

# **SteelFusion™ Core Management Console User's Guide**

Version 4.3

March 2016

**riverbed®**

© 2015 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 2012 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 © 1998-2013 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

Phone: 415-247-8800  
Fax: 415-247-8801  
Web: <http://www.riverbed.com>

Part Number  
712-00077-09

# Contents

<b>Preface.....</b>	<b>7</b>
About This Guide .....	7
Audience .....	8
Document Conventions.....	8
Documentation and Release Notes .....	8
Contacting Riverbed.....	8
 <b>Chapter 1 - Overview of the SteelFusion Core Management Console .....</b>	<b>11</b>
Product Dependencies and Compatibility.....	11
Hardware and Software Dependencies.....	12
Firewall Requirements .....	12
SNMP-Based Management Compatibility.....	12
How SteelFusion Works.....	12
Using the SteelFusion Core Management Console.....	14
Connecting to the SteelFusion Core Management Console .....	14
The Dashboard .....	15
Navigating in the SteelFusion Core Management Console.....	20
Getting Help .....	21
Displaying Online Help.....	21
Downloading Documentation .....	22
 <b>Chapter 2 - Configuring Storage .....</b>	<b>23</b>
Understanding Basic Procedures.....	23
Saving Your Configuration .....	23
Printing Pages and Reports.....	24
Restarting the Core Service .....	24
Logging Out .....	24
Performing the Initial Setup.....	24
Before You Begin.....	24
Configuring Network Settings .....	25
Mapping LUNs to an Edge .....	26

Importing a Saved Configuration .....	28
Exporting a Configuration .....	29
Configuring iSCSI Settings .....	30
Configuring LUNs .....	35
Fibre Channel LUN Configuration Steps in ESXi .....	35
Block-Disk LUN Considerations .....	36
Configuring Branch Recovery .....	42
Troubleshooting Branch Recovery .....	44
Configuring Edges .....	45
Configuring Replication .....	47
Base Requirements .....	48
Before You Begin .....	48
Basic Steps .....	49
Setting Up the Data Centers for Replication .....	49
Pairing the Cores .....	52
Configuring the Witness .....	53
Configuring Edges and LUNs for Replication .....	55
Suspending Replication .....	57
Initiating Failover (Recovering from Primary Data Center Failure) .....	58
Failing Back to the Primary Data Center .....	59
Terminating Replication .....	61
Configuring CHAP Users .....	63
Configuring Snapshots and Proxy Backup .....	65
Understanding Crash Consistency and Application Consistency .....	66
Configuring Snapshots .....	66
Configuring Snapshots for Storage Arrays .....	67
Configuring Handoff Hosts .....	69
Configuring Snapshot Schedule Policies .....	70
Defining Branch and Proxy Hosts .....	72
Configuring Snapshots for LUNs .....	73
Applying a Snapshot Schedule Policy to a LUN .....	74
Configuring Application-Consistent Snapshots for a LUN .....	75
Configuring Proxy Backup for a LUN .....	77
Configuring Failover .....	80
Configuring Pool Management .....	82
Configuring REST API Access .....	85
Disabling REST API Access .....	86
Best Practices .....	87
Troubleshooting .....	87
<b>Chapter 3 - Modifying Host and Network Settings .....</b>	<b>89</b>
Configuring Host Settings .....	89
Configuring the Management Interfaces .....	92
Configuring the Data Interfaces .....	95

Configuring the Core for Jumbo Frames.....	96
<b>Chapter 4 - Configuring System Settings .....</b>	<b>97</b>
Creating Announcements.....	97
Setting Alarm Parameters.....	98
Configuring Date and Time.....	103
Current NTP Server Status.....	104
NTP Servers.....	104
NTP Authentication .....	105
Setting SNMP Parameters and Trap Receivers.....	106
Creating SNMPv3 Users .....	108
Configuring SNMP Authentication and Access Control .....	110
Setting Up Email Notifications .....	114
Configuring Logging.....	114
Setting Up System Logging.....	115
Configuring Remote Log Servers.....	116
Filtering Logs by Application or Process .....	117
Managing Configuration Files .....	118
<b>Chapter 5 - Configuring Security Settings .....</b>	<b>121</b>
Configuring General Security Settings .....	121
Managing User Permissions.....	122
Accounts .....	123
Managing Password Policy .....	126
Selecting a Password Policy .....	126
Configuring RADIUS Server Authentication .....	130
Configuring TACACS+ Server Authentication .....	132
Unlocking the Secure Vault .....	135
Configuring Web Settings.....	136
Managing Web SSL Certificates.....	137
<b>Chapter 6 - Maintaining Your System .....</b>	<b>141</b>
Starting, Stopping, and Restarting the Service.....	141
Displaying Scheduled Jobs and Job Status.....	142
Managing Licenses .....	143
Removing a License.....	144
Fetching a License.....	144
Upgrading the Software.....	144
Rebooting and Shutting Down the Core .....	146
Changing the Administrative Password .....	146

<b>Chapter 7 - Displaying and Customizing Reports .....</b>	<b>149</b>
Viewing Storage Reports .....	149
Accessing Settings from Storage Reports.....	149
Viewing the SteelFusion Edge Stats.....	150
Viewing the SteelFusion Edge Trends .....	153
Viewing the LUN I/O Metrics Report.....	155
Viewing the SAN I/O Metrics Report.....	157
Viewing the Replication Data Sync - Remaining Bytes Report.....	158
Viewing the Replication Journal I/O Report.....	159
Viewing the Replication Write I/O Report.....	161
Building Custom Reports with the Report Builder .....	163
Viewing the Networking Interface Counters Report.....	164
What This Report Tells You.....	164
Viewing Diagnostic Reports.....	165
Viewing Alarm Status Reports .....	165
Viewing CPU Utilization Reports .....	170
Viewing Memory Paging Reports.....	171
Viewing Logs.....	172
Viewing User Logs .....	172
Viewing System Logs.....	174
Downloading Log Files.....	175
Generating Dumps .....	177
Generating System Dumps .....	178
Viewing Process Dumps.....	179
Capturing and Uploading TCP Dumps .....	179
Viewing a TCP Dump .....	185
<b>Appendix A - SteelFusion Core MIB .....</b>	<b>187</b>
Accessing the Core MIB.....	187
SNMP Traps.....	188

# Preface

Welcome to the *SteelFusion Core Management Console User's Guide*. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, additional reading, and contact information. This preface includes the following sections:

- [“About This Guide” on page 7](#)
- [“Documentation and Release Notes” on page 8](#)
- [“Contacting Riverbed” on page 8](#)

---

## About This Guide

This guide describes how to use the SteelFusion Core Management Console to configure product features, view reports, and modify host and network settings.

This guide includes information relevant to the following products:

- Riverbed SteelFusion Core (Core)
- Riverbed SteelFusion Core Virtual Edition (Core-v)
- Riverbed SteelFusion Edge (Edge)
- Riverbed Optimization System (RiOS)
- Riverbed SteelHead EX (SteelHead EX)
- Riverbed Virtual Services Platform (VSP)

This guide is intended to be used together with the following documentation:

- *SteelFusion Core Installation and Configuration Guide*
- *SteelFusion Command-Line Interface Reference Manual*
- *SteelFusion Design Guide*
- *Fibre Channel on SteelFusion Core Virtual Edition Solution Guide*

## Audience

This guide is written for storage and network administrators familiar with administering and managing storage arrays, snapshots, backups, virtual machines (VMs), Fibre Channel, and Internet Small Computer System Interface (iSCSI).

## Document Conventions

This guide uses this standard set of typographical conventions.

Convention	Meaning
<i>italics</i>	Within text, new terms and emphasized words appear in <i>italic</i> typeface.
<b>boldface</b>	Within text, CLI commands, CLI parameters, and REST API properties appear in <b>bold</b> typeface.
Courier	Code examples appear in Courier font: <pre>amnesiac &gt; enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: <b>interface</b> <ip-address>
[ ]	Optional keywords or variables appear in brackets: <b>ntp peer</b> <ip-address> [version <number>]
{ }	Elements that are part of a required choice appear in braces: {<interface-name>   <b>ascii</b> <string>   <b>hex</b> <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: { <b>delete</b> <filename>   <b>upload</b> <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](mailto:proserve@riverbed.com) or go to <http://www.riverbed.com/services-training/Services-Training.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](mailto:techpubs@riverbed.com).



## CHAPTER 1 Overview of the SteelFusion Core Management Console

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

- [“Product Dependencies and Compatibility” on page 11](#)
- [“How SteelFusion Works” on page 12](#)
- [“Using the SteelFusion Core Management Console” on page 14](#)
- [“Getting Help” on page 21](#)

---

### Product Dependencies and Compatibility

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

- [“Hardware and Software Dependencies” on page 12](#)
- [“Firewall Requirements” on page 12](#)
- [“SNMP-Based Management Compatibility” on page 12](#)

## Hardware and Software Dependencies

The following table summarizes the hardware and software requirements for the Core.

Riverbed Component	Hardware and Software Requirements
Core appliance	19-inch (483-mm) two-post or four-post rack.
SteelFusion Core Management Console	<p>Any computer that supports a web browser with a color image display.</p> <p>The SteelFusion Core Management Console has been tested with Mozilla Firefox Extended Support Release version 31.0, Google Chrome, and Microsoft Internet Explorer 9.0.</p> <p>JavaScript and cookies must be enabled in your web browser.</p>
SteelFusion Edge Management Console	<p>Any computer that supports a web browser with a color image display.</p> <p>The SteelFusion Edge Management Console has been tested with Mozilla Firefox Extended Support Release version 31.0, Google Chrome, and Microsoft Internet Explorer 9.0.</p> <p>JavaScript and cookies must be enabled in your web browser.</p>

## Firewall Requirements

Riverbed recommends that you deploy the Core behind your firewall. Ports 22, 443, 25, 7970, 7950, 7951, 7952, 7953, and 7954 must be open.

Virus scanning, deep packet inspection (IDS) and other features that scan the traffic should be disabled.

When specifying IP addresses for firewall rules, Riverbed recommends that you add all the interface IP addresses of all Edge devices including in-path IPs. If the Edges are set up for high availability, ensure that you also add the standby Edge IP address.

## SNMP-Based Management Compatibility

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

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

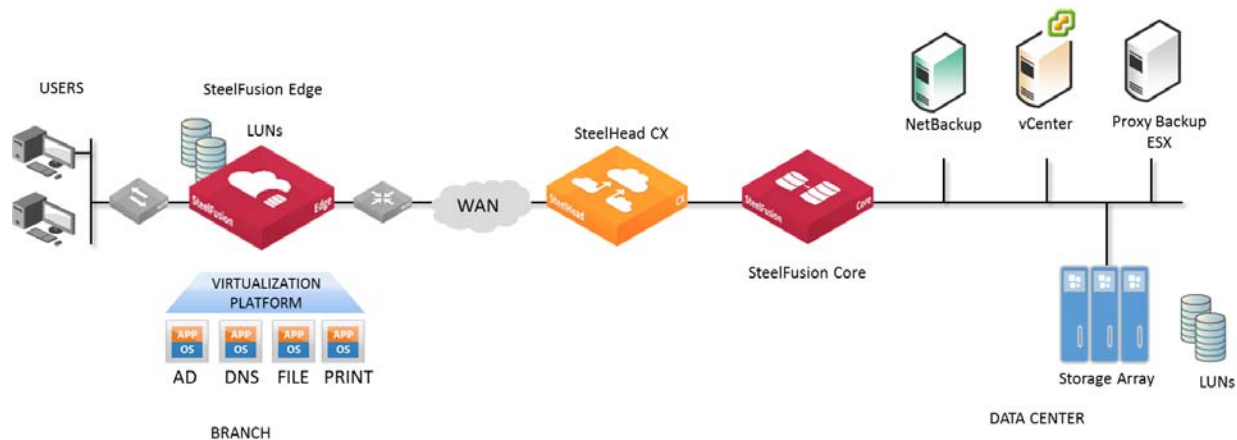
## How SteelFusion Works

SteelFusion enables branch office server systems to efficiently access storage arrays over the WAN. It is typically deployed in conjunction with SteelHeads and is composed of the following components:

- **SteelFusion Core** - A physical or virtual Core is deployed in the data center alongside SteelHeads and a storage array. The Core mounts logical unit numbers (LUNs) provisioned for the branch offices from a storage array, and manages block transfers between these LUNs and Edges. Additionally, SteelFusion Core-v can mount LUNs via Fibre Channel. At the data center, the Core integrates with existing storage systems and SteelHead implementations and connects dedicated LUNs with each Edge at the branch office.
- **SteelFusion Edge** - A standalone or high-availability Edge is deployed in the branch and presents the LUNs mounted on the Core at the branch through iSCSI targets. The Edge can also host local LUNs without a Core connection.
- **LUN** - A unique identifier associated with individual storage devices or collections of storage devices used with SCSI, iSCSI, or Fibre Channel interfaces.
- **Branch server** - The branch-side servers access data from the SteelFusion system instead of a local storage device. This server can also run as a Virtual Services Platform (VSP) VM on the local Edge.
- **Blockstore** - A persistent and authoritative local cache of storage blocks on the Edge that caches the writes on LUNs presented by the Edge. Blockstore also stores cached portions of the LUNs to quickly serve read requests.
- **Data Center SteelHead** - The data center-side SteelHead peer for WAN optimization.

The following diagram shows a typical SteelFusion deployment.

**Figure 1-1. Typical SteelFusion Deployment**



SteelFusion initially populates the blockstore using the following methods:

- **First request** - Blocks are added to the blockstore when first requested. The first request is subject to standard WAN latency because it is the first time the blocks are being served. Subsequent traffic is optimized. Reactive prefetch and policy-based prefetch minimizes these first requests.
- **On-demand prefetch** - The Core observes block requests, intelligently predicts the blocks most likely to be requested in the near future, and then requests those blocks from the data center LUN in advance.
- **Policy-based prefetch** - Configured policies identify the blocks that are likely to be requested at a given branch office site in advance. The Edge then requests those blocks from the data center LUN in advance.

Blocks are transferred between Edges and Cores through an internal protocol. The Core then writes the updates to the data center LUNs through the iSCSI or Fibre Channel protocol. Optionally, you can further optimize traffic between the branch offices and the data center by implementing SteelHeads.

When the branch office server requests blocks, those blocks are served locally from the blockstore (unless they are not present, in which case the Edge retrieves them from the data center LUN via the Core). Similarly, newly written blocks are spooled to the local cache, acknowledged by the Edge to the branch office server, and then asynchronously propagated to the data center. Because each Edge implementation is linked to one or more dedicated LUNs at the data center, the blockstore is authoritative for both reads and writes and can tolerate WAN outages without affecting cache coherency.

Riverbed strongly recommends that you read the SteelFusion Interoperability Matrix at <https://splash.riverbed.com/docs/DOC-4204>.

For information on compatibility between RiOS, Edge, Core, and vSphere releases, see the Knowledge Base article RiOS, SteelFusion Edge, SteelFusion Core and vSphere Release Matrix at <https://supportkb.riverbed.com/support/index?page=content&id=S:S27472>.

---

## Using the SteelFusion Core Management Console

This section describes how to connect to and navigate in the Management Console. If you prefer, you can use the CLI to perform configuring and monitoring tasks. For details, see the *SteelFusion Command-Line Interface Reference Manual*.

### Connecting to the SteelFusion Core Management Console

You can connect to the SteelFusion Core Management Console through any supported web browser.

To connect to the SteelFusion Core Management Console, you must know the IP address and URL for the Core primary interface and the administrator password that you specified during the initial setup of the Core.

---

**Note:** Cookies and JavaScript must be enabled in your web browser.

---

#### To connect to the SteelFusion Core Management Console

1. Enter the URL for the SteelFusion Core Management Console in the Web Address area of your web browser:

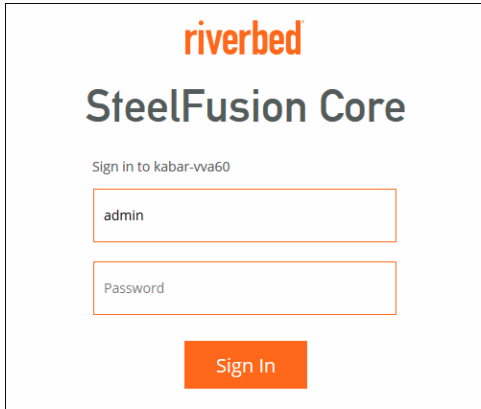
<protocol>://<ip-address>

<protocol> is http or https. HTTPS uses the Secure Sockets Layer (SSL) protocol to ensure a secure environment. If you use HTTPS to connect, you are prompted to inspect and verify the SSL key.

<ip-address> is the IP address for the primary interface of the Core.

The SteelFusion Core Management Console appears, displaying login controls.

**Figure 1-2. Login Page**



2. Specify the user login: admin, monitor, or a login from a RADIUS or a TACACS+ database.

The default login is admin. Users with administrator (admin) privileges may configure and administer the Core. Users with monitor privileges may display SteelFusion Core reports and system logs. A monitor user cannot make any configuration changes. For details about user permissions, see [“Managing User Permissions” on page 122](#).

3. Specify the password you assigned in the Initial Setup wizard.
4. Click **Sign In** to display the Dashboard.

## The Dashboard

The Dashboard displays a dynamic, at-a-glance view of how your entire Edge and storage infrastructure is performing, organized in various panels. Data is refreshed every 60 seconds.

If you are upgrading to version 4.3, the Dashboard will appear after you clear your browser cache.

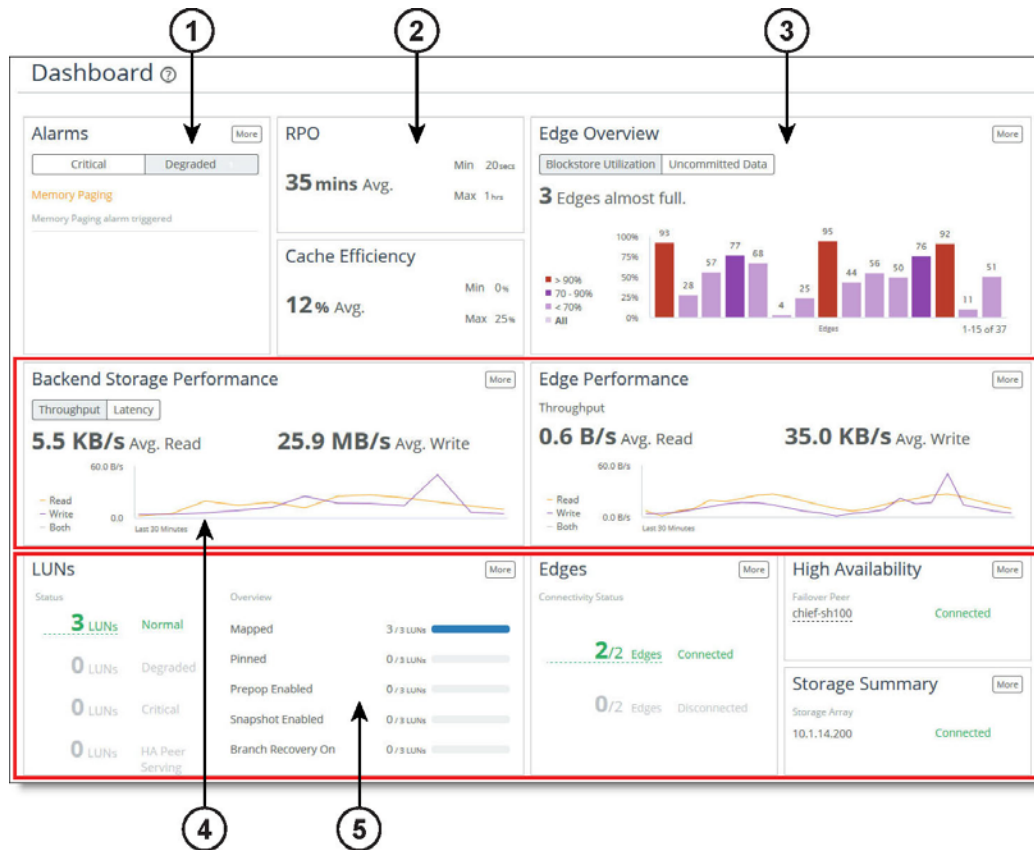
---

**Note:** On the RPO and Cache Efficiency panels, data is only available on Edges running version 4.3 and later.

---

The system health status icon and the system hostname are located in the upper-left corner of the Dashboard page. Hover the mouse over the icon to view the current system health status: Healthy, Degraded, or Critical.

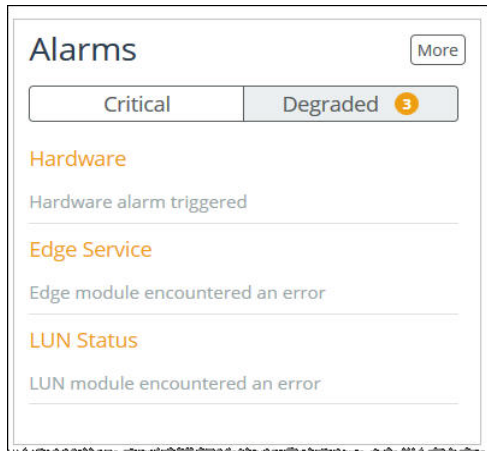
**Figure 1-3. Dashboard**



## ① Alarms

- **Alarms** - Notifies you of any triggered Core alarms that require attention. There are two types of alarms: Critical or Degraded, and the number of triggered alarms appears on each tab. When the page is loaded, the Alarms panel displays the tab with the highest priority alarm by default. If both tabs contain active alarms, the Critical tab is shown by default. Click an alarm to view details on the Alarm Status page. Click **More** to view details on the Alarm Status page.

Figure 1-4. Alarms panel

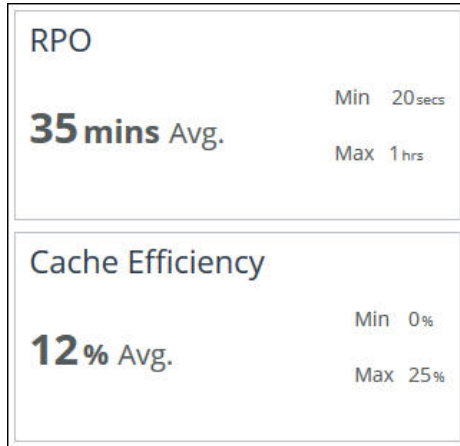


## ② RPO and Cache Efficiency

- **RPO (Recovery Point Objective)** - Summarizes the length of time since data from the Edges was last committed to the backend storage. For example, if the RPO is 3 hours this means that there is a 3 hour commit delay for data from the Edge to the Core, including any HA peers. These Average, Minimum, and Maximum values are taken across all connected Edges running version 4.3 and later. The Average value is useful as an overview of the average delay between the Edges and the Core, while the Maximum value can alert you to a specific Edge that is particularly behind. A lower RPO is desired, and the actual number depends on the infrastructure. Click **More** to view details on the SteelFusion Edge Stats page.

- **Cache Efficiency** - Displays how efficiently SteelFusion's prefetch mechanism reduces the traffic between the Edge and the Core. This value is the percentage of data that is served locally from the cache compared to data read from the data center storage array. These Average, Minimum, and Maximum percentages are calculated across all connected Edges in the infrastructure (including any HA peers).

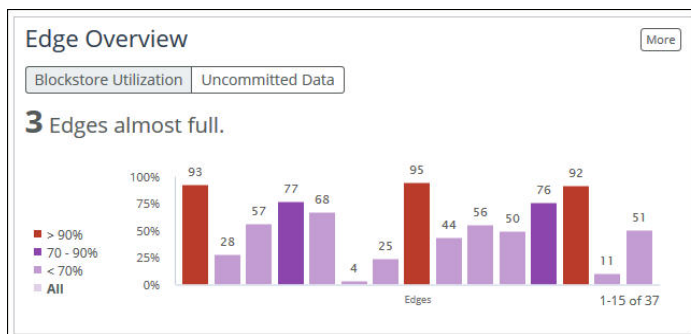
Figure 1-5. RPO and Cache Efficiency panels



### 3 Edge Overview

- **Edge Overview** - Displays an overview of the Edges connected to the Core. Select the Blockstore Utilization tab to display the percentage of the blockstore currently being used at each Edge. Select the Uncommitted Data tab to display the percentage of blockstore data that is pending to be committed to the Core. Click **More** to view details on the SteelFusion Edges page.

Figure 1-6. Edge Overview panel

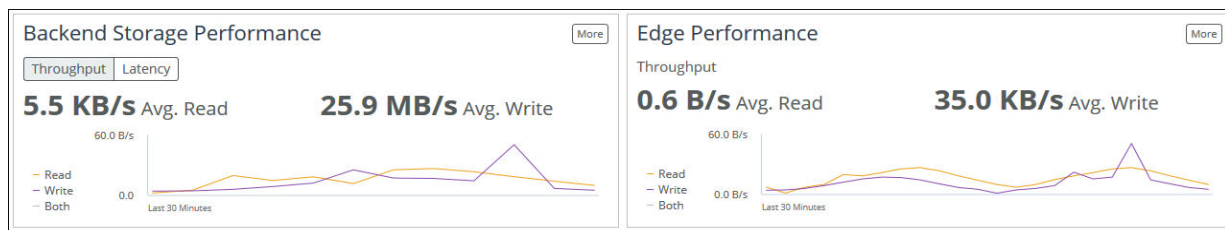


### 4 Backend Storage and Edge Performance

- **Backend Storage Performance** - Displays the performance statistics for all LUNs connected through this Core (and the peer Core, if HA is configured). This is an aggregate value across all LUNs and includes both iSCSI and Fibre Channel. Select the Latency tab to view the average read/write latencies for all LUNs in the last 30 minutes. Select the Throughput tab to view how much data has been written to and read from the LUNs in the last 30 minutes. In both tabs, select **Read**, **Write**, or **Both** to customize the view. Click **More** to view details on the LUN I/O Metrics page.

- **Edge Performance** - Displays the last 30 minutes of performance, measured in total read and write throughput, for all Edges connected to this Core (and the peer Core, if HA is configured). Click **Read**, **Write**, or **Both** to customize the view. Click **More** to view details on the SteelFusion Edges page.

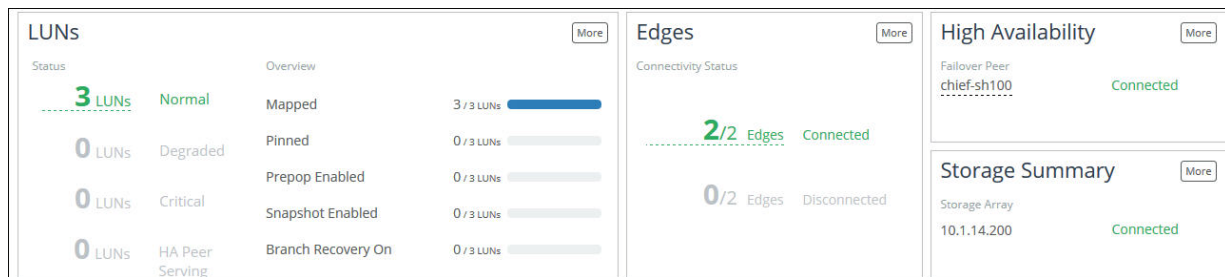
Figure 1-7. Backend Storage Performance and Edge Performance panels



## 5 LUNs, Edges, High Availability, and Storage Summary

- **LUNs** - Displays the total number of LUNs and their current health state: Normal, Degraded, Critical, or HA Peer Serving. In addition, this panel displays the features currently enabled or configured on these LUNs, such as prepopulation, snapshots, branch recovery, or whether the LUNs are mapped or pinned. If LUNs are not configured, click **Add a LUN** to add a LUN to this Core. Click **More** to view details on the LUNs page.
- **Edges** - Displays the number of Edges and their connection status. The data shown is calculated across all Edges in the infrastructure (including the HA peer). If Edges are not configured, click **Add an Edge** to add an Edge to this Core. Click **More** to view details on the SteelFusion Edges page.
- **High Availability** - Displays the failover peer's IP address or name and its failover state (if high availability is configured). If high availability is not configured, click **Add Failover Peer** to add a peer. Click **More** to view details on the Failover Configuration page.
- **Storage Summary** - Displays an overview of storage arrays connected to this Core, as well as their connection status. Click **More** to view details on the iSCSI, Initiators, MPIO page. The information in this panel will change depending on the type and number of connected storage arrays:
  - **iSCSI and fibre channel** - The name or IP address of one iSCSI storage array and the number of fibre channel LUNs.
  - **iSCSI only** - The name or IP address of up to two currently configured iSCSI storage arrays.
  - **Fibre channel only** - The number of fibre channel LUNs.
  - **No iSCSI or fibre channel** - Click **Add a Storage Array** to configure a storage array.

Figure 1-8. LUNs, Edges, High Availability, and Storage Summary panels



## Navigating in the SteelFusion Core Management Console

You can access the tools and reports in the Core using cascading menus.

### To display cascading menus

1. Click the name of the type of information that you want to access in the menu bar to display the submenus.

For example, click **Reports** to display the Storage and Diagnostics submenus. The menu item that is currently active is highlighted.

2. To go to a page, select the menu name you want to display.

For example, choose Configure > Manage: SteelFusion Edges to display the SteelFusion Edges page.

The following table summarizes the cascading menus.

Menu	Description
Dashboard	Displays the Dashboard.
Configure	<p><b>Storage Array</b> - Configure iSCSI settings (initiators and MPIO) and CHAP user settings. For details, see <a href="#">Chapter 2, "Configuring Storage."</a></p> <p><b>Networking</b> - Configure host settings (hostname, DNS servers, hosts, proxies, date, and time), management interfaces (primary, auxiliary, and routing), and data interfaces. For details, see <a href="#">"Configuring Host Settings" on page 89.</a></p> <p><b>Wizards</b> - Use the wizards to quickly perform initial setup, map LUNs to Edges, and import/export Core configurations. For details, see <a href="#">"Performing the Initial Setup" on page 24.</a></p> <p><b>Replication</b> - Configure replication between data centers. For details, see <a href="#">"Configuring Replication" on page 47.</a></p> <p><b>Manage</b> - View currently configured LUNs and Edges or add new ones to the configuration. For details, see <a href="#">"Configuring LUNs" on page 35</a> and <a href="#">"Configuring Edges" on page 45.</a></p> <p><b>Pool Management</b> - Manage up to 32 Cores from a single interface. For details, see <a href="#">"Configuring Pool Management" on page 82.</a></p> <p><b>Backups</b> - Configure snapshots for backup. For details, see <a href="#">"Configuring Snapshots and Proxy Backup" on page 65.</a></p> <p><b>Failover</b> - Configure another device for high-availability in case of failure. For details, see <a href="#">"Configuring Failover" on page 80.</a></p>

Menu	Description
Reports	<b>Storage</b> - Display and download Core storage reports such as Edge connectivity and statistics for LUNs, network traffic, SANs, and replication. For details, see <a href="#">“Viewing Storage Reports” on page 149</a> .
	<b>Report Builder</b> - Select and display two reports on the same page. For details, see <a href="#">“Building Custom Reports with the Report Builder” on page 163</a> .
	<b>Networking</b> - Display statistics for each configured interface. For details, see <a href="#">“Viewing the Networking Interface Counters Report” on page 164</a> .
	<b>Diagnostics</b> - Display and download Core diagnostic reports such as alarm status, CPU utilization, and memory paging. For details, see <a href="#">“Viewing Diagnostic Reports” on page 165</a> .
	<b>Logs</b> - Display and download Core user and system logs. For details, see <a href="#">“Viewing Logs” on page 172</a> .
Settings	<b>Dumps</b> - Display, download, or remove Core system and process dumps. View stored and currently running TCP dumps, add a new TCP dump, or use the TCP Dump scanner. For details, see <a href="#">“Generating System Dumps” on page 178</a> and <a href="#">“Capturing and Uploading TCP Dumps” on page 179</a> .
	<b>Maintenance</b> - Start and stop system services, schedule jobs, upgrade software, back up configurations, and reboot or shut down the appliance. For details, see <a href="#">“Starting, Stopping, and Restarting the Service” on page 141</a> .
	<b>Security</b> - Configure general security parameters, user permissions, RADIUS, TACACS+, secure vault, and web settings. For details, see <a href="#">“Configuring General Security Settings” on page 121</a> .
	<b>System Settings</b> - Configure alarm settings, announcements, email settings, log settings, monitored ports, and SNMP settings. For details, see <a href="#">“Configuring System Settings” on page 97</a> . Modify the administrator user password. For details, see <a href="#">“Changing the Administrative Password” on page 146</a> . Manage configuration files for the system. For details, see <a href="#">“Managing Configuration Files” on page 118</a> .
Help	Display online help and Core documentation; contact information for Riverbed Support; appliance details such as model number, revision type, serial number, software version; and appliance MIB files from this menu. For details, see <a href="#">“Getting Help” on page 21</a> .

## Getting Help

The Support page provides you with the following options:

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

## Displaying Online Help

The SteelFusion Core Management Console provides page level help. You can also display an online help book containing all the help, including an index and table of contents.

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

**To display online help in the Management Console**

- Click the question mark (?) icon next to the page title. The help for the page appears in a new browser window.

## Downloading Documentation

The Riverbed Support site contains PDF versions of the *SteelFusion Core Management Console User's Guide*, *SteelFusion Core Installation and Configuration Guide*, and the *SteelFusion Command-Line Interface Reference Manual*.

**To download the PDF versions of the User's Guide**

1. Go to the Software & Documentation section of the Riverbed Support site at <https://support.riverbed.com/content/support/software.html>.
2. Select the product name.
3. Select the product version from the Display Version drop-down list.  
Select PDF or HTML next to the document name to download the document.

## CHAPTER 2     **Configuring Storage**

This chapter describes how to configure storage settings in the SteelFusion Core Management Console. It includes the following sections:

- [“Understanding Basic Procedures” on page 23](#)
- [“Performing the Initial Setup” on page 24](#)
- [“Configuring iSCSI Settings” on page 30](#)
- [“Configuring LUNs” on page 35](#)
- [“Configuring Branch Recovery” on page 42](#)
- [“Configuring Edges” on page 45](#)
- [“Configuring Replication” on page 47](#)
- [“Configuring CHAP Users” on page 63](#)
- [“Configuring Snapshots and Proxy Backup” on page 65](#)
- [“Configuring Failover” on page 80](#)
- [“Configuring Pool Management” on page 82](#)
- [“Configuring REST API Access” on page 85](#)

---

### **Understanding Basic Procedures**

This section describes the following basic procedures you can perform on a Core:

- [“Saving Your Configuration” on page 23](#)
- [“Printing Pages and Reports” on page 24](#)
- [“Restarting the Core Service” on page 24](#)
- [“Logging Out” on page 24](#)

### **Saving Your Configuration**

The Save icon on the menu bar saves modifications to the configuration. When it is orange, there are unsaved changes to the configuration. For details, see [“Managing Configuration Files” on page 118](#).

## Printing Pages and Reports

You can print pages and reports using the print option on your web browser.

### To print pages and reports

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

## Restarting the Core Service

Some configuration settings apply to the Core service. The Core service is a daemon that runs in the background, performing operations when required. For details, see [“Starting, Stopping, and Restarting the Service” on page 141](#).

---

**Note:** Typically, you restart the Core service whenever your configuration changes affect network IP addresses or interface configurations, or when you add a failover peer. Otherwise, you do not need to restart the service.

---

## Logging Out

Above the menu bar, click **Sign out** to end your session.

---

## Performing the Initial Setup

This section describes how to use the Management Console wizards to quickly configure your network, time zone, and initiator settings. You can also use these wizards to quickly perform basic configuration tasks, such as adding Edges and iSCSI portals.

After initial setup, you can further modify the Core configuration using the relevant pages in the SteelFusion Core Management Console.

This section includes the following topics:

- [“Before You Begin” on page 24](#)
- [“Configuring Network Settings” on page 25](#)
- [“Mapping LUNs to an Edge” on page 26](#)
- [“Importing a Saved Configuration” on page 28](#)
- [“Exporting a Configuration” on page 29](#)

## Before You Begin

To map LUNs to an Edge, you must first complete the following tasks:

- Determine the hostname of your Edge.
- Configure an iSCSI portal with iSCSI targets and LUNs.

- Add the appropriate initiator to the iSCSI Portal Initiator list.

---

**Note:** If you do not know the hostname, iSCSI portal, or initiator name, you can obtain this information from the Start page of the LUN Mapping Wizard.

---

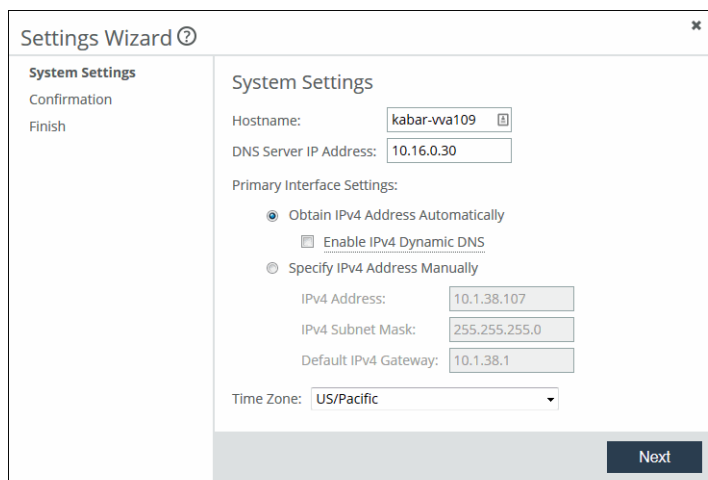
## Configuring Network Settings

In the Settings Wizard, you can quickly configure basic network and time zone settings.

### To set the initial network configuration

1. Choose Configure > Wizards: Initial Setup to display the Settings Wizard.

**Figure 2-1. Settings Wizard**



The screenshot shows the 'Settings Wizard' dialog box with a sidebar on the left containing 'System Settings', 'Confirmation', and 'Finish'. The main area is titled 'System Settings' and contains the following fields and options:

- Hostname:
- DNS Server IP Address:
- Primary Interface Settings:
  - ☒ Obtain IPv4 Address Automatically
  - ☐ Enable IPv4 Dynamic DNS
  - ☐ Specify IPv4 Address Manually
    - IPv4 Address:
    - IPv4 Subnet Mask:
    - Default IPv4 Gateway:
- Time Zone:

A 'Next' button is located at the bottom right of the dialog box.

2. Configure the system settings using the controls described in this table.

Control	Description
Hostname	Specify the IP address or hostname of the Edge.
DNS Server IP Address	Specify the IP address for the primary name server.
Primary Interface Settings	<p>Select and configure one of the following options:</p> <ul style="list-style-type: none"> <li>• <b>Obtain IPv4 Address Automatically</b> - Specify this option to automatically obtain the IPv4 address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it. <ul style="list-style-type: none"> <li>– <b>Enable IPv4 Dynamic DNS</b> - Select to send the Core's hostname with the DHCP request.</li> </ul> </li> </ul> <p><b>Note:</b> The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.</p> <ul style="list-style-type: none"> <li>• <b>Specify IPv4 Address Manually</b> - Specify this option if you do not use a DHCP server to set the IP address. Specify the following settings: <ul style="list-style-type: none"> <li>– <b>IPv4 Address</b> - Specify an IPv4 address.</li> <li>– <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> <li>– <b>Default IPv4 Gateway</b> - Specify the primary gateway IP address. The primary gateway must be in the same network as the primary interface. You must set the primary gateway for in-path configurations.</li> </ul> </li> </ul>
Time Zone	Select the time zone from the drop-down list. The default is US/Pacific.

3. Click **Next** to save the new configuration.

The Confirmation page displays. This page summarizes the settings you configured in the preceding pages.

4. Click **Save and Apply**.

The Finish page displays.

5. Click **Finish** to exit the Settings Wizard.

## Mapping LUNs to an Edge

The LUN Mapping Wizard enables you to quickly add LUNs, Edges, and iSCSI portals and targets to the Core configuration.

## To map LUNs to Edges

1. Choose **Configure > Wizards: LUN Mapping** to display the LUN Mapping Wizard.

**Figure 2-2. LUN Mapping Wizard**

2. In the iSCSI Initiator Name field, specify the initiator name to use when accessing your iSCSI storage device.
3. Click **Next** to open the Start page of the LUN Mapping Wizard.  
This page lists prerequisites for mapping LUNs, including the specific initiator that must be added to the iSCSI Portal Initiator list.
4. Click **Next** to display the Specify Portal page.
5. Specify the iSCSI portal using the controls described in this table.

Control	Description
Select from known Portals	<ul style="list-style-type: none"> <li>• <b>Hostname or IP address</b> - Select a previously configured iSCSI portal. Devices that appear in this list are independently configured in the <b>Configure &gt; Storage Array: iSCSI, Initiators, MPIO</b> page.</li> <li>• <b>Port</b> - Specify the port number of the iSCSI portal. The default is 3260.</li> <li>• <b>Authentication</b> - Select an authentication method (<b>None</b> or <b>CHAP</b>) from the drop-down list. If you select CHAP, the CHAP User field displays where you specify (or create) the CHAP username.</li> </ul>
Add new Portal	Specify the IP address and port for the intended iSCSI portal.

6. Click **Next** to display the Manage Targets page.
7. Select and add targets from the Discovered Targets List, which displays the targets discovered on the portal selected in the previous screen.

8. Click **Next** to display the Mount LUNs page.
9. Select and add LUNs from the Discovered LUNs List, which displays the LUNs discovered on the targets selected in the previous page.
10. Click **Next** to open the Specify SteelFusion Edge page.

---

**Note:** The term “SteelFusion Edge” refers to both Edge and xx60 model appliances in dedicated SteelFusion target mode.

---

11. Complete the Edge configuration using the controls described in this table.

Control	Description
Select from known SteelFusion Edges	From the Hostname drop-down list, select a previously configured SteelFusion Edge. Devices that appear in this list are independently configured in the Configure > Manage: SteelFusion Edges page. For details, see <a href="#">“Configuring Edges” on page 45</a> .
Add new SteelFusion Edge	Specify the hostname for the intended Edge.

12. Click **Next** to display the Map LUNs to Edge page.
13. Select and map LUNs from the top panel, which displays the unmapped LUNs, to the Edge selected in one of the preceding screens.
14. Click **Next** to display the Summary page.
15. Click **Exit** to complete the LUN mapping procedure.

## Importing a Saved Configuration

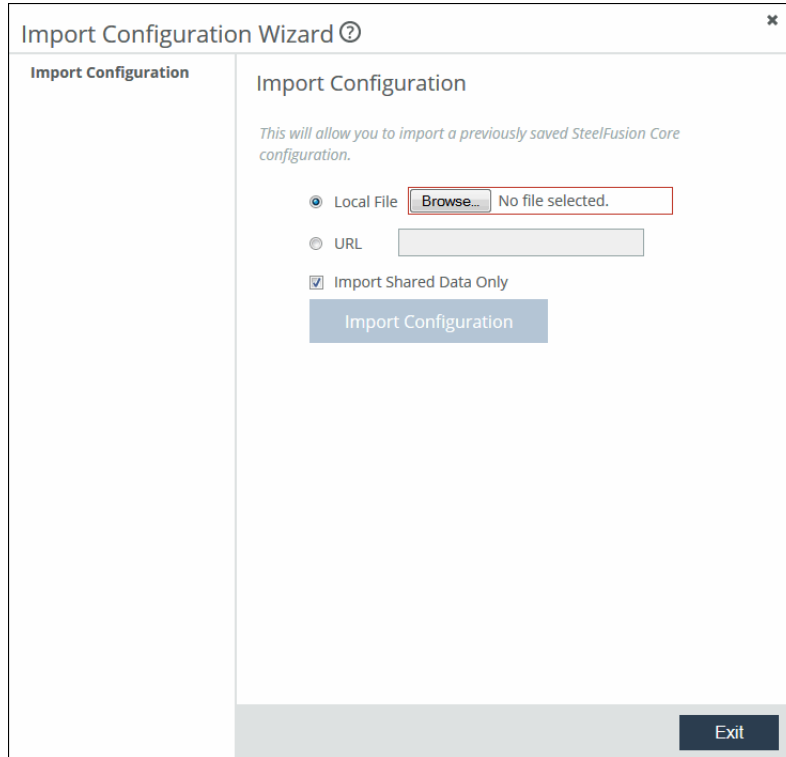
The Import Configuration Wizard enables you to quickly import a previously saved Core configuration from either a local file or URL.

