SteelHead™ SD User Guide

Models 570-SD, 770-SD, 3070-SD, SDI-2030

Version SteelHead SD 2.0, SteelConnect 2.11.1

December 2018

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®

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

Part Number
712-00328-01

Contents

Welcome to SteelHead SD

Welcome to the *SteelHead SD User Guide*. SteelHead SD combines SD-WAN and cloud networking capabilities (powered by SteelConnect) with Riverbed WAN optimization (powered by RiOS) into a single appliance.

This guide describes how to configure the Riverbed SteelHead SD (SteelHead SD) when used in conjunction with SteelConnect SDI-2030 and SDI-5030 gateways.

This guide is written for network administrators familiar with administering and managing WANs.

For a high-level look at how SteelConnect works, see the *SteelConnect Manager User Guide*.

Documentation and release notes

The most current version of all Riverbed documentation can be found on the Riverbed Support site at https://support.riverbed.com.

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 list new features, known issues, 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** - 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** - Need help with planning a migration or implementing a custom design solution? Contact Riverbed Professional Services. Email proserve@riverbed.com or go to http://www.riverbed.com/services/index.htm.
- **Documentation** - Have suggestions about the online documentation or printed materials? Send comments to techpubs@riverbed.com.

1

Introducing SteelHead SD

This topic provides an overview of the SteelHead SD. It includes these sections:

- "Introducing SteelHead SD" on page 9
- "SteelHead SD and SteelConnect feature compatibility by model" on page 11
- "SD-WAN feature restrictions for SteelHead SD 2.0" on page 12
- "SteelHead feature changes after upgrading to SteelHead SD 2.0" on page 14
- "Hardware and software requirements" on page 18

This guide doesn't provide detailed information about configuring and managing SD-WAN or WAN optimization features. For details, see the *SteelConnect Manager User Guide* and the *SteelHead User Guide*.

Introducing SteelHead SD

SteelHead SD combines SD-WAN and cloud networking capabilities (powered by SteelConnect) with Riverbed WAN optimization (powered by RiOS) into a single appliance. SteelHead SD seamlessly integrates advanced SD-WAN functionality with industry-leading WAN optimization, security, and visibility services all in one streamlined appliance. SteelHead SD WAN optimization reduces bandwidth utilization and accelerates application delivery and performance, while providing SteelConnect integration in the SteelOS environment.

SteelHead SD provides you with the ability to quickly provision branch sites and deploy applications remotely. At the same time, applications are optimized to ensure performance and reduce latency with zero touch provisioning.

Typically, SteelHead SD appliances and the SteelConnect SDI-2030 gateway are located in the branch office in conjunction with SteelConnect SDI-5030 gateways at the data center. The SteelConnect SDI-2030 gateway can also be deployed inline as a 1-Gbps data center gateway with active-active HA. The SteelConnect SDI-2030 gateway can also serve as a very large branch office box with high throughput requirements. The SteelConnect SDI-2030 gateway doesn't support WAN optimization capabilities.

SteelHead SD 2.0 advanced routing and high availability (HA) features are supported on the SteelHead SD 570-SD, 770-SD, and 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch. For details, see the *SteelHead SD User Guide* and the *SteelConnect Manager User Guide*.

Figure 1-1. SteelHead SD deployment

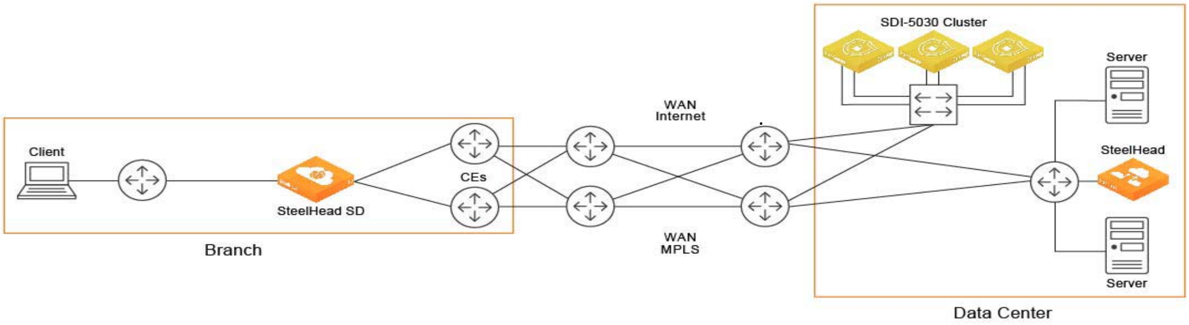

SteelHead SD supports these configuration modes:

- **SD-WAN and WAN optimization** - In this configuration, WAN optimization runs as a service on top of SD-WAN. The SteelCentral Controller for SteelHead (SCC) or the SteelHead Management Console handles management and configuration of the WAN optimization features. Also, SteelHead CLI-based management is supported for WAN optimization settings. You connect to the Management Console via the primary port, which also uses DHCP to acquire its IP address. For details on configuring WAN optimization features, see the *SteelCentral Controller for SteelHead User Guide* and the *SteelHead User Guide*.
- **SD-WAN only** - In this configuration, WAN optimization is not required. SCM handles the management and configuration of SD-WAN features. SCM connectivity requires one of the WAN ports that are used as uplink ports. Only the SD-WAN service can be enabled or disabled via SCM. The SD-WAN service upgrades are managed via SCM. SCM pushes the new software version according to the schedule that you set up. For details on configuring SD-WAN features, see the *SteelConnect Manager User Guide* and the *SteelHead SD User Guide*.

Note: In SteelConnect 2.11, SteelHead SD appliances do not perform source NATing on underlay traffic exiting via the Internet uplink if it is destined for a private address, regardless of the configured outbound NAT setting. This is a change from the previous behavior for SteelConnect 2.10 SteelHead SD appliances, if NAT was enabled for an uplink, NAT was performed for all traffic exiting via the Internet uplink. For details on configuring NAT, see in the *SteelConnect Manager User Guide*.

For details on SteelHead SD software architecture and new features for SteelHead SD 2.0, see the *SteelHead SD Installation Guide*.

SteelHead SD and SteelConnect feature compatibility by model

Feature	SteelHead 570-SD, 770-SD, 3070-SD	SDI-2030	SDI-130	SDI-330	SDI-1030	SDI-5030	Virtual GW	Cloud GW
eBGP	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
iBGP	Yes	Yes	No	No	No	No	No	No
OSPF single area	Yes	Yes	Yes	Yes	Yes	No	No	—
OSPF multi-area ABR	Yes	Yes	No	No	No	No	No	—
ASBR	Yes	Yes	Yes* (Underlay routing inter-working solution)	Yes* (Underlay routing inter-working solution)	Yes* (Underlay routing inter-working solution)	No	Yes* (Underlay routing inter-working solution)	No
Route retraction	Yes	Yes	No	No	No	Yes	No	No
Default route originate	OSPF/BGP	OSPF /BGP LAN and WAN	OSPF-only LAN	OSPF-only LAN	OSPF-only LAN	BGP only	OSPF-only LAN	No
Overlay route injection in LAN	Yes	Yes	No	No	No	Yes	No	No
Local subnet discovery	Yes	Yes	No	No	No	Yes	No	No
Static routes	Yes	Yes (LAN and WAN)	Yes (3rd-party routes)	Yes (3rd-party routes)	Yes (3rd-party routes)	Yes	Yes (3rd-party routes)	Yes (3rd-party routes)
VLAN support (LAN side)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	—
1:1 Active-Active High Availability	Yes	Yes	No (Active-Passive HA)	No (Active-Passive HA)	No (Active-Passive HA)	No (HA cluster)	No (Active-Passive HA)	No (Active-Passive HA AWS)

Feature	SteelHead 570-SD, 770-SD, 3070-SD	SDI-2030	SDI-130	SDI-330	SDI-1030	SDI-5030	Virtual GW	Cloud GW
Brownfield transit for internet-only branch	Yes (As an edge device only)	Yes	Yes (As an edge device only)	Yes (As an edge device only)	Yes	Yes	Yes (As an edge device only)	Yes (As an edge device only)
Native VLAN support	No	No	Yes	Yes	No	No	Yes	—

*SCM 2.9 and later support an underlay routing interworking solution that bridges BGP and OSPF. For details, see the *SteelConnect Manager User Guide*.

SD-WAN feature restrictions for SteelHead SD 2.0

This table summarizes the SDWAN feature restrictions for SteelHead SD 2.0.

SD-WAN feature	Description
Static uplinks on the WAN	If you have static uplinks on the WAN, a default static route is not added automatically in SteelConnect. On SCM, you must manually add static routes to reach networks that aren't present on the SteelConnect overlay network in order to send packets on those WANs. For details, see the Knowledge Base article, S32693.
WAN AutoVPN memberships	WAN AutoVPN memberships for zones are not supported on SteelHead SD 2.0 and SteelConnect 2.11 appliances.
Redirection of UDP traffic through the virtual SteelHead	Redirection of UDP traffic through the virtual SteelHead is not supported in SteelHead SD 2.0. You will not be able to optimize UDP traffic using the SteelHead IP blade.
Classic VPN	Classic VPN is not supported on SteelHead SD 2.0 and SteelConnect 2.11 appliances.
Flow distribution	Flow distribution for internet traffic across similar uplinks is not supported on SteelHead SD 570-SD, 770-SD, and 3070-SD appliances
General SD-WAN features	The following general SD-WAN features are not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030 appliances: ▪ PPPoE ▪ LTE uplinks ▪ USB port for tethering (initial ZTP/SCM via USB tethering) ▪ Cloudifi ▪ Agents tab under Sites

SD-WAN feature	Description
LAN-side settings	The following LAN-side settings are not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030 appliances: ▪ Multiple physical ports in a single zone. ▪ Spanning tree on LAN side. ▪ Multiple physical ports in a zone. ▪ Native VLANs. ▪ zones Import configuration at the Site level. ▪ xLAN option under Site configuration.
Path preference/path selection restrictions	When WAN optimization is enabled and the application target of a traffic rule is set to SSL, SteelConnect doesn't correctly classify SSL traffic and the traffic will not travel across the SteelHead optimized path. For details, see the Knowledge Base article, S32180.
Traffic path rule restrictions	When the SteelHead is located out-of-path, application-based path preference rules are not honored for deployments using WAN optimization with fixed target in-path rule to the SteelHead. You have these configuration options: ▪ Convert your deployment to an in-path or virtual in-path and adjust SteelHead SD WAN optimization in-path rules to remove the fixed target setting. ▪ Adjust the SteelHead SD WAN optimization in-path rules to pass-through and disable WAN optimization for application types you want to have follow the path preference rules.
Static uplinks on the WAN	If you have static uplinks on the WAN, a default static route is not added automatically in SteelConnect. On SCM, you must manually add static routes to reach networks that aren't present on the SteelConnect overlay network in order to send packets on those WANs. For details, see the Knowledge Base article, S32693.
Source NAT on underlay traffic	Source NAT on underlay traffic is not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-203 appliances. SteelHead SD appliances do not perform source NATing on underlay traffic exiting via the Internet uplink if it is destined for a private address, regardless of the configured outbound NAT setting. This is a change from the previous behavior for SteelHead SD 1.0 appliances, if NAT was enabled for an uplink, NAT was performed for all traffic exiting via the Internet uplink. For details on configuring NAT, see the *SteelConnect Manager User Guide*.
RADIUS/Authentication server under Sites configuration in SCM	RADIUS/Authentication server under Sites configuration in SCM is not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030 appliances. Consult with your Riverbed sales engineer or Riverbed Professional Services at http://www.riverbed.com/services/index.html.

SteelHead feature changes after upgrading to SteelHead SD 2.0

These tables summarize the SteelHead features after you upgrade SteelHead appliances to SteelHead SD 2.0 appliances. For details on upgrading SteelHead appliances to SteelHead SD 2.0 appliances, see the *SteelHead SD In-Field Upgrade Guide*.

Note: These tables do not summarize the feature changes that occur after you upgrade from SteelHead SD 1.0 to SteelHead SD 2.0. For details on feature changes that occur after upgrading from SteelHead SD 1.0 to SteelHead SD 2.0, see the *SteelHead SD Installation Guide*.

SteelHead features unchanged after upgrading to SteelHead SD 2.0

This table summarizes the SteelHead features that do not change when you upgrade to SteelHead SD 2.0.

SteelHead feature	Feature after upgrading to SteelHead SD 2.0
Layer 7 optimization blades	All Layer 7 SteelHead optimization blades are supported. For example, HTTP, SSL, CIFS/SMB, MAPI, Oracle Forms, NFS, Lotus Notes, and storage replication (for example, SnapMirror) all operate normally and are unchanged. The Citrix optimization blade is supported but the ability to support the optimization of Multi-Stream ICA within the blade is no longer possible because the QoS functionality is taken care of by the service virtual machine (SVM) in SteelHead SD. You cannot optimize UDP traffic using the SteelHead IP blade as traffic is not redirected through the virtual SteelHead.
SteelHead SaaS and the new SaaS Accelerator	SteelHead SD 2.0 supports SteelHead SaaS and the SaaS Accelerator are both supported. The SaaS Accelerator is not availble for SteelConnect 2.11 gateways.
Web proxy	SteelHead SD supports SteelHead Web proxy.
CIFS prepopulation	SteelHead SD supports SteelHead CIFS prepopulation.
Active Directory integration	SteelHead SD supports SteelHead Active Directory integration. Because the virtual SteelHead instance has full control of the primary interface, it supports Active Directory integration and server-side out-of-path deployments.
Data store synchronization	SteelHead SD supports SteelHead data store synchronization on the primary interface with an adjacent SteelHead appliance.
Caching DNS service	SteelHead SD supports the SteelHead caching DNS service. With the caching DNS service, because the AUX port is not available to the virtual SteelHead, caching DNS is limited to the primary interface only.
Transport performance features	SteelHead SD supports SteelHead high speed TCP and bandwidth estimation, satellite features such as SCPS, and single-ended connections.
Management, reporting, and diagnostics	SteelHead SD supports SteelHead domain, host, and port labels, as well as in-path and peering rules.
Secure vault	SteelHead SD supports SteelHead secure vault. The secure vault password is retained when you upgrade from SteelHead to SteelHead SD.

SteelHead feature	Feature after upgrading to SteelHead SD 2.0
Management access controls	SteelHead SD supports SteelHead management access controls including Radius and TACACS, and role-based access.
TCP dump export	SteelHead SD supports SteelHead export of TCP dumps.

SteelHead features changed after upgrading to SteelHead SD 2.0

This table summarizes the features that do change after upgrading to SteelHead SD 2.0.

SteelHead feature	Feature after upgrading to SteelHead SD 2.0
WAN-optimization only mode	WAN-optimization only mode is not supported on SteelHead SD.
Hybrid networking services (path selection, secure transport, QoS)	Hybrid networking services (path selection, secure transport, QoS) are not supported on SteelHead SD. The network services of QoS, path selection and secure transport replaced by SteelConnect SD-WAN counterparts. Any QoS feature configuration on the original SteelHead must be converted to the new QoS in SCM. MX-TCP, because it was part of QoS, is not supported on SteelHead SD. Citrix Multistream ICA is not supported on SteelHead SD.
Multiple in-path interfaces for WAN optimization	SteelHead SD doesn't support multiple in-path interfaces for WAN optimization. Given that SteelHead SD is a Layer 3 gateway, multiple LAN ports and segments can be mapped to a single in-path interface. There is no longer a need for multiple in-path interfaces on an SteelHead SD appliance. After upgrading from SteelHead to SteelHead SD you must reconfigure your multiple in-path interfaces to a single in-path configuration.
Virtual in-path or WCCP/PBR	Virtual in-path or WCCP/PBR is not supported on SteelHead SD. The concept of virtual in-path is not relevant for the WAN optimization of SteelHead SD. Thus, there is no need for WCCP or PBR.
Simplified Routing and VLAN transparency	Simplified Routing and VLAN transparency is not supported on SteelHead SD. Because the in-path interface on the virtual SteelHead instance within SteelHead SD doesn't sit physically in-path on the network, there is no need for Simplified Routing or VLAN transparency.
IPSec, subnet side rules, MXTCP and link state propagation	IPSec, subnet side rules, MXTCP and link state propagation are not supported on SteelHead SD.
Serial high availability (HA)	After upgrading, serial HA is not supported on SteelHead SD 2.0. SteelHead appliances in an HA pair must be individually shut down and upgraded separately. Active-active (1:1) HA is supported on SteelHead SD 2.0.

