



SteelCentral Controller™ for SteelHead™ Deployment Guide

Version 9.5

December 2016



© 2017 Riverbed Technology, Inc. All rights reserved.

Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed. 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 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, Inc. in the United States and in other countries.

This product includes Windows Azure Linux Agent developed by the Microsoft Corporation (<http://www.microsoft.com/>). Copyright 2016 Microsoft Corporation.

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.

The SteelHead Mobile Controller (virtual edition) includes VMware Tools. Portions Copyright © 1998-2016 VMware, Inc. All Rights Reserved.

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. This documentation may not be copied, modified or distributed without the express authorization of Riverbed 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 assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.



Riverbed Technology
680 Folsom Street
San Francisco, CA 94107
www.riverbed.com

Part Number
712-0086-06

Contents

Welcome	7
About this guide	7
Audience	7
Document Conventions	8
Documentation and release notes	8
Contacting Riverbed	9
What's new	9
 1 - Overview of the SCC.....	 11
Overview of the SCC	11
Overview of the SCC-VE	12
AppResponse support	12
 2 - Registration and Configuration	 13
SCC registration and configuration	13
Setting up an appliance for SCC management.....	13
Manual registration	13
Automatic registration	14
HTTPS communication channel.....	14
Automatic configuration.....	16
Locally managed SteelHeads	16
Enabling local changes on a SteelHead	17
 3 - Groups and Policies	 19
Groups.....	19
Policies	20
Policy overview	20
Inherited policies	20
Policy settings.....	24
Pushing policy	25

4 - SCC and Appliance Operations	29
Backing up and restoring an appliance	30
Migrating SteelHead settings with the SCC	33
SCC backups and high availability	34
Using the secure vault	36
REST API access	37
Using management ACLs	37
Managing licenses using the SCC	38
Using the SCC for Interceptor management	39
SCC bandwidth consumption	40
Application statistics	40
Using the SCC upgrade and downgrade appliance wizard	43
Configuring and Using the SCC as a Certificate Authority Service	45
Using the SCC CA service to issue proxy certificates for SSL optimization	55
5 - Sites and Networks, Path Selection, and QoS	61
Overview of sites and networks, path selection, and QoS	61
Configuring sites and networks	62
Configuring path selection	69
Configuring global QoS	73
6 - Web Proxy	79
Overview of the web proxy feature	79
Web proxy fundamental properties	80
Supported features	80
IP addressing support	81
TCP port support	81
Web proxy and SteelHead SaaS	81
Video caching	81
Configuring basic web proxy features (HTTP)	82
Advanced configurations for web proxy	83
Configuring web proxy for HTTPS	84
SSL decryption and TCP proxy for HTTPS	84
Using web proxy and certificate management	84
Using the global whitelist	87
Using parent proxy (proxy chaining) configurations	88
Using in-path rules	91
Troubleshooting web proxy	93
SCC to SteelHead communications	93
HTTP caching	93

HTTPS decryption.....	94
YouTube video caching	94
7 - Secure Transport	97
Overview of secure transport.....	98
Management plane	100
Control plane.....	101
Data plane.....	101
Secure transport sizing	102
Group sizing	102
Group member sizing	102
Controller sizing	103
Functional operations overview	103
Firewall considerations	104
Network interface used for SteelHead-to-controller communication	104
Maximum segment size.....	104
Network address translation	105
Encryption key operations	107
Disconnected mode.....	108
Fail closed.....	108
Secure transport concentrator	108
Secure transport configuration workflow	109
Deployment example of a hybrid network backhaul scenario	111
Deployment example of a dual in-path interface with split tunnel.....	119
Reporting	127
8 - Appliance Clusters	129
Overview of appliance clusters	129
Configuring the SCC to manage appliance clusters.....	129

Welcome

Welcome to the *SteelCentral Controller for SteelHead Deployment Guide*. This guide describes how to deploy the SteelCentral Controller for SteelHead (SCC).

About this guide

The *SteelCentral Controller for SteelHead Deployment Guide* describes why and how to configure the SteelCentral Controller for SteelHead.

This guide includes information relevant to the following products and product features:

- Riverbed SteelCentral Controller for SteelHead (SCC)
- Riverbed SteelCentral Controller for SteelHead (virtual edition) (SCC-VE)
- Riverbed Optimization System (RiOS)
- Riverbed SteelHead (SteelHead)
- Riverbed SteelHead CX (SteelHead CX)
- Riverbed SteelHead EX (SteelHead EX)
- Riverbed Virtual SteelHead (SteelHead-v)
- Riverbed Cloud SteelHead (SteelHead-c)
- Riverbed SteelCentral Controller for SteelHead Mobile (Mobile Controller)
- Riverbed SteelHead Mobile (SteelHead Mobile)
- Riverbed SteelHead Interceptor (Interceptor)
- Riverbed Virtual Services Platform (VSP)
- Riverbed SteelFusion Core (Core)
- Riverbed SteelFusion Edge (Edge)
- SteelCentral AppResponse (AppResponse)

Audience

This guide is written for storage and network administrators familiar with administering and managing WANs using common network protocols.

You must also be familiar with:

- the SCC Management Console. For details, see the *SteelCentral Controller for SteelHead User's Guide*, the *SteelCentral Controller for SteelHead Installation Guide*, and the *SteelCentral Controller for SteelHead (Virtual Edition) Installation Guide*.
- the SteelHead Management Console. For details, see the *SteelHead Management Console User's Guide*, the *SteelHead Deployment Guide*, and the *SteelHead Deployment Guide - Protocols*.
- connecting to the RiOS CLI. For details, see the *Riverbed Command-Line Interface Reference Manual*.
- the installation and configuration process for the SteelHead. For details, see the *SteelHead Installation and Configuration Guide* and the *SteelHead (Virtual Edition) Installation Guide*.
- the SteelHead Interceptor. For details, see the *SteelHead Interceptor User's Guide* and the *SteelHead Interceptor Deployment Guide*.
- the Mobile Controller. For details, see the *SteelCentral Controller for SteelHead Mobile User's Guide*.

Document Conventions

Convention	Meaning
<i>italics</i>	Within text, new terms and emphasized words appear in <i>italic</i> typeface.
boldface	Within text, CLI commands, CLI parameters, and REST API properties appear in bold typeface.
Courier	Code examples appear in Courier font: <pre>amnesiac > enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: interface <ip-address>
[]	Optional keywords or variables appear in brackets: ntp peer <ip-address> [version <number>]
{ }	Elements that are part of a required choice appear in braces: {<interface-name> ascii <string> hex <string>}
	The pipe symbol separates alternative, mutually exclusive elements of a choice. The pipe symbol is used in conjunction with braces or brackets; the braces or brackets group the choices and identify them as required or optional: {delete <filename> upload <filename>}

Documentation and release notes

To obtain the most current version of all Riverbed documentation, go to the Riverbed Support site at <https://support.riverbed.com>.

If you need more information, see the Riverbed Knowledge Base for any 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>.

Each software release includes release notes. The release notes identify new features in the software as well as known and fixed problems. To obtain the most current version of the release notes, go to the Software and Documentation section of the Riverbed Support site at <https://support.riverbed.com>.

Examine the release notes before you begin the installation and configuration process.

Contacting Riverbed

This section describes how to contact departments within Riverbed.

- **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 or go to <http://www.riverbed.com/services/index.html>.
- **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.

What's new

Since the last release of the *SteelCentral Controller for SteelHead Deployment Guide*, the following changes have been made:

- **Updated** - “Configuring and Using the SCC as a Certificate Authority Service” on page 45
- **Updated** - “Using parent proxy (proxy chaining) configurations” on page 88

Overview of the SCC

This chapter describes the SCC and includes the following sections:

- “Overview of the SCC” on page 11
- “Overview of the SCC-VE” on page 12
- “AppResponse support” on page 12

Overview of the SCC

The SCC facilitates administration tasks for groups of SteelHeads, Interceptors, Mobile Controller, Cores, and Edges. The SCC is designed to work in a large network with many devices. The SCC allows you to manage the following aspects of the appliances:

- **Configuration** - You can automatically configure SteelHead and Interceptors in your network. The SCC uses policies and groups to facilitate centralized configuration and reporting. The SCC enables consistent configuration of different types of appliances from a centralized management console. You can also perform automatic configuration backup and rollback for managed appliances.
- **Monitoring** - The SCC provides both high-level status and detailed statistics about the performance of appliances and the SCC itself.

The SCC collects statistics from SteelHeads every five minutes, and it aggregates these statistics for every hour and every day. The SCC stores the five-minute data points for a maximum of 30 days, the one-hour data points for a maximum of 90 days, and the one-day data points for a maximum of three years. The SCC reports provide a view into the performance of the optimization network, either from an aggregated perspective (groups) or per individual element.

You can also configure event notification for managed appliances.

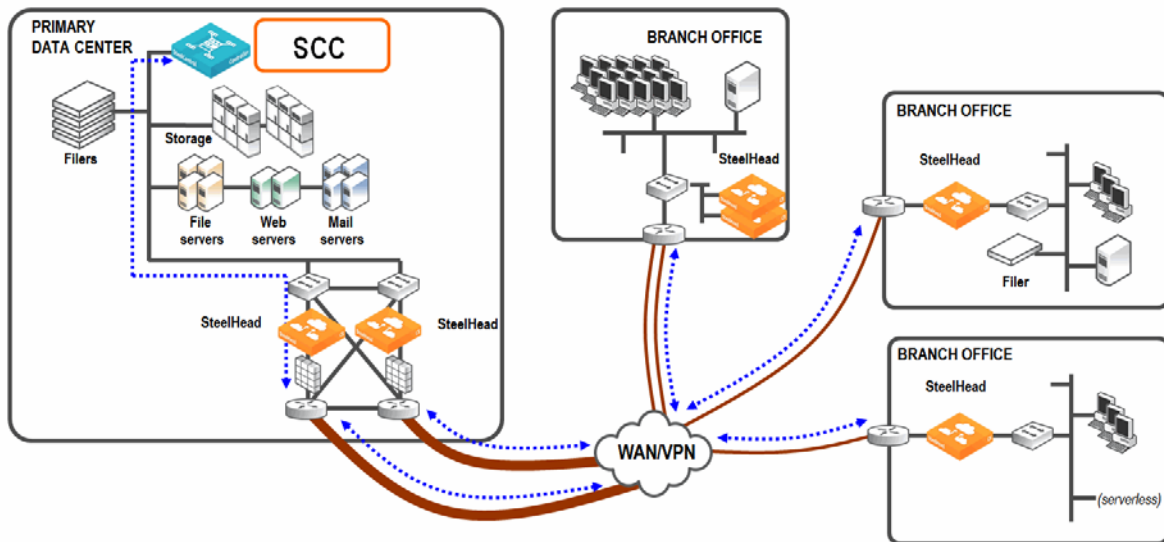
- **Maintenance** - The SCC supports maintenance tasks for SteelHeads, Interceptors, Mobile Controllers, and Cores, such as image update and appliance reboot.
- **Troubleshooting** - SCC 8.5 or later can generate a system dump on the managed SteelHeads and upload to a case number or a URL. You can also generate TCP dump files in the managed appliances and upload the file to a case number.

For details on all SCC features, see the *SteelCentral Controller for SteelHead User's Guide*.

For information on how many SteelHeads an SCC can manage, see the *SteelCentral Controller for SteelHead User's Guide* and <https://supportkb.riverbed.com/support/index?page=content&id=S14106>.

The following illustration shows a typical SCC deployment.

Figure 1-1. SCC deployment



Overview of the SCC-VE

You can also install a SCC-VE, which is the virtual edition of the SCC. The SCC-VE installs and runs on VMware ESXi 5.5.

For detailed information about the SCC-VE, see the *SteelCentral Controller for SteelHead (Virtual Edition) Installation Guide*.

AppResponse support

You can use SCC 9.0 to configure SteelFlow Web Transaction Analysis (WTA) between the SteelHeads running RiOS 9.0 or later and the AppResponse running 9.5 or later.

This integration enables you to monitor and troubleshoot the end-user experience for SteelHead optimized and nonoptimized web and SaaS applications, including Salesforce.com. AppResponse provides the locations of breakdowns in your network, application delays for a web page, and the individual web objects that make up the page. With this information, you can pinpoint whether the SaaS server or your network is the root cause of a performance problem. You can also monitor page time, server busy time, and network busy time for optimized SaaS applications to make sure you're delivering the best end-user experience.

For more information about AppResponse, see the *SteelCentral Controller for SteelHead User's Guide* and the AppResponse documentation listed on the Riverbed Support site at <https://support.riverbed.com/content/support/software/steelcentral-npm/appresponse-appliance.html>.

Registration and Configuration

This chapter describes the basic management operations for the SCC. It includes the following sections:

- [“SCC registration and configuration” on page 13](#)
- [“Locally managed SteelHeads” on page 16](#)
- [“Enabling local changes on a SteelHead” on page 17](#)

SCC registration and configuration

This section describes how to register and configure SteelHeads, Interceptors, Mobile Controller, and Cores with the SCC. This section includes the following topics:

- [“Setting up an appliance for SCC management” on page 13](#)
- [“Manual registration” on page 13](#)
- [“Automatic registration” on page 14](#)
- [“HTTPS communication channel” on page 14](#)
- [“Automatic configuration” on page 16](#)

Setting up an appliance for SCC management

All SteelHeads, Interceptors, Mobile Controllers, Cores, and Edges are shipped with SCC management and automatic registration enabled.

To disable automatic registration, see the *SteelCentral Controller for SteelHead User’s Guide* or the *SteelHead Management Console User’s Guide*.

Manual registration

To complete registration from the SCC, enter the following information about the appliance into the Appliances page of the SCC:

- Serial number
- IP address
- Username

■ Password

With this information, the SCC can connect to and manage the SteelHead. If you enable Common Administration Login, the SCC uses a single username and password to manage all appliances. The appliance-specific username and password are ignored. A single username and password is useful if the account used by the SCC to authenticate is the same across the organization or the SCC account is authenticated through TACACS+ or RADIUS.

For more information about using a single username and password, see the *SteelCentral Controller for SteelHead User's Guide*. For details on automatic registration, which doesn't require the SteelHead serial number, see ["Automatic registration" on page 14](#).

Automatic registration

Automatic registration allows a new appliance installed in an existing SCC-managed environment to automatically contact the SCC and register itself. Automatic registration is useful in environments in which you want to ship an appliance to a remote office and have a nonadministrator operator connect the appliance into the network.

To use automatic registration, prepare the appliance for registration with the SCC by entering the IP address or DNS-registered hostname of the SCC when you run the configuration wizard on the appliance. We recommend that you set up your DNS servers to resolve `riverbedcmc` to the primary interface IP address of your SCC. If you configure your DHCP servers to forward the correct DNS settings, any appliance added into the network and powered on should automatically register with the SCC.

Alternatively, you can use the **cmc enable** and **cmc hostname** commands. If you do not change the default hostname, it remains `riverbedcmc`.

If the appliance password has changed from the default, the SCC can automatically register the appliance, but it can't connect to nor manage it. The appliance appears as disconnected on the Appliances page. To continue the automatic registration process, from the Edit Appliance page, select Edit Appliance and enter the correct password.

If automatic registration is not working, use the **show cmc** command to verify that automatic registration is enabled. If the **show cmc** command indicates that the hostname is correct, but autoregistration is not working, check the SteelHead name server search-path configuration using the **show host** command. Use the **cmc hostname <hostname>** command to change the SCC hostname (by default, `riverbedcmc` in DNS).

Registration doesn't change the appliance configuration. It allows the SCC to communicate with the appliance. Configuration of the appliance is completed either automatically or manually through a policy push.

For more details on automatic registration, see the *SteelCentral Controller for SteelHead User's Guide*.

HTTPS communication channel

In SCC 9.0 and later and SteelHeads running RiOS 9.0 and later, there's a second, full duplex-persistent communication channel established between the SteelHead and the SCC. This channel is:

- HTTPS based.

- initiated by the SteelHead to the SCC.
- established using a certificate-based authentication mechanism.

- independent of and coexists with the current SSH-based communication channel.
- used to manage the new features introduced with RiOS 9.0—specifically, path selection, QoS, and secure transport.

Note: For information about upgrading to SCC 9.1, see <http://supportkb.riverbed.com/support/index?page=content&id=S26956>.

You can view the status of these two communication channels on the SteelHeads using the following CLI command:

```
amnesiac > show scc
Auto-registration:           Enabled
HTTPS connection (to the CMC):
  Status:                   Connected
  Hostname:                 bravo-123
SSH connection (from the CMC):
  Status:                   Connected
  Hostname:                 bravo-123 (10.0.0.7)
```

You must establish these two communication channels to the same SCC for a SteelHead to be fully manageable by the SCC. By default, a SteelHead tries to connect to the SCC using the hostname `riverbedcmc`. Make sure that your DNS system points the hostname `riverbedcmc` to the correct SCC that's managing the appliance.

To change the SCC hostname, enter the following command while in configuration mode:

```
scc hostname <hostname>
```

You can establish the SSH-based communication channel through a manual registration of the appliance in SCC while the HTTPS-based communication channel is disconnected because the default hostname `riverbedcmc` is unresolvable. For more details on the impact of such a situation, see <https://supportkb.riverbed.com/support/index?page=content&id=S25613&actp=REPORT>.

Automatic configuration

Automatic configuration allows a SteelHead, Interceptor, Mobile Controller, Core, and Edge that's registered with the SCC to automatically download its configuration when you add it as a new appliance and it connects to the SCC.

To confirm that automatic configuration is enabled, check the Auto Configure column on the Appliances page before you add the SteelHead or SteelHead Interceptor. The Auto Configuration column indicates whether a SteelHead and SteelHead Interceptor has the Enable Auto Configure check box selected on the Edit Appliance page. Checking the Enable Auto Configure check box pushes the configuration out one time. After the configuration is pushed, the check box is automatically cleared.

If automatic configuration is not enabled, the appliance configuration, including nonpolicy configuration settings found on the Appliance Pages tab on the Appliances page, is updated only when policies are pushed to the appliance.

Locally managed SteelHeads

The SCC has very little knowledge of a change that you apply directly to the SteelHead. You want to change settings directly on the SteelHead only:

- when the SCC software version is earlier than the SteelHead release version, and SteelHead settings haven't been integrated into the SCC.
- if the SCC is a monitoring-only appliance, rather than an appliance you use for configuration. In this deployment, to prevent accidental configuration changes, you can manage the remote appliances with the permissions of a *monitor* account or a role-based user in place of the *administrator* account.

Enabling local changes on a SteelHead

This section describes how to enable local changes on a SteelHead that's under SCC management.

To enable local changes on a SteelHead that's under SCC management

1. Open the Appliances page.
2. Select the name of the SteelHead about which you want to view information.
3. Select Edit Appliance.
4. Select Enable Branch Managed. This selection ensures that the SCC doesn't change the SteelHead configuration, including pushing policies, upgrading, or rebooting.

Figure 2-1. Enable branch managed option

The screenshot shows the SteelHead configuration interface. At the top, there's a header bar with a dropdown menu showing 'main-vsh93 / 10.1.42.97 (VC1KW00004e7d)', a 'SteelHead' icon, the identifier 'VCX555M', and a status indicator 'Connected: Healthy'. Below this is a tabbed interface with five tabs: 'Policies', 'Appliance Pages', 'Edit Appliance' (which is selected), 'Appliance Utilities', and 'Inherited Policies'. The 'Edit Appliance' tab contains the following fields and options:

- Serial Number: VC1KW00004e7d
- Hostname or IP Address: [text box] Autoregistered address: 10.1.42.97
- Comment: [text box]
- Group: Global (dropdown menu)
- Branch Managed: ☒ Enable Branch Managed
- Trusted: ☐ Trusted (Only available to manually untrust trusted appliances.)
- Auto Configure: ☐ Enable Auto Configure (Only available when disconnected. Used only when the policies are ready.)
- User Name: admin
- Password: [password box]
- Confirm Password: [password box]
- [Apply button]

A common mistake is to make a configuration change on a SteelHead locally, and then have the changes *disappear* the next time policies are pushed from the SCC. We recommend that you make all configuration changes for an SCC-managed SteelHead through the SCC. SteelHead settings are available in the SCC as a policy or on the Appliance Pages. You can change any settings not covered by a policy in the Appliance Pages tab.

You can also manage appliances as a cluster. For more information, see [“Appliance Clusters” on page 129](#).

Groups and Policies

This chapter describes the groups and policies that the SCC uses to facilitate centralized configuration and reporting of remote appliances. This chapter includes the following topics:

- “Groups” on page 19
- “Policies” on page 20

Groups

You use groups to organize and arrange sets of managed appliances that share a common configuration policy into a logical container. Common methods of group organization include:

- **Geography** - for example, by region and location
- **Business function** - for example, by office function, such as a branch office or a data center

The SCC uses a hierarchical group model. The SCC defines the group hierarchy with the default Global group as the root. All user-defined groups and managed appliances are contained within the Global group. A managed appliance can be a member of only one group.

Figure 3-1 shows an example of a group organized by business function.

Figure 3-1. Manage appliances page

The screenshot displays the 'Appliances' management interface. At the top, there's a breadcrumb 'Topology > Appliances' and a 'Save' button. Below this is a toolbar with actions like 'New Group', 'New Appliance', 'Remove Selected', 'Move Selected', 'Filter', and 'Appliance Operations'. A 'Migrate Appliances to Sites' button is also present. The main content area shows a tree view of groups: 'Global' (expanded), 'Branch', and 'Data Center'. Under 'Branch', an appliance is listed with details: 'amnesiac / 172.16.210.142 (VC1SM000D58EF)', a 'SteelHead' icon, and a status of 'Connected: Healthy'. The appliance is associated with the 'Perth' site and 'Etc/GMT' time zone. Below the tree view, there are instructions on enabling 'Common Administration' and 'Strict Key Verification', and a link to 'Related Topics: SCC Security'.

Policies

Policies are sets of common configuration options that can be shared among different appliances independently or through group membership. You can specify a policy to a single SteelHead, or it can represent settings for all of the appliances in your enterprise environment.

This section includes the following topics:

- [“Policy overview” on page 20](#)
- [“Inherited policies” on page 20](#)
- [“Policy settings” on page 24](#)
- [“Pushing policy” on page 25](#)

Policy overview

The following policy types are available:

- **Policy** - A configuration you can apply as a common configuration template to multiple appliances.
- **Appliances Specific Pages** - A configuration you can create on a per-appliance basis: for example, interface IP addresses.

Policies consist of one or more *policy pages*. Policy pages generally correspond to a feature (or part of a feature). You must enable a policy page to push the settings you configure.

You can assign each group of appliances, or single appliance, any number of policies, as long as there's no conflicting configuration among those policies.

For more information about working with policies, see the *SteelCentral Controller for SteelHead User's Guide*.

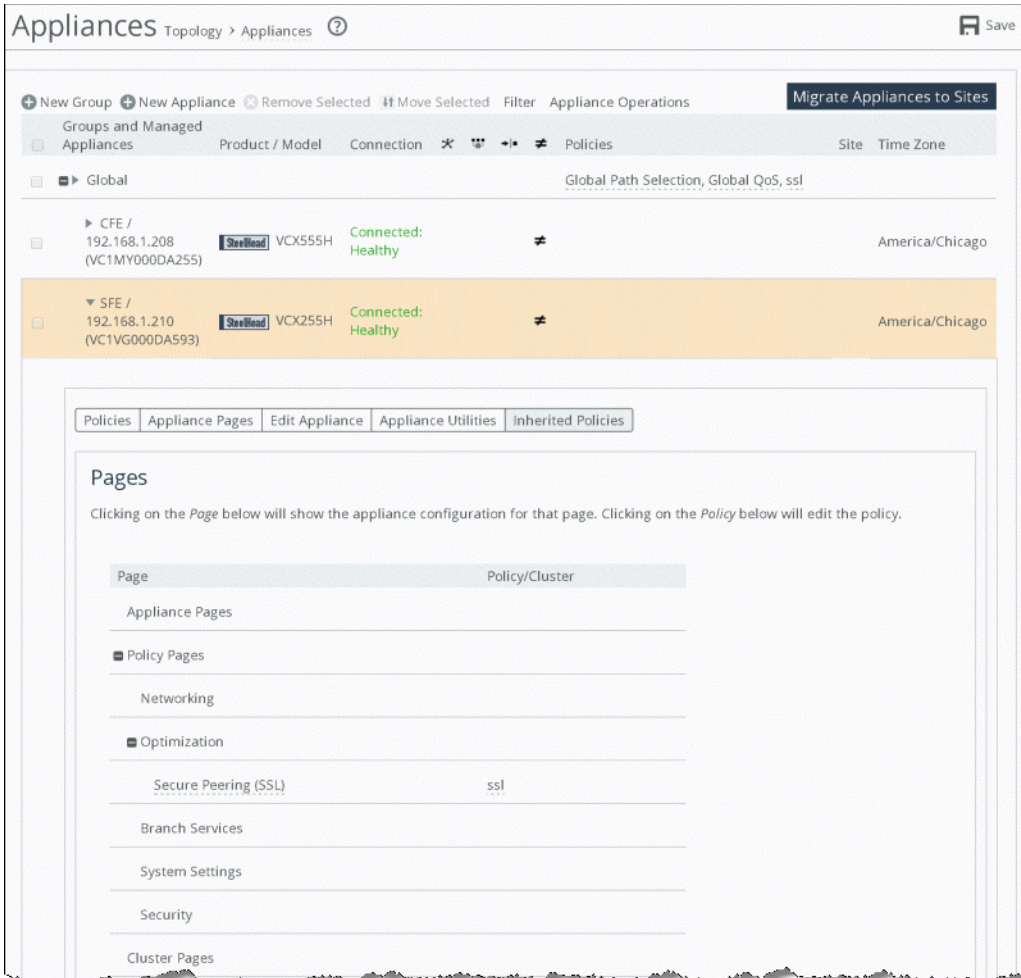
Inherited policies

You can apply policy pages to appliances as follows:

- Assigned and applied directly to appliances in a group
- Inherited from a policy assignment in an ancestor group
- Not assigned, in which case the settings on the appliance remain unchanged

The Inherited Policies view ([Figure 3-2](#)) shows all the policy pages that are scheduled to be pushed to an appliance on the next policy push, and which policy defines the configuration on that page. To view policies inherited by appliance, choose **Manage > Appliance > [the appliance you want to view] > Inherited Policies**.

Figure 3-2. Inherited policies view



[Figure 3-2](#) shows a SteelHead in the New York site (CFE) that doesn't have a policy specifically assigned, so it inherits properties from the SSL policy, which is assigned to the group Global. The other SteelHead in a different location (SFE) can have different policies assigned. Individual settings in this policy override settings in the Global policy.

Inherited policies enable you to:

- control the configuration of a specific appliance in great detail.
- identically configure a large number of appliances.

For example, a networking policy can configure Web Cache Communication Protocol (WCCP). It is unlikely that the WCCP configuration is the same across all data centers, so a Global DC Networking policy doesn't need to include the WCCP settings. You can have the same Network Time Protocol (NTP) servers across all sites and belong to the same global policy.

Figure 3-3 shows that the WCCP page check box is not selected and the Host Settings page check box is selected. These selections indicate that WCCP settings aren't configured by this policy, but that Domain Name System (DNS)—and other host settings—are configured by this policy.