### To import a saved configuration

1. Choose Settings > Maintenance: Service to display the Service page.
2. Click **Stop** to stop the current SteelFusion Core service.

3. Choose **Configure > Wizards: Import** to display the Import Configuration Wizard.

**Figure 2-3. Import Configuration Wizard**



4. Complete the configuration import using the controls described in this table.

Control	Description
Local File	Specify this option to import the configuration file from a local source. Click <b>Browse</b> to navigate to the file.
URL	Specify this option to import the configuration file from a URL. Specify the URL in the field provided.
Import Shared Data Only	Select this option to import only the following common settings: interface, protocols, CLI, web, statistics, NTP, SNMP, cloud settings, and alarm settings. The system does not automatically copy the other settings.
Import Configuration	Imports the configuration from the specified source.

5. Click **Exit**.
6. Choose **Settings > Maintenance: Service** to display the Service page.
7. Click **Start** to start the service using the imported configuration.

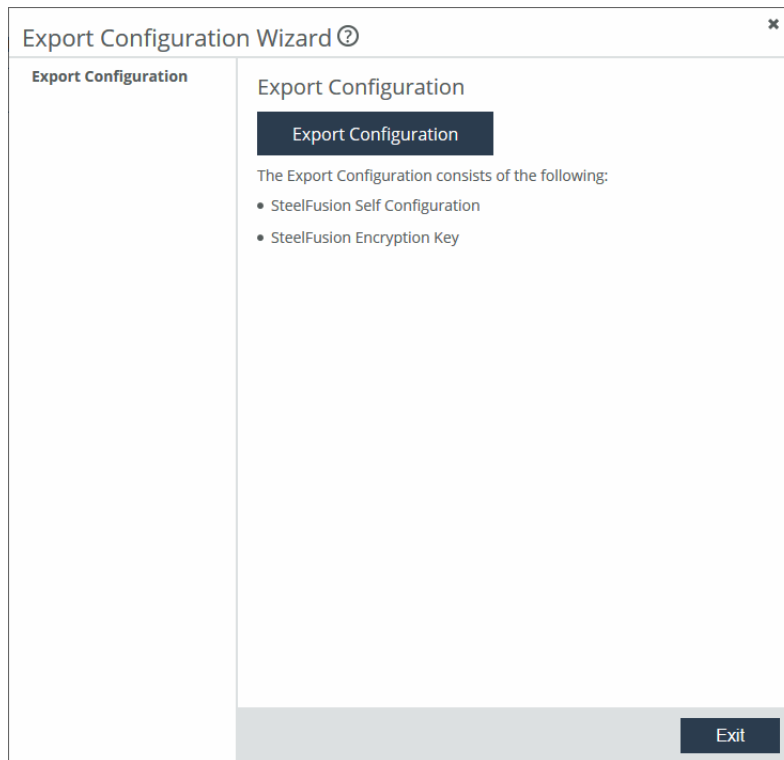
## Exporting a Configuration

The Export Configuration Wizard enables you to quickly export the current configuration. The export file consists of the SteelFusion self-configuration and encryption key.

## To export a configuration

1. Choose Configure > Wizards: Export to display the Export Configuration Wizard.

**Figure 2-4. Export Configuration Wizard**



2. Click **Export Configuration** to download or save the configuration as a .tgz file.

---

## Configuring iSCSI Settings

You can view and configure the iSCSI initiator, local interfaces for MPIO, portals, and targets in the iSCSI, Initiators, MPIO page.

The iSCSI initiator settings configure how the Core communicates with one or more storage arrays (targets) through the specified portal configuration.

After configuring the iSCSI portal, you can open the portal configuration to configure targets.

## To configure the iSCSI initiator

1. Choose Configure > Storage Array: iSCSI, Initiators, MPIO to display the iSCSI, Initiators, MPIO page.

**Figure 2-5. iSCSI, Initiators, MPIO Page**

**iSCSI, Initiators, MPIO** Storage Array > iSCSI, Initiators, MPIO ? Save

### iSCSI Initiator Configuration

Initiator Name:

- ☐ Enable Header Digest
- ☐ Enable Data Digest
- ☐ Enable Mutual CHAP Authentication
- ☒ Enable MPIO (Multi-Path I/O)
- ☐ Enable Standard Routing for MPIO

**Apply**

### Local Interfaces for MPIO (Multi-Path I/O)

MPIO Interfaces: primary - 10.5.128.48   
eth0\_0 - 10.5.132.192   
eth0\_1 - 10.5.140.9   
eth0\_2 - 10.5.142.173   
eth0\_3 - 10.5.142.189   
[Add interface](#)

### iSCSI Portal Configuration

**+ Add an iSCSI Portal**

Portal	Port	Status	Remove
▶ 10.5.132.233	3260	Connected	

2. Under iSCSI Initiator Configuration, configure authentication using the controls described in this table.

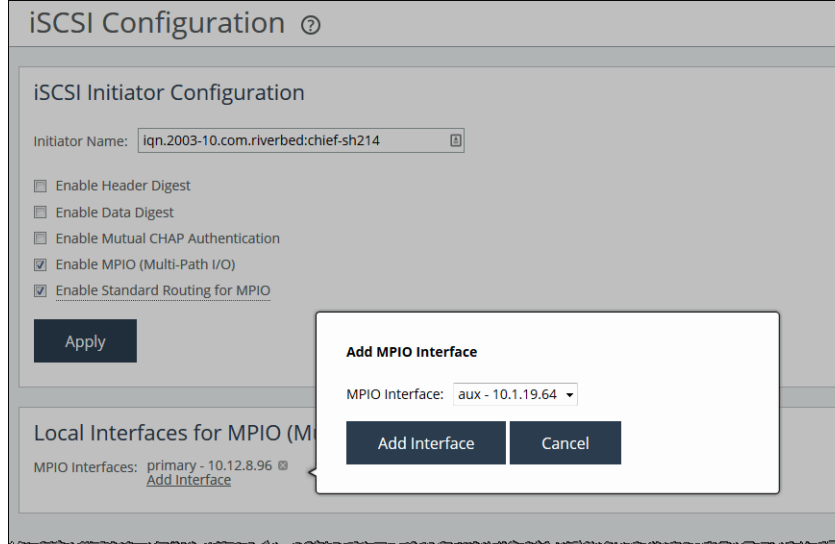
Control	Description
Initiator Name	Specify the name of the initiator to be configured.
Enable Header Digest	Includes the header digest data from the iSCSI PDU.
Enable Data Digest	Includes the data digest data from the iSCSI PDU.
Enable Mutual CHAP Authentication	<p>Enables bidirectional Challenge-Handshake Authentication Protocol (CHAP) authentication. This option adds yet another level of security where the target and the initiator authenticate each other. A separate secret is set for each target and each initiator in the storage array.</p> <p>CHAP provides protection against playback attack by the peer through the use of an incrementally changing identifier and a variable challenge value.</p> <p>For details about setting up mutual CHAP, see <a href="#">"To configure mutual CHAP" on page 64</a>.</p>
Enable MPIO (Multi-Path I/O)	<p>Enables the Multi-Path I/O (MPIO) feature.</p> <p>MPIO enables you to connect the appliance to the network and to the storage system through multiple physical I/O interfaces, so the data connection can continue in the event of a cable, switch, or other physical failure.</p>
Enable Standard Routing for MPIO	Enables the connection to be established through standard routing if the iSCSI portal is not in the same subnet as the MPIO interfaces.
Apply	Applies the changes to the running configuration.

### To add local interfaces for MPIO

1. Choose Configure > Storage Array: iSCSI, Initiators, MPIO to display the iSCSI, Initiators, MPIO page.
2. Under Local Interfaces for MPIO (Multi-Path I/O), click **Add Interface**.
3. Select an available interface from the drop-down list.
4. Click **Add Interface**.

To remove an MPIO interface configuration, click the X icon next to the interface name.

**Figure 2-6. iSCSI, Initiators, MPIO Page - Local Interfaces for MPIO (Multi-Path I/O)**

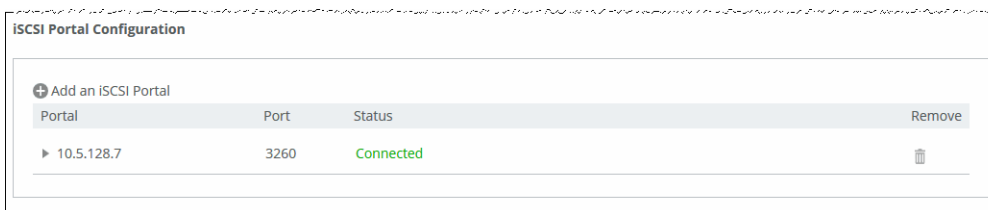


**Note:** When an MPIO interface is removed from the Edge, the existing initiator connections over that interface are not disconnected. You still see active data paths for a LUN over the old interfaces. This prevents you from accidentally losing access to the target if the new interfaces are not working or there are path issues. When the Edge target MPIO interfaces change, ESXi will continue showing old portals and old MPIO paths. Therefore, you will need to remove old portals manually and doing this will disconnect the old MPIO paths for the LUNs.

## To configure an iSCSI portal

1. Choose Configure > Storage Array: iSCSI, Initiators, MPIO to display the iSCSI, Initiators, MPIO page.

**Figure 2-7. iSCSI, Initiators, MPIO Page - iSCSI Portal Configuration**



2. Under iSCSI Portal Configuration, add or modify iSCSI portal configurations using the controls described in this table.

Control	Description
Add an iSCSI Portal	Displays controls for configuring and adding a new iSCSI portal.
Hostname or IP Address	Specify the IP address of the iSCSI portal.
Port	Specify the port number of the iSCSI portal. The default is 3260.
Authentication	Select an authentication method ( <b>None</b> or <b>CHAP</b> ) from the drop-down list. <b>Note:</b> If you select CHAP, the CHAP User field displays where you specify (or create) the CHAP username.

Control	Description
Add iSCSI Portal	Adds the defined iSCSI portal to the running configuration.

3. To view or modify portal settings, click the portal IP address in the portals list to access the following set of controls.

Control	Description
Status/Settings	<p>Configure or confirm the following settings:</p> <ul style="list-style-type: none"> <li>• <b>Portal Status</b> - Indicates whether the portal is connected.</li> <li>• <b>MPIO Portals</b> - Click <b>Rescan for Portals</b> to discover additional portals available on the storage system. You can then add them for redundancy, to prevent disconnection.</li> <li>• <b>Port</b> - Specify the port setting for the selected iSCSI portal.</li> <li>• <b>Authentication</b> - Specify either <b>None</b> or <b>CHAP</b> from the drop-down list.</li> <li>• <b>Update iSCSI Portal</b> - Updates the portal settings configuration.</li> </ul>
Offline LUNs	<p>Click <b>Offline LUNs</b> to take all LUNs offline that are serviced by this selected iSCSI portal.</p> <p><b>Note:</b> The process of taking a LUN offline requires you first to power off the Windows server at the affected branch and then to unmount the LUN from ESX (if necessary). For details, see the <i>SteelFusion Design Guide</i>.</p>

4. To add a target to the newly configured portal:
- Click the portal IP address in the list to expand the set of controls.
  - Under Targets, add a target for the portal using the controls described in this table.

Control	Description
Add a Target	Displays controls for adding a target.
Target Name	<p>Enter the target name or choose from available targets.</p> <p><b>Note:</b> You can also rescan for available targets using this field.</p>
Port	Specify the port number of the target.
Add Target	Adds the newly defined target to the current iSCSI portal configuration.

5. To modify an existing a target configuration:

- Click the portal IP address in the list to expand the set of controls.
- Under Targets, click the target name in the Targets table to expand the target settings using the controls described in this table.

Control	Description
Status/Settings	Open this tab to modify the port and snapshot configuration settings.
Offline LUNs	<p>Open this tab to access the <b>Offline LUNs</b> button.</p> <p>Clicking this button takes all configured LUNs offline that are serviced by the current target.</p> <p><b>Note:</b> The process of taking a LUN offline requires you first to power off the Windows server at the affected branch and then to unmount the LUN from ESX (if necessary). For details, see the <i>SteelFusion Design Guide</i>.</p>

## Configuring LUNs

You configure block disk (Fibre Channel), Edge Local, and iSCSI logical unit numbers (LUNs) in the LUNs page.

Typically, block disk and iSCSI LUNs are used to store production data. They share the space in the blockstore cache of the associated Edges, and the data is continuously replicated and kept synchronized with the associated LUN in the data center. The Edge blockstore caches only the working set of data blocks for these LUNs; additional data is retrieved from the data center when needed.

Edge Local LUNs are used to store transient and temporary data. Local LUNs also use dedicated space in the blockstore cache of the associated Edges, but the data is not replicated to the data center LUNs.

**Note:** Prefetch does not optimize access for VMs that contain any SE SPARSE (Space Efficient Sparse) format snapshots. These snapshots are created for virtual machine disks (VMDKs) greater than 2 TB in size.

## Fibre Channel LUN Configuration Steps in ESXi

To configure storage arrays using Fibre Channel, Core-v leverages raw device mapping (RDM) functionality through VMware ESXi. To configure a Fibre Channel LUN, you must do the following tasks in the ESXi server where Core-v is deployed:

- Expose the Fibre Channel LUNs to the ESXi server.
- Assign the discovered Fibre Channel LUNs as raw device mappings (RDM).
- Ensure that NPIV is disabled.

You can then discover and configure the Fibre Channel LUN as a block-disk LUN in the SteelFusion Core Management Console. For details, see [“To configure a LUN” on page 38](#). After configuring the block-disk LUN in the SteelFusion Core Management Console, you must return to the ESXi server to configure high-availability for the new LUN:

- In ESXi, modify the storage group to which the LUN belongs to expose the LUN to both ESXi servers running the Core failover peers.

- Ensure that the LUN is raw-data mapped (RDM) to both Core-v appliances.
- Ensure that the SCSI bus number of the RDM LUN is the same on both ESXi servers.

For details about all preconfiguration and postconfiguration steps for Fibre Channel LUNs, see the *Fibre Channel on SteelFusion Core Virtual Edition Solution Guide*.

## Block-Disk LUN Considerations

Block-disk LUNs for Fibre Channel are distinct from iSCSI LUNs in several important ways:

- **No MPIO configuration** - There is no MPIO configuration for block-disk LUNs. Multipathing is instead performed by the ESXi system.
- **No Volume Snapshot Service (VSS) snapshot support** - While all other snapshot features are supported, Windows VSS is not.
- **SCSI reservations** - Not taken on block-disk LUNs.
- **Additional HA configuration required** - Configuring high-availability for Core-v failover peers requires that each appliance be deployed on a separate ESXi system. Similarly, block-disk LUNs must be exposed to both systems. For details, see [“Fibre Channel LUN Configuration Steps in ESXi” on page 35](#).
- **Maximum of sixty Fibre Channel LUNs per ESXi system** - ESXi allows a maximum of four SCSI controllers, each of which can support up to 15 SCSI devices.

---

**Note:** Block-disk LUN configuration pertains to Fibre Channel support. Fibre Channel is supported only in Core-v deployments. For details, see the *SteelFusion Design Guide*. For important distinctions between iSCSI and block-disk LUNs, see [“Block-Disk LUN Considerations” on page 36](#).

---

### To add a new LUN

1. Choose Configure > Manage: LUNs to display the LUNs page.

**Figure 2-8. LUNs Page**

LUNs Manage > LUNs ?								
Status of all configured LUNs.								
+ Add a LUN <span style="float: right;">Filter Table: <input type="text" value="Enter search text"/></span>								
LUN	Type	Status	Size	SteelFusion Edge	Portal	Pinned	Prepop	Remove
▶ lun1 (hoiqdp3AkafV)	iSCSI	N/A	503.00 MB	None	10.1.13.190	No	Disabled	
▶ lun2 (hoiqdp3Ajs2V)	iSCSI	N/A	502.03 MB	None	None Connected	No	Disabled	
▶ lun3 (P3KRP502GuZw)	iSCSI	N/A	503.00 MB	edge1	10.1.13.190	No	Disabled	

2. Click **Add a LUN** and configure it using the controls described in this table.

Control	Description
Add a LUN	Displays controls for adding a LUN to the current configuration.

Control	Description
iSCSI tab	<p>Configure the iSCSI LUN with the following settings:</p> <ul style="list-style-type: none"> <li>• <b>LUN Serial Number</b> - Select from the drop-down list of discovered LUNs. The LUNs listed are shown using the following format: &lt;serial-number&gt;(&lt;portal&gt;/&lt;target&gt;). If the desired LUN does not appear in the drop-down list, click <b>Rescan background storage for new LUNs</b>.</li> <li>• <b>LUN Alias</b> - Specify an alias for the LUN.</li> <li>• <b>Add iSCSI LUN</b> - Adds the new LUN to the running configuration.</li> </ul>
Block Disk tab	<p>Configure the block-disk LUN with the following settings:</p> <ul style="list-style-type: none"> <li>• <b>LUN Serial Number</b> - Select from the drop-down list of discovered LUNs. The LUNs are listed by their serial numbers. If the desired LUN does not appear in the drop-down list, click <b>Rescan background storage for new LUNs</b>.</li> <li>• <b>LUN Alias</b> - Specify an alias for the LUN.</li> <li>• <b>Add Block Disk LUN</b> - Adds the new LUN to the running configuration.</li> </ul> <p>Additional steps are required before and after you configure a block-disk (Fibre Channel) LUN in the Core. For details, see <a href="#">“Configuring Branch Recovery” on page 42</a>.</p>
Edge Local tab	<p>Configure the Edge Local LUN with the following settings:</p> <ul style="list-style-type: none"> <li>• <b>SteelFusion Edge</b> - Select a LUN from the drop-down list. This list displays configured Edges. For details, see <a href="#">“Configuring Edges” on page 45</a>.</li> <li>• <b>Size</b> - Specify the LUN size, in megabytes, gigabytes, or terabytes, as specified by the drop-down list.</li> <li>• <b>Alias</b> - Specify the alias for the LUN.</li> <li>• <b>Add Local LUN</b> - Adds the new LUN to the running configuration.</li> </ul>

To configure a LUN

- 1. Click the LUN you want to configure.

Figure 2-9. LUN Configuration Tabs

LUN	Type	Status	Size
▼ lun_0 (oak-cs7_200)	iSCSI	Connected (no snapshot settings)	2 GB

DetailsEdge MappingFailoverMPIOSnapshotsPin/PrepopBranch Recovery

LUN is OnlineOfflineOnline

LUN Alias: lun\_0Update Alias

Connection Status: Ready

Accessibility: Edge Initiators configuredEdit Edge

Snapshot Status: This LUN is not configured for taking Snapshots

Locally Assigned LUN Serial: oak-cs7\_200

Origin LUN Serial: oak-cs7\_200

Origin Portal: 10.5.128.7

Origin Target: iqn.2003-10.com.riverbed:default-target

Size: 2 GB

LUN I/O Report: Click here

## 2. Configure the LUN using the controls described in this table.

Control	Description
Details	<p>Displays online or offline status:</p> <ul style="list-style-type: none"> <li>Click <b>Online</b> to bring the LUN online.</li> <li>Click <b>Offline</b> to take the LUN offline.</li> </ul> <p><b>Note:</b> The process of taking a LUN offline requires you first to power off the Windows server at the affected branch and then to unmount the LUN from ESX (if necessary). For details, see the <i>SteelFusion Design Guide</i>.</p> <hr/> <p>Displays the LUN alias, if applicable. Optionally, you can modify the value and click <b>Update Alias</b>.</p> <hr/> <p>Displays the following information about the LUN:</p> <ul style="list-style-type: none"> <li>Connection Status</li> <li>Accessibility (by iSCSI initiators at the Edge)</li> <li>Snapshot Status</li> <li>Locally Assigned LUN Serial</li> <li>Origin LUN Serial</li> <li>Origin Portal (Windows LUNs only)</li> <li>Origin Target (Windows LUNs only)</li> <li>Size</li> <li>LUN I/O Report (links to report page)</li> </ul>
Edge Mapping	<p>Displays the target name and Edge to which the LUN is mapped.</p> <ul style="list-style-type: none"> <li>To unmap, click <b>Unmap</b>.</li> </ul> <p><b>Note:</b> You must take the LUN offline before you can unmap it.</p>
Failover	<p>Displays whether the LUN is configured for failover.</p> <ul style="list-style-type: none"> <li>To enable or disable failover, click <b>Enable</b> or <b>Disable</b>.</li> </ul> <p><b>Note:</b> LUN failover is in effect only when failover is also configured for the current Core. For details, see <a href="#">“Configuring Failover” on page 80</a>.</p>

Control	Description
MPIO (iSCSI LUNs only)	<p data-bbox="492 254 1130 281">Applies the MPIO policy you select from the drop-down list:</p> <ul data-bbox="492 294 1414 441" style="list-style-type: none"> <li data-bbox="492 294 1414 373">• <b>Round-Robin</b> - The Core uses multiple paths for read I/O on the LUN in a round-robin pattern: for example, 100 reads on path 1 followed by 100 reads on path 2, and so on. I/O writes still always use a single path.</li> <li data-bbox="492 386 1414 441">• <b>Fixed-Path</b> - The Core uses a single path for I/O on the LUN. This path is chosen from all the connected paths based on the priorities and preferences detailed above.</li> </ul> <p data-bbox="492 453 1393 508">For the selected MPIO policy, the MPIO path table displays the following information about the iSCSI path to the backend:</p> <ul data-bbox="492 520 1406 970" style="list-style-type: none"> <li data-bbox="492 520 1086 548">• Source Interface: The IP address and port on the Core.</li> <li data-bbox="492 560 1318 588">• Destination Interface: The IP address and port on the backend storage array.</li> <li data-bbox="492 600 1406 655">• For the Fixed Path policy, the Path Included check box: Select to include this path in the list of paths for I/O.</li> <li data-bbox="492 667 1318 722">• For the Fixed Path policy, the User Preferred button: Click to make this path preferred. Preferred paths receive higher priority.</li> <li data-bbox="492 735 1344 789">• Array Preferred: Whether the array indicates that the I/O initiator should prefer this path when available. If so, the Core gives this path higher priority.</li> <li data-bbox="492 802 1235 829">• Connected: network status of this path (connected or disconnected).</li> <li data-bbox="492 842 1370 970">• Status: MPIO status of this path, as reported by the array. The status can be:               <ul data-bbox="516 877 1370 970" style="list-style-type: none"> <li data-bbox="516 877 1370 905">– Active Optimized: Active (I/O possible) and optimized (the best path for I/O).</li> <li data-bbox="516 917 1370 970">– Active Unoptimized: Active (I/O possible) but unoptimized (not the best path for IO).</li> </ul> </li> </ul>

Control	Description
Snapshots	<p>Displays the following sets of controls for snapshot configuration:</p> <ul style="list-style-type: none"> <li>• <b>History</b> - Displays a detailed list of snapshots taken of the LUN.</li> <li>• <b>Scheduler</b> - Use the controls in this tab to apply snapshot schedule policies to the current LUN. For details see <a href="#">“Applying a Snapshot Schedule Policy to a LUN” on page 74</a>.</li> <li>• <b>Configuration</b> - Use the controls in this tab to enable and configure application-specific snapshots and data protection options.</li> </ul> <p>Configure the following settings:</p> <ul style="list-style-type: none"> <li>– <b>Storage Array or Handoff Host</b> - For storage arrays from qualified vendors (Dell EqualLogic, EMC CLARiiON, EMC VNX, NetApp, or IBM v7000), specify Storage Array and select from the drop-down list the preconfigured array where the snapshot is stored. You can configure storage arrays for snapshots on the Configure &gt; Backups: Snapshots page. For details, see <a href="#">“Configuring Snapshots for Storage Arrays” on page 67</a>.</li> <li>– For nonqualified storage arrays, specify the Handoff Host option and select the host from the drop-down list. You can configure handoff hosts for snapshots on the Configure &gt; Backups: Snapshots page. For details, see <a href="#">“Configuring Handoff Hosts” on page 69</a>.</li> </ul> <p>After you specify the storage array or handoff host, click <b>Update Settings</b> to update the changes. You can also test the connection by clicking <b>Test Storage Array</b>.</p> <ul style="list-style-type: none"> <li>– <b>Snapshots Static Name</b> - Specify a string to be prepended to the names of snapshots taken of this LUN.</li> <li>– <b>Storage Group</b> - (EMC CLARiiON and IBM only) From the drop-down list, select the storage group/host for the proxy backup server.</li> </ul> <p>Click <b>Update Vendor-Specific Settings</b> to update the changes.</p> <ul style="list-style-type: none"> <li>– <b>Client Type</b> - Select <b>VMware</b>, <b>Windows</b>, or <b>Other</b>. The type you select determines the settings in the Application Consistent Snapshots and Proxy Backup panels.</li> </ul> <p>Click <b>Update Client Type</b> to update the changes.</p> <p>For details about application-consistent snapshots and proxy backup configuration, see <a href="#">“Configuring Application-Consistent Snapshots for a LUN” on page 75</a> and <a href="#">“Configuring Proxy Backup for a LUN” on page 77</a>.</p> <p><b>Note:</b> To configure application-consistent snapshots and proxy backup settings for the current LUN, you must specify either <b>VMware</b> or <b>Windows</b> for the client type.</p>
Pin/Prepop	<p>Displays the pin status (Pinned or Unpinned) and provides controls for changing the status.</p> <p>When a LUN is pinned, the data is reserved and not subject to the normal blockstore eviction policies. For details about pinning, see the <i>SteelFusion Design Guide</i>.</p> <p>This tab also contains controls for enabling or disabling the prepopulation service and for configuring a prepopulation schedule.</p> <p><b>Note:</b> You can create a prepopulation schedule only when the pin status is set and updated to Pinned.</p>

Control	Description
Branch Recovery	<p>The Core displays the following controls for configuring branch recovery for the selected LUN. For details, see <a href="#">“Configuring Branch Recovery” on page 42</a>.</p> <ul style="list-style-type: none"> <li>Click <b>Enable</b> to enable branch recovery, or click <b>Disable</b> to disable it.</li> <li>Click <b>Change Schedule</b> to set a date and time for the branch recovery operation. <ul style="list-style-type: none"> <li>From the scheduler, you can also select <b>Now</b> to start the operation immediately.</li> <li>Click <b>Submit Schedule</b> to finalize the schedule.</li> <li>Click <b>Hide Scheduler</b> to close the calendar controls.</li> </ul> </li> <li>Click <b>Add VM</b> to add one or more VMs to prepopulate for the operation. <ul style="list-style-type: none"> <li>From the drop-down list, select a VM, or select all.</li> <li>Set the VM capacity percentage. This percentage is the maximum data capacity prepopulated by branch recovery per virtual disk.</li> </ul> </li> <li>After you have added a VM, you can remove it by clicking the X icon to the right and confirming the removal.</li> </ul> <p><b>Note:</b> The VM functionality does not apply to NTFS LUNs. For these LUNs, simply enable the operation and schedule it.</p>

## Configuring Branch Recovery

The Branch Recovery Agent is a Windows service installed on VSP or external VMs or hosts, and enables prepopulation of the working set when LUNs on an Edge are being restored after an irrecoverable failure. After installation, it is always running, listing, and maintaining a list of accessed blocks in a file that composes the working set.

The Branch Recovery Agent is compatible with both Windows VMs and physical hosts. You can configure branch recovery on LUNs hosting the Virtual Machine File System (VMFS), for all or some VMs, or the New Technology File System (NTFS), for all drives on the LUN. The Windows VM comprises multiple VMDKs on a VMFS LUN. Each of the VMDKs is formatted as NTFS.

As of version 3.6, for each VM on a LUN where you have set up branch recovery, the Core cycles through each VM in a round-robin fashion, recovering the working set in phases starting with the most recently accessed blocks and working backwards. For example, if you have set up branch recovery on two VMs, the Core recovers the most recent day of data for VM1 first, and then for VM2. The Core then recovers a week of data on VM1, and then does the same for VM2. This process ensures that the VMs have the latest working set.

Branch recovery can also be used in conjunction with prefetch and prepopulation. During recovery, use branch recovery to quickly restore the working set. Later, use prepopulation or rely on prefetch to optimize cold files and folders.

The Branch Recovery Agent has been qualified on Windows 7, Windows 2008 Release 2, and Windows 2012 Release 2.

---

**Note:** The Branch Recovery Agent cannot distinguish between virtual disks on SteelFusion storage or local storage.

---

### To set up branch recovery

1. Install the Branch Recovery Agent on each host running Windows (virtual or physical) at the branch.

You can find the Riverbed Host Tools installer at <https://support.riverbed.com>.

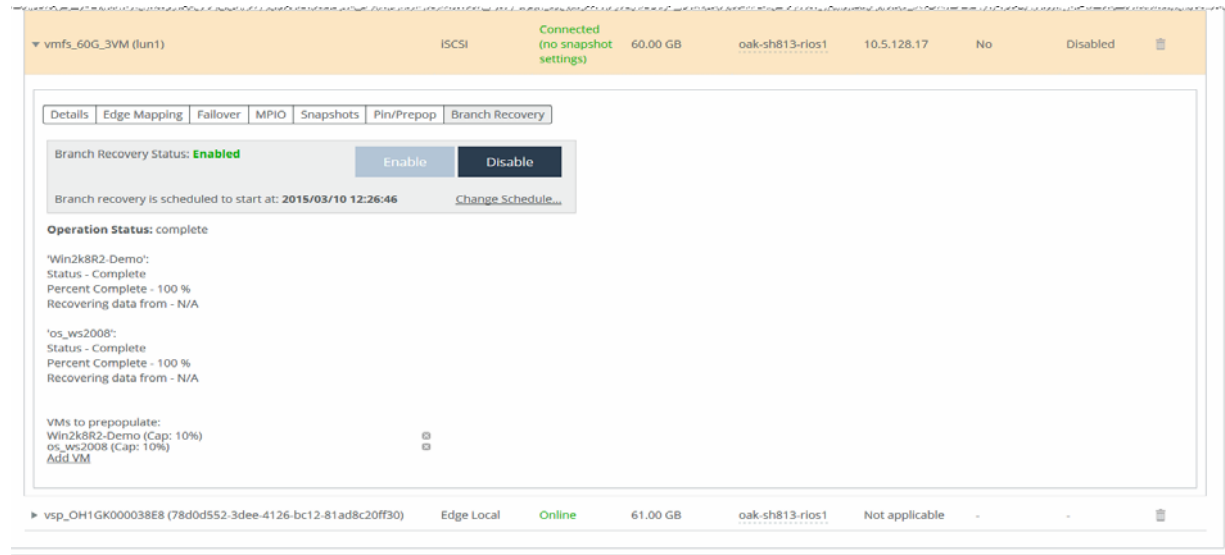
---

**Note:** You must run the agent installer as an administrator.

---

2. Choose Configure > Manage: LUNs to display the LUNs page.
3. Click the LUN you want to set up with branch recovery.
4. Select the Branch Recovery tab.

**Figure 2-10. LUNs Page - Branch Recovery**



5. Click **Add VM**.

Any number of VMs can be selected for the operation.

---

**Note:** The backend must be connected (and the prefetch stack instantiated on it) for VMs to be selected for recovery.

---

6. Select a VM to add to the branch recovery operation.
7. Specify a separate VM Cap for each VM.

The VM Cap is the maximum amount of data to be prepopulated by branch recovery per disk. To add more VMs at a later date, you must disable Branch Recovery, add the new VMs, and enable it again.

---

**Note:** Either all VMDKs of a VM are recovered, or none are. In particular, if a VM has one RDM disk, that VM is skipped.

---

8. Click **Add VM**.
9. Click **Change Schedule**.

The recovery operation can be scheduled immediately (Now) or for a future date. When the operation runs, progress is displayed.

10. Create a schedule for branch recovery on this LUN and click **Submit Schedule**.

11. Once you have finished setting up branch recovery on this LUN, click **Enable**.

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

---

**Note:** After Branch Recovery is set up and enabled, no settings can be changed until it is disabled again.

---

The Branch Recovery Agent uses the Event Tracing for Windows framework. This framework uses data supplied by the OS on all I/O (read and write) on each virtual disk. It writes all this data periodically to a file on the same disk.

For example, if a VM has two VMDKs, Disk0 (C:, 40 GB) and Disk1 (D:, 10 GB, E:, 20 GB),

- All disk-access history for Disk0 is written into C:\Riverbed\BranchRecovery\lru\*.log
- All disk-access history for Disk1 is written into E:\Riverbed\BranchRecovery\lru\*.log

All these files are transmitted to the LUN hosting the VMFS at the data center. After a failure, the VMFS LUN (or a snapshot) is mapped to a new Edge, and a prefetch technique crawls the LUN, parsing all the lru\*.log files it discovers and pushing the blocks recorded there.

## Troubleshooting Branch Recovery

The Branch Recovery Agent is incompatible with xperf or any other application that uses it, such as the Turbo Boot Agent. The agent fails to start while an xperf application is running, but it retries periodically until the other application exists.

Example messages:

```
[Mar 4 2014 16:20:06.315 9012 - /etw_controller ERR] The NT Kernel Logger session is already in use
[Mar 4 2014 16:20:06.315 9012 - /daemon NOTICE] Failed to start ETW Controller: sleeping for 60
seconds
```

This table lists useful directory locations on the Windows host:

Location Type	Path
Event logs	C:\ProgramData\Riverbed\BranchRecovery\log
Crash dumps	C:\ProgramData\Riverbed\BranchRecovery\debugging
Installation	C:\Program Files\Riverbed\BranchRecovery (also contains the uninstaller)

## Configuring Edges

You can configure and modify connectivity with Edges in the SteelFusion Edges page.

### To configure an Edge

1. Choose Configure > Manage: SteelFusion Edges to display the SteelFusion Edges page.

Figure 2-11. SteelFusion Edges Page

SteelFusion Edge	Connection	Duration	IP Address	Mapped LUNs	LUN Capacity	Remove
▶ oak-sh38	Connected	1h 1m 44s	10.5.128.201	2 LUNs	120.02 GB	🗑️
▶ oak-sh546	Connected	1h 1m 46s	10.5.140.197	4 LUNs	1.24 TB	🗑️

2. Configure the Edge using the controls described in this table.

Control	Description
Show/Hide Preferred Interfaces for SteelFusion Edge Connections	Click to show or hide the preferred interfaces for Edge connections. For details, see <a href="#">“To modify an existing Edge configuration” on page 46</a> .
Add a SteelFusion Edge	Displays controls for adding an Edge to the current configuration.
SteelFusion Edge Identifier	Specify the identifier for the Edge. This value must match the value configured on the Edge. <b>Note:</b> Edge identifiers are case sensitive.
Blockstore Encryption	Select one of the following encryption types from the drop-down list. The encryption types are listed from the least to the most secure. <ul style="list-style-type: none"> <li>• <b>No encryption</b> - Disables data encryption.</li> <li>• <b>Using AES_128 bit key</b> - Encrypts data using the AES cryptographic key length of 128 bits.</li> <li>• <b>Using AES_192 bit key</b> - Encrypts data using the AES cryptographic key length of 192 bits.</li> <li>• <b>Using AES_256 bit key</b> - Encrypts data using the AES cryptographic key length of 256 bits.</li> </ul>
Add SteelFusion Edge	Adds the new Edge to the running configuration. The newly added device appears in the list.

3. To remove an existing Edge configuration, click the trash icon in the Remove column.

## To modify an existing Edge configuration

1. Click the Edge name in the SteelFusion Edge table to expand a set of additional controls for configuring the selected Edge.

Control	Description
Status	<p>This panel displays the following information about the selected Edge configuration:</p> <ul style="list-style-type: none"> <li>• <b>IP Address</b> - The IP address of the selected Edge.</li> <li>• <b>Connection Status</b> - Indicates whether the selected Edge is connected to the Core.</li> <li>• <b>Connection Duration</b> - Duration of the current connection.</li> <li>• <b>Total LUN Capacity</b> - The total storage capacity of the LUN dedicated to the selected Edge.</li> <li>• <b>Blockstore Encryption</b> - Shows the type of encryption of the blockstore data (if any).</li> <li>• <b>Read IO Latency</b> - Shows the average read I/O time (in milliseconds) for the specified Edge.</li> <li>• <b>Write IO Latency</b> - Shows the average write I/O time (in milliseconds) for the specified Edge.</li> <li>• <b>SteelFusion Edge Data I/O report</b> - The graph is displayed in reduced scale. Click <b>View Full Report</b> to go to the SteelFusion Edge Data I/O report page. For details, see <a href="#">“Viewing the SteelFusion Edge Trends” on page 153</a>.</li> <li>• <b>Blockstore Metrics</b> <ul style="list-style-type: none"> <li>• <b>Commit Delay</b> - Estimated time (in seconds) to drain the uncommitted blockstore data back to Core.</li> <li>• <b>Space Utilization</b> - Percentage of the blockstore is currently being used by uncommitted data.</li> <li>• <b>Time To Commit</b> - Delay (in seconds) to commit the blockstore data.</li> <li>• <b>Uncommitted Bytes</b> - Percentage of the blockstore is currently being used by uncommitted data.</li> </ul> </li> </ul> <p>For details about blockstore metrics, see <a href="#">“Viewing Storage Reports” on page 149</a>.</p>
Target Settings	<p>This panel displays the following controls for configuring the target settings:</p> <ul style="list-style-type: none"> <li>• <b>Target Name</b> - Displays the system name of the selected Edge.</li> <li>• <b>Require Secured Initiator Authentication</b> - Requires CHAP authorization when the selected Edge is connecting to initiators.</li> <li>• <b>Enable Header Digest</b> - Includes the header digest data from the iSCSI PDU.</li> <li>• <b>Enable Data Digest</b> - Includes the data digest data from the iSCSI PDU.</li> <li>• <b>Update Target</b> - Applies the changes you make to these settings.</li> </ul> <p><b>Note:</b> If the <b>Require Secured Initiator Authentication</b> setting is selected, you must set authentication to CHAP in the adjacent Initiators tab.</p>

Control	Description
Initiators	<p>Displays controls for adding and managing initiator configurations:</p> <ul style="list-style-type: none"> <li>• <b>Initiator Name</b> - Specify the name of the initiator to be configured. As of version 4.3, you can click <b>Add Discovered Initiator</b> to display a list of initiator names connected to the Edge in a drop-down list to choose from instead of having to cut and paste the iSCSI Qualified Name (IQN) from ESXi. As long as an initiator is connected to an Edge, it appears in this list.</li> <li>• <b>Add to Initiator Group</b> - Select a group name from the drop-down list or click <b>New Group</b> to add a new group.</li> <li>• <b>Authentication</b> - Select the authentication method from the drop-down list: <ul style="list-style-type: none"> <li>• <b>None</b> - No authentication required.</li> <li>• <b>CHAP</b> - Only the target authenticates the initiator. The secret is set just for the target; all initiators that want to access that target must use the same secret to begin a session with the target.</li> <li>• <b>Mutual CHAP</b> - The target and the initiator authenticate each other. A separate secret is set for each target and for each initiator in the storage array.</li> </ul> </li> </ul> <p>For details about CHAP and mutual CHAP, see <a href="#">“To configure one-way CHAP” on page 63</a>.</p> <p><b>Note:</b> If the <b>Require Secured Initiator Authentication</b> setting is selected for the Edge in the Target Settings tab, authentication must be configured for a CHAP option.</p> <ul style="list-style-type: none"> <li>• <b>Add Initiator</b> - Adds the new initiator to the running configuration.</li> </ul>
Initiator Groups	<p>Displays controls for adding and managing initiator group configurations:</p> <ul style="list-style-type: none"> <li>• <b>Group Name</b> - Specifies a name for the group.</li> <li>• <b>Add Group</b> - Adds the new group. The group name displays in the Initiator Group list.</li> </ul> <p>After this initial configuration, click the new group name in the list to display additional controls.</p>
LUNs	<p>Displays LUNs to the selected Edge.</p> <p>To map LUNs to the selected Edge, choose Configure &gt; Manage: LUNs.</p> <p>To manage group and initiator access, click the name of the LUN to access additional controls.</p>
Prepopulation	<p>Displays controls for configuring prepopulation tasks.</p> <p><b>Note:</b> This prepopulation schedule is applied to all virtual LUNs mapped to the currently selected Edge if no LUN-specific schedule is configured.</p> <ul style="list-style-type: none"> <li>• <b>Schedule Name</b> - Specify a task name.</li> <li>• <b>Start Time</b> - Select the start day and time from the respective drop-down lists.</li> <li>• <b>Stop Time</b> - Select the stop day and time from the respective drop-down lists.</li> <li>• <b>Add Prepopulation Schedule</b> - Adds the task to the Task list.</li> </ul> <p><b>Note:</b> To delete an existing task, click the trash icon in the Task list.</p>

## Configuring Replication

This section describes how to configure replication on the Core. As of version 4.3, you can configure physical or virtual Cores for seamless failover and recovery between data centers without any data loss. In this environment, your primary and secondary data centers are always synchronized in case of large scale failures such as power loss, natural disasters, or hardware failure. You connect two separate Cores that are each connected to their own storage array in two separate data centers.

The primary data center receives all the reads and writes from the Edge, and is connected to the secondary data center. The secondary data center hosts the replica LUNs, which are copies of the primary LUNs located at the primary data center.

Replication can interoperate with all existing physical or virtual Cores, and all existing Edges, as long as the software version is compatible (version 4.0 and later).

This section describes the following topics:

- [“Base Requirements” on page 48](#)
- [“Before You Begin” on page 48](#)
- [“Basic Steps” on page 49](#)
- [“Setting Up the Data Centers for Replication” on page 49](#)
- [“Pairing the Cores” on page 52](#)
- [“Configuring the Witness” on page 53](#)
- [“Configuring Edges and LUNs for Replication” on page 55](#)
- [“Suspending Replication” on page 57](#)
- [“Initiating Failover \(Recovering from Primary Data Center Failure\)” on page 58](#)
- [“Failing Back to the Primary Data Center” on page 59](#)
- [“Terminating Replication” on page 61](#)

## Base Requirements

You must meet the following requirements to set up replication:

- The backend storage array must be configured for each Core that will be included in the replication configuration.
- The primary data center should be able to reach the secondary data center through the chosen interfaces.
- The secondary Core cannot have any Edges or LUNs.
- Each Core must have the same configuration. For example, if high-availability is configured on the primary data center’s Cores, it must also be configured on the secondary Cores.
- The Edges should be able to reach the secondary data center.

## Before You Begin

1. Create discoverable LUNs on storage arrays in both data centers. The LUNs you create on the primary data center will serve the Edges during normal operation. The LUNs you create on the secondary data center are replicas of the primary data center LUNs and should not be added to the Core. The secondary data center LUNs should be exactly the same size as those in the primary data center so they can serve the Edges in the event of a primary data center failure.

