0
Facebook Video Chat	The Facebook video chat service.	v8.5
Facebook Video	The Facebook streaming video and video upload service.	v8.5
FaceTime	A video conferencing service between supported Apple mobile devices.	v8.5
Freeetv	A streaming media Web site that provides free access to TV.	v8.5
Google Video	A free video sharing Web site and also a video search engine from Google Inc.	v6.5
Grooveshark	An online music search engine and streaming service.	v8.0
H.225	A VoIP call signalling and control protocol.	v8.0
H.245	A control channel protocol used with H.323 and H.324 communication sessions.	v8.0
H.248	An implementation of the Media Gateway Control Protocol architecture for controlling VoIP gateways.	v8.0
H.323	An H.323 VoIP call signalling and control protocol.	v6.5
Hulu	The Hulu online video streaming.	v8.0
iTunes	The Apple Computer, Inc. media player and online store.	v6.5
Last.fm	A social networking music streaming site.	v8.0
Live365	An Internet radio network that allows members to create their own online radio station or listen to others.	v8.5
MagicJack	A USB device that enables any phone to make free calls within the US and Canada.	v6.5
mck-ivpip	A VoIP extender ipvip protocol.	v8.0
Metacafe	A community-based video-sharing site that specializes in short-form original entertainment, where users upload, view, and share video clips.	v6.5

Streaming Media Applications	Description	First Available In
Movie2k	A Web site that allows visitors to stream files without requiring registration.	v8.5
MUZU.TV	An Irish-owned interactive music video site.	v8.5
Netflix site	A subscription-based video streaming service.	v8.0
Netflix video stream	A video streaming service.	v8.0
Paltalk Video	An instant messaging video chat traffic service.	v8.0
Paltalk Voice	An instant messaging audio chat traffic service.	v8.0
Pandora	A free Internet music site.	v8.5
Pandora-Audio	An Internet radio and audio streaming music site.	v8.5
Pandora.tv	A South Korean Web site that specializes in user-generated video sharing.	v8.5
PPStream	A Chinese peer-to-peer streaming video software. The software is available through Web, mobile, and Windows client application.	v8.5
PPTV	An online TV service offering both live streaming and video on demand of TV programs and shows, movies, and sports.	v8.5
PPTV-p2p	PPTV peer-to-peer traffic.	v8.5
Quicktime	An extensible proprietary multimedia framework developed by Apple Computer Inc., capable of handling various formats of digital video, picture, sound, panoramic images, and interactivity. It is available for Mac OS classic (System 7 and later), Mac OS X and Microsoft Windows operating systems.	v8.0
RTCP	A sister protocol of the Real-time Transport Protocol (RTP). RTCP provides out-of-band control information for an RTP flow.	v6.5
RTMP	A protocol commonly used for streaming flash video.	v6.5
RTP	A protocol primarily used to deliver real-time audio and video.	v6.5
RTSP	A protocol used for establishing and controlling media sessions between end points.	v6.5
RTSPS	A secure network control protocol designed for use in entertainment and communications systems to control streaming media servers. The protocol is used for establishing and controlling media sessions between end points.	v8.0
Shockwave	A multimedia platform used to add animation and interactivity to Web pages.	v6.5
SHOUTcast	A cross-platform proprietary protocol for streaming media over the Internet.	v6.5
SIP	A common control protocol for setting up and controlling voice and video calls.	v6.5
SoundCloud	An online audio distribution platform.	v8.5
T-Mobile	A carrier of VoIP services.	v8.0
UStream	A live, interactive broadcast platform that enables anyone with an Internet connection and a camera to create webcasts.	v8.0