SteelHead feature	Feature after upgrading to SteelHead SD 2.0
NIC bypass (fail-to-wire)	Currently, NIC level bypass or fail-to-wire is not supported in SteelHead SD. If at any point the status of the virtual SteelHead instance shows a failure condition, for example a reboot or a crash, the system stops sending traffic that was destined for the virtual SteelHead. Instead, it bypasses the SteelHead thereby ensuring the traffic is not black-holed. You can compare this behavior with a physical SteelHead entering bypass mode. The traditional SteelHead bypass functionality doesn't apply 1:1 to a SteelHead SD appliance because it is now an SD-WAN appliance that acts as a Layer 3 hop (or a custom edge router in some cases). Enabling NIC bypass mode without proper routing architecture support can lead to unintended traffic path behavior and can have security implications.
Fail-to-block	If a SteelHead SD appliance fails, the appliance goes into fail-to-block mode. If only the SteelHead WAN optimization service fails, then traffic is passed through unoptimized and the SteelConnect SD-WAN service remains fully operational. If only the SteelConnect SD-WAN service fails, then all traffic on the gateway is blocked.
Data store synchronization	Data store synchronization is supported only on the primary interface because the AUX interface is not available to the virtual SteelHead. (The AUX port is the dedicated port used in HA configurations; it can also be used as an additional WAN uplink.)
RADIUS/Authentication server under Sites	RADIUS/Authentication server under Sites configuration in SCM is not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030 appliances. Consult with your Riverbed sales engineer or Riverbed Professional Services at http://www.riverbed.com/services/index.html.
Redirection of UDP traffic through the virtual SteelHead	Redirection of UDP traffic through the virtual SteelHead is not supported in SteelHead SD 2.0. You cannot optimize UDP traffic using the SteelHead IP blade.

SteelHead feature	Feature after upgrading to SteelHead SD 2.0
Source NAT on underlay traffic	Source NAT on underlay traffic is not supported on SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030. SteelHead SD appliances do not perform source NATing on underlay traffic exiting via the Internet uplink if it is destined for a private address, regardless of the configured outbound NAT setting. This is a change from the previous behavior for SteelHead SD 1.0 appliances, if NAT was enabled for an uplink, NAT was performed for all traffic exiting via the Internet uplink. For details on configuring NAT, see the *SteelConnect Manager User Guide*.
SteelHead Management Console GUI pages	These SteelHead Management Console GUI elements are not supported in SteelHead SD 2.0: ▪ QoS reports. ▪ Flow export settings: Export QoS and application statistics to Cascade Flow Collectors. ▪ Subnet side rules. ▪ WCCP settings. ▪ Connection forwarding settings. ▪ Failover settings. ▪ In-Path Settings: Enabling Link State Propagation. ▪ IPSec settings. ▪ AUX interface setting in the Base Interfaces page. ▪ Caching DNS: Listen on AUX interface check box.

Hardware and software requirements

Riverbed component	Hardware and software requirements
SteelHead SD appliance	The SteelHead SD 570-SD and 770-SD appliances are desktop models. The SteelHead SD 3070-SD appliance requires a 19-inch (483 mm) four-post rack. For details, see the *Rack Installation Guide*.
SteelHead SD Management Console	The Management Console has been tested with all versions of Chrome, Mozilla Firefox Extended Support Release version 38, and Microsoft Internet Explorer 11. JavaScript and cookies must be enabled in your web browser.
SteelConnect and SteelConnect Manager (SCM)	SteelHead SD requires SteelConnect 2.11. SCM supports the latest version of the Chrome browser. SCM requires a minimum screen resolution of 1280 x 720 pixels. We recommend a maximum of 1600 pixels for optimal viewing.
SteelCentral Controller for SteelHead (SCC)	We recommend you have SCC 9.7.1 installed.

NIC support

Network interface card (NICs) are supported on the SteelHead SD 3070-SD appliances for nonbypass traffic. SteelHead SD 570-SD and 770-SD appliances do not support NICs.

Note: For SteelHead SD 3070-SD appliances, bypass NICs are not required for SteelConnect gateway deployments since LAN traffic requires network address translation (NAT) before it reaches the service provider network.

You can install these NICs in the SteelHead SD 3070-SD for nonbypass traffic.

NICs	Size (*)	Manufacturing part #	Orderable part #
Two-Port 10-GbE Fiber SFP+	HHHL	410-00036-02	NIC-1-010G-2SFPP
Four-Port 10-GbE Fiber SFP+	HHHL	410-00108-01	NIC-1-010G-4SFPP

*HHHL = Half Height, Half Length

For details on NICs, see the *Network and Storage Card Installation Guide*.

Firewall requirements

The SteelHead SD 570-SD, 770-SD, 3070-SD, and SDI-2030 support stateful application-based firewalls at the network edge. For details on SteelConnect firewall and security features, see the *SteelConnect SD-WAN Deployment Guide*.

All communication is sourced from the site out to the SteelConnect management service. There’s no need to set up elaborate firewall or forwarding rules to establish the dynamic full-mesh VPN or to gain connectivity to the cloud. After you register an appliance, it receives its assigned configuration automatically. For details on SteelConnect firewall requirements, see the *SteelConnect Manager User Guide*.

Make sure the firewall ports 80 and 443 are open so that software installation and SCM operations aren’t blocked. For details on SteelConnect default ports, see the *SteelConnect Manager User Guide*.

Ethernet network compatibility

The SteelHead SD appliance supports these Ethernet networking standards.

Ethernet standard	IEEE standard
Ethernet Logical Link Control (LLC)	IEEE 802.2 - 1998
Fast Ethernet 100BASE-TX	IEEE 802.3 - 2008
Gigabit Ethernet over Copper 1000BASE-T (All copper interfaces are autosensing for speed and duplex.)	IEEE 802.3 - 2008
Gigabit Ethernet over Fiber 1000BASE-SX (LC connector)	IEEE 802.3 - 2008
Gigabit Ethernet over Fiber 1000BASE-LX	IEEE 802.3 - 2008
Gigabit Ethernet over Fiber 10GBASE-LR Single Mode	IEEE 802.3 - 2008
Gigabit Ethernet over 10GBASE-SR Multimode	IEEE 802.3 - 2008

SNMP-based management compatibility

SteelConnect SD-WAN service supports proprietary MIBs accessible through SNMPv2 and SNMPv3. For detailed information about the SD-WAN service MIB, see the *SteelConnect Manager User Guide*.

The SteelHead WAN optimization supports proprietary MIBs accessible through SNMP, SNMPv1, SNMPv2c, and SNMPv3, although some MIB items might only be accessible through SNMPv2 and SNMPv3. For detailed information about the WAN optimization service MIB, see the *SteelHead User Guide*.

For detailed information on SteelConnect SNMP support, see the *SteelConnect Manager User Guide* and the *SteelHead SD User Guide*.

2

Configuring WAN Optimization

This topic describes how to enable WAN optimization for SteelHead SD 2.0. It includes these sections:

- "Overview" on page 21
- "Assigning the in-path IP address and default gateway in SCM" on page 22
- "Enabling WAN Optimization in SCM" on page 23
- "Enabling WAN optimization on the virtual SteelHead instance" on page 25
- "Troubleshooting" on page 27

For detailed information on installing and configuring SteelHead SD, see the *SteelHead SD Installation Guide*.

These instructions assume you have created an organization, site, and LAN zone for the SteelHead SD appliance. For details, see the *SteelConnect Manager User Guide* and the *SteelHead SD User Guide*.

Overview

When you enable WAN optimization for SteelHead SD, you perform the initial configuration within SCM. You must also configure settings on the virtual SteelHead instance itself, using the SteelHead Management Console, the CLI, or the SteelCentral Controller for SteelHead (SCC).

When enabling WAN optimization, keep these guidelines in mind:

- To enable WAN optimization, the location where the SteelHead SD is installed must have at least one LAN zone. The in-path IP address in the virtual SteelHead instance must match the address in SCM.
- The LAN port must be configured as a single-zone uplink for the SteelHead WAN optimization service. If you do not enable the LAN port, SteelConnect won't see either the SteelHead WAN optimization service or the clients on the LAN side of the network.
- The default gateway in the virtual SteelHead instance must be the IP address for the LAN zone in SCM.
- If the LAN port of the SteelHead SD appliance is on a VLAN trunk, make sure to configure the VLAN ID on the virtual SteelHead in-path interface.
- Depending on the in-path rules you have defined, the virtual SteelHead instance optimizes any traffic received from SteelConnect via the LAN interface.
- When WAN optimization is enabled in SCM, there is a momentary interruption to service as the SteelConnect is reconfigured with its SteelHead LAN and WAN interfaces.

When WAN optimization is enabled, a virtual SteelHead instance is automatically provisioned by the system. The primary port on the SteelHead SD appliance is connected directly to the primary interface of the virtual SteelHead instance.

The virtual SteelHead instance is configured with a single in-path interface pair (LAN/WAN). Unlike physical SteelHead appliances or their virtual equivalents that exist outside of an SD-WAN service-chained deployment, the number of in-path interfaces is irrelevant. For consistency and compatibility, the in-path interface pair is configured as LAN0_0 and WAN0_0.

Assigning the in-path IP address and default gateway in SCM

The first step is to assign an in-path IP address within the LAN zone for the site. You choose an IP address for the LAN zone in which the SteelHead SD is installed. You will use this IP address to configure the in-path interface and default gateway on the virtual SteelHead instance.

These instructions assume that you have configured the primary and LAN ports on the SteelHead SD appliance in SCM:

- The LAN port is configured as a single-zone uplink for the SteelHead WAN optimization service. By default, the LAN port is disabled on SteelHead SD appliances unless it is explicitly enabled. If you don't enable the LAN port, SteelConnect won’t see either the SteelHead WAN optimization service or the clients on the LAN side of the network.
- The primary port is configured as SteelHead Primary mode for the SteelHead SD appliance.

Note: If the LAN port attached to the SteelHead SD appliance is in a VLAN trunk, the virtual SteelHead instance must be given an IP address from one of the zones that is part of the trunk, and the virtual SteelHead in-path IP address must also be configured with the corresponding VLAN ID.

For details on configuring the primary and LAN ports, see the *SteelHead SD Installation Guide*.

To assign the in-path IP address and the default gateway in SCM

1. In SCM, choose Network Design > Zones.

2. Select the zone with the SteelHead SD appliance to expand the pane. The IP tab is displayed.

3. Under IPv4 Network, specify the LAN zone subnet. Write down this IP address. You will use this address when you configure the inpath0_0 interface for WAN optimization on the virtual SteelHead instance.

Figure 2-1. Assigning the IP address for the in-path IP address and default gateway

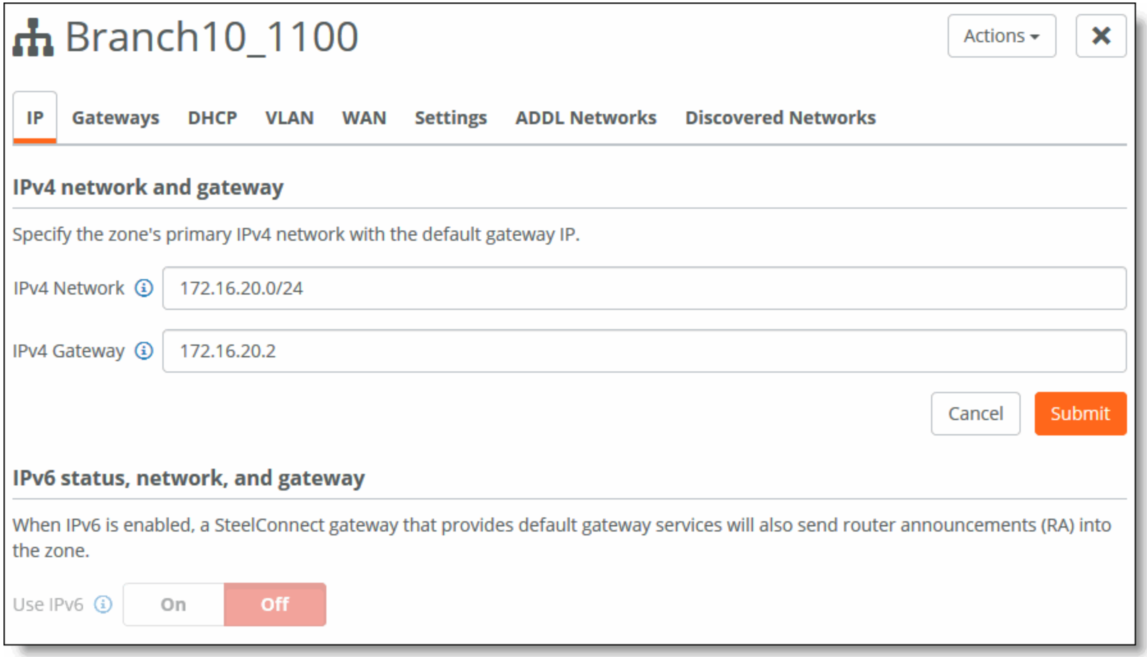

IPv4 network and gateway

Specify the zone's primary IPv4 network with the default gateway IP.

IPv4 Network: 172.16.20.0/24

IPv4 Gateway: 172.16.20.2

IPv6 status, network, and gateway

When IPv6 is enabled, a SteelConnect gateway that provides default gateway services will also send router announcements (RA) into the zone.

Use IPv6: On [✓ Off]

For example, if the network IP address is 172.16.20.0/24, you can assign any IP address from 172.16.20.1 to 172.16.20.254 for the SteelHead in-path interface.

4. Under IPv4 Gateway, specify the default gateway. Write down this IP address. You will use this address when you configure the default gateway for WAN optimization on the virtual SteelHead instance.

Enabling WAN Optimization in SCM

You enable WAN optimization in SCM in the Appliances page under the Services tab. You also specify the virtual SteelHead instance in-path IP address. The in-path IP address must be within the LAN zone subnet that you have defined.

The WAN optimization service is disabled by default. When disabled, the WAN optimization service will not participate in any WAN optimization functionality. If disabled, any configuration related to WAN optimization service on this appliance will not be applied.

Important: Only zones that are attached to a physical port can be used to configure the SteelHead SD IP address. Choose Appliances > Port to attach a zone to a port.