Storage array vendors and models can vary between data centers, as well as LUN size. For this reason, Riverbed permits a size leeway of 1 GB to allow the LUN created on the secondary data center storage array to be at most 1 GB larger (but not smaller) than the primary data center LUN. For example, a 4 GB LUN on the primary data center may have a 5 GB replica LUN on the secondary data center. The larger part of the LUN in secondary data center is not used.

2. Create the Journal LUN you will be using on the backend. The Journal LUN is a dedicated LUN on the backend storage that is used to temporarily journal writes for replicated LUNs when replication is suspended. If you have Cores set up for high-availability, share the same Journal LUN between the Cores in case one of the Cores fails.

The sizing of the Journal LUN is dependent on the number of LUNs you want to replicate. Riverbed recommends a thinly provisioned LUN of 500 GB or more. If the Journal LUN is not large enough, an alarm will be triggered with a reminder of the minimum required size (calculated based on the currently configured replica LUNs).

3. Ensure that the Journal LUN is accessible by the Core, and that the Journal LUN is not added to the configuration.
4. If you plan to configure Cores for high-availability, ensure that they are set up before you start configuring replication.

## Basic Steps

This table describes the basic steps needed to set up replication across data centers, followed by detailed procedures.

Task	Reference
1. Configure replication between the two data centers (roles, replication interfaces, and Journal LUN).	<a href="#">“Setting Up the Data Centers for Replication” on page 49</a> and <a href="#">“Pairing the Cores” on page 52</a>
2. Configure the Witness.	<a href="#">“Configuring the Witness” on page 53</a>
3. Configure the Edges and LUNs for replication and start first sync.	<a href="#">“Configuring Edges and LUNs for Replication” on page 55</a>

## Setting Up the Data Centers for Replication

First, set up the primary data center, where the Edges and LUNs will reside. The secondary data center is on standby and should not have any configured LUNs or Edges, and it is the data center that the primary will be replicating to.

### To set up the primary data center for replication

1. Log in to the management console on the primary Core.

- Choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

**Figure 2-12. Replication (Set Up / Monitor) Page - Replication Configuration Panel**

Replication (Set Up / Monitor) Replication > Replication (Set Up / Monitor) ?

Replication Configuration

Role: ☒ Primary ☐ Secondary

Data Center Name:

Replication Interface: primary - 10.5.62.105 ▼

Journal LUN: lun\_9G (9.00 GB) ▼  
*(journal LUN must be larger than 3.00 GB. However, a thinly-provisioned LUN of 500 GB or more is recommended.)*

**Set Configuration**

- Select Primary in the Replication Configuration panel.
- Complete the configuration using the controls described in this table.

Setting	Description
Data Center Name	Specify a unique identifier for the primary data center.
Replication Interface	Specify the replication interface (for example, primary, auxiliary) to be used for routing replication traffic between the two data centers. Ensure that the specified interface can reach the secondary data center and vice-versa. <b>Note:</b> If possible, use a dedicated interface for replication.
Journal LUN	Specify the dedicated Journal LUN you provisioned on the storage array from the drop-down list. <b>Note:</b> If the Journal LUN you created does not appear in the drop-down list, select Rescan LUNs.

- Click **Set Configuration**.

To modify any of these settings, click **Clear Replication Settings**.

You can add more replication interfaces by clicking **Add Interface**. To remove a replication interface, click the X icon next to the interface.

Next, set up the secondary data center.

### To set up the secondary data center for replication

- On the secondary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

The Replication Configuration panel shows the current replication settings.

**Figure 2-13. Replication (Set Up / Monitor) Page - Replication Configuration Panel**

The screenshot shows the 'Replication (Set Up / Monitor)' page. The 'Replication Configuration' panel is active. It contains the following elements:

- Role:** Radio buttons for 'Primary' and 'Secondary'. 'Secondary' is selected.
- Data Center Name:** A text input field.
- Replication Interface:** A dropdown menu showing 'primary - 10.1.39.70'.
- Journal LUN:** A dropdown menu showing 'No LUNs found'.
- Set Configuration:** A blue button at the bottom.
- Help:** A question mark icon in the top right corner.

Below the Journal LUN dropdown, there is a note: *(Journal LUN must be larger than 3.00 GB. However, a thinly-provisioned LUN of 500 GB or more is recommended.)*

2. Select Secondary in the Replication Configuration panel.
3. Complete the configuration using the controls described in this table.

Setting	Description
Data Center Name	Specify the name for the secondary data center.
Replication Interface	<p>Specify the replication interface (for example, primary, auxiliary) to be used for routing replication traffic between the two data centers. If you have Cores that are set up for high-availability, assign the same interface to both.</p> <p>Ensure that the specified interface can reach the primary data center and vice-versa.</p> <p><b>Note:</b> If possible, use a dedicated interface for replication.</p>
Journal LUN	<p>Select the dedicated Journal LUN you provisioned on the storage array from the drop-down list. If you have Cores that are set up for high-availability, choose the same Journal LUN for both.</p> <p><b>Note:</b> If the Journal LUN that you created does not appear in the drop-down list, select Rescan LUNs.</p>

**Note:** If the Cores are configured for high-availability, each pair of Cores in each data center must share the same Data Center Name, Role, and Journal LUN. A warning appears if you attempt to assign different values to each Core.

#### 4. Click **Set Configuration**.

The secondary data center Core is now ready for communication with the primary Core.

Once you have set up the data centers, they are prepared to connect to each other. If the Cores are set up for high-availability, repeat these steps on the failover peer.

The next step is to peer the Cores together from the primary Core.

## Pairing the Cores

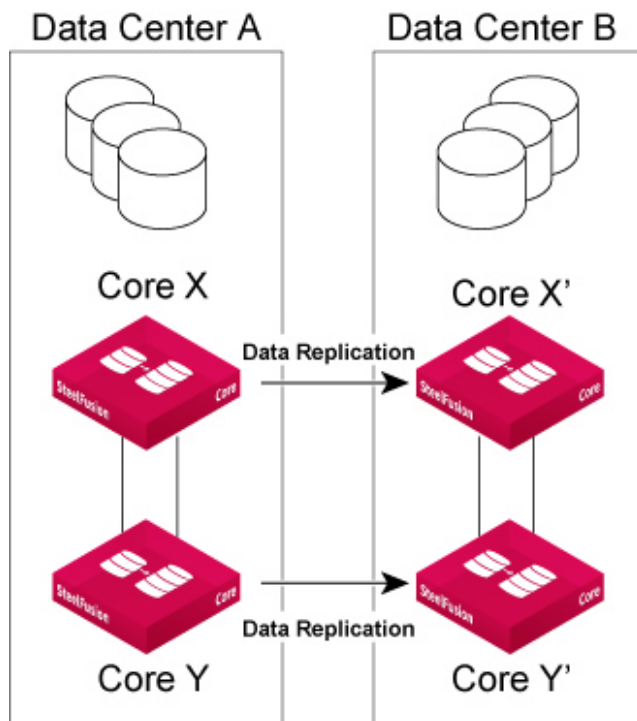
After you have specified the roles of the primary and secondary Cores for replication, pair them together from the primary data center to establish connections for replication. If you have two Cores in a high-availability configuration, you pair both of them to the secondary data center. For example, you could pair Core X to Core X' and Core Y to Core Y'.

---

**Note:** If you are setting up Cores for replication and also plan on setting them up for high-availability, ensure that you configure high-availability first. For more information about high-availability, see [“Configuring Failover” on page 80](#).

---

Figure 2-14. Replication in a High-Availability Setup



## To pair the Cores

1. From the primary Core, choose **Configure > Replication: Set Up / Monitor** to display the Replication (Set Up / Monitor) page.

**Figure 2-15. Replication (Set Up / Monitor) Page - Replication Pair Connection Panel**

The screenshot shows the 'Replication (Set Up / Monitor)' page. The top navigation bar includes 'Replication' and 'Replication (Set Up / Monitor)'. The main content area is divided into two sections:

**Replication Configuration**

Role:	Primary
Data Center Name:	site2
Replication Interface:	eth0.0 - 10.5.137.240
Journal LUN:	B8N0U+DzsfnD (50.01 GB)

Below the configuration table is a button labeled 'Clear Replication Settings'.

**Replication Pair Connection**

Below this heading are two text input fields:

Secondary Data Center Name:

Secondary IP:

Below these fields is a button labeled 'Connect to Secondary'.

2. In the Replication Pair Connection panel, specify the name of the secondary data center you chose in [Step 3](#) in the Secondary Data Center Name text box.
3. In the Secondary IP text box, specify the IP address of the secondary Core's replication interface.
4. Click **Connect to Secondary**.

Repeat Steps 1 to 4 if you have configured high-availability Cores in the primary data center. You only have to peer each high-availability Core with one Core in the secondary data center.

After you have paired the Cores, the next step is to set up the Witness.

## Configuring the Witness

The Witness is an Edge that you choose to be the authoritative source on each data center's state in case they enter a "split-brain" scenario, in which both Cores attempt to journal writes at the same time. Any requests to suspend replication and start journaling are approved by the Witness, which ensures that only one data center is approved for journaling at any given time.

You can only configure the Witness from the primary data center; however, once it is set up, it is available and visible on all Cores across both the primary and secondary data centers.

These requirements must be met for the Witness to work:

- The Edge must be running version 4.3 or later.
- The Witness must be reachable from both data centers when it is configured.

## Leader and Follower Roles in High-Availability

If the Cores are set up for high-availability, one Core has the role of the leader and the other has the role of the follower. Some configuration changes are only possible on the leader Core. If the leader is down, the follower Core assumes the role of leader. When the original leader comes back up again, it resumes its role.

Even though you can add Edges to both high-availability Cores, only the Edges on the leader Core can be configured as a Witness.

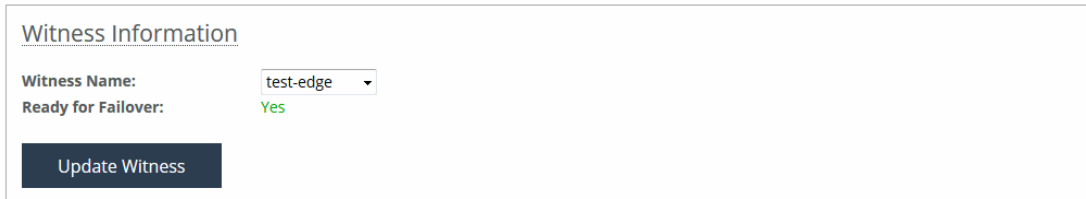
## Witness Recommendations

- Choose a high-availability Edge on the leader Core as the Witness in order to protect against having a single point of failure on the Edge.
- Choose a Witness in a location that would not be affected by disasters that could potentially bring down the data center.
- Set up multiple redundant paths from the Witness to the primary and secondary data centers. For more information about configuring WAN redundancy, see the *SteelFusion Design Guide*.

### To set up the Witness

1. From the primary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.
2. In the Witness Information panel, select an Edge from the Witness Name drop-down list.

**Figure 2-16. Replication (Set Up / Monitor) Page - Witness Information Panel**



---

**Note:** If the Cores are configured for high-availability, select the Edge that is connected to the primary data center Leader Core.

---

3. Click **Update Witness**.

Once the Ready for Failover status changes to Yes, the Edge is ready to assume the role of Witness.

Now that the Cores are aware of each other and the Witness is configured, add the replica LUNs on the primary data center and start first sync.

## Changing the Witness

If the Witness is down or experiencing another type of failure, you can change it to a different Edge in the Witness Information panel on the primary Core. Any issues relating to the Witness appear in the Reports > Diagnostics: Alarm Status page under Edge Service.

---

**Note:** You can change the Witness at any time, as long as replication is not suspended.

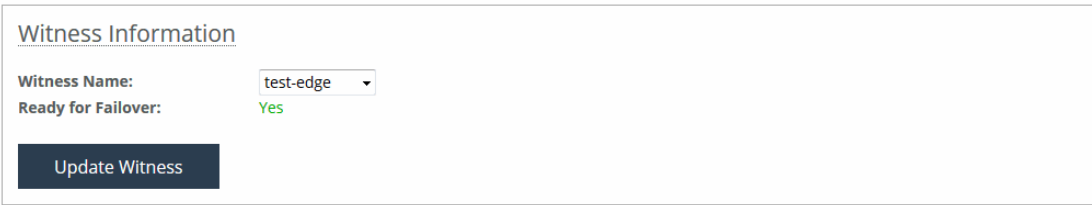
---

### To change the Witness

1. From the primary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

2. In the Witness Information panel, select the Edge from the Witness Name drop-down list.

**Figure 2-17. Replication (Set Up / Monitor) Page - Witness Information Panel**



3. Click **Update Witness**.

## Configuring Edges and LUNs for Replication

After you have set up the Witness, you will map the replica LUNs. A replica LUN is a LUN on the secondary data center that acts as a mirror to the primary LUN in the primary data center. The replica LUN can be either iSCSI or Fibre Channel, and must be equal in size or larger than the primary LUN (a 1 GB leeway exists). To start replication for an Edge, each of its mapped LUNs must have an associated replica LUN. When you start replication for an Edge, the Core replicates all the LUNs that are mapped to that Edge.

Once the replication for an Edge is enabled, the first sync of data from the primary LUN to the replica LUN starts. During this process, the Core continues to run normally without interruption. There are two options for first sync:

- **Full Sync** - The Core performs a full block-by-block copy of the primary LUN to the replica LUN until they are exact copies of each other and are fully synchronized. This mode is selected by default.
- **None** - The Core does not copy any blocks from the primary LUN to the replica LUN. You may want to select this option if your LUNs are not formatted yet or if you plan to use a third-party replication tool to synchronize the LUNs.

---

**Note:** Enabling and starting replication is not supported for individual LUNs.

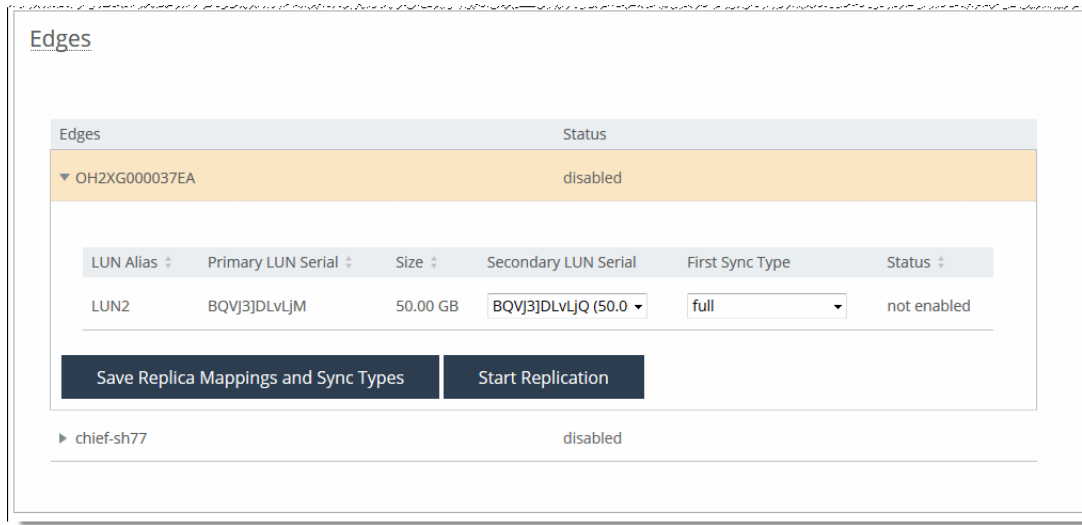
---

### To map the replica LUNs and start replication

1. From the primary Core, choose **Configure > Replication: Set Up / Monitor** to display the Replication (Set Up / Monitor) page.
2. In the Edges panel, click the Edge you want to configure.

- Click the LUN you want to map.

**Figure 2-18. Replication (Set Up / Monitor) Page - Edges Panel**



- Select a replica LUN from the Secondary LUN Serial drop-down list.

The drop-down list is prepopulated with the replica LUNs that you previously created on the storage array in the secondary data center. The size of the replica LUN is listed next to each one.

To verify whether all LUNs configured on the secondary data center storage array are being displayed on the primary Core for replication, you can use the **show storage coredr sec-site-luns** command.

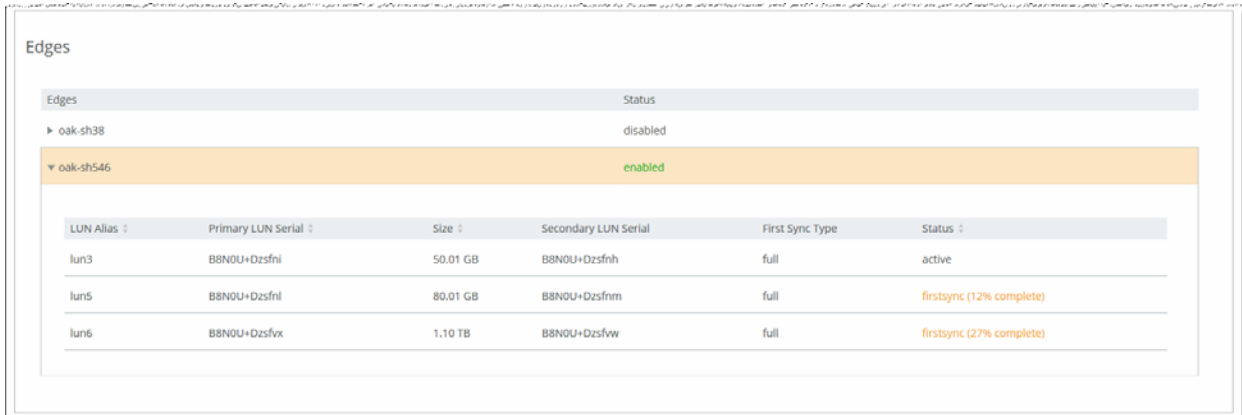
- Select the First Sync Type from the drop-down list.
- Click **Save LUN Replica Mappings and Sync Types**.

If there are multiple LUNs mapped to the Edge, the **Start Replication** button is not enabled until all LUNs have mapped replicas.

- Click **Start Replication**.

The Edge Status changes to enabled. If you selected Full from the First Sync Type drop-down list, the Status indicates the first sync percentage of completion.

**Figure 2-19. Replication (Set Up / Monitor) Page - Edges Panel**

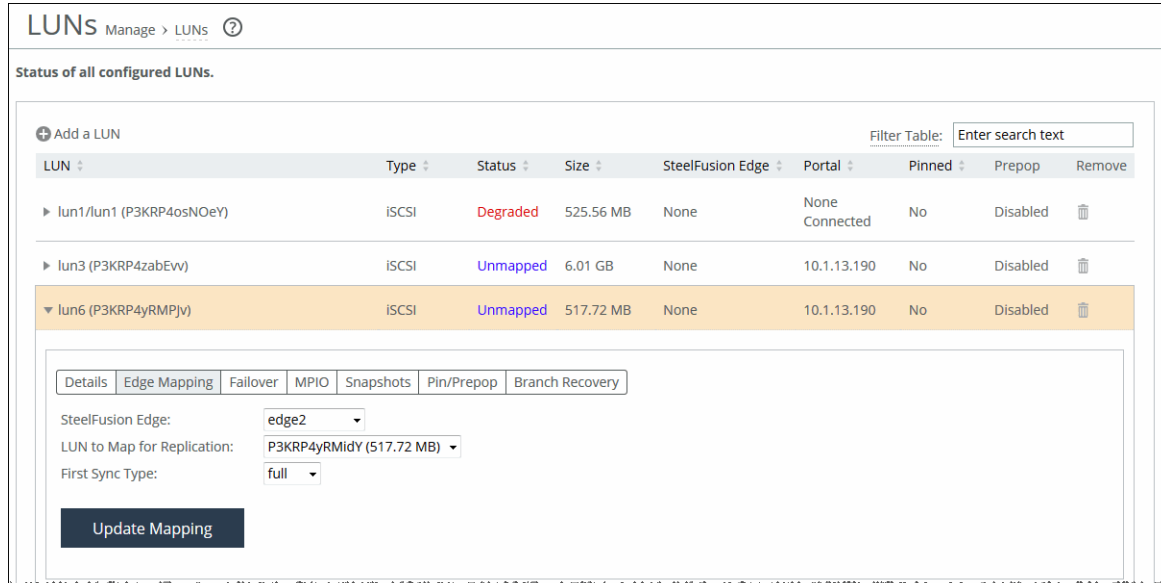


After the LUN is copied to the secondary data center during the first sync, the primary LUNs and replica LUNs are in a consistent state and replication is active. You can choose a different sync type for each LUN.

### To add a new LUN to a replicating Edge

1. Choose Configure > Manage: LUNs to display the LUNs page and select the Edge Mapping tab.

**Figure 2-20. LUNs Page - Edge Mapping Tab**



2. Select the Edge to map the LUN to from the SteelFusion Edge drop-down list.
3. Select the replica LUN from the LUN to Map for Replication drop-down list.
4. Select the First Sync Type from the drop-down list.
5. Click **Update Mapping**.

## Suspending Replication

During an extended data center failure or planned network event such as a network upgrade on the secondary data center, you can suspend replication on the primary data center to prevent the Edge's blockstore from accumulating excessive data. Suspending replication will not clear any replication settings.

While replication is suspended, the primary data center Core uses the Journal LUN to track the writes from the Edge until the connection is restored. When the secondary data center is restored, the Cores automatically reconnect and journaled writes are synced to the secondary data center to make the replica LUN and primary LUN consistent. Failback is possible if this data center is preferred. For more information about failback, see [“Failing Back to the Primary Data Center”](#) on page 59.

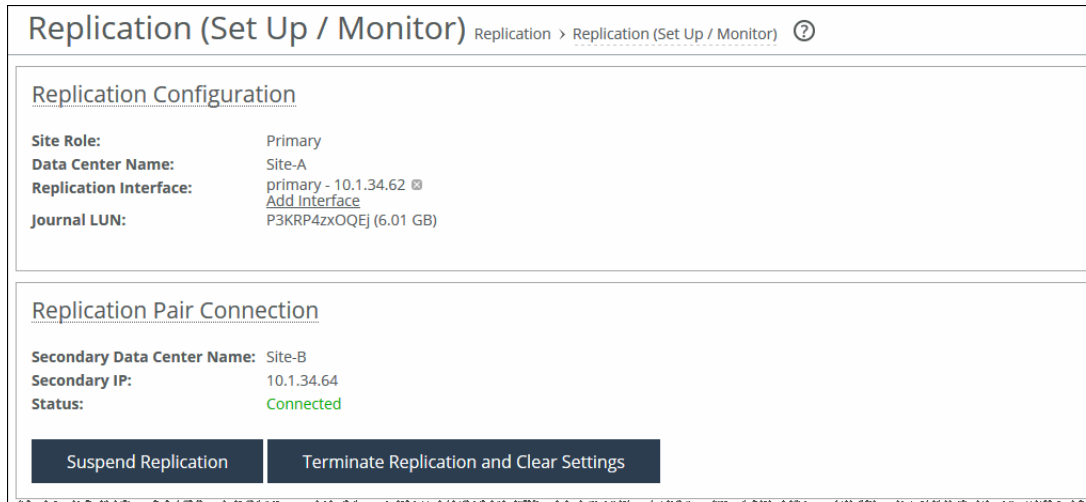
If the Cores are set up for high-availability, replication must be suspended from the leader Core in the primary data center. Once the leader approves the change, the follower also suspends replication, and they are both in the suspended state. If the leader fails for any reason, the follower automatically becomes the leader.

## To suspend replication on the primary data center

1. On the primary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

Replication is still active, so the primary Core is still attempting to replicate to the secondary data center before acknowledging writes to the Edge.

**Figure 2-21. Replication (Set Up / Monitor) Page - Replication Pair Connection Panel**



2. Click **Suspend Replication** in the Replication Pair Connection panel.

The Core now contacts the Witness to verify whether it is possible to suspend replication to the secondary data center. Because the Witness knows the state of each data center, it is the authoritative source for this information.

Once the Core successfully suspends replication, the Status in the Replication Pair Connection and Edges panels will change to Suspended. The primary data center now writes to the storage array and journal all the incoming writes to the Journal LUN.

To resume replication, click **Resume Replication**.

## Initiating Failover (Recovering from Primary Data Center Failure)

In the event of primary data center failure, where the data center is decommissioned for an extended period of time, you can easily have the primary data center's LUNs and Edges fail over to the secondary data center to continue Edge operations. For the Edges, this process is entirely transparent. A failover is not automatic; it is initiated by the administrator and is only permitted when the following criteria are met:

- The primary data center is down.
- All live and replica LUNs are synchronized across data centers (to prevent data loss).
- The failover is approved by the Witness.
- All LUNs on Edges are in an *active* state.

After failover is complete, the Secondary Core's role in Replication Configuration panel will change to Primary. All Edge connections and data commits will move to the new Primary, Edge commits will be "In Progress", and replication will be suspended.

The original secondary data center will now act as primary and the original primary data center will act as secondary whenever it recovers from the failure.

**Note:** Any Edge that you did not set up for replication does not have its LUNs available during failover because they were not being replicated on the secondary data center.

### To initiate failover to the secondary data center

1. On the secondary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

Figure 2-22. Replication (Set Up / Monitor) Page - Initiate Failover Panel

The screenshot shows the 'Replication (Set Up / Monitor)' page. The breadcrumb trail is 'Replication > Replication (Set Up / Monitor)'. The page is divided into four main sections:

- Replication Configuration:**
  - Role: Secondary
  - Data Center Name: sec
  - Replication Interface: primary - 10.1.48.52 (with an 'Add Interface' link)
  - Journal LUN: C4hsKK1tjsye (8.01 GB)
- Replication Pair Connection:**
  - Secondary Data Center Name: pri
  - Status: Disconnected
- Witness Information:**
  - Witness Name: kabar-wa65
  - Ready for failover: Yes
- Initiate Failover:**
  - A button labeled 'Initiate Failover'.

2. Click **Initiate Failover** in the Initiate Failover panel.

Core contacts the Witness to verify if the primary data center is actually down.

If the secondary Core has successfully taken over and is serving the LUNs, the Role in the Replication Configuration panel appears as "Primary."

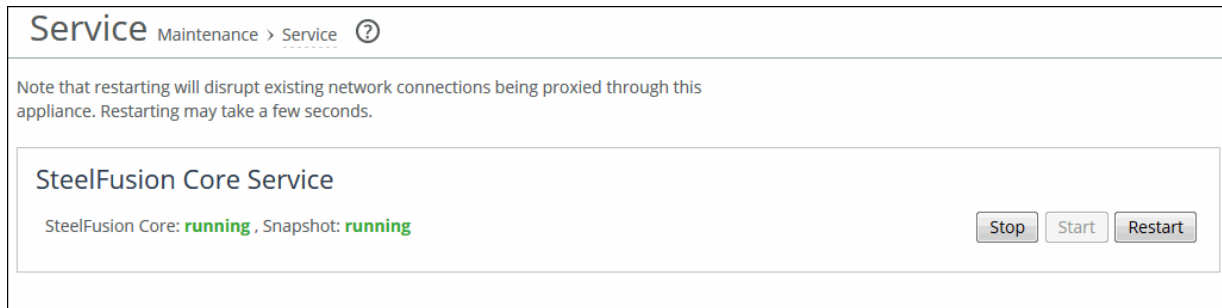
## Failing Back to the Primary Data Center

Once a failed data center has recovered, you can restore its role to primary if it is a preferred data center (due to hardware configuration or geographical location, for example). To failback to the original primary data center, you initiate failover from the preferred data center, which is currently in the secondary role.

## To failback to the original primary data center

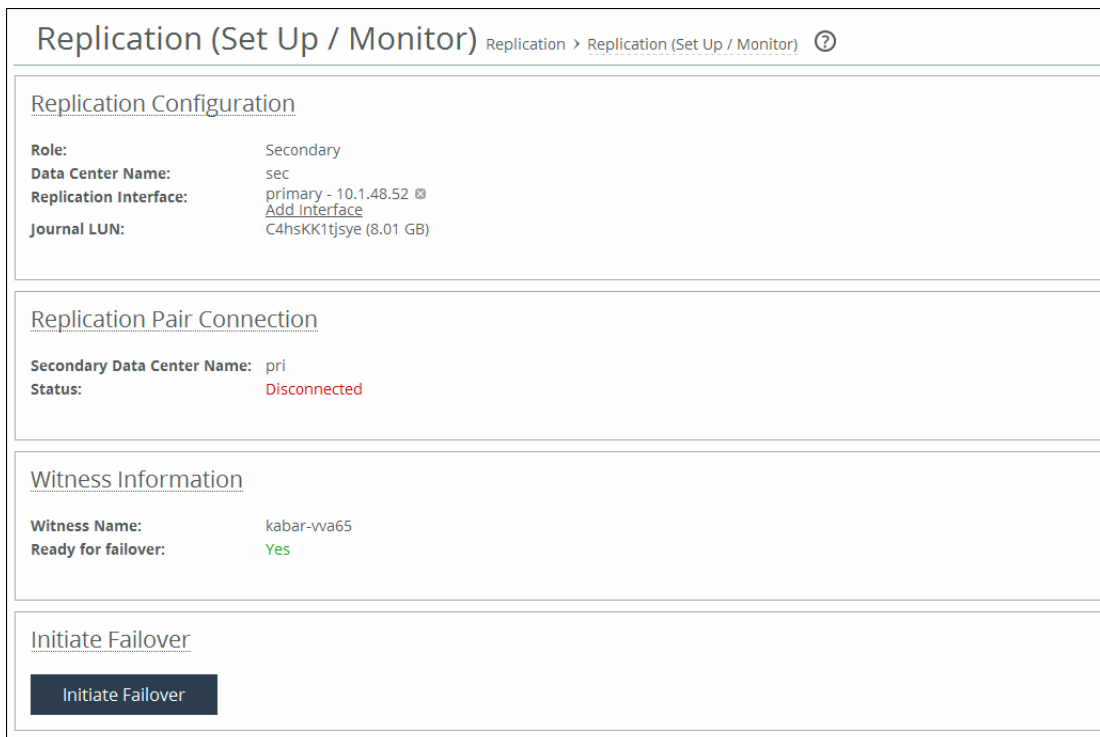
1. On the current primary Core, choose Settings > Maintenance: Service.

Figure 2-23. Service Page



2. Click **Stop** to stop the SteelFusion Core service.
3. On the secondary Core (the original primary Core), choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

Figure 2-24. Replication (Set Up / Monitor) Page



4. Click **Initiate Failover** in the Initiate Failover panel.
5. On the Core where you stopped the SteelFusion Core service, choose Settings > Maintenance: Service to display the Service page.
6. Click **Restart** to restart the SteelFusion Core service.  
The data centers return to their original primary and secondary roles.

## Terminating Replication

To stop replication for an Edge, you must first terminate replication for the Edge, and then unmap all replica LUNs associated with that Edge. To completely terminate replication for a Core, you must disable replication for all the replicating Edges.

---

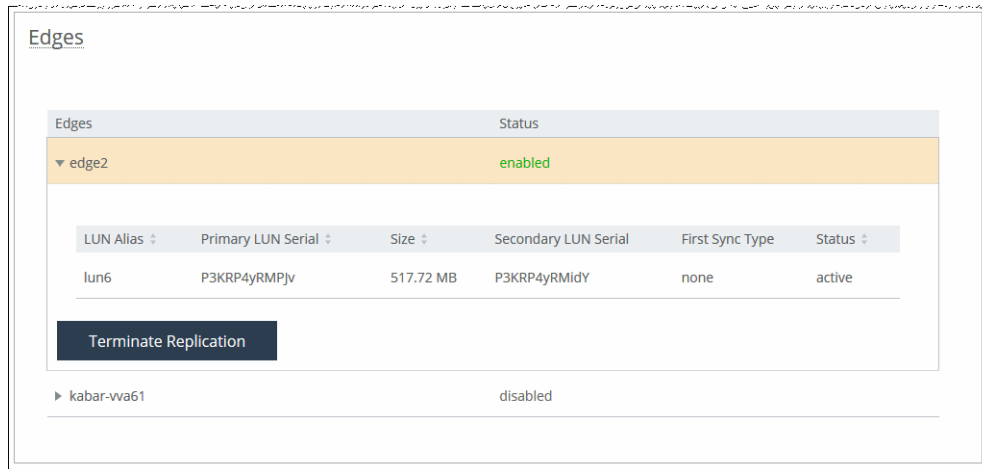
**Note:** If the Cores are set up for high-availability, terminating replication from the primary data center leader also terminates it on all four nodes.

---

### To terminate replication for the Edge

1. On the primary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

**Figure 2-25. Replication (Set Up / Monitor) page - Edges panel**



2. In the Edges panel, click **Terminate Replication**.

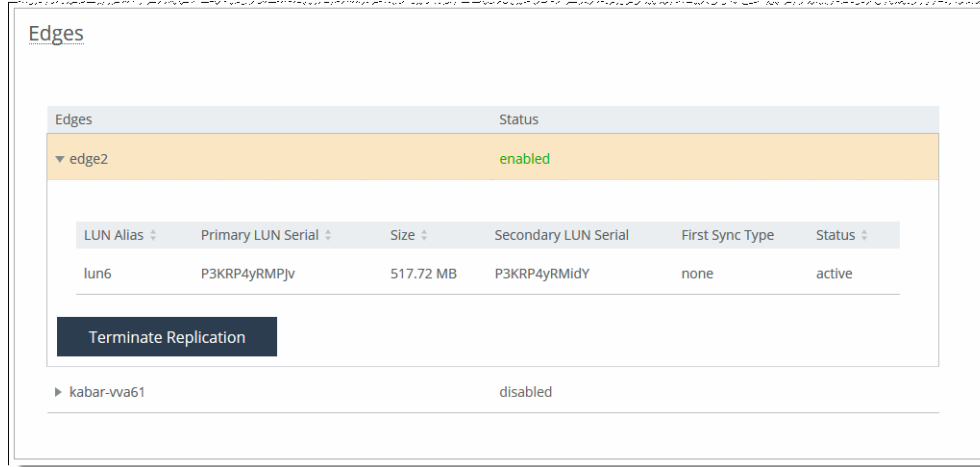
The status for the Edge changes to *disabled*.

Next, you will terminate replication for the Core.

## To terminate replication for the Core

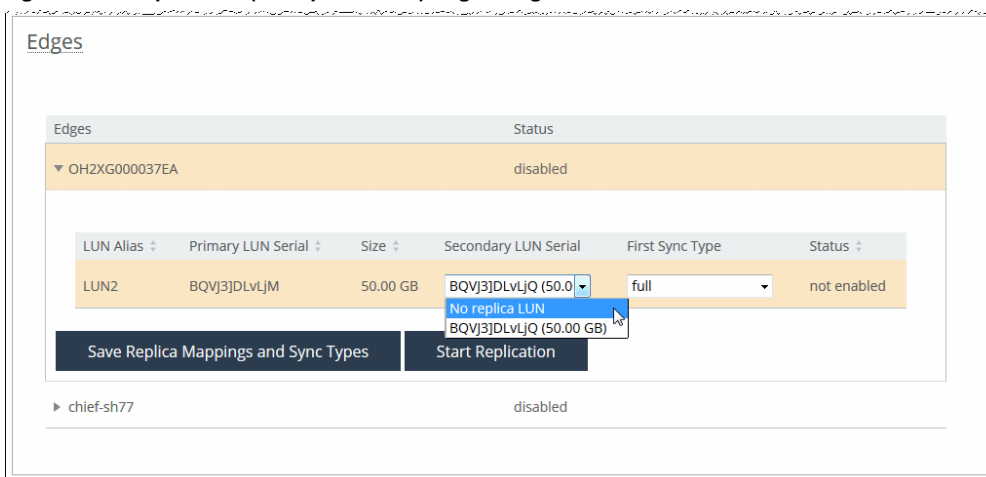
1. On the primary Core, choose Configure > Replication: Set Up / Monitor to display the Replication (Set Up / Monitor) page.

Figure 2-26. Replication (Set Up / Monitor) page - Edges panel



2. In the Edges panel, click **Terminate Replication**.  
The status for the Edge changes to *disabled*.
3. Repeat Step 2 for each Edge with an *enabled* status.  
Replication for all Edges must be in a *disabled* state in order to terminate Core replication.
4. For each LUN in the Edges panel, select No Replica LUN from the Secondary LUN Serial drop-down list.

Figure 2-27. Replication (Set Up / Monitor) Page - Edges Panel



5. Click **Save Replica Mappings and Sync Types**.

The status for the current mapped LUN changes from *active* to *deleting*. Once the Core has finished unmapping the replica LUN, the status changes to *unmapped*.

Repeat these steps for each LUN under the Edge.

6. Click **Terminate Replication and Clear Settings** in the Replication Pair Connection panel.
7. Click **Clear Replication Settings** in the Replication Configuration panel.

## Configuring CHAP Users

Challenge-Handshake Authentication Protocol (CHAP) validates the identity of remote clients by periodically verifying the identity of the client using a three-way handshake. This validation happens at the time of establishing the initial link and might happen again at any time. CHAP bases verification on a user password and transmits an MD5 sum of the password from the client to the server.

CHAP can be one-way or mutual. In one-way CHAP, the target (server) authenticates the initiator (Core). In mutual CHAP, the target authenticates the initiator and additionally the initiator authenticates the target.

You can configure CHAP users and passwords in the CHAP Users page. You will be using these CHAP credentials for authentication when you log in to the storage array from the Core.

**Note:** You can also configure CHAP users dynamically in the iSCSI Configuration page. For details, see [“Configuring iSCSI Settings” on page 30](#).

### To configure one-way CHAP

1. Choose **Configure > Storage Array: CHAP Users** to display the CHAP Users page.

**Figure 2-28. CHAP Users Page**

CHAP User Table			
+ Add a CHAP User			
User	Enabled	In Use	Remove
Chap Users ▶ chap_core	Yes	No	
Chap Users ▶ chap_granitecore	Yes	No	
Chap Users ▶ chap_netapp	Yes	No	

2. Add new CHAP users using the controls described in this table.

Control	Description
Add a CHAP User	Displays controls for adding a new CHAP user to the running configuration.
CHAP Username	Specify a descriptive CHAP username or the IQN of the Core.
Password/Confirm Password	Specify the password for the new CHAP user that you configured on the backend.

Control	Description
Add CHAP User	Adds the new CHAP user to the running configuration.

- To modify an existing CHAP user configuration, click the username in the User table to expand a set of additional controls.

New CHAP users are enabled by default.

- To disable a CHAP user:

- Click the username to expand the set of additional controls.
- Clear the Enable check box.
- Click **Update CHAP User**.

- To remove an existing CHAP user configuration, click the trash icon in the Remove column.

### To configure mutual CHAP

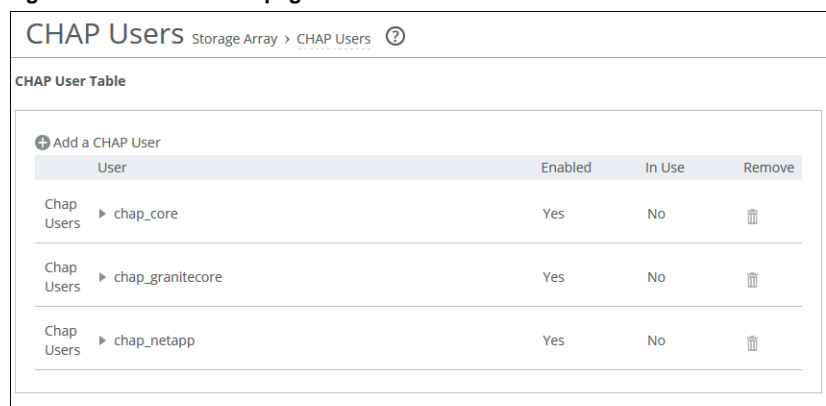
To configure mutual CHAP, you will create two CHAP users: one on the storage array and one on Core.

- Set up a target secret on the storage array.

For more information, refer to the documentation that came with your storage array.

- On Core, choose **Configure > Storage Array: CHAP Users** to display the CHAP Users page.

**Figure 2-29. CHAP Users page**



- Click **Add a CHAP User** and create the first CHAP username.

For example, if you are using a NetApp storage array, you could specify **chap\_netapp** as the username.

This username represents the CHAP user that the Core uses while connecting to the storage array.

- In the Password field, specify the secret you created on the storage array in Step 1.

The Core uses these credentials to connect to the storage array.

5. Choose Configure > Storage Array: iSCSI, Initiators, MPIO to display the iSCSI Configuration page.

**Figure 2-30. iSCSI, Initiators, MPIO Page**

6. Select Enable Mutual CHAP Authentication in the iSCSI Initiator Configuration panel.
7. Click **Add new Mutual CHAP User** and create a second CHAP username and password.  
For example, you could specify **chap\_core** as the username.  
The Core requires any storage array to provide these credentials before it can be authenticated.
8. On the storage array, enter the credentials you created in Step 7.  
The storage array uses these credentials when replying to the Core as part of the authentication process.

## Configuring Snapshots and Proxy Backup

The Core integrates with the snapshot capabilities of the storage array and enables you to configure application-consistent snapshots through the SteelFusion Core Management Console.

By bringing backups back to the data center, the Core eliminates the requirement and infrastructure of running backups at the remote branches. Data from the branch is always synced back to the data center so you can schedule VADP backups (application consistent) of the branch VMs. Once the snapshot reaches the data center, the Core triggers the application-consistent snapshot on the supported array and automatically mounts the snapshot to a proxy host. The backup software on the storage array can then perform the backups through the proxy host in the data center.

If you have configured your data centers for replication, you must schedule snapshots on the current primary data center. Keep in mind that snapshots taken on one data center are only present on that particular data center, and snapshots will not work after failover if the snapshot configuration is not present on the failover data center. For more information about replication, see [“Configuring Replication” on page 47](#).

This section describes the following topics:

- [“Understanding Crash Consistency and Application Consistency” on page 66](#)

- [“Configuring Snapshots” on page 66](#)
- [“Configuring Snapshots for Storage Arrays” on page 67](#)
- [“Configuring Handoff Hosts” on page 69](#)
- [“Configuring Snapshot Schedule Policies” on page 70](#)
- [“Defining Branch and Proxy Hosts” on page 72](#)
- [“Configuring Snapshots for LUNs” on page 73](#)
- [“Applying a Snapshot Schedule Policy to a LUN” on page 74](#)
- [“Configuring Application-Consistent Snapshots for a LUN” on page 75](#)
- [“Configuring Proxy Backup for a LUN” on page 77](#)

## Understanding Crash Consistency and Application Consistency

In the context of snapshots and backups and data protection in general, there are two types or states of data consistency:

- **Crash consistency** - A backup or snapshot is *crash consistent* if all of the interrelated data components are as they were (write-order consistent) at the instant of the crash. To better understand this type of consistency, imagine the status of the data on your PC’s hard drive after a power outage or similar event. A crash-consistent backup is usually sufficient for nondatabase operating systems and applications like file servers, DHCP servers, print servers, and so on.
- **Application consistency** - A backup or snapshot is *application consistent* if, in addition to being write-order consistent, running applications complete all their operations and flush their buffers to disk (application quiescing). Application-consistent backups are recommended for database operating systems and applications such as SQL, Oracle, and Exchange.

SteelFusion ensures continuous crash consistency at the branch and at the data center by using journaling and by preserving the order of write operations across all the exposed LUNs. For application-consistent backups, you can directly configure and assign hourly, daily, or weekly snapshot policies. Edges interact directly with both VMware ESXi and Microsoft Windows servers, through VMware Tools and VSS, to quiesce the applications and generate application-consistent snapshots of both VMFS and NTFS data drives.