Streaming Media Applications	Description	First Available In
Videobb	A video sharing Web site.	v8.5
Vonage	A VoIP company that provides telephone service over a broadband connection.	v6.5
Windows Media	A multi-media player and library application by Microsoft. Users can play audio and video, view images, burn recordable discs with music or data, and purchase music from a number of online music stores.	v6.5
YouTube	A video-sharing Web site on which users can upload, share, and view videos.	v6.5
VPN & Tunneling Applications	Description	First Available In
AH	A member of the IPSec protocol suite that guarantees connectionless integrity and data origin authentication of IP packets.	v8.0
BEETPH	A mode for IPSec ESP that augments the existing ESP tunnel and transport modes. For end-to-end tunnels, the mode provides limited tunnel mode semantics without the regular tunnel mode overhead. The mode is intended to support new uses of ESP, including mobility and multiaddress multihoming.	v8.0
DynGate	A firewall router that allows TeamViewer to route a TCP/IP connection over an HTTP tunnel.	v8.0
ESP	A member of the IPSec protocol suite that provides origin authenticity, integrity, and confidentiality protection of packets.	v8.0
GRE	A tunneling protocol developed by Cisco Systems that can encapsulate a wide variety of network layer protocol packet types inside IP tunnels, creating a virtual point-to-point link to various brands of routers at remote points over an Internet Protocol (IP) internetwork.	v8.0
Hamachi	A hosted VPN service that lets you securely extend LAN-like networks to distributed teams.	v8.0
IPComp	A low-level compression protocol for IP datagrams defined in RFC 3173.	v8.0
IPIP	An IP tunneling protocol that encapsulates one IP packet in another IP packet.	v8.0
IPSec	An end-to-end security scheme commonly used for VPNs.	v6.5
L2TP	A tunneling protocol used to support virtual private networks (VPNs). It does not provide any encryption or confidentiality by itself; it relies on an encryption protocol that it passes within the tunnel to provide privacy.	v8.0
OpenVPN	A free and open-source software application that implements virtual private network (VPN) techniques for creating secure point-to-point or site-to-site connections.	v6.5

VPN & Tunneling Applications	Description	First Available In
PPTP	A method for implementing virtual private networks. PPTP uses a control channel over TCP and a GRE tunnel operating to encapsulate PPP packets.	v8.0
RSVP Tunnel	A new RSVP-based tunnel protocol establishes packet tunnels between a tunnel source point (TSP) and a tunnel destination point (TDP) in such a way that guaranteed services to aggregated packet flows is provided.	v8.0
Web Service Applications	Description	First Available In
12306.cn	The official web of The Ministry of Railways of The People's Republic of China.	v8.5
2345COM	A Chinese web navigation site.	v8.5
39.net	A leading Chinese health portal.	v8.5
About	A source for original information and advice.	v8.5
Adobe	The Adobe applications and updates.	v8.0
Acrobat	An Adobe Web site that provides online PDF services.	v8.5
Adfly	A general Web site browsing and URL shortening service.	v8.5
Adrive	An online cloud storage service.	v8.5
Admin5	A general Web site browsing of Chinese webmaster information.	v8.5
Aizhan	A Chinese Web site that assists webmasters.	v8.5
Amazon	An online retail service.	v6.5
Answers	An Internet-based knowledge exchange.	v8.5
Apple	The Apple Web site.	v8.5
Archive	A non-profit digital library.	v8.5
Atom	Atom is a web content syndication system similar to RSS.	v8.0
Audio	Any audio files or streams delivered over HTTP that were not detected as being part of a more specific application.	v8.0
Barnes&Noble	A Web site that offers books, ebooks, DVDs, music, and so on.	v8.5
Bildde	A German tabloid Web site.	v8.5
Bing	A search engine by Microsoft. Bing can be integrated with Hotmail or Facebook to message friends Bing search results, or with Facebook to have the option to send messages to friends in the search results.	v6.5
Bingbot	A Microsoft Web crawler for the Bing search engine.	v8.0
Blogger	A blog publishing service owned by Google, formerly known as BlogSpot.	v8.0
Bloomberg	A general business news and financial information Web site.	v8.5
Brighttalk	An online webinar and video provider.	v8.5
Brothersoft	A free software download Web site.	v8.5

Web Service Applications	Description	First Available In
CBS	The Web site for the American commercial broadcasting service.	v8.5
CBSinteractive	An online content network for news, sports, entertainment, technology and business.	v8.5
CNET	An online forum for technology news, reviews, teaching videos, product pricing, free downloads, and newsletters.	v8.5
CNN	An American online news site.	v8.5
Conduit	An online platform that allows Web publishers to create free custom toolbars, Web applications, and mobile applications.	v8.5
Craigslist	A Web site providing local classifieds and forums, moderated by the community and largely for free.	v6.5
Dell	The Web site for Dell Inc.	v8.5
Domaintools.com	An Internet domain name intelligence service.	v8.5
DynamicIntranet	An on-demand Web site for intranet applications.	v8.5
eBay	An online auction and shopping Web site.	v6.5
Enet.com.cn	A Chinese IT information portal.	v8.5
Envato	An Australian-based company that provides digital marketplaces for images, templates, project files, and education tutorials.	v8.5
Facebook Search	The Facebook Web site search.	v8.5
Flickr	An image-hosting and video-hosting site, including a Web services suite, and online community.	v8.0
Fogbugz	A hosted bug tracking software and project management system by Fog Creek. Users can manage, filter, sort, and navigate tasks related to a particular issue.	v6.5
Forbes	A national American online business magazine.	v8.5
FOX Sports	A FOX sports and affiliates providing sports news, images, blogs, and videos.	v8.5
Free	A French Internet service provider for general browsing on the Web site free.fr.	v8.5
Google	The traffic generated by the Google search engine or one of the other Internet services provided by Google, Inc.	v6.5
Google Analytics	A Google service that tracks and generates detailed web statistics.	v8.0
Google APIs	The application programming interfaces that support the development of Web applications that leverage Google services.	v8.0
Google App Engine	An application that enables users to build and host Web applications that run on Google's application infrastructure, eliminating the need for additional hardware, patches, and back-ups.	v8.0
Google Calendar	A free time-management Web application offered by Google.	v6.5
Google Desktop	A desktop search and desktop gadget software.	
Google Docs	A free, Web-based word processor, spreadsheet, presentation, form, and data storage service offered by Google.	v6.5