To enable WAN optimization

1. Choose Appliances > Overview.

2. Select the SteelHead SD appliance to expand the page.

3. Select the Services tab.

Figure 2-2. Enabling WAN optimization in SCM

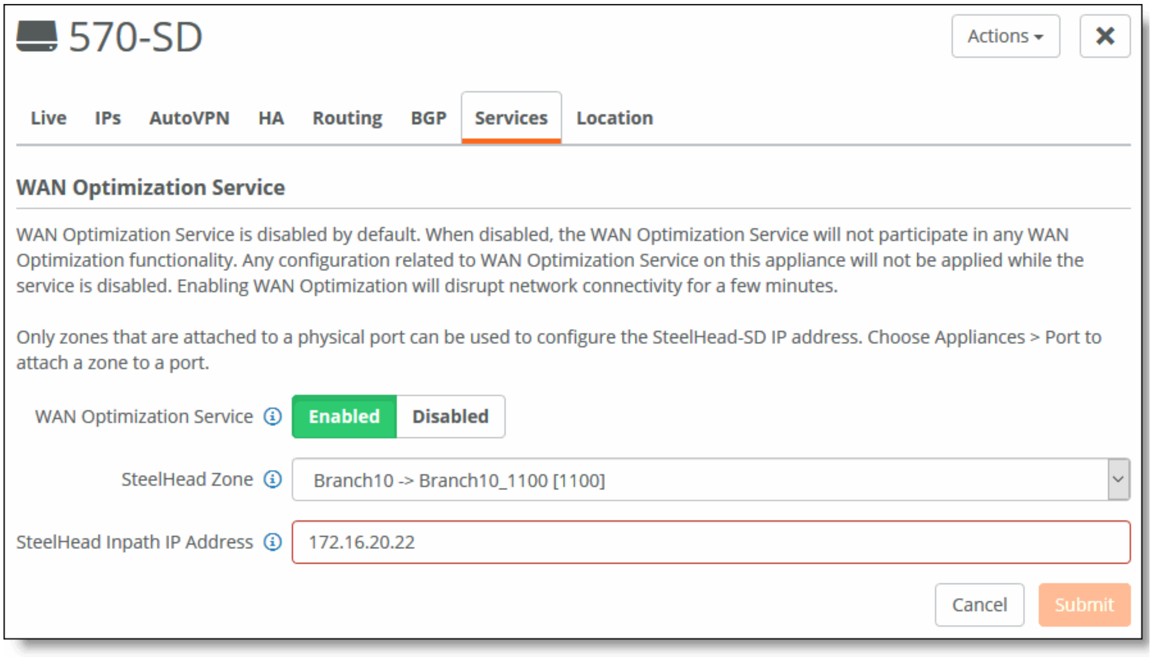

WAN Optimization Service

WAN Optimization Service is disabled by default. When disabled, the WAN Optimization Service will not participate in any WAN Optimization functionality. Any configuration related to WAN Optimization Service on this appliance will not be applied while the service is disabled. Enabling WAN Optimization will disrupt network connectivity for a few minutes.

Only zones that are attached to a physical port can be used to configure the SteelHead-SD IP address. Choose Appliances > Port to attach a zone to a port.

WAN Optimization Service: [✓ Enabled] Disabled

SteelHead Zone: Branch10 -> Branch10_1100 [1100]

SteelHead Inpath IP Address: 172.16.20.22

4. Under WAN Optimization Service, fill out these required session attributes:

- **WAN Optimization Service** - Click **Enabled** to enable the WAN optimization service for the selected SteelHead SD appliance. When disabled, the WAN optimization service will not participate in any WAN optimization functionality. If disabled, any configuration related to WAN optimization service on this appliance will not be applied.
- **SteelHead Zone** - Select the zone to which this SteelHead SD appliance belongs. Only zones that are attached to a physical port can be used to configure the SteelHead SD IP address. Choose Appliances > Port to attach a zone to a port.
- **SteelHead Inpath IP Address** - Specify the SteelHead in-path IP address. The IP address must be within the LAN zone subnet. This value tells SCM what in-path IP address you are using for the virtual SteelHead instance.

5. Click **Submit.**

After the WAN optimization service has been enabled within SCM, the SteelHead SD triggers the orchestration and provisioning of the virtual SteelHead instance. This action causes a momentary interruption to operations within SteelConnect because it is reconfigured with the SteelHead LAN and WAN interfaces.

As the virtual SteelHead instance boots within SteelHead SD, its primary interface tries to obtain an IP address via DHCP. It is important to ensure the SteelHead SD primary port is attached to a network where a DHCP service is available.

Identifying the primary IP address of the SteelHead

You use the primary IP address to connect to the virtual SteelHead instance. You can identify the primary IP address of the SteelHead in one of the following ways:

- **When SteelConnect acts as the DHCP server -** You can set the SteelConnect virtual gateway to act as a DHCP server and identify the primary IP address for the SteelHead in SCM. To view the SteelHead primary IP address in SCM, choose Appliances > Overview and select the SteelHead SD appliance. The primary IP address is listed under the IPs tab. For details on configuring SteelConnect to act as a DHCP server, see the *SteelHead SD Installation Guide*.
- **When the SCC is used to manage SteelHeads -** If you are using the SCC to manage the WAN optimization service, you can obtain the primary IP address for each appliance in your network. SCC automatically registers all SteelHeads it detects in your network and provides the primary IP address for each in the Appliances page. For details on connecting to SCC, see the *SteelCentral Controller for SteelHead User Guide*.
- **When an external server acts as the DHCP server -** You can obtain the MAC address from the appliance and search for the primary IP address on the DHCP server console. You can find the MAC address on the appliance label or you can view it in SCM. To view the MAC address in SCM, choose Ports and select the primary port for the appliance. The MAC address is listed under the Info-Mode tab.

After you have discovered the primary IP address that has been leased to the virtual SteelHead instance, you simply log in to the management console user interface and complete the configuration of the virtual SteelHead instance.

Enabling WAN optimization on the virtual SteelHead instance

To enable WAN optimization for SteelHead SD, you must configure the inpath0_0 interface and default gateway for each appliance in your network using the SCC or the SteelHead Management Console.

Configuring the in-path interface and default gateway

These instructions describe how to configure the in-path interface and default gateway using the SteelHead Management Console.

Tip: In the SCC, choose Manage: Appliances > Appliance Pages > In-Path Interfaces to modify the inpath0_0 interface and default gateway. You can push the policy to the selected appliance.

To configure the in-path interface and the default gateway in the SteelHead

1. Using the Primary IP address you obtained from SCM, SCC, or the DHCP server, enter it in the location box of your web browser using HTTPS. The login page for the SteelHead Management Console is displayed.

2. Specify the default user login (**admin**) and password (**password**).

3. Click **Sign In** to display the Dashboard.

4. Choose Networks > Networking: In-Path Interfaces.

Figure 2-3. In-Path Interfaces page

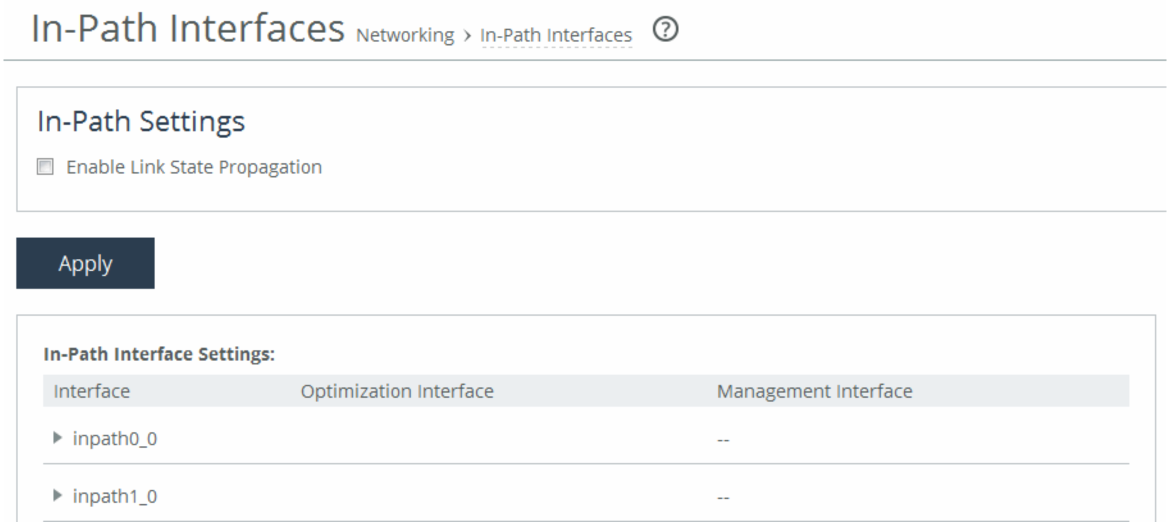

In-Path Interfaces Networking > In-Path Interfaces

In-Path Settings

□ Enable Link State Propagation

In-Path Interface Settings:

Interface	Optimization Interface	Management Interface
inpath0_0		--
inpath1_0		--

5. Select the interface to expand the page.

Figure 2-4. Configuring the in-path interface

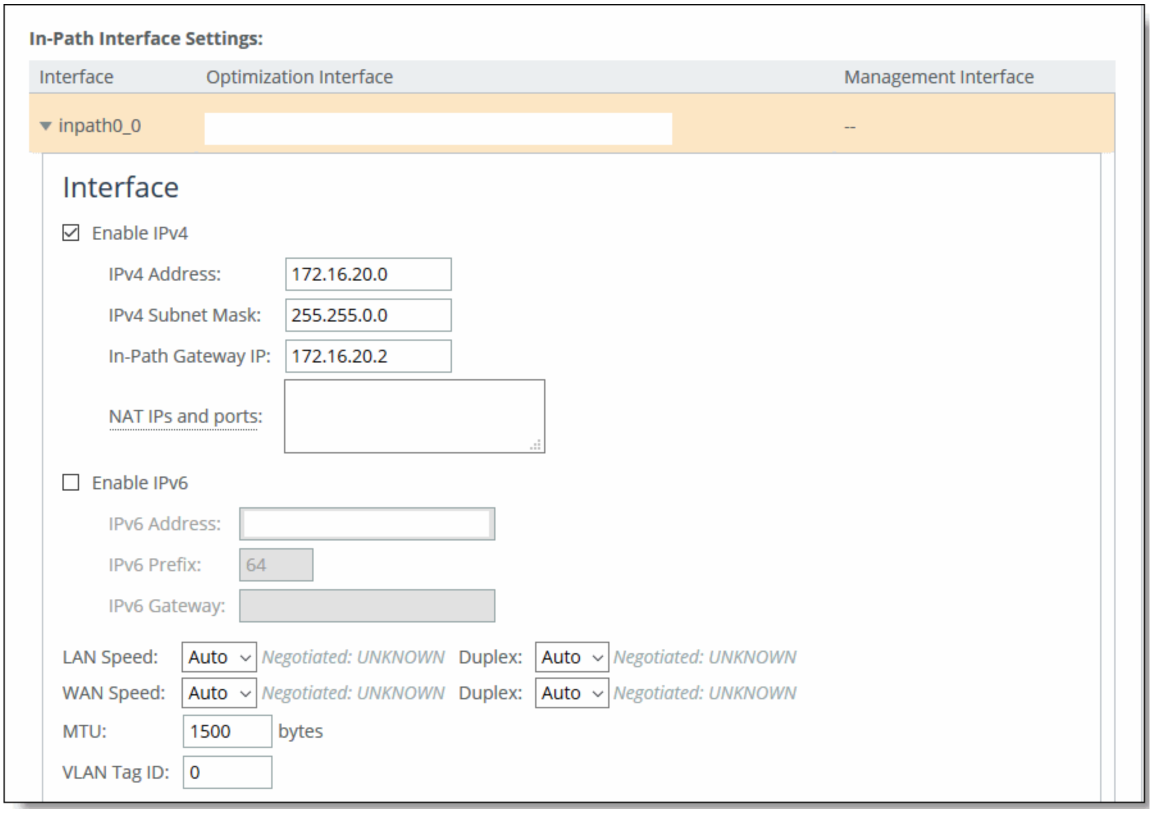

In-Path Interface Settings:

Interface	Optimization Interface	Management Interface
inpath0_0		--

Interface

☑ Enable IPv4

IPv4 Address: 172.16.20.0

IPv4 Subnet Mask: 255.255.0.0

In-Path Gateway IP: 172.16.20.2

NAT IPs and ports:

□ Enable IPv6

IPv6 Address:

IPv6 Prefix: 64

IPv6 Gateway:

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

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

MTU: 1500 bytes

VLAN Tag ID: 0

6. Type the IP address that you assigned in SCM. For details, see "To assign the in-path IP address and the default gateway in SCM" on page 22.

7. Type the subnet mask address. The subnet mask on the in-path must match the subnet mask on the zone (typically /24, but it can be whatever you specified in the zone settings).

8. Type the IP address that you assigned in SCM for the default gateway. For details, see "To assign the in-path IP address and the default gateway in SCM" on page 22.

9. If the LAN port is part of a VLAN trunk, enter the correct VLAN ID for the in-path.

10. Click **Apply**.

11. You can refine your in-path WAN optimization settings using the SteelHead Management Console. For details, see the *SteelHead User's Guide*.

Troubleshooting

The virtual SteelHead instance is accessible for management and diagnostics via primary and in-path interface.

You cannot ping the in-path interface for the virtual SteelHead instance.

You can ping the primary interface for virtual SteelHead instance.

TCP dumps can be taken to verify and examine traffic flows on following interfaces:

– In-path interface of virtual SteelHead instance

– Knet interfaces of the service virtual machine.

To gather and verify information, check these SteelHead reports:

- Current Connections
- In-path Rule Counters
- Throughput

3

Configuring Zscaler on SteelHead SD

This topic describes how to integrate Zscaler on SteelHead SD 2.0. It includes these sections:

- "Zscaler overview" on page 29
- "Key features" on page 29
- "SteelHead SD Restrictions" on page 29
- "Basic steps" on page 30

Zscaler overview

Zscaler is a cloud-based security provider that distributes components of a standard proxy to create a giant global network that acts as a single virtual proxy.

Zscaler cloud security solution is supported on SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway. The configuration procedures are the same for SteelConnect gateway appliances. For details on Zscaler, see the *SteelConnect Manager User Guide*.

Key features

SteelHead SD supports these features:

- Availability and management of ZEN lists.
- Automatic selection of ZEN nodes based on latency, and the ability to select Zen nodes manually.
- The ability to create two VPN tunnels per site to Zscaler that act as active and passive.
- Failover based on IPsec dead per detection (DPD) method, which takes about 120 seconds.
- Zscaler support for high availability deployments.

SteelHead SD Restrictions

Currently, Zscaler for SteelHead SDs has these restrictions:

- Only tunnel status is reported.
- You can't enable traffic path rules for the Zscaler WAN.

Basic steps

Perform these basic steps to configure Zscaler. For details, see the *SteelConnect Manager User Guide*.

1. On SCM, enable Zscaler by selecting the Zscaler Cloud.

2. On SCM, select the ZEN lists either automatically or manually.

3. On Zscaler, download the VPN credentials and locations on the SCM Zscaler page.

4. Import the credentials in the Zscaler portal.

5. On the site, set the Zscaler WAN as the Internal breakout preference at the organization, site, or zone level.

4

Configuring Local Subnet Discovery on SteelHead SD

This topic describes how to configure SteelHead SD 2.0 to discover global and local subnets on the LAN side of the network. It includes these topics:

- "Introducing local subnet discovery" on page 31
- "Routing criteria" on page 31
- "Defining global subnet discovery at the organization level" on page 32
- "Defining local subnet discovery" on page 33

These procedures describe local and global subnet autodiscovery for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch. For additional information, see the *SteelConnect Manager User Guide*.

Introducing local subnet discovery

SteelHead SD 2.0 provides the ability to discover subnets at the zone and site level in a branch. Local subnet discovery identifies routes that are local to a particular branch. These routes can be reached from other sites or branches using the overlay tunnels.

Ideally, all routes learned over the LAN interfaces of an appliance, on a particular branch, should be qualified as routes local to that branch. However, this qualification isn't always straight forward. Consider the case where OSPF is configured with both the LAN zones and the WAN uplinks attached to it. In this case, OSPF will not be able to differentiate the routes that it learns over the LAN zones from the ones that it learns over the WAN uplinks. A similar problem can arise when BGP is the chosen protocol where an IBGP neighbor is established with the LAN router and eBGP neighbor is established with the provider-edge router over the WAN uplink. Here BGP will not be able to call out the local subnets implicitly. Another case to consider is when the appliance is placed behind a branch router, it loses the notion of LAN zones and WAN up-links. The local subnet autodiscovery feature provides a means for identifying subnets that are local to a branch.