## Configuring Snapshots

The general process for setting up snapshots is as follows:

### 1. Determine snapshot support for your storage array.

The Core supports snapshot configuration for storage arrays from specific providers (Dell EqualLogic, EMC CLARiiON, EMC VNX, NetApp, or IBM Storwize v7000). However, if your storage array is different, you can use the handoff hosts feature to point to a host that takes the snapshot using scripts. For details, see [“Configuring Handoff Hosts” on page 69](#).

For details about storage array support in general, see the *SteelFusion Design Guide*.

### 2. Define the storage array details for the snapshot configuration.

Before you can configure snapshot schedules, application-consistent snapshots, or proxy backup servers for specific LUNs, you must specify for the Core the details of the storage array, such as IP address, type, protocol, and so on.

You add storage array definitions in the Storage Arrays tab in the Configure > Backups: Snapshots page. For details, see [“Configuring Snapshots for Storage Arrays” on page 67](#).

**3. Configure the proxy server and storage array for snapshots.**

In conjunction with the Core configuration, you must separately configure the proxy server and storage array for backup. For details, see the *SteelFusion Design Guide*.

**4. Define snapshot schedule policies.**

Snapshot schedule policies are predefined schedules that can later be applied to snapshot configurations for specific LUNs.

You define snapshot schedule policies in the Snapshot Schedule Policies tab in the Configure > Backups: Snapshots page. For details, see [“Configuring Snapshot Schedule Policies” on page 70](#).

**5. Define snapshot host credentials.**

Host credentials are predefined storage host configurations that can later be applied to snapshot configurations for specific LUNs.

You define storage array definitions in the Host Credentials tab in the Configure > Backups: Snapshots page. For details, see [“Defining Branch and Proxy Hosts” on page 72](#).

## Configuring Snapshots for Storage Arrays

You can implement block-level snapshots for specified storage arrays, configure snapshot schedule policies (which apply to configured LUNs), and configure host information for data protection in the Snapshots page.

When you add a storage array in this page, you are associating a set of credentials with individual VLUNs. This association enables the Core to take snapshots against the specified storage array.

## To implement block-level snapshots for specified storage arrays

1. Choose Configure > Backups: Snapshots to display the Snapshots page and select the Storage Arrays tab.

**Figure 2-31. Storage Arrays Tab in the Snapshots Page**

The screenshot shows the 'Snapshots' page with the 'Storage Arrays' tab selected. The page header includes 'Snapshots' and a breadcrumb 'Backups > Snapshots'. Below the header, there are four tabs: 'Storage Arrays', 'Handoff Hosts', 'Snapshot Schedule Policies', and 'Branch and Proxy Hosts (Data Protection)'. The main content area states: 'These are Storage Arrays on which snapshots of LUNs can be taken.' Below this is a section titled 'Add a Storage Array' with a form containing the following fields: 'Hostname or IP Address' (text input), 'Type' (dropdown menu with 'EMC CLARiiON' selected), 'Username' (text input), 'Password' (password input), 'Confirm Password' (password input), 'Protocol' (dropdown menu with 'HTTP' selected), and 'Port' (text input). An 'Add Storage Array' button is at the bottom of the form. Below the form is a table with the following columns: 'Hostname', 'Type', 'Status', 'In Use By', and 'Remove'. The table contains one entry: '10.1.13.190', 'Dell EqualLogic', 'Ready', '0 LUNs', and a trash icon.

2. Add new storage arrays for snapshots using the controls described in this table.

Control	Description
Add a Storage Array	Displays controls for adding a storage array to the running snapshot configuration.
Hostname or IP Address	Specify the hostname or IP address of the storage array. If you are configuring an EMC VNX storage array, you must use the Storage Processor IP address.
Type	From the drop-down menu, select the type of array by provider (Dell EqualLogic, EMC CLARiiON, EMC VNX, NetApp, or IBM V7000). If you are configuring a NetApp storage array, fields for setting protocol (HTTP or HTTPS) and port become active. Configure accordingly.
Username	Specify a username for the storage array access.
Password/Confirm Password	Specify and confirm a password for the storage array access.
Protocol/Port	Specify the communication protocol and port number.
Add Storage Array	Adds the specified storage array snapshot implementation to the running configuration.

3. To modify the username and password configuration of an existing storage array, or to view which LUNs are associated with the current storage array, click the storage array in the Hostname column.

4. To remove an existing snapshot configuration, click the trash icon in the Remove column.
5. Click **Update Snapshot Credentials** to save your settings permanently.

## Configuring Handoff Hosts

This section describes how to configure application-consistent snapshots for storage arrays from nonqualified providers. The Core can interoperate with any iSCSI-compliant storage array, but support for application-consistent snapshots is limited to storage arrays from qualified vendors (Dell EqualLogic, EMC CLARiiON, EMC VNX, NetApp, or IBM v7000). The handoff host feature enables you to configure external hosts and scripts to take the snapshots on other, nonqualified storage arrays.

The following procedure assumes that you have already installed the handoff host in your network and installed the snapshot program (Windows executable) and script on the host.

---

**Note:** Windows Server 2012 R2 is supported as of Core version 3.6 for application-consistent snapshots.

---

### To configure a handoff host

1. Choose **Configure > Backups: Snapshots** to display the Snapshots page, and select the Handoff Hosts tab.

**Figure 2-32. Handoff Hosts Tab in the Snapshots Page**

The screenshot shows the 'Snapshots' page with the 'Handoff Hosts' tab selected. The page title is 'Snapshots' with a breadcrumb 'Backups > Snapshots'. Below the title are four tabs: 'Storage Arrays', 'Handoff Hosts', 'Snapshot Schedule Policies', and 'Branch and Proxy Hosts (Data Protection)'. A message states: 'These are the hosts which manage taking snapshots of LUNs (if a LUN is configured to do so).' Below this is a section 'Add a Handoff Host' with a table listing the configured host. The table has columns: Hostname, Status, In Use By, and Remove. The host 'chief-cs123.lab.nbttech.com' is listed with a status of 'Ready' and '0 LUNs'. Below the table, the 'Configuration Status' is 'Ready'. There are two sub-sections: 'Configuration' and 'LUNs using this Handoff Host'. The 'Configuration' section contains fields for 'Username' (chief-cs123\user), 'New Password', 'Confirm Password', 'Script Path' (C:\Python27\python.exe C:\handoff\_scripts\sample\_script.py.txt), and 'Script Arguments' (-option1 value1). An 'Update Handoff Host' button is at the bottom.

2. Add a new handoff host for snapshots using the controls described in this table.

Control	Description
Add a Handoff Host	Displays controls for adding a handoff host to the running snapshot configuration.
Hostname or IP Address	Specify the hostname or IP address of the handoff host.

Control	Description
Username	Specify a username to give the Core access to the handoff host.
Password/Confirm Password	Specify and confirm a password for the Core to access the handoff host.
Script Path	Specify the paths to both the executable and the script on the handoff host. In both cases, you must provide the absolute path. For example: <code>C:\Python27\python.exe C:\handoff_scripts\sample_script.py.txt</code>
Script Arguments	Supply the arguments that conform to your script to configure the snapshot settings. Delimit arguments and their values with a space and two dashes, for example: <code>--username root --password corepass --storage-array 10.6.72.86</code>
Add Handoff Host	Adds the specified handoff host snapshot implementation to the running configuration.

After a snapshot configuration has been added, it appears in the list on the same page. This list displays the type (NetApp, and so on) as well as one of the following status conditions:

- **Ready** - This status indicates that a target has been configured to point to the LUN and that the LUN is reporting that it is connected.
- **Standby** - This status indicates that the LUN is available but no targets are currently configured to connect to it.
- **Inactive** - This status indicates that no target has been configured for the LUN and the LUN is not available.

3. To modify an existing configuration, click the link in the Hostname column to display the settings.
4. To remove an existing configuration, click the trash icon in the Remove column.
5. Click **Save** to save your settings permanently.

## Configuring Snapshot Schedule Policies

You can configure snapshot schedule policies that create snapshots on a scheduled basis when applied to a configured LUN.

All the snapshot schedules are added in the Snapshots page on the Core, from where they then get pushed to the Edge. The Core has a default snapshot policy called `default_policy` that is triggered every four hours starting at midnight.

To apply a policy to a LUN configuration, see [“Applying a Snapshot Schedule Policy to a LUN” on page 74](#).

## To configure policies for a snapshot schedule

1. Choose Configure > Backups: Snapshots to display the Snapshots page and select the Snapshot Schedule Policies tab.

**Figure 2-33. Snapshot Schedule Policies Tab in the Snapshots Page**

Storage Arrays Handoff Hosts **Snapshot Schedule Policies** Branch and Proxy Hosts (Data Protection)

These policies are used when configuring snapshot schedules for a LUN.

**Add a Snapshot Policy**

Snapshot Policy Name:

Automatically take snapshots:

☐ Every week on Sunday at 12AM keeping 5 snapshots.

☒ Every Sun, Mon, Tue, Wed at 12AM keeping 5 snapshots.

☐ Every day at 12AM, 1AM, 2AM, 3AM keeping 5 snapshots.

Policy Name	In Use	Remove
▶ default_policy	No	

2. Add new snapshot schedule policies using the controls described in this table.

Control	Description
Add a Snapshot Policy	Displays controls for adding a snapshot schedule policy to the running configuration. A maximum of five snapshots are held by default.
Snapshot Policy Name	Type a descriptive name for the policy: for example, <code>snap_policy_early</code> .
Automatically take snapshots	<p>Specify the interval for the snapshot policy:</p> <ul style="list-style-type: none"> <li>• <b>Every week...</b> - Select this option to specify a day and time for the snapshot.</li> <li>• <b>Every [specific weekday]...</b> - Select this option to specify one or more week days and time for the snapshot.</li> <li>• <b>Every day...</b> - Select this option to specify one or more specific times for a daily snapshot.</li> </ul> <p>You can specify multiple options. In all three options you can specify how many snapshots to keep.</p> <p><b>Note:</b> Multiple schedules can be combined. For example, you can have hourly, daily and weekly schedules, all as part of the same policy.</p>
Add Snapshot Policy	<p>Adds the specified snapshot schedule policy to the running configuration.</p> <p>In the Configure &gt; Manage: LUNs page, you can apply a snapshot schedule policy to a LUN configuration.</p>

After a snapshot policy configuration has been added, it appears in the list on the same page.

3. To modify the configuration of an existing snapshot schedule policy, click the policy name in the list to display these settings.

4. To remove an existing snapshot schedule policy, click the trash icon in the Remove column.
5. Click **Update Snapshot Policy**.
6. Click **Save** to save your settings permanently.

## Defining Branch and Proxy Hosts

You can define branch and proxy host information for data protection in the Snapshots page. You can then use the host credentials when configuring snapshot settings and data protection for a LUN configuration.

For details, see [“Configuring Application-Consistent Snapshots for a LUN” on page 75](#).

### To configure branch and proxy hosts

1. Choose **Configure > Backups: Snapshots** to display the Snapshots page and select the **Branch and Proxy Hosts (Data Protection)** tab.

**Figure 2-34. Branch and Proxy Hosts (Data Protection) Tab in the Snapshots Page**

The screenshot shows the 'Snapshots' page with the 'Branch and Proxy Hosts (Data Protection)' tab selected. The page has a 'Save' button in the top right. Below the tabs, there is a description: 'These hosts are used when configuring Application Consistent Snapshots, or Proxy Backup for a LUN.' Below this is an 'Add a Host' section with a form containing the following fields: 'Hostname or IP Address' (text input), 'Type' (dropdown menu with 'VMWare Host/vCenter for Branch' selected), 'Username' (text input), 'Password' (text input), and 'Confirm Password' (text input). An 'Add Host' button is at the bottom of the form. Below the form is a table with columns 'Hostname', 'Type', and 'Remove'. The table contains one entry: '10.1.12.341' for Hostname and 'VMWare Host/vCenter for Branch' for Type. A trash icon is in the 'Remove' column for this entry.

2. Add new configurations for hosts at the branch or proxy at the data center using the controls described in this table.

Control	Description
Hostname or IP Address	Displays controls for adding a snapshot host to the running configuration. Snapshot hosts store snapshots until they expire.
Type	Specify one of the following types of hosts: <ul style="list-style-type: none"> <li>• VMware Host/vCenter for Branch</li> <li>• Windows Host (Proxy Backup) for Data center</li> <li>• VMware Host/vCenter (Proxy Backup) for Data center</li> <li>• VMware Host/vCenter (Proxy Backup) for Branch and Data center</li> </ul>

Control	Description
Username	Specify a username for host access.
Password/Confirm Password	Specify and confirm a password for host access.
Add Host	<p>Adds the specified host configuration to the running configuration.</p> <p>In the Configure &gt; Manage: LUNs page, you can specify a snapshot host for a LUN configuration.</p>

After a snapshot host configuration has been added, it appears in the list on the same page.

3. To modify the configuration of an existing snapshot host configuration, click the name in the list to display these settings.
4. To remove an existing snapshot host configuration, click the trash icon in the Remove column.
5. Click **Update Host** to save your settings permanently.

## Configuring Snapshots for LUNs

The following steps describe the general process for applying specific snapshot configurations to LUNs through the Core:

1. **Select the LUN for the snapshot and access the snapshot settings.**

You can access the snapshot settings for a specific LUN in the Configure > Manage: LUNs page. Click the desired LUN to display controls that include the Snapshots tab. The Snapshots tab itself has three tabs: History, Scheduler, and Configuration. For details, see [“Configuring LUNs” on page 35](#).

2. **Apply a snapshot schedule policy to the current LUN.**

The controls in the Scheduler tab enable you to apply a previously configured policy to the current LUN. You can also create a new schedule directly in this panel. For details, see [“Applying a Snapshot Schedule Policy to a LUN” on page 74](#).

3. **Specify the storage array where the LUN resides.**

The controls in the Configuration tab enable you to specify the storage array where the current LUN resides and to apply a static name that is prepended on the names of snapshots. For details, see [“Configuring Application-Consistent Snapshots for a LUN” on page 75](#).

4. **Specify the client type.**

The controls in the Configuration tab enable you to specify the client type. To configure application-consistent snapshots and a proxy backup, you must set this value to **VMware** or **Windows**. For details, see [“Configuring Application-Consistent Snapshots for a LUN” on page 75](#).

5. **Enable and configure application-consistent snapshots.**

In the Configuration tab you can enable and configure application-consistent snapshots. The settings vary depending on which client type is selected. For details, see [“Configuring Application-Consistent Snapshots for a LUN” on page 75](#).

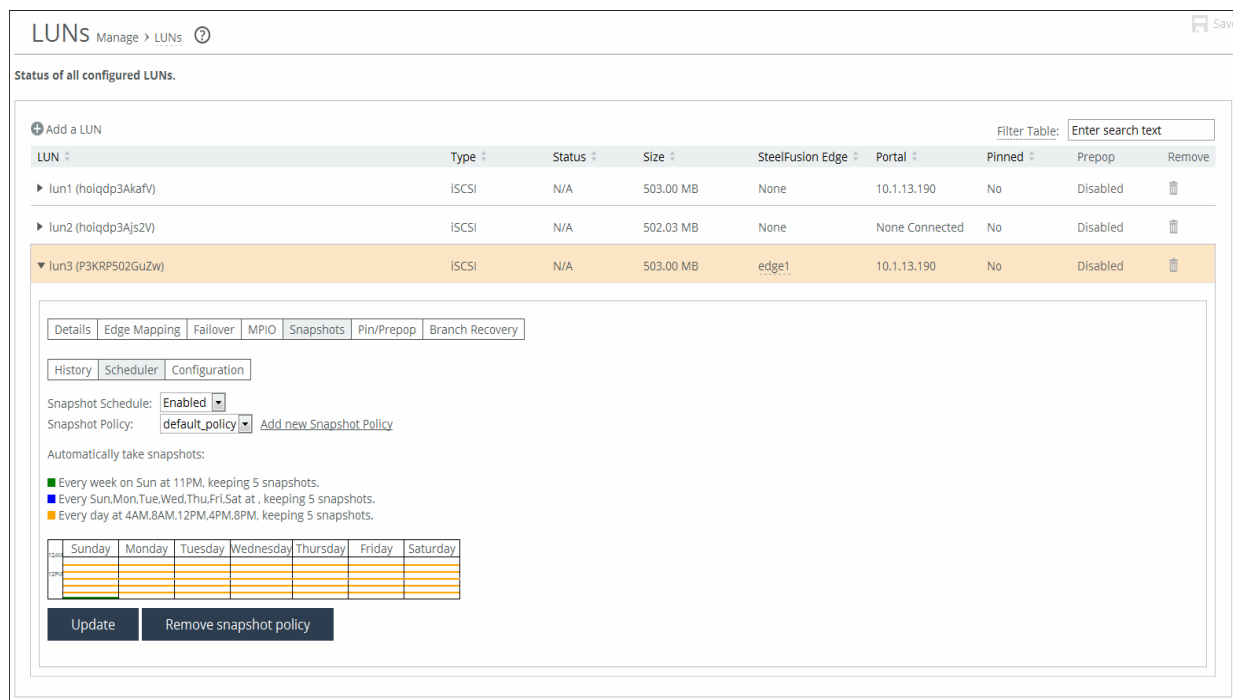
## Applying a Snapshot Schedule Policy to a LUN

When configuring specific LUNs, you can apply snapshot schedule policies to the current LUN in the Schedules tab within the Snapshots tab of the Configure > Manage: LUNs page.

### To enable and apply snapshot schedule policies

1. Choose Configure > Manage: LUNs to open the LUNs page.
2. From the list of configured LUNs, click the desired LUN to display controls and settings.
3. Select the Snapshots tab.
4. Within the Snapshots tab, select the Scheduler tab.

**Figure 2-35. Scheduler Tab in the LUNs Page**



5. Apply snapshot schedule policies to the current LUN using the controls in this table.

Control	Description
Snapshot Schedule	Select Enabled or Disabled from the drop-down list. If you select Enabled, additional controls appear that enable you to select from previously configured policies. If you select Disabled, the additional controls are disabled.
Snapshot Policy	Select the policy from the drop-down list. For details about configuring snapshot schedule policies, see <a href="#">“Configuring Snapshot Schedule Policies” on page 70</a> .
Update	Updates the changes made to the snapshot policy configuration.

6. Click **Update**.
7. To remove an existing snapshot schedule policy, click **Remove snapshot policy**.

## Configuring Application-Consistent Snapshots for a LUN

When configuring the snapshot schedule of a specific LUN, you can configure hourly, daily, or weekly application-consistent snapshots.

The procedure varies slightly, depending on the client type (**VMware** or **Windows**).

---

**Note:** To configure application-consistent snapshots for the current LUN, you must specify either **VMware** or **Windows** for the client type.

---

### To configure application-consistent snapshot settings

1. Choose Configure > Manage: LUNs to open the LUNs page.
2. From the list of configured LUNs, click the desired LUN to display controls and settings.
3. Select the Snapshots tab.
4. Within the Snapshots tab, select the Configuration tab.

**Figure 2-36. LUNs Page - Snapshots Tab - Configuration Tab**

The screenshot shows the SteelFusion Core Management Console interface. At the top, the breadcrumb navigation is "LUNS Manage > LUNS". Below this, a message states "Device Failover is enabled. (Click [here](#) for Failover details)". A dropdown menu indicates "You are currently viewing configurations for: Self".

The main section is titled "Status of all configured LUNs." and contains a table with the following columns: LUN, Type, Status, Size, SteelFusion Edge, Portal, Pinned, Prepop, and Remove. The table lists one LUN: "test\_lun2 (P3KRPS3dQqbZ)" with Type "iSCSI", Status "Connected (no snapshot settings)", Size "5.00 GB", SteelFusion Edge "oak-vva50", Portal "10.1.13.190", Pinned "No", and Prepop "Disabled".

Below the table, there are tabs for "Details", "Edge Mapping", "Failover", "MPIO", "Snapshots", "Pin/Prepop", and "Branch Recovery". The "Snapshots" tab is selected, and within it, the "Configuration" sub-tab is active.

The configuration panel is divided into several sections:

- General Settings:** Includes a "Storage Array" dropdown set to "10.5.31.78" with an "Add new Storage Array" link, and a "Handoff Host" dropdown set to "None". There are "Update Settings" and "Test Storage Array" buttons.
- Vendor-Specific Settings:** Includes a "Snapshots Static Name" text input field and an "Update Vendor-Specific Settings" button.
- Client Type:** Includes a "Client Type" dropdown set to "VMware" and an "Update Client Type" button.
- Application Consistent Snapshots (at the Branch):** Includes a checked "Enable" checkbox, a "Branch Hostname" dropdown set to "10.5.12.31" with an "Add new" link, a "Data center Name" text input field with a note "(Applicable for vCenter)", "Include VMs" and "Exclude VMs" text input fields with "(regex)" placeholders, a "Quiesce Guest VM" dropdown set to "No", and a checked "Fail on snapshot error" checkbox. There are "Update Settings" and "Reset Settings" buttons at the bottom.

5. In the General Settings panel, configure the settings using the controls in this table.

Control	Description
Add new Storage Array	Click to add a new storage array.
Handoff Host	Click to add a new handoff host.
Update Settings	Click to apply the modified settings.
Test Storage Array	Click to test the selected storage array to ensure that the specified credentials are correct and that the storage array is licensed for snapshots.
Client Type	Specify which kind of operating system mounts the LUN. <ul style="list-style-type: none"> <li>• <b>VMware</b> - The LUN is mounted by the Microsoft Windows OS and is formatted using NTFS.</li> <li>• <b>Windows</b> - The LUN is mounted by the VMware ESXi OS and is formatted using VMFS.</li> <li>• <b>Other</b> - All other cases.</li> </ul>

6. In the Application Consistent Snapshots panel, configure the settings using the controls in this table.

**Note:** Hints are available to help you specify regular expressions for VM names. To view the hints, hover over the phrase **(regex)** after the name field.

Control	Description
Enable	Select to enable application-consistent snapshots.
Branch Hostname (VMware vCenter clients only)	From the drop-down menu, select one of the previously configured snapshot hosts. If no host has been configured, click <b>Add New Host</b> to access controls for doing so.
Data center Name (VMware clients only)	Optionally, specify a data center for VMware-based application-consistent snapshots.
Include VMs (VMware clients only)	Optionally, specify by regular expression the names of the virtual machines (VMs) to be included in the snapshot. For example: <ul style="list-style-type: none"> <li>• To include VMs whose names begin with a specific prefix: <code>^2k3</code></li> <li>• To include VMs whose names end with a specific suffix: <code>ova1\$</code></li> <li>• To include VMs whose names start with a digit: <code>^\d</code></li> <li>• To include VMs with a specific name: <code>2k3vm_15g_ova1</code></li> <li>• To include VMs whose names have a particular format: <code>2k3v._15g*.ova</code></li> <li>• To include VMs whose names have a particular format: <code>2k3v.*ova</code></li> </ul>
Exclude VMs (VMware clients only)	Optionally, specify by regular expression the names of the VMs to be excluded from the snapshot. For example: <ul style="list-style-type: none"> <li>• To exclude VMs whose names end with a specific suffix: <code>ova1\$</code></li> <li>• To exclude VMs that have a specific name: <code>2k3vm_15g_ova1</code></li> </ul>
Quiesce Guest VM (VMware clients only)	Enables or disables quiescing VMs in the application-consistent snapshots. This setting must be set to Yes for snapshots to be application consistent.

7. Click **Update Settings** to apply the modified settings.

8. Click **Reset Settings** to restore the previous settings.

## Configuring Proxy Backup for a LUN

The proxy backup feature enables backup integration directly from the data center. The Core clones the snapshot specified for data protection and mounts the cloned LUN on the proxy server. The backup software then communicates with the proxy server to perform the backup.

You can configure data protection settings for the current LUN in the Setup tab within the Snapshots tab of the Configure > Manage: LUNs page.

The procedure varies slightly, depending on the type of LUN (Windows or VMware).

### Prerequisites

To configure proxy backup settings for the current LUN, you must specify either **VMware** or **Windows** for the client type in the General Settings panel in the same tab as the Proxy Backup panel.

---

**Caution:** Because each backup operation takes approximately two minutes to complete, enabling the hourly option for more than 30 LUNs can result in an increasing number of nonreplicated snapshots on Edges.

---

### To configure proxy backup options

1. Choose Configure > Manage: LUNs to display the LUNs page.
2. From the list of configured LUNs, click the desired LUN to display controls and settings.
3. Select the Snapshots tab.

4. Within the Snapshots tab, select the Configuration tab.

**Figure 2-37. LUNs Page - Snapshots Tab - Configuration Tab**

**LUNs** Manage > LUNs ? Save

Status of all configured LUNs.

Add a LUN

LUN	Type	Status	Size	SteelFusion Edge	Portal	Pinned	Prepop	Remove
test_jun2 (P3KRP53dQqbZ)	iSCSI	Connected (no snapshot settings)	5.00 GB	oak-vva50	10.1.13.190	No	Disabled	

Filter Table: Enter search text

Details | Edge Mapping | Failover | MPIO | Snapshots | Pin/Prepop | Branch Recovery

History | Scheduler | Configuration

**General Settings**

Storage Array: 10.5.31.78 Add new Storage Array

Handoff Host: None

Update Settings Test Storage Array

**Vendor-Specific Settings**

Snapshots Static Name:

Update Vendor-Specific Settings

**Client Type**

Client Type: VMware Update Client Type

**Application Consistent Snapshots (at the Branch)**

Enable

Branch Hostname: 10.5.12.31 Add new

Data center Name: (Applicable for vCenter)

Include VMs: .\* (regex)

Exclude VMs: (regex)

Quiesce Guest VM: No

Fail on snapshot error

Update Settings Reset Settings

**Proxy Backup (Snapshot mount at the Data center)**

Enable

Snapshot Mount Category: Daily

Proxy Hostname: 10.5.17.42 Add new

Proxy Host Storage Group:

Data center Name: (Applicable for vCenter)

Include VMs: .\* (regex)

Exclude VMs: (regex)

Include ESXi hosts: .\* (regex)

Exclude ESXi Hosts: (regex)

Update Settings Reset Settings

5. In the Proxy Backup panel, configure the settings using the controls in this table.

**Note:** Hints are available to help you specify regular expressions for VM names. To view the hints, hover over the phrase **(regex)** after the name field.

Control	Description
Enable	Select to enable this feature.
Snapshot Mount Category	<p>Select from the drop-down menu the snapshot mount category: <b>Weekly</b>, <b>Daily</b>, or <b>Hourly</b>.</p> <p><b>Note:</b> To successfully mount a snapshot, you must schedule a policy that includes a matching category.</p>
Proxy Hostname	From the drop-down list, select the hostname of the machine that will mount the cloned LUNs on the storage array to perform the backup the data. To add a new Proxy Hostname, click <b>Add New</b> .
Proxy Host Storage Group	<p>Specify the storage or initiator group for the snapshot (backup) host. This is the group that you defined on the backend storage array that grants access to the snapshot host. Once Core exposes a snapshot as a cloned LUN, it makes the cloned LUN available for reads and writes to the proxy host in the data center. This proxy may host the registered VMs in the case of VMware LUNs or may have the cloned LUN added as a mount point for NTFS LUNs. The backup software communicates with the proxy host to back up the data.</p> <p>Proxy Host Storage Group is different than a storage group in order to separate LUNs from cloned LUNs being used for snapshots. This way, the Core can only see and project live LUNs to the Edge.</p> <p><b>Caution:</b> If you are using a Dell EqualLogic storage array, you must configure the snapshot access policy on the storage array itself to ensure access for the proxy host.</p> <p>For details about creating snapshots of LUNs using VMware vStorage APIs for Data Protection (VADP), see <i>Protecting LUNs Using VMware vStorage APIs for Data Protection</i>, at <a href="https://splash.riverbed.com/docs/DOC-3588">https://splash.riverbed.com/docs/DOC-3588</a>.</p> <p>For details about creating snapshots of NTFS LUNs, see <i>Protecting SteelFusion NTFS LUNs</i>, at <a href="https://splash.riverbed.com/docs/DOC-4666">https://splash.riverbed.com/docs/DOC-4666</a>.</p>
Data center Name (VMware clients only)	Optionally, specify a data center for VMware-based backups.
Include VMs (VMware clients only)	<p>Optionally, specify by regular expression the names of the VMs included in the snapshot to be included in the backup. For example:</p> <ul style="list-style-type: none"> <li>• To include VMs whose names begin with a specific prefix: <code>^2k3</code></li> <li>• To include VMs whose names end with a specific suffix: <code>ova1\$</code></li> <li>• To include VMs whose names start with a digit: <code>^\d</code></li> <li>• To include VMs with a specific name: <code>2k3vm_15g_ova1</code></li> <li>• To include VMs whose names have a particular format: <code>2k3v._15g*.ova</code></li> <li>• To include VMs whose names have a particular format: <code>2k3v.*ova</code></li> </ul>
Exclude VMs (VMware clients only)	<p>Optionally, specify by regular expression the names of the VMs included in the snapshot to be excluded from the backup. For example:</p> <ul style="list-style-type: none"> <li>• To exclude VMs whose names end with a specific suffix: <code>ova1\$</code></li> <li>• To exclude VMs that have a specific name: <code>2k3vm_15g_ova1</code></li> </ul>

Control	Description
Include ESXi Hosts (VMware clients only)	Optionally, specify by regular expression the ESXi hosts to be included. For example, to include an ESXi host with a specific name: <code>dcESXi1c_backup</code>
Exclude ESXi Hosts (VMware clients only)	Optionally, specify by regular expression the ESXi hosts to be excluded. For example, to exclude an ESXi host whose name has a particular format: <code>dcESXi1.*_backup</code>
Update Settings	Applies the modified settings.
Reset Settings	Clears all proxy backup settings.

## Configuring Failover

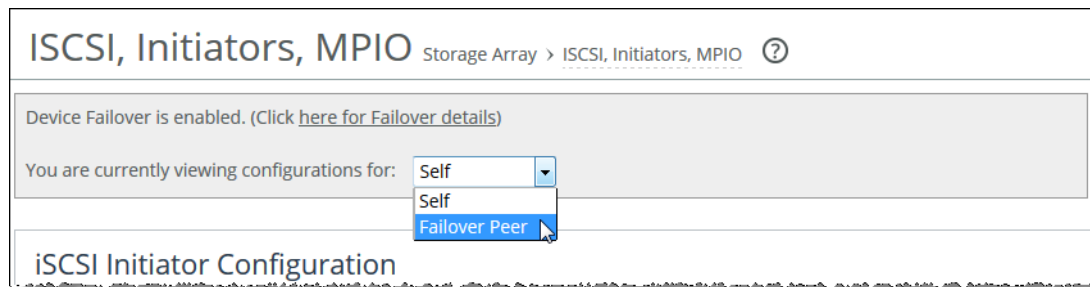
You can configure another device for failover in the Failover Configuration page.

For failover, Riverbed recommends connecting both failover peers directly with cables through two interfaces. If direct connection is not an option, Riverbed recommends that each failover connection use a different local interface and reach its peer IP address through a completely separate route.

If you configure the current Core for failover, all Storage configuration pages include an area with a link to failover details. Click the link to display the Configure > Failover: Failover Configuration page and the current settings.

Figure 2-38 shows a sample Storage configuration page with Failover enabled.

**Figure 2-38. Sample Storage Page with Failover Area**



This area appears on all Storage configuration pages except for Failover Configuration. The area appears below the page title, and it also includes a drop-down list from which you can select Self (the current appliance), Failover Peer, or another pool member. Changing the selection displays the configuration for those Cores.

**Note:** If you are planning to set up Cores for replication as well as high-availability, ensure that you configure high-availability first. For more information about replication, see [“Configuring Replication” on page 47](#).

## To configure a high-availability peer appliance for the first time

1. Choose Configure > Failover: Failover Configuration to display the Failover Configuration page.

**Figure 2-39. Failover Configuration Page - Configure Device Failover**

**Failover Configuration** Failover > Failover Configuration ?

*The recommended setup for failover is two direct cables from two interfaces on this SteelFusion Core, connected to a pair of interfaces on the Peer. If that is not possible, it is recommended that each failover connection use a different local interface, and reach its Peer IP address via a completely separate route.*

**Configure Device Failover**

Peer IP Address:

Local Interface: primary - 10.1.44.114 ▼

Second Peer IP Address:  *Must be a different IP address from the one above.*

Second Local Interface: aux - 169.254.1.1 ▼

**Enable Failover**

**Note:** If a peer interface has already been configured for high-availability, the Device Failover Settings panel displays the name of the host and a Disable/Enable button (depending on the current state). Otherwise, the page displays the Configure Device Failover panel containing the required fields to set up failover for the first time.

2. Configure the failover peer by specifying two connection paths using the controls described in this table.

Control	Description
Peer IP Address	Specify the IP address of the primary peer appliance.
Local Interface	Specify a local interface for connection to the above peer IP address.
Second Peer IP Address	Specify the IP address of the failover peer appliance. It must be different than the Peer IP Address.
Second Local Interface	Specify a different local interface for connection to the Second Peer IP address.
Enable Failover	Enables the new failover configuration.

3. Click **Enable Failover**.

The Failover Configuration page now displays the current failover settings in the Device Failover Settings panel.

**Figure 2-40. Failover Configuration - Device Failover Settings**

The screenshot shows the 'Failover Configuration' page. At the top, there is a breadcrumb trail: 'Failover > Failover Configuration' with a help icon. Below this is a note: 'The recommended setup for failover is two direct cables from two interfaces on this SteelFusion Core, connected to a pair of interfaces on the Peer. If that is not possible, it is recommended that each failover connection use a different local interface, and reach its Peer IP address via a completely separate route.'

The 'Device Failover Settings' section is highlighted in a light gray box. It contains the following information:

- Currently in Failover configuration with host: **oak-sh626**
- Currently Serving: Self Configuration
- High Availability State: Active Self

Below this information is a button labeled 'Disable Failover'.

The 'Add Failover Interface' section is below the settings. It has a title 'Add Failover Interface' with a plus icon. Inside this section, there are two input fields:

- 'Peer IP Address:' with an empty text input field.
- 'Local Interface:' with a dropdown menu showing 'primary - 10.5.64.53'.

At the bottom of this section is a button labeled 'Add Failover Interface'.

## Configuring Pool Management

You can configure a pool for managing up to 32 (physical or virtual) Cores from a single UI. You no longer have to log in to each individual Core to manage configuration of LUNs, Edges, or storage arrays. In addition, you can view reports on Edge data from the pool manager UI.

**Note:** All Cores must be running version 3.0 or later.

Pool management enables you to create and provide a consolidated access point for Cores, organized along logical, departmental, or geographic lines. In addition, the Core that is the pool manager can view and edit the configuration of any other Core in its pool.

**Note:** You can view the details of only one Core at a time from the manager's UI.

The pool is a single-level hierarchy with a flat structure, in which all members of the pool except the manager have equal priority and cannot themselves be managers of pools. Any Core can be the manager of the pool, but the pool manager cannot be a member of any other pool.

The pool has a *loose* membership, in which pool members are not aware of one another, except for the manager. The pool is dissolved when the manager is no longer available (unless the manager has an HA peer). Only a manager can dissolve a pool. However, a pool member can release itself from a pool.

Management of a pool can be taken over by a failover peer. However, a member's failover peer cannot be managed by the member's pool manager through the member, even if the failover peer is down.

The SteelFusion Edge Trends, LUN I/O, and SAN I/O graphs show data from pool members on the pool manager's pages. When a pool manager tries to select one of its members from the drop-down list, if that pool member is unavailable, a warning header appears, explaining that the member cannot be reached.

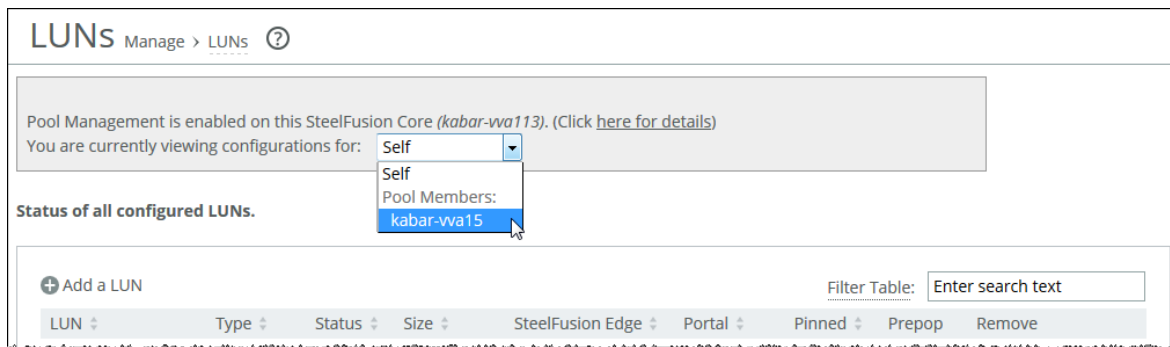
The pool management feature is implemented using the SteelHead REST API framework. For detailed information about enabling REST API access, see [“Configuring REST API Access” on page 85](#).

Pool management is available only in the UI (not in the CLI). This implementation is a one-directional on-demand query/action system. It does not maintain a persistent connection between pool members and the manager: that is, a member cannot proactively send data to its manager.

If you configure pool management, all Storage configuration and Reports pages include a link to pool management details. Click the link to jump to the Configure > Pool Management: Edit Pool page and the current settings.

Figure 2-41 shows a sample Storage configuration page with Pool Management enabled.

**Figure 2-41. Sample Storage Page - Pool Management Area**



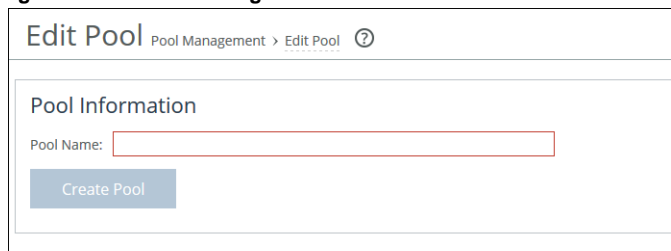
## To create a pool

### 1. Choose a manager.

You can choose any manager you like—the role requires minimal to no resource overhead.

### 2. Choose Configure > Pool Management: Edit Pool to display the Edit Pool page.

**Figure 2-42. Edit Pool Page**



### 3. Enter the pool name and click **Create Pool**.

### 4. Configure REST API access for each pool member. For information, see [“Configuring REST API Access” on page 85](#).

- On the pool manager's Edit Pool page, click **Add a Pool Member**.

**Figure 2-43. Edit Pool Page**

**Edit Pool** Pool Management > Edit Pool ?

**Pool Information**  
Pool name: **Test**  
**Dissolve Pool**

**Add a Pool Member**

Filter Table: Enter search text

Hostname or IP address:

API Access Code:

**Add Pool Member**

Hostname	Health	Total LUNs	Total Edges	Total Storage Mounted	Total Storage Mapped	Failover Status	Model	Remove
▶ kabar-va15	Healthy	1	1	2155023360	2155023360	Active Self	Virtual	

- Add the Hostname and API Access Code for each pool member.
  - Click **Add Pool Member**.
  - Continue adding members to the pool until you are done.
- The pool manager shows key statistics on the Edit Pool page about each member added.

### To manage a pool member

- Choose **Configure > Pool Management: Edit Pool** to display the Edit Pool page.  
You can also access the Edit Pool page by choosing a **Configure > Manage** or **Configure > Storage Array** page on the pool manager's web UI: for example, **Configure > Manage: LUNs**.
- Choose the pool member from the drop-down list.

**Figure 2-44. LUNs Page - Pool Management List**

**LUNs** Manage > LUNs ?

Pool Management is enabled on this SteelFusion Core (kabar-va113). (Click [here for details](#))

You are currently viewing configurations for: **Self**

**Status of all configured LUNs.**

**Add a LUN**

Filter Table: Enter search text

LUN	Type	Status	Size	SteelFusion Edge	Portal	Pinned	Prepop	Remove
-----	------	--------	------	------------------	--------	--------	--------	--------

The configuration for the selected pool member appears.

- Make the desired changes and click **Save**.

### To dissolve a pool (as a manager)

1. Choose Configure > Pool Management: Edit Pool.

You then have two options:

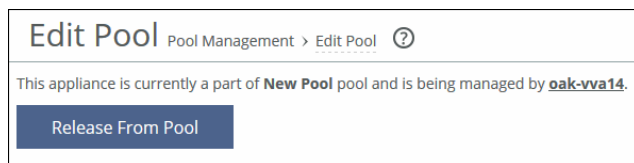
- Click the trash icon under Remove to remove a member.
- Click **Dissolve Pool** to dissolve the entire pool.

The pool manager sends a *reset pool* message to any removed pool member. However, if the member is down or the authorization code has been revoked, you must reset the pool membership on the member's UI to allow management again.

### To request release from a pool (as a member)

1. Choose Configure > Pool Management: Edit Pool.

Figure 2-45. Edit Pool Page



2. Click **Release From Pool**.

---

## Configuring REST API Access

You can access the Riverbed REST API framework to generate a REST API access code for use in Core pool management.

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 generate the REST API access code on the Core that you want to be a member of the pool, and supply the access code on the Core that is the manager of the pool. The access code generates a token that is valid for 60 minutes.

The access code generates a token that is valid for 60 minutes. The pool manager must request access to the member within that time period. The token is used for any REST API access request—for example, adding a member to the pool, changing a member's configuration, or viewing the member's status.

## To set up REST API access

1. On the Core pool member, choose **Configure > Pool Management: REST API Access** to display the REST API Access page.

**Figure 2-46. REST API Access Page**

Rest API Access Pool Management > Rest API Access ? Save

REST API Access Settings

☒ Enable REST API Access

Apply

Access Codes:

+ Add Access Code - Remove Selected

Access Code Description	Creator
kabar-vva109	admin

Related Topics: [Web Settings](#)

2. Select the Enable REST API Access check box and click **Apply**.
3. Click **Add Access Code**.
4. In the Description of Use field, enter a meaningful description for the access code.  
For example, include the hostname of the Core in the description.
5. Select Generate New Access Code and click **Add**.  
The access code description appears on the Access Code Description list, along with its creator.
6. To view the access code, click the Access Code Description.
7. To apply the access code to the pool member, click **Apply**.
8. Copy the access code to the clipboard for use in the Core pool manager.  
You use the hostname and the access code for the pool member to add the member to the pool. For the procedure, see [“To create a pool” on page 83](#).

## Disabling REST API Access

You can disable REST API at any time. To remove REST API access by a pool manager, you have two options:

- Clear the Enable REST API Access check box.
- From the Access Code Description list, select the access code description and click the Remove Selected tab.

When a pool manager sends REST API requests to a member, it uses the access code to get an access token from the member. The token is then used to send the actual REST API requests to the member. The access token is valid for 60 minutes.

If you delete a REST API access code, the access token is not invalidated immediately but expires after the 60 minute time period. For this reason, Riverbed recommends that you disable REST API access by clearing the Enable REST API Access check box if you have any security concerns.

## Best Practices

- Use one pool manager for all Cores in a logical cluster; make the cluster a single pool.
- Consider using different pool managers for different clusters, such as Finance, Legal, and so on.
- If hostnames are used for manager and members, DNS resolution must work. (Hostname entries can be added to manager and members through Configure > Networking: Host Settings to ensure that hostnames resolve correctly.)
- Be sure to give the pool a recognizable name. When reviewing a member's Edit Pool page, the user sees the pool name, and thereby knows what the pool is used for.
- Riverbed recommends that one failover peer be the manager of the other peer. It gives the manager more control over the peer than the failover connection does when both peers are up and running.

## Troubleshooting

Error messages provide information about why a manager cannot contact a member. This event occurs when trying to add the member to a pool, or after it has been added and another issue arises (for example, the member is not up, the REST API authentication code has been removed, and so on).