Web Service Applications	Description	First Available In
Google Earth	A Google virtual globe, map, and geographical information program.	v6.5
Google Maps	A Google Web mapping services application and technology.	v8.0
Google Safe Browsing	A Google service that protects users from known phishing and malware sites.	v8.0
Google Translate	A Google Web service that provides language translation for web pages and text for any selected language.	v8.0
Googlebot	A Googlebot searches Web sites for new pages and updated content and adds it to the Google index.	v8.0
Goo.ne.jp	A Japanese Web portal.	v8.5
Gopher	A TCP/IP application layer protocol designed for distributing, searching, and retrieving documents over the Internet.	v8.0
Groupon	A dial-of-the-day Web site that features discounted gift certificates.	v8.5
HP Website	The Hewlett-Packard Web site that provides company news and information.	v8.5
HTTP	The HyperText Transfer Protocol, the principal transport protocol for the World Wide Web.	v6.5
HTTP Audio	Any audio files or streams delivered over HTTP that were not detected as being part of a more specific application.	v8.0
HTTP Video	Any video files or streams delivered over HTTP that were not detected as being part of a more specific application.	v8.0
Hupu	A Chinese sports news Web site.	v8.5
IBM	The International Business Machines corporate Web site.	v8.5
Ikea	A Scandinavian furniture and accessories online store.	v8.5
Image-Venue	A free image hosting and uploading Web site for bloggers, message board users, and eBay sellers.	v8.5
Imgur	A free image-sharing application that is a hosted service.	v8.0
In.com	An Indian Web portal where users have access to news, blogs, feeds, streaming music and video, and mail and classified service.	v8.5
InfoSeek	A popular search engine originally operated by the Infoseek Corporation. Infoseek was bought by The Walt Disney Company in 1998, and the technology was merged with that of the Disney-acquired Starwave to form the Go.com network. It has been replaced with Yahoo! search and is no longer in use.	v8.0
Isohunt	A bit torrent and peer-to-peer index and search engine for finding fully categorized BitTorrent files.	v8.5
IT168	A Chinese social media site.	v8.5
Lebncoin	An online sales site, that allows the publication of advertisements for sale of any items, accessible by everyone.	v8.5
LiveJournal	A virtual community where Internet users can keep a blog, journal, or diary.	v8.0
Mozilla	The Mozilla corporation Web site, that includes downloads and updates to Mozilla Firefox.	v8.5

Web Service Applications	Description	First Available In
MS Online	A hosted software suite that includes Exchange Online, SharePoint Online, Office Communications Online, Microsoft Forefront, and Microsoft Office Live Meeting.	v8.0
MSN	A collection of Internet site and services by Microsoft. The MSN portal provides access to Windows Live services such as Messenger, Hotmail, and SkyDrive, as well as News, Sports, Financial and Entertainment Services.	v8.0
Mywebsearch	A search engine site powered by Google.	v8.5
Netease	A popular Chinese portal.	v8.5
Newegg	An online hardware and software retailer.	v8.5
NFL	The national football league Web site.	v8.5
pchome	A Chinese IT information Web site.	v8.5
Photobucket	An image-hosting, video-hosting, slideshow-creation, and photo-sharing Web site.	v8.0
Picasa	A Google image organizer and image viewer for organizing and editing digital photos, plus an integrated photo-sharing Web site.	v8.0
Quote	A financial market information and trading resource site.	v8.5
Raging-Bull	A financial message board hosted by Quote.com.	v8.5
RSS	A Really Simple Syndication feed format for Web feeds.	v8.0
Salesforce	An online CRM and cloud computing service.	v6.5
Samsung	The Web sites associated with South Korean multinational conglomerate company Samsung.	v8.5
Schmedley	An online desktop start page that provides widgets for browsing around the Web.	v8.5
Sina	A Chinese infotainment Web portal.	v8.5
Slideshare	A slide hosting service.	v8.5
Soku	A Chinese search engine from Youku.	v8.5
SSL	A Secure Sockets Layer cryptographic protocol that provides security over the Internet.	v6.5
StumbleUpon	A Web browser plugin that allows users to discover and rate Web pages, photos, videos, and news articles.	v8.0
Squidoo	A community Web site that allows users to create pages, called lenses, about subjects of interest.	v8.5
SurveyMonkey	A private American company that enables users to create their own survey.	v8.5
Target	An American online retail store.	v8.5
TechInline	A Web-based service for remote support, remote control, desktop sharing, remote training, and file transfer between two computers.	v8.5
Theme Forest	An Envato marketplace where users can buy and sell site templates and themes to skin CMS products like WordPress, Drupal, and Joomla.	v8.5