Routing criteria

Local subnet discovery allows you to define a set of routing criteria so that routes that match the criteria are qualified as subnets local to the branch. The routing criteria are:

- **Zone inclusion list** - You select one or more of the configured LAN zones. Routes whose next-hop interface matches one of the selected zones are qualified as local subnets. Preexisting zones that are directly connected to a site are added to the list automatically. You can also manually add other zones. Zones deleted from a site are automatically removed from the list.
- **Uplink inclusion list** - You select one or more of the configured WAN uplinks. Routes whose next-hop interface matches one of the selected WAN uplinks are qualified as local subnets.
- **Prefix inclusion list** - You configure a list of prefixes. If a route for one of the prefixes in the list is received, it is qualified as a local subnet.
- **Next-hop inclusion list** - You configure a list of next-hop prefixes. All routes whose next-hop matches one of the entries in the list are qualified as a local subnet.
- **Prefix exclusion list** - You configure a list of prefixes. If a route for one of the prefixes in the list is received, it's not qualified as a local subnet.

Note: SteelConnect SDI-2030 and SDI-5030 gateways do not allow you to define routing criteria based on zones and uplinks.

For SteelHead SD, you can create inclusion and exclusion lists at the organization, zone, and site level. For example, you could create an umbrella subnet 10.0.0.0/8 inclusion list at the organization level and then drill down to a particular site to exclude 10.0.0.0/16.

Inclusion lists are applied first, and then exclusions lists will be applied.

Defining global subnet discovery at the organization level

Users can add an organization level subnet discovery rule under the Global Subnet Discovery tab. This rule will be applied to all sites, unless they are specifically overridden by the site-level subnet discovery rule.

To define global subnet discovery for an organization

1. Choose Organizations.

2. Select the Global Subnet Discovery tab.

Figure 4-1. Defining organization level subnet discovery

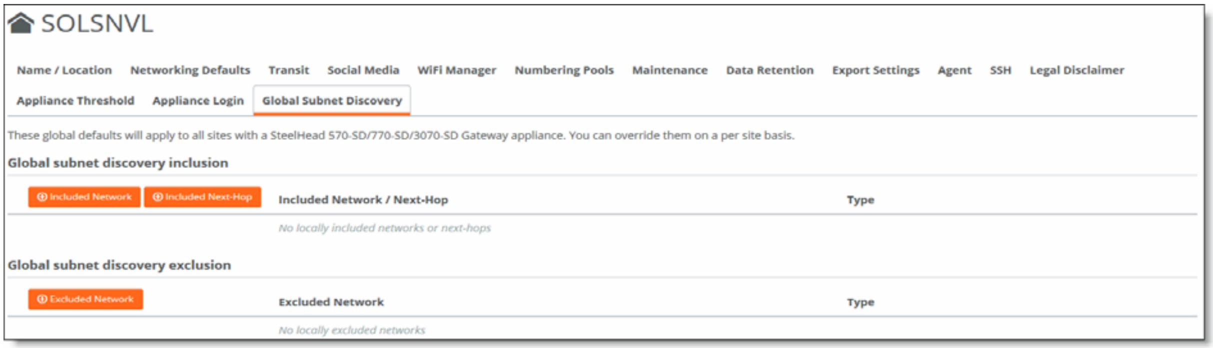

These global defaults will apply to all sites with a SteelHead 570-SD/770-SD/3070-SD Gateway appliance. You can override them on a per site basis.

Global subnet discovery inclusion

Included Network / Next-Hop	Type
No locally included networks or next-hops	

Global subnet discovery exclusion

Excluded Network	Type
No locally excluded networks	

3. Click **New Included Network.**

Figure 4-2. Defining included networks

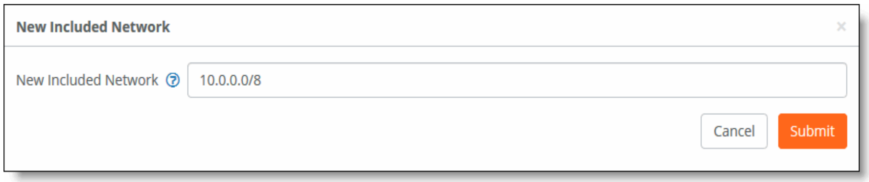

New Included Network: 10.0.0.0/8

4. Specify a Classless Inter-Domain Routing (CIDR) IPv4 address, including the network prefix to be included in local subnet autodiscovery.

5. Click **Submit.**

6. Click **Included Next Hop.**

Figure 4-3. Defining the next hop

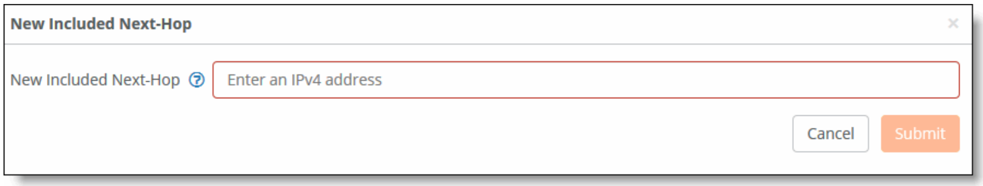

7. Specify the IPv4 IP address for the local-subnet next hop from SteelConnect appliance in this site.

8. Click **Submit.**

9. Click **Excluded Networks.**

Figure 4-4. Defining excluded networks

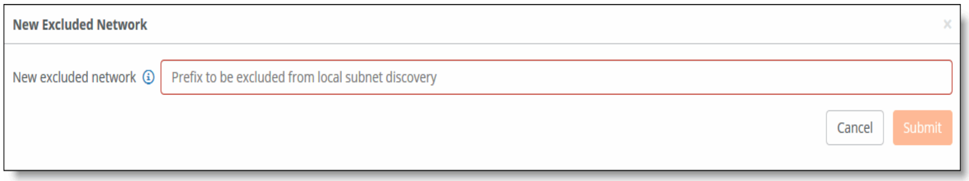

10. Specify the IP address, including the network prefix, to be excluded from local subnet autodiscovery.

11. Click **Submit.**

Defining local subnet discovery

After you have defined subnet discovery at the organization level, you can drill down to particular sites to define inclusion and exclusion lists. For OSPF routes, make sure that your branch has the LAN zone and WAN uplink attached to it before you begin. For BGP routes, make sure that your branch has the iBGP neighbor defined for the LAN router and the eBGP neighbor defined for the WAN router.

To define local subnet discovery

1. Choose Network Design to display the sites for the organization.

2. Select the site for which you want to define local subnet discovery.

3. Select the Local Subnet Discovery tab.

4. Under Inherit Global, click **On** to globally include subnets and next hops. Whatever subnets were configured for inclusion or exclusion at the organization level can be inherited at the site level.

Figure 4-5. Defining site level subnet discovery

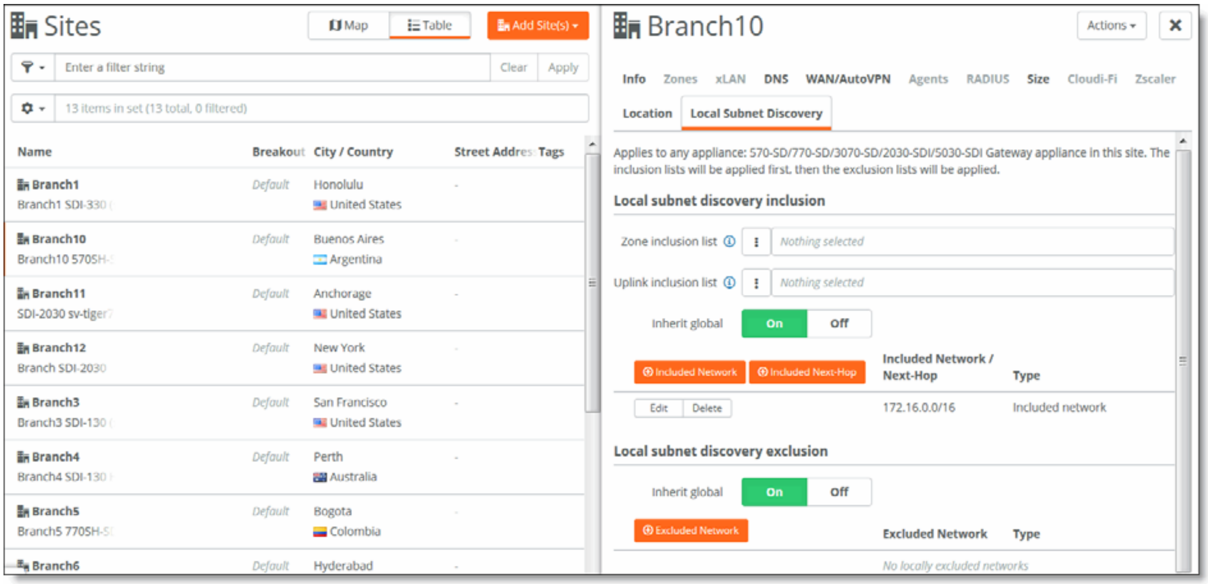

Enter a filter string

13 items in set (13 total, 0 filtered)

Name	Breakout	City / Country	Street Addres	Tags
Branch1 Branch1 SDI-330 (	Default	Honolulu United States	-	
Branch10 Branch10 570SH-	Default	Buenos Aires Argentina	-	
Branch11 SDI-2030 sv-tiger	Default	Anchorage United States	-	
Branch12 Branch SDI-2030	Default	New York United States	-	
Branch3 Branch3 SDI-130 (	Default	San Francisco United States	-	
Branch4 Branch4 SDI-130	Default	Perth Australia	-	
Branch5 Branch5 770SH-S	Default	Bogota Colombia	-	
Branch6	Default	Hyderabad	-	

Applies to any appliance: 570-SD/770-SD/3070-SD/2030-SDI/5030-SDI Gateway appliance in this site. The inclusion lists will be applied first. then the exclusion lists will be applied.

Local subnet discovery inclusion

Zone inclusion list: Nothing selected

Uplink inclusion list: Nothing selected

Inherit global: On Off

Included Network / Next-Hop	Type
172.16.0.0/16	Included network

Local subnet discovery exclusion

Inherit global: On Off

Excluded Network	Type
No locally excluded networks	

5. Select the zone to discover all of the LAN-side subnets routed through the zone's peers. The list includes automatically populated and manually added zones. Preexisting zones that are directly connected to a site are added to the list automatically. You can manually add other zones. Zones deleted from a site are automatically removed from the list.

Tip: Click the trash can to remove a zone from the inclusion list and add its prefix to the exclusion list.

6. Select the uplink from the Uplink Inclusion list.

7. Click **Included Network**.

Figure 4-6. Defining included networks

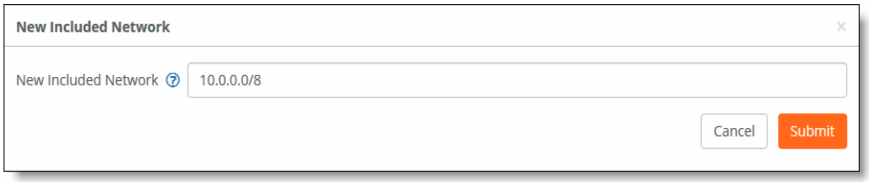

New Included Network

New Included Network: 10.0.0.0/8

8. Specify a Classless Inter-Domain Routing (CIDR) IP address, including the network prefix, and click **Submit**.

9. Click **Included Next Hop**.

Figure 4-7. Defining the next hop

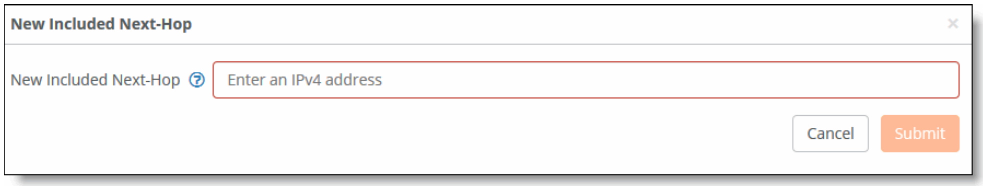

New Included Next-Hop

New Included Next-Hop: Enter an IPv4 address

10. Enter the IPv4 IP address for the next hop, and click **Submit**.

11. Click **On** to globally include subnets and next hops. Whatever subnets were configured for inclusion or exclusion at the organization level can be inherited at the site level.

12. Click **Excluded Networks**.

Figure 4-8. Defining excluded networks

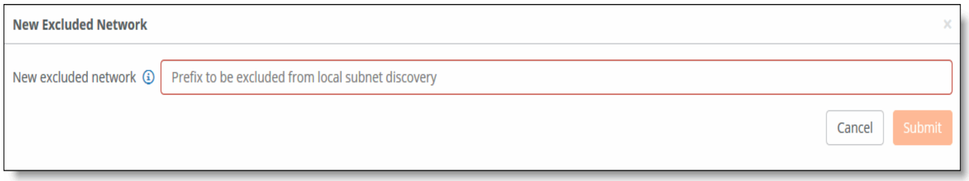

New Excluded Network

New excluded network: Prefix to be excluded from local subnet discovery

13. Specify a CIDR IP address, including the network prefix.

14. Click **Submit**.

To exclude subnets from local subnet discovery

1. Choose Network Design to display the sites for the organization.

2. Select the site for which you want to define local subnet discovery.

3. Select the Local Subnet Discovery tab.

4. Under Local subnet discovery exclusion, click **On** to globally exclude subnets and next hops. Whatever subnets were configured for inclusion or exclusion at the organization level can be inherited at the site level.

Figure 4-9. Defining site level subnet discovery

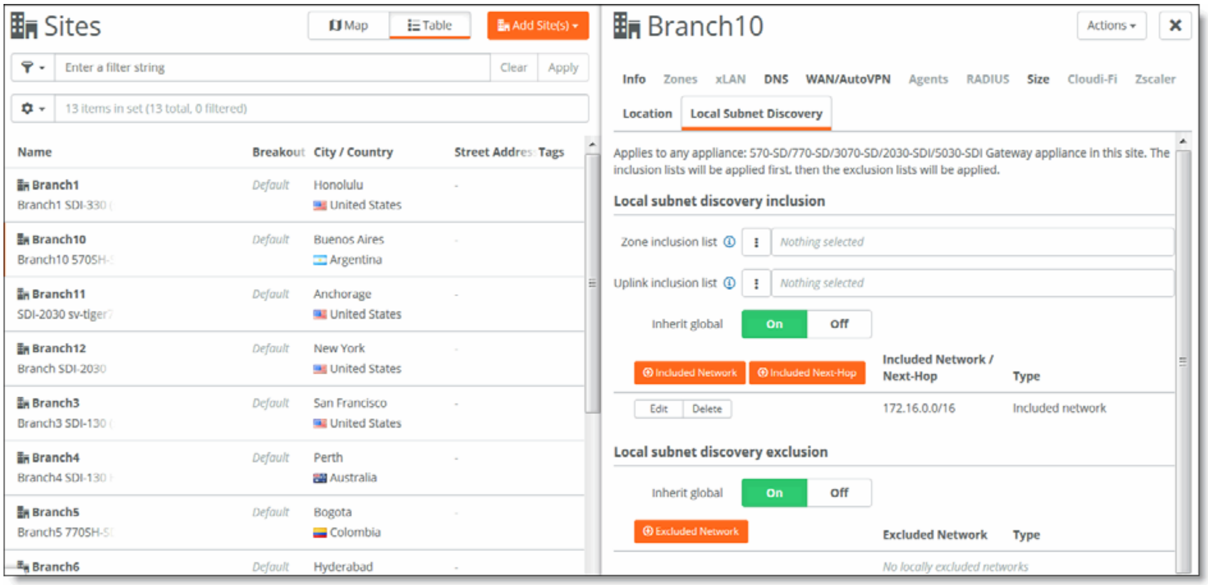

Name	Breakout	City / Country	Street Addres	Tags
Branch1 Branch1 SDI-330	Default	Honolulu United States	-	
Branch10 Branch10 570SH-	Default	Buenos Aires Argentina	-	
Branch11 SDI-2030 sv-tiger	Default	Anchorage United States	-	
Branch12 Branch SDI-2030	Default	New York United States	-	
Branch3 Branch3 SDI-130	Default	San Francisco United States	-	
Branch4 Branch4 SDI-130	Default	Perth Australia	-	
Branch5 Branch5 770SH-S	Default	Bogota Colombia	-	
Branch6	Default	Hyderabad	-	

Applies to any appliance: 570-SD/770-SD/3070-SD/2030-SDI/5030-SDI Gateway appliance in this site. The inclusion lists will be applied first, then the exclusion lists will be applied.