The system log for pool management is at the INFO level. Managers log messages before and after sending requests to members (but not the contents of the request). Members log requests from managers. You can filter for Pool Management to get relevant logs.

## Example Logs

On a manager, a failed call to add a member to a pool:

```
Mar 13 17:48:13 kabar-vva39 webasd[6791]: [web.ERR]: web: restd: Pool Management API call -
addMember. Result - failed. Error - This member is currently a part of a different pool.
```

A failed call to remove a member from a pool (the user had already manually removed the member):

```
Mar 13 18:21:31 kabar-vva39 webasd[5490]: [web.ERR]: web: restd: Pool Management API call -
removeMember. Result - failed. Error - Error attempting to get data - Request sent from 'kabar39',
but this box's pool name is 'Not Available'. Wrong pool.
```

A successful REST request to get the member's status:

```
Mar 13 18:23:10 kabar-vva39 webasd[5490]: [web.INFO]: web: restd: Pool Management API call -
memberStatus. Result - success
```

On a member, successful call examples:

```
Successful config query: Mar 17 15:48:40 kabar-vva15 restd[4218]: [web.INFO]: web: restd: Pool
Management API Request Received. Request - getMgmtLocalChildrenNames. Request successful.
```

```
Successful config query: Mar 17 15:48:40 kabar-vva15 restd[4218]: [web.INFO]: web: restd: Pool
Management API Request Received. Request - present. Request successful.
```



## CHAPTER 3    Modifying Host and Network Settings

This chapter describes how to configure network settings in the SteelFusion Core Management Console. It includes the following sections:

- [“Configuring Host Settings” on page 89](#)
- [“Configuring the Management Interfaces” on page 92](#)
- [“Configuring the Data Interfaces” on page 95](#)
- [“Configuring the Core for Jumbo Frames” on page 96](#)

---

### Configuring Host Settings

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

When you initially run the Setup Wizard, you set required network host settings for the Core. You can configure or modify the following settings:

- **Name** - Modify the hostname only if your deployment requires it.
- **DNS Settings** - Riverbed recommends that you use DNS resolution.
- **Hosts** - If you do not use DNS resolution, or if the host does not have a DNS entry, you can add hosts to the system.
- **Web/FTP Proxy** - Configure proxy addresses for web or FTP proxy access to the Core.

## To change the hostname

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

Figure 3-1. Host Settings Page

**Host Settings** Networking > Host Settings ? Save

**Name**

Hostname:

**DNS Settings**

Primary DNS Server:

Secondary DNS Server:

Tertiary DNS Server:

DNS Domain List:

2. Under Name, create or modify the hostname value.
3. Click **Apply** to apply the settings to the current configuration.
4. Click **Save** to save your settings permanently.

## To specify DNS settings

1. Choose Configure > Networking: Host Settings to display the Host Settings page.
2. Under DNS Settings, configure the settings using the controls described in this table.

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

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

### To add a new host

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

**Figure 3-2. Hosts Panel in the Host Settings Page**

The screenshot shows the 'Hosts' panel. At the top, there are two buttons: 'Add a New Host' (with a plus icon) and 'Remove Selected' (with a minus icon). Below these is a form with two input fields: 'IP Address:' and 'Hostname:'. An 'Add' button is positioned below the 'Hostname' field. Underneath the form is a table with two columns: 'IP Address' and 'Hostname'. The table contains one entry: '127.0.0.1' under IP Address and 'localhost' under Hostname.

2. Under Hosts, add and manage host configurations using the controls described in this table.

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

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

### To add a proxy

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

**Figure 3-3. Web/FTP Proxy Panel in the Host Settings Page**

The screenshot shows the 'Web/FTP Proxy' panel. It starts with a checkbox labeled 'Enable Web Proxy' which is checked. Below this is a section for 'Web/FTP Proxy:' with an input field and a 'Port:' field set to '1080'. There is another checkbox labeled 'Enable Authentication' which is unchecked. Below this are fields for 'User Name:' and 'Password:'. At the bottom, there is a dropdown menu for 'Authentication Type:' with 'Basic' selected.

2. Under Configure How this Appliance Connects to the Network, complete the configuration using the controls described in this table.

Control	Description
Enable Proxy Settings	<p>Select the check box to provide web proxy access to the Core. Specify the IP address and port for the web/FTP proxy.</p> <p>Enables the Core to use a web proxy to contact the Riverbed licensing portal and fetch licenses in a secure environment. You can optionally require user credentials to communicate with the proxy, and you can specify the method used to authenticate and negotiate user credentials.</p> <p>Web proxy access is disabled by default.</p>
Enable Authentication	<p>Optionally, select this option to enable authentication. Specify the following settings:</p> <ul style="list-style-type: none"> <li>• Username</li> <li>• Password</li> <li>• Authentication Type - Specify one of the following authentication types: <ul style="list-style-type: none"> <li>– Basic</li> <li>– NTLM</li> <li>– Digest</li> </ul> </li> </ul>

3. Click **Apply** to apply the settings to the current configuration.
4. Click **Save** to permanently save the settings.

## Configuring the Management Interfaces

You can view and modify settings for the primary interface in the Management Interfaces page.

When you initially ran the Setup Wizard, you set required settings for the data interfaces for the Core. Use the following groups of controls on this page only if modifications or additional configuration is required:

- **Primary Interface** - On the device, the primary interface is the port you connect to the LAN switch. The primary interface is the device management interface. You connect to the primary interface to use the web UI or the CLI.
- **Auxiliary Interface** - On the device, the auxiliary interface is an optional port you can use to connect the device to a non-Riverbed network management device. The IP address for the auxiliary interface must be on a subnet different from the primary interface subnet.
- **Main IPv4 Routing Table** - Displays a summary of the main routing table for the device. If necessary, you can add static routes that might be required for out-of-path deployments or particular device management subnets.

## To configure management interface settings

1. Choose **Configure > Networking: Management Interfaces** to display the Management Interfaces page.

**Figure 3-4. Management Interfaces Page**

**Management Interfaces** Networking > Management Interfaces ?

---

**Primary Interface**

☒ Enable Primary Interface

☐ Obtain IPv4 Address Automatically  
☐ Enable IPv4 Dynamic DNS

☒ Specify IPv4 Address Manually

IPv4 Address:   
 IPv4 Subnet Mask:   
 Default IPv4 Gateway:

Speed:  Negotiated: 1000Mb/s (auto)  
 Duplex:  Negotiated: full (auto)  
 MTU:  bytes

---

**Auxiliary Interface**

☒ Enable Aux Interface

☐ Obtain IPv4 Address Automatically  
☐ Enable IPv4 Dynamic DNS

☒ Specify IPv4 Address Manually

IPv4 Address:   
 IPv4 Subnet Mask:

Speed:  Negotiated: 100Mb/s (auto)  
 Duplex:  Negotiated: full (auto)  
 MTU:  bytes

---

**Apply**

---

**Main IPv4 Routing Table:**

☒ Add a New Route ☒ Remove Selected

Destination	Subnet Mask	Gateway	Interface	Status
default	0.0.0.0	10.1.15.1	aux	User Configured
10.1.15.0	255.255.255.0	0.0.0.0	aux	
10.12.0.0	255.255.0.0	0.0.0.0	primary	

2. Under **Primary Interface**, complete the configuration using the controls described in this table.

Control	Description
Enable Primary Interface	<p>Enables the primary interface.</p> <p><b>Caution:</b> The remaining Primary Interface settings are activated only when this box is selected.</p>
Obtain IPv4 Address Automatically	<p>Specify this option to automatically obtain the IPv4 address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it.</p> <p><b>Caution:</b> The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.</p>

Control	Description
Specify IPv4 Address Manually	Specify this option if you do not use a DHCP server to set the IP address. Specify the following settings: <ul style="list-style-type: none"> <li>• <b>IPv4 Address</b> - Specify an IPv4 address.</li> <li>• <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> <li>• <b>Default IPv4 Gateway</b> - Specify the primary gateway IP address. The primary gateway must be in the same network as the primary interface. You must set the primary gateway for in-path configurations.</li> </ul>
Speed and Duplex	<b>Speed</b> - Select a speed from the drop-down list. The default value is Auto. <b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them manually.
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.

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

Control	Description
Enable Auxiliary Interface	Enables the auxiliary interface. <b>Caution:</b> The remaining auxiliary interface settings are activated only when this box is selected.
Obtain IPv4 Address Automatically	Specify this option to automatically obtain the IPv4 address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it. <b>Note:</b> The primary and in-path interfaces can share the same subnet. The primary and auxiliary interfaces cannot share the same network subnet.
Specify IPv4 Address Manually	Specify this option if you do not use a DHCP server to set the IP address. Specify the following settings: <ul style="list-style-type: none"> <li>• <b>IPv4 Address</b> - Specify an IPv4 address.</li> <li>• <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> </ul>
Speed and Duplex	<b>Speed</b> - Select a speed from the drop-down list. The default value is Auto. <b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them manually.
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.

4. Click **Apply** to apply the settings to the current configuration.

5. Under Main IPv4 Routing Table, you can configure routing for interfaces that share the same subnet. You can add or remove routes from the list using the controls described in this table.

Control	Description
Add a New Route	Displays the controls for adding a new route.

Control	Description
Destination IPv4 Address	Specify the destination IP address for the out-of-path appliance or network management device.
IPv4 Subnet Mask	Specify the subnet mask.
Gateway IPv4 Address	Specify the IP address for the gateway.
Interface	From the drop-down list, select the interface.
Add	Adds the route to the table list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

- Repeat for each interface that requires routing.
- Click **Save** to save your changes permanently.

## Configuring the Data Interfaces

You can view and configure available data interfaces in the Data Interfaces page.

### To view and configure data interfaces

- Choose **Configure > Networking: Data Interfaces** to display the Data Interfaces page.

**Figure 3-5. Data Interfaces Page**

**Data Interfaces** Networking > Data Interfaces ?

**Data Network interfaces:**

Network Interface	IP Configuration
▼ eth0_0	10.12.5.201/16
▶ eth0_1	10.12.5.203/16
▶ eth0_2	10.12.5.202/16
▶ eth0_3	10.12.5.204/16

**eth0\_0**

☒ Enable Data Interface

☐ Obtain IPv4 Address Automatically  
☒ Specify IPv4 Address Manually

IPv4 Address:

IPv4 Subnet Mask:

Speed:  Negotiated: 1000Mb/s (auto)

Duplex:  Negotiated: full (auto)

MTU:  bytes

**Apply**

2. To enable, configure, or modify a data interface, click the name in the Data Network Interfaces column to display the following controls.

Control	Description
Enable Data Interface	Enables the current data interface. <b>Caution:</b> The remaining Primary Interface settings are activated only when this box is selected.
Obtain IPv4 Address Automatically	Specify this option to automatically obtain the IPv4 address from a DHCP server. A DHCP server must be available so that the system can request the IP address from it.
Specify IPv4 Address Manually	Specify this option if you do not use a DHCP server to set the IP address. Specify the following settings: <ul style="list-style-type: none"> <li>• <b>IPv4 Address</b> - Specify an IPv4 address.</li> <li>• <b>IPv4 Subnet Mask</b> - Specify a subnet mask.</li> </ul>
Speed and Duplex	<b>Speed</b> - Select a speed from the drop-down list. The default value is Auto. <b>Duplex</b> - Select Auto, Full, or Half from the drop-down list. The default value is Auto. If your network routers or switches do not automatically negotiate the speed and duplex, be sure to set them manually.
MTU	Specify the MTU value. The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.

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

## Configuring the Core for Jumbo Frames

If your network infrastructure supports jumbo frames, Riverbed recommends configuring the connection between the Core and the storage system as described here.

In addition to configuring Core for jumbo frames, you must also configure the storage system, as well as any switches, routers, or other network devices between the Core and the storage system.

### To configure the Core for jumbo frames

1. Log in to the SteelFusion Core Management Console.
2. Choose Configure > Networking: Management Interfaces to open the Management Interfaces page.
3. Under Primary Interface, set the configuration as follows:
  - Select the Enable Primary Interface check box.
  - Select the Specify IPv4 Address Manually option and configure appropriately.
  - For the MTU setting, specify 9000 bytes.
4. Click **Apply** to apply the settings to the current configuration.
5. Click **Save** to save your settings permanently.

## CHAPTER 4      **Configuring System Settings**

This chapter describes how to configure settings to manage the system. It includes the following sections:

- [“Creating Announcements” on page 97](#)
- [“Setting Alarm Parameters” on page 98](#)
- [“Configuring Date and Time” on page 103](#)
- [“Setting SNMP Parameters and Trap Receivers” on page 106](#)
- [“Creating SNMPv3 Users” on page 108](#)
- [“Configuring SNMP Authentication and Access Control” on page 110](#)
- [“Setting Up Email Notifications” on page 114](#)
- [“Configuring Logging” on page 114](#)
- [“Managing Configuration Files” on page 118](#)

---

### **Creating Announcements**

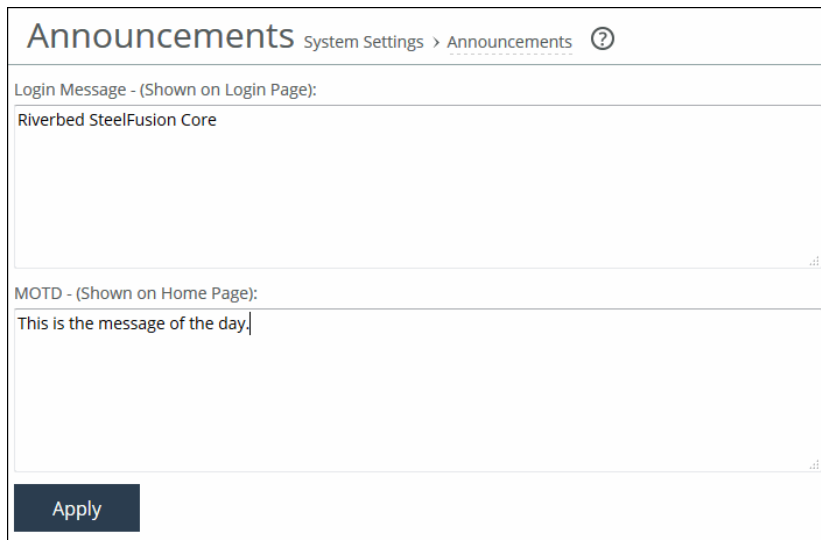
You can create or modify a login message or a message of the day in the Announcements page.

The login message appears in the SteelFusion Core Login page. The message of the day appears in the Home page and when you first log in to the CLI.

### To set an announcement

1. Choose Settings > System Settings: Announcements to display the Announcements page.

Figure 4-1. Announcements Page



Announcements System Settings > Announcements ?

Login Message - (Shown on Login Page):  
Riverbed SteelFusion Core

MOTD - (Shown on Home Page):  
This is the message of the day

Apply

2. Complete the configuration using the controls described in this table.

Control	Description
Login Message	In the text field, type a message to appear in the Login page.
MOTD	In the text field, type a message to appear in the Home page.

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

---

## Setting Alarm Parameters

Alarms have rising and reset thresholds. When an alarm reaches the rising threshold, it is activated; when it reaches the lowest (or reset) threshold, it is reset. After an alarm is triggered, it is not triggered again until it has fallen below the reset threshold.

You set the alarm thresholds in the Alarms page. Enabling this feature is optional.

## To set alarm parameters

1. Choose Settings > System Settings: Alarms to display the Alarms page.

Figure 4-2. Alarms Page

2. Complete the configuration using the controls described in this table.

Control	Description
Backup Integration	<p>Enables an alarm and sends an email notification if the backup-integration module encounters an error.</p> <p>By default, this alarm is enabled.</p>
Block-disk	<p>Enables an alarm and sends an email notification if the block-disk module encounters an error.</p> <p>By default, this alarm is enabled.</p> <p>This alarm applies only to Core-v implementations.</p>
Core Disaster Recovery	<p>Enables an alarm and sends an email notification if the system encounters any of the following issues with replication:</p> <ul style="list-style-type: none"> <li>• The Journal LUN size is not large enough to support the configured replica LUNs.</li> <li>• Replication latency to the secondary data center is over 150 ms.</li> <li>• Replication status for one or more LUNs is “suspended.”</li> <li>• Any type of corruption is detected on the Journal LUN for any of the replica LUNs.</li> <li>• No connections to the peer data center.</li> <li>• The Journal LUN is not found on the storage array because it was accidentally unmapped from the backend, or the connection to the storage array was lost.</li> </ul> <p>By default, this alarm is enabled.</p>
CPU Utilization	<p>Enables an alarm and sends an email notification if the average and peak threshold for the CPU utilization is exceeded.</p> <ul style="list-style-type: none"> <li>• <b>Rising Threshold</b> - Specify a whole number to represent a percent of CPU utilization.</li> <li>• <b>Reset Threshold</b> - Specify a whole number to represent a percent of CPU utilization.</li> </ul> <p>This alarm is enabled by default, with a rising threshold of 90 percent and a reset threshold of 70 percent.</p>

Control	Description
Disk Full	<p>Enables an alarm and sends an email notification if the disk space is full.</p> <p>Select one or more of the following system partitions:</p> <ul style="list-style-type: none"> <li>• Partition <code>"/boot"</code> Full</li> <li>• Partition <code>"/bootmgr"</code> Full</li> <li>• Partition <code>"/config"</code> Full</li> <li>• Partition <code>"/data"</code> Full</li> <li>• Partition <code>"/var"</code> Full</li> </ul> <p>By default, all Disk Full alarms are enabled.</p>
Edge Service	<p>Enables an alarm and sends an email notification if Core loses connection with one of its configured Edges.</p> <p>By default, this alarm is enabled.</p>
Hardware	<p>Enables an alarm and sends an email notification if one or more hardware failures occur.</p> <p>This alarm setting also enables you to select one or more types of hardware failure (fan error, memory error, and so on), including:</p> <ul style="list-style-type: none"> <li>• <b>Fan Error</b> - Enables an alarm and sends an email notification if a fan is failing or has failed and needs to be replaced. By default, this alarm is enabled.</li> <li>• <b>Flash Error</b> - Enables an alarm when the system detects an error with the flash drive hardware. By default, this alarm is enabled.</li> <li>• <b>IPMI</b> - Enables an alarm and sends an email notification if an Intelligent Platform Management Interface (IPMI) event is detected.</li> <li>• <b>Memory Error</b> - Enables an alarm and sends an email notification if a memory error is detected, for example, when a system memory stick fails.</li> <li>• <b>Other Hardware Error</b> - This alarm indicates that the system has detected a problem with the hardware. The alarm clears when you add the necessary hardware, remove the nonqualified hardware, or resolve other hardware issues. The following issues trigger the hardware error alarm: <ul style="list-style-type: none"> <li>• The appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration.</li> <li>• The appliance is using a dual in-line memory module (DIMM), a hard disk, or a NIC that is not qualified by Riverbed.</li> <li>• DIMMs are plugged into the appliance but the system cannot recognize them because the DIMM modules are in the wrong slot. You must plug DIMM modules into the black slots first and then use the blue slots when all of the black slots are in use.</li> <li>• A DIMM module is broken and you must replace it.</li> <li>• Other hardware issues.</li> </ul> </li> </ul> <p>By default, all Hardware alarms are enabled.</p> <ul style="list-style-type: none"> <li>• <b>Power Supply</b> - Enables an alarm and sends an email notification if an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted. By default, this alarm is enabled.</li> </ul>
High-Availability	<p>Enables an alarm and sends an email notification if one of the Cores in a high-availability environment fails.</p> <p>By default, this alarm is enabled.</p>
iSCSI Service	<p>Enables an alarm and sends an email notification if the iSCSI configuration fails.</p> <p>By default, this alarm is enabled.</p>

Control	Description
Licensing	<p>Enables an alarm and sends an email notification if the appliance is unlicensed, if there is an issue with the autolicense, the licenses have expired, the licenses are about to expire, or the model is unlicensed.</p> <p>By default, all Licensing alarms are enabled.</p>
Link Duplex	<p>Enables an alarm and sends an email notification when an interface was not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex alarm.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable an alarm, choose Settings &gt; System Settings: Alarms and select or clear the check box next to the link name.</p>
Link I/O Errors	<p>Enables an alarm and sends an email notification when the link error rate exceeds 0.1% while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured LAN connection experiences very few errors.</p> <p>The alarm clears when the rate drops below 0.05%.</p> <p>You can change the default alarm thresholds by entering the <b>alarm link_io_errors err-threshold &lt;threshold-value&gt;</b> CLI command at the system prompt. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p> <p>By default, this alarm is enabled.</p> <p>You can enable or disable the alarm for a specific interface: for example, you can disable the alarm for a link after deciding to tolerate the errors. To enable or disable an alarm, choose Settings &gt; System Settings: Alarms and select or clear the check box next to the link name.</p>
Link State	<p>Triggers an alarm if network interface link errors are detected. If you receive this alarm, check the status of the interface to begin diagnosing the problem.</p>
LUN Status	<p>Enables an alarm and sends an email notification if the connection to the LUN fails or there is an issue with LUN resizing.</p> <p>By default, this alarm is enabled.</p>
Memory Paging	<p>Triggers an alarm if the Core detects extended memory paging activity.</p> <p>If 100 pages are swapped every couple of hours, the appliance is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support.</p> <p>By default, this alarm is enabled.</p>
Process Dump Creation Error	<p>Enables an alarm and sends an email notification if the system detects an error while trying to create a process dump.</p> <p>This alarm indicates an abnormal condition in which the system cannot collect the core file after three retries. It can be caused when the /var directory is reaching capacity or by other conditions. When the alarm is raised, the directory is blacklisted.</p> <p>By default, this alarm is enabled.</p>
Secure Vault	<p>Enables an alarm and sends an email notification if the system encounters a problem with the secure vault:</p> <p><b>Secure Vault Locked</b> - Indicates that the secure vault is locked. To optimize SSL connections or to use RiOS data store encryption, the secure vault must be unlocked. Go to Settings &gt; Security: Secure Vault and unlock the secure vault.</p> <p>For details, see <a href="#">“Unlocking the Secure Vault” on page 135</a>.</p>

Control	Description
Snapshot	<p>Enables an alarm and sends an email notification if the connection to any of the snapshot storage arrays fails.</p> <p>By default, this alarm is enabled.</p>
SSL	<p>Enables an alarm if an error is detected in your SSL configuration. By default, this alarm is enabled.</p>
SteelFusion Core configuration status	<p>Enables an alarm and sends an email notification if the Core configuration has been reverted to a previous version.</p> <p>By default, this alarm is enabled.</p>
Temperature	<ul style="list-style-type: none"> <li>• <b>Critical Temperature</b> - Enables an alarm and sends an email notification if the CPU temperature exceeds the rising threshold. When the CPU returns to the reset threshold, the critical alarm is cleared. The default value for the rising threshold temperature is 70°C; the default reset threshold temperature is 67°C.</li> <li>• <b>Warning Temperature</b> - Enables an alarm and sends an email notification if the CPU temperature approaches the rising threshold. When the CPU returns to the reset threshold, the warning alarm is cleared.</li> <li>• <b>Rising Threshold</b> - Specifies the rising threshold. The alarm activates when the temperature exceeds the rising threshold. The default value is 70 percent.</li> <li>• <b>Reset Threshold</b> - Specifies the reset threshold. The alarm clears when the temperature falls below the reset threshold. The default value is 67 percent.</li> </ul> <p>After the alarm triggers, it cannot trigger again until after the temperature falls below the reset threshold and then exceeds the rising threshold again.</p>

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

# Configuring Date and Time

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

You can either set the system date and time by entering it manually or assigning an NTP server to the Core.

## To set the date and time manually

1. Choose Settings > System Settings: Date and Time to display the Date and Time page.

**Figure 4-3. Date and Time Page**

Date and Time
System Settings > Date and Time
?

### Date and Time

Time Zone: US/Pacific

Set Time Manually

Change Date: 2014/09/20
Change Time: 22:31:23

Use NTP Time Synchronization

Apply

Requested NTP Servers:

Add a New NTP Server
Remove Selected

Server	Version	Enabled	Key ID
0.riverbed.pool.ntp.org	4	Enabled	
1.riverbed.pool.ntp.org	4	Enabled	
2.riverbed.pool.ntp.org	4	Enabled	
3.riverbed.pool.ntp.org	4	Enabled	
208.70.196.25	4	Enabled	

Connected NTP Servers:

Active	Server	Auth Status	Key ID	Ref ID
Yes	247.conarusp.net	None		216.218.254.202
	clock.team-cymru.org	None		140.142.16.34
	ftp.riverbed.com	None		10.16.0.15
	lithium.constant.com	None		206.246.122.250
	mirror	None		130.173.91.58

NTP Authentication Keys:

Add a New NTP Authentication Key
Remove Selected

Key ID	Key Type	Encrypted Text
No NTP Authentication Keys.		

## 2. Under Date and Time, click **Set Time Manually**.

Complete the configuration as described in this table.

Control	Description
Time Zone	Select a time zone from the drop-down list. The default value is GMT. <b>Note:</b> If you change the time zone, log messages retain the previous time zone until you reboot.
Change Date	Specify the date in this format: yyyy/mm/dd.
Change Time	Specify military time in this format: hh:mm:ss.

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

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

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

## Current NTP Server Status

NTP server state information appears in these server tables:

- **Requested NTP server table** - Displays all of the configured NTP server addresses.
- **Connected NTP server table** - Displays all of the servers to which the Core is actually connected.

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

This information appears after an NTP server name:

- Authentication information; unauthenticated appears after the server name when it is not using authentication.
- When the system has no NTP information about the current server, nothing appears.

## NTP Servers

By default, the Core uses the Riverbed-provided NTP server IP address 208.70.196.25 and these public NTP servers:

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

Riverbed recommends synchronizing the Core to an NTP server of your choice.

### To add an NTP server

1. Choose Settings > System Settings: Date and Time to display the Date and Time page.
2. Under Requested NTP servers, complete the configuration as described in this table.

Control	Description
Add a New NTP Server	Displays the controls to add a server.
Hostname or IP Address	Specify the hostname or IP address for the NTP server. You can connect to an NTP public server pool. For example, 0.riverbed.pool.ntp.org.  When you add an NTP server pool, the server is selected from a pool of time servers.
Version	Select the NTP server version from the drop-down list: 3 or 4.
Enabled/Disabled	Select Enabled from the drop-down list to connect to the NTP server. Select Disabled from the drop-down list to disconnect from the NTP server.
Key ID	Specify the MD5 or SH1 key identifier to use to authenticate the NTP server. The valid range is from 1 to 65534. The key ID must appear on the trusted keys list.
Add	Adds the NTP server to the server list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

## NTP Authentication

NTP authentication verifies the identity of the NTP server sending timing information to the Core. RiOS 8.5 and later supports MD5-based Message-Digest Algorithm symmetric keys and Secure Hash Algorithm (SHA1) for NTP authentication. MD5 is a widely used cryptographic hash function that produces a 128-bit (16-byte) hash value. SHA1 is a set of related cryptographic hash functions. SHA1 is considered to be the successor to MD5.

NTP authentication is optional.

Configuring NTP authentication involves these tasks that you can perform in any order:

- Configure a key ID and a secret pair.
- Configure the key type.
- Configure the NTP server with the key ID.

### NTP Authentication Keys

NTP authentication uses a key and a shared secret to verify the identity of the NTP server sending timing information to the Core. The Core encrypts the shared secret text using MD5 or SHA1, and uses the authentication key to access the secret.

### To add an NTP authentication key

1. Under NTP Authentication Keys, choose Settings > System Settings: Date and Time to display the Date and Time page.

2. Complete the configuration as described in this table.

Control	Description
Add a New NTP Authentication Key	Displays the controls to add an authentication key to the key list. Both trusted and untrusted keys appear on the list.
Key ID	Optionally, specify the secret MD5 or SHA1 key identifier for the NTP server. The valid range is from 1 to 65534.
Key Type	Select the authentication key type: MD5 or SHA1.
Secret	<p>Specify the shared secret. You must configure the same shared secret for both the NTP server and the NTP client.</p> <p>The MD5 shared secret:</p> <ul style="list-style-type: none"> <li>is limited to 16 alphanumeric characters or less, or exactly 40 characters hexadecimal.</li> <li>cannot include spaces or pound signs (#)</li> <li>cannot be empty</li> <li>is case sensitive</li> </ul> <p>The SHA1 shared secret:</p> <ul style="list-style-type: none"> <li>is limited to exactly 40 characters hexadecimal</li> <li>cannot include spaces or pound signs (#)</li> <li>cannot be empty</li> <li>is case sensitive</li> </ul> <p>The secret appears in the key list as its MD5 or SHA1 hash value.</p>
Add	Adds the authentication key to the trusted keys list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

### ***NTP Key Information***

NTP keys appear in a list that includes the key ID, type, secret (displays as the MD5 or SHA1 hash value), and whether the system trusts the key for authentication.

You can only remove a key from the trust list using the CLI command **ntp authentication trustedkeys**. For details, see the *Riverbed Command-Line Interface Reference Manual*.

## **Setting SNMP Parameters and Trap Receivers**

You can set the SNMP server settings and set up SNMP traps in the SNMP Basic page.

SNMP traps are messages sent by an SNMP agent that indicate the occurrence of events. The text of the Core MIB is available on the Support page by choosing Help > Customer Service: Support.

**Note:** By default, SNMP is disabled.

## To set general SNMP parameters

1. Choose Settings > System Settings: SNMP Basic to display the SNMP Basic page.

**Figure 4-4. SNMP Basic Page**

The screenshot shows the 'SNMP Basic' configuration page. The 'SNMP Server Settings' section includes a checkbox for 'Enable SNMP Traps' which is checked. Below it are text input fields for 'System Contact', 'System Location', and 'Read-Only Community String' (containing 'riverbed'). An 'Apply' button is located below these fields. The 'Trap Receivers' section features a table with columns: Receiver, Version, Port, Community / User, and Enabled. One receiver, 'chief-cs57', is listed with version 'v1', port '162', and community 'community: riverbed', with an 'Enabled' status. Below the table is a 'SNMP Trap Test' section with a 'Run' button.

2. Under SNMP Server Settings, complete the configuration using the controls described in this table.

Control	Description
Enable SNMP Traps	Click to enable traps.
System Contact	Specify the username for the SNMP contact.
System Location	Specify the physical location of the SNMP system.
Read-Only Community String	Specify a password-like string to identify the read-only community: for example, public. This community string overrides any VACM settings. <b>Note:</b> This string cannot contain a pound sign (#).

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

## To add or remove a trap receiver

1. Under Trap Receivers, complete the configuration using the controls described in this table.

Control	Description
Add a New Trap Receiver	Displays the controls to add a new trap receiver.
Receiver	Specify the destination IP address for the SNMP trap.
Destination Port	Specify the destination port.
Receiver Type	Select v1, v2c, or v3 (User-Based Security Model) to specify the receiver version.

Control	Description
Community	For v1 or v2 trap receivers, specify the SNMP community name: for example, public or private v3 trap receivers need a remote user with an authentication protocol, as well as a password and security level.
Enable Receiver	Enables the trap receiver.
Add	Adds a new trap receiver to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

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

### To test an SNMP trap

1. Choose Settings > System Settings: SNMP Basic to display the SNMP Basic page.
2. Under SNMP Trap Test, click **Run**.

---

## Creating SNMPv3 Users

You create SNMP v3 users in the SNMPv3 page.

SNMPv3 provides additional authentication and access control for message security. For example, you can verify the identity of the SNMP entity (manager or agent) sending the message.

Using SNMPv3 is more secure than SNMPv1 or v2; however, it requires more configuration steps to provide the additional security features.

### To create SNMP users

1. Create the SNMP server users. Users can be authenticated using either a password or a key.
2. Configure SNMP server views to define which part of the SNMP MIB tree is visible.
3. Configure SNMP server groups.  
SNMP server groups map users to views, enabling you to control who can view what SNMP information.
4. Configure the SNMP server access policies that contain a set of rules defining access rights. Based on these rules, the entity decides how to process a given request.

## To create users for SNMPv3

1. Choose Settings > System Settings: SNMP v3 to display the SNMP v3 page.

Figure 4-5. SNMP v3 Page

2. Under Users, complete the configuration using the controls described in this table.

Control	Description
Add a New User	Displays the controls to add a new user.
User Name	Specify the username.
Authentication Protocol	Select an authentication method from the drop-down list: <ul style="list-style-type: none"> <li>• <b>MD5</b> - Specifies the message-digest 5 algorithm, a widely used cryptographic hash function with a 128-bit hash value. This is the default value.</li> <li>• <b>SHA</b> - Specifies the secure hash algorithm, a set of related cryptographic hash functions. SHA is considered to be the successor to MD5.</li> </ul>
Authentication	Optionally, select either Supply a Password or Supply a Key to use while authenticating users.
Password/Password Confirm	Specify a password. The password must have a minimum of eight characters. Retype the password in the Password Confirm text field.
MD5 Key	(Appears only when you select Supply a Key.) Specify a unique authentication key. The key is an MD5 or SHA digest created using md5sum or sha1sum.
Use Privacy Option	Optionally, select this setting to configure privacy using encryption. Configure the following: <ul style="list-style-type: none"> <li>• <b>Privacy Protocol</b> - Specify one of the following encryption standards: AES or DES.</li> <li>• <b>Privacy</b> - Specify one of the following methods: <ul style="list-style-type: none"> <li>– <b>Same as Authentication Key.</b></li> <li>– <b>Supply a Password</b> - If you select this option, you are prompted to set and confirm a password.</li> <li>– <b>Supply a Key</b> - If you select this option, you are prompted to supply a key.</li> </ul> </li> </ul>
Add	Adds the user.

Control	Description
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

- Click **Add** to add the user.

## Configuring SNMP Authentication and Access Control

You configure SNMP authentication and access control on the SNMP ACLs page.

The features on this page apply to SNMPv1, v2c, and v3 unless noted otherwise:

- **Security Names** - Specify an individual user (v1 or v2c only).
- **Groups** - Specify a security name, security model by a group, and referred to by a group name.
- **Views** - Create a custom view using the VACM that controls who can access which MIB objects under agent management by including or excluding specific OIDs: for example, some users have access to critical read-write control data, while other users have access only to read-only data.
- **Access Policies** - Defines who gets access to which type of information. An access policy is composed of Group Name, Security Level, and Read View name.

### To set secure usernames

- Choose Settings > System Settings: SNMP ACLs to display the SNMP ACLs page.

**Figure 4-6. SNMP ACLs Page**

- Under Security Names, complete the configuration using the controls described in this table.

Control	Description
Add a New Security Name	Displays the controls to add a security name.

Control	Description
Security Name	<p>(v1 and v2c only) Specify a name to identify a requestor allowed to issue gets and sets. The security name might make changes to the VACM security name configuration.</p> <p><b>Note:</b> This control does not apply to SNMPv3 queries. To restrict v3 USM users from polling from a particular subnet, use the ACL feature, located on the Settings &gt; System Settings: SNMP ACLs page.</p> <p><b>Note:</b> Traps for v1 and v2c are independent of the security name.</p>
Community String	<p>Specify the password-like community string to control access. Use a combination of uppercase, lowercase, and numerical characters to reduce the chance of unauthorized access to the Core.</p> <p>Community strings do not allow printable 7-bit ASCII characters, except for white spaces. Also, the community strings cannot begin with '#' and '-'. If you specify a read-only community string (located in the SNMP Basic page under Settings &gt; System Settings), it takes precedence over this community name and allows users to access the entire MIB tree from any source host. If this is not desired, delete the read-only community string.</p> <p>To create multiple SNMP community strings on a Core, leave the default public community string and then create a second read-only community string with a different security name. Or, you can delete the default public string and create two new SNMP ACLs with unique names.</p> <p><b>Note:</b> If you specify a read-only community string (located on the SNMP Basic page under SNMP Server Settings), it takes precedence over this community name and enables users to access the entire MIB tree from any source host. If this is not desired, delete the read-only community string.</p>
Source IP Address and Mask Bits	Specify the host IP address and mask bits to which you permit access using the security name and community string.
Add	Adds the security name.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

## To set secure groups

1. Choose Settings > System Settings: SNMP ACLs to display the SNMP ACLs page.

**Figure 4-7. SNMP ACLs Page - Groups**

A group is one or more entries of the form security-model:security-name.

**Groups:**

Specify the group name and select the security models. For v1 and v2c security models, select the security name. For usm security models, select the user name.

Group Name:

Security Model and Name Pairs: v1

Group Name	Security Models, Names
No Groups.	

2. Under Groups, complete the configuration using the controls described in this table.

Control	Description
Add a New Group	Displays the controls to add a new group.
Group Name	Specify a group name.
Security Model and Name Pairs	Click the <b>+</b> button and select a security model from the drop-down list: <ul style="list-style-type: none"> <li><b>v1</b> or <b>v2c</b> - Displays another drop-down menu; select a security name.</li> <li><b>usm</b> - Displays another drop-down menu; select a user.</li> </ul> To add another security model and name pair, click the <b>+</b> button. To remove a pair, click the <b>—</b> button.
Add	Adds the group name and security model and name pairs.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

### To set secure views

1. Choose Settings > System Settings: SNMP ACLs to display the SNMP ACLs page.

**Figure 4-8. SNMP ACLs Page - Views**

Add OIDs that should be included or excluded from this view

**Views:**

☒ Add a New View ☐ Remove Selected

View Name:

Includes:

(one .x.y.z per line)

Excludes:

(one .x.y.z per line)

View Name	Includes	Excludes
No Views.		

2. Under Views, complete the configuration using the controls described in this table.

Control	Description
Add a New View	Displays the controls to add a new view.
View Name	Specify a descriptive view name to facilitate administration.

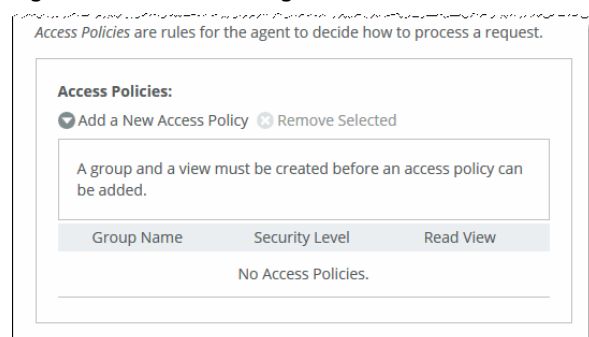
Control	Description
Includes	Specify the object identifiers (OIDs) to include in the view, separated by commas: for example, .1.3.6.1.2.1.1. By default, the view excludes all OIDs.  You can specify .iso or any subtree or subtree branch.  You can specify an OID number or use its string form: for example, .iso.org.dod.internet.private.enterprises.rbt.products.steelhead.system.model.
Excludes	Specify the OIDs to exclude in the view, separated by commas. By default, the view excludes all OIDs.
Add	Adds the view.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

### To add an access policy

1. Choose Settings > System Settings: SNMP ACLs to display the SNMP ACLs page.

**Note:** To be able to add an access policy, first you must create a new group and a view. For details, see [“To set secure groups” on page 111](#) and [“To set secure views” on page 112](#).

**Figure 4-9. SNMP ACLs Page - Access Policies**



2. Under Access Policies, complete the configuration using the controls described in this table.

Control	Description
Add a New Access Policy	Displays the controls to add a new access policy.
Group Name	Select a group name from the drop-down list.
Security Level	Determines whether a single atomic message exchange is authenticated. Select one of the following options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>No Auth</b> - Does not authenticate packets and does not use privacy. This is the default setting.</li> <li>• <b>Auth</b> - Authenticates packets but does not use privacy.</li> </ul> <b>Note:</b> A security level applies to a group, not to an individual user.
Read View	Select a view from the drop-down list.
Add	Adds the policy to the policy list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

## Setting Up Email Notifications

You set email notification parameters for events and failures in the Email page.

By default, no email addresses are specified for event and failure notification.

### To set event and failure email notification

1. Choose Settings > System Settings: Email to display the Email page.

**Figure 4-10. Email Page**

2. Under Email Settings, complete the configuration using the controls described in this table.

Control	Description
SMTP Server	Specify a valid SMTP server. External DNS and external access for SMTP traffic are required for this feature to function.
SMTP Port	Specify the port on the SMTP server.
Report Events via Email	Specify this option to report events to the specified email address. Specify a space-separated list of email addresses to which to send notification messages. To complete SNMP settings, see <a href="#">“Setting SNMP Parameters and Trap Receivers” on page 106</a> .
Report Failures via Email	Specify this option to report failures to the specified email address. Specify a space-separated list of email addresses to which to send notification messages.
Report Failures to Technical Support	Specify this option to report failures to Riverbed Support. Riverbed recommends that you activate this feature so that problems are promptly corrected.

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

## Configuring Logging

This section describes how to modify local logging and how to set remote logging for the Core. It includes the following sections:

- [“Setting Up System Logging” on page 115](#)
- [“Configuring Remote Log Servers” on page 116](#)
- [“Filtering Logs by Application or Process” on page 117](#)

## Setting Up System Logging

You configure system logging at the top of the Logging page.

### To set up logging

1. Choose Settings > System Settings: Logging to display the Logging page.

**Figure 4-11. Logging Page**

2. Under Logging Configuration, complete the configuration using the controls described in this table.

Control	Description
Minimum Severity	<p>Select the minimum severity level for the system log messages. The log contains all messages with this severity level or higher.</p> <p>Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Emergency</b> - The system is unusable.</li> <li>• <b>Alert</b> - Action must be taken immediately.</li> <li>• <b>Critical</b> - Conditions that affect the functionality of the Core.</li> <li>• <b>Error</b> - Conditions that probably affect the functionality of the Core.</li> <li>• <b>Warning</b> - Conditions that could affect the functionality of the Core, such as authentication failures.</li> <li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change.</li> <li>• <b>Info</b> - Messages that provide general information about system operations.</li> </ul> <p><b>Note:</b> This control applies to the system log only. It does not apply to the user log.</p>
Maximum Number of Log Files	Specify the maximum number of logs to store. The default value is 10.
Lines Per Log Page	Specify the number of lines per log page. The default value is 100.

Control	Description
Rotate Based On	<p>The automatic rotation of system logs deletes your oldest log file, labeled as Archived log #10, pushes the current log to Archived log # 1, and starts a new current-day log file.</p> <p>By default, the system rotates each log file every 24 hours or if the file size reaches 16 MBs (uncompressed). You can change this setting to rotate every week or month and you can rotate the files based on file size.</p> <p>Specify one of the following rotation options:</p> <ul style="list-style-type: none"> <li>• <b>Time</b> - Select Day, Week, or Month from the drop-down list.</li> <li>• <b>Disk Space</b> - Specify how much disk space, in megabytes, the log uses before it rotates.</li> </ul>

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

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

## Configuring Remote Log Servers

You configure remote log servers on the Logging page.

### To add or remove a log server

1. Choose Settings > System Settings: Logging to display the Logging page.

**Figure 4-12. Logging Page - Remote Log Servers**

2. Under Remote Log Servers, complete the configuration using the controls described in this table.

Control	Description
Add a New Log Server	Displays the controls for configuring new log servers.
Server IP	Specify the server IP address.

Control	Description
Minimum Severity	<p>Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Emergency</b> - The system is unusable.</li> <li>• <b>Alert</b> - Action must be taken immediately.</li> <li>• <b>Critical</b> - Conditions that affect the functionality of the Core.</li> <li>• <b>Error</b> - Conditions that probably affect the functionality of the Core.</li> <li>• <b>Warning</b> - Conditions that could affect the functionality of the Core, such as authentication failures.</li> <li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change.</li> <li>• <b>Info</b> - Messages that provide general information about system operations.</li> </ul>
Add	Adds the server to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

- To rotate the logs immediately, under Log Actions click **Rotate Logs**.