Figure 3-3. Editing page for global dc networking policy

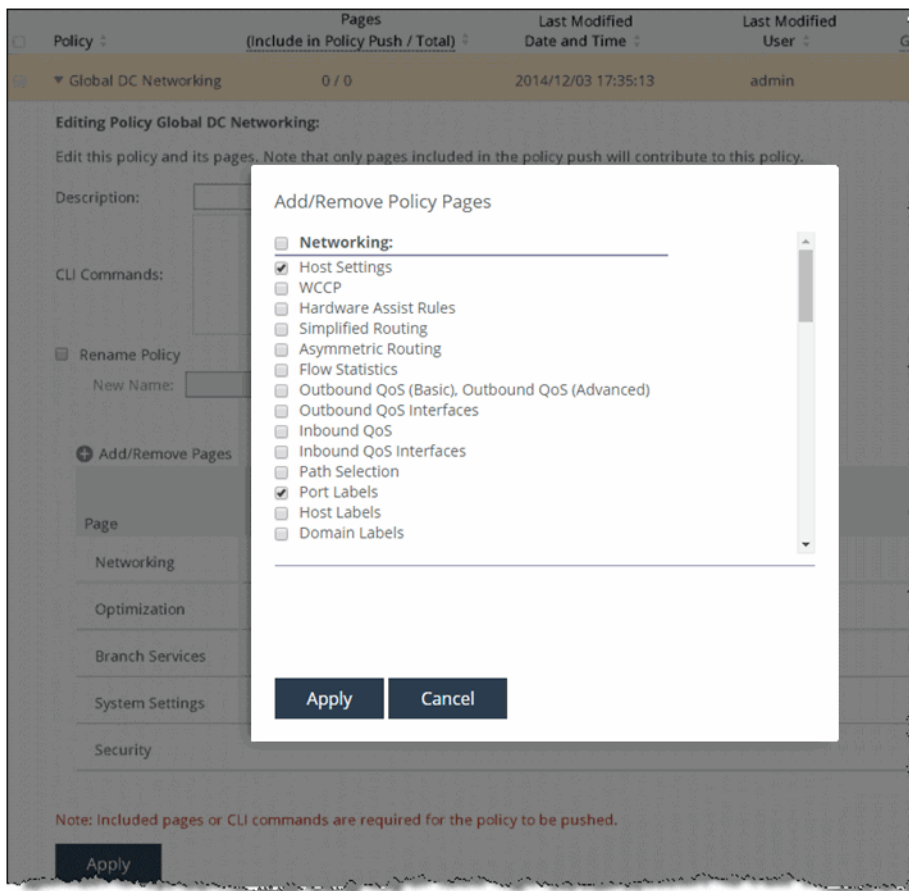
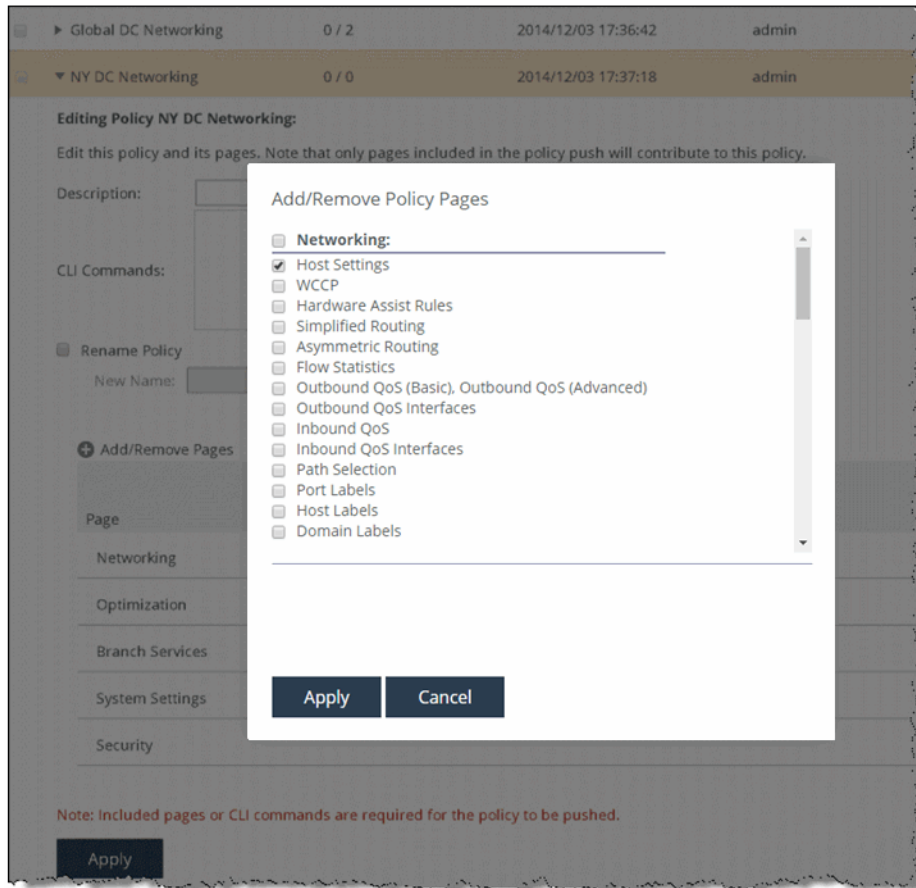


Figure 3-4 shows the Host Settings page check box is selected in the NY DC Networking policy, but the Port Labels page check box is not selected. A SteelHead doesn't inherit host settings from a global networking policy, but it does inherit other pages, if available. The enabled pages in NY DC Networking override inheritance from parent policies.

Figure 3-4. Editing page for NY DC networking policy



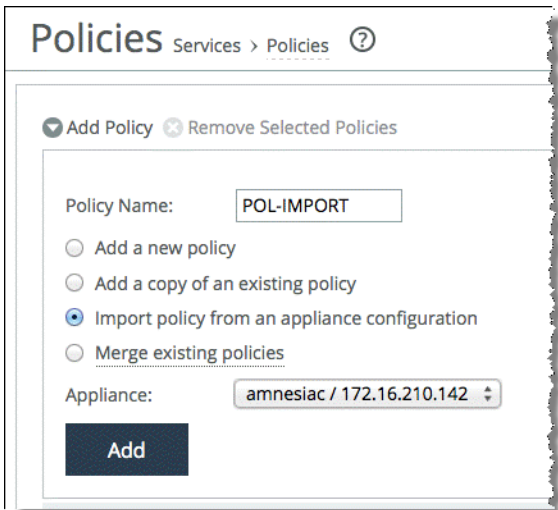
Because of policy inheritance, policies you apply at the Global group provide the default settings for all managed appliances in your environment.

For more information about inherited policies, see the *SteelCentral Controller for SteelHead User's Guide*.

Policy settings

You create policies on the Policies page. In environments in which there are preexisting supported Riverbed appliances deployed, you can use SCC to import the settings of these appliances into a policy. Importing allows you to quickly reuse and apply these settings to another appliance rather than creating them again in a new policy. To import the settings from a managed appliance, choose Manage > Services: Policies > [the appliance that you want to import from] and select Import policy from an appliance configuration (Figure 3-5).

Figure 3-5. Import configuration from managed appliance



Note: When you perform an import from a SteelHead running RiOS 9.0 to SCC 9.0, the SCC only imports the legacy configuration. The new RiOS 9.0 QoS, path selection, and secure transport configurations aren't imported. You've to configure these features on the SCC and push them into the SteelHeads. The SCC maintains the global configuration for these features.

You can use the SCC-managed appliance group hierarchy and policy inheritance to streamline your policy design and configuration push. For example, if you want to apply port labels universally to every SteelHead in the environment, but NTP servers are specific to different locations, you can create two different networking policies. Apply the first policy for the port labels at the Global group level, and apply the second policy for the NTP servers at the data center group level or to the managed appliance.

Although creating a policy from a running appliance configuration is useful, you can end up with as many policies as there are appliances. Consider a group and policy plan before beginning to import SteelHead configurations into policies. Advance thinking about which settings you want to apply at which level prevents mistakes that require pages to be disabled in one policy and reenabled in a different policy at a different layer of the hierarchy. Advance planning about which common settings are shared among appliances in a group avoids the situation in which there's a 1:1 ratio of policies to SteelHeads.

For more information about importing a running configuration into a new policy, see the *SteelCentral Controller for SteelHead User's Guide*.

Pushing policy

Existing settings in a SteelHead do not change until the new policy is pushed. The Push Required symbol is shown in [Figure 3-6](#).

Figure 3-6. Push required symbol



This symbol indicates that there have been policy or configuration changes on the SCC that haven't been pushed to the SteelHead. The symbol also indicates that the configuration of the SteelHead is different from the expected configuration of the SCC for that SteelHead: for example, if you made a local change to a SteelHead. You only see the changes if the SteelHead is managed by the SCC at the times you made the change. You can fetch the appliance configuration from the Appliance Utilities tab in the Appliances page, but there's no automatic comparison of this configuration against the expected policy configuration.

The symbol doesn't persist after an SCC reboots.

In [Figure 3-7](#), SteelHead EX is connected, but the Push Required symbol indicates that policies haven't been pushed. The configuration doesn't necessarily reflect the current policy settings until the policy has been pushed.

Figure 3-7. Example of an SCC-managed SteelHead requiring policy push



To push a policy to a SteelHead

1. Choose Manage > Topology: Appliances.
2. Select Appliance Operations.
3. Select Push Policies from the drop-down list.
4. Select the appliances or groups or both where the policies are to be pushed.

5. Click **Push**.

Figure 3-8. Appliance operations

Appliances Topology > Appliances ?

+ New Group + New Appliance ✕ Remove Selected ⇅ Move Selected Filter Appliance Operations

Choose an operation to perform against the selected groups and appliances:

Push Policies ▼

Push policies and appliance pages to the selected appliances.

☐ Include Path Selection, QoS and Applications configuration for RIOS 9.0 and later.

☐ Restart Optimization Service If Required

☐ Restart QoS Service If Required

☐ Schedule Deferred Push

Date and Time: 2014/11/28 11:46:22 YYYY/MM/DD HH:MM:SS

(Note that this operation only applies to SteelHeads and interceptors.)

Push

6. To see the success or failure of the push, choose Manage > Operations: Operations History.
- Select the date and time to see pushed details.

Figure 3-9. Operations history page

▼ 2014/11/21 12:34:45 Policy Push success admin Successfully pushed to all (1) attempted appliance(s).

Operation Details:

Comment:

Apply

Selected appliances: chief-int8 / 10.1.14.71 (R54NQ0000AFD0).

Appliance	Product / Model	Status	Message
▶ chief-int8 / 10.1.14.71 (R54NQ0000AFD0)	SteelHead Interceptor	9350 success	Push Completed successfully

When the push is successful, the Appliances page shows an empty Push Required column. Figure 3-10 shows that the SteelHead is operating with the new policy.

Figure 3-10. Appliances page showing an SCC-managed SteelHead not requiring a policy push

main-vsh93 / 10.1.42.97 (VC1KW00004e7d) **SteelHead** VCX555M Connected: Healthy

It is a best practice to make a copy of the policy before applying major changes to it. If the changes do not work as expected, you can reassign the previous policy to the affected SteelHeads and repush it to roll back the changes. You can delete the previous policy after changes are successfully applied.

The policy push operation from the SCC is atomic; that is, configuration nodes on the SteelHead aren't left in a partially configured state.

SCC and Appliance Operations

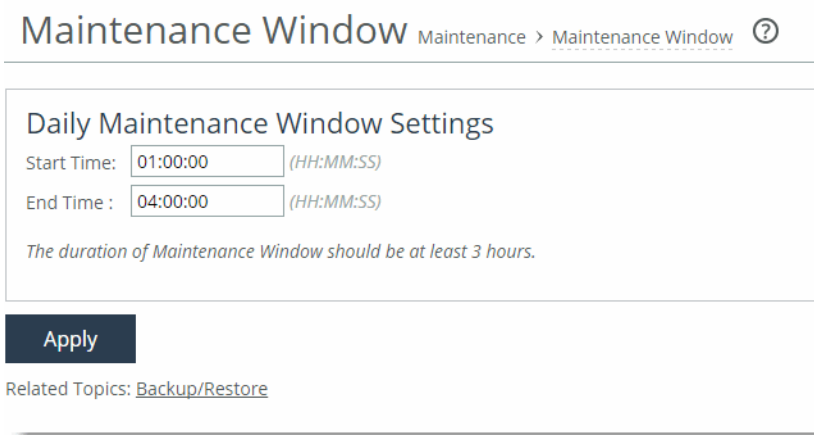
This chapter describes the basic management operations for the SCC. It includes the following sections:

- [“Backing up and restoring an appliance” on page 30](#)
- [“Migrating SteelHead settings with the SCC” on page 33](#)
- [“SCC backups and high availability” on page 34](#)
- [“Using the secure vault” on page 36](#)
- [“REST API access” on page 37](#)
- [“Using management ACLs” on page 37](#)
- [“Managing licenses using the SCC” on page 38](#)
- [“Using the SCC for Interceptor management” on page 39](#)
- [“SCC bandwidth consumption” on page 40](#)
- [“Application statistics” on page 40](#)
- [“Using the SCC upgrade and downgrade appliance wizard” on page 43](#)
- [“Configuring and Using the SCC as a Certificate Authority Service” on page 45](#)

Backing up and restoring an appliance

The SCC automatically performs a configuration backup (*snapshot*) of the managed appliances. The backup occurs during the maintenance window (1:00 A.M. to 4:00 A.M. by default). To display the Maintenance Window page, choose Administration > Maintenance: Maintenance Window.

Figure 4-1. Maintenance window



The screenshot shows the 'Maintenance Window' configuration page. At the top, the title 'Maintenance Window' is followed by a breadcrumb 'Maintenance > Maintenance Window' and a help icon. Below this is a section titled 'Daily Maintenance Window Settings'. It contains two time input fields: 'Start Time' set to '01:00:00' and 'End Time' set to '04:00:00', both with '(HH:MM:SS)' format hints. A note below the fields states: 'The duration of Maintenance Window should be at least 3 hours.' At the bottom of the settings section is a dark blue 'Apply' button. Below the button, it says 'Related Topics: [Backup/Restore](#)'.

To back up the configuration of a managed appliance

1. Choose Manage > Operations: Backup/Restore (Figure 4-2).
2. Select a Source Appliance and enter a snapshot name.
3. Click **Backup**.

The backup is stored on the SCC.

Figure 4-2. Managed appliance back up

Backup/Restore Operations > Backup/Restore ?

Appliance Snapshot Rollover Strategy

Rollover Strategy: **Age-Limited**

Keep only the most recent backups. All backups for the last 30 days are kept, and for anything older only the first backup of the month is kept.

Apply

Source Appliance: **amnesiac / 172.16.210.142**

Backup Operation

To backup this appliance to the SCC, provide a snapshot name and press the *Backup* button. Note: Backup for RiOS 9.0+ does not backup Path Selection, QoS or Secure Transport configuration.

Name for Snapshot: **Snapshot-Tue**

Backup

Note: When you back up and restore a SteelHead running RiOS 9.0 with SCC 9.0, the SCC only backs up the legacy configuration. The new RiOS 9.0 QoS, path selection, and secure transport configurations aren't backed up to or restored from a SteelHead. For details, see the *SteelCentral Controller backup/restore of SteelHead for 9.0.0* article at <https://supportkb.riverbed.com/support/index?page=content&id=S25391>.

To restore from a backup

1. Choose Manage > Operations: Backup/Restore.
2. Select a source appliance from the drop-down menu.
3. Enter the snapshot name to restore from and the target appliance.
4. Click **Restore**.

Note: We recommend that you also perform a policy push to restore the configuration stored in the SteelHead secure vault.

Figure 4-3. Managed appliance restore

Restore Operation

To restore a snapshot from the SCC to an appliance, click one of the snapshots below.

Snapshot Name	Type	Creation Date
▶ 2015.02.10-T1	timed	2015/02/10 03:36:08
✓ ▼ Snapshot-Mon	user	2015/02/25 02:14:28

Snapshot Configuration Version: 9.0.0

```
##
## Network interface configuration
##
interface aux shutdown
interface inpath0_0 description ""
no interface inpath0_0 dhcp
no interface inpath0_0 dhcp dynamic-dns
no interface inpath0_0 force-mdi-x enable
interface inpath0_0 ip address 172.16.1.10 /24
```

☐ Include defaults and hidden commands

Restore Operation

Restore this snapshot to the source appliance, or to another appliance. Note: Restoring configuration to RIOS 9.0+ does not restore or modify Path Selection, QoS or Secure Transport configuration.

Source serial: VC15M000D58EF

☐ Restore to a Different Appliance

Target Appliance:

☐ Restore Primary and Aux network interfaces

Restore

Restoring the backup to the source appliance is useful for rolling back configuration changes. As a best practice, create a manual backup before any major configuration changes. You can restore a backup from one appliance to another; the restore operation only restores the appliance configuration.

As a best practice, use the migrate operation when moving a configuration between compatible appliances (SteelHeads, Interceptors, Mobile Controller, and Cores). The migrate operation moves statistics and an SCC configuration from one appliance to another, including snapshots. The source appliance is deleted after a successful migration.

For details on migration, see the *SteelCentral Controller for SteelHead User's Guide*.

You can restore SteelHead CX and SteelHead xx50 configurations among each other. For example, you can restore a SteelHead 1050 configuration to a SteelHead CX 1555.

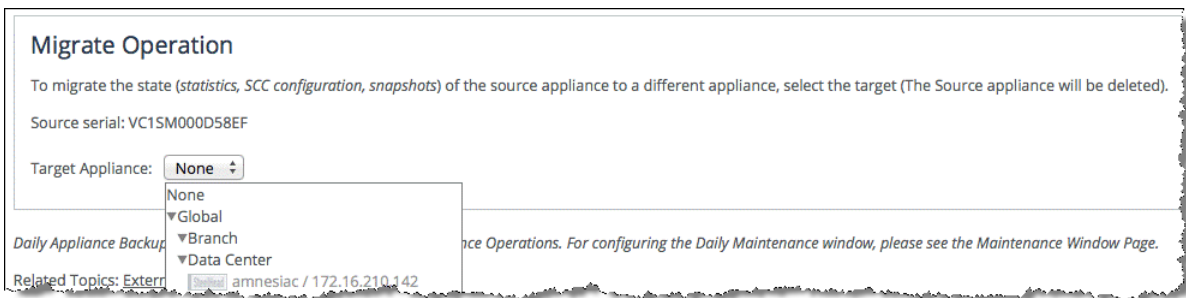
The following caveats apply to for backup and restore operations:

- The destination appliance RiOS version must be equal to or greater than the source appliance RiOS version.
- You can't back up and restore between SteelHead EX and other SteelHeads.
- You can't back up and restore SteelHead-v (xx50) and hardware-based or SteelHead-v SteelHead CX.
- You can't restore configurations from a SteelHead EX to a SteelHead CX or a SteelHead xx50.
- You can't restore from a SteelHead CX or a SteelHead xx50 to a SteelHead EX.
- You can't back up or restore VSP or SteelFusion Edge configurations on a SteelHead EX.

Migrating SteelHead settings with the SCC

You generally use the migrate operation to move the state of an existing appliance to the replacement appliance. The migrate operation moves the state statistics, SCC configuration, and snapshots from one SteelHead (referenced by serial number) to another. After you've a successful migrate operation, the new SteelHead is managed by the SCC and the source settings are deleted.

Figure 4-4. Migrate operation on the backup/restore page



You can't perform a migrate operation between SteelHead CX and SteelHead EX.

For complete migration procedures, see the *SteelCentral Controller for SteelHead User's Guide*.

SCC backups and high availability

The SCC can back up its own configuration, appliance snapshots, and statistics to an external file server over NFS, SSH, or CIFS. Choose Administration > Maintenance: External Backup and configure the external file server to use for backup.

Note: You can also use the external backup functionality to migrate data and configuration settings from one SCC to another. For more information, go to <https://supportkb.riverbed.com/support/index?page=content&id=S14439>.

Figure 4-5. External backup page

The screenshot shows the 'External Backup' configuration page. At the top, the breadcrumb is 'Maintenance > External Backup'. Below the title, a subtitle reads 'Configure SCC and Appliance backups to an external server.' The main section is titled 'Backup Server'. It contains several fields: 'Protocol' is set to 'SSH' with a dropdown arrow; a note below it says '(Backups and Restores using SSH can be slow. Consider using CIFS or NFS instead.)'; 'Hostname or IP Address' is '10.10.10.10'; 'Remote Path' is '\\sccbackuppath\scc'; 'CIFS Domain' is empty; 'User Name' is 'admin'; 'SSH Authentication Type' has two radio buttons, 'Password' and 'Public Key', with 'Public Key' selected; 'Password' and 'Password Confirm' fields are present but empty; 'Public Key' has a large text area and a 'Generate Key' button; 'Time Limit for Statistics Backup' is '300' with the unit 'Minutes (0 for no limit)'; 'Disk Space Limit' is '0' with the unit 'MBytes (0 for no limit)'. At the bottom, a message states 'Backup space usage information is unavailable.' with a progress bar below it.

In SCC 9.2 and later, you can use a public security key associated with the user instead of a password. Configuring external backups using a public key eliminates the need for password authentication. You can generate a 2048-bit encryption key from the SCC Management Console and then cut and paste the key to your home directory target machine, for example, /u/test/.ssh/authorized_keys.

To restore an SCC from an external backup, choose Administration > Maintenance: External Backup and select a restore type from the Backup Operation (Figure 4-6).

Figure 4-6. Restore an SCC from an external backup

The screenshot shows the 'Backup Operations' page. It contains a table of status information and a configuration section for a backup operation.

SCC Configuration Backup Status:	success at 2015/02/25 03:12:09 (duration: 0:00:02)
SCC Snapshot Restore Status:	idle
Appliance Snapshots Backup Status:	idle
Appliance Snapshots Restore Status:	idle
Statistics Backup Status:	idle
Statistics Restore Status:	idle

Below the table, the 'Backup Operation' is set to 'Restore SCC Snapshot'. The 'Restore Snapshot Name' is 'SNAP-SCC-01 (9.0.0a, this SCC)'. There are two checkboxes: 'Restore Secure Vault' (checked) and 'Restore Primary and Aux network interfaces' (unchecked). A 'Vault Password' field is present with a hint '(leave blank if factory password)'. A 'Start' button is at the bottom.

If you want to restore the SCC management IP addresses with those from the external backup configuration, select Restore Primary and Aux network interfaces. This selection is useful if the SCC you want to restore is located in the same data center and on the same subnet as the original SCC. For more information, see *Changing the IP address of a SteelCentral Controller for SteelHead (SCC)* at https://supportkb.riverbed.com/support/index?page=content&id=S13134&actp=search&viewlocale=en_US&searchid=1368563114732.

For details on the SCC external backup process, see the *SteelCentral Controller for SteelHead User's Guide*.

You can maintain a second SCC offline as a cold, standby device. If the primary SCC fails and is no longer available, then you can power on the second SCC and configure it through a restore operation from the external backup configuration and data. You can manage all appliances with one SCC.

Note: The SCC doesn't support true active-active high availability failover. If you use SCC-VE, you can restore a failed host with VMware high-availability features such as vMotion and Distributed Resource Scheduler (DRS). In addition, you can routinely restore external backups from a primary SCC to a secondary SCC to keep their state consistent. This routine restoration allows for a faster switchover in the event that the primary SCC fails.

Use one of the following methods to switch the managed SteelHeads from the primary to the secondary SCC:

- Assign the primary IP address of the primary SCC to the secondary SCC. If the secondary SCC was restored from an external backup of the primary SCC, select Restore Primary and Aux network interfaces. The secondary SCC has the list of managed appliances from the external backup.
- If you're using automatic registration, update the SCC hostname in the DNS accordingly. The default SCC hostname configured in SteelHeads is **riverbedcmc**.
- Use the **cmc hostname** and **scc hostname** commands on each managed SteelHead to set it to the correct SCC.

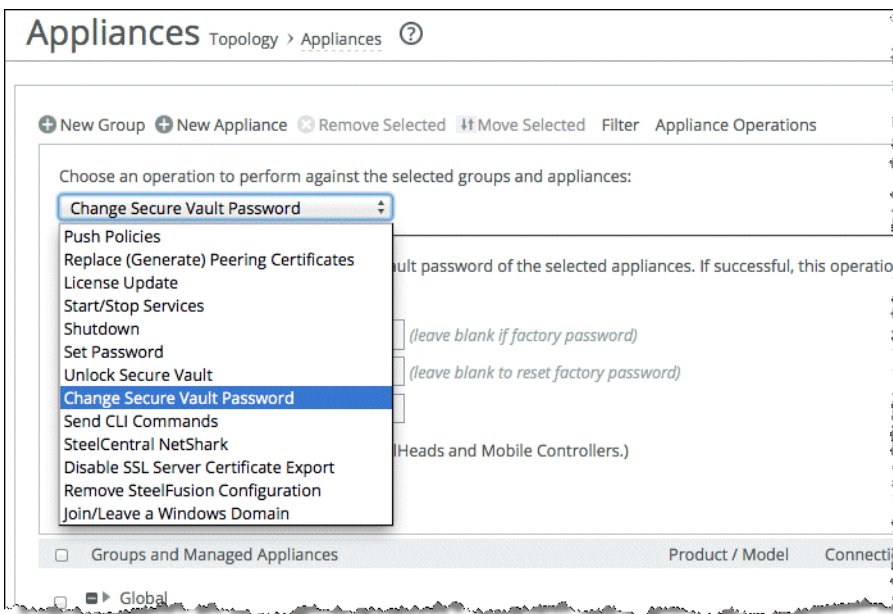
Multiple SCCs can't manage the same SteelHead. When performing a full restore on an SCC on warm standby, there's potential risk that the standby SCC will take over the management for some SteelHead if there's a temporary network disruption to the active (primary) SCC connectivity.

Note: Each SCC requires its own set of licenses. You can't share licenses among difference SCC appliances.

Using the secure vault

You can change user passwords, change the secure vault password, and unlock the secure vault on your managed appliances from the Appliance Operation tab on the Topology > Appliances page (Figure 4-7).

Figure 4-7. Appliance operations



The SteelHead configuration contains sensitive information such as SSL private keys, authentication tokens, and the RiOS data store encryption key. These configuration settings are encrypted on the disk in the secure vault, using AES 256-bit encryption.

Initially, the secure vault is keyed with a default password known only to RiOS. This default password allows the SteelHead to automatically unlock the vault during system startup.

You can change the secure vault password. The SteelHead reboots when you change the secure vault password. The secure vault doesn't automatically unlock on startup if you've changed the default password.

Certain operations such as SSL optimization and RiOS data store encryption, require you to unlock the secure vault.

If you configure the secure vault password through the SCC Change Secure Vault Password page or manage it through the Unlock Secure Vault page, the SCC automatically unlocks the SteelHead secure vault through a secure channel when the SteelHead reboots.

If you want to use the SCC to manage a SteelHead that has a nondefault secure vault password, use the unlock secure vault operation on the Appliances Operations tab on the Appliances page to unlock it. If this operation is successful, SCC automatically stores the appliance secure vault password in the SCC secure vault.

The SteelHead secure vault password is stored in the SCC secure vault. If you change the SCC secure vault password from the default, the vault must be manually unlocked when the SCC boots up.

Certificates stored in the SteelHead secure vault aren't backed up by the nightly SCC backups. You must create a separate policy to back up the certificates.

Certificates you configure in the SCC policy are stored in the secure vault of the SCC.

REST API access

Choose Administration > Security: REST API Access to display the REST API Access page.

Representational State Transfer (REST) 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 SteelCentral NetProfiler communicating with a SteelCentral NetShark
- A SteelCentral NetProfiler retrieving a QoS configuration from a SteelHead

For all uses you must preconfigure an access code to authenticate communication between parties and to authorize access to protected resources. You can use an SCC policy to configure the Rest API Access Code on a managed appliance.

For more information about the SteelHead REST API, see the *SteelHead Management Console User's Guide* and the *SteelHead REST API Guide*.

Using management ACLs

The SCC supports the Management ACL. The Management ACL enables you to restrict access to the SCC management functions by limiting network access. You configure rules to accept or reject packets that match specific source IP addresses, subnets, or ports.

For details on configuring the Management ACL, see the *SteelCentral Controller for SteelHead User's Guide*. For details on SSL and SCC, see the *SteelHead Deployment Guide - Protocols*.

Managing licenses using the SCC

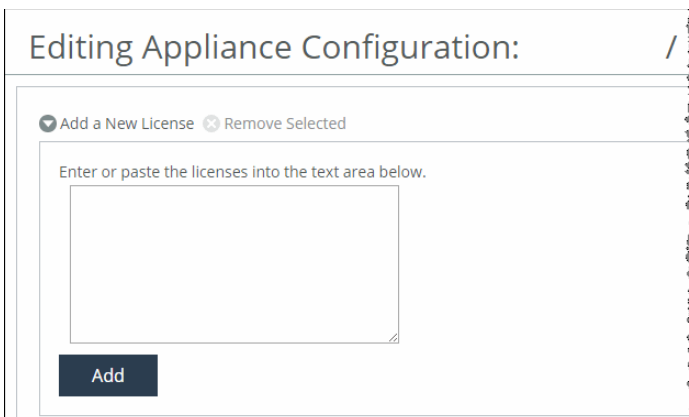
The SCC can manage licenses on SteelHead and SteelHead Interceptors.

To view licenses for an appliance on the SCC

1. Open the Appliances page.
2. Select the desired appliance.
3. Select the Appliance Pages tab.
4. Select Licenses (it is highlighted).

You can add or remove licenses from this page (Figure 4-8). License information doesn't populate until you perform the fetch operation on the Appliance Pages. Alternatively, you can configure the SCC to fetch licenses from the Riverbed licensing portal.

Figure 4-8. Licenses page

The screenshot shows a web interface titled "Editing Appliance Configuration:". Below the title bar, there are two buttons: "Add a New License" (with a plus icon) and "Remove Selected" (with a minus icon). Below these buttons is a text input area with the placeholder text "Enter or paste the licenses into the text area below." and a large empty text box. At the bottom left of the text box is a dark blue button labeled "Add".