Web Service Applications	Description	First Available In
Tumblr	A microblogging platform that allows users to post text, images, videos, links, quotes, and audio to their tumblelog, a short-form blog.	v8.0
USA Today	An American online daily newspaper Web site.	v8.5
Video	An HTTP video application that delivers video files and streams that were not detected as belonging to a more specific application.	v8.5
W3Schools	A Web site developer's portal.	v8.5
Webs.com	A space for individuals, groups, or small businesses to share photos and videos, open a store, and build a community.	v8.5
Weebly	A free Web site creator that uses a widget-style format, enabling users to create a site by dragging-and-dropping the page elements.	v8.5
Wetpaint	A social publishing and distribution company that focusses on coverage of TV shows, stars, and fashion.	v8.5
Wikidot	A free and paid Wiki hosting site.	v8.5
Wikipedia	An online editable encyclopedia.	v6.5
Wordpress	An online blogging community.	v8.0
Wretch	A general browsing and streaming media Web site from Taiwan.	v8.5
Xanga	A Web site that hosts weblogs, photoblogs, and social networking profiles.	v8.0
Yahoo	An online service.	v6.5
Yahoo Slurp	A Web crawler that obtains content for the Yahoo search engine.	v8.0
Zoho	A Web-based online office suite containing word processing, spreadsheets, presentations, databases, note-taking, wikis, CRM, project management, invoicing, and other applications developed by ZOHOO Corporation.	v6.5