After the logs are rotated, the following message appears:

logs have been successfully rotated

You can also schedule a log rotation based on time or the amount of disk space the log uses, described in Step 2 in [“Setting Up System Logging” on page 115](#).

## Filtering Logs by Application or Process

You can filter a log by one or more applications or one or more processes. This is particularly useful when capturing data at a lower severity level, at which the Core might not be able to sustain the flow of logging data the service is committing to disk.

### To filter a log

- Choose Settings > System Settings: Logging to display the Logging page.

**Figure 4-13. Logging Page - Per-Process Logging**

**Per-Process Logging:**

☒ Add a New Process Logging Filter ☐ Remove Selected

Process:

Minimum Severity:  (applies only to system log)

<input type="checkbox"/>	Description	Process	Minimum Severity
<input type="checkbox"/>	Alarm Manager	alarmd	Emergency

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

Control	Description
Add a New Process Logging Filter	Displays the controls for adding a process level logging filter.
Process	<p>Select a process to include in the log from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>alarmd</b> - Alarm Manager</li> <li>• <b>cmcfc</b> - SCC Autoregistration Utility</li> <li>• <b>rgpd</b> - SCC Connection Manager</li> <li>• <b>rgp</b> - SCC Connector</li> <li>• <b>cli</b> - Command-Line Interface</li> <li>• <b>mgmtd</b> - Device Control and Management, which directs the entire device management system. It handles messages passing between various management daemons, managing system configuration and general application of system configuration on the hardware underneath through the Hardware Abstraction Layer Daemon (HALD).</li> <li>• <b>hald</b> - Hardware Abstraction Layer Daemon, which handles access to the hardware.</li> <li>• <b>pm</b> - Process Manager, which handles launching of internal system daemons and keeps them running.</li> <li>• <b>sched</b> - Process Scheduler, which handles one-time scheduled events.</li> <li>• <b>statsd</b> - Statistics Collector, which handles queries and storage of system statistics.</li> <li>• <b>wdt</b> - Watchdog Timer, the motherboard watchdog daemon.</li> <li>• <b>webasd</b> - Web Application Process, which handles the web user interface.</li> </ul>
Minimum Severity	<p>Select the minimum severity level for the log messages. The log contains all messages with this severity level or higher. Select one of the following levels from the drop-down list:</p> <ul style="list-style-type: none"> <li>• <b>Emergency</b> - The system is unusable.</li> <li>• <b>Alert</b> - Action must be taken immediately.</li> <li>• <b>Critical</b> - Conditions that affect the functionality of the Core.</li> <li>• <b>Error</b> - Conditions that probably affect the functionality of the Core.</li> <li>• <b>Warning</b> - Conditions that could affect the functionality of the Core, such as authentication failures.</li> <li>• <b>Notice</b> - Normal but significant conditions, such as a configuration change.</li> <li>• <b>Info</b> - Messages that provide general information about system operations.</li> </ul>
Add	Adds the filter to the list, after which the process logs at the selected severity and higher levels.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> to remove the filter.

## Managing Configuration Files

You save, activate, and import configurations in the Configurations page.

Each Core has an active, running configuration and a written, saved configuration. When you apply your settings in the Core, the values are applied to the active running configuration, but the values are not written to disk and saved permanently.

When you save your configuration settings, the values are written to disk and saved permanently. The settings take effect after you restart the Core service.

Each time you save your configuration settings, they are written to the current running configuration and a backup is created. For example, if the running configuration is myconfig and you save it, myconfig is backed up to myconfig.bak and myconfig is overwritten with the current configuration settings.

The Configuration Manager is a utility that enables you to save configurations as backups or to activate configuration backups.

---

**Note:** Some configuration settings require that you to restart the Core service for the settings to take effect. For details about restarting the Core service, see ["Starting, Stopping, and Restarting the Service" on page 141](#).

---

## To manage configurations

1. Choose Settings > System Settings: Configurations to display the Configurations page.

**Figure 4-14. Configurations Page**

**Configurations** System Settings > Configurations ?

Current Configuration: working  
[View Running Config](#)  
Save Revert

Save Current Configuration  
New Configuration Name:   
Save As

**Configurations:**  
+ Import a New Configuration - Remove Selected

<input type="checkbox"/> Configuration	Date
<input type="checkbox"/> cold	2013/07/15 16:12:53
<input type="checkbox"/> initial	2013/07/15 16:12:53
<input type="checkbox"/> initial.bak	2013/07/15 16:11:39
<input type="checkbox"/> working (active)	2014/09/17 20:48:18
<input type="checkbox"/> working.bak	2014/09/16 12:48:18

Change Active Configuration  
working

- Under Current Configuration: <name>, use the following controls to view, save, or revert configurations.

Control	Description
View Running Config	Displays the running configuration settings in a new browser window.
Save	Click to save settings that have been applied to the running configuration.
Revert	Reverts your settings to the running configuration.

- Under Save Current Configuration, specify a new filename to save settings that have been applied to the running configuration as a new file, and then click **Save**.
- To import a configuration from another device, click **Import a New Configuration** and use the controls described in this table.

Control	Description
IP/Hostname	Specify the IP address or hostname of the Core from which you want to import the configuration.
Remote Admin Password	Specify the administrator password for the remote Core.
Remote Config Name	Specify the name of the configuration you want to import from the remote Core.
New Config Name	Specify a new, local configuration name.
Import Shared Data Only	This value is enabled by default. Copies only the in-path and out-of-path interface, protocols, CLI, web, statistics, NTP, SNMP, and alarm settings. The system does not automatically copy the failover, SNMP (contact and location), log, and network settings.
Import	Imports the configuration. The imported configuration appears in the Configuration list but does not become the active configuration until you click <b>Activate</b> .
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

**Note:** Click the configuration name to display the configuration settings in a new browser window.

- To change the currently active configuration, select another configuration from the drop-down list under Change Active Configuration, and click **Activate**.

**Caution:** You must restart the SteelFusion Core Management Console for a new configuration to take effect.

## CHAPTER 5      Configuring Security Settings

This chapter describes how to configure security settings for the system. It includes the following sections:

- [“Configuring General Security Settings” on page 121](#)
- [“Managing User Permissions” on page 122](#)
- [“Managing Password Policy” on page 126](#)
- [“Configuring RADIUS Server Authentication” on page 130](#)
- [“Configuring TACACS+ Server Authentication” on page 132](#)
- [“Unlocking the Secure Vault” on page 135](#)
- [“Configuring Web Settings” on page 136](#)

---

### Configuring General Security Settings

You can prioritize local, RADIUS, and TACACS+ authentication methods for the system and set the authorization policy and default user for RADIUS and TACACS+ authorization systems in the Settings > Security: General Settings page.

---

**Caution:** Make sure to put the authentication methods in the order in which you want authentication to occur. If authorization fails on the first method, the next method is attempted, and so on, until all of the methods are attempted.

---

---

**Note:** To set TACACS+ authorization levels (admin or read-only) to allow certain members of a group to log in, add the following attribute to users on the TACACS+ server:

```
service = rbt-exec {  
    local-user-name = "monitor"  
}
```

where you replace *monitor* with *admin* for write access.

---

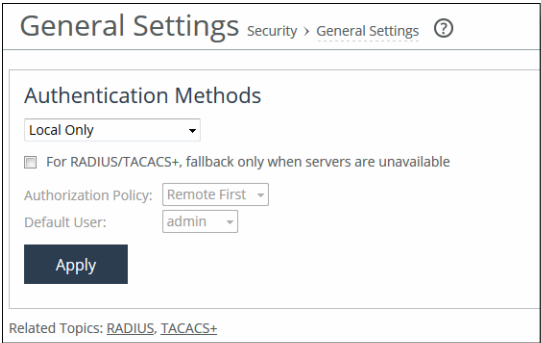
For details about setting up RADIUS and TACACS+ servers, see the *SteelHead Deployment Guide*.

Beta Draft

To set general security settings

- 1. Choose Settings > Security: General Settings to display the General Settings page.

Figure 5-1. General Settings Page



- 2. Under Authentication Methods, complete the configuration using the controls described in this table.

Control	Description
Authentication Methods	Specifies the authentication method. Select an authentication method from the drop-down list. The methods are listed in the order in which they occur. If authorization fails on the first method, the next method is attempted, and so on, until all of the methods have been attempted.
For RADIUS/TACACS+, fallback only when servers are unavailable.	Specifies that the Core falls back to a RADIUS or TACACS+ server only when all other servers do not respond. This is the default setting.  When this feature is disabled, the Core does not fall back to the RADIUS or TACACS+ servers. If it exhausts the other servers and does not get a response, it returns a server failure.
Authorization Policy	Appears only for some Authentication Methods. Optionally, select one of these policies from the drop-down list: <ul style="list-style-type: none"><li>• <b>Remote First</b> - Check the remote server first for an authentication policy, and only check locally if the remote server does not have one set. This is the default behavior.</li><li>• <b>Remote Only</b> - Only checks the remote server.</li><li>• <b>Local Only</b> - Only checks the local server. All remote users are mapped to the user specified. Any vendor attributes received by an authentication server are ignored.</li></ul>

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

Managing User Permissions

This section describes how to change the administrator or monitor passwords and define users in the Settings > Security: User Permissions page.

**Beta Draft**

## Accounts

The system uses these accounts based on what actions the user can take:

- **Admin** - The administrator user has full privileges: for example, as an administrator you can set and modify configuration settings, add and delete users, restart and reboot Core services, and create and view performance and system reports. The system administrator role allows you to add or remove a system administrator role for any other user, but not for yourself.
- **Monitor** - A monitor user may view reports, user logs, and change their password. A monitor user cannot make configuration changes, modify private keys, view logs, or manage cryptographic modules in the system.

---

**Note:** The default administrator password is *password*.

---

You can also create users, assign passwords to the user, and assign varying configuration roles to the user.

An administrator role configures a system administrator role. Read-only permission is not allowed for this role. This role allows permission for all other RBM roles, including creating, editing and removing user accounts.

A user role determines whether the user has permission to:

- **Read-only** - With read-only privileges you can view current configuration settings but you cannot change them.
- **Read/Write** - With read and write privileges you can view settings and make configuration changes for a feature.
- **Deny** - With deny privileges you cannot view settings or save configuration changes for a feature.

As an example, you might have user Jane who can make configuration changes to storage settings whereas user John can only view these configuration settings; and finally, user Joe cannot view, change, or save the storage settings.

Available menu items reflect the privileges of the user. For example, any menu items that a user does not have permission to use are unavailable. When a user selects an unavailable link, the User Permissions page appears.

**Beta Draft****To configure user permissions**

1. Choose Settings > Security: User Permissions to display the User Permissions page.

**Figure 5-2. User Permissions Page**

**User Permissions** Security > User Permissions ? Save

**Accounts:**  
 + Add a New Account - Remove Selected Accounts

Account	Enabled	Role	Permanent	AAA Default User
▼ admin	Yes	Administrator	✓	✓
▶ monitor	No	Special	✓	

**Login Failure Details**

Login Failure Count: 0  
 Login Failure Lockout: No  
 Login Password Status: Never Expire  
 Last Login Failure: None

Clear

☐ Change Password

New Password:   
 New Password Confirm:

☒ Enable Account

☒ Make this the AAA Default User (for RADIUS/TACACS+ logins)

**Roles and Permissions**

This is an administration account with full access to configurations and reports on this Appliance. This account can also be used to create/edit/remove user accounts.

Apply

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

Control	Description
admin/monitor	Click the right arrow to change the password or to create a default user account.
	<p><b>Change Password</b> - Enables password protection.</p> <p>Password protection is an account control feature that allows you to select a password policy for more security. When you enable Account Control on the Password Policy page, a user must use a password.</p> <p>When a user has a null password to start with, the administrator can still set the user password with account control enabled. However, once the user or administrator changes the password, it cannot be reset to null as long as account control is enabled.</p>
	<b>Password</b> - Specify a password in the text box.
	<b>Password Confirm</b> - Retype the new administrator password.
	<p><b>Enable Account</b> - Select to enable or clear to disable the administrator or monitor account.</p> <p>When enabled, you may make the account the default user for Radius and TACACS+ authorization. You may only designate one account as the default user. Once enabled, the default user account may not be disabled or removed. The Accounts table displays the account as permanent.</p>

**Beta Draft**

In the Accounts panel of the User Permissions page, you can create user accounts, assign them passwords, and assign varying configuration roles and access restrictions.

**To configure user permissions**

1. Choose Settings > Security: User Permissions to display the User Permissions page.

**Figure 5-3. User Permissions Page**

2. Under Accounts, complete the configuration using the controls described in this table.

Control	Description
Add a New User	Click to display the controls for creating a new account.
Account Name	Specify a name for the account.
Password/New Password Confirm	Specify a password in the text box, and then retype the password for confirmation.
Enable Account	Select the check box to enable the new account.
System Settings	Any nonstorage settings.
System Diagnostics and Reports	All nonstorage reports such as CPU and memory utilization, system logs, system dumps, or TCP dumps. You can set permissions to either Deny or Read/Write. This user role has restricted configuration capabilities.
Storage Settings	Any settings relating to the Core configuration, such as configuring LUNs, Edges, and Failover.
Storage Diagnostics and Reports	Storage-specific graphs and statistics, such as memory logs. You can set permissions to either Deny or Read/Write. This user role also has Read/Write access to System Diagnostics and Reports to assist you with troubleshooting.

**Beta Draft**

Control	Description
Permissions	Configures a role that determines whether the user: <ul style="list-style-type: none"> <li>Has permission to view current configuration settings but not change them (Read-Only).</li> <li>Has permission to view settings and make configuration changes for a feature (Read/Write).</li> <li>Cannot view or save settings or configuration changes for a feature (Deny).</li> </ul>
Add	Click to add the new user to the system. The new user appears in the User table.
Remove Selected Users	Select the check box next to the name and click <b>Remove Selected Users</b> .

## Recommended Permissions by Role

This table shows possible combinations of permissions listed by role.

Role	Permission
Storage Administrator	<ul style="list-style-type: none"> <li>System Settings - Read-Only</li> <li>Storage Settings - Read/Write</li> <li>Storage Diagnostics and Reports - Read/Write</li> </ul>
Storage Monitor	<ul style="list-style-type: none"> <li>System Settings - Read-Only</li> <li>Storage Settings - Read-Only</li> <li>Storage Diagnostics and Reports - Read/Write</li> </ul>
Network Administrator	<ul style="list-style-type: none"> <li>System Settings - Read/Write</li> <li>System Diagnostics and Reports - Read/Write</li> <li>Storage Settings - Deny</li> <li>Storage Diagnostics and Reports - Deny</li> </ul>
Network Monitor	<ul style="list-style-type: none"> <li>System Settings - Read-Only</li> <li>System Diagnostics and Reports - Read/Write</li> <li>Storage Settings - Deny</li> <li>Storage Diagnostics and Reports - Deny</li> </ul>

## Managing Password Policy

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

### Selecting a Password Policy

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

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

**Beta Draft****To set a password policy**

1. Choose Settings > Security: User Permissions to display the User Permissions page.
2. Click the Password Policy link at the bottom of the User Permissions page to display the Password Policy page.

**Figure 5-4. Password Policy Page**

3. Select the Enable Account Control check box to set a password policy. Enabling account control makes password use mandatory.

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

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

The Core does not allow empty passwords when account control is enabled.

4. Optionally, select either the Basic or Strong Security template. When you select the basic template, the system prepopulates the page with the secure settings. Also, the system prompts a user logging in to the Core after 60 days to change their password. By default, the Core locks out a user logging in after 300 days without a password change. After the system locks out a user, an administrator must unlock the account. After the system locks them out, an administrator must unlock the user account. For more details on unlocking user accounts, see [“Unlocking an Account” on page 129](#).

**Beta Draft**

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

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

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

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

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

**Beta Draft**

## Unlocking an Account

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

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

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

### To unlock an account

1. Log in as an administrator (admin).
2. Choose Settings > Security: User Permissions page to display the User Permissions page.
3. Select the account from the User column and click **Clear Login Failure Details**.

When the user logs in to their account successfully, the Core resets the login failure count.

## Resetting an Expired Password

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

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

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

### To reset the password and unlock the account

1. Log in as an administrator (admin).
2. Choose Settings > Security: User Permissions to display the User Permissions page.
3. Select the account from the User column and click **Clear Login Failure Details**.
4. Select Change Password.
5. Type and confirm the new password and click **Change Password**.

---

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

---

## Configuring RADIUS Server Authentication

You set up RADIUS server authentication in the Settings > Security: RADIUS page.

RADIUS is an access control protocol that uses a challenge and response method for authenticating users. Setting up RADIUS server authentication is optional.

You can prioritize local, RADIUS, and TACACS+ authentication methods for the system and set the authorization policy and default user for RADIUS and TACACS+ authorization systems in the Settings > Security: General Settings page.

For detailed information about setting up RADIUS and TACACS+ servers, see the *SteelHead Deployment Guide*.

### To set RADIUS server authentication

1. Choose Settings > Security: RADIUS to display the RADIUS page.

Figure 5-5. RADIUS Page

**RADIUS** Security > RADIUS ?

**Default RADIUS Settings**

☐ Set a Global Default Key

Global Key:  (leave unchanged to leave the global key unchanged)

Confirm Global Key:

Timeout (seconds):  (1 - 60)

Retries:  (0 - 5)

**Apply**

**RADIUS Servers:**

+ Add a RADIUS Server - Remove Selected

Server	Port	Type	Key	Timeout	Retries	Status
No RADIUS servers.						

Related Topics: [General Settings](#)

2. Under Default RADIUS Settings, complete the configuration using the controls described in this table.

Control	Description
Set a Global Default Key	Enables a global server key for the RADIUS server.
Global Key	Specify the global server key.
Confirm Global Key	Confirm the global server key.
Timeout (seconds)	Specify the time-out period in seconds (1 - 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. The default value is 1.

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

**Beta Draft****To add a new RADIUS server**

1. Choose Settings > Security: RADIUS to display the RADIUS page.

**Figure 5-6. RADIUS Page**

**RADIUS** Security > RADIUS ?

**Default RADIUS Settings**

☐ Set a Global Default Key

Global Key:  (leave unchanged to leave the global key unchanged)

Confirm Global Key:

Timeout (seconds):  (1 - 60)

Retries:  (0 - 5)

**Apply**

**RADIUS Servers:**

[+ Add a RADIUS Server](#) [- Remove Selected](#)

Server	Port	Type	Key	Timeout	Retries	Status
No RADIUS servers.						

Related Topics: [General Settings](#)

**Beta Draft**

2. Complete the configuration using the controls described in this table.

Control	Description
Add a RADIUS Server	Displays the controls for defining a new RADIUS server.
Hostname or IP Address	Specify the hostname or server IP address. RiOS doesn't support IPv6 server IP addresses.
Authentication Port	Specify the port for the server.
Authentication Type	Select one of these authentication types: <ul style="list-style-type: none"> <li>• <b>PAP</b> - Password Authentication Protocol (PAP), which validates users before allowing them access to the RADIUS server resources. PAP is the most flexible protocol but is less secure than CHAP.</li> <li>• <b>CHAP</b> - Challenge-Handshake Authentication Protocol (CHAP), which provides better security than PAP. CHAP validates the identity of remote clients by periodically verifying the identity of the client using a three-way handshake. This validation happens at the time of establishing the initial link and might happen again at any time. CHAP bases verification on a user password and transmits an MD5 sum of the password from the client to the server.</li> </ul>
Override the Global Default Key	Overrides the global server key for the server. <b>Server Key</b> - Specify the override server key. <b>Confirm Server Key</b> - Confirm the override server key.
Timeout	Specify the time-out period in seconds (1 to 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are from 0 to 5. The default value is 1.
Enabled	Enables the new server.
Add	Adds the RADIUS server to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

**Note:** If you add a new server to your network and you do not specify these settings at that time, the global settings are applied automatically.

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

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

## Configuring TACACS+ Server Authentication

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

Enabling this feature is optional.

TACACS+ is an authentication protocol that enables a remote access server to forward a login password for a user to an authentication server to determine whether access is allowed to a given system.

For detailed information about configuring RADIUS and TACACS+ servers to accept login requests from a Riverbed appliance, see the *SteelHead Deployment Guide*.

**Beta Draft****To modify TACACS+ settings**

1. Choose Settings > Security: TACACS+ to display the TACACS+ page.

**Figure 5-7. TACACS+ Page**

2. Under Default TACACS+ Settings, complete the configuration using the controls described in this table.

Control	Description
Set a Global Default Key	Select this option to enable a global server key for the server.
Global Key	Specify the global server key.
Confirm Global Key	Confirm the global server key.
Timeout	Specify the time-out period in seconds (1 to 60). The default value is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are from 0 to 5. The default is 1.

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

**Beta Draft****To add a TACACS+ server**

1. Choose Settings > Security: TACACS+ to display the TACACS+ page.

**Figure 5-8. TACACS+ Page - TACACS+ Servers**

2. Under TACACS+ Servers, click **Add a TACACS+ Server** and complete the configuration using the controls described in this table.

Control	Description
Add a TACACS+ Server	Displays the controls for defining a new TACACS+ server, as described in this table.
Server IP Address	Specify the server IP address.
Authentication Port	Specify the port for the server. The default value is 49.
Authentication Type	Select either <b>PAP</b> or <b>ASCII</b> as the authentication type.
Override the Global Default Key	Select this option to override the global server key for the server.
Server Key	Specify the override server key.
Confirm Server Key	Confirm the override server key.
Timeout	Specify the time-out period in seconds (1 to 60). The default is 3.
Retries	Specify the number of times you want to allow the user to retry authentication. Valid values are from 0 to 5. The default is 1.
Enabled	Enables the new server.
Add	Adds the TACACS+ server to the list.
Remove Selected	Select the check box next to the name and click <b>Remove Selected</b> .

**Note:** If you add a new server to your network and you do not specify these fields at that time, the global settings are applied automatically.

**Beta Draft**

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

## Unlocking the Secure Vault

You can unlock and change the password for the secure vault in the Settings > Security: Secure Vault page.

The secure vault contains sensitive information from your Core configuration, including SSL private keys and the data store encryption key. These configuration settings are encrypted on the disk at all times, using AES 256-bit encryption.

Initially, the secure vault is keyed with a default password known only to the system. This password enables the system to automatically unlock the vault during system startup. You can change the password, but the secure vault does not automatically unlock upon startup. To manage SSL configuration on the Core and to unlock the secure stores on Edges, you must unlock the secure store.

### To unlock and change the password of the secure vault

1. Choose Settings > Security: Secure Vault to display the Secure Vault page.

**Figure 5-9. Secure Vault Page**

2. Under Unlock Secure Vault, specify the password.

**Note:** To optimize SSL connections or to use datastore encryption, the secure vault must be unlocked.

3. Click **Unlock Secure Vault**.
4. Under Change Secure Vault Password, complete the configuration using the controls described in this table.

Control	Description
Current Password	Specify the current password. If you are changing the default password that ships with the product, leave the text field blank.

**Beta Draft**

Control	Description
New Password	Specify a new password for the secure vault.
New Password Confirm	Retype the new password for the secure vault.
Change Password	Changes the password to the new value.

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

## Configuring Web Settings

You can modify the SteelFusion Core Management Console web user interface and certificate settings in the Settings > Security: Web Settings page.

### To modify web settings

- Choose Settings > Security: Web Settings to display the Web Settings page.

**Figure 5-10. Web Settings Page**

- Under Web Settings, complete the configuration using the controls described in this table.

Control	Description
Default Web Login ID	Specify the username that appears on the authentication page. The default value is admin.
Web Inactivity Timeout (minutes)	Specify the number of idle minutes before time-out. The default is 15. A value of 0 disables time-out.

**Beta Draft**

Control	Description
Allow Session Timeouts When Viewing Auto-Refreshing Pages	By default, session time-out is enabled, which stops the automatic updating of the report pages when the session times out. Clear this check box to disable the session time-out, remain logged in indefinitely, and automatically refresh the report pages. <b>Caution:</b> Disabling this feature poses a security risk.

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

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

## Managing Web SSL Certificates

You can manage SSL certificates for the web user interface in the SteelFusion Core Management Console. In this page, you can:

- generate the certificate and key pairs on the Core. This overwrites the existing certificate and key pair regardless of whether the previous certificate and key pair was self-signed or user added. The new self-signed certificate lasts for one year (365 days).
- create certificate signing requests from the certificate and key pairs.
- replace a signed certificate with one created by an administrator or generated by a third party certificate authority.

### To modify web certificates

1. Choose Settings > Security: Web Settings to display the Web Settings page.

**Figure 5-11. Web Settings Page - Web Certificate Details Tab**



**Beta Draft**

2. Under Web Certificate, select the Details tab to display the Core identity certificate details.

Detail	Description
Issued To/Issued By	<b>Common Name</b> - Displays the common name of the certificate authority. <b>Email</b> - Displays the email of the appliance administrator. <b>Organization</b> - Displays the organization name (for example, the company). <b>Locality</b> - Displays the city. <b>State</b> - Displays the state. <b>Country</b> - Displays the country.
Validity	<b>Issued On</b> - Displays the date the certificate was issued. <b>Expires On</b> - Displays the date the certificate expires.
Fingerprint	Displays the SSL fingerprint.
Key	<b>Type</b> - Displays the key type. <b>Size</b> - Displays the size, in bytes.

3. Under Web Certificate, select the Replace tab.

**Figure 5-12. Web Settings Page - Web Certificate Replace Tab**

4. Complete the configuration using the controls described in this table.

Control	Description
Import Certificate and Private Key	Imports the certificate and key. The page displays controls for browsing to and uploading the certificate and key files. Or, you can use the text box to copy and paste a PEM file. The private key is required regardless of whether you are adding or updating the certificate.
Certificate	<b>Upload</b> - Browse to the local file in PKCS-12, PEM, or DER formats. <b>Paste it here (PEM)</b> - Copy and then paste the contents of a PEM file.

**Beta Draft**

Control	Description
Private Key	<p>Select the private key origin.</p> <ul style="list-style-type: none"> <li>• <b>The Private Key is in a separate file (see below)</b> - You can either upload it or copy and paste it.</li> <li>• <b>This file includes the Certificate and Private Key</b></li> <li>• <b>The Private Key for this Certificate was created with a CSR generated on this appliance</b></li> </ul>
Separate Private Key	<p><b>Upload (PEM or DER formats)</b> - Browse to the local file in PEM, or DER formats.</p> <p><b>Paste it here (PEM only)</b> - Paste the contents of a PEM file.</p> <p><b>Decryption Password</b> - Specify the decryption password, if necessary. Passwords are required for PKCS-12 files, optional for PEM files, and never needed for DER files.</p>
Generate Self-Signed Certificate and New Private Key	<p>Select this option to generate a new private key and self-signed public certificate. The page displays controls to identify and generate the new certificate and key.</p> <p><b>Organization Name</b> - Specify the organization name (for example, the company).</p> <p><b>Organization Unit Name</b> - Specify the organization unit name (for example, the section or department).</p> <p><b>Locality</b> - Specify the city.</p> <p><b>State (no abbreviations)</b> - Specify the state.</p> <p><b>Country (2-letter code)</b> - Specify the country (2-letter code only).</p> <p><b>Email Address</b> - Specify the email address of the contact person.</p> <p><b>Validity Period (Days)</b> - Specify how many days the certificate is valid.</p>
Private Key	<b>Cipher Bits</b> - Select the key length from the drop-down list. The default is 1024.
Generate Certificate and Key	Click to generate the certificate and key.

**Beta Draft****To generate a CSR**

1. Under Web Certificate, select the Generate CSR tab.

**Figure 5-13. Web Settings Page - Web Certificate Generate CSR Tab**

Web Certificate:

Details PEM Replace **Generate CSR**

Common Name: oak-sh395.example.com

Organization Name: Riverbed Technology, Inc.

Organization Unit Name: SteelFusion

Locality: San Francisco

State: California

Country: US (two-letter code)

Email Address:

**Generate CSR**

2. Complete the configuration using the controls described in this table.

Control	Description
Organization Name	Specify the organization name (for example, the company).
Organization Unit Name	Specify the organization unit name (for example, the section or department).
Locality	Specify the city.
State	Specify the state. Do not abbreviate.
Country	Specify the country (two-letter code only).
Email Address	Specify the email address of the contact person.
Generate CSR	Click to generate the Certificate Signing Request.

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

## CHAPTER 6 Maintaining Your System

This chapter describes how to view job status, how to upgrade your software, and how to shut down and reboot the system. It includes the following sections:

- [“Starting, Stopping, and Restarting the Service” on page 141](#)
- [“Displaying Scheduled Jobs and Job Status” on page 142](#)
- [“Managing Licenses” on page 143](#)
- [“Upgrading the Software” on page 144](#)
- [“Rebooting and Shutting Down the Core” on page 146](#)
- [“Changing the Administrative Password” on page 146](#)

---

### Starting, Stopping, and Restarting the Service

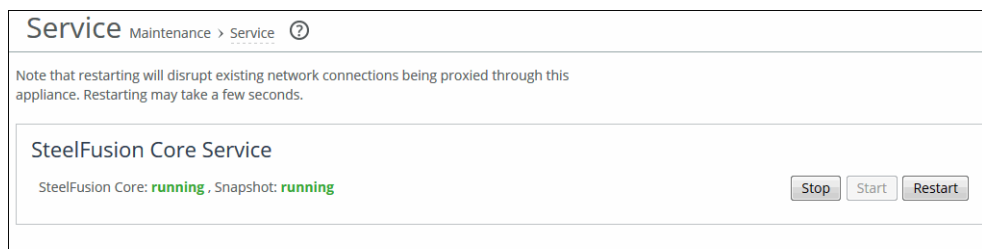
You can start, stop, and restart the Core service in the Settings > Maintenance: Service page. You can also use this page to reset the service alarm after it has been triggered.

The SteelFusion Core service is a daemon that runs in the background, performing operations when required. Many of the SteelFusion Core service commands are initiated at startup. It is important to restart the SteelFusion Core service when you have made changes to your configuration.

#### To start, stop, or restart services

1. Choose Settings > Maintenance: Service to display the Service page.

**Figure 6-1. Service Page**



2. Under SteelFusion Core Service, click **Stop**, **Start**, or **Restart**.

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

### To reset the service alarm

1. Choose Settings > Maintenance: Services to display the Service page.  
The option to reset the service alarm appears only after the service triggers the Reset Service Alarm.
2. Under Reset Service Alarm, click **Reset Service Alarm**.
3. Click **Save** to save your settings permanently.

## Displaying Scheduled Jobs and Job Status

You can view completed, pending, or inactive jobs, as well as jobs that were not completed because of an error, in the Settings > Maintenance: Scheduled Jobs page.

Jobs are CLI commands that execute at a time you specify.

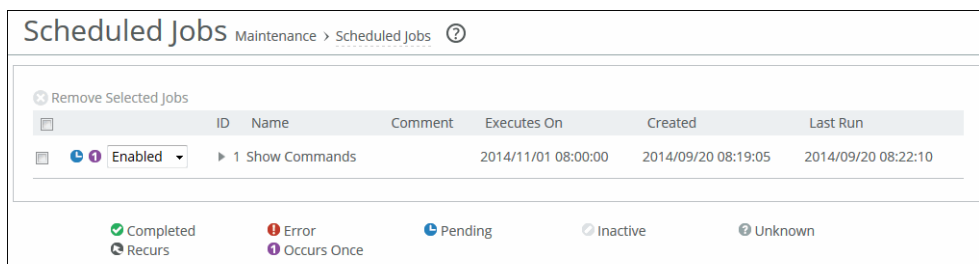
The only jobs you can schedule using the Core are software upgrades and configuration pushes; for all other jobs, you must use the CLI.

For details about scheduling jobs using the CLI, see the *SteelFusion Command-Line Interface Reference Manual*.

### To modify scheduled jobs

1. Choose Settings > Maintenance: Scheduled Jobs to display the Scheduled Jobs page.

**Figure 6-2. Scheduled Jobs Page**



2. To cancel a job or to remove a completed job from the list, select the entry and click **Remove Selected Jobs**.
3. Click the job number to display the details about the job.
4. Under Details for Job <#>, complete the configuration using the controls described in this table.

Control	Description
ID	Specify an ID number for the job.
Name	Specify a name for the job.
Comment	Specify a comment.
Interval (seconds)	Specify how often you want the job to run.

Control	Description
Executes On	Specify the date on which the job runs.
Enable/Disable Job	Select the check box to enable the job or clear the check box to disable the job.
Apply Changes	Click to apply changes.
Cancel/Remove This Job	Click to remove the currently displayed job.
Execute Now	Click to execute the currently displayed job.

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

## Managing Licenses

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

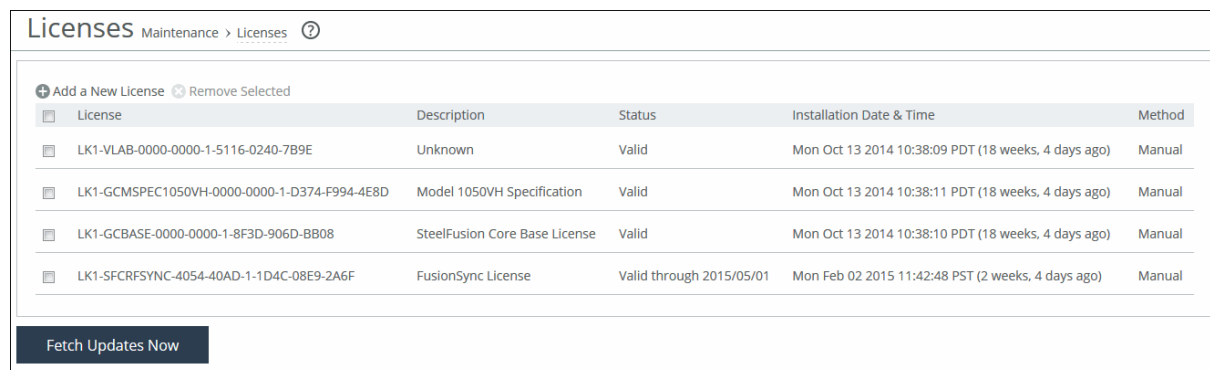
You can install licenses and update or remove expired licenses on the Licenses page.

For more information on managing licenses and model upgrades, see the *Upgrade and Maintenance Guide*.

### To install a license

- Choose Settings > Maintenance: Licenses to display the Licenses page.

**Figure 6-3. Licenses Page**



Licenses Maintenance > Licenses ?				
<input type="button" value="Add a New License"/> <input type="button" value="Remove Selected"/>				
<input type="checkbox"/> License	Description	Status	Installation Date & Time	Method
<input type="checkbox"/> LK1-VLAB-0000-0000-1-5116-0240-7B9E	Unknown	Valid	Mon Oct 13 2014 10:38:09 PDT (18 weeks, 4 days ago)	Manual
<input type="checkbox"/> LK1-GCMSPEC1050VH-0000-0000-1-D374-F994-4E8D	Model 1050VH Specification	Valid	Mon Oct 13 2014 10:38:11 PDT (18 weeks, 4 days ago)	Manual
<input type="checkbox"/> LK1-GCBASE-0000-0000-1-8F3D-906D-BB08	SteelFusion Core Base License	Valid	Mon Oct 13 2014 10:38:10 PDT (18 weeks, 4 days ago)	Manual
<input type="checkbox"/> LK1-SFCRFSYNC-4054-40AD-1-1D4C-08E9-2A6F	FusionSync License	Valid through 2015/05/01	Mon Feb 02 2015 11:42:48 PST (2 weeks, 4 days ago)	Manual
<input type="button" value="Fetch Updates Now"/>				

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

- Complete the configuration using the controls described in this table.

Control	Description
Add a New License	Displays the controls to add a new license.
Licenses Text Box	Copy and paste the license key provided by Riverbed Support or Sales into the text box. <b>Tip:</b> Separate multiple license keys with a space, Tab, or Enter.
Add	Adds the license.

Control	Description
Fetch Updates Now	Contacts the Riverbed license portal and downloads all applicable licenses for the Core.

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

## Removing a License

Riverbed recommends that you keep previously used licenses in case you want to downgrade to an earlier software version; however, in some situations you might want to remove a license.

### To remove a license

1. Choose Settings > Maintenance: Licenses to display the Licenses page.
2. Select the license you want to delete.
3. Click **Remove Selected**.
4. Click **Save** to save your settings permanently.

## Fetching a License

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

### To fetch a license on demand

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

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

---

## Upgrading the Software

You can upgrade or revert to a backup version of the software in the Settings > Maintenance: Software Upgrade page. The top of the page displays the current version number and the backup version. The Core version histories appear at the bottom of the page.

---

**Note:** As of version 4.2, software upgrade and downgrade restrictions are in place to enforce supported upgrade/downgrade paths. You can upgrade or downgrade the Core software up to a maximum of two versions beyond the current version. If you attempt to upgrade or downgrade to an unsupported version, a warning will display.

---

## To upgrade your software

1. Download the software image from the Riverbed Support site to a location such as your desktop. Optionally, you can download the image directly from a specified URL to the Core.
2. Log in to the Management Console using the Administrator account (admin).
3. Choose Settings > Maintenance: Software Upgrade to display the Software Upgrade page.

**Figure 6-4. Software Upgrade Page**

**Software Upgrade** Maintenance > Software Upgrade ⓘ

**Software Upgrade**

**Booted Version:**  
rbt\_dva 4.0.0 2015-03-04 18:03:00 x86\_64

**Backup Version:**  
rbt\_dva 4.0.0 2015-03-04 18:03:00 x86\_64

[Switch to Backup Version](#)

**Install Upgrade**

☒ From URL

☐ From Local File  
[Browse...](#) No file selected.

☐ Schedule Upgrade for Later  
Date:  (YYYY/MM/DD) Time:  (HH:MM:SS)

[Install](#)

**Software Version History**

4.0.0 (Thu Mar 5 05:14:03 UTC 2015)

Related Topics: [Scheduled Jobs](#)

4. Select one of the following options under Install Upgrade:
  - **From URL** - Type the URL that points to the software image in the text box. You can use HTTP, HTTPS, FTP, or SCP formats for the URL. Use one of the following formats:
    - http://host/path/to/file
    - https://host/path/to/file
    - ftp://user:password@host/path/to/file
    - scp://user:password@host/path/to/file
  - **From Local File** - Browse your file system and select the software image. Select this option and specify the path, or click Browse to go to the local file directory. The image is uploaded immediately; however the image is installed and the system is rebooted at the time you specify.
  - **Schedule Upgrade for Later** - Schedules the upgrade process. Specify the date and time to run the upgrade using the following format: yyyy/mm/dd hh:mm:ss.
5. Click **Install**.

The system installs the image in the backup partition and sets the option to load the backup partition version on reboot.

6. Reboot the Core.

### To switch to the backup version

1. Log in to the management console using the Administrator account (admin).
2. Go to the Settings > Maintenance: Software Upgrade page and click **Switch to Backup Version**.
3. Reboot the Core or click **Cancel Version Switch** to cancel.

---

## Rebooting and Shutting Down the Core

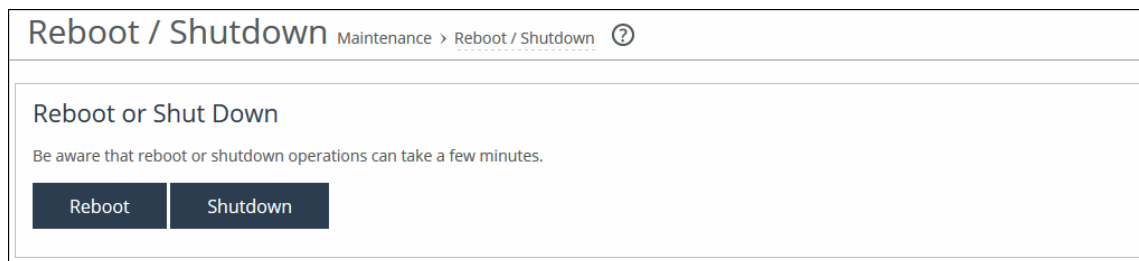
You can reboot or shut down the system in the Reboot/Shutdown page. Rebooting the system disrupts existing network connections and can take a few minutes.

To restart the system, you must manually power on the Core.

### To reboot or shut down the system

1. Choose Settings > Maintenance: Reboot/Shutdown to display the Reboot/Shutdown page.

Figure 6-5. Reboot/Shutdown Page



2. To reboot the system, click **Reboot**.  
After you click **Reboot**, you are logged out of the system and it is rebooted.
3. To shut down the system, click **Shutdown**.  
After you click **Shutdown**, the system is powered off.

---

## Changing the Administrative Password

You can change the administrative password in the My Account page.

This page includes a way to clear the preferences for the current user if any user settings result in an unsafe state and the Management Console cannot display the page.

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

---

**Note:** You must be logged in as the admin user to change the administrative password.

---

### To change the administrative password

1. Choose Settings > System Settings: My Account to display the My Account page.

**Figure 6-6. My Account Page**

**My Account** System Settings > My Account ?

**Password**

☐ Change Password

New Password:

Confirm New Password:

**Apply**

**User Preferences**

User preferences are used to remember the state of the management console across sessions on a per-user basis. They do not affect the configuration of the appliance.

**Restore Defaults**

**Administrator**

This is an administration account with full access to configurations and reports on this Appliance. This account can also be used to create/edit/remove user accounts.

2. Under Password, complete the configuration using the controls described in this table.

Control	Description
Change Password	Select this option to change the password.
New Password	Specify a new password.
Confirm New Password	Confirm the new password.

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

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

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



## CHAPTER 7    Displaying and Customizing Reports

This chapter describes how to display and customize Core storage and diagnostic reports. It includes the following sections:

- [“Viewing Storage Reports” on page 149](#)
- [“Building Custom Reports with the Report Builder” on page 163](#)
- [“Viewing the Networking Interface Counters Report” on page 164](#)
- [“Viewing Diagnostic Reports” on page 165](#)
- [“Viewing Logs” on page 172](#)
- [“Generating Dumps” on page 177](#)

---

### Viewing Storage Reports

This section describes how to create storage reports and logs. It includes the following sections:

- [“Accessing Settings from Storage Reports” on page 149](#)
- [“Viewing the SteelFusion Edge Stats” on page 150](#)
- [“Viewing the SteelFusion Edge Trends” on page 153](#)
- [“Viewing the LUN I/O Metrics Report” on page 155](#)
- [“Viewing the SAN I/O Metrics Report” on page 157](#)
- [“Viewing the Replication Data Sync - Remaining Bytes Report” on page 158](#)
- [“Viewing the Replication Journal I/O Report” on page 159](#)
- [“Viewing the Replication Write I/O Report” on page 161](#)

If the selected host is the manager of a pool, these reports can show graphs from the members of the pool. For details about pool management, see [“Configuring Pool Management” on page 82](#).

### Accessing Settings from Storage Reports

If you configure the current Core for failover with another device, all Storage report pages include an area with a link to failover details. Click the link to jump to the [Configure > Failover: Failover Configuration](#) page and the current settings.

If you configure the current Core for pool management, all Storage configuration and reports pages include a similar link to pool management details. Click the link to jump to the Configure > Pool Management: Edit Pool page and the current settings.

This area appears below the page title, and it also includes a drop-down list from which you can select Self (the current device), Failover Peer, or another pool member. Changing the selection displays the reports for those Cores.