Local subnet discovery inclusion

Zone inclusion list: Nothing selected

Uplink inclusion list: Nothing selected

Inherit global: On | Off

Included Network | Included Next-Hop

	Included Network / Next-Hop	Type
Edit Delete	172.16.0.0/16	Included network

Local subnet discovery exclusion

Inherit global: On | Off

Excluded Network

Excluded Network	Type
No locally excluded networks	

5. Click **Excluded Network.**

Figure 4-10. Defining excluded networks

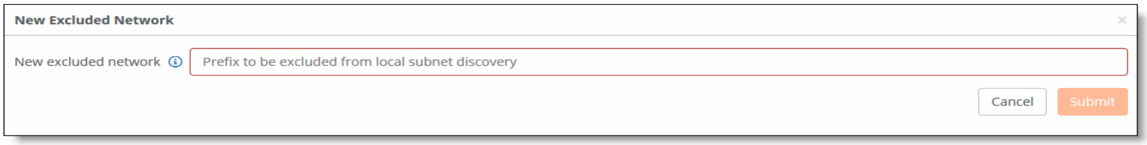

New Excluded Network

New excluded network: Prefix to be excluded from local subnet discovery

6. Specify the CIDR IP address, including the network prefix.

7. Click **Submit.**

5

Configuring BGP, OSPF, Static Routing, and Route retraction on SteelHead SD

This topic describes how to configuring SteelHead SD 2.0 Border Gateway Protocol (BGP), open shortest path first (OSPF) with an area border router (ABR), static routing, and route retraction. It includes these sections:

- "Configuring BGP on SteelHead SD" on page 37
- "Configuring OSPF with ABR on SteelHead SD" on page 43
- "Defining static routing on SteelHead SD" on page 53
- "Route retraction for SteelHead SD" on page 54

Important: Before you begin configuring BGP and OSPF for SteelHead SD, we recommend you read the *SteelConnect Manager User Guide*. The procedures here provide the basic steps for configuring SteelHead SD appliances.

Configuring BGP on SteelHead SD

This section describes how to configuring BGP on SteelHead SD. It includes these sections:

- "Introducing BGP on SteelHead SD" on page 37
- "Configuring BGP on SteelHead SD" on page 38
- "Configuring inbound and outbound prefixes, AS paths, and route maps for BGP neighbors" on page 40
- "Configuring BGP route redistribution" on page 41
- "Configuring BGP route summarization" on page 42
- "Viewing BGP status" on page 43

Introducing BGP on SteelHead SD

SteelHead SD provides full BGP support for local autonomous system (AS) numbers and neighbor configurations (including router ID, password, keepalive time, and hold time) for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

SteelHead SD 2.0 provides support for both exterior Border Gateway Protocol (eBGP) and interior Border Gateway Protocol (iBGP). SteelHead SD doesn't restrict BGP to the LAN or the WAN; it can communicate with its associated neighbors regardless of whether it is on the LAN or WAN.

You can configure BGP regardless of whether it is a zone or an uplink.

Important: Before you begin configuring BGP for SteelHead SD, we recommend you consult the *SteelConnect Manager User Guide*.

Configuring BGP on SteelHead SD

This section describes how to enable BGP and configure BGP neighbors on branch SteelHead SD appliances. By default, BGP is disabled.

Note: For SteelHead SD appliances, you can only add BGP neighbors under the Appliances > BGP tab. You can't add BGP neighbors from the Routing > BGP page.

To enable BGP and configure a BGP neighbor

1. Choose Appliances and select the appliance to expand the pane.
2. Select the BGP tab.

Figure 5-1. Enabling BGP

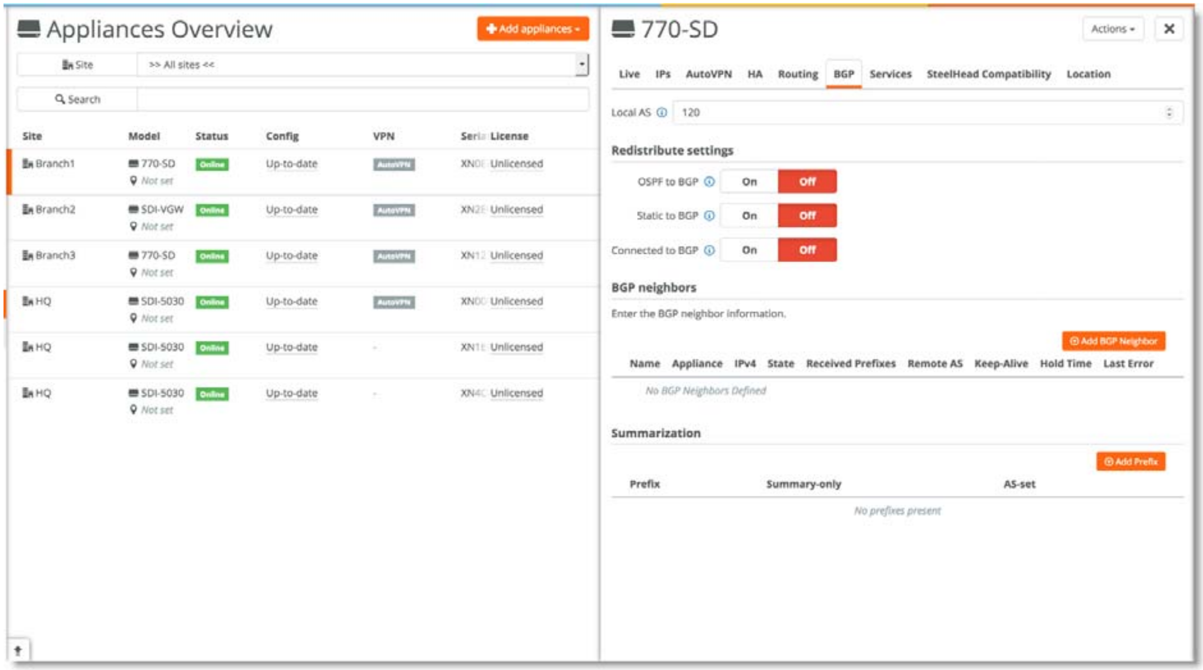

Site: >> All sites <<

Search

Site	Model	Status	Config	VPN	Seri	License
Branch1	770-SD Not set	Online	Up-to-date	AutoVPN	XN0	Unlicensed
Branch2	SDI-VGW Not set	Online	Up-to-date	AutoVPN	XN2	Unlicensed
Branch3	770-SD Not set	Online	Up-to-date	AutoVPN	XN1	Unlicensed
HQ	SDI-5030 Not set	Online	Up-to-date	AutoVPN	XN0	Unlicensed
HQ	SDI-5030 Not set	Online	Up-to-date	-	XN1	Unlicensed
HQ	SDI-5030 Not set	Online	Up-to-date	-	XN4	Unlicensed

Local AS 120

Redistribute settings

OSPF to BGP: On [✓ Off]

Static to BGP: On [✓ Off]

Connected to BGP: On [✓ Off]

BGP neighbors

Enter the BGP neighbor information.

Name	Appliance	IPv4	State	Received Prefixes	Remote AS	Keep-Alive	Hold Time	Last Error
No BGP Neighbors Defined								

Summarization

Prefix	Summary-only	AS-set
No prefixes present		

3. Specify an AS number in Local AS to start a BGP session. The range is from 1 to 4294967295.
4. Under BGP neighbors, click **Add BGP Neighbor**.

5. Specify a name for the BGP neighbor.

6. Specify the IP address of the BGP neighbor.

7. Specify the remote AS number that the BGP peer belongs to: for example, 200. The range is from 1 to 4294967295.

8. The remainder of the BGP attributes are provided by default. They can be changed based on your administrator settings. Optionally, fill out these BGP neighbor attributes:

- **Password** - Optionally, type a password to enable MD5 authentication. You must use the same password on both BGP neighbors. If you do not require MD5 authentication, you can leave this field blank.

  Click the eye icon to see the password as you type. The view persists until you click the eye icon again to hide the password.

- **Keep Alive Time** - Optionally, specify the amount of time, in seconds, that the eBGP neighbors exchange keepalive messages to determine whether a link has failed or is no longer available. The neighbors exchange keepalive messages often enough so that the hold time doesn't expire. The default setting is 60.

- **Hold Time** - Optionally, specify the amount of time, in seconds, that a gateway neighbor waits for an incoming keepalive, update, or notification message from a neighbor before it assumes its neighbor is down. If the gateway doesn't receive a keepalive, update, or notification message from its neighbor within the period specified, it closes the connection and routing through that neighbor becomes unavailable.

  A 0 value means that no keepalive messages are sent and the connection will never close. The hold-time range is from 0 to 65535. The default setting is 180.

  The hold-time value is three times the interval at which keepalive messages are sent. Using the default values for the keepalive time of 60 and the hold time of 180, the settings work together like this: after two neighbors establish an eBGP session, 60 seconds later they'll each send a keepalive message. When a gateway receives a keepalive message from its neighbor, that gateway's hold time for the session will have counted down from 180 to 120, but it's then reset to 180. This process continues every 60 seconds. However, should neighbor A lose power, then neighbor B won't receive any keepalives. So after 180 seconds, neighbor B determines that neighbor A is down and closes the session.

- **Weight** - When multiple routes to the same destination exist, BGP will prefer the route with the highest weight. The default value is 0 and the range is from 0 to 65535.

- Click **On** to distribute the default route (0.0.0.0/0) to the specified BGP neighbor. The default setting is **Off**.

  – **Route map** - Click the search selector and select the route use case. The routing policies defined by the selected route map are applied while accepting routes to the BGP neighbor.

9. Click **Submit**.

10. Repeat Step 4 through Step 9 if you have two MPLS providers that need to do BGP peering with the current appliance. You need to create a BGP configuration for each one.

11. Repeat this process for other SteelHead SDs behind other routers.

Note: BGP redistribution and summarization can only be configured after you have defined route maps and prefixes.

Configuring inbound and outbound prefixes, AS paths, and route maps for BGP neighbors

Optionally, you can configure inbound and outbound prefixes, AS paths, and route policies for BGP neighbors. By specifying these options, you can define what inbound and outbound routes are allowed or denied for BGP neighbors.

We recommend you define route policies, AS lists, and prefix lists before you configure the inbound and outbound settings. For details on configuring routing policies, see "What are routing policies?" on page 58.

Note: Outbound fields are disabled if a cluster site is selected as a transit hub for SteelHead SD 570-SD, 770-SD, and 3070-SD appliances. For details on configuring transit hubs, see the *SteelConnect Manager User Guide*.

To configure inbound and outbound BGP route settings

1. Choose Appliances and select the appliance to expand the pane.

2. Select the BGP tab.

Figure 5-2. Configuring inbound and outbound settings

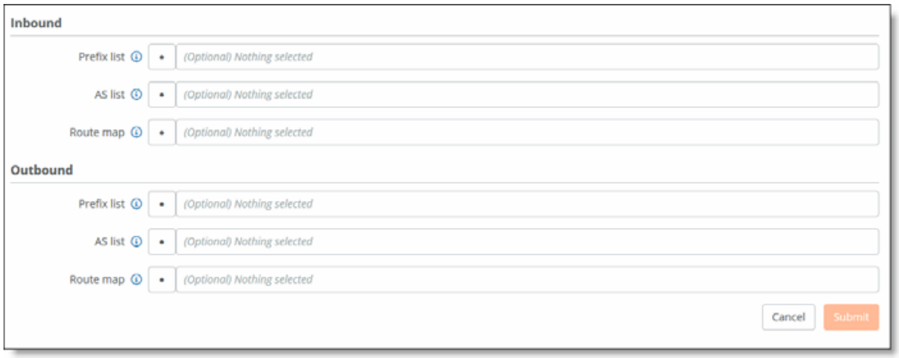

Inbound

Prefix list (Optional) Nothing selected

AS list (Optional) Nothing selected

Route map (Optional) Nothing selected

Outbound

Prefix list (Optional) Nothing selected

AS list (Optional) Nothing selected

Route map (Optional) Nothing selected

3. Under Inbound, fill out these attributes:

- **Prefix list** - Specify the prefixes to be allowed or denied for route advertisements from the BGP neighbor to the appliance.
- **AS list** - Specify the AS paths. The route from the neighbor is permitted if the AS path matches the regular expression in the AS path list.
- **Routemap** - Specify route policy for the BGP neighbor. The routing policies defined by the selected route map are applied while accepting routes from the BGP neighbor.

4. Under outbound, fill out these attributes:

- **Prefix list** - Specify the prefixes to be allowed or denied for route advertisements to the BGP neighbor from the appliance.

- **AS list** - Specify the AS paths. The route to the neighbor is permitted if the AS path matches the regular expression in the AS path list.
- **Routemap** - Specify the route policy for the BGP neighbor. The routing policies defined by the selected route map are applied while accepting routes to the BGP neighbor.

5. Click **Submit.**

Configuring BGP route redistribution

SteelHead SD includes options to globally configure:

- redistribution of OSPF routes into BGP.
- redistribution of static and overlay routes into BGP.
- redistribution of traffic using the route map with the use case for static and connected route injection in BGP.

We recommend you define route policies, AS lists, and prefix lists before you configure BGP route redistribution. For details on configuring routing policies, see "Creating routing IPv4 prefix lists" on page 60.

To configure BGP route redistribution

1. Choose Appliances and select the appliance to expand the pane.

2. Select the BGP tab.

Figure 5-3. Configuring BGP route redistribution

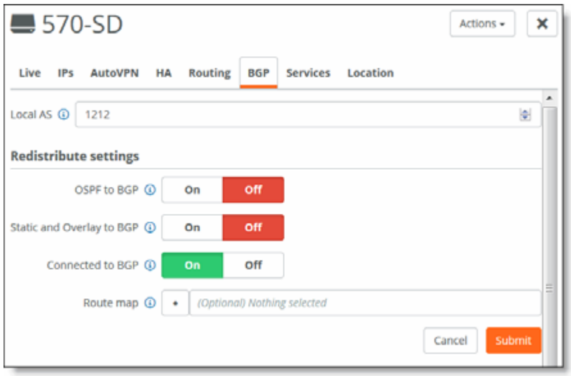

Local AS 1212

Redistribute settings

OSPF to BGP: On [✓ Off]

Static and Overlay to BGP: On [✓ Off]

Connected to BGP: [✓ On] Off

Route map: (Optional) Nothing selected

3. Specify an AS number in Local AS. The range is from 1 to 4294967295.

4. Specify your BGP redistribution settings:

- **OSPF to BGP** - Click **On** to enable redistribution of OSPF routes into BGP. By default, redistribution is disabled.
- **Static and Overlay BGP** - Click **On** enable redistribution of static and overlay routes into BGP. By default, redistribution is disabled.
- **Connected to BGP** - Click **On** or **Off** to enable redistribution of connected routes into BGP. By default, redistribution is disabled.

5. Click the search selector to select the route map. This option only applies to those route maps with the use case of static and connected route injection in BGP. This option redistributes static and connected routes in BGP using a list of IPv4 prefixes.