# Index

## Numerics

10 Gigabit Ethernet Fiber NICs 3

## A

A SYN or RST packet contains data 263

Access Control List 207

Account control 200

Account lockout 201

Accounts

    capability-based 197

    privileges 197

    role-based 197

Active peer in a Steelhead failover pair 48

Add a New TCP Dump 332

Administrator password 197

Admission Control alarm status 308

AES 98

AES256 98

AFE

    application signatures recognized 365

    classifying encrypted applications 109

    overview 109

Alarm status

    admission control 308, 309

    asymmetric routing 308

    connection forwarding 309

    disk full 309

    fan error 219

    hardware error 310

    licensing 312

    link duplex error 313

    link I/O error 313

    link state 313

    memory error 313

    memory paging 313

    neighbor incompatibility 313

    network bypass 314

    optimization service 314, 354

    process dump creation 315

    software version mismatch 315, 316

    SSL 316

    temperature 316

Alarm thresholds, setting 217

Analyzer for NetFlow 101

Announcement, setting on home page 224

Appliance failure bypass 48

Application Flow Engine 109

Application signatures list 365

Application Statistics reports 289

Application Visibility reports 291

AS/400 host environments 79

Assignment scheme 164

Asymmetric routing

    alarm status 308

    auto-detection 91

Asymmetric routing pass-through

    reason 262

Authentication

    encryption setting 97

    setting general security 196

Authentication methods

    Local 196

    MD5 98

    SHA-1 98

    TACACS+ 196

Auto-detection of asymmetric routing 91

Auto-discover

    self-signed certificates 172

Auto-discover rules, overview of 37, 209

Auto-discover, in-path rule 37

Automatic peering, see Enhanced auto-discovery 54

## B

Backup Steelhead appliance 48

Bad RST log entry 61

Bandwidth

    oversubscription in QoS 117

    setting for the default QoS site 116

    setting the link rate in Basic QoS 116

    sharing between remote sites 117

Bandwidth estimation algorithm 68

Bandwidth limit reached 266

Bandwidth Optimization report 299

Basic password template 200

Blue Coat Proxy SG 17

Branch services 10

Bypass card naming conventions 50

## C

Cable troubleshooting 324

Capability-based accounts 197

Cascade Pilot 332

Cascade Shark

    enabling 332

    process 229

- Cascading menus
  - displaying and using 10
  - summary of 10
- Certificate Authorities
  - configuring in SSL 169
- Class name in QoS 130, 149
- Class parent in H-QoS 131
- Classification and shaping network traffic. 108, 109
- Client connections, resetting upon startup 49
- Clocks
  - synchronizing client and server 351
- Cloud proxy is down 265
- Collect traffic flow data 101
- Configuration files, managing 194
- Configuration, saving 11
- Congestion control
  - error-tolerant 68
  - for SCPS-capable connections 68
  - methods 68
  - using with rate pacing 70
- Connection
  - resetting 267
  - resetting a single 257
  - resetting globally 49
  - resetting per in-path rule 41
- Connection count above QOS limit 262
- Connection for local host 262
- Connection forwarding 94
- Connection forwarding alarm status 309
- Connection from proxy target 262
- Connection history, viewing 268
- Connection initiated by neighbor 262
- Connection limit reached 266
- Connection paused 261
- Connection pool
  - report, viewing 301
  - setting size for 51
- Connection tracking for pass-through traffic 109
- Console
  - connecting to 8
  - navigating 10
- Continuous log 327
- Correct addressing 43
- CPU settings 58
- CPU utilization
  - alarm status 309
  - report 317
- CSV file, exporting statistics to 339
- Current connections
  - viewing 248
  - viewing details 257
- D**
  - Data reduction levels 82
  - Data replication over GigE 80
  - Data store
    - CPU settings 58
  - Default password protection 200
  - Deny in-path rules, overview of 38, 209
  - Deny privileges 198
  - Diagnostic test
    - Network health 323
  - DIF 77, 81
    - data block size 79, 80
  - DIF headers
    - isolating from SRDF traffic 84
    - within VMAX-to-VMAX traffic 78
  - Discard in-path rules, overview of 38, 209
  - Disk drive failure email notification 226
  - Disk Error 219, 310
  - Disk full alarm status 309
  - Documentation, contacting 6
  - DPI similarities with AFE 109
  - DSCP
    - marking in path selection 153
    - marking in RiOS 136
    - using a global marking 107
    - using in packet-mode optimization 35
    - using to prioritize SnapMirror traffic 107
  - Duplex setting 22
- E**
  - Email notification, setting 224
  - Enabling
    - encryption 97
    - NetFlow 99, 101
    - Virtual in-path on a client-side appliance 99
  - Encapsulation scheme 163
  - Encryption
    - 3DES 98
    - AES 98
    - AES256 98
    - enabling on a Steelhead appliance 97
  - Enhanced auto-discovery 54
  - Enhanced automatic peering 54
  - Enterprise MIB
    - accessing 341
  - Error connecting to server 262
  - Error on neighbor side 263
  - Error on SSL inner channel
    - handshake 263
  - Ethernet compatibility, summary of 3
  - Ethernet network compatibility 3
  - Event and failure notification, setting 224
  - Expired Password page 200
  - Exporter for NetFlow 101
- F**
  - Failed to append CP code 265
  - Failed to cache sock pointer 264
  - Failed to create sport outer channel 264
  - Failed to discover SCPS device 263
  - Failover 48
  - Fan Error 219, 310
  - Fan Error alarm status 219
  - FCIP optimization 77
  - Fibre Channel over TCP/IP 77
  - FIFO queue in QoS 132
  - Filter logs 326
  - Fixed-target rules 37
  - Flash Error 219, 310
  - Flexible licensing 192

- Flow export
  - configuring 101
  - configuring subnet side rules for a collector 100
- Flows not matching in-path rule 264
- FTP channels, setting optimization policies for 40
- FTP proxy access 15
- Full Transparency 43

## G

- Generic Flow error 264
- GigE-based SRDF traffic between VMAX arrays 78
- Global
  - application list in QoS 109
  - automatic kickoff 41, 49
  - DSCP marking 107
- Granite
  - common ports used by the system 359
  - default ports 359
  - port label 75
- Gray list, SSL 178
- Guaranteed bandwidth, in QoS 133

## H

- Half open connections above limit 262
- Half-opened connections
  - restricting 51
- Hardware
  - alarm status 310
  - assist rules 166
  - dependencies 3
- Hash assignment 164
- Health, checking the Steelhead 323
- High availability
  - using mask assignment 164
- Home page announcement, setting 224
- Home Welcome page, overview of 9
- Host labels
  - creating 73
  - overview of 73
  - resolving hostnames 74