Note: The SCC doesn't delete any licenses on the appliance. You can push licenses to the SteelHead even if you haven't fetched appliance information.

Figure 4-9 shows the fetch operation. The SCC doesn't automatically fetch the licenses from the SteelHeads. Click **Fetch Appliance Configuration** to fetch licenses (and other appliance-specific information).

Figure 4-9. Fetch operation

Policies Appliance Pages Edit Appliance Appliance Utilities Inherited Policies

Changes made to the appliance configuration pages will modify the appliance after a policy push.

Appliance Configuration Pages	Include in Policy Push
Host Settings	<input type="checkbox"/>
Base Interfaces	<input type="checkbox"/>
In-Path Interfaces	<input type="checkbox"/>
Subnet Side Rules	<input type="checkbox"/>
SSL	<input type="checkbox"/>
Licenses	<input type="checkbox"/>
Web Settings	<input type="checkbox"/>
Outbound QoS Interfaces	<input type="checkbox"/>
Inbound QoS Interfaces	<input type="checkbox"/>
Path Selection	<input type="checkbox"/>
Connection Forwarding	<input type="checkbox"/>

Apply

Fetch Appliance-Specific Configuration

Copy the appliance-specific configuration (network interfaces, licenses, etc.) into the Appliance Pages.

☐ Set the fetched pages to be included in a policy push. By default none will be included.

Fetch Appliance Configuration

Using the SCC for Interceptor management

SteelHead Interceptor management is supported by SCC 8.5 or later. From a topology and functional point of view, SteelHead Interceptors work in conjunction with SteelHeads. The topology you choose for your Interceptor and SteelHeads defines how you configure the function of your SteelHead Interceptor. The SCC enables you to manage all SteelHead Interceptor configurations in one centralized location. You can manage SteelHead Interceptors as individual appliances or as part of a SteelHead and SteelHead Interceptor cluster.

We recommend that you use the SCC Cluster Configuration Wizard to configure SteelHead and SteelHead Interceptor clusters instead of individually configuring each appliance. When you use the Cluster Configuration Wizard, you streamline and simplify the configuration and avoid mistakes.

If there are configuration parameters of the SteelHead Interceptors that aren't directly related to the cluster, you can use the SCC to configure these as policies. You configure SteelHead Interceptors that aren't in clusters the same way you configure SteelHeads—define policies and apply them to a SteelHead Interceptor or groups of SteelHead Interceptors.

The SCC doesn't support VLAN segregation in Interceptor 4.0 or later.

For more details on SteelHead and SteelHead Interceptor clusters, see [“Appliance Clusters” on page 129](#). For more general information about the SteelHead Interceptor, see the *SteelHead Interceptor User's Guide* and the *SteelHead Interceptor Deployment Guide*. For details on how to configure the SteelHead Interceptor and the SCC, see the *SteelCentral Controller for SteelHead User's Guide*.

SCC bandwidth consumption

In a steady state, the SCC is continuously gathering statistics and state information from managed appliances. Anywhere from 25 to 150 KB of data once every five minutes is gathered from the appliance.

Operations that are bandwidth intensive include configuration pushes and upgrades. A configuration push can send up to 2 to 3 MB to each appliance. When performing upgrades on SteelHeads, the SCC has to transfer the software image. The images are typically around 85 MB in size.

Application statistics

SCC 9.1 and later include application statistics. Application statistics provide an additional application visibility component. With application statistics you can view a report based on the top ten used applications traversing a SteelHead. The report shows you the number of bits per second that a certain application has incurred over a one-week period in one-hour increments. The report displays optimized traffic, pass-through traffic, and combined totals for the observed period.

For more information about application statistics, see the *SteelCentral Controller for SteelHead User's Guide*.

Note: Application statistics data collection from SteelHeads might lag a few hours for an SCC managing more than 500 SteelHeads running RiOS 9.1 or later.

Application statistics leverages the Riverbed Application Flow Engine (AFE) for determining the top applications instead of relying on raw destination IP addresses and port numbers. The report displays the actual WAN traffic as reported by the independent SteelHeads.

Note: Application statistics doesn't track custom applications. Only applications that are a part of AFE are shown.

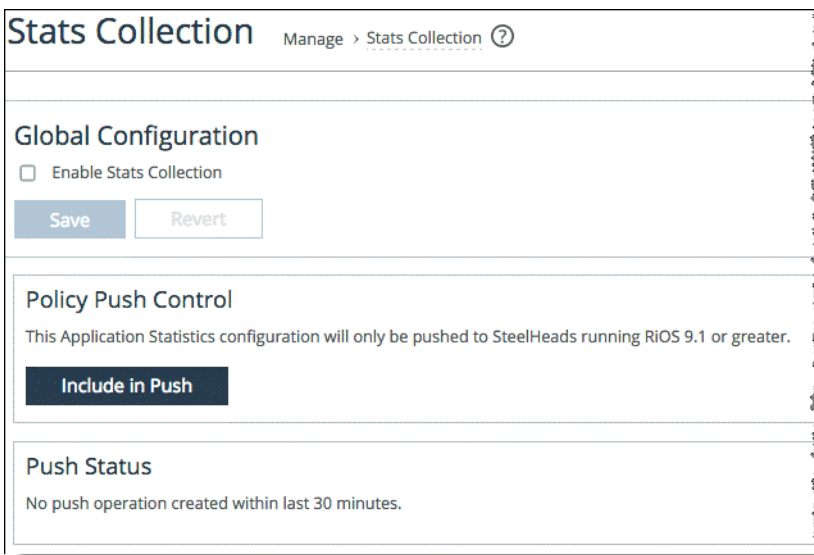
Before you configure application statistics, you must define the topology—in that the SteelHeads must already belong to a predefined site. Multiple SteelHeads, as part of a single site have their data combined and displayed on a per-site perspective as opposed to a per-SteelHead perspective.

For information about topologies, see [“Sites and Networks, Path Selection, and QoS” on page 61](#).

Application statistics configuration requires SCC 9.1 or later and RiOS 9.1 or later.

To configure application statistics

1. Configure the topology and sites accordingly.
2. From the SCC Management Console, choose Manage > Applications: Stats Collection (Figure 4-10).

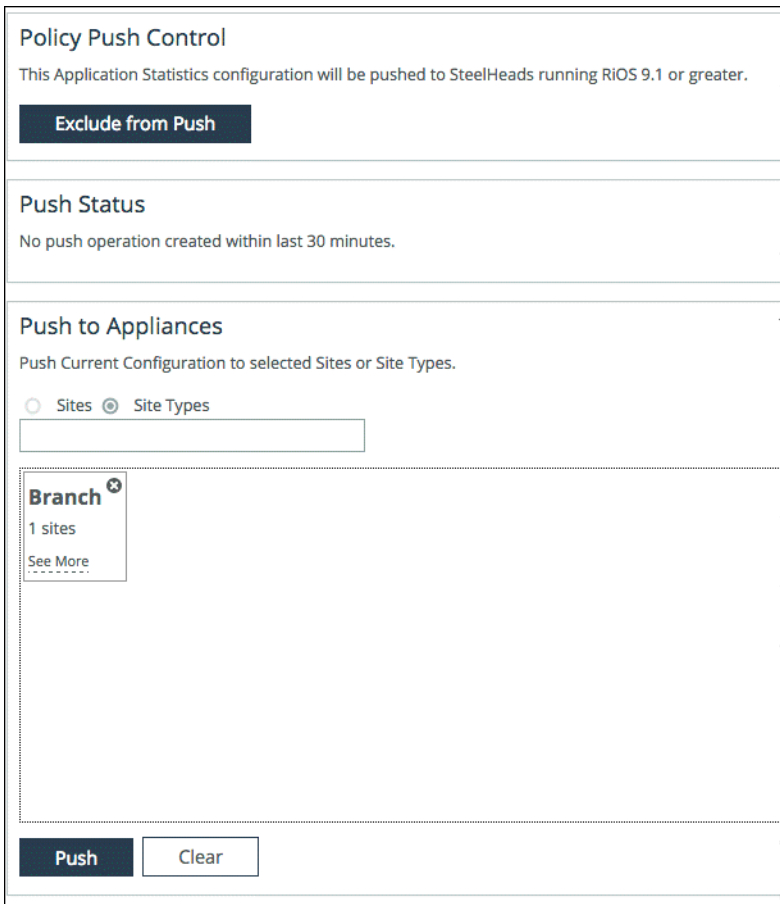
Figure 4-10. Stats collection page

The screenshot shows the 'Stats Collection' configuration page. At the top, there is a breadcrumb trail 'Manage > Stats Collection' with a help icon. The page is divided into three main sections: 'Global Configuration', 'Policy Push Control', and 'Push Status'. In the 'Global Configuration' section, there is a checkbox labeled 'Enable Stats Collection' which is currently unchecked. Below this checkbox are two buttons: 'Save' and 'Revert'. The 'Policy Push Control' section contains a message stating 'This Application Statistics configuration will only be pushed to SteelHeads running RiOS 9.1 or greater.' and a button labeled 'Include in Push'. The 'Push Status' section shows a message 'No push operation created within last 30 minutes.'

3. Select Enable Stats Collection.
4. Click **Save**.
5. Select the Site or Site Types from which you want to gather application statistics.
6. Click **Push**.

You can exclude certain sites ([Figure 4-11](#)).

Figure 4-11. Select site or site types



Policy Push Control
This Application Statistics configuration will be pushed to SteelHeads running RiOS 9.1 or greater.

Exclude from Push

Push Status
No push operation created within last 30 minutes.

Push to Appliances
Push Current Configuration to selected Sites or Site Types.

☐ Sites ☒ Site Types


Branch ✕
1 sites
[See More](#)

Push **Clear**

The SteelHead starts to collect application statistics immediately after it is enabled with a policy push. These statistics are rolled up into an hourly average reading on the SteelHead and polled hourly by the SCC. You must allow for a short period of time to pass after configuration before you can see the statistics appear on the SCC.

Figure 4-12 shows an example report.

Figure 4-12. Top applications report



You can verify if application statistics is enabled on a SteelHead by running the following command on individual SteelHeads:

```
sh2 > show appstats
Application Statistics:
  Enabled:    yes
  Resolution: 300
  Rollup:     AVERAGE
```

Consider the following restrictions when configuring application statistics:

- Application statistics is not supported in out-of-path deployments.
- In virtual in-path deployments, subnet side rules are required for proper pass-through traffic statistic collection.
- Application statistics supports IPv4 only.
- There's a 10 to 15% CPU penalty in high connection setup rate conditions.

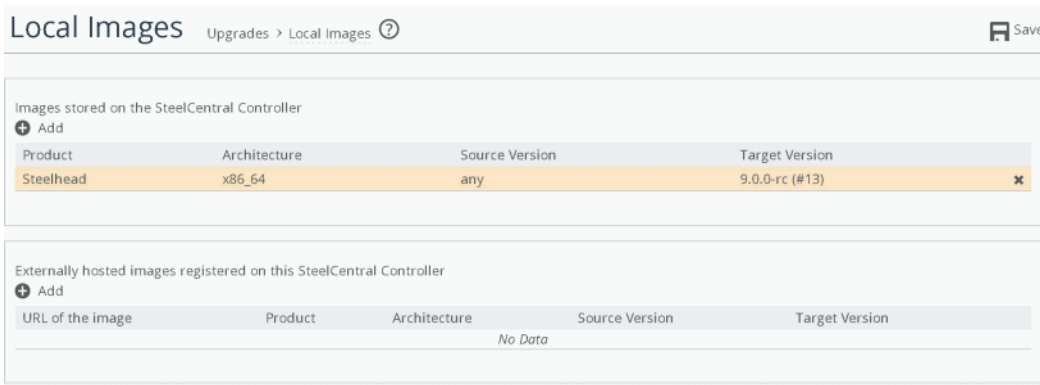
Using the SCC upgrade and downgrade appliance wizard

SCC 8.6 and later include a wizard to upgrade and downgrade SteelHead RiOS images. Several tasks are involved to upgrade and downgrade:

1. Specify the image library.
2. Configure upgrade or downgrade job.
3. Configure the reboot job parameter, which enables that image to run on the specific appliance.

Each specific task has its own dashboard as shown in [Figure 4-13](#).

Figure 4-13. Specify the image library

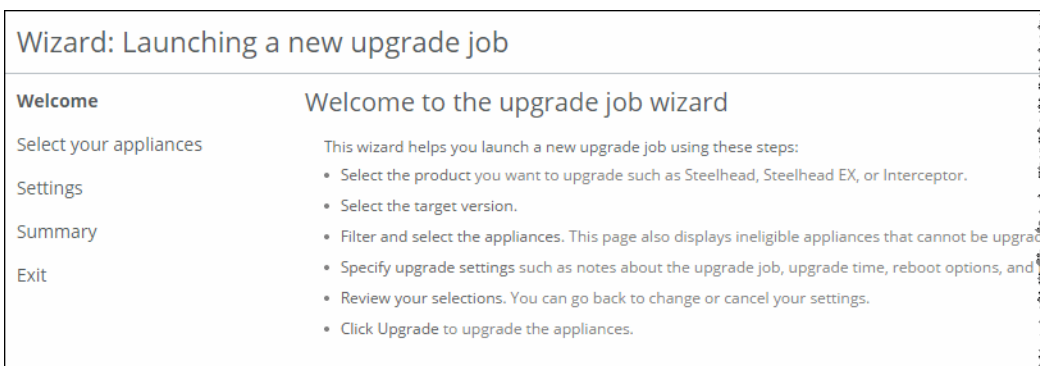


Use the following methods to identify the image libraries available for use:

- **Host the images stored locally** - Local image storage is a process in which you can upload various RiOS images to the local storage of the SCC. You either browse a directory to the image or use a URL to the image and upload it to the SCC. When you use this method, the SCC holds the various images and makes them available for distribution to various other appliances according to the upgrade job parameters.
- **Customer-defined URL links** - Instead of locally storing the images on the SCC itself, you can provide a URL link to which various images are located. This method doesn't upload the images to the SCC locally but serves as a trusted pointer to an available remote link.
- **Integrate available images directly from Riverbed Support** - The SCC validates (according to the software entitlement program) available images for use with any of the managed devices in its inventory. This method provides all the images available that the SteelHead can access from the Riverbed Support portal. You can download the chosen image directly to the SteelHead; however, you can't store the image locally on the SCC.

Use the Jobs Wizard to guide you through the process. Consider the outcome you want to achieve prior to starting the wizard: for example, upgrading a certain set of appliances to a certain code version. To open the Jobs Wizard, choose **Manage > Upgrades: Upgrade Appliances (or Downgrade Appliances) > Launch new upgrade job....**

Figure 4-14. Jobs wizard



The Jobs Wizard guides you to:

1. choose the product to upgrade: for example, SteelHead, SteelHead EX, or SteelHead Interceptor.
2. target the product version (this information is provided from the image library result of the previous task).
3. select the appliances (the Jobs Wizard displays available licenses compatible with choices selected).
4. choose the upgrade settings: for example, notes about the upgrade job, upgrade time, reboot options, and so on.
5. view the summary page to review your selections.
6. initiate the upgrade or downgrade job.

The downgrade image must exist locally on the SCC or referenced as a customer-defined URL. You can't point to referenced images on the Riverbed Support site for downgrade jobs.

Configuring and Using the SCC as a Certificate Authority Service

You can enable a certificate authority (CA) service in SCC 8.6 and later. You can configure the SCC CA as a private root CA or an intermediate CA that is trusted within your organization. The SCC CA service enables you to issue the following certificates to SteelHeads:

- Secure peering certificates
- Proxy certificates for SSL optimization (starting from SCC 9.5 and later)
- Web proxy certificates for HTTPS traffic proxy. For details, see [Chapter 6, "Web Proxy."](#)

As a SteelHead deployment size increases, independently managing certificates for secure protocol optimization and HTTPS web proxy can be a daunting and time-consuming task. The SCC CA service offers you a method to simplify, streamline and automate this task from a central console. Using the SCC CA service, you would be able to:

- Easily manage and issue secure peering certificates to SteelHeads.
- Simplify the task of configuring secure peering trust relationships between SteelHeads for secure protocol optimization.

When you replace the secure peering certificate on a SteelHead with one issued by the SCC CA, and use the SCC to configure the secure peering trust relationship, the SCC CA is configured as a trusted entity on the SteelHead. This allows the SteelHead to automatically trust all peers that have a secure peering certificate issued by the same SCC CA, and eliminates the need to configure the secure peering trust by trusting a SteelHead peer one at a time

- Automate the issue of proxy certificates for SSL optimization.
- Automate the issue of web proxy certificates for HTTPS traffic proxy.

- Increase operational efficiency by centrally managing and issuing certificates for secure protocol optimization and HTTPS web proxy

Note: You can't submit a Certificate Signing Request (CSR) to have a certificate signed by the SCC CA through the SCC Management Console. The SCC CA can only be used to issue certificates and implicitly signs all certificates that it issues. Using and trusting only CA-signed certificates increases the security of your SteelHead installation.

This section describes how to use the SCC to complete these tasks:

- “To enable the SCC CA service” on page 46
- “To use the SCC CA service to issue a secure peering certificate” on page 48
- “To use SCC to configure secure peering relationships” on page 51
- “To use the SCC CA service to issue a proxy certificate for bypassed servers” on page 55

To enable the SCC CA service

1. From the Management Console, choose Administration > Security: Certificate Authority.
2. Select Enable/disable the certificate authority and click **Apply** (Figure 4-15).

Figure 4-15. SCC CA

Certificate Authority Security > Certificate Authority ?

Configuration of the authority itself

☒ Enable/disable the certificate authority

Cipher bits: 2048

Signing algorithm: SHA256withRSA

Apply

Note: Many vendors are phasing out the support of SHA-1 certificates and it is recommended that you do so as well. A cipher bits of 2048 or higher setting is also recommended.

3. If you're using the SCC CA as a root CA, generate the root CA certificate using the next step.

4. Select Generate New Private Key and Self-Signed Public Certificate (Figure 4-16).

Figure 4-16. Generate keys and certificate

The authority's keys

Details
PEM
Replace

☐ Import Existing Private Key and CA-Signed Public Certificate (One File in PEM or PKCS12 formats)
☐ Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM or DER formats)
☒ Generate New Private Key and Self-Signed Public Certificate

Private Key

Cipher: RSA Cipher Bits: 2048

Self-Signed Certificate

Common Name: cmc-ca.example.com

Organization Name: Riverbed Technology, Inc

Organization Unit Name: SteelHead

Locality: Singapore

State: SG

Country: SG (two-letter code)

Email Address: admin@example.com

Validity Period: 1024 Days (60 to 3650 days)

Generate Key And Certificate

If you're using the SCC CA as an intermediate CA, import the root CA certificate into the Trusted CA Store. The Trusted CA Store is used only by the SCC CA service for the purpose of certificate chaining to establish itself as an intermediate CA.

5. To import the root CA into the Trusted CA Store choose Administration > Security: Trusted CA Store.

Figure 4-17. Trusted CA Store



Next, import the intermediate CA certificate and private key in PEM format into the SCC CA service.

6. To import the intermediate CA and private key, choose Administration > Security: Certificate Authority and click **Replace**.

Figure 4-18. Importing the intermediate CA

The authority's keys

Details PEM Replace

☒ Import Existing Private Key and CA-Signed Public Certificate (One File in PEM format)
☐ Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM format)
☐ Generate New Private Key and Self-Signed Public Certificate

Import Single File

☐ Local File
Browse... No file selected.

☒ Text

```
-----BEGIN CERTIFICATE-----
MIIECjCCAnKgAwIBAgIBATANBgkqhkiG9w0BAQwFADARMQ8
wDQYDVQQDDAZTQ0Mt
Q0EwHhcNMTYwOTI5MDgzODM2WhcNMTgwOTI5MDgzODg
-----
```

Decryption Password:

Import Key And Certificate

Note: You can only enable the SCC CA service as an intermediate CA starting from SCC 9.5 and later.

Note: You can also import the entire certificate chain through the SCC CA service. The intermediate CA certificate must be the first certificate in the chain. The root CA certificate will not be automatically installed into the Trusted CA Store.

To use the SCC CA service to issue a secure peering certificate

1. From the Management Console, choose Manage > Topology: Appliances.
2. Select Appliance Operations.

3. Select Replace (Generate) Peering Certificates from the drop-down menu (Figure 4-19).

Figure 4-19. Appliance operation page

The screenshot shows the 'Appliances' page in the SteelCentral Controller. The breadcrumb is 'Topology > Appliances'. The page has a toolbar with 'New Group', 'New Appliance', 'Remove Selected', 'Move Selected', 'Filter', and 'Appliance Operations'. Below the toolbar, there is a section for 'Choose an operation to perform against the selected groups and appliances:'. A dropdown menu is open, showing 'Replace (Generate) Peering Certificates'. Below the dropdown, there is a description: 'Replace the peering certificate (used to secure the inner channel between appliances) by generating new certificates. Otherwise, a policy push must be initiated to all affected appliances (i.e. each appliance for which the peering certificate excludes any affected appliances, SSL optimization to those appliances will not work correctly.)'. Below this is a 'Certificate Signing Request' form. The form has the following fields: 'Common Name:' with radio buttons for 'Appliance Hostname / IP Address' and 'Custom: vcx101.example.com'; 'Organization Name:' with the value 'Riverbed Technology, Inc'; 'Organization Unit Name:' with the value 'Steelhead'; 'Locality:' with the value 'Singapore'; 'State:' with the value 'SG'; 'Country:' with the value 'SG' and a note '(two-letter code)'; 'Email Address:' with the value 'admin@example.com'; and 'Validity Period:' with the value '730' and a note 'Days (60 to 3650 days)'. At the bottom of the form, there is a note: '(Note that this operation only applies to SteelHeads.)' and a 'Replace' button.

4. Specify the certificate details.
5. Select the appliance in which you want to replace the secure peering certificate and click **Replace**.
The secure peering certificate on the SteelHead is not replaced until a Push operation is completed.
6. Select Push Policies from the drop-down menu.

7. Select the same appliance and click **Push** (Figure 4-20).

Figure 4-20. Push policy page

Appliances Topology > Appliances ?

+ New Group + New Appliance - Remove Selected - Move Selected Filter Appliance Operations

Choose an operation to perform against the selected groups and appliances:

Push Policies

Push policies and appliance pages to the selected appliances.

☐ Include Path Selection, QoS and Applications configuration for RiOS 9.0 and later.

☐ Restart Optimization Service If Required

☐ Restart QoS Service If Required

☐ Schedule Deferred Push

Date and Time: 2014/11/10 06:57:30 YYYY/MM/DD HH:MM:SS

(Note that this operation only applies to SteelHeads and Interceptors.)

Push

8. From the Management Console, choose Manage > Operations: Operations History to verify that the operations are successful.
9. From the Management Console, choose Administration > Security: Certificate Authority to view the peering certificate issued by the SCC CA (Figure 4-21).

Figure 4-21. Certificate history

History of certificates issued

timestamp	subject
2014/11/10 06:56:01	vcx101.example.com

Certificate Details

Issued To

Common Name: vcx101.example.com

Email: admin@example.com

Organization: Riverbed Technology, Inc

Organization Unit: Steelhead

Locality: Singapore

State: SG

Country: SG

Issued By

Common Name: cmc-ca.example.com

Email: admin@example.com

Organization: Riverbed Technology, Inc

Organization Unit: SteelHead

Locality: Singapore

State: SG

Country: SG

Validity

Expires On: Nov 9 06:56:01 2016 GMT

Fingerprint

SHA1: 52418FCD7AC318744E289A91AD3E71FD

Key

Type: RSA

Size: 3072

10. From the SteelHead Management Console, verify that the secure peering certificate has been replaced.

Note: The SCC CA is automatically configured as a Trusted Entity only when an SCC secure peering policy is pushed to a SteelHead. The policy must include a peer SteelHead that has a secure peering certificate issued by the SCC CA.

For more information about SSL, see the *SteelHead Deployment Guide - Protocols*.

To use SCC to configure secure peering relationships

1. From the Management Console, choose Manage > Services: Policies.
2. Select Add Policy and specify a policy name (Figure 4-22).
3. Click Add.

Figure 4-22. Add policy page

Policies Services > Policies ?

☒ Add Policy ☐ Remove Selected Policies

Policy Name:

☒ Add a new policy
☐ Add a copy of an existing policy
☐ Import policy from an appliance configuration
☐ Merge existing policies

Description:

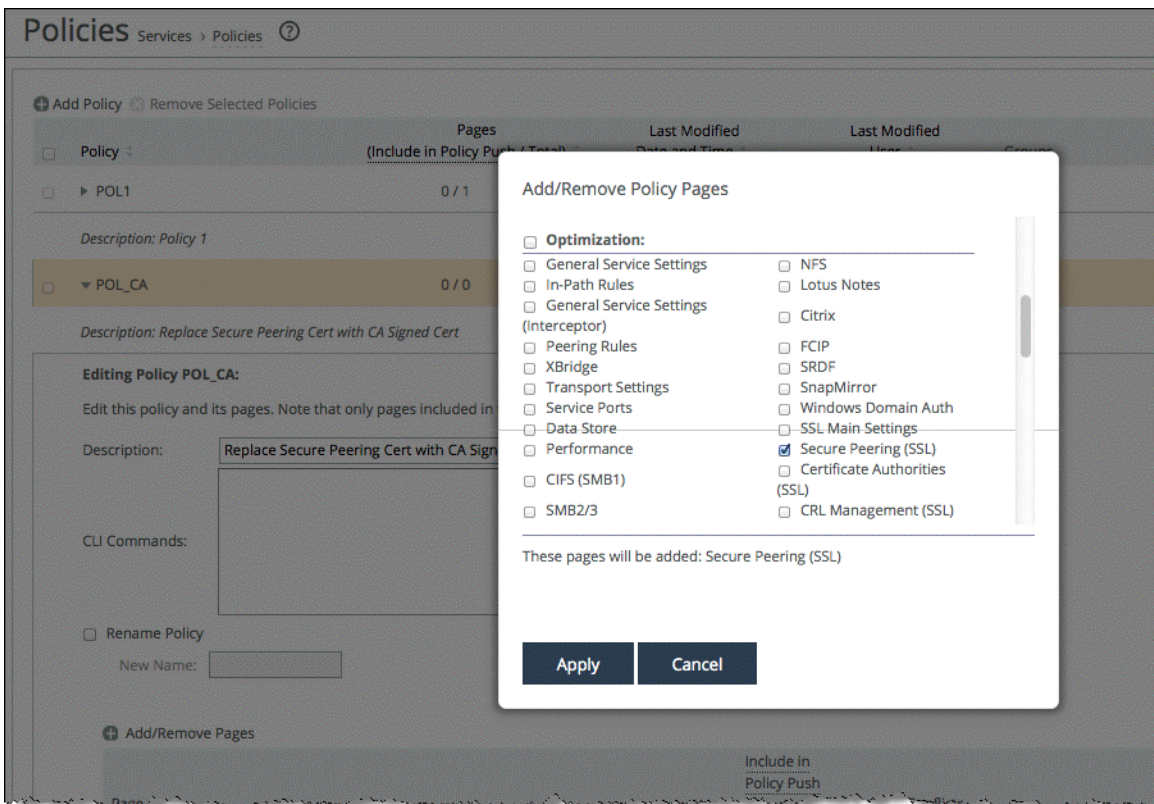
CLI Commands:

Add

4. Select Add/Remove Pages.

5. Select Secure Peering (SSL) in the Add/Remove Policy Pages and click **Apply** (Figure 4-23).

Figure 4-23. Add/Remove policy



6. Edit the Secure Peering (SSL) policy to configure a Trusted Peer SteelHead and click **Apply** (Figure 4-24).

Figure 4-24. Editing policy

Editing Policy: POL_CA , Secure Peering (SSL) ⓘ

SSL Secure Peering Settings

Traffic Type:

☒ Fallback to No Encryption (Does not apply to SSL only traffic)

Apply

Trusted Peering CAs & Peer Certificates:

+ Add a New Trusted Entity - Remove Selected

Trusted Entity : Issued To : No trusted entities.