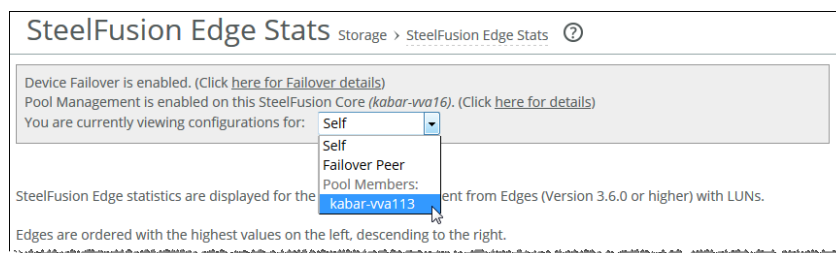
---

**Note:** A message in this area clarifies that viewing reports for “All” LUNs, Edges, or storage arrays includes data from both Self and Peer, regardless of the selection from the drop-down list.

---

Figure 7-1 shows a sample Storage report page with both failover and pool management enabled.

**Figure 7-1. Sample Storage Report Page with Failover and Pool Management Area**



## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the SteelFusion Edge Stats

The SteelFusion Edge Stats report gives you a high-level performance overview of all configured Edges running version 3.6.0 and later that are connected to the current Core. The configured Edges send their statistics every 5 minutes to the Core. The Core then displays this point-in-time information in the report, giving you a network-wide view of valuable details about how they are performing.

The highest values are displayed on the left and descend to the right, enabling you to easily identify any Edges with performance problems. Each Edge is assigned a color so you can compare statistics between several Edges at a glance.

If you have set up the Core for high-availability, the SteelFusion Edge Stats report only displays Edges that are served by individual Cores. If one Core is down, the report shows Edges served by both Cores, however on separate pages.

This report displays the following statistics:

- Blockstore Uncommitted Bytes
- Estimated Time to Drain Blockstore Uncommitted Data
- Blockstore Commit Delay
- Blockstore Space Utilization

- Read I/O Latency (All LUNs)
- Write I/O Latency (All LUNs)

## What This Report Tells You

The SteelFusion Edge Stats report answers these questions:

- How many bytes of uncommitted data is currently in the blockstore?
- What is the estimated time (in seconds) to drain the uncommitted blockstore data back to the Core?
- How long is the delay (in seconds) to commit the blockstore data?
- What percentage of the blockstore is currently being used by uncommitted data?
- What is the average read and write I/O latency (in milliseconds) over the last hour for all LUNs?

## To view the SteelFusion Edge Stats report

1. Choose Reports > Storage: SteelFusion Edge Stats to display the SteelFusion Edge Stats page.

Figure 7-2. SteelFusion Edge Stats Page



**Note:** If the Edge becomes disconnected, the report shows last received statistics until it is reconnected. If the Edge becomes idle, it continues to report statistics to the Core even if individual values do not change.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the SteelFusion Edge Trends

The SteelFusion Edge Trends report summarizes the standard I/O data traffic read from and written by the selected Edge for the specified period of time with a granularity of 5 minutes. The report also summarizes time series information for average commit delay and uncommitted data for all connected Edges.

You can choose from several graphs in the Edge Report drop-down list:

- SteelFusion Edge Read/Write Throughput
- Average Commit Delay at All Edges
- Uncommitted Data at All Edges

If you have set up the Core for high-availability, the SteelFusion Edge Trends report shows the average across all Edges that are served by individual Cores. If one Core is down, the report shows the average across all Edges served by both Cores.

### What This Report Tells You

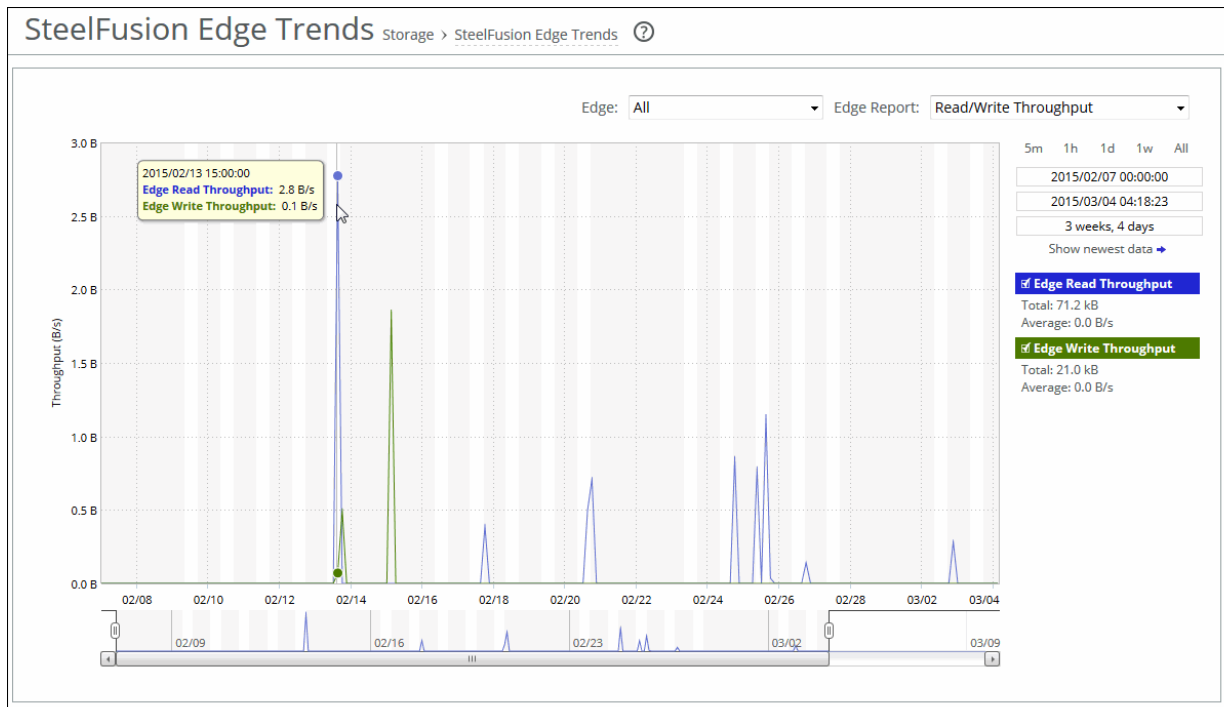
The SteelFusion Edge Trends report answers these questions:

- How many megabytes have been prefetched to, written by, and read from the selected Edge for the specified period?
- What is the average delay (in seconds) for committing data at all Edges?
- How much uncommitted data (in bytes) is accumulated at all Edges?

## To view the SteelFusion Edge Trends report

1. Choose Reports > Storage: SteelFusion Edge Trends to display the SteelFusion Edge Trends page.

Figure 7-3. SteelFusion Edge Trends Page



2. Customize the report using the controls described in this table.

Control	Description
Edge	Select the Edge whose statistics you want to see from the drop-down list.
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Select the type of report to display: <ul style="list-style-type: none"> <li>• Read/Write/Prefetch Throughput</li> <li>• Average Commit Delay at All Edges</li> <li>• Uncommitted Data at All Edges</li> </ul>

**Note:** If an Edge becomes disconnected, the report continues to show the aggregated time statistics. The values are averaged across all Edges that are currently connected to the Core. If one Edge becomes idle, the page also continues to show the aggregated time statistics. If all Edges become idle, the curve is flat.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the LUN I/O Metrics Report

The LUN I/O Metrics report summarizes the standard I/O data traffic read from and written to the selected LUN for the specified period of time.

This report can display three graphs:

- I/O
- I/O Operations Per Second
- I/O Latency

## What This Report Tells You

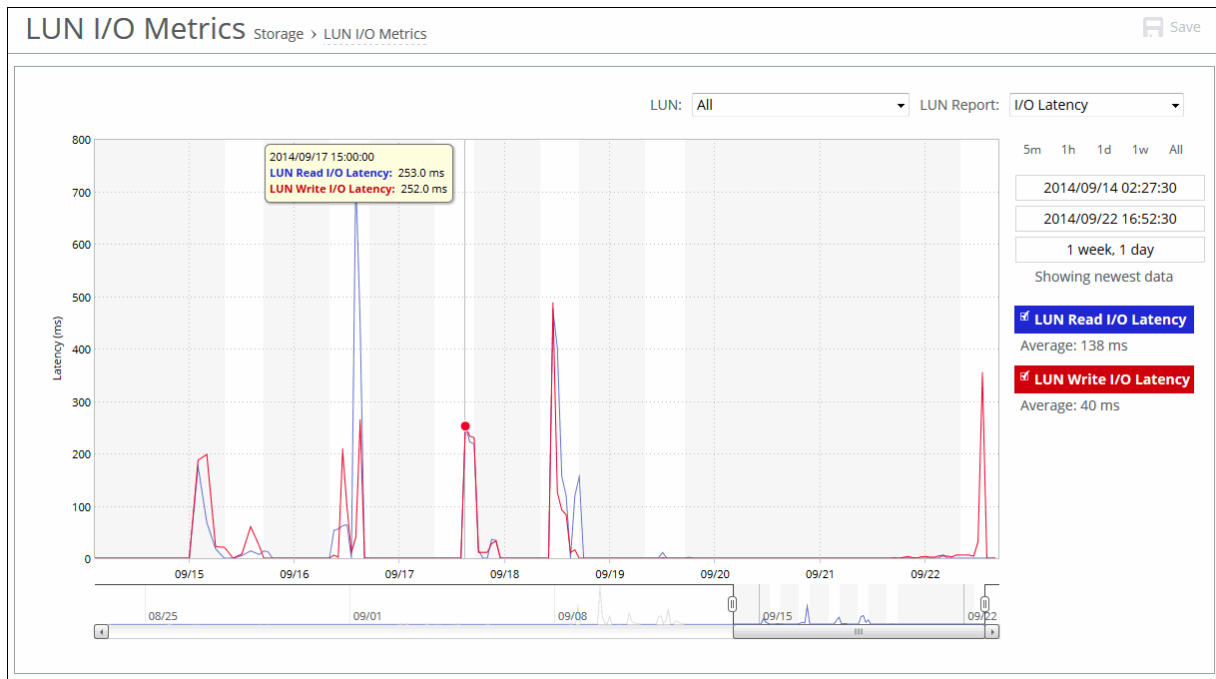
The LUN I/O Metrics report answers these questions:

- How many megabytes have been prefetched, written to, and read from the selected LUN for the specified period?
- How frequently were I/O operations written to and read from the selected LUN for the specified period?
- What are the average read and write latencies for the selected LUN for the specified period?

## To view the LUN I/O Metrics report

1. Choose Reports > Storage: LUN I/O Metrics to display the LUN I/O Metrics page.

Figure 7-4. LUN I/O Metrics Page



2. Customize the report using the controls described in this table.

Control	Description
LUN	Select the LUN whose statistics you want to see from the drop-down list, or choose All.
LUN Report	Select the LUN report to display from the drop-down list: <ul style="list-style-type: none"> <li>• I/O - read and write throughput</li> <li>• I/O Operations Per Sec - read and write I/O operations per second.</li> <li>• I/O Latency - read and write I/O latency</li> </ul>
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Toggle the reports to be displayed by clicking the type heading. The types of reports available depends on the type of LUN report selected above.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval that you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the SAN I/O Metrics Report

The SAN I/O Metrics report summarizes the standard I/O data traffic read from and written to the selected storage area for the specified period of time.

---

**Note:** The information in this report pertains only to iSCSI and Local Edge LUNs.

---

This report contains three graphs:

- I/O
- I/O Operations Per Second
- I/O Latency

### What This Report Tells You

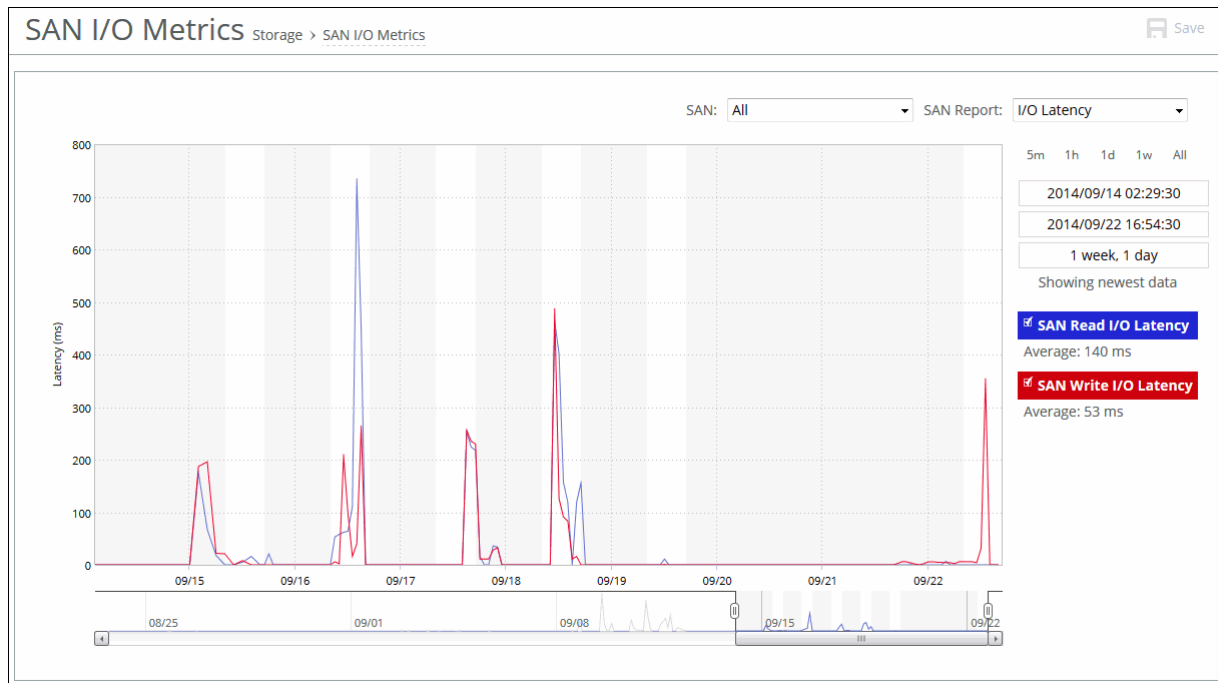
The SAN I/O Metrics report answers these questions:

- How many megabytes have been written to and read from the selected storage array for the specified period?
- How frequently were I/O operations written to and read from the selected storage array for the specified period?
- What are the average read and write latencies for the selected storage array for the specified period?

### To view the SAN I/O Metrics report

1. Choose Reports > Storage: SAN I/O Metrics to display the SAN I/O Metrics page.

**Figure 7-5. SAN I/O Metrics Page**



2. Customize the report using the controls described in this table.

Control	Description
SAN	Select the storage array whose statistics you want to see from the drop-down list.
SAN Report	Select the storage array report to display from the drop-down list: <ul style="list-style-type: none"> <li>• I/O - read and write throughput</li> <li>• I/O Operations Per Sec - read and write I/O operations per second</li> <li>• I/O Latency - read and write I/O latency</li> </ul>
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Toggle the reports to be displayed by clicking the type heading. The types of reports available depends on the type of storage array report selected above.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the Replication Data Sync - Remaining Bytes Report

The Replication Data Sync - Remaining Bytes report summarizes the amount of data remaining to be replicated to the secondary data center once replication has been resumed. This data was recorded to the Journal LUN when replication was suspended. This report also displays the amount of data remaining to be replicated to the secondary data center during first sync (first sync creates the initial copy of the primary data center LUN in the secondary data center).

### What This Report Tells You

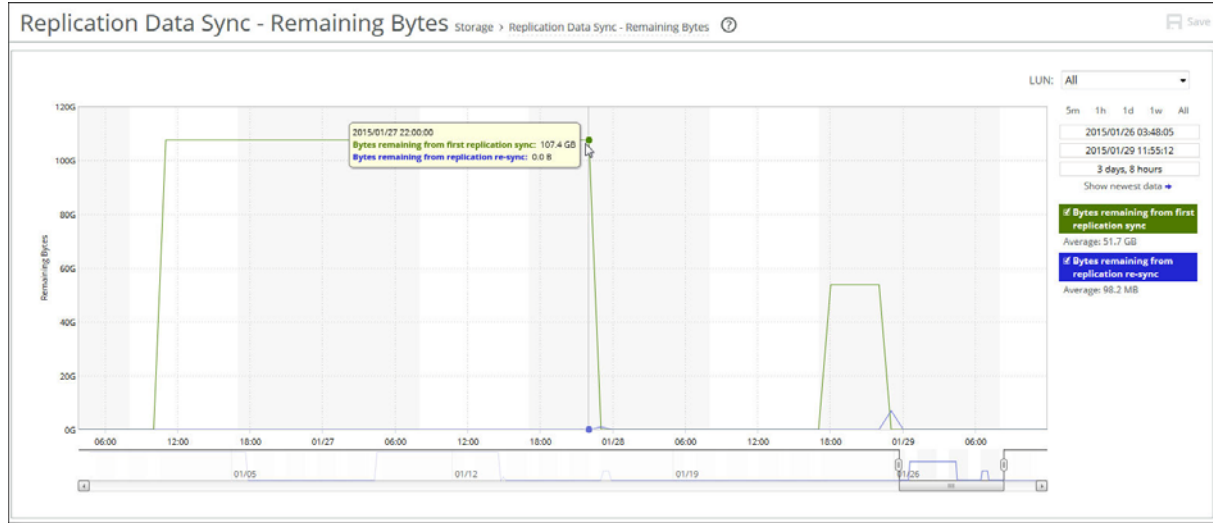
The Replication Data Sync - Remaining Bytes report answers these questions:

- How much data remains to be copied to the secondary data center once replication has restarted?
- How much data remains to be copied to the secondary data center during first sync?

## To view the Replication Data Sync - Remaining Bytes report

1. Choose Reports > Storage: Replication Data Sync to display the Replication Data Sync - Remaining Bytes page.

Figure 7-6. Replication Data Sync - Remaining Bytes Page



2. Customize the report using the controls described in this table.

Control	Description
LUN	Select the LUN whose statistics you want to see from the drop-down list.
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Toggle the reports to be displayed by clicking the type heading. The types of reports available depends on the type of report you selected above.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the Replication Journal I/O Report

The Replication Journal I/O report summarizes the number of bytes recorded in the journal for the specified LUN while replication has been suspended. When replication is resumed, all of these bytes have to be written to the secondary data center to bring its LUN in sync.

This report also displays the number of bytes that have been received from the Edge that are in the process of being written to the secondary data center. A high value indicates possible issues in the secondary data center write path.

**Note:** The Replication Journal Bytes Pending (on disk) report only shows information while replication is suspended or is being re-synchronized. The Replication Journal Bytes Pending (in memory) report only shows information during active replication.

## What This Report Tells You

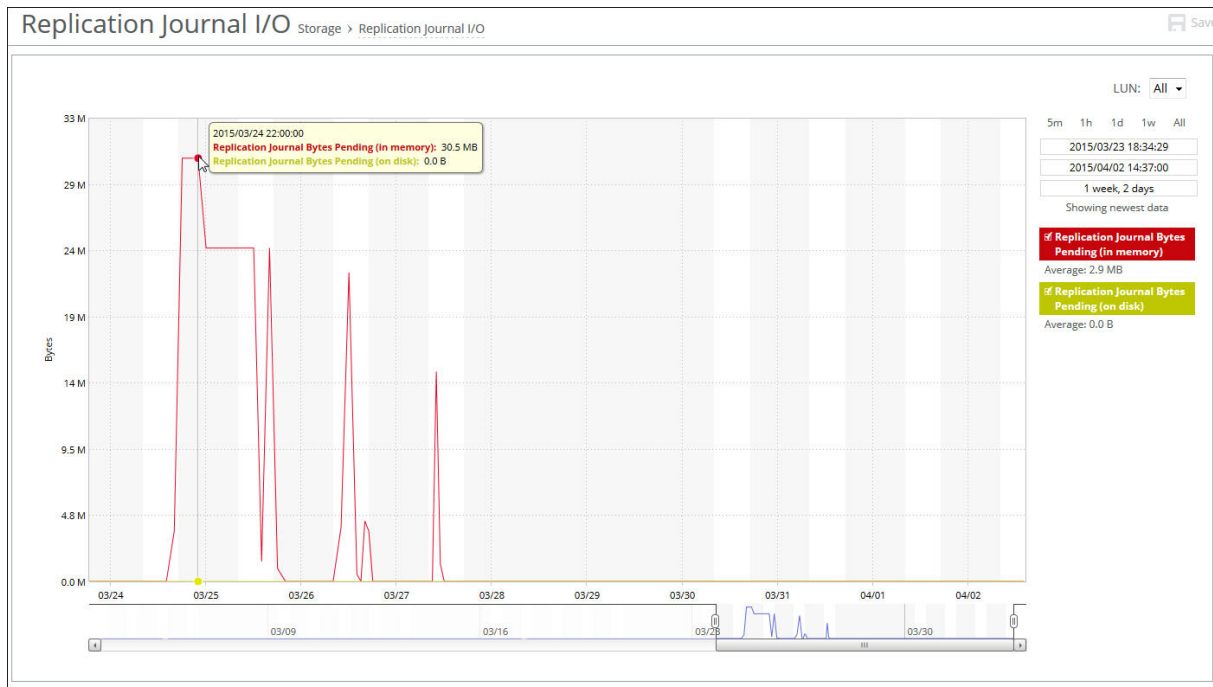
The Replication Data Journal I/O report answers these questions:

- How many bytes are pending in memory during active replication for the specified LUN?
- How many bytes are pending on disk in the journal for the specified LUN?
- Is there an excessive amount of data pending to be replicated to the secondary data center?

### To view the Replication Journal I/O report

1. Choose Reports > Storage: Replication Journal I/O to display the Replication Journal I/O page.

**Figure 7-7. Replication Journal I/O Page**



2. Customize the report using the controls described in this table.

Control	Description
LUN	Select the LUN whose statistics you want to see from the drop-down list.
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Toggle the reports to be displayed by clicking the type heading. The types of reports available depends on the type of report you selected above.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

## Viewing the Replication Write I/O Report

The Replication Write I/O report summarizes the amount of data that has been written to the secondary data center for the specified LUN, and the average latency (in milliseconds) of write operations to the secondary data center for the specified LUN. This report includes the latency of the inter-data center network and the latency of the storage array being used on the secondary data center.

### What This Report Tells You

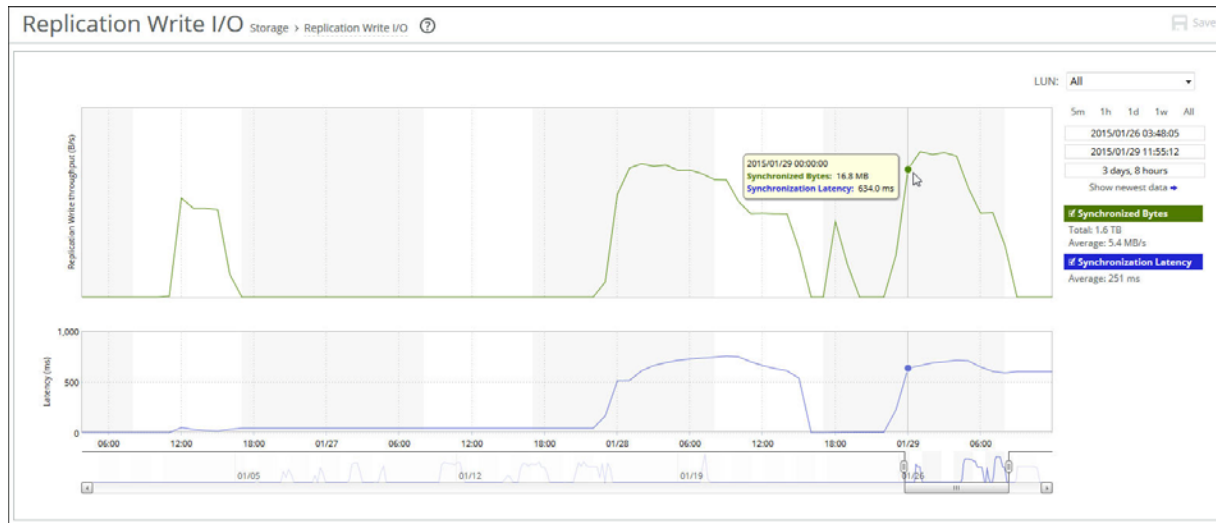
The Replication Write I/O report answers these questions:

- How much data has been synchronized between data centers?
- How much data is being replicated per second between data centers?
- What is the average replication latency between the primary and secondary data center?

## To view the Replication Write I/O report

1. Choose Reports > Storage: Replication Write I/O to display the Replication Write I/O page.

Figure 7-8. Replication Write I/O Page



2. Customize the report using the controls described in this table.

Control	Description
LUN	Select the LUN whose statistics you want to see from the drop-down list.
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Report type	Toggle the reports to be displayed by clicking the type heading. The types of reports available depends on the type of report you selected above.

## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

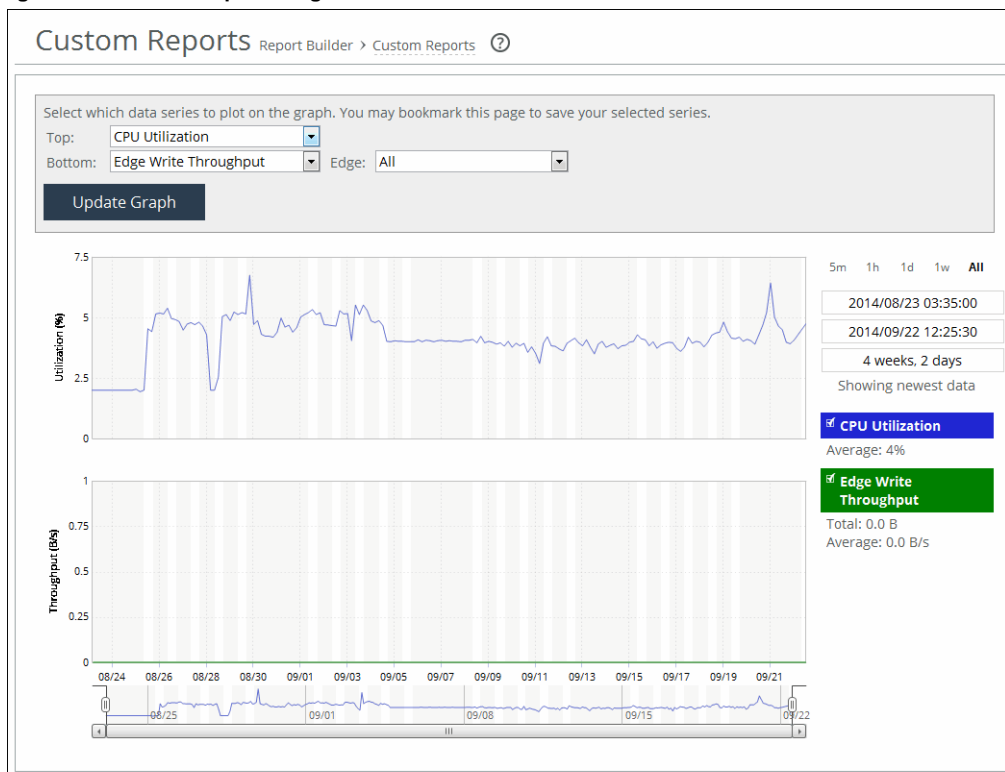
## Building Custom Reports with the Report Builder

The Report Builder feature enables you to display two different reports of your selection on a single page.

### To use the Report Builder

1. Choose Reports > Report Builder: Custom Reports to display the Custom Reports page.

**Figure 7-9. Custom Reports Page**



2. Customize the report using the controls described in this table.

Control	Description
Top	Select from the drop-down list the data report to be displayed in the top panel. Additional fields might appear depending on your selection. For example, if you select Edge Read Throughput, another drop-down list displays from which you can specify the Edge.
Bottom	Select from the drop-down list the data report to be displayed in the bottom panel.
Update Graph	Click to display the graph based on the selections.

## Viewing the Networking Interface Counters Report

The Networking Interface Counters report displays statistics for the Core's network interfaces.

### What This Report Tells You

The Interface Counters report answers these questions:

- What interfaces are active?
- What is the IP address of each interface?
- What is the Ethernet configuration for each interface?
- What number and types of packets have been received over each interface?
- What number and types of packets have been transmitted over each interface?

### To view the Interface Counters report

1. Choose Reports > Networking: Interface Counters to display the Interface Counters page.

**Figure 7-10. Interface Counters Page**

Interface Counters <span>Networking &gt; Interface Counters</span> <span>?</span>					
The network interface statistics have been collected since the system was booted 21:46:41 ago (or since the statistics were last cleared).					
Interface Statistics:					
Interface	IP	Ethernet	Link	Receive Packets	Transmit Packets
primary	10.5.53.61/24	MAC: 00:50:56:81:66:66 Speed: 1000Mb/s (auto) Duplex: full (auto)	true	73765 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 0 Multicast	106366 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions
aux	169.254.1.1/16	MAC: 00:50:56:81:66:67 Speed: 1000Mb/s (auto) Duplex: full (auto)	true	1887 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 0 Multicast	3 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions
eth0_0	10.5.134.20/17	MAC: 00:50:56:81:66:68 Speed: unknown Duplex: unknown	false	0 Packets 0 Discards 0 Errors 0 Overruns 0 Frames 0 Multicast	0 Packets 0 Discards 0 Errors 0 Overruns 0 Carriers 0 Collisions

2. To clear the current statistics, click **Clear All Interface Statistics**.

## Viewing Diagnostic Reports

This section describes how to display Core system files to help diagnose problems. It includes the following sections:

- [“Viewing Alarm Status Reports” on page 165](#)
- [“Viewing CPU Utilization Reports” on page 170](#)
- [“Viewing Memory Paging Reports” on page 171](#)

## Viewing Alarm Status Reports

The Alarm Status report provides status for the Core alarms.

For details about configuring alarm settings, see [“Setting Alarm Parameters” on page 98](#).

### What This Report Tells You

The Alarm Status report answers these questions:


- What is the current status of the system?
- What is the status of the system alarms?

### To view the alarm status report

1. Choose Reports > Diagnostics: Alarm Status to display the Alarm Status page.

Alternately, you can click the current alarm status that appears in the status bar of each page (Healthy, Admission Control, Degraded, or Critical) to display the Alarm Status page.

**Figure 7-11. Alarm Status Page**



Alarm	Status
<a href="#">Backup Integration</a>	OK
<a href="#">Block-disk</a>	OK
<a href="#">Core Disaster Recovery</a>	OK
<a href="#">CPU Utilization</a>	OK
<a href="#">Disk Full</a>	OK

The following table summarizes the alarms shown in the report.

Control	Description
Backup Integration	Indicates that the backup-integration module has failed.
Block-disk	Indicates that the block-disk module has failed.
Failover	Indicates that the Core has failed and the failover peer is in operation.

Control	Description
Core Disaster Recovery	<p>Indicates that there is an issue with replication:</p> <ul style="list-style-type: none"> <li>• <b>Journal LUN</b> - This alarm triggers if the Journal LUN size is not large enough to support the configured replica LUNs. To resolve this issue, increase the size of the Journal LUN.</li> <li>• <b>Replication synchronization latency</b> - This alarm triggers if the replication latency to the secondary data center is over 150 ms. To troubleshoot this issue, check the storage array latency on the secondary data center, inter-data center WAN latency, or if there are any MPIIO errors on the secondary data center.</li> <li>• <b>Replication state</b> - This alarm triggers if the status of one or more replicating LUNs is <i>suspended</i>. This acts as a reminder that data is not being replicated for the LUNs and therefore not protected from data center failure. This alarm is automatically resolved when data is synchronized and replication is <i>active</i>.</li> <li>• <b>Journal corruption</b> - This alarm triggers when any type of corruption is detected on the Journal LUN. Refer to the alarm email for details, including any possible resolutions (these will vary depending on the type of corruption).</li> <li>• <b>Data center connection</b> - This alarm triggers when connections are lost to the peer data center. To troubleshoot, check whether the inter-data center WAN and all the added interfaces are functioning correctly.</li> <li>• <b>Journal LUN missing</b> - This alarm triggers when the Journal LUN is not found on the storage array because it was accidentally unmapped from the backend, or the connection to the storage array was lost. To resolve this issue, check the backend settings, ensure that the LUN is exposed correctly, and check iSCSI logs for any issues.</li> </ul>
CPU Utilization	<p>Indicates that the system has reached the CPU threshold for one or more of the CPUs in the Core. If the system has reached the CPU threshold, check your settings.</p> <p>If your alarm thresholds are correct, reboot the Core.</p> <p><b>Note:</b> If more than 100 MB of data are moved through the Core while performing PFS synchronization, the CPU utilization might become high and result in a CPU alarm. This CPU alarm is not cause for concern.</p>
Disk Full	<p>Indicates that one or more of the following partitions on the disk is full:</p> <ul style="list-style-type: none"> <li>• Partition <code>"/boot"</code> Full</li> <li>• Partition <code>"/bootmgr"</code> Full</li> <li>• Partition <code>"/config"</code> Full</li> <li>• Partition <code>"/data"</code> Full</li> <li>• Partition <code>"/var"</code> Full</li> </ul>
Edge Service	<p>Indicates that the Core has lost connection with one of the configured Edges.</p>

Control	Description
Hardware	<p>Indicates that one or more hardware failures have occurred.</p> <p>This alarm setting also enables you to select one or more types of hardware failure (fan error, memory error, and so on), including:</p> <ul style="list-style-type: none"> <li>• <b>Fan Error</b> - Enables an alarm and sends an email notification if a fan is failing or has failed and needs to be replaced. By default, this alarm is enabled.</li> <li>• <b>Flash Error</b> - Enables an alarm when the system detects an error with the flash drive hardware. By default, this alarm is enabled.</li> <li>• <b>IPMI</b> - Enables an alarm and sends an email notification if an Intelligent Platform Management Interface (IPMI) event is detected.</li> <li>• <b>Other Hardware Error</b> - This alarm indicates that the system has detected a problem with the hardware. The alarm clears when you add the necessary hardware, remove the nonqualified hardware, or resolve other hardware issues. The following issues trigger the hardware error alarm: <ul style="list-style-type: none"> <li>• The appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration.</li> <li>• The appliance is using a dual in-line memory module (DIMM), a hard disk, or a NIC that is not qualified by Riverbed.</li> <li>• DIMMs are plugged into the appliance but the system cannot recognize them because the DIMM modules are in the wrong slot. You must plug DIMM modules into the black slots first and then use the blue slots when all of the black slots are in use.</li> <li>• A DIMM module is broken and you must replace it.</li> <li>• Other hardware issues.</li> </ul> </li> </ul> <p>By default, all Hardware alarms are enabled.</p> <ul style="list-style-type: none"> <li>• <b>Power Supply</b> - Enables an alarm and sends an email notification if an inserted power supply cord does not have power, as opposed to a power supply slot with no power supply cord inserted. By default, this alarm is enabled.</li> <li>• <b>RAID</b> - Indicates an error with the RAID array (for example, missing drives, pulled drives, drive failures, and drive rebuilds). An audible alarm might also sound. To see if a disk has failed, enter this CLI command from the system prompt: <pre>show raid diagram</pre> </li> <li>• For drive rebuilds, if a drive is removed and then reinserted, the alarm continues to be triggered until the rebuild is complete. Rebuilding a disk drive can take 4-6 hours. This alarm applies only to the SteelHead RAID Series 3000, 5000, and 6000.</li> </ul>
High-Availability	Indicates that the High-Availability feature is degraded.
iSCSI Service	Indicates that the iSCSI initiators are not accessible. Review the iSCSI configuration in Core. The iSCSI initiators might have been removed.
Licensing	<p>Enables an alarm and sends an email notification if the appliance is unlicensed, if there is an issue with the autolicense, the licenses have expired, the licenses are about to expire, or the model is unlicensed.</p> <p>By default, all Licensing alarms are enabled.</p>

Control	Description
Link Duplex	<p>Indicates that an interface was not configured for half-duplex negotiation but has negotiated half-duplex mode. Half-duplex significantly limits the optimization service results.</p> <p>The alarm displays which interface is triggering the duplex error.</p> <p>Choose Configure &gt; Networking: Data Interfaces and examine the Core link configuration. Next, examine the peer switch user interface to check its link configuration. If the configuration on one side is different from the other, traffic is sent at different rates on each side, causing many collisions.</p> <p>To troubleshoot, change both interfaces to automatic duplex negotiation. If the interfaces do not support automatic duplex, configure both ends for full duplex.</p> <p>You can enable or disable the alarm for a specific interface. To disable an alarm, choose Settings &gt; System Settings: Alarms and select or clear the check box next to the link alarm.</p>
Link I/O Errors	<p>Indicates that the error rate on an interface has exceeded 0.1 percent while either sending or receiving packets. This threshold is based on the observation that even a small link error rate reduces TCP throughput significantly. A properly configured LAN connection experiences very few errors. The alarm clears when the error rate drops below 0.05 percent.</p> <p>You can change the default alarm thresholds by entering the <b>alarm error-threshold</b> CLI command at the system prompt. For details, see the <i>SteelFusion Command-Line Interface Reference Manual</i>.</p> <p>To troubleshoot, try a new cable and a different switch port. Another possible cause is electromagnetic noise nearby.</p> <p>You can enable or disable the alarm for a specific interface: for example, you can disable the alarm for a link after deciding to tolerate the errors. To enable or disable an alarm, choose Settings &gt; System Settings: Alarms and select or clear the check box next to the link name.</p>
Link State	<p>Indicates that the system has lost one of its Ethernet links due to an unplugged cable or dead switch port. Check the physical connectivity between the appliance and its neighbor device. Investigate this alarm as soon as possible. Depending on which link is down, the system might no longer be optimizing and a network outage could occur.</p> <p>You can enable or disable the alarm for a specific interface. To enable or disable the alarm, choose Settings &gt; System Settings: Alarms and select or clear the check box next to the link name.</p>

Control	Description
LUN Status	<p>Indicates that a LUN is having any of these issues:</p> <ul style="list-style-type: none"> <li>• A LUN is deactivated and unavailable. A LUN will be deactivated if the blockstore has a critical amount of low space and this particular LUN has a high rate of new writes.</li> <li>• Initialization of the blockstore for the LUN fails, making the LUN unavailable.</li> </ul> <p>Check if the data center LUN was offlined on the Core while IO operations were in progress. Reactivate the LUN through the Management Console or the CLI to troubleshoot this issue.</p> <ul style="list-style-type: none"> <li>• A Resize alarm will be triggered for a LUN if its size is changed on the storage array and the Core is not able to make the new size available to the branch client. Some reasons why a resize may not be propagated to the branch are: <ul style="list-style-type: none"> <li>• The size of the LUN on the storage array is reduced.</li> <li>• The increased size of a pinned LUN cannot be accommodated in the Edge blockstore.</li> <li>• In the FusionSync (replication) configuration, the replica LUN size is smaller than the primary LUN size.</li> </ul> </li> </ul> <p>In FusionSync, due to the allowed replica leeway when configuring replication, the replica LUN on the secondary data center can be larger than the LUN in the Core configuration (which is the size of the LUN on the primary data center). If the primary data center goes down and you fail over to the secondary data center, the LUN size on secondary will show as larger than the configured LUN size, causing the Resize alarm to be triggered.</p>
Memory Paging	<p>Indicates extended memory paging activity.</p> <p>If 100 pages are swapped every couple of hours, the appliance is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support.</p>
Process Dump Creation Error	<p>Indicates that the system detected an error while trying to create a process dump.</p> <p>This alarm indicates an abnormal condition in which the system cannot collect the core file after three retries. This condition can be caused when the /var directory reaches capacity. When the alarm is raised, the directory is blacklisted.</p>
Secure Vault	<p><b>Secure Vault Locked</b> - Indicates that the secure vault is locked. To optimize SSL connections or to use RiOS data store encryption, the secure vault must be unlocked. Go to Settings &gt; Security: Secure Vault and unlock the secure vault.</p>
Snapshot	<p>Indicates that the connection to one or more of the snapshot storage arrays has failed.</p>
SSL	<p>Indicates that the system detected an error in your SSL configuration.</p>
SteelFusion Core configuration status	<p>Indicates that the Core configuration has been reverted to a previous version and all connections to the Edges are lost. Contact Riverbed Support at <a href="https://support.riverbed.com">https://support.riverbed.com</a>.</p>
SteelFusion Core Service	<p>Indicates that the Core service is not running.</p>
Temperature	<ul style="list-style-type: none"> <li>• <b>Critical Temperature</b> - Indicates that the CPU temperature exceeds the rising threshold. When the CPU returns to the reset threshold, the critical alarm is cleared. The default value for the rising threshold temperature is 70°C; the default reset threshold temperature is 67°C.</li> <li>• <b>Warning Temperature</b> - Indicates that the CPU temperature is approaching the rising threshold. When the CPU returns to the reset threshold, the warning alarm is cleared.</li> </ul> <p>After the alarm triggers, it cannot trigger again until after the temperature falls below the reset threshold and then exceeds the rising threshold again.</p>

## Viewing CPU Utilization Reports

The CPU Utilization report summarizes the percentage of the CPU used within the time period specified.

Typically, the Core operates on approximately 30 to 40 percent of total CPU capacity during nonpeak hours and approximately 60 to 70 percent capacity during peak hours. No single Core CPU usage should exceed 90 percent.

### What This Report Tells You

The CPU Utilization report answers these questions:

- How much of the CPU is being used?
- What is the average and peak percentage of the CPU being used?

### About Report Graphs

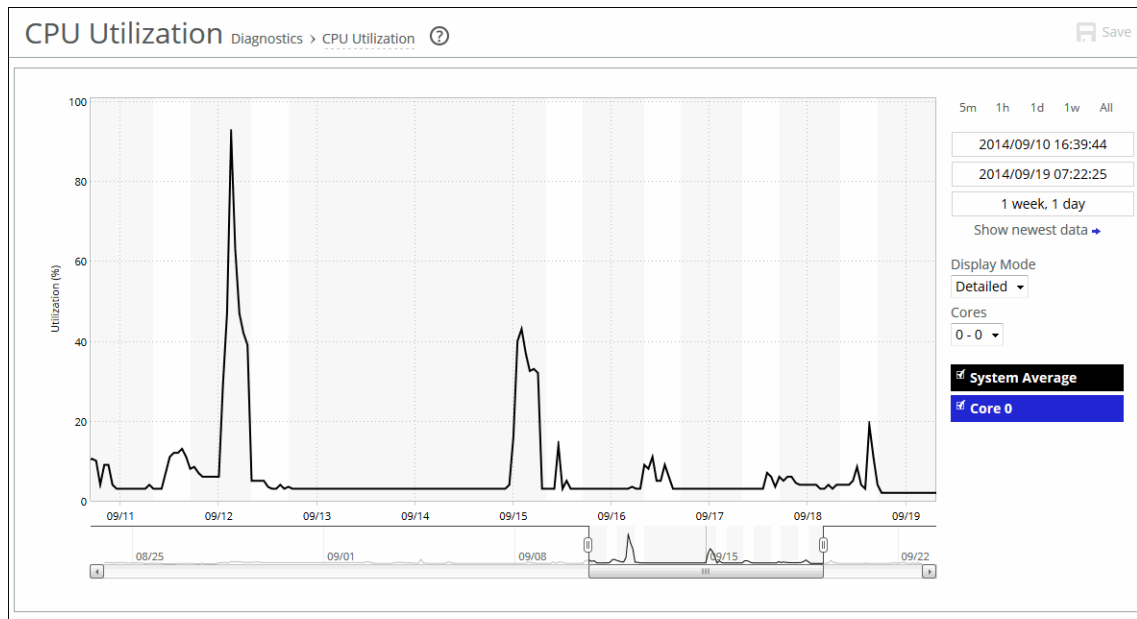
In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

The right margin of the graph points to the value on the y-axis (for example, the percent) that is the average value for the time period selected.