## I

- Incompatible probe version 262
- Inner failed to establish 261
- In-path
  - physical, overview of 48
  - setting optimization policies for FTP channels 40
  - virtual, overview of 48
- In-path rule pass-through reason 261
- In-path rules
  - auto-discover 37
  - configuring 33
  - fixed-target 37
  - pass-through 38
  - type 37
  - VLAN identification number 67
- Installing
  - license 192
  - SSL license 170

- Interactive ports
  - forwarding traffic on 35, 75
  - list of 361
- Interface naming conventions 50
- Interface statistics, viewing 294
- Invalid customer ID 266
- Invalid Entitlement code 266
- Invalid ESH ID 266
- Invalid SaaS ID 266
- Invalid timestamp 266
- IPMI error 219, 310
- IPSec encryption
  - enabling 97
- IPv6
  - static routing table 28
  - support summary 18
- IPv6 connection forwarding requires multi-interface support 265

## J

- JavaScript 3
- Job scheduling 188
- Jobs, viewing details 188
- Jumbo frames 4

## K

- Keep-alive
  - for an optimized connection 267
- Kickoff
  - reset an existing pass-through or optimized connection 257
  - reset existing client connections upon startup 49
  - reset existing connections that match an in-path rule 41
- Known issues 5

## L

- Labelling traffic in reports 234
- LAN port trace 93
- Layer-4 switch support, overview of 50
- LFN, optimizing traffic for 143
- Licenses
  - fetching automatically 192
  - managing 191
- Licensing alarm status 312
- Link duplex alarm status 313
- Link I/O error alarm status 313
- Link rate
  - imposing a global data transmit limit 70
- Link share weight 134, 150
- Link state alarm status 313
- Link State Propagation 25
- Local logging, setting 226, 230
- Logs
  - customizing 327
  - downloading 328
  - filtering 326
  - viewing 326
  - viewing continuous 328
- Long fat networks, optimizing traffic for 143

**M**

- Management ACL 207
- Management Console
  - navigation of 10
  - overview of 7
- Management In-Path (MIP) interface 28
- Mask assignment 164
- Master Steelhead appliance 48
- McAfee Web Gateway 17
- Memory error 219, 310
- Memory error alarm status 313
- Memory Paging
  - alarm status 313
  - reports 319
- Message of the day
  - See MOTD
- MIB file
  - accessing 341
  - SNMP traps sent 343
- Microsoft WebSense 17
- Middle Steelhead 262
- MIP interface 28
- Mobile trust 178
- Monitor password 197
- Monitored ports, setting 234
- MOTD, setting 224
- MTU value, setting 22, 23, 27
- Multicore balancing in data store 58
- MX-TCP 143
  - basic steps for 143
  - queue in QoS 132

**N**

- Neighbor does not support IPv6 265
- Neighbor Incompatibility alarm status 313
- Neighbor Statistics report 271
- NetApp ONTAP 85
- NetFlow
  - adding a collector 105
  - enabling 99, 101
  - in-path deployment 101
  - troubleshooting 106
- Netflow only
  - probe packet of optimized connection 265
- Network Bypass alarm status 314
- Network status report 323
- No matching client/server IPv6 scope 264
- No proxy port for probe response 262
- No room for more TCP options 261
- No route for probe response 261
- No Steelhead on path to server 261
- Not a supported SaaS destination 265
- NTP 177, 351

**O**

- Object identifiers
  - viewing through SNMP 240
- Online documentation 5
- Online help 12
- Open System environments 79
- Optimization
  - controlling with peering rules 54
  - FCIP 77
  - service alarm status 314, 354

- SnapMirror 85
- SRDF 80
- optimization service, restarting 11
- Optimized connection 265
- Optimized Throughput report 296
- Optimizing local connections only 265
- Other hardware error 219, 310
- Out of memory 261
- Out-of-Band connection 48
- Out-of-path
  - adding static routes to deployments 18
  - overview of 48
- Overview
  - of asymmetric routing auto-detection 91
  - of port labels 75
  - of the Management Console 7

**P**

- Packet-mode channel setup pending 264
- Packet-mode optimization 34
- Packet-mode optimization disabled 264
- Packet-order queue 133
- Passing to downstream Steelhead 266
- Pass-through
  - enable for traffic on interactive ports 75
  - IPsec encryption 97
  - QoS 110
  - reasons 260, 265
  - resetting a connection 267
  - secure inner channel connections 182
  - traffic on secure ports 75
  - traffic on system ports 75
- Passthrough SYN retransmit 266
- Password policy 200
- Path selection
  - configuring 154
  - limits 158
  - monitoring the path state 153
  - overview 153
  - path down alarm 156
  - use cases 156
- PBR, overview of 50
- PCoIP 107, 108
- Peer does not support packet-mode optimization 264
- Peer in fixed-target rule down 261
- Peer IPv6 Incompatible 309
- Peering
  - Automatic 53
  - regular auto-discovery 53
- Peering rule pass-through reason 261
- Peering rules 53, 54
- PEP device 60
- Performance page 58
- Perl regular expression 336
- Permissions 197
- Physical in-path, overview of 78
- Port label handling in QoS 109
- Port labels overview of 75
- Port Transparency 43
- Ports
  - commonly excluded 360
  - default Granite 359