Mobile Trust:

+ Add a New Mobile Entity - Remove Selected

Trusted Entity : Issued To : No current Mobile Entity

Trusted Peers

☒ Trust Selected Peers (only SSL-capable or disconnected appliances are shown):

Appliance	Product/Model	Version	<input type="checkbox"/> Trusted	Expiration Date	Certificate
VCOX / 172.16.220.129	VCOX55SH	9.0.0-mainline#258	<input type="checkbox"/>	Nov 9 07:07:57 2016 GMT	vcx201.example.com

☐ Trust All Peers

Update

7. Ensure that this policy page is included in the policy push.

8. Click **Apply**. (Figure 4-25).

Figure 4-25. Include setting in policy push

Editing Policy:

Page:

Copy Page Contents From Policy: **Copy**

This page is **included** in the policy push. **Exclude**

[Go to "POL_CA1" Policy Page](#)

Editing Policy: POL_CA , Secure Peering (SSL) ⓘ

SSL Secure Peering Settings

Traffic Type:

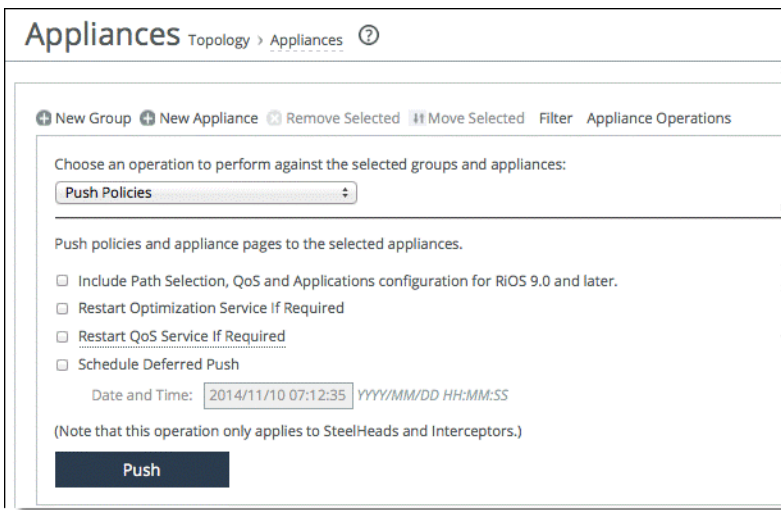
☒ Fallback to No Encryption (Does not apply to SSL only traffic)

Apply

9. Choose Manage > Topology: Appliances to associate the policy with the SteelHead.

10. From the Appliance Operation page, select Push Policies from the drop-down menu, and push the policy to the SteelHead (Figure 4-26).

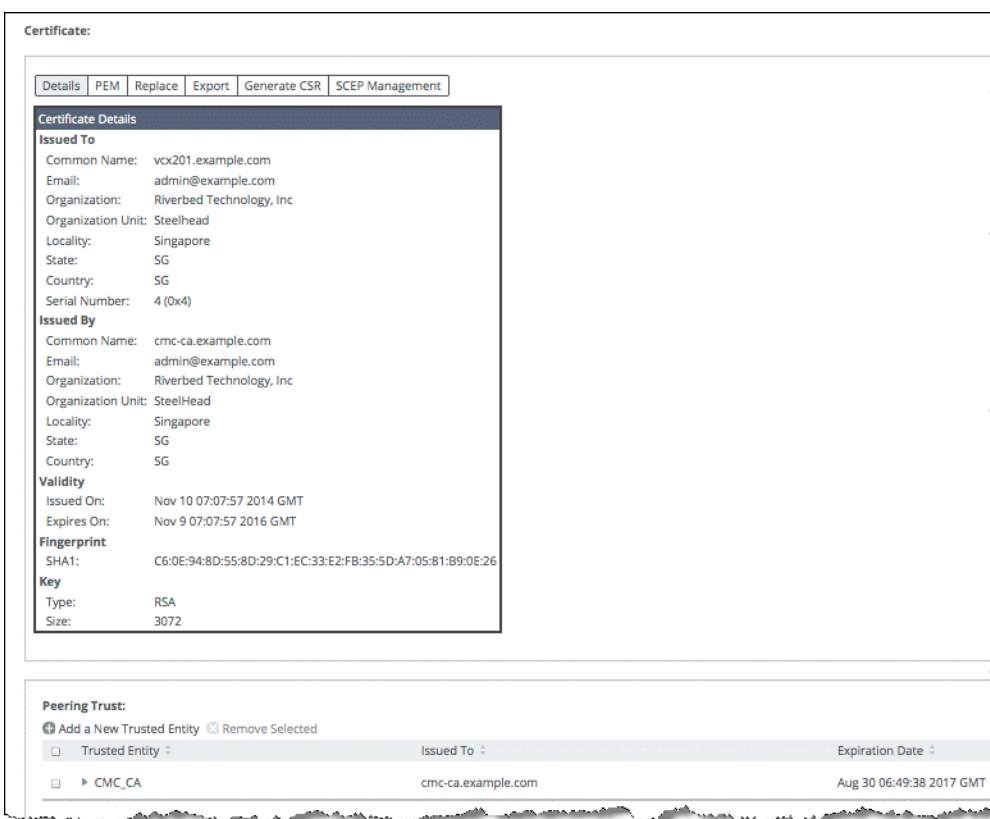
Figure 4-26. Pushing policy



11. Choose Manage > Operations: Operations History to verify that the operation is successful.
12. Verify from the SteelHead that the SCC CA is configured as a Trusted Entity.

This SCC CA is configured as a Trusted Entity on the SteelHead only if the peer SteelHead has a secure peering certificate that's issued by the SCC CA (Figure 4-27).

Figure 4-27. SCC CA on the SteelHead



Note: We recommend that you maintain a uniform and consistent operational policy by using only one type of secure peering certificates across all your SteelHeads.

Using the SCC CA service to issue proxy certificates for SSL optimization

You can use the CA service of SCC 9.5 or later to perform these two functions for SSL optimization:

- Automate the issue of proxy certificates for bypassed servers.
- Issue proxy certificates.

Before SCC 9.5, you can only use the SCC to import a proxy certificate or to generate a self-signed proxy certificate.

To use the SCC CA service to issue a proxy certificate for bypassed servers

1. From the Management Console, choose Manage > Policies.
2. Add a new policy.
3. Add the SSL Main Settings page to the policy.

4. Edit the SSL Main Settings page and include it in the policy push.

Figure 4-28. Editing SSL Main Settings page

Editing Policy:

Page:

Copy Page Contents From Policy:

This page is **included** in the policy push.

[Go to "pol1" Policy Page](#)

Editing Policy: pol1, SSL Main Settings ?

General SSL Settings

☐ Enable SSL Optimization

SSL Server Certificates:

Name	Issuer	Issued To
No current SSL		

5. From the Management Console, choose Manage > Appliances.

6. Associate the SteelHead appliances to the policy.

Figure 4-29. Associating appliance to a policy

Appliances Topology > Appliances ?

<input type="checkbox"/>	Groups and Managed Appliances	Product / Model	Connection	
<input type="checkbox"/>	Global			
<input type="checkbox"/>	SteelHead / VCX555M (10.6.7.101)	SteelHead	VCX555M	Connected: Needs Attention

Policy	Description
pol1	

7. From the Management Console, choose Manage > Policies and go back to the policy.
8. Edit the SSL Main Settings page of the policy and select Generate Certs for Bypassed Servers.

Figure 4-30. Selecting Generate Certs for Bypassed Servers in the policy

Editing Policy:

Page:

Copy Page Contents From Policy:

This page is **included** in the policy push.

[Go to "pol1" Policy Page](#)

Editing Policy: pol1, SSL Main Settings ?

General SSL Settings

☐ Enable SSL Optimization

SSL Server Certificates:

Name	Issuer	Issued To
No current SSL		

Note: A new service on SCC queries the SteelHead once every 6 hourly for the list of bypassed SSL servers. This list is displayed when you select Generate Certs for Bypassed Servers.

9. Select the bypassed server(s) that you would like to issue a proxy certificate for and click **Generate Certs**.

Figure 4-31. Selecting bypassed servers

Editing Policy:

Page:

Copy Page Contents From Policy:

This page is **included** in the policy push.

[Go to "pol1" Policy Page](#)

Editing Policy: pol1, SSL Main Settings ?

General SSL Settings

☐ Enable SSL Optimization

SSL Server Certificates:

Name	Issuer	Issued To
No current SSL Certificates		

Generate SCC CA Signed Certificates for Bypassed Servers

Bypassed Servers:

<input type="checkbox"/> Server Name	Appliance Name	Covered By
<input checked="" type="checkbox"/> oak-vcs1890	oak-vsh229	
<input type="checkbox"/> oak-vcs1891	oak-vsh245	
<input type="checkbox"/> oak-vcs1892	oak-vsh245	

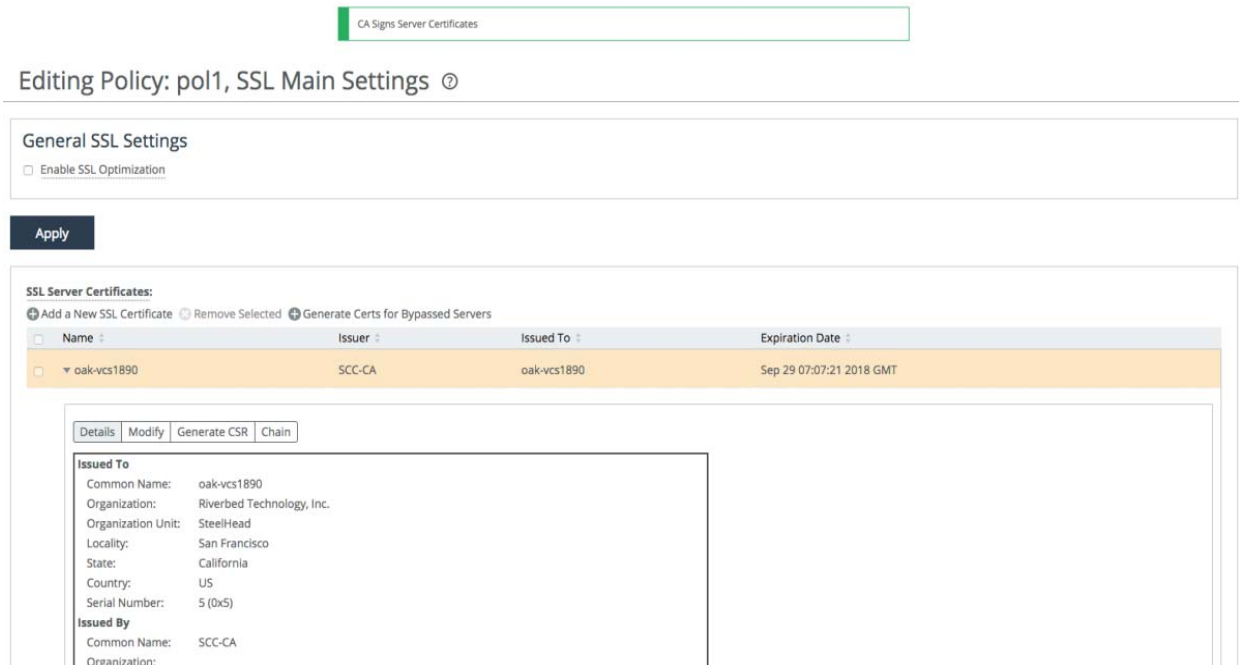
#Items are disabled if:

1. Wild card certificate exists for the bypassed server
2. Certificates are generated and yet to be pushed to the appliance
3. Certificates have been pushed recently

Note: To generate proxy certificate using Generate Certs for Bypassed Servers wizard, the SCC CA certificate must have a validity of greater than 60 days.

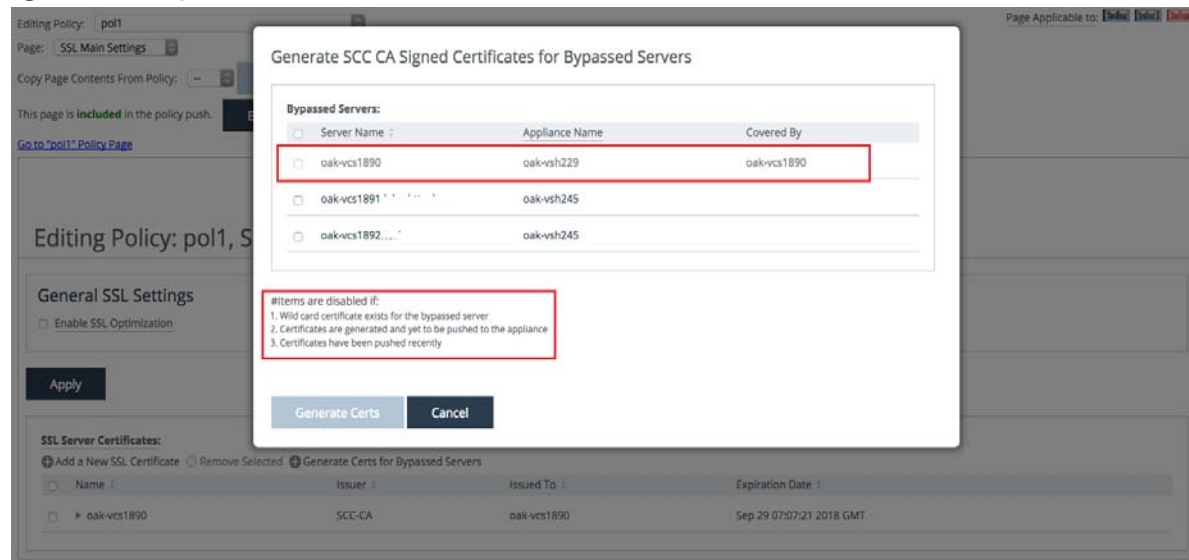
10. Verify that the Proxy Certificate is successfully generated.

Figure 4-32. Verifying the proxy certificate is generated



11. At this point, if you click **Generate Certs for Bypassed Servers** once again, you will see that the server that you have just generated a proxy certificate is grayed out. This means that the proxy certificate is awaiting a push operation to the SteelHead.

Figure 4-33. Greyed-out servers



12. From the Appliance Operation page, select Push Policies from the drop-down menu, and push the policy to the SteelHead.

13. Choose Manage > Operations: Operations History to verify that the operation is successful.

To use the SCC CA service to issue a proxy certificate for SSL optimization

1. It is also possible to use the SCC CA service to simply issue a proxy certificate.

Figure 4-34. SCC CA signed proxy certificate

The screenshot shows the 'SSL Server Certificates' configuration page. At the top, there are three radio buttons: 'Add a New SSL Certificate' (selected), 'Remove Selected', and 'Generate Certs for Bypassed Servers'. Below this, the 'Name' field is set to 'oak555'. There are three radio buttons for the certificate type: 'Import Certificate and Private Key', 'Generate Self-Signed Certificate and New Private Key', and 'Generate SCC CA Signed Certificate and New Private Key' (selected). The 'Generate SCC CA Signed Certificate' section contains the following fields: 'Common Name' (oak555.example.com), 'Organization Name' (Riverbed Technology, Inc.), 'Organization Unit Name' (Steelhead), 'Locality' (Singapore), 'State' (SGP), 'Country' (SG, with a note '(two-letter code)'), 'Email Address' (admin@example.com), and 'Validity Period' (730 Days (60 to 3650 days)). An 'Add' button is located at the bottom left of the form.

Note: As with all SCC policy push operations, the existing settings of the managed SteelHead appliance are replaced with that of the SCC policy. When a policy is used to manage the proxy certificates on a SteelHead, any proxy certificates that are installed locally on the SteelHead but not in the SCC policy will be lost after the policy push. We recommend that you review the proxy certificates that are already installed in the SteelHead and copy them to the SCC, if necessary, before you perform a proxy certificate push operation from the SCC.

Sites and Networks, Path Selection, and QoS

This chapter describes how to configure sites and networks, path selection, and QoS for the SCC. It includes the following sections:

- “Overview of sites and networks, path selection, and QoS” on page 61
- “Configuring sites and networks” on page 62
- “Configuring path selection” on page 69
- “Configuring global QoS” on page 73

This chapter requires you be familiar with path selection, Quality of Service (QoS), topology, sites, and networks. For more information, see the *SteelHead Deployment Guide* and the *SteelCentral Controller for SteelHead User’s Guide*.

Overview of sites and networks, path selection, and QoS

The new Sites & Networks configuration page combines a set of parameters that enables SteelHeads to build their view of the WAN. The network configuration defines the WAN clouds, and sites include the representation of the IP address network and definition of how the appliances connect to the various defined networks.

With the architectural model of *topology*, SteelHeads can automatically build the various paths to other remote sites (a function needed for path selection and secure transport) and calculate the bandwidth available on these paths—a function needed for QoS.

You must be running SCC 9.0 or later to configure sites and networks on SteelHeads running RiOS 9.0 or later.

In RiOS 9.0, Riverbed introduces a new process to configure path selection and QoS on SteelHeads. This new process greatly simplifies the configuration and administration efforts, compared to earlier releases. As such, this new process relies heavily on a common central configuration, which is best managed and configured using a central management: the SCC.

We recommend that you use the SCC to manage Riverbed QoS, path selection, and secure transport configurations for the following reasons:

- **QoS and path selection with RiOS 9.0** - Starting with SCC 9.0, you can deploy a single global QoS and path selection configuration to all the managed SteelHeads running RiOS 9.0 and later. This new configuration is in contrast to earlier RiOS releases in which you configured multiple SteelHead policies for QoS and path selection.

- **Secure transport** - The secure transport feature requires you to use SCC 9.0 or later and RiOS 9.0 or later.

For more information about secure transport, see [“Secure Transport” on page 97](#).

Configuring sites and networks


This section discusses a high-level configuration of sites and networks. For more information about sites and networks, see the *SteelHead Deployment Guide* and the *SteelCentral Controller for SteelHead User's Guide*.

The topology process in SCC 9.0 focuses on ease of deployment and manageability. As such, SCC focuses on a central configuration that you can deploy across the entire SteelHead enterprise—a *configure once, deploy to many* concept. This configuration process relies on making effective use of templates when possible to ease the repeatability of configuration steps.

When you start a configuration for path selection, including secure transport or QoS, you must have a well-planned design vision of the overall network. This section uses the example network shown in [Figure 5-1](#).

[Figure 5-1](#) shows multiple dual-homed sites with an MPLS WAN link provided by a carrier, labeled MPLS, and the second link is a private VPN circuit. A third site, RemoteBranch3, has a single connection back to the MPLS cloud and a secondary link through an internet-based firewall. All traffic, including public internet, is backhauled through the main headquarters site and egresses directly through the firewall connection.

Figure 5-1. Example network scenario



The following steps are a high-level configuration for sites and networks:

1. Identify the networks you want to configure: MPLS, VPN, and Internet.
2. Configure the multiple sites.

Site configuration contains basic information such as site name, contact information, SteelHead (if any), and network information.

For the example shown in [Figure 5-1](#), you need to add the following sites: RemoteBranch1, RemoteBranch2, RemoteBranch3, and Headquarters. As part of the configuration, notice that RemoteBranch1 and RemoteBranch2 have similar uplinks, and you can use a connectivity template to ease the configuration.

You can initiate site configuration by:

- manual single site creation ([page 65](#)).
- bulk site migration using the CSV template ([page 68](#)).

Manual site creation is the simplest, but you've to configure one site at a time. The bulk site migration requires you to complete a CSV file entry and then upload to the SCC as a one-time push.

Site creation is a required task even if you've an existing SteelHead/SCC deployment. Current appliances aren't added to sites as part of an upgrade from earlier releases, and they aren't automatically discovered as part of a new deployment. SCC maintains a separate inventory of the appliances for both management and topology features. SteelHeads need to be identified under sites inventory as well as the management if using path selection or QoS.

For more details about migration, see the *SteelCentral Controller for SteelHead Installation Guide*.

After you've defined the sites and networks, you can continue to configure QoS, secure transport, and path selection independently.

For more information about configuring path selection, see [“Configuring path selection” on page 69](#). For more information about configuring QoS, see [“Configuring global QoS” on page 73](#). For more information about secure transport, see [“Secure Transport” on page 97](#).

Note: The following configuration steps apply only for SCC 9.0 and later and SteelHead RiOS 9.0 and later. For earlier RiOS versions, configuration is still managed through classic policies. SCC 9.0 enables you to maintain both configuration features simultaneously to manage your appliances across the enterprise. If you're running versions earlier than 9.0 for RiOS and SCC, see the documentation for the appropriate release on the Riverbed Support site.

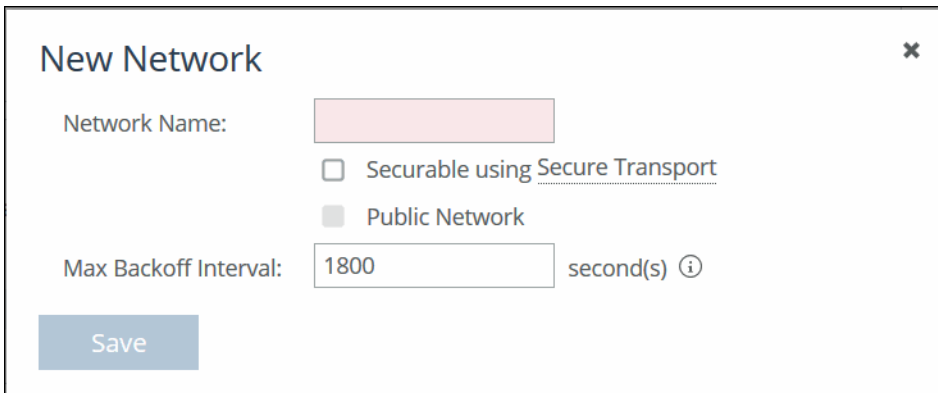
Note: You must enable REST API on the SteelHeads prior to configuring Topology configuration. REST API is enabled by default. For more information on REST API, see the *SteelHead Deployment Guide*.

To configure the SCC to manage networks

1. Choose Manage > Topology: Sites & Networks.
2. Select Add a Network.

3. Specify the information for all known WAN clouds (Figure 5-2).

Figure 5-2. Add a network



The screenshot shows a 'New Network' dialog box with a close button (X) in the top right corner. It contains the following fields and options:

- Network Name:** A text input field.
- Securable using Secure Transport:** A checkbox.
- Public Network:** A checkbox.
- Max Backoff Interval:** A text input field containing the value '1800', followed by the text 'second(s)' and an information icon (i).
- Save:** A blue button at the bottom left.

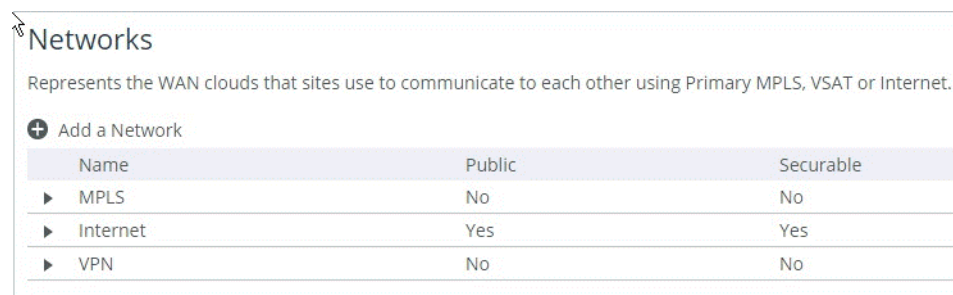
- **Network Name** - Specify the name of your network. A network represents a shared communication domain. In other words, a network is any site with uplinks to a network, and can communicate to other sites on the same network. SCC is prepopulated with two networks: MPLS and Internet.
- **Securable using Secure Transport** - Select this check box if you want to secure data communication by enabling encryption of traffic as it traverses this WAN cloud.
- **Public Network** - Select this check box to specify to the SteelHead that the network is a public network. If the network is a private network, do not select the check box. If you define your network as a public network by selecting this check box, the SteelHead assumes that the traffic sent to this network traverses a device that performs NAT. The public network option only takes effect if you also configure secure transport.

For more information about secure transport, see [“Secure Transport” on page 97](#).

- **Max Backoff Interval** - Specify the number of seconds of the probing frequency. This setting reduces the number of probes on links that aren't frequently used. This value indicates the maximum probing frequency if there's no traffic detected on the path and reduces the number of probes initiated by the SteelHead. If the path experiences any traffic, the configured uplink probe frequency value is assumed (default value is 2 seconds). As traffic lessens, a gradual exponential backoff in probing frequency begins and continues until the maximum value of probing with no traffic is reached. If there's a path failover, the probe timeout value (default of 2 seconds) is assumed.

To configure networks for the example shown in [Figure 5-1](#), you need to configure the following networks: MPLS, VPN, and Internet. [Figure 5-3](#) shows how the network table looks upon completion.

Figure 5-3. Configured networks



Name	Public	Securable
MPLS	No	No
Internet	Yes	Yes
VPN	No	No

Note: RiOS 9.2 and later support up to 500 configured sites as part of the overall topology.

To manually configure the SCC to manage a single site

1. Choose Manage > Topology: Sites & Networks.
2. Select Add a Site.
3. Specify the information for the site:
 - **Site Name** - Specify a name to replace the Local site field on the SteelHead belonging to that site.
 - **Site Type** - Used to identify the site's operational purpose and used for organizational purposes. By default, three types are constructed: Branch, Data Center, and Headquarters. To create a new site type, enter the new site type name.
 - **Region** - Specify a region for organizational purposes. A region enables you to group appliances based on location.
 - **Description and Contact information** - Specify optional information for site identification.
 - **Network Information** - Specify the subnets local to this site.
 - **Internet Traffic** - Describes how this site accesses the internet (for example: directly through a local gateway or through a backhauled connection to another site, such as a data center or hub site).

You must configure each site correctly depending on the type of deployment you've. If Direct-to-Internet, SCC automatically creates a rule to relay and not use a path selection rule on Internet-bound traffic from that site. Normal routing is then expected to provide access to the internet.

For sites that access the internet through a backhauled remote location, SCC prompts for that remote site. SCC uses the peer of that remote site to monitor for path availability. You can configure a path selection policy for internet-bound traffic between sites. Note that backhauled sites are considered Direct-to-Internet and we recommend that you configure them as such.

- **Riverbed appliances** - Specify the SteelHead local to this site, if any. The appliance must be registered (that is, added as an appliance) in SCC. This field, when selected, lists all available managed appliances to select from.
- **Custom Probe Endpoint** - Specify an IP address to use as the probe point for site availability. If you do not specify an IP address, the SCC automatically assigns the in-path IP addresses of the SteelHeads that are part of the specific site you're configuring. You can enter as many IP addresses as needed.

The endpoint IP address is not required as part of the LAN-side IP addressing. An example use case for a custom probe endpoint is when you want your spoke to probe only the hub and no other spokes (that is, a true hub-and-spoke design without the need for full mesh connectivity). You can point all the spokes to probe a single IP address in the data center to ensure that no probing traverses the WAN between the spokes.

- **Uplinks** - Specify a new uplink.

An uplink connects a site to a network. A site can have one or more uplinks to the same network and can connect to multiple networks. If using QoS you must specify, per uplink, the bandwidth available for uploading and downloading data. The bandwidth statement serves for QoS throughput enforcement.

Uplinks are used for path selection traffic steering. You can configure uplinks manually or from a preconfigured uplink template. If the uplink you're configuring also shares the physical interface with another uplink, you must select one as Default for inpathX_X so that this uplink is considered the default gateway for the link. This latter configuration doesn't match a direct configuration on SteelHeads.

Configuring an IP address on the gateway IP setting enables the SteelHead to use the IP address you enter as the next hop gateway for traffic to be sent down a specified path (for path selection purposes). If you leave the field blank and enable the default for inpathX_X setting, the SteelHead uses the configured default gateway IP address of the associated in-path interface. We recommend you use the latter for parallel deployments in which the SteelHeads might share the same in-path interface slot number but with differing gateway IP addresses. Make sure to have the in-path default gateway IP address pointing to the WAN-side IP address instead of the LAN-side gateway.

In SCC 9.2 or later, you've the option to set a bandwidth limit on the number of probes consumed per uplink. This field sets the maximum rate limit of the probes this uplink initiates. The limit is only in the outbound direction and doesn't enforce inbound probes. Be cautious using this setting because when you limit the throughput for probing, you essentially limit the number of probes. Limiting probes can create a longer than expected failover time.

For more information about uplinks, see the *SteelHead Deployment Guide 9.0* or later and the *SteelCentral Controller for SteelHead User's Guide*.

You can configure the primary interface as an uplink for sites that can be used in QoS (path selection doesn't apply).

- **Secure Transport Concentrator** - Select the SteelHead at the site that you want to use as a secure transport concentrator, and provide the uplinks to connect to securable networks.

A secure transport concentrator serves as an appliance to provide encryption services.

For more information about secure transport concentrator, see [“Secure transport concentrator” on page 108](#).

Figure 5-4 shows how the Sites table looks upon completion for the example shown in Figure 5-1. We recommend that you start with the headquarters (data center) that's the backhauled site, because you need to reference it for the branch sites.