6. Click **Submit**.

Configuring BGP route summarization

With route summarization, a new network prefix with a shorter prefix length is advertised into BGP. Summarizing prefixes conserves router resources and accelerates best path calculation by reducing the size of the BGP table. Summarization also provides increased stability by reducing routing loops.

You can configure BGP route summarization using one of these modes:

- **Manual** - Creates a static route and advertises the network via a network statement. The summary route will always be advertised even if the networks are not available.
- **Automatic** - Creates a network range. When viable routes that match the network range enter the BGP table, an aggregate route is created. On the originating router, the aggregated prefix sets the next hop to Null 0. The route to Null 0 is automatically created by BGP to prevent routing loops.

When configured, the routing policy advertises a summary address only and not the individual prefixes to a BGP neighbor.

Note: Routing policies only impact the underlay routing. They do not impact the overlay routing orchestrated by SCM.

To configure BGP route summarization

1. Choose Appliances and select the appliance to expand the pane.

2. Select the BGP tab.

3. Under Summarization, click **Add Prefix**.

   You can configure one or more summary addresses matching the individual addresses to advertise to a BGP neighbor. You can also advertise individual addresses. By default, only summary addresses are advertised.

Figure 5-4. Adding AS summarization prefixes

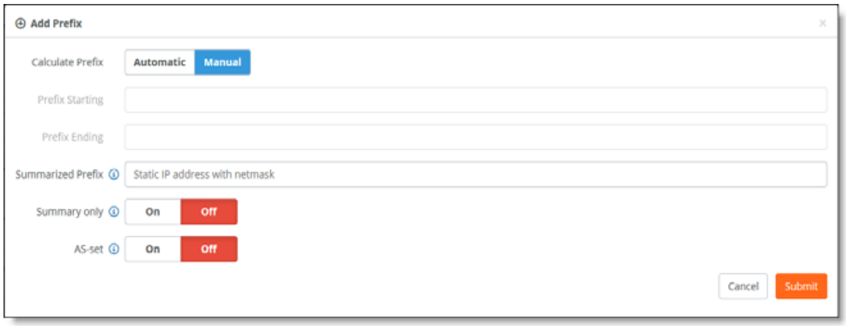

Add Prefix

Calculate Prefix: Automatic [✓ Manual]

Prefix Starting

Prefix Ending

Summarized Prefix: Static IP address with netmask

Summary only: On [✓ Off]

AS-set: On [✓ Off]

   The system default is to calculate the prefix manually.

4. Click **Automatic** to have the system calculate the prefixes automatically, or click **Manual** to specify the prefix.

- For automatic prefix calculation, specify a starting and an ending address, and SteelConnect provides the summarized prefix. For example, entering the starting address 160.0.1.0 and the ending address 160.0.2.0 results in the automatic prefix 160.0.0.0/22.
- For manual prefix calculation, after Summarized Prefix, enter a static IP address with a netmask.

5. Specify the prefix starting and ending point.

6. Specify an IP address with the netmask in the Summarized Prefix text box.

7. Specify your summary and AS-set settings:

- **Summary Only** - Click **On** to advertise both summary and individual prefix advertisements to an eBGP peer.
- **AS-set** - Click **On** to provide an AS-set to use to detect and avoid routing loops. An AS-set summarizes the path attributes of all the BGP individual routes that the aggregate summarizes to help detect and avoid BGP routing loops.

8. Click **Submit**.

Viewing BGP status

SCM displays the advertised and learned network routes and peering session state information. To filter the list, type a search filter in the search box; for example, type IPv6 to narrow the search to all IPv6 networks.

To view BGP neighbors

1. Choose Appliances and select the appliance you want to view.

2. Click the BGP tab to view the BGP neighbors state, received prefixes, remote AS, keep-alive value, hold time, and last error. You can edit BGP values as well.

To view BGP routing tables

1. Choose Health Check > Routing Tables.

2. Select the BGP tab and select the appliance to display the BGP learned and advertised routes.

Configuring OSPF with ABR on SteelHead SD

This section describes how to configure OSPF with ABR on SteelHead SD. It includes these sections:

- “Introducing OSPF with ABR” on page 44
- “Creating an OSPF network” on page 44
- “Configuring OSPF interfaces” on page 47

- "Creating OSPF areas" on page 48
- "Redistributing OSPF settings" on page 50
- "Configuring OSPF route summarization" on page 52
- "Viewing OSPF status" on page 53

Introducing OSPF with ABR

SteelHead SD 2.0 provides single and multiple area OSPF with ABR and route redistribution between OSPF zone interfaces and ABRs on the LAN side of the network. You can configure OSPF regardless of whether it is a zone or an uplink.

SteelHead SD supports OSPF for a branch site with one or two MPLS providers, where each provider is connected to a customer edge (CE) router. A SteelConnect branch gateway is deployed in front of the CE routers. The provider edge (PE) routers on the MPLS WAN side are using BGP and the CE routers on the LAN side are using OSPF.

Figure 5-5. OSPF single area

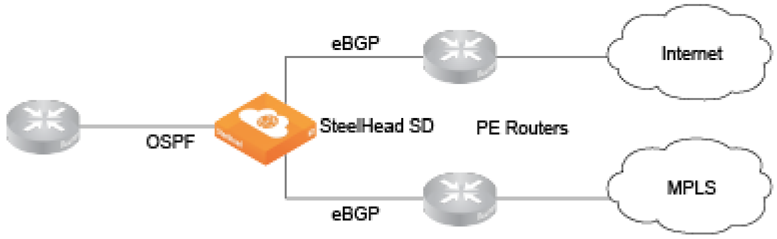

Figure 5-6. OSPF multiple area

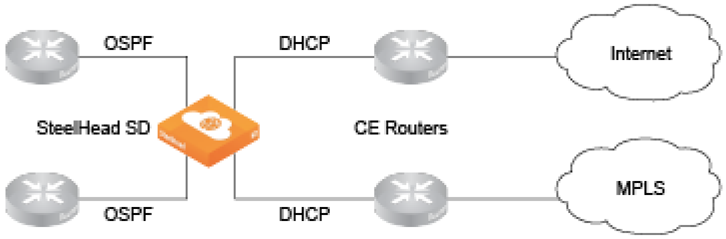

Important: Before you begin configuring OSPF for SteelHead SD, we recommend you consult the *SteelConnect Manager User Guide*.

Creating an OSPF network

You create OSPF networks at the site level. Before configuring OSPF, make sure the appliance is registered in SCM and that OSPF is enabled on any routing device that will peer with the appliance. You create an OSPF network based on a site location that includes one area and then you attach one or more interfaces to the OSPF area.

To create an OSPF network

1. Deploy the SteelHead SD and assign a zone and uplink to a port.

2. Choose Routing > OSPF.

3. Click **New OSPF Network**.

4. Select the site for the new OSPF network.

   After you select a site for an OSPF network, the system automatically populates all the remaining fields based on the default settings. You can simply click **Submit** to create a network using the default settings. You do not have to explicitly configure the settings.

Figure 5-7. Creating an OSPF network on the branch SteelHead SD

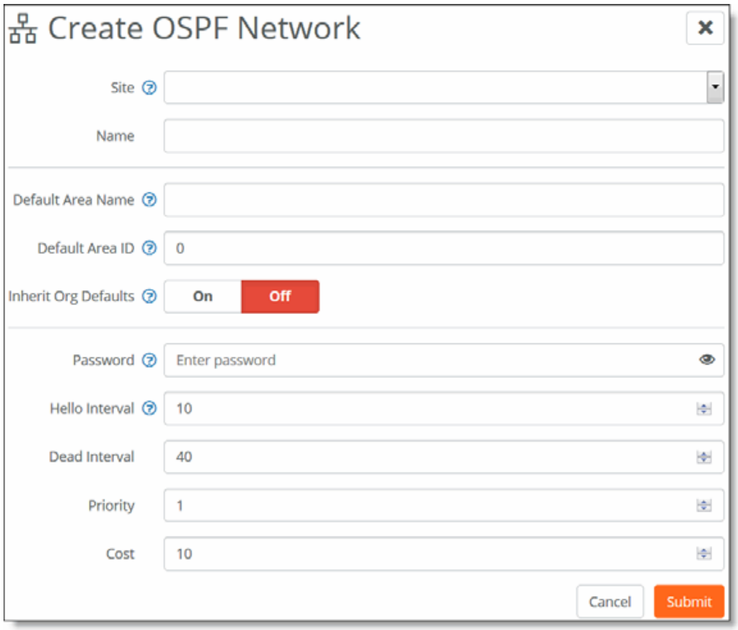

Create OSPF Network

Site	
Name	
Default Area Name	
Default Area ID	0
Inherit Org Defaults	On [✓ Off]
Password	Enter password
Hello Interval	10
Dead Interval	40
Priority	1
Cost	10

5. Manually fill out the network attributes that you do not want to inherit:

- **Site** - Select the site where the OSPF network is located. Optionally, leave the site selection blank to select the first site in the list shown on the Network Designs > Sites page. Use this method to save time by quickly creating OSPF networks based on the order in which the sites appear in the site list. Creating another network and leaving the site selection blank again selects the second site in the list, and so on.

- **Name** - Specify a network name.

- **Default Area Name** - Specify a name for the area.

- **Default Area ID** - Specify the area in which the zone resides. This ID will typically be one of the already existing areas in the branch. If not, either specify a 32-bit unsigned number from 0 to 4294967295 or an IPv4 address in dotted decimal notation (x.x.x.x). The default setting is the backbone area ID 0; however, you can change the value to your existing area ID. For small LANs, area 0 might be all you need, but as a network grows, you will need more than one area connecting to area 0.

For a routing device to become an OSPF neighbor with another device, both devices must belong to the same area ID and their passwords and authentication methods must match.

- **Inherit Org Defaults** - Click **On** to allow the OSPF network and area to automatically inherit the settings when an organization's default network settings are updated. This OSPF network's settings will change to match the new values.

  Click **Off** to define unique settings for the network and to lock the network configuration so any changes do not overwrite the settings.

- **Password** - Specify a password. The authentication methods appear when typing a password. All OSPFv2 exchanges between routing devices can be authenticated using one of these methods:

  - **MD5** - Select this tab to use the Message Digest 5 algorithm as the authentication method. MD5 authentication enables routing devices to securely identify one another before they establish adjacency. MD5 is a cryptographic hash function with a 128-bit hash value derived from the contents of the OSPF packet and a key and key ID. This method doesn't send the password but instead calculates and includes an encoded MD5 checksum in the transmitted packet. The receiving routing device uses the key and key ID to verify the packet.

    The MD5 key doesn't have to be the same within the area, but it must be exactly the same between two OSPF neighbors.

Tip: Click the eye icon to see the password as you type. The view persists until you click the eye icon again to hide the password.

  - **Simple** - Select this tab to include an unencrypted plain text password with the packet. The receiving routing device uses the password to verify the packet. The simple password can be from one to eight characters and can include ASCII strings. If you include spaces, enclose the password in quotation marks. Use this authentication method when devices within an area do not support the more secure MD5 authentication, as Simple is the least secure setting.

- **MD5 Key ID** - (Appears when you select MD5.) Specify a value to associate with the MD5 key. The ID is used by the receiver of the OSPF packet to determine which key to use for authentication.

  To change your MD5 key, specify a new key and key ID. When both OSPF neighbors have a new key and key ID, the old key is deleted and the current MD5 key and key ID become active.

- **Hello Interval** - Specify how often, in seconds, to send a hello packet. Initially the gateway sends a hello packet to all OSPF-enabled interfaces to form an adjacency as a neighbor. The routing devices become neighbors and exchange link-state advertisements. After the gateway learns the common network topology, it sends the hello to check if an OSPF neighbor is alive. The range is from 1 to 65535. The default is 10. The hello interval must be exactly the same between two OSPF neighbors.

- **Dead Interval** - Specify how many seconds to wait for a hello packet before declaring an OSPF neighbor out of service, triggering a refresh of the link-state database and routing information. The range is from 1 to 65535. The default is 40. The dead interval must be exactly the same between two OSPF neighbors.

- **Priority** - Specify the priority for becoming the network's designated routing device. The designated router originates network link advertisements on behalf of the network, and it establishes adjacencies with all routing devices on the network.

The routing device that has the highest priority value on the logical IP network or subnet is elected as the designated router. A priority value of 0 means that the routing device never becomes the designated router; it doesn't even participate in the election process. A value of 1 means that the routing device participates in the election process but has the least chance of becoming a designated router. A priority of 255 means the routing device is always the designated router.

To ensure that a routing device is elected as the designated routing device, configure the priority value to a higher value than any other interface on the Ethernet network. The range is from 0 to 255. The default value is 1.

- **Cost** - Specify a routing metric used in the link-state calculation. OSPF selects ideal routes by locating destination routes with the least cost. Routes with lower total path metrics are preferred to those with higher path metrics. This setting controls the cost calculation of OSPF network segments. The default formula to calculate the cost for the OSPF metric is dividing the reference bandwidth (100 Mbps by default) by the interface bandwidth. For example, in the case of Ethernet, it is 100 Mbps / 10 Mbps = 10.

  You can manipulate the cost by specifying a number within the range of 1 to 65535. 10 is the default setting.

The OSPF network needs a zone and, optionally, one or more uplinks to report OSPF learned routes to SCM.

Tip: If you modify the Default Area settings, keep in mind the impact the changes will have on new and existing OSPF networks. Changes to the Default Area Name, Default Area ID, and Inherit Org Defaults impact only new OSPF networks. Changes to the Password, Hello Interval, Dead Interval, Priority, and Cost impact new OSPF networks as well as existing OSPF networks with Inherit Org Defaults enabled. For details on editing OSPF networks, see the *SteelConnect Manager User Guide*.

6. Click **Submit.** The OSPF network appears with the available interfaces listed.

Figure 5-8. Created OSPF network

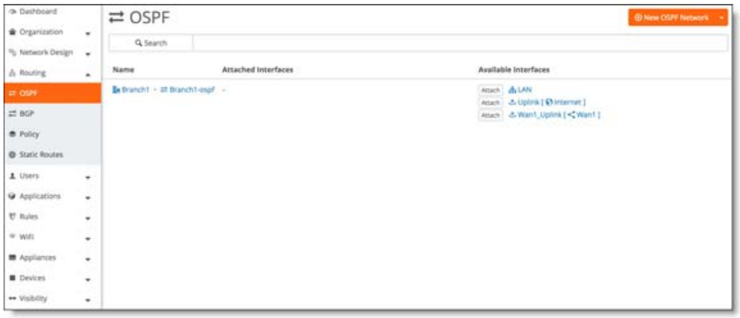

Name	Attached Interfaces	Available Interfaces
Branch1 · Branch1-ospf	-	Attach LAN Attach Uplink [Internet] Attach Wan1_Uplink [Wan1]

Configuring OSPF interfaces

After you define your OSPF network, you must attach interfaces on which you want to run OSPF.

To configure OSPF interfaces

1. Choose Routing and select the OSPF network for which you want to attach an interface.

2. Select the OSPF Interfaces tab.

3. Click **Attach Interface**.

Figure 5-9. Attaching an OSPF interface

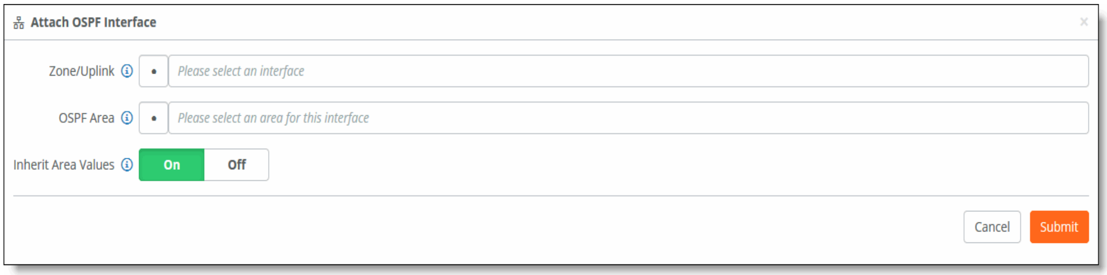

Attach OSPF Interface

Zone/Uplink: Please select an interface

OSPF Area: Please select an area for this interface

Inherit Area Values: On Off

4. Fill out these interface attributes:

- **Zone Uplink** - Select the zone or uplink to attach to the OSPF area.
- **OSPF Area** - Select the OSPF area associated with the interface from the drop-down list.
- **Inherit Area Values** - Click **On** to allow the interface to automatically inherit the area settings. When enabled and this interface area is updated, this interface settings will change to match the containing OSPF area.