- default listening 359, 360
- interactive ports forwarded 361
- secure automatically forwarded 361
- Power supply error 219, 310
- Preexisting connection 261
- Primary gateway IP address 21
- Primary interface
  - on the Steelhead appliance 18
  - setting 21
- Printing pages and reports 11
- Priorities, QoS 110
- Privileges
  - deny 198
  - read 198
  - write 198
- Process dump creation alarm status 315
- Process dumps, displaying and
  - downloading 330
- Professional services, contacting 6
- Protect access to a Steelhead 207
- Protocol
  - SRDF 80
  - SSL, basic steps for 171
  - SSL, verifying optimization for 173
- Proxy
  - addresses for Web access 15
  - certificate for SSL back-end server 171
  - setting an IP for Web/FTP 17

## Q

### QoS

- basic steps for 129, 147
- class name 130, 149
- configuring inbound 144
- enabling on a WAN interface 116
- FIFO queue 132
- Inbound QoS report 275
- LAN bypass 113
- latency priority 133, 149
- maximum classes by Steelhead
  - appliance 130, 148
- MX-TCP enabling 143
- MX-TCP queue 132
- Outbound QoS 273
- priorities 110
- service class 118, 123, 139
- service ports for multiple mappings 72
- setting the interface bandwidth link
  - rate 116, 129, 148
- SFQ queue 132
- Statistics report 281
- QoS global application list 365
- QoS policies, port transparency 43
- Queue
  - capture file 333
  - FIFO 132
  - in advanced outbound QoS 111
  - MX-TCP 143
  - packet-order 133
  - SFQ 132
  - specifying the trace dump size 334
  - using to prioritize TCP/IP traffic 143
  - viewing messages in Intercept
    - module 323

## R

- RADIUS authentication method,
  - setting 196
- Rate pacing
  - configuring an MX-TCP QoS rule 70
  - setting a global data transmit limit on the link rate 70
- RBT-Proto, common ports used by the system 360
- RDF groups
  - setting custom data reduction levels for 82
  - specifying the number 83
- Reached maximum TTL 262
- Read privileges 198
- Reboot 191
- Redundancy 48
- Rejected by cloud proxy 266
- Related reading 4, 6
- Release notes 4
- Remote data facility (RDF) groups 82
- Remote packet capture analysis 332
- Reports
  - Application Statistics 289
  - Application Visibility 291
  - Bandwidth Optimization 299
  - Connection History 268
  - Connection Pooling 301
  - CPU Utilization 317
  - Current Connections 248
  - Current Connections Details 257
  - Export Performance Statistics 339
  - Health Status 323
  - Inbound QoS 275
  - Interface Statistics 294
  - Memory Paging 319
  - Neighbor Statistics 271
  - New format in v8.0 244
  - Optimized Throughput 296
  - Outbound QoS 273
  - QoS Statistics 281
  - SnapMirror 304
  - SRDF 302
  - TCP Memory 320
  - TCP Statistics 295
  - TCP trace dump 331
  - Top Talkers 281
  - Traffic Summary 284
  - WAN Throughput 287
- Reset
  - an individual connection 257
  - existing client connections at start up 49
  - existing client connections matching an in-path rule 41
- Restarting the optimization service 187
- Restrict inbound IP access to a Steelhead appliance 207
- Revert to a backup version 189
- Ricochet or probe 263
- RiOS and SCPS connection 60
- Riverbed Application Flow Engine 109
- Riverbed, contacting 6
- Role-based
  - accounts 197