Figure 5-4. Configured sites

Name	Type	Region	Appliances	Uplinks	
Headquarter	Data Center	Americas	0	3	View Site Edit Site
RemoteBranch1	Branch	Americas	0	2	View Site Edit Site
RemoteBranch2	Branch	EMEA	1	2	View Site Edit Site
RemoteBranch3	Branch	APAC	0	2	View Site Edit Site

If the headquarters site is backhauling all internet traffic, you select the site as a Direct-to-Internet site. In the example shown in Figure 5-1, the headquarters is the data center in which all the servers are located. All branch offices are set to backhaul through the headquarter site.

The Direct-to-Internet designation signals to the SCC to create a top-level path selection rule to be applied at that specific site you're configuring. This automatically generated path selection rule denotes to relay any traffic with a destination to the default 0.0.0.0/0 site, and it is placed at the start of the path selection rule set.

The rule is needed to trigger on Internet-bound traffic that doesn't have a specific site destination, hence the 0.0.0.0/0 (catch all) site matching. This rule allows the traffic to be processed across the SteelHead without exerting any path selection logic on it. You want Internet-bound traffic to traverse to a local internet connection instead of being backhauled through the main data center access (Figure 5-5).

Figure 5-5. Data center internet traffic setting

Internet Traffic

Designate how Internet access is achieved at this Site. Only applies when Path Selection is enabled.

☐ Direct to Internet
☒ Backhaul through Site

headquarters

As part of the site configuration, you can construct a site connectivity template (Manage > Topology: Sites & Networks), which you can reuse across multiple sites that share the same configuration. For the example shown in [Figure 5-1](#), you can create a template ([Figure 5-6](#)) to reuse for RemoteBranch1 and RemoteBranch2 because they have a common uplink configuration.

Figure 5-6. Site connectivity templates

+ Add New Uplink

Uplink 1

Network: Internet

Uplink Type: untrusted ⓘ

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface: inpath0_1

☒ Is Default for inpath0_1 ⓘ

☐ Enable GRE Tunneling ⓘ

Bandwidth Up: 1

Mbps

Bandwidth Down: 1

Mbps

Uplink bandwidth only affects QoS

► Probe Settings ⓘ

Uplink 2

Network: MPLS

Uplink Type: trusted ⓘ

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface: inpath0_0

☒ Is Default for inpath0_0 ⓘ

☐ Enable GRE Tunneling ⓘ

Bandwidth Up: 1

Mbps

Bandwidth Down: 1

Mbps

Uplink bandwidth only affects QoS

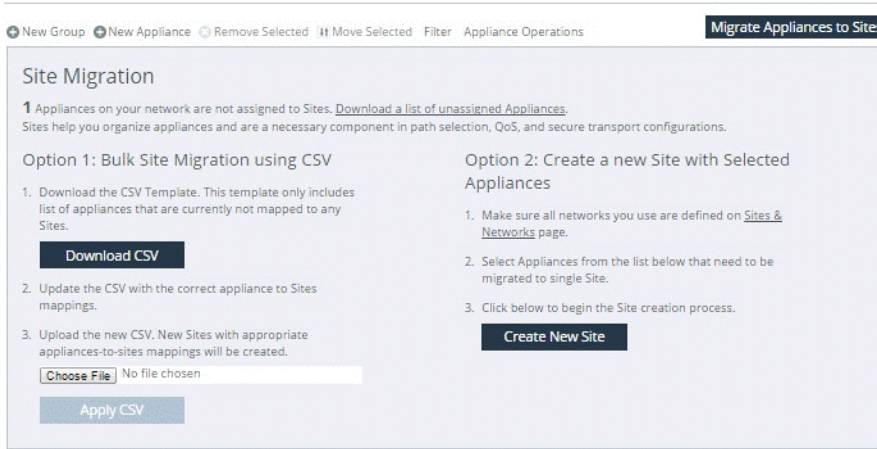
► Probe Settings ⓘ

To use CSV for bulk site migration

1. Choose Manage > Topology: Appliances.
2. Click **Migrate Appliances to Sites**.

The Site Migration page opens (Figure 5-7).

Figure 5-7. Download CSV



3. Click **Download CSV**.
4. Open the downloaded file and complete the required entries to add the various sites. The information you provide is identical to the site information requested in “To manually configure the SCC to manage a single site” on page 65.
5. After you’ve completed the information, upload the CSV file and click **Apply CSV**.
If there’s an error in the CSV file, those rows are silently ignored and the migration continues at the next correct row.

After you’ve configured the sites and network, you can complete your path selection and QoS configurations, referencing the topology you’ve just built.

Configuring path selection

The path selection configuration process involves choosing applications to steer onto uplinks as defined in sites and networks and then enabling the path selection feature. You can select the sites that receive the common set of configured path selection rules. Be aware that there can only be one global path selection to be pushed out to the SteelHeads. You can’t configure a path selection rule based on different destinations or different path selection rules to different SteelHeads.

Note: If you’re using path selection in a release earlier than SCC 9.0, see earlier versions of the *SteelCentral Controller for SteelHead Deployment Guide* on the Riverbed Support site at <https://support.riverbed.com>.

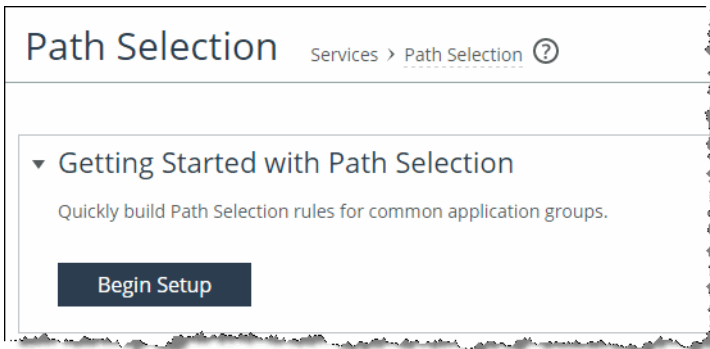
Note: After upgrading from a 9.x version of RiOS to 9.2, the first policy push from SCC can cause pre-existing path selected connections to be blocked and/or QoS shaped connections to be misclassified. For more information, go to <https://supportkb.riverbed.com/support/index?page=content&id=S28250>.

To configure path selection on the SCC

1. Choose Manage > Services: Path Selection.

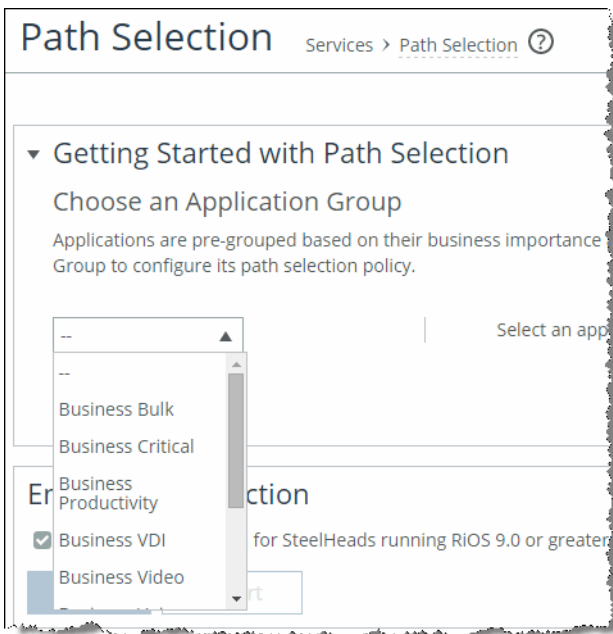
2. Click **Begin Setup** (Figure 5-8).

Figure 5-8. Begin path selection setup



3. Select an application group to steer, and click **Configure** (Figure 5-9).

Figure 5-9. Path selection applications



4. In the Uplink Priority page, select the uplinks you want to steer this application group onto (Figure 5-10).

You can manipulate the order of the uplinks.

Figure 5-10. Path selection uplinks

▼ Getting Started with Path Selection

Uplink Priority for Business Bulk

Select and re-order the desired uplinks.

- ☒ Uplink Type
- ☐ MPLS
- ☐ MPLS (secured)
- ☐ Internet
- ☐ Internet (secured)

If all checked uplinks are down: Relay ▼

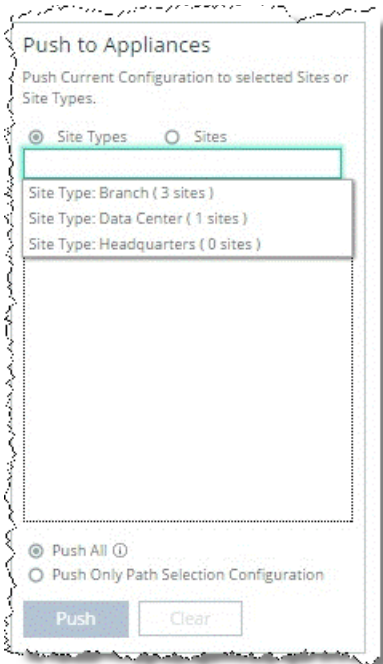
Save Rule Back

5. Click **Save Rule**.
6. Repeat **Step 4** and **Step 5** to add additional rules if needed.
7. In the Enable Path Selection page, select Enable Path Selection for SteelHeads running RiOS 9.0 or later and click **Save**.
8. Select the SteelHeads to which you want to push the path selection rules.
9. Click **Include in Push**.
10. Select the SteelHeads you want to include as part of the configuration push.

You can select SteelHeads based on site type or per site.

You've the option to push only path selection rules or include the sites and networks configuration. Push the entire topology (sites and networks) configuration if there was a change made to the topology or application section ([Figure 5-11](#)).

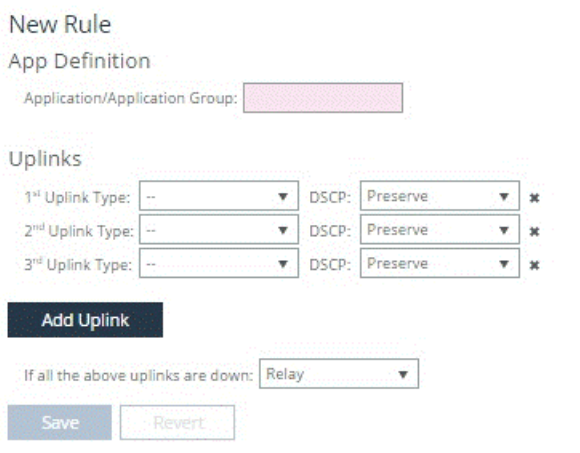
Figure 5-11. Push to appliances



11. Optionally, to use path selection for a specific or custom application, you can select Add a Rule to create a new rule ([Figure 5-12](#)), or select the rule itself to edit it.

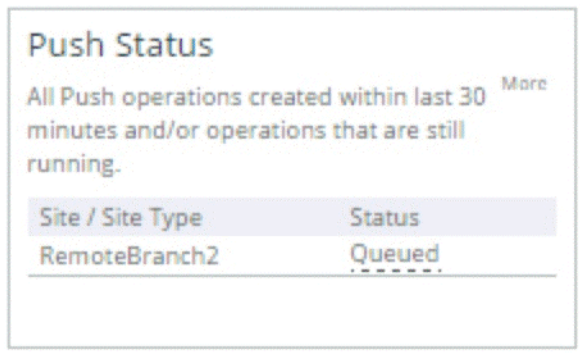
You can select up to 16 uplink types for a path selection rule. SCC pushes the first three uplinks to the SteelHead that apply to that site.

Figure 5-12. New rule



After you've pushed the changes, the Push Status window shows additional detail on the operation (Figure 5-13).

Figure 5-13. Push status



Push Status	
All Push operations created within last 30 minutes and/or operations that are still running. More	
Site / Site Type	Status
RemoteBranch2	Queued

Configuring global QoS

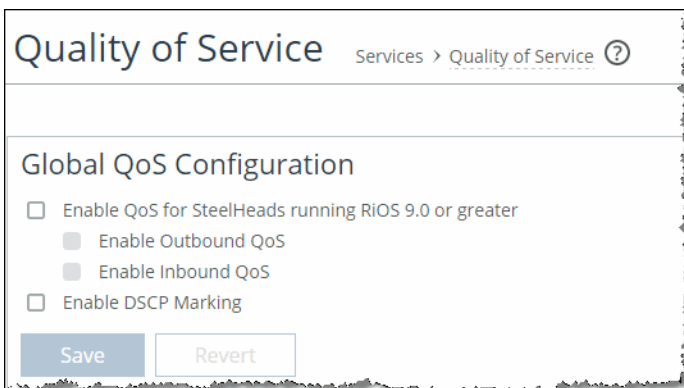
To configure QoS, you need to enable QoS as a feature and build QoS profiles based on source and destination sites. The profile includes QoS rule configurations. The final task is to choose certain SteelHeads to obtain this common configuration.

To configure global QoS on the SCC

1. Choose Manage > Services: Quality of Service.
2. Enable QoS for SteelHeads running RiOS 9.0 or later.

You can also select to enable outbound QoS, inbound QoS, and DSCP marking (Figure 5-14).

Figure 5-14. Global QoS configuration



Quality of Service [Services > Quality of Service](#) [?](#)

Global QoS Configuration

☐ Enable QoS for SteelHeads running RiOS 9.0 or greater

- ☐ Enable Outbound QoS
- ☐ Enable Inbound QoS

☐ Enable DSCP Marking

[Save](#) [Revert](#)

You've the option to override QoS interface settings on a per SteelHead basis. You might need to disable QoS configuration on an interface level because certain deployments do not support QoS across all SteelHead interfaces.

3. Select Add a QoS Profile.

Note: The following step is specific to SCC only and is not available directly on the SteelHead.

4. When you configuring a new profile, you must indicate the source and destination sites (Figure 5-15).

You can configure the source and destination sites per site or by types of site that have already been defined in the sites and networks. You also have the option to select Any Site to indicate that the profile is applied to all known predefined sites.

This configuration example doesn't have an equivalent configuration directly on a SteelHead. Unlike the path selection single global configuration, the SCC can deploy a different QoS configuration based on the site or the site types. A global QoS configuration is not capable of differentiating between different types of sites, requiring different QoS parameters. With these settings, SCC only applies the proper profile that matches the correct sites and site types the profile is tied to. During an SCC push, the site profile is matched to the SteelHead first, and if that profile is not present, then the site types profile is matched. If neither is present, then Any Site profile is matched.

Figure 5-15 indicates the source and destination sites as part of the QoS profile, and SCC then applies the proper profile based on the sites already configured in sites and networks.

Figure 5-15. New QoS profile


5. Choose either to start with new QoS classes and rules or copy from another existing profile. New installations generally require starting with new QoS classes. By default, SCC creates an Any Site-to-Any Site configuration; therefore, you can't create a new profile based upon a similar site delineation, although you're able to edit the default profile if needed.

6. Click **Create Profile**.

7. Select Edit Profile to manage the associated QoS classes and rules.

8. Click **Edit** to create the classes (Figure 5-16).

Figure 5-16. QoS classes



9. Click **add class** (Figure 5-17).

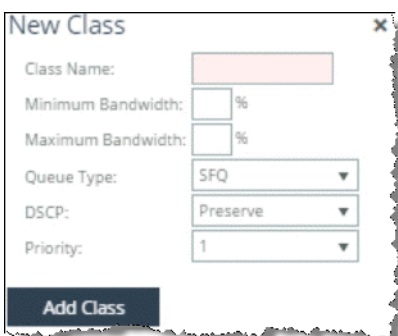
Figure 5-17. Adding a class



10. Specify a name and desired percentage for a minimum and maximum bandwidth range (Figure 5-18).

You can also alter the queue type, change the DSCP mark, and manipulate the priority level from 1 to 6 with 1 being highest.

Figure 5-18. New class



We recommend that you start with the default catch-all class parameters (as shown in [Figure 5-19](#)), and then configure the other additional classes.

Figure 5-19. Default class

New Class

Class Name:

Minimum Bandwidth: %

Maximum Bandwidth: %

Queue Type:

DSCP:

Priority:

Add Class

- After you've created all the classes, you must create the QoS rules and associate them to the classes you've just created ([Figure 5-20](#)).

QoS rules are application centric. You select an application or application group, choose the class from the list you had just completed, and optionally manipulate the DSCP marking for this traffic. The DSCP marking overrides whatever you set at the class level.

Figure 5-20. New rule

New Rule

Application or Application Group:


QoS Class:

DSCP Mark:

Add **Revert**

[Figure 5-21](#) shows a typical example of a QoS class structure and select rules.

Figure 5-21. QoS class structure and rules



Important: Figure 5-21 is not intended to be used as a reference guide for a QoS configuration.

12. Click Include in Push.

Push the configuration to the various SteelHeads in the enterprise. This process is similar to the Path Selection configuration push as described in [“To configure path selection on the SCC” on page 69](#).

13. Select the SteelHeads for which you want to obtain the QoS configuration based on site type or per site.

14. You’ve the option to push only QoS profiles or include the sites and networks configuration. Push the entire topology (sites and networks) and application configuration if there was a change made to the topology or application sections.

After you’ve pushed the changes, the Push Status window shows additional detail on the operation.

Web Proxy

This chapter describes the web proxy feature. Web proxy transparently intercepts HTTP traffic bound to the internet and provides acceleration services such as web caching, caching of eligible YouTube video content, Secure Sockets Layer (SSL) decryption and encryption services (for example, HTTPS) to enable encrypted content caching, and logging services through audit trails.

This chapter includes the following sections:

- [“Overview of the web proxy feature” on page 79](#)
- [“Web proxy fundamental properties” on page 80](#)
- [“Supported features” on page 80](#)
- [“Configuring basic web proxy features \(HTTP\)” on page 82](#)
- [“Advanced configurations for web proxy” on page 83](#)
- [“Troubleshooting web proxy” on page 93](#)

Overview of the web proxy feature

RiOS 9.1 and later include the web proxy feature. Web proxy uses the traditional single-ended internet HTTP-caching method enhanced by Riverbed for web browsing methodologies of today. The web proxy feature enables SteelHeads to provide a localized cache of web objects (or files). The localized cache alleviates the cost of repeated downloads of the same data. Using the web proxy feature in the branch office provides a significant overall performance increase due to the localized serving of this traffic and content from the cache. Furthermore, multiple users accessing the same resources receive content at LAN speeds while freeing up valuable bandwidth.

RiOS has embedded features such as strip-compression and parse-and-prefetch, which provide dual-ended SteelHead deployments with HTTP and HTTPS optimization abilities. For more information about strip-compression and parse-and-prefetch, see the *SteelHead Deployment Guide - Protocols*.

Riverbed enhanced the traditional HTTP caching feature on RiOS by including HTTPS data caching. Use of the web proxy feature differs from traditional SteelHead optimization because you do not need a server-side SteelHead for web proxy to intercept and optimize traffic. The web proxy feature is a true single-ended interception because you need only a branch-side SteelHead for accelerating HTTP-bound internet traffic.

Note: You can use web proxy while simultaneously optimizing different internal-bound traffic using any of the other available optimizations, such as the SteelHead HTTP optimization.

Web proxy fundamental properties

To use web proxy, you need the following appliances:

- **SCC** - The SCC is used to configure, manage, and maintain the web proxy feature on each SteelHead on which you've web proxy. Additionally, you can use the SCC to centrally view and monitor the cache-hit data collected across sites in which you've deployed web proxy.
- **SteelHead** - The SteelHead is usually located at the branch location and hosts the configurations created on the SCC. The SteelHead provides the proxy and cache services for each independent location.

The web proxy feature is currently only supported in a physical in-path deployment or a virtual in-path deployment (using WCCP or PBR) model. Web proxy is not supported on the xx50 models, xx60 models, or SteelHead-v.

Note: Web proxy is critically dependent on DNS resolution, specifically Reverse DNS lookups sourced from the Primary interface, for appropriate HTTP/HTTPS proxy services to occur. Because the SteelHead must successfully resolve hostnames to be cached and proxied the Primary interface of the SteelHead must be configured with valid IP address and DNS information. In addition, the interface must be in an active state (even when it is not used by your supported deployment model). Make sure that the SteelHead DNS configuration and the Primary interface on the SteelHead are both configured and active.

You can deploy a basic web proxy running on the branch office SteelHead specifically as a transparent forward proxy. In this deployment the client connections have no knowledge of the existence of the proxy. Because of this implementation, the client machines do not require any additional configuration like a proxy auto-config (PAC) file addition or the need to change the gateway address to point at the SteelHead (or to configure a specific proxy server address in their browser).

Beginning in SCC and RiOS 9.5 Riverbed can now support Proxy Chaining configurations to additional upstream transparent (Manual mode) or explicit (Automatic mode) proxy services (for example, Squid, Zscaler, etc.). Alternative proxy functionality such as reverse proxy services (for example, many inbound connections being proxied to few data center hosts) are not supported.

The SteelHead houses a separate logical data store to hold cache data for the HTTP and HTTPS content that the web proxy caches. In SCC 9.2 and later, web proxy caching is RFC 2616 compliant and persistent in that the cache data services a SteelHead reboot as a server restart. While the total cache data store size varies based on the model of SteelHead you deploy, the maximum single cacheable file size for SCC 9.2 and later web proxy releases is set as unlimited. Unlimited means that a single cache-eligible file can be as large as the amount of available cache.

The basic configuration for web proxy is to enable the SCC for the web proxy service and then choose which supportable branch locations to enable web proxy on for the configuration update. You can additionally choose to enable HTTPS optimization and define a global whitelist of HTTPS domains that you can access from the HTTPS configured locations.

Note: HTTPS optimization assumes that you've configured the SCC for certificate authority (CA) service. For more information, see ["Configuring and Using the SCC as a Certificate Authority Service" on page 45](#).

Supported features

The web proxy feature supports the following features:

- [“IP addressing support” on page 81](#)
- [“TCP port support” on page 81](#)
- [“Web proxy and SteelHead SaaS” on page 81](#)
- [“Video caching” on page 81](#)

IP addressing support

Web proxy, by default, supports proxying connections that use public IPv4 addresses (non-RFC 1918 IP addresses, or private reserved IP addresses). RFC 1918 addressing standard uses the SteelHead HTTP optimization and doesn't service the web proxy feature without additional configuration. You can choose to create specific in-path rules to force RFC 1918 IP addresses to use the web proxy service. For more information, see [“Using in-path rules” on page 91](#).

TCP port support

While TCP ports 80 (HTTP) and 443 (HTTPS) are the only ports supported by the web proxy default configurations, you can also configure nonstandard TCP ports supporting HTTP/HTTPS ports (8080/8000) to use web proxy. You configure these ports on the branch office SteelHead to bypass the HTTP optimization by creating additional in-path rules. For more information, see [“Using in-path rules” on page 91](#).

Web proxy and SteelHead SaaS

SteelHead SaaS doesn't use web proxy but instead uses the HTTP optimization within the SteelHead to access the external SaaS service infrastructure. Riverbed doesn't recommend that you configure SteelHead SaaS traffic to specifically use the web proxy feature; however, both the web proxy feature and the SteelHead SaaS can coexist on the same SteelHead. For more information, see [“Using in-path rules” on page 91](#).

Video caching

Some internet video services, specifically YouTube, can take advantage of the caching features of web proxy. Other cacheable video content that's static in nature (for example, video on-demand training and other nonstreaming video services) are also potentially able to be cached as they are usually presented to the web proxy as a cache-eligible file.

YouTube video typically relies on HTTPS. Make sure that you configure the global whitelist to include entries for both *.YouTube.com and *.googlevideo.com. Most of the actual content for YouTube is housed under *.googlevideo.com.

For more information about the whitelist, see [“Using the global whitelist” on page 87](#).

YouTube video content is automatically cached when using standard internet browsers, with the following exceptions:

- Clients using Firefox can't use the video caching features of web proxy. Firefox makes some header manipulations that the web proxy feature can't identify.
- Use of browsers in which users have selected to implement the use of the Google Quick UDP Internet Connections (QUIC) protocol multistreaming might have performance issues.
- Mobile browsers haven't been validated in this initial release of web proxy.

Configuring basic web proxy features (HTTP)

This section provides a simple overview about how to configure web proxy for HTTP traffic. For complete instructions, see the *SteelCentral Controller for SteelHead User's Guide*. Web proxy caching for HTTP is automatically enabled when you enable the feature.

To configure web proxy for HTTP

1. From the SCC Management Console, choose **Manage > Optimization: Web Proxy**.
2. Select **Enable Web Proxy**.
3. Click **Save**.

Figure 6-1. Web proxy - global configuration

The screenshot displays the 'Web Proxy' configuration page. At the top, there's a breadcrumb 'Manage > Web Proxy' with a help icon. The main section is titled 'Global Profile' and includes a descriptive text: 'The following Web Proxy configuration applies to all sites. More profiles can be added below to have distinct configuration at specific sites'. Below this, the 'Web Proxy Configuration' section has two checkboxes: 'Enable Web Proxy' (checked) and 'Enable HTTPS Optimization' (unchecked). There are 'Save' and 'Revert' buttons. The 'HTTPS Whitelist' section has an 'Add Domain' button and a table with one header 'Domain Name' and one row containing 'No Data'. At the bottom, there's a collapsed 'Parent Proxy Configuration' section. A note at the very bottom states: 'Note: For manual parent proxying of HTTPS connections, importing the Certificate Authority of the parent proxy into the client-side SteelHead is required if the parent proxy is intercepting and decrypting HTTPS connections.'

4. Select Site or Site Types.
5. Click **Push** to send the configuration to the SteelHead.

Figure 6-2. Web proxy push

Policy Push Control

This Web Proxy configuration will be pushed to SteelHeads running RiOS 9.1 or greater.

Exclude from Push

Push Status [More](#)

No push operation created within last 30 minutes.

Push to Appliances

Push Current Configuration to selected Sites or Site Types.

☒ Sites ☐ Site Types

branch1 ✕

1 appliances
1 uplinks
[Show Details](#)

Push **Clear**

Advanced configurations for web proxy

This section provides a detailed overview of the elements of web proxy and how to configure for HTTPS traffic optimization. It includes the following topics:

- [“Configuring web proxy for HTTPS” on page 84](#)
- [“SSL decryption and TCP proxy for HTTPS” on page 84](#)
- [“Using web proxy and certificate management” on page 84](#)
- [“Using the global whitelist” on page 87](#)
- [“Using in-path rules” on page 91](#)

Configuring web proxy for HTTPS

This section describes a simple overview about how to configure web proxy for HTTPS traffic. For complete instructions, see the *SteelCentral Controller for SteelHead User's Guide*.

To configure web proxy for HTTPS

1. From the SCC Management Console, choose Administration > Security: Certificate Authority.
Make sure that the certificate authority is enabled and that you've configured the associated key. For more information, see ["Using web proxy and certificate management" on page 84](#).
2. Choose Manage > Optimization: Web Proxy.
3. Select Enable Web Proxy and Enable HTTPS Optimization in the Global Configuration box.
4. Click **Save**.
5. Under Global HTTPS Whitelist, select Add Domain and populate the required domains.
For more information, see ["Using the global whitelist" on page 87](#).
6. Choose Manage > Optimization: Web Proxy.
7. Select the Site or Site Types.
8. Click **Push** to send the configuration to the SteelHeads.
9. From the SteelHead Management Console, choose Optimization > Network Services: In-Path Rules.
10. Configure the required in-path rules to support your implementation.
For more information, see ["Using in-path rules" on page 91](#).

SSL decryption and TCP proxy for HTTPS

Web proxy can decrypt HTTPS traffic for TLSv1.0, v1.1, and v1.2 and additionally supports TCP proxying for SSLv3 connections. For web proxy to leverage the HTTP SSL decryption feature, the SSL handshake must include the Server Name Indication (SNI). For more information about SNI, see the *SteelHead Deployment Guide - Protocols*.

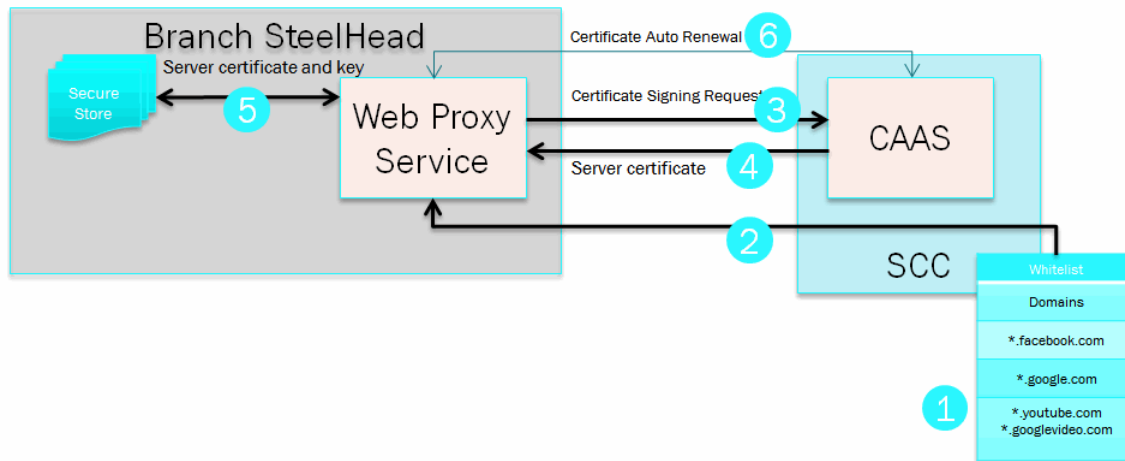
Note: To use the SSL caching features of web proxy for HTTPS traffic, you must make sure that the remote SteelHeads are running RiOS 9.1 or later with the SSL licensing installed.