### To view the CPU Utilization report

1. Choose Reports > Diagnostics: CPU Utilization to display the CPU Utilization page.

Figure 7-12. CPU Utilization Page



2. Customize the report using the controls described in this table.

Control	Description
Period settings	Click <b>5m</b> (last five minutes), <b>1h</b> (last hour), <b>1d</b> (last day), <b>1w</b> (last week), or <b>All</b> (last month). For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss. You can also use the adjustable slider at the bottom of the graph.
Display Mode	Select a display mode from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Brief</b> - Displays the CPU utilization percentage of all cores as a system average.</li> <li>• <b>Detailed</b> - Displays the CPU utilization percentages for individual cores. The individual cores appear with a number and a color in the data series. To hide or display a core in the plot area, select or clear the check box next to the core name.</li> </ul>
Cores	Displays which cores are available for detailed side-by-side comparison in the graph. <b>Note:</b> This option is only available if you are in Detailed display mode.

## Viewing Memory Paging Reports

The Memory Paging report provides the total number of memory pages, per second, utilized in the time period specified. It displays the following statistics for the time period you specify.

Field	Description
Total Pages Swapped Out	Specifies the total number of pages swapped. If 100 pages are swapped approximately every two hours, the Core is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support at <a href="https://support.riverbed.com">https://support.riverbed.com</a> .
Average Pages Swapped Out	Specifies the average number of pages swapped. If 100 pages are swapped every two hours, the Core is functioning properly. If thousands of pages are swapped every few minutes, contact Riverbed Support at <a href="https://support.riverbed.com">https://support.riverbed.com</a> .
Peak Pages Swapped Out At <time> on <date>	Specifies the date and time that the peak number of pages were swapped.

## What This Report Tells You

The Memory Paging report answers these questions:

- How much memory is being used?
- What is the average and peak number of memory pages swapped?

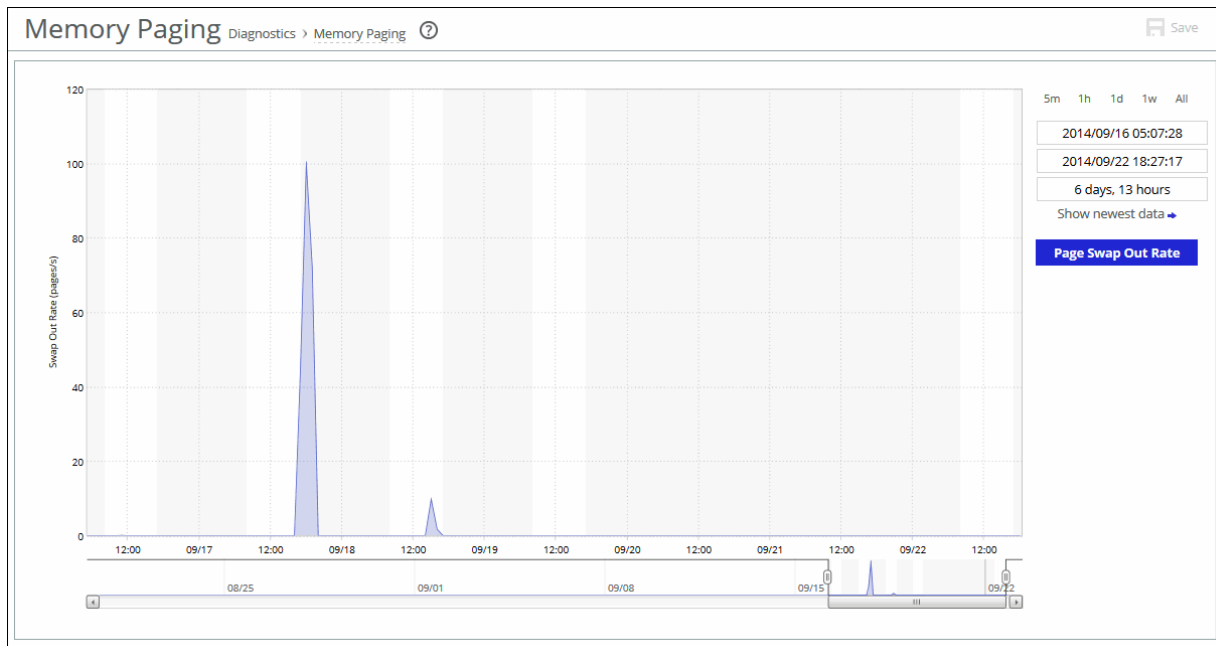
## About Report Graphs

In bar-graph and line-graph reports, the x-axis (or tick mark) plots time, according to the interval you select. The y-axis plots the metric of interest, such as gigabytes of bandwidth, percent of data reduction, connection counts, and so on.

## To view the Memory Paging report

1. Choose Reports > Diagnostics: Memory Paging to display the Memory Paging page.

Figure 7-13. Memory Paging Page



2. Customize the report by modifying the period settings:

- To the right of the table, click **5m** (last five minutes), **1h** (last hour), **1d** (last day), **1w** (last week), or **All** (last month).

For a custom time period, specify the start and end times in the fields immediately below these settings. Use the following format: yyyy/mm/dd hh:mm:ss.

- Use the adjustable slider at the bottom of the graph to modify the time period.

## Viewing Logs

This section describes how to view both user and system logs, and how to download log files. Core log reports provide a high-level view of network activity, and system files help diagnose problems. This section describes how to perform the following tasks:

- [“Viewing User Logs” on page 172](#)
- [“Viewing System Logs” on page 174](#)
- [“Downloading Log Files” on page 175](#)

## Viewing User Logs

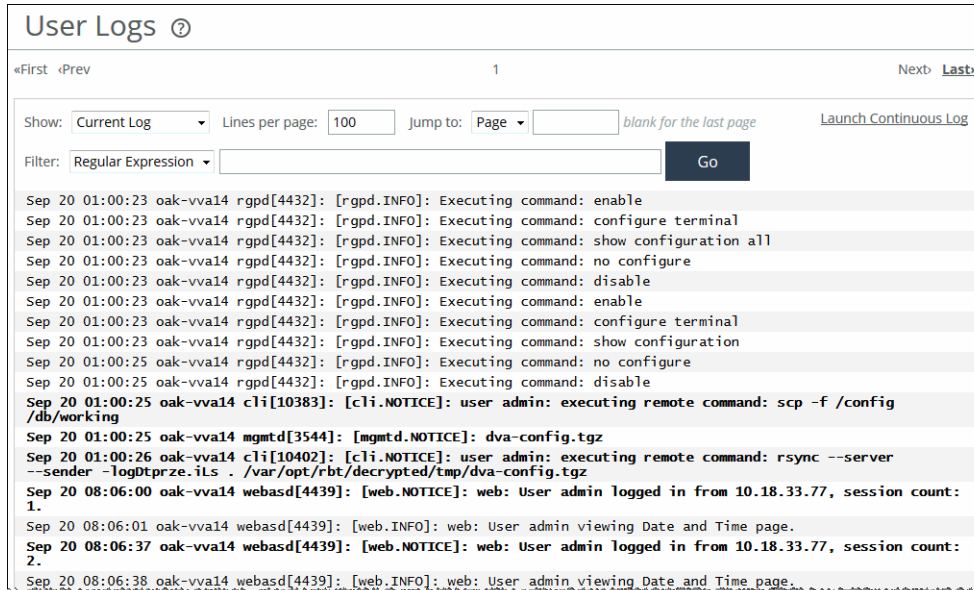
You can view user logs in the Reports > Logs: User Logs page. The user log filters messages from the system log to display messages that are of immediate use to the system administrator.

View user logs to monitor system activity and to troubleshoot problems: for example, you can monitor who logged in, who logged out, and who entered particular CLI commands or caused particular alarms. The most recent log events are listed first.

### To view and customize user logs

1. Choose Reports > Logs: User Logs to display the User Logs page.

Figure 7-14. User Logs Page



2. Customize the log as described in this table.

Control	Description
Show	Select one of the archived logs or Current Log from the drop-down list.
Lines per Page	Specify the number of lines that you want to display in the page.
Jump to	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Page</b> - Specify the number of pages that you want to display.</li> <li>• <b>Time</b> - Specify the time for the log that you want to display.</li> </ul>
Filter	Select one of the following filtering options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Regular expression</b> - Specify a regular expression on which to filter the log.</li> <li>• <b>Error or higher</b> - Displays Error-level logs or higher.</li> <li>• <b>Warning or higher</b> - Displays Warning-level logs or higher.</li> <li>• <b>Notice or higher</b> - Displays Notice-level logs or higher.</li> <li>• <b>Info or higher</b> - Displays Info-level logs or higher.</li> </ul>
Go	Displays the report.

You can continuously display new lines as the log grows and appends new data.

### To view a continuous log

1. Choose Reports > Logs: User Logs to display the User Logs page.

2. Customize the log using the controls described in [“To view and customize user logs” on page 173](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

If the continuous log does not appear, a pair of Cores might be optimizing HTTP traffic between the user's web browser and the primary or auxiliary interface of the Core for which the user is viewing the log, and they are buffering the HTTP response.

To display the continuous log, you can switch to HTTPS, because the Cores do not optimize HTTPS traffic. Alternatively, you can configure the other Cores to pass through traffic on the primary or auxiliary interfaces for port 80.

## Viewing System Logs

You can view system logs in the Reports > Logs: System Logs page. View system logs to monitor system activity and to troubleshoot problems. The most recent log events are listed first.

### To customize system logs

1. Choose Reports > Logs: System Logs to display the System Logs page.

Figure 7-15. System Logs Page

2. Customize the report using the controls described in this table.

Control	Description
Current Log	Select one of the archived logs or Current Log from the drop-down list.
Lines per page	Specify the number of lines that you want to display in the page.
Jump to	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Page</b> - Specify the number of pages that you want to display.</li> <li>• <b>Time</b> - Specify the time for the log that you want to display.</li> </ul>
Filter	Select one of the following filtering options from the drop-down list: <ul style="list-style-type: none"> <li>• <b>Regular expression</b> - Specify a regular expression on which to filter the log.</li> <li>• <b>Error or higher</b> - Displays Error-level logs or higher.</li> <li>• <b>Warning or higher</b> - Displays Warning-level logs or higher.</li> <li>• <b>Notice or higher</b> - Displays Notice-level logs or higher.</li> <li>• <b>Info or higher</b> - Displays Info-level logs or higher.</li> </ul>
Go	Displays the report.

### To view a continuous log

1. Choose Reports > Logs: System Logs to display the System Logs page.
2. Customize the log using the controls described in [“To customize system logs” on page 174](#).
3. Click **Launch Continuous Log** in the upper-right corner of the page.

If the continuous log does not appear, a pair of Cores might be optimizing the HTTP traffic between the user's web browser and the primary or auxiliary interface of the Core for which the user is viewing the log, and they are buffering the HTTP response.

To display the continuous log, you can switch to HTTPS, because the Cores do not optimize HTTPS traffic. You might want to configure the other Cores to pass through traffic on the primary or auxiliary interfaces for port 80.

## Downloading Log Files

This section describes how to download user and system log files.

- [“Downloading User Log Files” on page 175](#)
- [“Downloading System Log Files” on page 176](#)

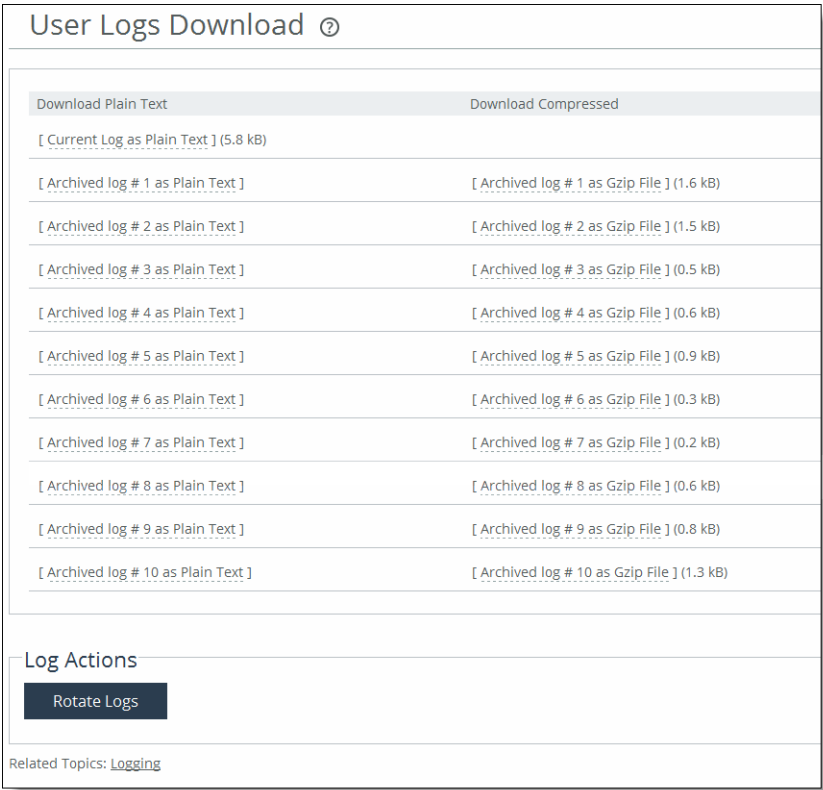
### Downloading User Log Files

You can download user logs in the User Logs Download page. Download user logs to monitor system activity and to troubleshoot problems.

To download user logs

- 1. Choose Reports > Logs: User Logs Download to display the User Logs Download page.

Figure 7-16. User Logs Download Page



- 2. Click the filename to open a file, or save the file to disk.
- 3. Click **Rotate Logs** to manually archive the current log to a numbered archived log file and then clear the log so that it is empty again.

Downloading System Log Files

You can download system logs in the System Logs Download page. Download system logs to monitor system activity and to troubleshoot problems.

## To download system logs

1. Choose Reports > Logs: System Logs Download to display the System Logs Download page.

Figure 7-17. System Logs Download Page

System Logs Download ⓘ	
Download Plain Text	Download Compressed
[ Current Log as Plain Text ] (3.2 MB)	
[ Archived log # 1 as Plain Text ]	[ Archived log # 1 as Gzip File ] (0.1 MB)
[ Archived log # 2 as Plain Text ]	[ Archived log # 2 as Gzip File ] (0.1 MB)
[ Archived log # 3 as Plain Text ]	[ Archived log # 3 as Gzip File ] (0.1 MB)
[ Archived log # 4 as Plain Text ]	[ Archived log # 4 as Gzip File ] (0.1 MB)
[ Archived log # 5 as Plain Text ]	[ Archived log # 5 as Gzip File ] (0.1 MB)
[ Archived log # 6 as Plain Text ]	[ Archived log # 6 as Gzip File ] (103.5 kB)
[ Archived log # 7 as Plain Text ]	[ Archived log # 7 as Gzip File ] (103.0 kB)
[ Archived log # 8 as Plain Text ]	[ Archived log # 8 as Gzip File ] (0.1 MB)
[ Archived log # 9 as Plain Text ]	[ Archived log # 9 as Gzip File ] (0.1 MB)
[ Archived log # 10 as Plain Text ]	[ Archived log # 10 as Gzip File ] (0.1 MB)
Log Actions	
Rotate Logs	

2. Click the log name in the Download Plain Text column or in the Download Compressed column.
3. Open or save the log (these procedures vary depending on which browser you are using).
4. Click **Rotate Logs** to manually archive the current log to a numbered archived log file and then clear the log so that it is empty again.

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

## Generating Dumps

This section describes how to generate and download system, process, and TCP dumps. It includes the following sections:

- [“Generating System Dumps” on page 178](#)
- [“Viewing Process Dumps” on page 179](#)
- [“Capturing and Uploading TCP Dumps” on page 179](#)
- [“Viewing a TCP Dump” on page 185](#)

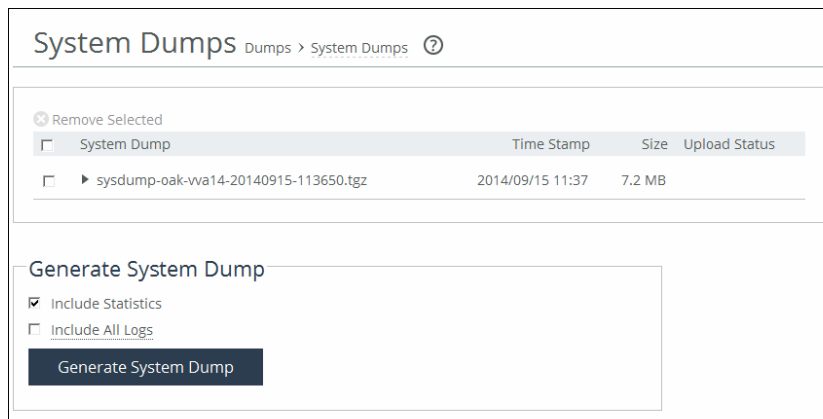
## Generating System Dumps

A system dump contains a copy of the kernel data on the system. System dump files can help you diagnose problems in the system.

### To view system dump files

1. Choose Reports > Dumps: System Dumps to display the System Dumps page.

Figure 7-18. System Dumps Page



2. Under Generate System Dump, select the type of information to include in the report:
  - **Include Statistics** - Select to collect and include CPU, memory, and other statistics in the system dump. This option is enabled by default. These statistics are useful while analyzing traffic patterns to correlate to an issue. The system adds the statistics to a file in the system dump called stats.tgz.
  - **Include All Logs** - Removes the 50 MB limit for compressed log files to include all logs in the system dump.
3. Click **Generate System Dump**.

When the system dump is complete, it appears in the list of links to download.

### To view a previously saved system dump

- Click the filename in the System Dump column to open a file or save the file to disk.

### To download a system dump file to Riverbed support

1. Choose Reports > Dumps: System Dumps to display the System Dumps page.
2. Click **Download** to receive a copy of the previously saved system dump.
3. Select the filename to open a file or save the file to disk.
4. To remove a log, check the box next to the name and click Remove Selected.

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

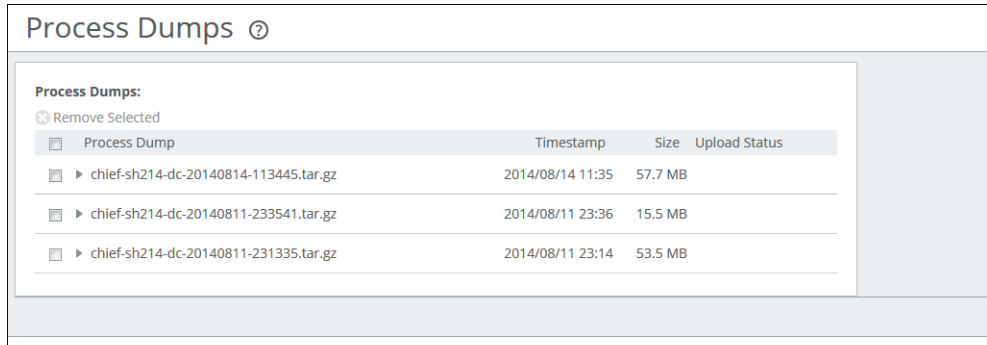
## Viewing Process Dumps

You can display and download process dumps in the Process Dumps page. A process dump is a saved copy of memory, including the contents of all memory, bytes, hardware registers, and status indicators. It is periodically performed to restore the system in the event of failure. Process dump files can help you diagnose problems in the system.

### To view process dump files

1. Choose Reports > Dumps: Process Dumps to display the Process Dumps page.

Figure 7-19. Process Dumps Page



Process Dump	Timestamp	Size	Upload Status
<input type="checkbox"/> ▶ chief-sh214-dc-20140814-113445.tar.gz	2014/08/14 11:35	57.7 MB	
<input type="checkbox"/> ▶ chief-sh214-dc-20140811-233541.tar.gz	2014/08/11 23:36	15.5 MB	
<input type="checkbox"/> ▶ chief-sh214-dc-20140811-231335.tar.gz	2014/08/11 23:14	53.5 MB	

2. Select the filename to open a file, or save the file to disk.

To remove an entry, check the box next to the name and click **Remove Selected**.

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

## Capturing and Uploading TCP Dumps

You can capture, download, and upload TCP dumps in the Reports > Dumps: TCP Dumps page. TCP trace dump files contain summary information for every Internet packet received or transmitted on the interface. TCP trace dump files can help diagnose problems in the system.

You can easily capture and retrieve multiple TCP trace dumps from the Management Console. You can generate trace dumps from multiple interfaces at the same time, limit the size of the trace dump, and schedule a specific date and time to generate a trace dump. Scheduling and limiting a trace dump by time or size allows unattended captures.

The top of the TCP Dumps page displays a list of existing TCP trace dumps and the bottom of the page displays controls to create a new trace dump. The bottom of the page also includes the trace dumps that are currently running. The Running Capture Name list includes TCP trace dumps running at a particular time. It includes TCP trace dumps started manually and also any dumps that were scheduled previously and are now running.

## To capture TCP trace dumps

1. Choose Reports > Dumps: TCP Dumps to display the TCP Dumps page.

Figure 7-20. TCP Dumps Page

### TCP Dumps

Dumps > TCP Dumps

Remove Selected

TCP Dump	Time Stamp	Size	Upload Status
No stored TCP dumps.			

TCP Dumps Currently Running:

Add a New TCP Dump Stop Selected Captures

Name

Capture Name:  
Dumpfile

Endpoints

Capture traffic between:

IPs: All

Ports: All

and:

IPs: All

Ports: All

Capture Interfaces

All Interfaces

Base Interfaces:  
☐ primary  
☐ aux

Data Interfaces:  
☐ eth0\_0  
☐ eth0\_1  
☐ eth0\_2  
☐ eth0\_3

Capture Parameters

☐ Capture Untagged Traffic Only

☐ Capture VLAN-Tagged Traffic Only

☒ Capture both VLAN and Untagged Traffic

Capture Duration: 30 seconds (Enter "0" or "continuous" for no time limit)

Maximum Capture Size: 100 MB

Buffer Size: 154 kB

Snap Length (bytes)

☐ 0

Use the default of 65535. This is recommended for CIFS, MAPI, and SSL traces.

☒ 1518

## 2. Configure the dump using the controls described in this table.

Control	Description
Add a New TCP Dump	Displays the controls for creating a TCP trace dump.
Name	<p>Specify the name of the capture file in the Capture Name field. The default filename uses the following format:</p> <p>&lt;hostname&gt;_&lt;interface&gt;_&lt;timestamp&gt;.cap</p> <p>where &lt;hostname&gt; is the hostname of the Core, &lt;interface&gt; is the name of the interface selected for the trace (for example, lan0_0, wan0_0), and &lt;timestamp&gt; is in the yyyy/mm/dd hh:mm:ss format.</p> <p>If this trace dump relates to an open Riverbed Support case, specify the capture filename case_&lt;number&gt;, where &lt;number&gt; is your Riverbed Support case number: for example, case_12345.</p> <p>The .cap file extension is not included with the filename when it appears in the capture queue.</p>
Endpoints	<p>Specify the source and destination endpoints for the dump:</p> <ul style="list-style-type: none"> <li>• <b>IPs</b> - Specify the source IP addresses. Separate multiple IP addresses with a comma to include all addresses bidirectionally. The default setting is all IP addresses.</li> <li>• <b>Ports</b> - Specify the source ports. Separate multiple ports with a comma. The default setting is all ports.</li> <li>• <b>IPs</b> - Specify the destination IP addresses. Separate multiple IP addresses with a comma to include all addresses bidirectionally. The default setting is all IP addresses.</li> <li>• <b>Ports</b> - Specify the destination ports. Separate multiple ports with a comma. The default setting is all ports.</li> </ul>
Capture Interfaces	Captures the TCP trace dump on the selected interface. You can select a physical, MIP, or RSP interface. Click only one interface per trace dump. The default setting is none. You must specify a capture interface.

Control	Description
Capture Parameters	<p>These parameters let you capture information about dot1q VLAN traffic. You can match traffic based on VLAN-tagged or untagged packets, or both. You can also filter by port number or host IP address and include or exclude ARP packets. Select one of these parameters for capturing VLAN packets:</p> <ul style="list-style-type: none"> <li>• <b>Capture Untagged Traffic Only</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– All untagged VLAN traffic.</li> <li>– Untagged 7850 traffic and ARP packets. You must also specify <b>or arp</b> in the custom flags field in this page.</li> <li>– Only untagged ARP packets. You must also specify <b>and arp</b> in the custom flags field in this page.</li> </ul> </li> <li>• <b>Capture VLAN-Tagged Traffic Only</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– Only VLAN-tagged traffic.</li> <li>– VLAN-tagged packets with host 10.11.0.6 traffic and ARP packets. You must also specify <b>10.11.0.6</b> in the IPs field, and specify <b>or arp</b> in the custom flags field in this page.</li> <li>– VLAN-tagged ARP packets only. You must also specify <b>and arp</b> in the custom flags field in this page.</li> </ul> </li> <li>• <b>Capture both VLAN and Untagged Traffic</b> - Select this option for the following captures: <ul style="list-style-type: none"> <li>– All VLAN traffic.</li> <li>– Both tagged and untagged 7850 traffic and ARP packets. You must also specify the following parameters in the custom flags field in this page: (port 7850 or arp) or (vlan and (port 7850 or arp))</li> <li>– Both tagged and untagged 7850 traffic only. You must also specify <b>7850</b> in one of the port fields in this page. No custom flags are required.</li> <li>– Both tagged and untagged ARP packets. You must also specify the following parameters in the custom flags field in this page: (arp) or (vlan and arp)</li> </ul> </li> </ul>

Control	Description
Capture Parameters	<p>Complete the following settings:</p> <ul style="list-style-type: none"> <li>• <b>Capture Duration (Seconds)</b> - Specify how long the capture runs, in seconds. The default value is 30. Leave this value blank to initiate a continuous trace. When a continuous trace reaches the maximum space allocation of 100 MB, the oldest file is overwritten.</li> <li>• <b>Maximum Capture Size (MB)</b> - Specify the maximum capture file size, in megabytes. The default value is 100. The recommended maximum capture file size is 1024 MB (1 GB).</li> <li>• <b>Buffer Size</b> - Optionally, specify the maximum number of packets allowed to queue while awaiting processing by the TCP trace dump. The default value is 154.</li> <li>• <b>Snap Length</b> - Optionally, select the snap length value for the capture file or specify a custom value. The snap length equals the number of bytes the report captures for each packet. Having a snap length smaller than the maximum packet size on the network enables you to store more packets, but you might not be able to inspect the full packet content. The default value is 1518 bytes.</li> </ul> <p>Select 0 for a full packet capture (recommended for CIFS, MAPI, and SSL captures).</p> <p>When using jumbo frames, Riverbed recommends selecting 9018. The default custom value is 16383 bytes.</p> <ul style="list-style-type: none"> <li>• <b>Number of Files to Rotate</b> - Specify how many TCP trace dump files to rotate. The default value is 5.</li> <li>• <b>Custom Flags</b> - Specify custom flags as additional statements within the filter expression. Custom flags are added to the end of the expression created from the Endpoints fields and the Capture Parameters radio buttons (pertaining to VLANs).</li> </ul> <p>If you require an “and” statement between the expression created from other fields and the expression that you are entering in the custom flags field, you must include the “and” statement at the start of the custom flags field.</p> <p>For complete control of your filter expression, use the CLI <b>tcpdump</b> command. For details, see the <i>Riverbed Command-Line Interface Reference Manual</i>.</p>
Schedule	<p>Schedules a trace dump to run at a later date and time.</p> <ul style="list-style-type: none"> <li>• <b>Schedule Dump</b> - Enables the scheduling feature.</li> <li>• <b>Start Date</b> - Specifies the start date for the scheduled trace dump.</li> </ul>
Add	Adds the TCP trace dump to the capture queue.

## Stopping a TCP Dump After an Event Occurs

Capture files offer visibility into intermittent network issues, but the amount of traffic they capture can be overwhelming. Also, because rotating logs is common, after a capture logs an event, the Core log rotation can overwrite debugging information specific to the event.

The Core makes troubleshooting easier because it provides a trigger that can stop a continuous capture after a specific log event occurs. The result is a smaller file to help pinpoint what makes the event happen.

The stop trigger continuously scans the system logs for a search pattern. When it finds a match, it stops all running captures.

### To stop a capture after a specific log event

1. Choose Reports > Dumps: TCP Dumps to display the TCP Dumps page.

## 2. Schedule a capture.

**Figure 7-21. TCP Dump Stop Trigger**

**TCP Dump Stop Trigger**

Continuously scan System Logs for a pattern and stop all running TCP Dumps when there's a match.

Status: Not Running

Last Pattern:

Last Triggered: Never

Pattern:  (Perl regex)

Delay:  seconds

When triggered, a notification is sent to the event notification email address specified on the [Email Page](#).

This is typically used with TCP Dumps with "Capture Duration" set to "continuous" seconds to keep the traces from stopping on their own before there's a match.

3. In the Pattern text box, enter a Perl regular expression (regex) to find in a log. The Core compares the Perl regex against each new line in the system logs and the trigger stops if it finds a match.

The simplest regex is a word or a string of characters. For example, if you set the pattern to Limit, the trigger matches the line Connection Limit Reached.

Notes:

- Perl regular expressions are case sensitive.
- Perl treats the space character like any other character in a regex.
- Perl reserves some characters, called metacharacters, for use in regex notation. The metacharacters are:

{ } [ ] ( ) ^ \$ . | \* + ? \

You can match a metacharacter by putting a backslash before it. For example, to search for a backslash in the logs, you must enter two backslashes (\\) as the pattern.

- The pattern follows Perl regular expression syntax. For details, go to:  
<http://perldoc.perl.org/perlre.html>
- You cannot change the pattern while a scan is running. You must stop the scan before changing a pattern.
- You do not need to wrap the pattern with the metacharacters to match the beginning or end of a line (^ \$) or with the wildcard character (\*).

4. Specify the amount of time to pause before stopping all running captures when the Core finds a match. This gives the system some time to log more data without abruptly cutting off the capture. The default is 30 seconds. Specify 0 for no delay; the capture stops immediately.

After a trigger has fired, the capture can stop by itself before the delay expires; for example, the capture duration can expire.

## 5. Click **Start Scan**.

When the scan stops, the Core sends an email to all email addresses on the Settings > System Settings: Email page appearing under Report Events via Email. The email notifies users that the trigger has fired.

The page indicates Last Triggered: Never if a TCP Dump stop trigger has never triggered on the Core. After the delay duration of the stop trigger, the Core displays the last triggered time.

Before changing the Perl regular expression or amount of delay, you must first stop the process.

### To stop a running scan

- Click **Stop Scan** to halt the background process that monitors the system logs. The Core dims this button when the stop trigger is idling.

## Stop Trigger Limitations

These limitations apply to the trigger:

- You cannot create a trigger to stop a specific capture; the trigger affects all running captures.
- If the search pattern contains a typo, the trigger might never find a match.
- Only one instance of a trigger can run at one time.

## Viewing a TCP Dump

The top of the TCP Dumps page displays a list of existing captures.

### To view TCP trace dump files

1. Choose Reports > Dumps: TCP Dumps to display the TCP Dumps page.
2. Under Download Link, click the trace dump name to open the file.

### To stop a running TCP trace dump

1. Choose Reports > Dumps: TCP Dumps to display the TCP Dumps page.
2. Click the trace dump filename in the Running Capture Name list.
3. Click **Stop Selected Captures**.

### To upload the trace to Riverbed Support

In continuous mode, after you complete the capture, perform the following steps (for timed TCP dumps, start with Step 2):

1. On the TCP Dumps page, select the running TCP dump and click **Stop Selected Captures**.  
The trace appears as a download link in the list of TCP dumps stored on the Core.
2. Click the top file in the TCP Dumps list and save it locally.  
This file should contain the current date.
3. Compress (zip) the file and follow the upload instructions to share it with Riverbed Support.  
Attach the files to your case at <https://support.riverbed.com/cases/viewcases.htm> or upload the file to <ftp://ftp.riverbed.com/incoming> (for FTP, be sure the file is prefixed with case\_#).  
ftp ftp.riverbed.com  
User: anonymous

```
Password: your_email@address
ftp> cd /incoming
ftp> bi
ftp> put case_12345-tcpdump.zip
```

## APPENDIX A SteelFusion Core MIB

This appendix provides a reference to the Core MIB and SNMP traps. These tools allow for easy management of Cores and straightforward integration into existing network management systems.

This appendix includes the following sections:

- [“Accessing the Core MIB” on page 187](#)
- [“SNMP Traps” on page 188](#)

---

### Accessing the Core MIB

The Core MIB monitors device status and peers, and it provides network statistics for seamless integration into network management systems such as Hewlett Packard OpenView Network Node Manager, PRTG, and other SNMP browser tools.

For details about configuring and using these network monitoring tools, consult their product documentation.

The following guidelines describe how to download and access the Core MIB using common MIB browsing utilities:

- You can download the Core MIB file (GC-MIB.txt) from the Support page of the Management Console or from the Riverbed Support site at <https://support.riverbed.com> and load it into any MIB browser utility.
- Some utilities might expect a file type other than a text file. If this occurs, change the file extension to the type required by the utility you have chosen.
- Some utilities assume that the root is **mib-2** by default. If the utility sees a new node, such as **enterprises**, it might look under **mib-2.enterprises**. If this occurs, use **.iso.org.dod.internet.private.enterprises.rbt** as the root.
- Some command-line browsers might not load all MIB files by default. If this occurs, find the appropriate command option to load the GC-MIB.txt file: for example, for NET-SNMP browsers, **snmpwalk -m all**.

---

## SNMP Traps

Every Core supports SNMP traps and email alerts for conditions that require attention or intervention. An alarm triggers for most, but not every, event, and the related trap is sent. For most events, when the condition clears, the system clears the alarm and also sends a clear trap. The clear traps are useful in determining when an event has been resolved.

This section describes the SNMP traps. It does not list the corresponding clear traps.

RiOS v6.0 and later includes support for SNMPv3.

You can view Core health at the top of each Management Console page, by entering the CLI **show info** command, and through SNMP (health, systemHealth).

The Core tracks key hardware and software metrics and alerts you of any potential problems so that you can quickly discover and diagnose issues. The health of an appliance falls into one of the following states:

- **Healthy** - The Core is functioning optimally.
- **Needs Attention** - Accompanies a healthy state to indicate management-related issues not affecting the ability of Core to perform.
- **Degraded** - The Core system has detected an issue.
- **Admission Control** - The Core is performing but has reached its connection limit.
- **Critical** - The Core might or might not be performing; you must address a critical issue.

The following table summarizes the SNMP traps sent from the system to configured trap receivers and their effect on the Core health state.

Trap and OID	Appliance State	Text	Description
procCrash (enterprises.17163.1.100.4.0.1)	Healthy	A procCrash trap signifies that a process managed by PM has crashed and left a core file. The variable sent with the notification indicates which process crashed.	A process has crashed and subsequently been restarted by the system. The trap contains the name of the process that crashed. A system snapshot associated with this crash has been created on the appliance and is accessible through the CLI or the Management Console. Riverbed Support might need this information to determine the cause of the crash. No other action is required on the appliance as the crashed process is automatically restarted.
procExit (enterprises.17163.1.100.4.0.2)	Healthy	A procExit trap signifies that a process managed by PM has exited unexpectedly, but not left a core file. The variable sent with the notification indicates which process exited.	A process has unexpectedly exited and been restarted by the system. The trap contains the name of the process. The process might have exited automatically or due to other process failures on the appliance. Review the release notes for known issues related to this process exit. If none exist, contact Riverbed Support to determine the cause of this event. No other action is required on the appliance as the crashed process is automatically restarted.
configChange (enterprises.17163.1.100.4.0.3)	Healthy	A change has been made to the system's configuration.	A configuration change has been detected. Check the log files around the time of this trap to determine what changes were made and whether they were authorized.
cpuUtil (enterprises.17163.1.100.4.0.4)	Degraded	The average CPU utilization in the past minute has gone above the acceptable threshold.	Average CPU utilization has exceeded an acceptable threshold. If CPU utilization spikes are frequent, it might be because the system is undersized. Sustained CPU load can be symptomatic of more serious issues. Consult the CPU Utilization report to gauge how long the system has been loaded and also monitor the amount of traffic currently going through the appliance. A one-time spike in CPU is normal but Riverbed recommends reporting extended high CPU utilization to Riverbed Support. No other action is necessary as the alarm clears automatically.
pagingActivity (enterprises.17163.1.100.4.0.5)	Degraded	The system has been paging excessively (thrashing).	The system is running low on memory and has begun swapping memory pages to disk. This event can be triggered during a software upgrade while the optimization service is still running but there can be other causes. If this event triggers at any other time, generate a debug <b>sysdump</b> and send it to Riverbed Support. No other action is required as the alarm clears automatically.

Trap and OID	Appliance State	Text	Description
linkError (enterprises.17163.1.100.4.0.6)	Degraded	An interface on the appliance has lost its link.	<p>The system has lost one of its Ethernet links, typically due to an unplugged cable or dead switch port. Check the physical connectivity between the Core and its neighbor device. Investigate this alarm as soon as possible. Depending on what link is down, the system might no longer be optimizing and a network outage could occur.</p> <p>This is often caused by surrounding devices, like routers or switches interface transitioning. This alarm also accompanies service or system restarts on the Core.</p>
powerSupplyError (enterprises.17163.1.100.4.0.7)	Degraded	A power supply on the appliance has failed.	A redundant power supply on the appliance has failed on the appliance and needs to be replaced. Contact Riverbed Support for an RMA replacement as soon as practically possible.
fanError (enterprises.17163.1.100.4.0.8)	Degraded	A fan has failed on this appliance.	A fan is failing or has failed and needs to be replaced. Contact Riverbed Support for an RMA replacement as soon as practically possible.
memoryError (enterprises.17163.1.100.4.0.9)	Degraded	A memory error has been detected on the appliance (not supported on all models).	A memory error has been detected. A system memory stick might be failing. Try reseating the memory first. If the problem persists, contact Riverbed Support for an RMA replacement as soon as practically possible.
ipmi (enterprises.17163.1.100.4.0.10)	Degraded	An IPMI event has been detected on the appliance. Please check the details in the alarm report on the web UI.	<p>An Intelligent Platform Management Interface (IPMI) event has been detected. Check the Alarm Status page for more detail. You can also view the IPMI events on the Core, by entering the CLI command:</p> <p><b>show hardware error-log all</b></p>
localFSFull (enterprises.17163.1.100.4.0.11)	Critical	The appliance local file system is full.	<p>The appliance local file system is full. You must create more space.</p> <p><b>Note:</b> The appliance local file system contains no block files from the LUNs.</p>
temperatureCritical (enterprises.17163.1.100.4.0.12)	Critical	The system temperature has reached a critical stage.	This trap/alarm triggers a critical state on the appliance. This alarm occurs when the appliance temperature reaches 90 degrees Celsius. The temperature value is not user-configurable. Reduce the appliance temperature.

Trap and OID	Appliance State	Text	Description
temperatureWarning (enterprises.17163.1.100.4.0.13)	Degraded	The system temperature has exceeded the threshold.	The appliance temperature is a configurable notification. By default, this notification is set to trigger when the appliance reached 70 degrees Celsius. Raise the alarm trigger temperature if it is normal for the device to get that hot, or reduce its temperature.
scheduledJobError (enterprises.17163.1.100.4.0.14)	Healthy	A scheduled job has failed during execution.	A scheduled job on the system (for example, a software upgrade) has failed. To determine which job failed, use the CLI or the Management Console.
confModeEnter (enterprises.17163.1.100.4.0.15)	Healthy	A user has entered configuration mode.	A user on the system has entered a configuration mode from either the CLI or the Management Console. A log in to the Management Console by user admin sends this trap as well. This is for notification purposes only; no other action is necessary.
confModeExit (enterprises.17163.1.100.4.0.16)	Healthy	A user has exited configuration mode.	A user on the system has exited configuration mode from either the CLI or the Management Console. A log out of the Management Console by user admin sends this trap as well. This is for notification purposes only; no other action is necessary.
secureVaultLocked (enterprises.17163.1.100.4.0.17)	Critical	Secure vault is locked. The secure datastore cannot be used.	You must unlock the secure vault. For details, see <a href="#">“Unlocking the Secure Vault” on page 135</a> .
testTrap (enterprises.17163.1.100.4.0.19)	Healthy	Trap test.	An SNMP trap test has occurred on the Core. This message is informational and no action is necessary.
temperatureNonCritical (enterprises.17163.1.100.4.0.1012)	Degraded	The system temperature is no longer in a critical stage.	This message is informational and no action is necessary.
temperatureNormal (enterprises.17163.1.100.4.0.1013)	Healthy	The system temperature is back within the threshold.	This message is informational and no action is necessary.
secureVaultUnlocked (enterprises.17163.1.100.4.0.1017)	Healthy	Secure vault is unlocked. The secure data store can be used now.	This message is informational and no action is necessary.
edgeError (enterprises.17163.1.100.4.0.10500)	Critical	Edge module encountered error.	Edge module encountered an error.

Trap and OID	Appliance State	Text	Description
highAvailabilityError (enterprises.17163.1.100.4.0.10501)	Degraded or Critical	High-Availability module encountered error.	<p>A degraded state indicates one of the following conditions:</p> <ul style="list-style-type: none"> <li>• Edge heartbeat channel failure</li> <li>• High availability heartbeat timed out</li> <li>• Edge blockstore connection failure</li> </ul> <p>A critical state indicates one of the following conditions:</p> <ul style="list-style-type: none"> <li>• Edge blockstore activation failure</li> <li>• Edge blockstore local write failure</li> <li>• Edge detected split-brain</li> <li>• Edge requires activation</li> </ul>
lunError (enterprises.17163.1.100.4.0.10502)	Degraded	LUN module encountered error.	LUN module encountered an error. Check if the data center LUN was offlined in Core while I/O operations were in progress.
iscsiError (enterprises.17163.1.100.4.0.10503)	Critical	iSCSI module encountered error.	An iSCSI initiator is not accessible. Review the iSCSI configuration in Core.
snapshotError (enterprises.17163.1.100.4.0.10505)	Critical	Snapshot module encountered error.	<p>A snapshot failed to be committed to the SAN, or a snapshot has failed to complete due to Windows timing out.</p> <p>Check the Core logs for details. Retry the Windows snapshot.</p>
applianceUnlicensedError (enterprises.17163.1.100.4.0.10506)	Critical	Appliance license expired/invalid.	Appliance license expired/invalid.
modelUnlicensedError (enterprises.17163.1.100.4.0.10507)	Critical	Model license expired/invalid.	Model license expired/invalid.
blkdiskError (enterprises.17163.1.100.4.0.10508)	Critical	Block-disk module encountered error.	<p>Block-disk module encountered an error.</p> <p><b>Note:</b> This alarm applies only to Core-v implementations.</p>

Trap and OID	Appliance State	Text	Description
backupIntegrationError (enterprises.17163.1.100.4.0.10509)	Critical	Backup-integration module encountered error.	Backup-Integration module encountered error.
otherHardwareError (enterprises.17163.1.100.4.0.10510)	Either Critical or Degraded, depending on the state	Hardware error detected.	<p>Indicates that the system has detected a problem with the hardware. These issues trigger the hardware error alarm:</p> <ul style="list-style-type: none"> <li>the appliance does not have enough disk, memory, CPU cores, or NIC cards to support the current configuration</li> <li>the appliance is using a memory Dual In-line Memory Module (DIMM), a hard disk, or a NIC that is not qualified by Riverbed</li> <li>other hardware issues</li> </ul> <p>The alarm clears when you add the necessary hardware, remove the unqualified hardware, or resolve other hardware issues.</p>