  Click **Off** to define unique settings for the area. This option locks the interface configuration so any changes to the area do not overwrite the interface parameters.

5. Click **Submit**.

After you attach the interface to the OSPF area, the gateway configures the zone or zones to run OSPF and establishes OSPF neighbors with LAN routers in the same network segment.

Creating OSPF areas

All of the networks learned from an OSPF zone interfaces are mapped to the OSPF area that the interface is connected to. For details on dynamic routing with OSPF, we recommend you consult the *SteelConnect Manager User Guide*.

A large OSPF domain is broken into separate areas to restrict the multiplication of routes and reduce the resources required by each router to maintain its link state database. Each area is connected to a central backbone, typically called area 0. OSPF uses different types of Link State Advertisements (LSAs) to communicate link state information between neighbors.

SteelHead SD supports these LSA types:

- **Standard** - Routers in this area accept default and autonomous system boundary router (ASBR) injected external routes. The backbone is considered a standard area.
- **Stub** - Routers in this area accept inter-area routes and the default route from their ABR. They do not accept ASBR injected external routes. A stub type can contain type 1, 2, and 3 LSAs.
- **Totally Stub** - This type of router is similar to a stub router. They accept inter-area routes and the default route from their ABR. They do not accept ASBR injected external routes. A totally stub type can only contain type 1 and 2 LSAs, and a single type 3 LSA. The type 3 LSA describes a default route, substituted for all external and inter-area routes.

To create an OSPF area

1. Choose Routing > OSPF and select the OSPF network for which you want to create an area.

2. Select the Areas tab and click **New Area**.

Figure 5-10. Creating OSPF areas

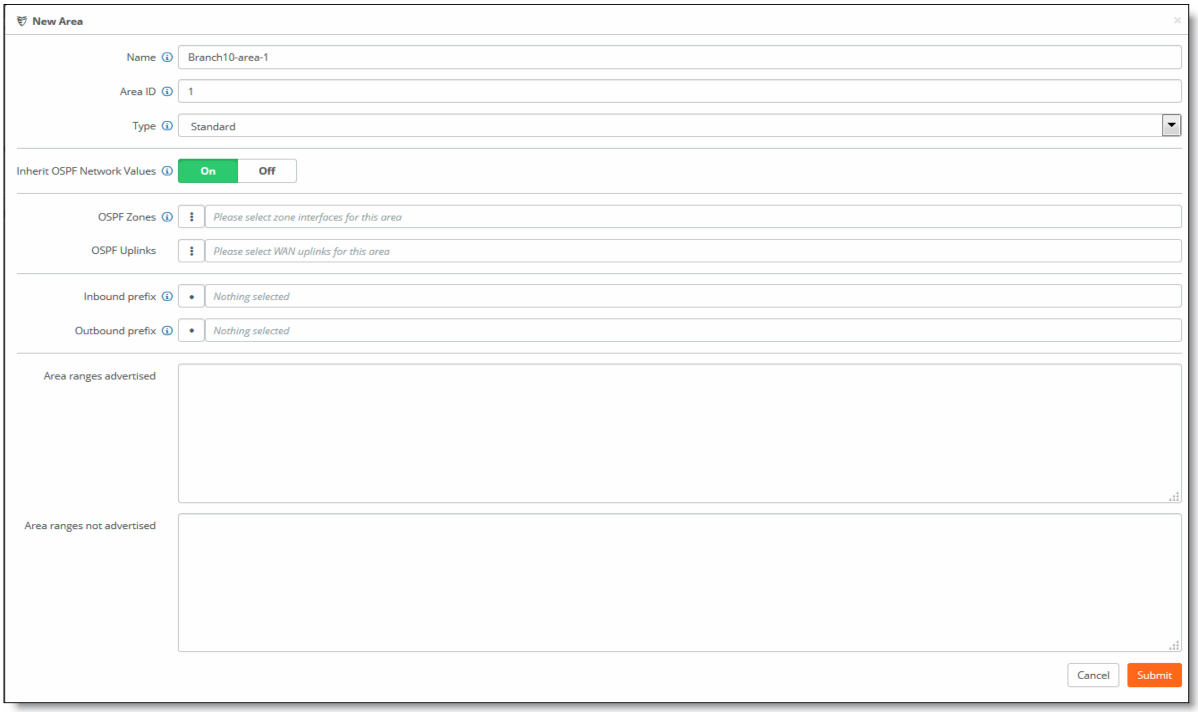

New Area

Name: Branch10-area-1

Area ID: 1

Type: Standard

Inherit OSPF Network Values: On Off

OSPF Zones: Please select zone interfaces for this area

OSPF Uplinks: Please select WAN uplinks for this area

Inbound prefix: Nothing selected

Outbound prefix: Nothing selected

Area ranges advertised

Area ranges not advertised

3. Fill out the attributes for the OSPF area:

- **Name** - Specify a descriptive name for the OSPF area.
- **Area ID** - Specify a valid area ID as either a 32-bit unsigned number from 0 to 4294967295 or an IPv4 address in dotted decimal notation (x.x.x.x). The default setting is the backbone area ID 0; however, you can change the value to your existing area ID. For small LANs, area 0 might be all you need, but as a network grows, you will need more than one area connecting to area 0.
- **Type** - Specify the OSPF LSA type:
  - **Standard** - Routers in this area accept default and autonomous system boundary router (ASBR) injected external routes. The backbone is considered a standard area.
  - **Stub** - Routers in this area accept inter-area routes and the default route from their ABR. They do not accept ASBR injected external routes. A stub type can contain type 1, 2, and 3 LSAs.
  - **Totally Stub** - This type of router is similar to a stub router. They accept inter-area routes and the default route from their ABR. They do not accept ASBR injected external routes. A totally stub type can only contain type 1 and 2 LSAs, and a single type 3 LSA. The type 3 LSA describes a default route, substituted for all external and inter-area routes.
- **Inherit OSPF Network Values** - Click **On** to allow the OSPF network to inherit the OSPF network values previously configured, such as password, hello interval, dead interval, priority, and cost.

Click **Off** to define unique settings for the network and to lock the network configuration so any changes do not overwrite the settings. This OSPF network's settings will change to match the new values.

- **OSPF Zone** - Select the zone from the list. These are the zones that are participating in OSPF for the area that is configured on this page. Only one zone interface per area is allowed.
- **OSPF Uplinks** - Select the uplinks from the list. These are the uplinks that will be participating in OSPF in the area that is configure on this page.
- **Inbound prefix** - Optionally, specify the inbound prefix. Any prefixes defined in this prefix list are used to filter networks sent to this area.
- **Outbound prefix** - Optionally, specify the outbound prefix. Any prefixes defined in this prefix list are used to filter networks advertised from this area.
- **Area ranges advertised** - Specify a set of advertised routes to be advertised. In order to aggregate routing information at area boundaries, area address ranges can be employed. Each address range is specified by an [address, mask] pair.
- **Area ranges not advertised** - Specify the set routes that will not be advertised. In order to aggregate routing information at area boundaries, area address ranges can be employed. Each address range is specified by an [address, mask] pair. In this case, Type 3 summary-LSA is suppressed and the component networks remain hidden from other areas.
- Click **Submit.**

Redistributing OSPF settings

The LAN/WAN routing interworking solution bridges eBGP and OSPF to redistribute underlay routing information between the protocols on a gateway.

For details on how redistribution works, see the *SteelConnect Manager User Guide*.

To redistribute OSPF settings

1. Choose Routing and select the OSPF network.

2. Select the Redistribute Settings tab.

Figure 5-11. Redistributing OSPF settings

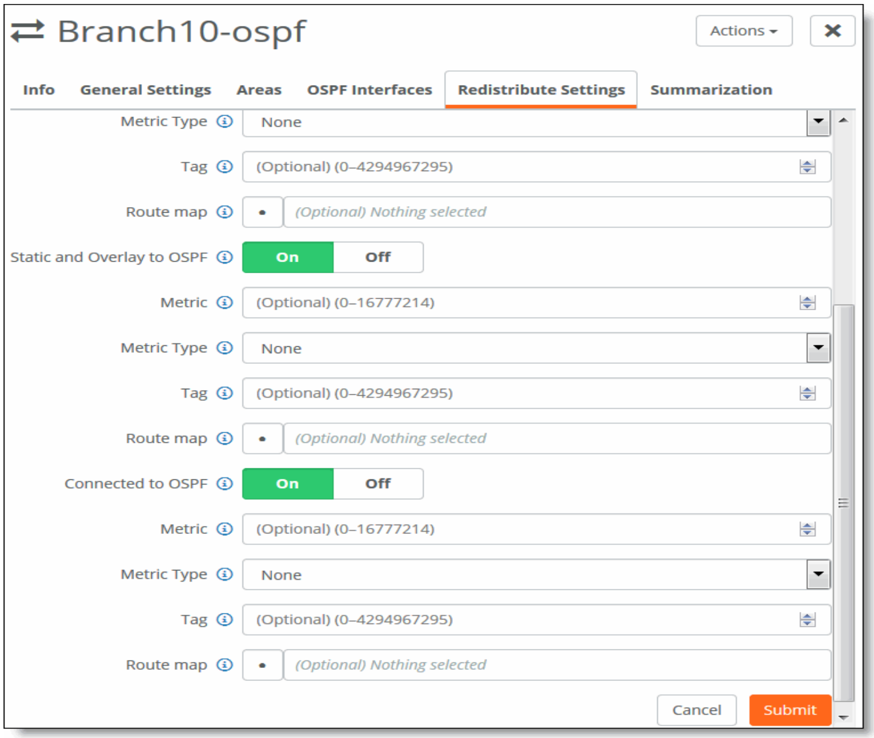

Metric Type: None

Tag: (Optional) (0–4294967295)

Route map: (Optional) Nothing selected

Static and Overlay to OSPF: On | Off

Metric: (Optional) (0–16777214)

Metric Type: None

Tag: (Optional) (0–4294967295)

Route map: (Optional) Nothing selected

Connected to OSPF: On | Off

Metric: (Optional) (0–16777214)

Metric Type: None

Tag: (Optional) (0–4294967295)

Route map: (Optional) Nothing selected

3. Optionally, specify the default metric with a range of 1 to 16777214. The ABR generates a default route with a specified metric into the stub area. The default route matches any destination that is not explicitly reachable from within the area. Routing protocols use default metrics to calculate the best path to a specified destination. The routes that are redistributed carry the specific value.

4. Specify whether you want the default route (0.0.0.0/0) injected in OSPF:

- **Default Route Origination** - Click **On** to enable default route origination. Enabling this option injects a default route into the participating areas in OSPF.

5. Specify your OSPF redistribution settings:

- **BGP to OSPF** - Click **On** to redistribute the routes learned from BGP into the OSPF protocol.
- **Static and Overlay to OSPF** - Click **On** to redistribute static and overlay routes into OSPF.
- **Connected to OSPF** - Click **On** to redistribute connected routes into OSPF.

6. Optionally, specify these settings if any of the above OSPF redistribution settings is enabled:

- **Metric** - Optionally, enter the cost metric that you want the route to be injected with into OSPF. The range is 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router.
- **Metric type** - The type of external route that you want the routes to be injected as. When the type matches the value specified, then that route is qualified to be distributed:

  - **Type 1 (EI)**- This type includes the external cost to the destination as well as the cost (metric) to reach the AS boundary router.
  - **Type 2 (E2)** - This type uses only the external cost to the destination and ignores the cost (metric) to reach the AS boundary router.
- **Tag** - Optionally, enter a value 32 bit value from 0 to 4294967295 that will be attached to the routes. When a tag in route matches this value, the route qualifies for distribution by the router.
- **Route map** - Click the search selector and select the route map. This option applies a routing policy based on which routes will be redistributed into OSPF.

7. Click **Submit.**

Configuring OSPF route summarization

For an OSPF area, you can filter intra-area prefixes. All routes that match the specified area range are filtered.

To add summarization for OSPF

1. Choose Routing > OSPF.
2. Select the OSPF network.
3. Select the OSPF Summarization tab.
4. Click **Add Prefix** to add prefixes.

You can configure one or more summary addresses matching the individual addresses to advertise to a OSPF peer. You can also advertise individual addresses. By default, only summary addresses are advertised.

Figure 5-12. Summarizing routes for OSPF

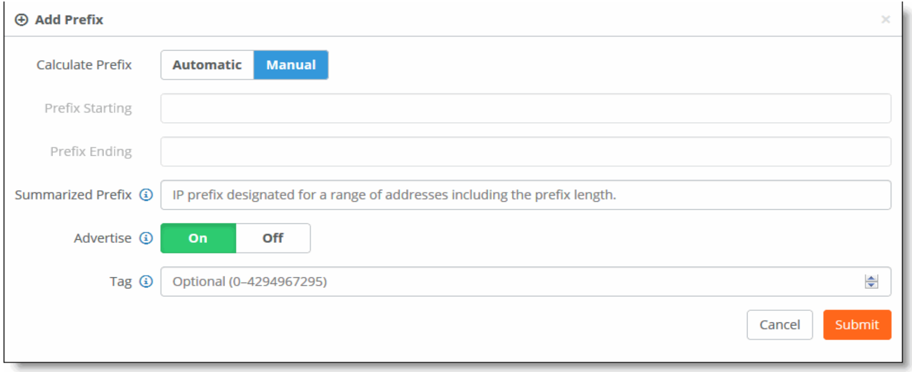

Add Prefix

Calculate Prefix: Automatic [✓ Manual]

Prefix Starting

Prefix Ending

Summarized Prefix ⓘ IP prefix designated for a range of addresses including the prefix length.

Advertise ⓘ [✓ On] Off

Tag ⓘ Optional (0–4294967295)

5. Click **Automatic** to have SCM calculate the prefixes automatically, or click **Manual** to specify the summarized prefix.

   For automatic prefix calculation, specify a starting and an ending IP address. The system provides the summarized prefix. For example, entering the starting address 160.0.1.0 and the ending address 160.0.2.0 results in the automatic prefix 160.0.0.0/22.

6. Fill out the these attributes for automatic or manual:

- **Summarized Prefix** - Specify the IP prefix designated for the range of addresses, including the prefix length.
- **Advertise** - Click **On** to advertise the summary prefix. Click **Off** to stop advertisements of the summary address.
- **Tag** - Specify a 32-bit value attached to the summary route. The specified value will be tagged to the advertised summary routes.

7. Click **Submit**.

Viewing OSPF status

There are multiple places where SCM provides visibility to OSPF and the state of routes.

To view OSPF routing tables

1. Choose Health Check > Routing Tables.

2. Select the OSPF tab and select the appliance to display the OSPF neighbors and learned routes.

To view the Forward Information Base (FIB) routing table

1. Choose Health Check > Routing Tables.

2. Select the FIB tab and select the appliance to display the FIB information, including destination, next hop, metric value, route type, and subroute type. This table is very useful and should be the first step in debugging if the expected routes are learned by the appliance.