Using web proxy and certificate management

This section describes some aspects that are important to consider regarding certificate management within your environment and its appropriate configuration for proper operation.

The web proxy HTTPS feature is critically dependent on the exchange of signed certificates between the SCC and the branch office SteelHead. **Figure 6-3** and the following steps show the certificate workflow of the web proxy feature.

Figure 6-3. Certificate workflow of the web proxy feature

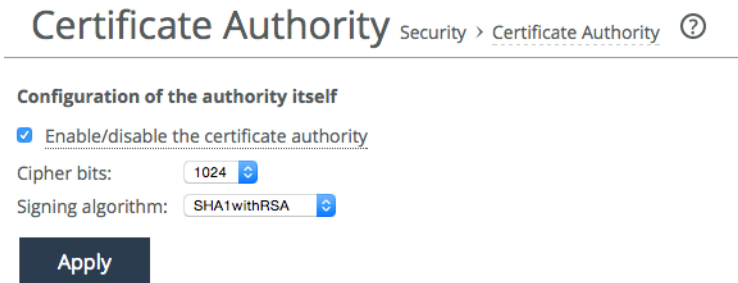


The following steps correlate to the numbers in **Figure 6-3**:

1. The whitelist is manually configured with the approved domain information.
2. The approved whitelist domains are pushed to the client-side SteelHead web proxy configuration.
3. Web proxy automatically sends a certificate signing request for the approved domain to the certificate authority service, which is configured on the SCC.
4. The SCC certificate authority (CA) responds with appropriately signed server certificate.
5. Web proxy stores the server certificate and the associated license key in the SteelHead secure vault for use when a client requests the approved domain.
6. Web proxy and the CA service automatically renew the server certificate as required.

You must enable the CA service feature of the SCC to generate server certificates and decrypt authorized content before optimizing HTTPS traffic with web proxy. The SCC CA service certificates must be trusted by clients using the service. [Figure 6-4](#) shows the CA configuration option on the SCC Management Console under Administration > Security: Certificate Authority.

Figure 6-4. Certificate authority page



The screenshot shows the 'Certificate Authority' configuration page. The breadcrumb navigation is 'Security > Certificate Authority' with a help icon. The section is titled 'Configuration of the authority itself'. There is a checked checkbox for 'Enable/disable the certificate authority'. Below this are two dropdown menus: 'Cipher bits' set to '1024' and 'Signing algorithm' set to 'SHA1withRSA'. An 'Apply' button is at the bottom.

Certificate Authority Security > Certificate Authority ⓘ

Configuration of the authority itself

☒ Enable/disable the certificate authority

Cipher bits: 1024

Signing algorithm: SHA1withRSA

Apply

If you already have an existing private key and CA-signed public certificate, you can import them (in PEM format only) by cutting and pasting the certificate into the SCC CA Service configuration page.

If you do not already own certificates and keys, you can generate a private key and self-signed certificate through the SCC CA Service. [Figure 6-5](#) shows an example on how to replace a self-signed trusted certificate under the CA.

Figure 6-5. PEM format

The authority's keys

Details
PEM
Replace

☐ Import Existing Private Key and CA-Signed Public Certificate (One File in PEM format)
☐ Import Existing Private Key and CA-Signed Public Certificate (Two Files in PEM format)
☒ Generate New Private Key and Self-Signed Public Certificate

Private Key

Cipher: RSA Cipher Bits: 1024

Self-Signed Certificate

Common Name: Example Self Signed CA

Organization Name: Example Organization

Organization Unit Name: Example Team Name

Locality: San Francisco

State: CA

Country: US (two-letter code)

Email Address: admin@exampleemail.com

Validity Period: 730 Days (60 to 3650 days)

Generate Key And Certificate

Select the PEM tab to view the certificate.

After you've configured the SCC CA service and have an SCC CA certificate created, we recommend that you follow your internal procedures to install the SCC CA certificate on your web client configurations as a trusted root certificate.

After you've configured the client-side SteelHead to support HTTPS web proxy, it automatically generates and renews the server certificates that the domain whitelist has allowed. Each client-side SteelHead contains its own secure vault and locally stores the generated keys and certificates within.

For more information, see [“Configuring and Using the SCC as a Certificate Authority Service” on page 45](#) and the *SteelCentral Controller for SteelHead User's Guide*.

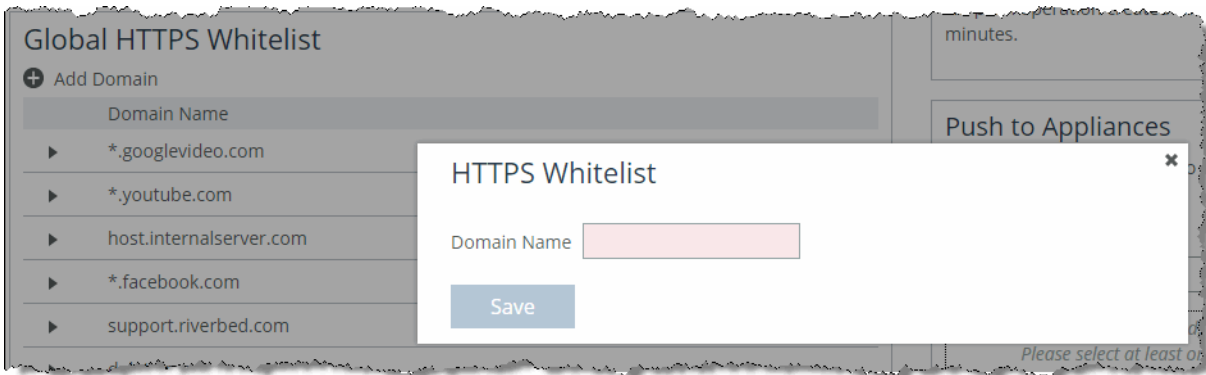
Using the global whitelist

You must configure the global HTTPS whitelist to contain the top-level and subdomain names for which the SCC permits the branch offices to proxy HTTPS. Choose Manage > Optimization: Web Proxy.

Be as specific as possible when you enter the whitelist domains; use the fully qualified domain name (FQDN) for each unique site requesting proxy service. In addition to using a specific FQDN, the whitelist accepts:

- wildcard domains (for example, *.facebook.com, *.YouTube.com, *.Riverbed.com).
- hostnames (for example, webserver.myinternaldomain.com).

Figure 6-6. Entering a name into the whitelist



Using parent proxy (proxy chaining) configurations

RiOS 9.5 introduces support for web proxy configurations requiring the integration into environments where additional proxy services reside upstream from the SteelHead—we refer to these upstream proxies as *parent proxies*. The web proxy service is able to now operate within a hierarchical chain of proxy servers that can pull content from the localized cache of each proxy up the chain—commonly referred to as *proxy chaining*.

There are two available configuration methods for using the parent proxy service within web proxy:

- **Manual mode** - utilized when clients need to access content transparently with no knowledge of a proxy servers' existence.
- **Automatic mode** - utilized when clients are required to explicitly access a specific proxy server as configured in the end-user browser or via a client PAC file locally.

Note: Parent proxy can only be deployed as a manual or automatic mode configuration but not as both configurations simultaneously.

Configuring manual mode parent proxy

To enable the manual parent proxy configuration options for HTTP and HTTPS first select Parent Proxy Configuration, then select the Manual radio button. Enter the upstream parent proxy server hostname, FQDN or IP address along with the specific server port in the following format:

```
<Parent Server>:<Service Port>
```


Figure 6-7 illustrates how to configure manual mode for a parent proxy service utilizing both the HTTP and HTTPS schemes.

Figure 6-7. Manual mode

▼ Parent Proxy Configuration ⓘ

☐ Off
☐ Automatic
☒ Manual

HTTP Servers Ordered List, Comma Separated
 HTTPS Servers Ordered List, Comma Separated

The ability to exclude specific domains from a manual mode parent proxy configuration is also available. This parent proxy exception option is only applicable to the manual mode server and not configurable for automatic mode. Figure 6-8 illustrates how to enter domain name exclusions that should bypass the parent proxy configured.

Figure 6-8. Entering the domain name exclusions

Parent Proxy Exceptions ⓘ

+ Add Domain

Domain
▶ google.com
▶ riverbed.com
▶ att.com
▶ riverbed.cc

Parent Proxy Exceptions ⓘ

Domain Name

Note: For manual parent proxying of HTTPS connections, importing the Certificate Authority of the parent proxy into the client-side SteelHead is required if the parent proxy is intercepting and decrypting HTTPS connections.

The parent proxy used (when multiple are configured) is selected based upon a combination of the traffic scheme, which is limited to five parents per scheme (HTTP as opposed to HTTPS) and the operational mode selected. Failover mode is the configured default and selects the configured parent proxy in order of entry. Load-balanced mode enables parent proxies to be selected round-robin based on client IP hash. For either resiliency mode, if no configured parent is available in the requested scheme then the parent will be marked as down for a five minute interval and traffic for that scheme will be blackholed, that is dropped.

Configuring automatic mode parent proxy

To enable the automatic parent proxy option for HTTP and HTTPS first select Parent Proxy Configuration, then select the Automatic radio button. No additional configuration is required on the web proxy for default operation. Clients need to be configured with a PAC file or explicit browser configurations prior to enabling for correct operation. **Figure 6-9** illustrates this selection in the SCC Management Console. By default all HTTPS-cache eligible content will be proxied and cached and HTTP traffic will be proxied. In order to cache HTTP content under automatic mode you need to manually add the parent proxy IP addresses via the CLI on the SteelHead using the following command:

```
[no] web-proxy parent automatic whitelist ip <IPv4 address of each parent proxy>
```

Figure 6-9. Automatic mode

▼ Parent Proxy Configuration ⓘ

☐ Off

☒ Automatic

☐ Manual

HTTP Servers Ordered List, Comma Separated

HTTPS Servers Ordered List, Comma Separated

Note: When configuring automatic or manual parent proxy modes, the SteelHead must trust the certificates issued by the parent proxy server or provider in order to properly proxy and cache HTTPS traffic when using the parent proxy configurations. For more information about adding Certificate Authorities for a proxy service on SteelHead, see the *SteelHead Deployment Guide - Protocols*.

Using in-path rules

Relative to web proxy, you configure in-path rules locally on the SteelHead, or you can alternatively configure them on the SCC and push them to the SteelHeads. These rules are used on the SteelHead to determine whether traffic has optimization applied or is passed through when the SteelHead detects a connection initiated by a client. The very basic implementation of web proxy (only enabled for public IP addresses using HTTP on TCP port 80) uses the default in-path rule. **Figure 6-10** shows settings for the default in-path rule (bottom of **Figure 6-10**) on the SteelHead. web proxy is set to Auto by default.

Figure 6-10. Default in-path rule on the SteelHead

Rule	Type	Source	Destination	VLAN	Protocol	Preoptimization Policy	Latency Policy	Data Reduction Policy	Cloud Acceleration	Web Proxy
1	Auto Discover	All-IPv4:*	74.0.0.0/8:443	All	--	None	Normal	Normal	Auto	Auto
Description: Manual Rule to allow HTTPS connections to utilize Web Proxy and not pass-through										
2	Pass Through	All-IP:*	All-IP:Secure	All	TCP	--	--	--	Auto	None
3	Pass Through	All-IP:*	All-IP:Interactive	All	TCP	--	--	--	Auto	None
4	Pass Through	All-IP:*	All-IP:RBT-Proto	All	TCP	--	--	--	Auto	None
5	Pass Through	10.5.0.0/16:*	10.12.0.0/16:80	All	TCP	--	--	--	Pass	Force
Description: In-Path rule configuration for Private or Intranet Addresses										
default	Auto Discover	All-IP:*	All-IP:*	All	--	None	Normal	Normal	Auto	Auto
Description: Default In-Path Rule										

For more information on in-path rules, see the *SteelHead Management Console User's Guide* and the *SteelHead Deployment Guide*.

Consider the following selections when you configure the web proxy options in the in-path rule table:

- **Auto** - All non-RFC 1918 IPv4 addresses on ports 80 and 443 matching this rule are forwarded to the web proxy service.
- **Force** - Any IP address and port matching this rule (even those in RFC 1918) are forwarded to the web proxy service. The Force option is a pass-through rule.
- **None** - Traffic matching this rule is not forwarded to the web proxy service.

Web proxy doesn't leverage SteelHead transparency settings, although you can select them under Auto Discover and Pass Through In-Path rule configurations. If you select these options on the in-path rule, either the Auto or Force Web Proxy rule option, they are ignored. When web proxy is set to the automatic setting, eligible traffic is optimized using the web proxy feature; however, if the traffic can't be optimized by the web proxy feature, then autodiscovery occurs and full transparency/port transparency options are preserved.

RFC 1918 (private IP address range) traffic bypasses the web proxy and uses the SteelHead HTTP optimization unless you add an in-path rule that specifies the RFC 1918 address you want to the web proxy to service. Configure this rule Type as Pass Through and select the Force option as shown in [Figure 6-11](#).

Figure 6-11. Example configuration of an in-path rule

Add a New In-Path Rule ✕ Remove Selected Rules ⇅ Move Selected Rules...

Type: Pass Through
 Web Proxy: Force
 Source: { Subnet: IPv4 10.5.0.0/16 (X.X.X.X/X)
 Port: All Ports
 Destination: { Subnet: IPv4 10.12.0.0/16 (X.X.X.X/X)
 Port: Number 80 (1 - 65535)
 Domain Label: n/a
 VLAN Tag ID: all
 Protocol: TCP
 Cloud Acceleration: Pass Through *Must be set to "Pass Through" if a Domain Label (see above) is selected*
 Position: End
 Description: In-Path rule configuration for Private or Intranet Addresses
 Enable Rule: ☒
 Add

To successfully optimize HTTPS connections, you must configure a new in-path rule for destination TCP port 443. You must configure this rule to negate the preconfigured secure content pass-through rule (secure PT rule) that ships with all SteelHeads. The secure PT rule takes precedence even with web proxy HTTPS optimization enabled. [Figure 6-12](#) shows a new rule added above the existing secure PT rule. Be sure to configure the rule with the Auto or Force Web proxy rule option while entering the destination port of 443.

Figure 6-12. Adding a new rule

Rule	Type	Source	Destination	VLAN	Protocol	Preoptimization Policy	Latency Policy	Data Reduction Policy	Cloud Acceleration	Web Proxy	Kickoff	Status
1	Auto Discover	All-IPv4*	74.0.0.0/8:443	All	--	None	Normal	Normal	Auto	Auto	No	Enabled
Description: Manual Rule to allow HTTPS connections to utilize Web Proxy and not pass-through												
2	Pass Through	All-IP*	All-IP:Secure	All	TCP	--	--	--	Auto	None	--	Enabled

SteelHead SaaS traffic to services such as Office365 and Salesforce.com bypass the web proxy feature and use the SteelHead SaaS HTTP optimization. Riverbed doesn't recommend that you use an in-path rule with the Force option to proxy specific SteelHead SaaS-associated address.

In SCC 9.2 and later, you can choose to use the host and domain labels within the in-path rule configuration. These labels allow for increased flexibility when there's a need to group IP addresses, IP ranges, FQDN, or wildcard domain names (for example, *.riverbed.com).

Use these guidelines when using host or domain labels with web proxy:

- Implement domain and host labels to allow or deny web proxy serviceable HTTP/HTTPS traffic, based on your unique traffic needs. By default all TCP port 80 traffic is serviced by the web proxy through the default rule, and TCP port 443 always requires an inclusive in-path rule for desired domains.
- The SSL whitelist is functioning as a CA encryption/decryption validator, not a HTTPS traffic policer. You need to create more than only a host and domain label or specific HTTPS domains or you might not receive the expected cached content.
- We recommend that you order a domain label in-path rule low on the rule list, and configure that rule with a destination host label or focused destination IP range. This ordering ensures that the rule engine is not inappropriately conflicting with other rules, such as fixed-target rules or SaaS pass-through rules.

Fixed-target and SaaS pass-through rules can also leverage domain labels using the same matching domains. These fixed-target and SaaS pass-through rules might never be used because these rules occur below a higher-ordered domain label rule.

Troubleshooting web proxy

This section describes some of the common issues that may arise within the deployment and use of web proxy and outlines some procedures to validate correct configurations. It includes the following topics:

- [“SCC to SteelHead communications” on page 93](#)
- [“HTTP caching” on page 93](#)
- [“HTTPS decryption” on page 94](#)
- [“YouTube video caching” on page 94](#)

SCC to SteelHead communications

Because of the interdependency of both the SCC and the branch office SteelHeads, when you use the web proxy feature, you must validate that the two appliances can communicate with each other through their HTTPS channel. You can validate this command using the **show scc** command. This command shows the current state of both HTTPS connections to and from the respective SteelHead and the SCC that manages it.

HTTP caching

If the browser client is reporting Service Unavailable or if some traffic is not showing as a web proxy connection in the Current Connections report on the SteelHead, this report might be indicating that the web server is not accessible from the in-path interface of the client-side SteelHead. In either case, make sure that the in-path rules for those web services are correctly configured for each SteelHead requiring that access.

Additionally, make sure that web proxy service is enabled on the branch office SteelHead with the **show web-proxy status** command.

If you're not seeing cache-hits register or increment, consider the following:

- When looking at cache utilization on the SCC, understand that the SCC polling on cache utilization occurs in 60-minute intervals. Make sure that the caching is enabled with the **show web-proxy cache status** command.
- Make sure that the cache-hit counters are incrementing on the SteelHead with the **show web-proxy stats cache** command.
- Make sure that the proxy content you're looking to validate as being cached is actually cacheable content as outlined in the RFC 2616 standard.

HTTPS decryption

If you observe that there's no HTTPS content serviced by web proxy being optimized, immediately validate the following elements:

- Check that HTTPS traffic has an in-path rule added on the SteelHead configured with web proxy for TCP port 443.
- Verify that SSL has been properly enabled on the SteelHead with the following commands:
 - **show web-proxy ssl**
 - **show web-proxy cache ssl**
- Make sure that the domain being proxied is in the HTTPS global whitelist configuration and that the configuration has been pushed to the SteelHead in question.
- Make sure that the trusted CA certificate is the one that's actually being presented to the client browser. You can check this certificate status by clicking on the lock icon within the URL field of most browsers.
- If the SCC trusted certificate is not being seen on the client machines, you need to make sure that the SCC issued the certificate on the Certificate Authority page on the SCC Management Console.

YouTube video caching

If you're not seeing any YouTube traffic being optimized in the Current Connections report on the SCC or the video playback quality is abnormally impaired, you can investigate the following reasons:

- If you do not see the CA certificate being used for any other proxied websites, there might be an issue with traffic not being serviced by the web proxy (for example, the port 443 in-path rule is not configured for web proxy or SSL service is not running).
- The correct domains might not be added to the global whitelist configuration on the SCC.
- Because YouTube content can't present certificate warnings, check other sites listed in the global white list for certificate warnings to make sure that there's not a certificate issue.
- You can try manually adding the CA certificate from the SCC directly to the client browser for test purposes.

If you believe that videos on YouTube aren't being serviced from the cache, you can investigate the following reasons:

- Validate that the global whitelist configuration on the SCC includes both *.youtube.com and *.googlevideo.com. The majority of YouTube content is served from the latter.
- Follow the troubleshooting suggestions for both HTTP and HTTPS content caching to validate.

Secure Transport

This chapter describes secure transport, which enables you to integrate and configure encryption services with path selection. Secure transport enables simple, secure, and manageable group encryption for inter-SteelHead communication for path selection deployments. It includes the following sections:

- [“Overview of secure transport” on page 98](#)
- [“Secure transport sizing” on page 102](#)
- [“Functional operations overview” on page 103](#)
- [“Secure transport configuration workflow” on page 109](#)
- [“Deployment example of a hybrid network backhaul scenario” on page 111](#)
- [“Deployment example of a dual in-path interface with split tunnel” on page 119](#)
- [“Reporting” on page 127](#)

This chapter requires you be familiar with path selection, topology, sites, and networks. For more information, see [“Sites and Networks, Path Selection, and QoS” on page 61](#), the *SteelHead Deployment Guide*, and the *SteelCentral Controller for SteelHead User’s Guide*.

Secure transport requires:

- SCC 9.0 or later
- RiOS 9.0 or later
- SteelHeads deployed physically in-path and in the natural flow of traffic
- SteelHeads at each end of the path to be secured
- SSL license
- Configured path selection policies
- Secured paths

Overview of secure transport

This section provides an overview of secure transport and includes the following topics:

- [“Management plane” on page 100](#)
- [“Control plane” on page 101](#)
- [“Data plane” on page 101](#)

A key feature for the hybrid enterprise is to leverage multiple paths between sites. This feature enables diversity across a variety of networks for resiliency. More importantly this feature enables you to more efficiently select the network delivery mechanism for an application.

However, not all network paths are equal in terms of exposure to security risks. A common technique to secure traffic is to encrypt traffic at the IP layer with IPSec, which prevents disclosure if traffic is intercepted in transit. This type of encryption also prevents replay attacks from an attacker in the middle of the chain.

Starting in RiOS 9.0 and SCC 9.0 Riverbed introduces the secure transport feature. Secure transport is integrated with path selection and provides a method to configure and enable encryption services for traffic over a path. Secure transport provides security services at the IP layer for path selection—you can add security for an application, application groups, or custom applications based on the path the traffic travels. By integrating encryption services into the SteelHead, the secure transport feature delivers a secure path between peers over a private WAN, public internet, or a hybrid network.


One key capability of secure transport with path selection managed by the SCC is that each peer automatically joins a security group with other SteelHeads whenever you mark a network as *securable* and assign the network to the site the peers belong to. This capability has benefits over other forms of encryption services that require configuration at both ends and policies to select the traffic that’s encrypted.

One benefit is that automatically joining a security group requires minimal configuration because it is easier to implement the network design changes when you want to include new sites or new networks that require encryption. Additionally, configuration is done centrally on the SCC, which minimizes the potential for configuration errors as opposed to configuring encryption separately on each device. Also, other virtual private network (VPN) technologies require traffic to match a policy, which then triggers negotiation between devices performing encryption before the encryption is actually performed. With secure transport, encryption is applied packet-by-packet using IPSec tunnel mode.

Secure transport uses a group model for encryption services. In a group model each appliance registers with a controller to learn network-reachability information, policies, and important material to use for encrypting traffic. Each appliance also performs reachability tests to ensure the path is available.

Secure transport uses standards-based IPSec using AES-256 and SHA-2 to secure traffic over a path. Secure transport uses IPSec tunnel mode of which the original IP addresses are part of the encrypted payload (Figure 7-1).

Figure 7-1. Tunnel mode

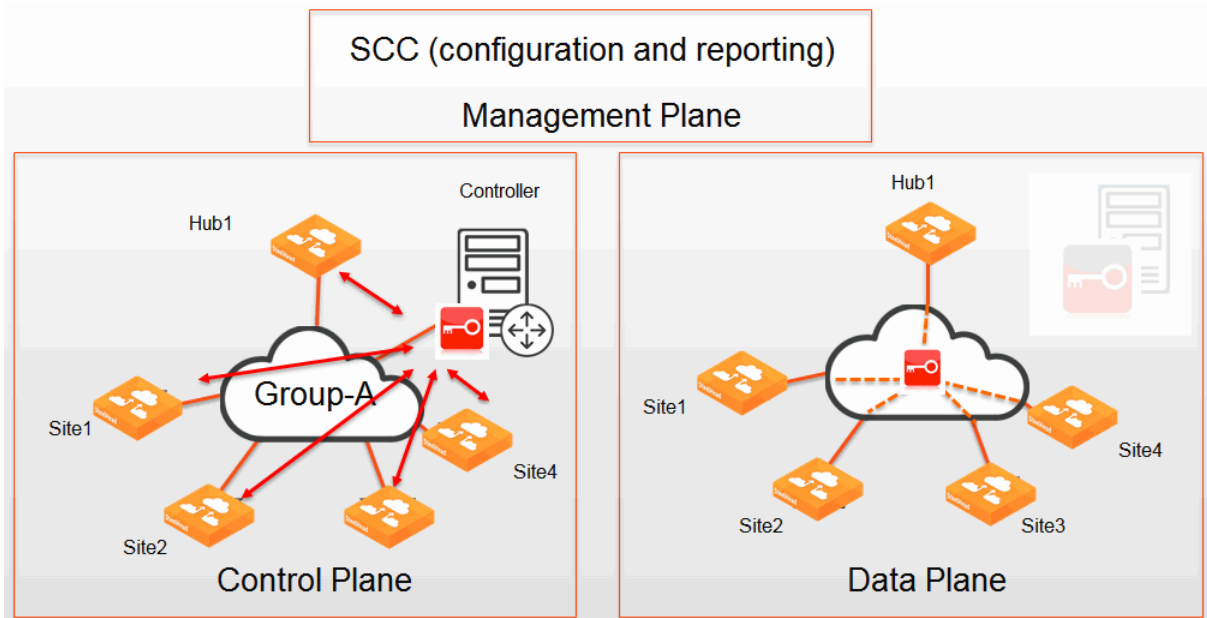


You manage the encryption services centrally through the SCC. Functionality of secure transport is separated into the:

- management plane (page 100).
- control plane (page 101).
- data plane (page 101).

Figure 7-2 shows how the functionality of secure transport performs at a high level. The SCC works with the management plane for configuration and reporting, a SteelHead in the group acting as the controller manages the control plane actions (such as distributing network-reachability information and encryption key information), and the SteelHeads in the group encrypt traffic for the data plane.

Figure 7-2. Secure transport functionality



Management plane

The management plane is configured on the SCC using SSL and SSH secured communications, in which you create and distribute the path selection policy. When a network is marked as securable, it indicates that the SteelHead will join a group of other SteelHeads that can encrypt traffic. The SCC pushes the path selection policy to all the SteelHeads.

The management plane also provides reports on the status and metrics associated with the secure transport feature.

The SCC is not intended to be an internet-facing device. We recommend that you connect over a network path, controlled by your organization: for example, a private WAN such as MPLS or a third-party VPN connection. After the SteelHead communicates with the SCC, it receives information on how to reach the controller, registers with the controller, and then can perform encryption services for path selection.

Currently, performing encryption services when a site doesn't have a network path, controlled by your organization, is not a current use case. However, after the SteelHead has registered with the SCC and subsequently with the secure transport controller (controller), it can provide encryption services during periods of time in which connectivity with the SCC is not available through a network path controlled by your organization. This behavior occurs because the SteelHead can continue operating with the path selection policy received from its last contact with the SCC and the key information with the SteelHead acting as the controller.

Control plane

The control plane is a secure communications channel between each SteelHead peer (group member) and a controller. The controller can run on a SteelHead that has been configured as a controller and activated using SCC. When a SteelHead receives a path selection configuration that has a securable path, the SteelHead completes a process of registering with the controller. The controller communicates with each SteelHead in the group and distributes network-reachability information and generates the encryption keys to use over the data plane.

For more information about encryption keys, see [“Encryption key operations” on page 107](#).

Remember that the control plane is providing network reachability information related to the site, such as the public IP address for a secured path.

For more information about sites, see [“Configuring sites and networks” on page 62](#).

For more information about the controller, see [“Controller sizing” on page 103](#).

Data plane

Each SteelHead in the group takes part in the data plane and secures traffic over any path marked as securable. The SteelHeads use the encryption keys received from the controller to perform encryption and use the path selection services policy received from the SCC to determine which traffic is encrypted.