- user permissions 197
- Routing
  - asymmetric, auto-detection of 91
  - enabling simplified 159
- RX probe from failover buddy 262
- S**
- Satellite
  - compatibility 59
  - WAN optimization 59
- Scanning system logs 336
- SCPS
  - compatibility 59
  - initiating a connection 67
  - role-based permission 59
  - rules 61
- SCPS-only connection 60, 250
- SDRF 49
  - setting a custom data reduction level 82
- Secure access by inbound IP address 207
- Secure ports
  - automatically forwarded 361
  - forwarding traffic on 35, 75
- Secure vault
  - unlocking and changing the password 206
- Securing passwords 200
- SEI 60
- Serial cascade deployment 53
- Serial cluster deployment 53
- Service ports, setting 72
- Services, starting, stopping, restarting 187
- Setting
  - alarm thresholds 217
  - email notification 224
  - event notification 224
  - failure notification 224
  - local logging 226, 230
  - monitored ports 234
  - SNMP trap receivers 235
- SFQ queue in QoS 132
- Shark
  - enabling 332
  - process logging 229
- Shut down 191
- Simplified routing, enabling 159
- Single-Ended Interception 60
  - interaction with in-path rules 61
  - network asymmetry 61
- SkipWare 59
- SnapMirror
  - optimization 85
  - report 304
- SnapMirror replication 85
- SNMP
  - access policies 240
  - access policy security 236, 241
  - adding groups 239
  - adding trap receivers 235
  - adding views 240
  - compatibility 4
  - creating users 237
  - including specific OIDs in a view 240
  - MIB, accessing 341
  - supported versions 235

- testing a trap 237
- traps, summary of sent 343
- SNMP compatibility, summary of 4
- Software dependencies, overview of 3
- Software version mismatch alarm
  - status 315, 316
- Software, upgrading 189
- Specification license 192
- Speed and duplex
  - avoiding a mismatch 22
  - setting 22
- Squid proxy 17
- SRDF
  - creating rules 84
  - report 302
  - viewing connections 82
- SRDF/A
  - optimization 80
- SSL
  - basic steps for 171
  - black list 177
  - cipher settings 184
  - gray list 178
  - peering list 178, 211
  - trusted entities 181
  - verifying 173
- SSL Alarm status report 316
- Storage optimization 77
- Strong password template 200
- Subnet
  - for aux interface 18
- Subnet side rules, configuring 99
- Symmetrix array 81
- SYN before SFE outer completes 262
- SYN on WAN side 261
- SYN retransmit (backhauded) 265
- SYN retransmit (direct) 266
- SYN/ACK at MFE not SFE 263
- SYN/ACK, but no SYN 263
- Synchronizing
  - clocks 177
  - Steelhead appliance to NTP server 231
- System details
  - alarm 316
  - report 323
- System is heavily loaded 263
- System snapshot report 323
- System, logging out of 12

- T**
- TACACS+ authentication method,
  - setting 196
- TCP dump 332
  - with Cascade Shark 332
- TCP dumps, displaying 331
- TCP passthrough 166
- TCP proxy only connection 61
- TCP statistics report, viewing 295
- TCP trace dump 331
- TCP-over-IPv6 34
- TCP-PEP device 60, 64, 70, 250
- Technical Publications, contacting 6
- Technical support, contacting 6
- Temperature alarm status 316

- Time zone
  - setting 230
- Too many retransmitted SYNs 262
- Top Talkers report 281
- Traffic Summary report 284
- Transparency packet from self 263
- Transparent addressing 43
- Transparent inner not for this host 263
- Transparent inner on wrong VLAN 262
- Transparent RST to reset firewall state 263
- Transport Settings page 61, 66
- Traps, summary of SNMP traps sent 343
- Troubleshooting
  - asymmetric routes 93
  - cables 324
  - flow export settings 106
  - gateway connection 324
  - IP port reachability 325
  - peer reachability 325
  - scanning system logs 336
  - speed and duplex settings 324
  - Steelhead network health 323
- Trust
  - between peer appliances 178
  - between Steelhead Mobile Controller entities 178
  - establishing for secure inner channel 176
- U**
- UDP passthrough 166
- Unknown reason 262
- Upgrading
  - software 189
- Uplink, sharing with remote sites 117
- User logs
  - downloading 328
  - viewing 326
- User permissions 197, 198, 200
- V**
- Vault
  - unlocking and changing the password 206
- View-Based Access Control Mechanism 235
- Virtual in-path on a client-side appliance, enabling 99
- Virtual in-path, overview of 48
- VLAN
  - identification number 40, 67
  - preserving tags 43
  - using tags with in-path rules 34
- VMAX array 77, 81
- VMAX-to-VMAX traffic 78, 84
- W**
- WAN
  - calculating buffer size for high-speed TCP 71
  - sharing available bandwidth between remote sites 117
  - top bandwidth consumers 102, 281
  - visibility modes 43
- WAN path selection 153
- WAN Throughput report 287
- WAN uplink speed, sharing with remote sites 117
- WCCP
  - multiple Steelhead interfaces 161
  - service groups 160
  - service groups, modifying settings 166
- Web proxy access to the Steelhead appliance 17
- Web/FTP proxy 17
- White list, SSL 177
- Windows branch mode detected 263
- Write privileges 198