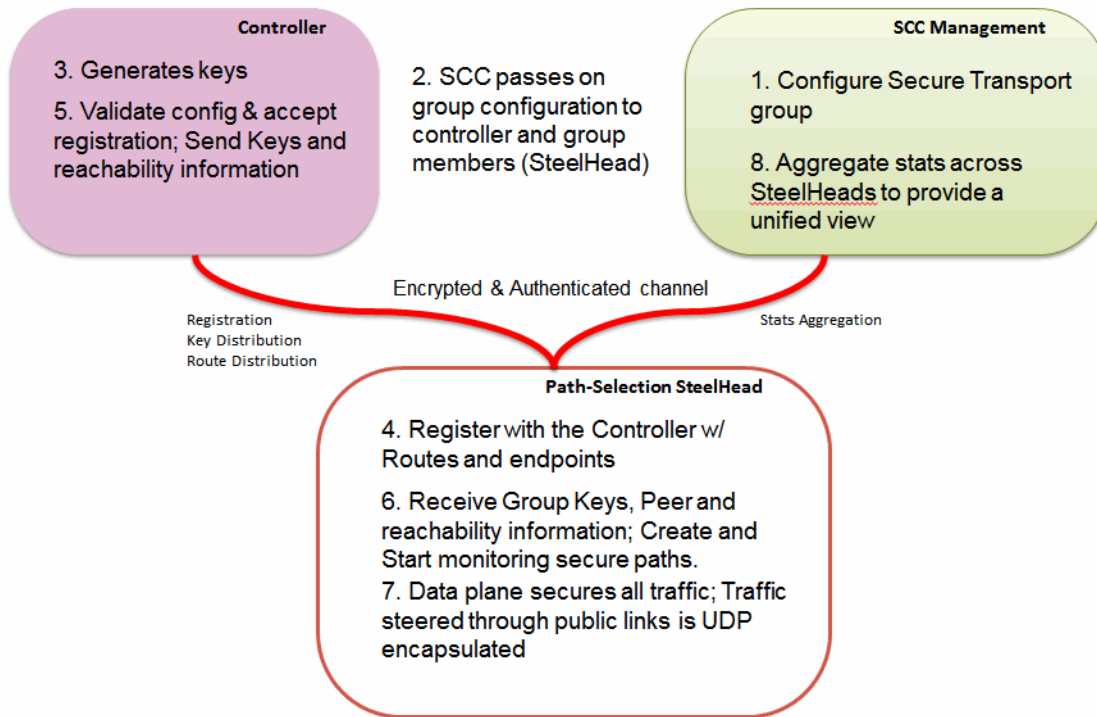
Remember that the control plane is providing network reachability information related to the site such as the public IP address for a secured path.

For more information about sites, see [“Configuring sites and networks” on page 62](#).

Note: In RiOS 9.0 and later, secure transport provides encryption services for only IPv4. IPv6 is not encrypted.

Figure 7-3 shows another view of the functionality of the management, control, and data plane for secure transport. The SCC performing the management plane is the central point for configuration of the secure transport group. The SCC passes the group configuration information to the controller and group members. The controller generates encryption key information and network-reachability information, and it also performs registration and validation of group members in the control plane.

Figure 7-3. Secure transport functionality



Secure transport sizing

This section has information about sizing. It includes the following topics:

- [“Group sizing” on page 102](#)
- [“Group member sizing” on page 102](#)
- [“Controller sizing” on page 103](#)

Group sizing

In RiOS 9.2 and later, the maximum number of supported secure transport sites is limited to 250. If you’re using secure transport in a RiOS version previous than 9.2, see that version of the *SteelCentral Controller for SteelHead Deployment Guide*.

Group member sizing

Secure transport is a network service. For more information about the throughput and connections, see the *SteelHead Product Family Specification Sheet*.

Consider the concentration point in most VPN networks (typically the data center) has the most throughput, and as a result requires higher capacity devices. Likewise with WAN optimization, the SteelHeads in the central part of the network are larger model appliances with more CPU, RAM, and data store capacity. Furthermore, you can configure a dedicated SteelHead as a secure transport concentrator to perform only encryption services. For more detailed questions, contact your Riverbed account team.

For more information on a secure transport concentrator, see [“Secure transport concentrator” on page 108](#).

Controller sizing

Sizing of a controller follows similar logic as discussed in [“Group member sizing” on page 102](#). We recommend that you consider the following factors when sizing and selecting the controller:

- The SteelHead you choose as the active controller must be available on all securable paths because this appliance is the control plane for any remote site SteelHead performing encryption services for a path. The active controller is responsible for rekey operations and must securely contact each SteelHead in the group in order to distribute encryption key information to use on the data plane.
- The active controller has ample resources such as bandwidth usage. The CPU is not completely consumed by optimization workloads. The operation is lightweight in terms of bandwidth and CPU usage, but selecting an active controller at a site that has to compete for resources, such as bandwidth, can impact the control plane.

For more information about rekeying, see [“Encryption key operations” on page 107](#).

- The active controller is ideally situated in a central part of the network to timely complete any control plane actions. Thus, the round-trip time from the active controller to each SteelHead in the group is a consideration for selecting any appliance.

Functional operations overview

This section describes functional operations of the secure transport feature. It contains the following topics:

- [“Firewall considerations” on page 104](#)
- [“Network interface used for SteelHead-to-controller communication” on page 104](#)
- [“Maximum segment size” on page 104](#)
- [“Network address translation” on page 105](#)
- [“Encryption key operations” on page 107](#)
- [“Disconnected mode” on page 108](#)
- [“Fail closed” on page 108](#)
- [“Secure transport concentrator” on page 108](#)

Firewall considerations

Secure transport requires communication on the management plane, control plane, and data plane. All communication is initiated by the SteelHead peer (group member), except for the management plane. Consider the following port usage:

- The management plane requires communication between the SteelHead and SCC on TCP port 9443 and TCP port 22.
- The control plane between the SteelHead acting as the controller and SteelHeads acting as group members is over TCP port 9443.
- Encryption services flows over ESP (IP protocol 50) or if the network is public UDP port 4500.

Network interface used for SteelHead-to-controller communication

The controller can communicate using any available SteelHead interface in which the controller is running on. By default, the SteelHead peer (group member) uses only the primary interface when communicating to the controller. This limited interface use presents a challenge if the primary interface can't communicate over the data plane.

To communicate from a separate interface, the SteelHead must run RiOS 9.1 or later, and you must enter the **stp-client controller in-path enable** command.

After you use this command to enable another interface, the SteelHead attempts to communicate first from the in-path interfaces (starting with the lowest number, such as 0/0) and then moving onto the primary and auxiliary interface, until communication is established. Remember, because the request travels to both the private and then public address of the controller, it can take up to 40 seconds before the next interface is attempted.

Consider the choice of the controller IP address based on the reliability of the networks the in-path interfaces are connected to. For example, you can configure only a private IP address for the controller. This configuration allows all the SteelHeads participating in the secure transport group to contact the controller only to that single private IP address.

Because secure transport was built taking hybrid networking into account, you can configure a private and public IP address. The public IP address relies on reachability over a network marked Public. This configurations allows the controller to first be reached over the private network and then the public network. The private and public IP addresses do not need to come from the same in-path interface. With the **stp-client controller in-path enable** command, the data center SteelHead (DC-SH) uses the private and public IP addresses from the same in-path interface, while the branch office SteelHead (BR-SH) uses the private and public IP addresses from different in-path interfaces.

Maximum segment size

Whenever you use encryption services or any encapsulation technique, packet size and fragmentation become an issue. The SteelHead is aware of which paths are using an encapsulation technique, such as firewall traversal or secure transport, and lowers the TCP maximum segment size (MSS) as part of the initial setup of the connection. This behavior occurs for pass-through and optimized traffic and minimizes the possibility for fragmentation.

The secure transport feature uses more overhead than the 24 bytes for the generic routing encapsulation (GRE) header used in the firewall traversal feature but operates in the same manner. For information about maximum transmission unit (MTU) considerations for the firewall traversal feature, see the *SteelHead Deployment Guide*.

Network address translation

The encryption techniques used for secure transport are standard-based IPsec that includes encapsulating security payload (ESP). As part of the secure transport feature, the SteelHead can further encapsulate ESP packets in UDP. This encapsulation allows the SteelHead to leverage the private to public address translation, commonly referred to as NAT, which occurs at the boundary device between the private LAN and public WAN.

By default, RiOS uses UDP port 4500. On the SCC, NAT traversal is employed when you mark that a network is public and securable (Figure 7-4).

Figure 7-4. Public and securable network

Networks

Represents the WAN clouds that sites use to communicate to each other using Primary MPLS, VSAT or Internet.

+ Add a Network

Name	Public	Securable
▶ MPLS	No	Yes
▼ Internet	Yes	Yes

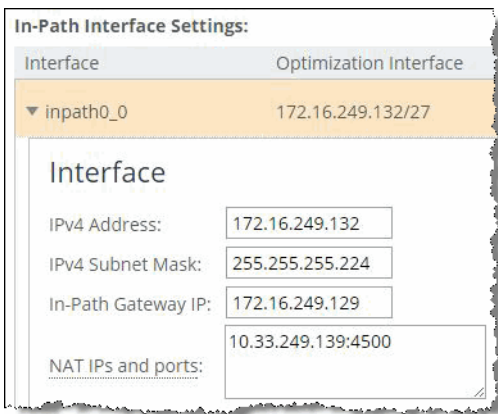
Network Name:

☒ Securable using Secure Transport

☒ Public Network

For each SteelHead to use a public and securable path, the SteelHead must register its public IP address to the controller. On the SteelHead, you configure the public IP address and port number as part of the in-path interface configuration (Figure 7-5).

Figure 7-5. Public IP address



Interface	Optimization Interface
▼ inpath0_0	172.16.249.132/27

Interface

IPv4 Address: 172.16.249.132

IPv4 Subnet Mask: 255.255.255.224

In-Path Gateway IP: 172.16.249.129

NAT IPs and ports: 10.33.249.139:4500

You must manually enter the public IP address because it is not automatically discovered. If the public IP address changes, you must change it on the SteelHead. We recommend that you use a static IP address from your service provider for a network configured as public.

If providing secure transport services over a public network in which NAT is used, the SteelHead acting as the controller must have its public IP address also assigned. This configuration is performed at the CLI for any SteelHead that's an available controller.

```
DC-SH (config) # stp-controller address private-ip 172.16.250.132 public-ip 10.33.249.140 port 4500
```

Figure 7-6 shows packet details from the Wireshark display for an encapsulated packet using NAT traversal.

Figure 7-6. Details of NAT traversal

No.	Time	Delta Time	Conv.	VLAN	Source	Src	Destination	Dest	S Port	D Port	Proto	Len
313	26.320170				00:50:56:8a:29:cd	10.33.249.140	00:50:56:8a:db:5d	172.16.249.132	4500	4500	ESP	142
<div> <div>Frame 313: 142 bytes on wire (1136 bits), 142 bytes captured (1136 bits)</div> <div>Ethernet II, Src: Vmware_8a:29:cd (00:50:56:8a:29:cd), Dst: Vmware_8a:db:5d (00:50:56:8a:db:5d)</div> <div>Internet Protocol Version 4, Src: 10.33.249.140 (10.33.249.140), Dst: 172.16.249.132 (172.16.249.132)</div> <div>User Datagram Protocol, Src Port: ipsec-nat-t (4500), Dst Port: ipsec-nat-t (4500)</div> <div>Source port: ipsec-nat-t (4500)</div> <div>Destination port: ipsec-nat-t (4500)</div> <div>Length: 108</div> <div>Checksum: 0x0000 (none)</div> <div>UDP Encapsulation of IPsec Packets</div> <div>Encapsulating Security Payload</div> <div>ESP SPI: 0x80000009 (2147483657)</div> <div>ESP Sequence: 3349</div> </div>												

Encryption key operations

In the world of data encryption, no key is considered secure for eternity. Given a period of time or given enough data, an attacker might be able to compromise the encrypted key. Secure transport supports the configuration of a rekey interval (time based) or rekey data size (volume based). By default the SCC has a rekey interval of 3600 seconds or 1 hour (**Figure 7-7**). However, for higher speed links a rekey can be triggered by the total amount of data transacted by the group, which is by default 4 TB. You can adjust these values by choosing Manage > Services: Secure Transport.

Figure 7-7. Rekeying settings

Group Settings

Rekey Interval

3600

seconds

Rekey Data Size

4194304

MBs

Disconnected Mode Timeout

300

seconds

Authentication Algorithm

SHA-256 (HMAC)

Encryption Algorithm

AES-256 (CBC)

Save and Apply

Revert

The rekey operation is performed without incurring a period of packet loss. Each SteelHead in the group is given a new encryption key by the controller. As each appliance in the group has the new encryption key, it uses the new key to perform path monitoring. After each group member on a path is using the new key for path monitoring, the SteelHeads switch to the new key.

Disconnected mode

The SteelHeads in the group contact the controller every 15 seconds as a means to verify if the controller is reachable. When a SteelHead in the group is unable to reach the controller after three attempts, that SteelHead enters into *disconnected mode*. Disconnected mode lasts for 300 seconds by default, after which time the SteelHead ceases performing encryption services. You can adjust the disconnected mode timeout to allow for more or less time. You can add more time to account for the manual process of selecting a new controller in the SCC. Less time can potentially run the risk of causing encryption services to stop during a momentary outage.

Fail closed

You can decide that a SteelHead can stop all traffic if encryption services for a path can't be performed and no alternative secure paths are available. This operation can be done in the path selection rules and applied to individual applications, application groups, custom applications, or for all traffic. However, this setting doesn't take effect if the SteelHead fails and goes into the default fail-to-wire mode. In fail-to-wire mode, all traffic is relayed through the SteelHead without encryption. As a result, configure both the in-path interface for fail-to-block and the path selection rules to drop traffic if the desired secure path is not available.

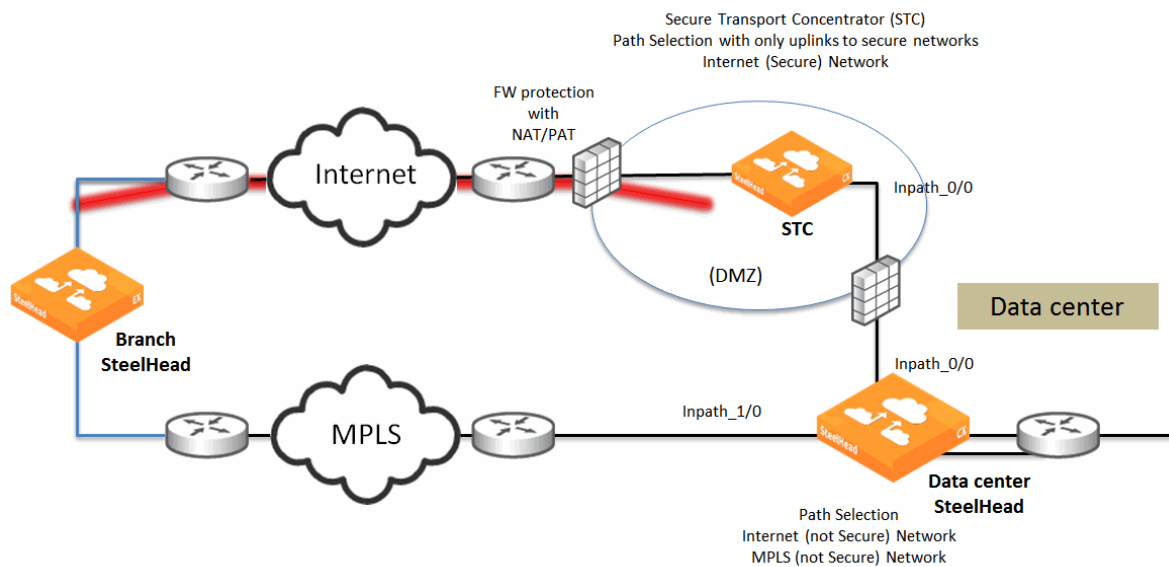
Secure transport concentrator

You can dedicate a SteelHead for secure transport functionality by configuring the SteelHead in the SCC as a secure transport concentrator. When configured as a secure transport concentrator, the SteelHead only performs encryption services. This configuration can be an ideal option for:

- SteelHeads positioned in the Demilitarized Zone (DMZ) in which additional security filtering is performed on the LAN side of the concentrator.
- when traffic might not flow through a SteelHead as expected when the SteelHead is positioned closer to a different WAN egress point at a site.
- increasing encryption bandwidth for higher throughput deployments.
- deployments with SteelHead clusters using the Interceptor.

Figure 7-8 shows a typical topology that's applicable to the secure transport concentrator capability.

Figure 7-8. Topology for secure transport concentrator



Use the SCC in order to configure a SteelHead as a secure transport concentrator on the sites page by adding the SteelHead and one or more secure uplinks.

For more details about the secure transport concentrator, see the *SteelCentral Controller for SteelHead User's Guide*.

Secure transport configuration workflow

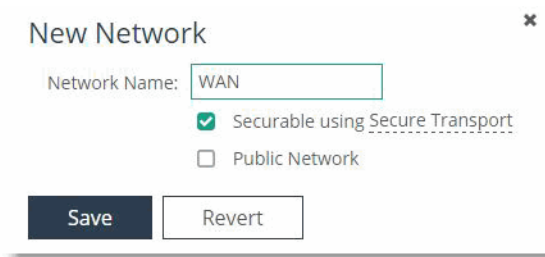
This basic secure transport workflow provides encryption services over a private WAN.

To configure secure transport

1. Configure sites and networks with the network subnets identified.
For details, see [“Sites and Networks, Path Selection, and QoS” on page 61](#).
2. Configure a site with a network that's marked as securable (Figure 7-9).

After making the network securable, the SteelHead registers with the controller to complete the actions over the control plane, such as identifying routes and key information.

Figure 7-9. Securable network



New Network

Network Name: WAN

☒ Securable using Secure Transport

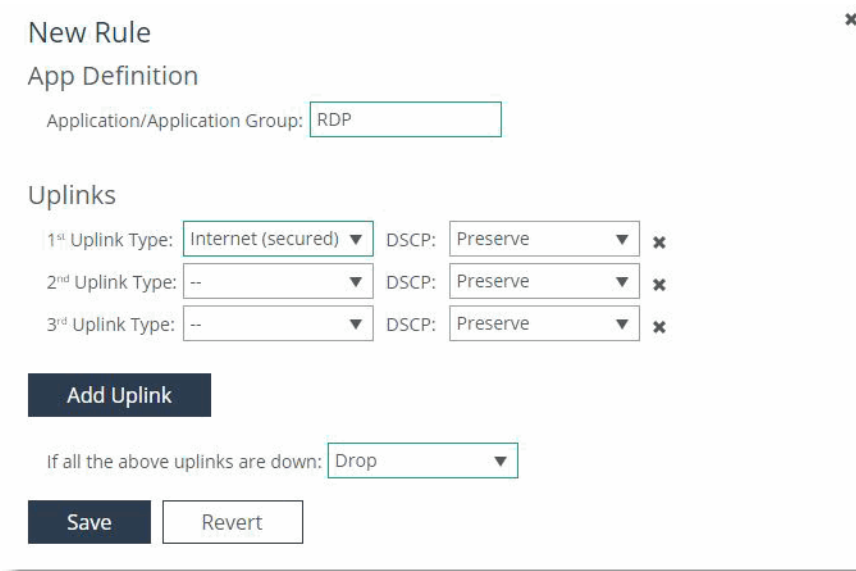
☐ Public Network

Save Revert

3. Create a path selection rule to mark traffic over a path you want to secure.

Figure 7-10 shows the RDP application as using the internet (secured) path. If the specified secured path is unavailable, then the SteelHead drops the traffic.

Figure 7-10. Secure path



New Rule

App Definition

Application/Application Group: RDP

Uplinks

1 st Uplink Type:	Internet (secured)	DSCP:	Preserve	x
2 nd Uplink Type:	--	DSCP:	Preserve	x
3 rd Uplink Type:	--	DSCP:	Preserve	x

Add Uplink

If all the above uplinks are down: Drop

Save Revert

4. Configure a SteelHead as an available controller through the SteelHead CLI.

This step is not mandatory on all SteelHeads. We recommend making at least two SteelHeads available as controllers so you can streamline operations when selecting a new active controller. Use the following commands to configure a SteelHead as an available controller:

```
DC-SH (config) # stp-controller enable
DC-SH (config) # stp-controller address private-ip 172.16.250.132
```

After you configure the SteelHead, it appears in the list of available controllers on the SCC.

5. To make one of the SteelHeads the active controller for the group, click **Make Active** (Figure 7-11).

Figure 7-11. Active controller

Active Secure Transport Controller			
Appliance	Public Address	Private Address	Active
• scw-DC1	10.33.249.140:4500	172.16.250.132	Active
• scw-BR1	10.33.249.139:4500	172.16.249.132	Make Active

Deployment example of a hybrid network backhaul scenario

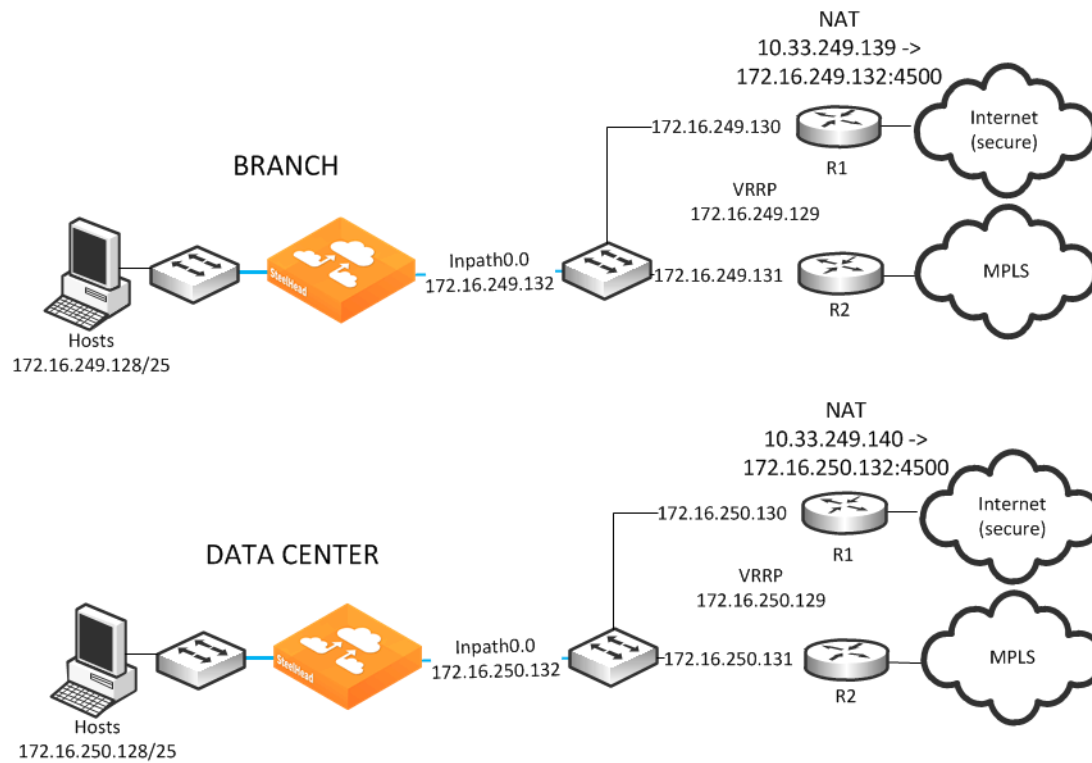
Figure 7-12 shows a single SteelHead located at the data center and the branch office. Each SteelHead has a single in-path interface. Each site has a separate Layer-3 device providing routing services for a private WAN (such as MPLS) and a public WAN (such as the internet). A general purpose VPN is not incorporated at the branch office for internet traffic encryption. Instead internet access is meant to offload corporate traffic from the private WAN onto the public WAN. This deployment is a hybrid network backhaul use case for the secure transport feature.

The example also includes encryption services for an application over the MPLS that demonstrates another use case.

Figure 7-12 has a path selection policy that can leverage this internet offload path and provides encryption services for traffic selected on that path. Because that path is on the public internet, other traffic doesn't use the internet path because this path is not available unless the secure transport path is operational.

Each site is configured in the SCC.

Figure 7-12. Basic deployment example



To configure a basic secure transport deployment example

1. Configure the branch office by associating the SteelHead with this site and the network (Figure 7-13).

Figure 7-13. Branch office site

Edit a Site

Basic information

Site Name: Description:

Site Type:

Region:

If site type and/or region do not exist, they will be created for you.

Contact information

Name: Address:

Job Title: City/Town:

Email: State/Province:

Phone: Country:

Riverbed Appliances

Add Appliance:

Enter hostname or Serial Number or IP address

Hostname: Branch
IP: 10.33.249.135

Network Information

Subnets:

Separate with comma ",", or semicolon ";", or new line

2. Configure the uplinks for the branch office (Figure 7-14).

Figure 7-14. Uplinks at branch office

Uplinks

You can add uplinks manually or select a connectivity template to save time.

+ Add New Uplink

Uplink 1

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☒ Is Default for inpath0_0

☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

Uplink 2

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☐ Is Default for inpath0_0

☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

3. Configure the public IP address and port of the in-path interface in the SteelHead Management Console (Figure 7-15) because the internet path uses NAT traversal.

Figure 7-15. Branch office NAT

In-Path Interface Settings:

Interface	Optimization Interface
▼ inpath0_0	172.16.249.132/27

Interface

IPv4 Address:

IPv4 Subnet Mask:

In-Path Gateway IP:

NAT IPs and ports:

4. Configure the data center by associating the SteelHead with this site and the network (Figure 7-16).

Figure 7-16. Data center site

Edit a Site

Basic information

Site Name	<input type="text" value="DC"/>	Description	<div></div>
Site Type	<input type="text" value="Data Center"/>		
Region	<input type="text" value="DC"/>		

If site type and/or region do not exist, they will be created for you.

Contact information

Name	<input type="text"/>	Address	<input type="text"/>
Job Title	<input type="text"/>	City/Town	<input type="text"/>
Email	<input type="text"/>	State/Province	<input type="text"/>
Phone	<input type="text"/>	Country	<input type="text"/>

Riverbed Appliances

Add Appliance:

Enter hostname or Serial Number or IP address

Hostname: Data-Center-1

IP: 10.33.249.131

Network Information

Subnets:

Separate with comma ",", or semicolon ";", or new line

5. Configure the uplinks for the data center (Figure 7-17).

Figure 7-17. Uplinks at data center

Uplinks

You can add uplinks manually or select a connectivity template to save time.

+ Add New Uplink

Uplink 1

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☒ Is Default for inpath0_0

☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

Uplink 2

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☐ Is Default for inpath0_0

☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

6. Configure the public IP address and port of the in-path interface in the SteelHead Management Console (Figure 7-18) because the internet path uses NAT traversal.

Figure 7-18. Data center NAT

In-Path Interface Settings:

Interface	Optimization Interface
▼ inpath0_0	172.16.250.132/25

Interface

IPv4 Address:

IPv4 Subnet Mask:

In-Path Gateway IP:

NAT IPs and ports:

Figure 7-19 shows how the network table looks upon completion. The internet path is configured as public and securable.

Figure 7-19. Internet path

The screenshot shows the 'Networks' configuration page in the SteelCentral Controller. The page title is 'Networks' and the subtitle is 'Represents the WAN clouds that sites use to communicate to each other using Primary MPLS, VSAT or Internet.' Below the subtitle is a '+ Add a Network' button. A table lists the configured networks:

Name	Public	Securable
▶ MPLS	No	Yes
▼ Internet	Yes	Yes

The 'Internet' network is selected, and its configuration details are shown below the table:

Network Name:

☒ Securable using Secure Transport

☒ Public Network

At the bottom are 'Apply' and 'Revert' buttons.

Figure 7-20 shows how the MPLS path is configured as securable.

Figure 7-20. MPLS path

The screenshot shows the 'Networks' configuration page in the SteelCentral Controller. The page title is 'Networks' and the subtitle is 'Represents the WAN clouds that sites use to communicate to each other using Primary MPLS, VSAT or Internet.' Below the subtitle is a '+ Add a Network' button. A table lists the configured networks:

Name	Public	Securable
▼ MPLS	No	Yes

The 'MPLS' network is selected, and its configuration details are shown below the table:

Network Name:

☒ Securable using Secure Transport

☐ Public Network

At the bottom are 'Apply' and 'Revert' buttons. Below these buttons, the 'Internet' network is visible in the table with 'Public' set to 'Yes' and 'Securable' set to 'Yes'.

To complete the path selection policy, you must create path selection rules on the SCC. In the example, traffic using TCP source port 5001 or TCP destination port 5001 is configured to use the MPLS-secure path and be dropped if that path is not available. Traffic using TCP source port 5002 or TCP destination port 5002 is configured to use the internet path and be dropped if that path is not available. All other traffic is relayed. **Figure 7-21** shows an example of specifying a custom application to take a securable path and drop if that path is not available. However, you can use different path selection rules to meet the business requirements.

Figure 7-21. Path selection rules

Enable Path Selection

☒ Enable Path Selection for SteelHeads running RIOS 9.0 or greater

Path Selection Rules

+ Add a Rule

Application	Type	Action	DSCP
Internet Bound Traffic	System Rule	1 st Relay	1 st --
▶ NC-5001	Application	1 st MPLS (secured) 2 nd Drop	1 st Preserve 2 nd --
▶ NC-5002	Application	1 st Internet (secured) 2 nd Drop	1 st Preserve 2 nd --
▶ NC-5001-SEND	Application	1 st MPLS (secured) 2 nd Drop	1 st Preserve 2 nd --
▶ NC-5002-SEND	Application	1 st Internet (secured) 2 nd Drop	1 st Preserve 2 nd --
▶ Any	System Rule	1 st MPLS 2 nd Relay	1 st Preserve 2 nd --

7. From the SteelHead CLI, configure each SteelHead as a controller using the following commands:

```
DC-SH (config) # stp-controller address private-ip 172.16.250.132 public-ip 10.33.249.140 port 4500
DC-SH (config) # stp-controller enable
BR-SH (config) # stp-controller address private-ip 172.16.249.132 public-ip 10.33.249.139 port 4500
BR-SH (config) # stp-controller enable
```

Because there are only two SteelHeads in this example, both are registered. However, we recommend as a best practice that you only configure those SteelHeads that are centrally located and have the resources to perform the controller functionality. Note that while the actual configuration is done on each SteelHead at the CLI, you select one of the available controllers as active on the SCC.

Deployment example of a dual in-path interface with split tunnel


Figure 7-22 shows a single SteelHead located at the data center and a single SteelHead at the branch office. Each SteelHead has two in-path interfaces in which each in-path interface sits between a LAN router and a WAN edge router.

Each site has a separate Layer-3 device providing routing services for a private WAN (such as MPLS) and a public WAN (such as the internet). A general purpose VPN is not incorporated for encryption over the internet. Instead, internet access is serving the following purposes in this example:

- Offload corporate traffic from the private WAN onto the public WAN
- Direct internet access (split tunnel)

This example also shows direct internet access, also known as a *split tunnel*, in which traffic that's configured as a secured path to the internet and traffic going to internet destinations, egresses through the same interface. For more information about internet traffic, see [Step 3 in "To manually configure the SCC to manage a single site" on page 65](#).

Figure 7-22. Split tunnel example



To configure a dual in-path interface with split tunnel

Note: Configure each site in the SCC.

1. Configure the branch office site name, administrative information, networks, and internet access.

When you configure a site Direct-to-Internet, traffic not destined to a subnet at another site is relayed. **Figure 7-23** shows a configuration of the appliance at the site, network, and that internet access is directly available.

Figure 7-23. Branch office direct-to-internet configuration

Riverbed Appliances

Add Appliance:

Enter hostname or Serial Number or IP address

Hostname: Branch

IP: 172.16.249.240

Network Information

Subnets:

Separate with comma "," or semicolon ";" or new line

Internet Traffic

Designate how Internet access is achieved at this Site. Only applies when Path Selection is enabled.

☒ Direct to Internet
☐ Backhaul through Site

2. Configure the uplinks at the branch office (**Figure 7-24**).

Figure 7-24. Configure the uplinks for the branch office

Uplinks

You can add uplinks manually or select a connectivity template to save time.

+ Add New Uplink

Uplink 1

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☒ Is Default for inpath1_0
☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

Uplink 2

Network:

Uplink Type:

New uplink type will be created automatically and associated with selected network.

Gateway IP:

Interface:

☒ Is Default for inpath0_0
☐ Enable GRE Tunneling

Bandwidth Up: kbps

Bandwidth Down: kbps

Uplink bandwidth only affects QoS

► Probe Settings

3. Configure the public IP address associated with the in-path interface connected to the internet WAN edge router (Figure 7-25) for the branch office.

Figure 7-25. Configure public ip address for the branch office

The screenshot shows the 'In-Path Interface Settings' window. At the top, there are two tabs: 'Interface' and 'Optimization Interface'. Below the tabs, a table lists the interface 'inpath0_0' with the IP address '172.16.249.19/28'. Below the table, the 'Interface' section contains the following fields:

IPv4 Address:	172.16.249.19
IPv4 Subnet Mask:	255.255.255.240
In-Path Gateway IP:	172.16.249.18
NAT IPs and ports:	10.33.249.139:4500

4. Configure the data center site name, administrative information, networks, and internet access. The data center also has Direct-to-Internet configured (Figure 7-26).

Figure 7-26. Data center direct-to-internet configuration

The screenshot shows the 'Riverbed Appliances' configuration window. On the left, under 'Add Appliance:', there is a text input field. Below it, a list of appliances is shown, with the first one having the hostname 'Data-Center-1' and IP '172.16.250.8'. On the right, the 'Network Information' section shows 'Subnets:' with the value '172.16.250.0/24'. Below this, the 'Internet Traffic' section has a radio button selected for 'Direct to Internet' and a text input field for 'N/A'.

5. Configure the uplinks at the data center (Figure 7-27).

Figure 7-27. Configure the uplinks for the data center

Uplinks
You can add uplinks manually or select a connectivity template to save time.

+ Add New Uplink

Uplink 1

Network: MPLS

Uplink Type: MPLS

New uplink type will be created automatically and associated with selected network.

Gateway IP: 172.16.250.34

Interface: inpath1_0

☒ Is Default for inpath1_0

☐ Enable GRE Tunneling

Bandwidth Up: 5000 kbps

Bandwidth Down: 5000 kbps

Uplink bandwidth only affects QoS

► Probe Settings

Uplink 2

Network: Internet

Uplink Type: Internet

New uplink type will be created automatically and associated with selected network.

Gateway IP: 172.16.250.18

Interface: inpath0_0

☒ Is Default for inpath0_0

☐ Enable GRE Tunneling

Bandwidth Up: 10000 kbps

Bandwidth Down: 10000 kbps

Uplink bandwidth only affects QoS

► Probe Settings

6. Configure the public IP address associated with the in-path interface connected to the internet WAN edge router (Figure 7-28).

Figure 7-28. Configure public IP address for the data office

In-Path Interface Settings:

Interface	Optimization Interface
▼ inpath0_0	172.16.250.19/28

Interface

IPv4 Address: 172.16.250.19

IPv4 Subnet Mask: 255.255.255.240

In-Path Gateway IP: 172.16.250.18

NAT IPs and ports: 10.33.249.140:4500

Figure 7-29 and Figure 7-30 show that the internet and MPLS networks are both marked secure while the internet network is also marked as public. The secure designations provide an internet-secure path and an MPLS-secure path to be created for applications to use. The public designation for the internet path allows for traffic to be UDP encapsulated and encrypted to traverse the NAT occurring at the internet WAN edge routers.

Figure 7-29. Internet network

The screenshot shows the 'Networks' configuration page. At the top, it says 'Represents the WAN clouds that sites use to communicate to each other using Primary MPLS, VSAT or Internet.' Below this is a table with columns 'Name', 'Public', and 'Securable'. The 'Internet' network is selected and highlighted in orange. Below the table, the 'Network Name' is set to 'Internet'. The 'Securable using Secure Transport' checkbox is checked, and the 'Public Network' checkbox is also checked. At the bottom are 'Apply' and 'Revert' buttons.

Name	Public	Securable
MPLS	No	Yes
Internet	Yes	Yes

Network Name:

☒ Securable using Secure Transport

☒ Public Network

Figure 7-30. MPLS network

The screenshot shows the 'Networks' configuration page. At the top, it says 'Represents the WAN clouds that sites use to communicate to each other using Primary MPLS, VSAT or Internet.' Below this is a table with columns 'Name', 'Public', and 'Securable'. The 'MPLS' network is selected and highlighted in orange. Below the table, the 'Network Name' is set to 'MPLS'. The 'Securable using Secure Transport' checkbox is checked, and the 'Public Network' checkbox is unchecked. At the bottom are 'Apply' and 'Revert' buttons.

Name	Public	Securable
MPLS	No	Yes

Network Name:

☒ Securable using Secure Transport

☐ Public Network

- From the SteelHead CLI, configure two SteelHeads as an available controller using the following commands:

```
DC-SH (config) # stp-controller address private-ip 172.16.250.19 public-ip 10.33.249.140 port 4500
DC-SH (config) # stp-controller enable
BR-SH (config) # stp-controller address private-ip 172.16.249.35 public-ip 10.33.249.139 port 4500
BR-SH (config) # stp-controller enable
```

Because there are only two SteelHeads in this example, both are registered. However, we recommend as a best practice that you only configure those SteelHeads that are centrally located and have the resources to perform the controller functionality. While the actual configuration is done on each SteelHead at the CLI, you select one of the available controllers as active on the SCC.

- Configure the relevant path selection rules to match the business objectives.

Figure 7-31 shows a custom application and business bulk applications are selected over the internet-secure path. Another application and all the remaining application groups are selected over the MPLS-secure path.

Figure 7-31. Path selection rules

Path Selection Rules				
+ Add a Rule				
Application	Type	Action	DSCP	
Internet Bound Traffic	System Rule	1 st Relay	1 st --	i
▶ NC-5001	Application	1 st Internet (secured) 2 nd Drop	1 st Preserve 2 nd --	
▶ NC-5002	Application	1 st MPLS (secured) 2 nd Drop	1 st Preserve 2 nd --	
▶ NC-5001-SEND	Application	1 st Internet (secured) 2 nd Drop	1 st Preserve 2 nd --	
▶ NC-5002-SEND	Application	1 st MPLS (secured) 2 nd Drop	1 st Preserve 2 nd --	
▶ Business Bulk	Application Group	1 st Internet (secured) 2 nd MPLS 3 rd Relay	1 st Preserve 2 nd Preserve 3 rd --	
▶ Any	System Rule	1 st MPLS 2 nd Internet (secured) 3 rd Relay	1 st Preserve 2 nd Preserve 3 rd --	i

9. Push the policy to the SteelHeads.

After the policy is pushed, the SteelHeads can use the secured paths. **Figure 7-32** shows a connection that was optimized and path selected over the internet-secure path. The connection was encrypted using secure transport.

Figure 7-32. Secure transport encryption

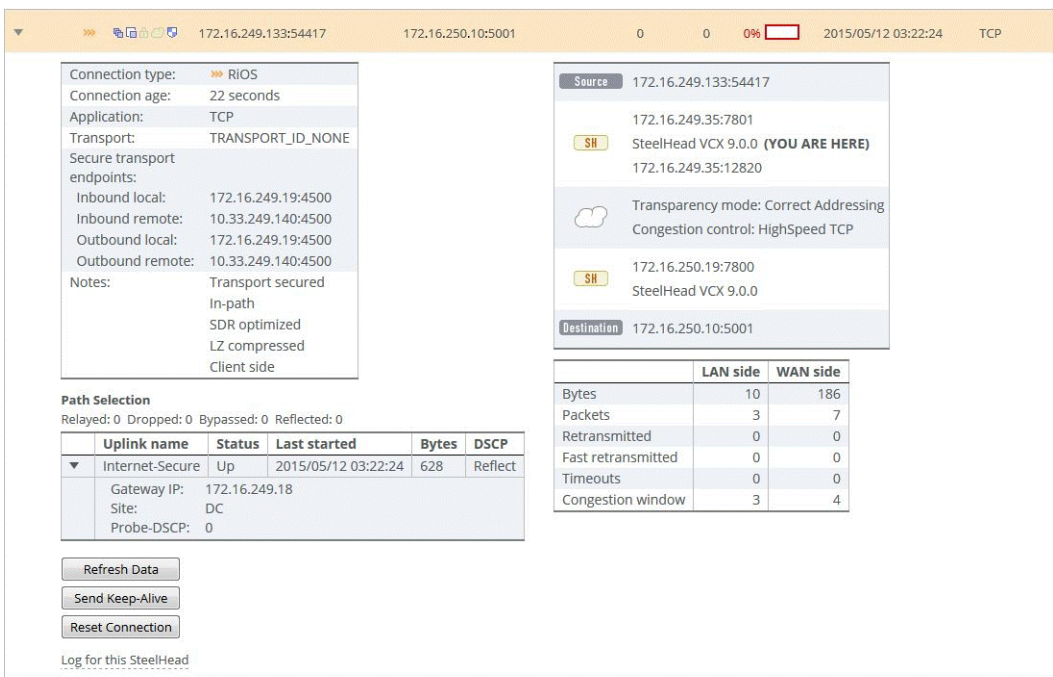
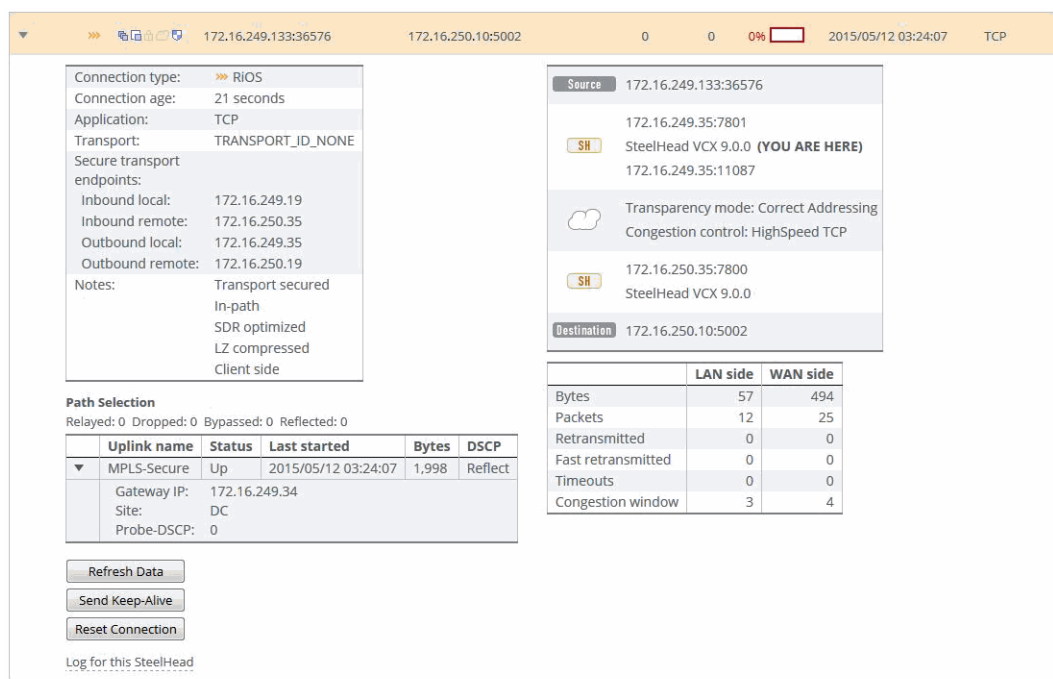


Figure 7-33 shows another connection that was optimized and path selected over the MPLS-secure path. This connection was also encrypted using secure transport but was not UDP encapsulated.

Figure 7-33. MPLS-secure path



Internet connections at the branch were identified and relayed as the site has Direct-to-Internet configured. **Figure 7-34** shows a connection that appears as Failed Terminated because the SteelHead attempted autodiscovery and there was no SteelHead on the path to the server. **Figure 7-35** shows how changing the in-path rules causes the connection to appear as intentional pass-through traffic and the same path selection action, relay, is performed.

Figure 7-34. Failed terminated

▼ ➔ 172.16.249.133:51737 205.251.215.188:80

Connection type:	➔ Failed terminated
Connection age:	17 seconds
Transport:	TRANSPORT_ID_NONE
Passthrough reason:	No Steelhead on path to server

Path Selection
Relayed: 3,124 Dropped: 0 Bypassed: 0 Reflected: 0

Uplink name	Status	Last started	Bytes	DSCP
(Relayed)		2015/05/12 02:32:52	3,124	

Refresh Data
Reset Connection

Log for this SteelHead

Figure 7-35. Traffic passing through)

▼ ⇄ 172.16.249.133:56995 74.125.239.96:443

Connection type:	⇄ Intentional Passthrough
Connection age:	1 minute, 38 seconds
Transport:	TRANSPORT_ID_NONE
Passthrough reason:	In-path rule

Path Selection
Relayed: 4,736 Dropped: 0 Bypassed: 0 Reflected: 0

Uplink name	Status	Last started	Bytes	DSCP
(Relayed)		2015/05/12 02:29:04	4,736	

Refresh Data
Reset Connection

Log for this SteelHead

Reporting

Figure 7-36 shows details in the SteelHead current connection report that indicate whether an individual connection is using secure transport and on which uplink traffic was sent.

Figure 7-36. Current connections report

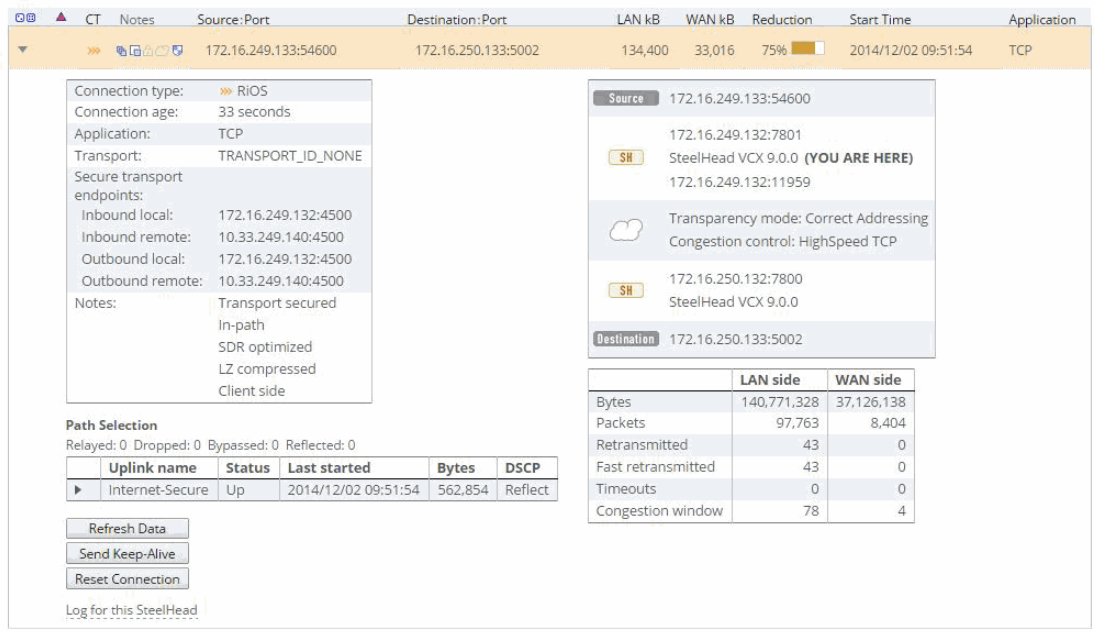


Figure 7-37 shows aggregate traffic information that's available in the SCC on the Secure Transport page.

Figure 7-37. Traffic information



Appliance Clusters

This chapter describes the basic appliance cluster operations for the SCC. It includes the following sections:

- [“Overview of appliance clusters” on page 129](#)
- [“Configuring the SCC to manage appliance clusters” on page 129](#)

Overview of appliance clusters

SteelHead and SteelHead Interceptor clusters are sets of appliances collaborating to provide optimization in complex architectures in such a way that they act as one large component.

You must be running SCC 8.5 or later to configure and manage clusters.

For more information about clusters, see the *SteelHead Interceptor User’s Guide* and the *SteelHead Interceptor Deployment Guide*.

We recommend that you use the SCC to manage appliance clusters for the following reasons:

- **Enables easier configuration, operation, and management** - You tend to make fewer errors because you create one rule in one place for all cluster members (load balancing rules, and so on). You can also manage all the IP addresses in the cluster, simplifying cluster modification when you add an appliance.
- **Provides graphical representation of a particular topology** - With the SCC Cluster Configuration Wizard, you create a graphical representation of your topology, which creates an easier environment for you to understand visually.

Configuring the SCC to manage appliance clusters

This section focuses on a high-level configuration of the cluster. For complete details on cluster configuration, see the *SteelCentral Controller for SteelHead User’s Guide*.

We recommend that you create a group for all clustered appliances, and inside that group create a subgroup for SteelHead Interceptors and another subgroup for SteelHeads. A graphical representation of your cluster is shown during the configuration process when you use the Cluster Configuration Wizard. Breaking appliances into subgroups enables you to view a tree that matches the elements of the cluster from the Manage > Appliances page.

After you create a cluster, the appliances composing the cluster are locked. For example, in the Appliances page, you can't delete a SteelHead that's a member of a cluster.

Figure 8-1. Cannot delete appliance error message



Cannot delete appliance VC1MY000DA255 because it belongs to the 'goatfood' cluster.

Because a cluster is a key element of the optimization layer, this locking down mechanism also serves as a security measure, avoiding deletion by mistake. You can add or remove SteelHeads from a defined cluster using the Cluster Configuration Wizard. After making any changes to the cluster, perform another cluster push to apply those changes to the appliances.

To configure the SCC to manage appliance clusters

1. Register SteelHead and SteelHead Interceptors.
2. Configure in-path interfaces for each appliance in the cluster.

You must correctly configure the in-path addresses of every appliance in the cluster, because the cluster protocol uses these IP addresses for communication.

Connection forwarding neighbors are automatically configured with the Cluster Configuration Wizard, shown in [Step 4 on page 130](#).

3. Configure the general policy settings.

Use policies to configure individual appliance options (because these configurations aren't a part of the cluster configuration).

For example, configure a SteelHead to specify a particular CIFS configuration. You can also apply a policy to the Cluster group: for example, if you want to use a port label for a particular application, you can configure this policy for all members of the cluster.

4. Use the Cluster Configuration Wizard to define the cluster.

The Cluster Configuration Wizard binds the appliances into a cluster and allows you to choose among the well-known SteelHead Interceptor cluster topologies.

By the time you complete the wizard, every cluster member has the proper configuration, including failover and connection forwarding. At the same time, the wizard creates a graphical representation of the cluster.

After you configure the cluster, you can check in the cluster by choosing Manage > Topology: Clusters. **Figure 8-2** shows an example cluster, which consists of the SteelHead Interceptors and two SteelHeads.

Figure 8-2. Example cluster

Cluster Name	Topology	Interceptors	SteelHeads
goatfood	1 Path	ica / 192.168.1.214	SH1 / 192.168.1.208, SH2 / 192.168.1.210

Figure 8-3 shows the cluster members in the Manage > Topology: Appliances page. Next to the cluster members is the cluster icon.

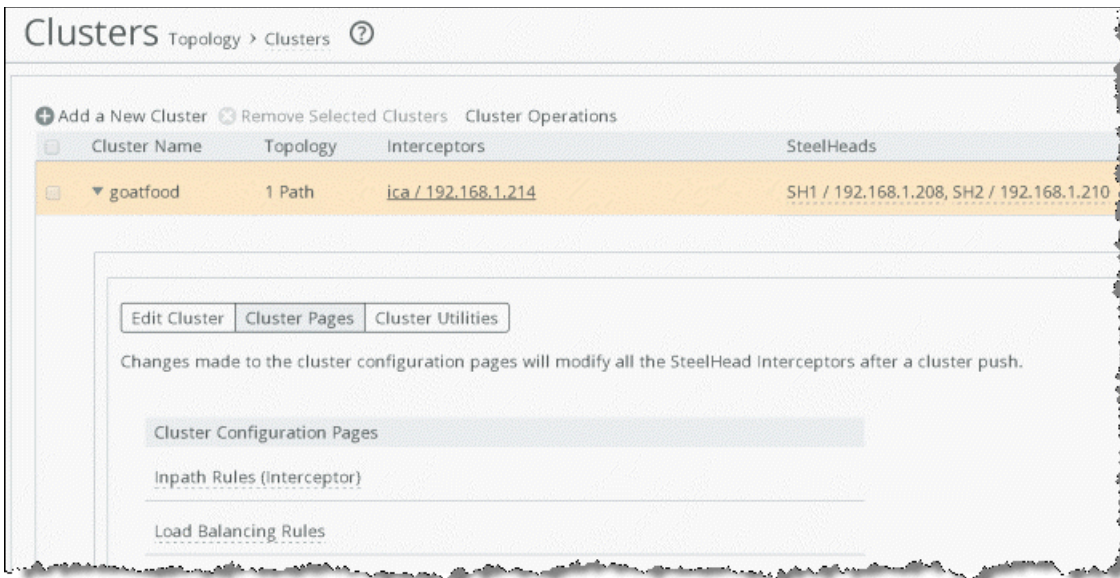
Figure 8-3. Single cluster members

Groups and Managed	Product / Model	Connection	Cluster Icon
Global			
SH1 / 192.168.1.208 (VC1MY000DA255)	SteelHead VCX555H	Connected: Healthy	
SH2 / 192.168.1.210 (VC1VG000DA593)	SteelHead VCX255H	Connected: Healthy	
ica / 192.168.1.214 (H18QW000185E8)	SteelHead Interceptor 500	Connected: Healthy	

5. Configure cluster rules.

You can configure the cluster rules from the Manage > Topology: Cluster page. You can apply all rules, avoiding human configuration errors because you configure once and apply that configuration to all appliances.

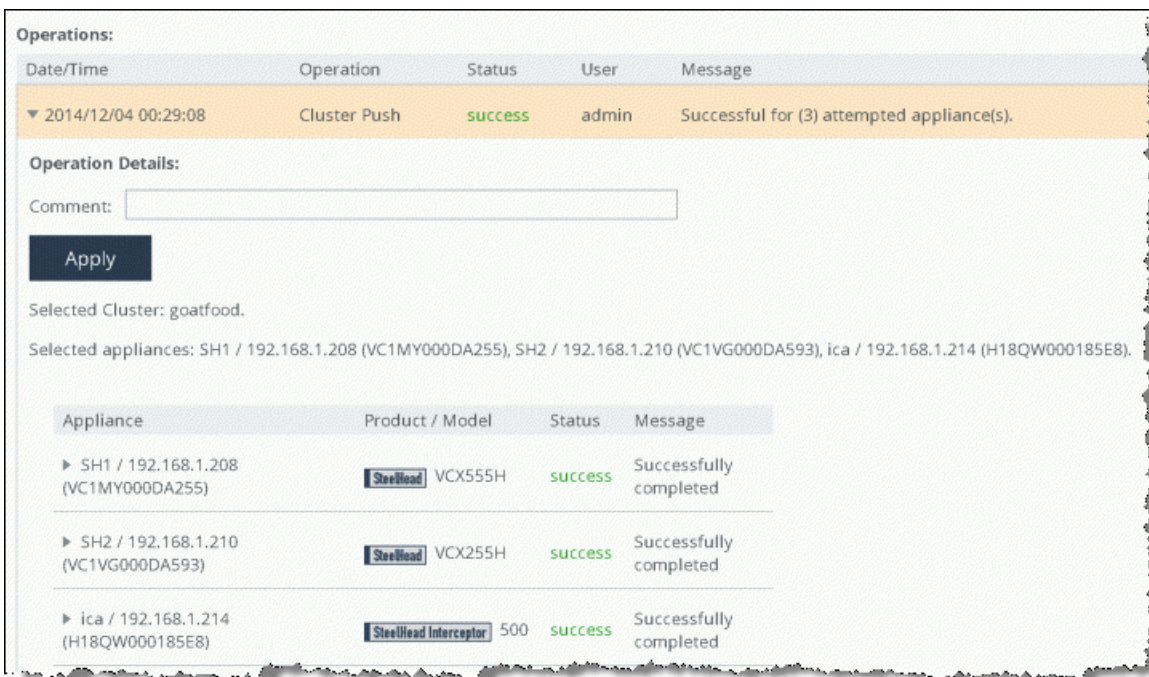
Figure 8-4. Cluster configuration



6. Push the general policy configuration you created in **Step 3** to the cluster members.

7. Push the cluster policy you created in **Step 4** to the cluster members.

Figure 8-5. Policy push success



If the policy push fails, read through the operation history. The history provides information about the policy push failure. **Figure 8-6** shows an example failure.

Figure 8-6. Policy push failure

The screenshot displays a web interface for a policy push operation. At the top, a status bar indicates the operation is 'failed' and 'Failed to all (3) appliance(s)'. Below this, the 'Operation Details' section includes a 'Comment' field and an 'Apply' button. The 'Selected Cluster' is 'goatfood', and the 'Selected appliances' are listed as SH1, SH2, and ica. A table below provides details for each appliance, showing that the policy push failed for all three.

Appliance	Product / Model	Status	Message
▶ SH1 / 192.168.1.208 (VC1MY000DA255)	VCX555H	failed	Applying a change failed
▶ SH2 / 192.168.1.210 (VC1VG000DA593)	VCX255H	failed	Applying a change failed
▶ ica / 192.168.1.214 (H18QW000185E8)	500	failed	Applying a change failed

Consider the following factors when you create clusters:

- Changes made to the cluster configuration pages modify all the SteelHead Interceptors in the cluster after a cluster push.
- The cluster push operation results in either all or nothing success.
- Use the Cluster page instead of individual Appliance pages. By doing so, you avoid mistakes and have a graphical representation of your design.