Defining static routing on SteelHead SD

SteelHead SD provides static routing at the appliance level where it is essentially acting as a router. The static route is not tied to a particular zone. Static routes:

- can only be defined on SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.
- support IPv4 destination networks and gateways.

To define static routes

1. Choose Routing > Static Routes.

2. Click **Add Static Route.**

Figure 5-13. Adding static routes

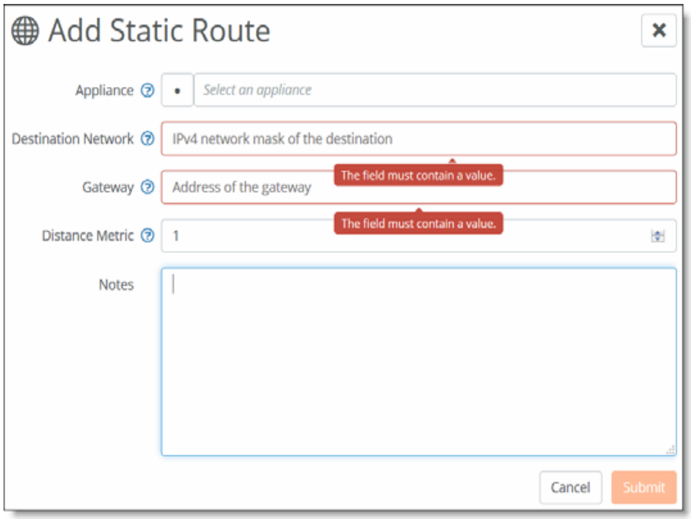

Add Static Route

Appliance: Select an appliance

Destination Network: IPv4 network mask of the destination

The field must contain a value.

Gateway: Address of the gateway

The field must contain a value.

Distance Metric: 1

Notes:

3. Select the Appliance to which you want to add the static route. Only SteelHead SD appliances are listed. Static routes can only be applied to SteelHead SD models; other non-SteelHead SD appliances are not listed.

4. Specify the IPv4 destination mask address.

5. Specify the IPv4 address for the gateway.

6. Specify the destination metric to prioritize the routing protocol where two routes provide the same route destination. The default value is 1.

7. Optionally, include any notes that will help you identify this static route.

8. Click **Submit.**

Route retraction for SteelHead SD

SteelHead SD advertises available routes and doesn't advertise unavailable routes. If a route becomes unavailable, route retraction withdraws this route and ensures it is no longer advertised.

Note: The behavior of route attraction and retraction is the same across all SteelConnect appliances.

To benefit from route retraction on a SteelHead SD, you need to meet the following requirements:

- You need to redistribute the overlay network into the internet gateway protocol on the LAN.
- For SteelHead SD appliances deployed in HA mode, you need to redistribute the overlay network and connected routes into iBGP.

Here are the detailed steps.

To redistribute the overlay into the internet gateway protocol on the LAN

1. In SCM, choose Routing > OSPF.
2. Select your OSPF network to edit the settings.
3. In the Redistribute settings tab, click **On** for Static to OSPF, and click **On** for Connected to OSPF.
4. Click **Submit.**

To redistribute the overlay and connected routes into iBGP for HA pairs

1. In SCM, choose Appliances > Overview.
2. Select an HA site.
3. Open the BGP tab and click **On** for Static and Overlay to BGP and click **On** for Connected to BGP.
4. Click **Submit.**

6

Configuring ASBR Routing Policies on SteelHead SD

This topic describes how to configure autonomous system boundary routers (ASBR) and route policies on SteelHead SD 2.0. It includes these sections:

- "Introducing ASBR-full route policies on SteelHead SD" on page 57
- "What are routing policies?" on page 58
- "Creating routing IPv4 prefix lists" on page 60
- "Creating routing community lists" on page 61
- "Creating routing AS path lists" on page 62
- "Configuring use case route maps" on page 63

Introducing ASBR-full route policies on SteelHead SD

SteelHead SD appliances act as a full ASBR when they are located at the branch. ASBR-full routing policies are supported on SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch. ASBR-full support is not available on SteelConnect SDI-130/330/1030 and virtual gateways.

An ASBR is a router that is connected to several autonomous systems (ASs) using multiple protocols. Typically, ASBRs are connected via an exterior routing protocol (for example, BGP). An ASBR can also connect LAN routers through an interior gateway protocol (IGP), such as OSPF within its own AS. Basically, with an ASBR you are distributing routes from BGP to OSPF and from OSPF to BGP. An ASBR can also distribute static and connected routes into these protocols.

For example, if you have a SteelHead SD on Branch-1 with eBGP configured on the WAN uplink and OSPF configured on the LAN zones. The SteelHead SD can distribute LAN routes to the WAN and WAN routes to the LAN. This method of redistributing routes received via one routing protocol through another protocol is called *route redistribution* or *route injection*.

Figure 6-1. ASBR deployment in branch 1

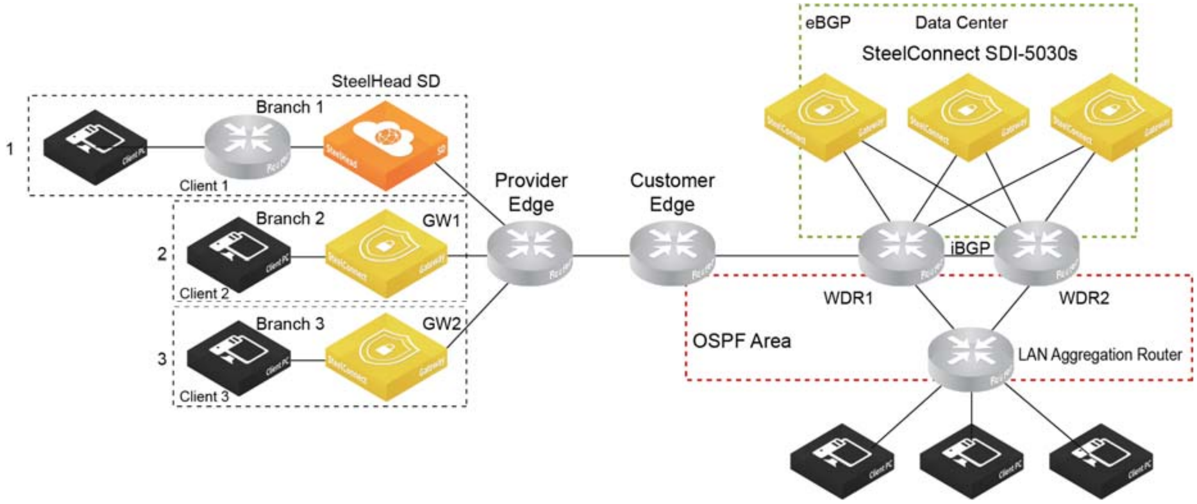

SteelHead SD provides:

- **BGP redistribution** - Support includes static and connected route redistribution, OSPF route redistribution, and default route redistribution to BGP neighbors.
- **OSPF redistribution** - Support includes static and connected route redistribution, OSPF route redistribution, and default route redistribution.

Important: ASBR routing policies are available only on underlay branch networks.

What are routing policies?

Routing polices are rules that are applied when routes are distributed between the routers. Creating routing policies enables you to redistribute BGP, OSPF, static, and connected routes.

Note: ASBR routing policies are not policy-based routing where routing decisions are made while directing the traffic.

Creating routing policies enables you to apply certain rules and route attributes while redistributing BGP, OSPF, static and connected routes. You can create route-maps for the following purposes:

- Route injection in OSPF.
- Default route origination in OSPF.
- Static and connected route injection BGP.
- OSPF route injection in BGP.
- Policies at the BGP neighbor level.

- Default route origination in BGP for a neighbor.

Each route map clause has two types of values:

- A match value selects routes to which the clause should be applied.
- A set value modifies information that will be redistributed into the target protocol.

For each route that is being redistributed, the router evaluates the match criteria of a clause in the route map. If the match criteria succeed, then the route is redistributed and some of the attributes may be modified by the set clause. If the route doesn't match any clause in a route map, then the route redistribution is denied.

Once configured, the route-maps can be applied to satisfy the needs of these use-cases:

- **Route Injection in OSPF** - OSPF redistributes static, connected, and BGP routes. This route-map category contains only a match criteria. The route map depends on the following objects:
  - IP prefix lists
  - Interface (all zones/uplinks associated with SteelHead SD appliances in the given organization)
- **Default Route Origination in OSPF** - Redistributes the default route in OSPF. This category of route-map contains both match and set criteria. This is the simplest route-map category that is not dependent on other objects.
- **Static and Connected Route Injection in BGP** - Redistributes static and connected routes in BGP using a list of IPv4 prefixes. This route-map category contains both match and set criteria. Also dependent on the following objects:
  - IPv4 prefix lists
  - Interface
- **OSPF Route Injection in BGP** - Redistributes OSPF routes in BGP using a list of IPv4 prefixes. This category of route map contains both match and set criteria.
- **Route-Policy for BGP Neighbor** - Redistributes routes for BGP neighbors using a community list and list of IP next-hop prefixes. This route-map category contains both match and set criteria. The match criteria in this use case is dependent on:
  - Community list
  - Prefix list
- **Default Route Origination in BGP** -Redistributes the route in BGP. This route-map category contains both match and set criteria. There are not any dependent objects for this type of route map.

Basic steps

Perform these basic steps to configure ASBR routing polices.

1. If you have a SteelConnect SDI-2030 gateway, configure a dynamic routing policy. For details, see "Creating routing IPv4 prefix lists" on page 60.

Note: You can't create dynamic routing policies for SteelHead SD 570-SD, 770-SD, and 3070-SD appliances located at the branch.

2. Configure IPv4 prefix lists. For details, see “Creating routing IPv4 prefix lists” on page 60.

3. Configure community lists. For details, see “Creating routing community lists” on page 61.

4. Configure AS prefix lists. For details, see “Creating routing AS path lists” on page 62

5. Configure route maps by specifying the available use cases. For example, create a route map for a routing policy to establish BGP neighbors. For details, see “Creating routing IPv4 prefix lists” on page 60

6. Configure inbound and outbound route maps and prefixes for BGP neighbors using the configured route maps. For details on configuring BGP neighbors, see “Configuring BGP on SteelHead SD” on page 38.

7. Configure BGP redistribution and BGP summarization settings using the configured route maps. For details, see “Configuring BGP route redistribution” on page 41 and “Configuring BGP route summarization” on page 42.

Creating routing IPv4 prefix lists

An IPv4 prefix list contains a list of IPv4 prefixes and a name that is associated for each list.

To create a IPv4 prefix list

1. Choose Routing > IPv4 Prefix Lists.

2. Click **New IPv4 Prefix List**.

Figure 6-2. Creating an IPv4 prefix list

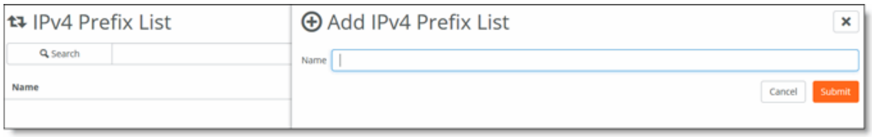

3. Specify the name of the IPV4 prefix list.

4. Click **Submit**.

5. To define the prefixes for the list, select the list in the IPv4 Prefix List page.

Figure 6-3. Defining IPv4 prefixes list

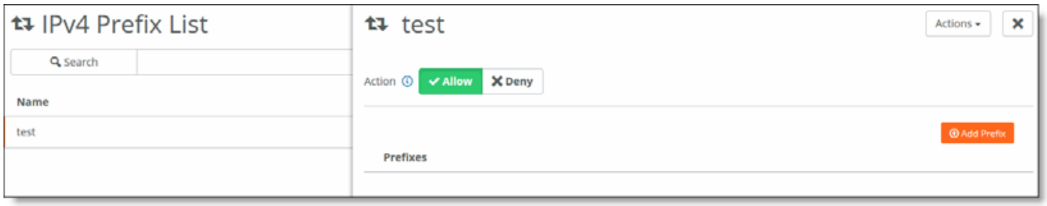

6. Click **Allow** to distribute only the specified prefixes and deny the rest. Click **Deny** to stop distribution of the prefixes specified and allow the rest.

7. Click **Add Prefix.**

Figure 6-4. Adding a prefix

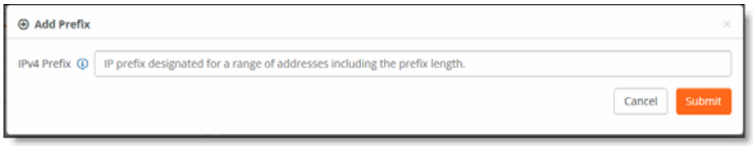

IP prefix designated for a range of addresses including the prefix length.

8. Enter the IP prefix designated for the range of addresses to distribute. Use the format: xxx.xxx.xxx.xxx/xx

9. Click **Submit.**

Tip: Click **Actions** to delete a list.

Creating routing community lists

A BGP community is a group of routes to which a BGP router applies the same policies. You specify the name of the community list and a string that contains values only from a predefined set of keywords and numbers.

To create a community list

1. Choose Routing > Community Lists.

2. Click **Add Community List.**

Figure 6-5. Creating a community list

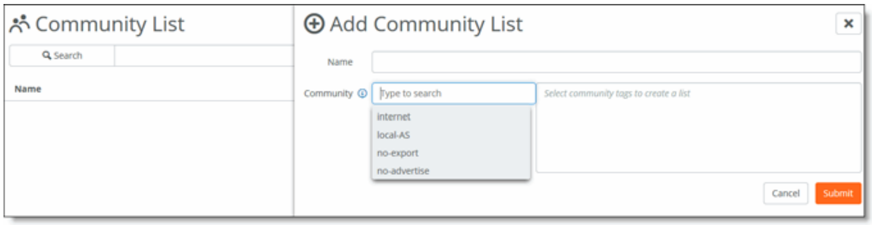

3. Enter a descriptive name for the community list.

4. Click the search selector for community list options. In addition to the keywords below, you can also configure numbers in the range from 1 to 65535 and numbers in AA:NN format where the range for AA and NN is 1 to 65535.

- **internet** - Advertises this route to the internet community; by default, all prefixes are members of the internet community.
- **local-AS** - Doesn't advertise the route to any external peers.

- **no-export** - Instructs routers not to export a prefix to eBGP neighbors. For instance, subnets of a larger block can be advertised to influence external AS best-path selection, and those not required for this traffic engineering purpose may be tagged NO-EXPORT to prevent them from being leaked to the internet (and thus contributing to unnecessary global routing table growth).
- **no-advertise** - Instructs a BGP router not to advertise the tagged prefix to any other neighbor, including other iBGP or eBGP routers.

5. Click **Submit.**

6. To edit a community list, select the list in the Community List page, edit the expressions, and click **Submit.**

Figure 6-6. Editing an AS path list

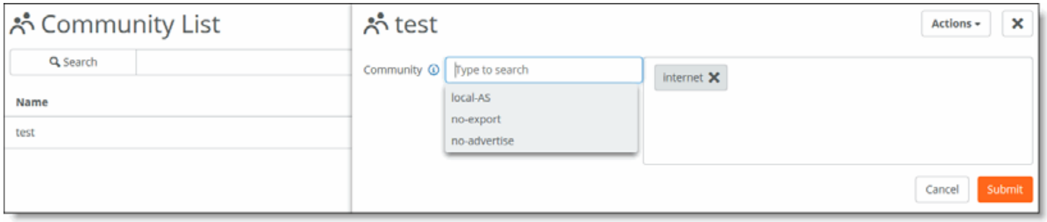

Tip: Click **Actions** to delete a list.

Creating routing AS path lists

You specify the name of the AS path list and define a regular expression that defines the attributes of the AS path.

The AS path list can be used while applying route policies at the BGP neighbor level.

To create an AS path list

1. Choose Routing > AS Path List.

2. Click **Add AS Path List** to expand the page.

Figure 6-7. Creating an AS path list

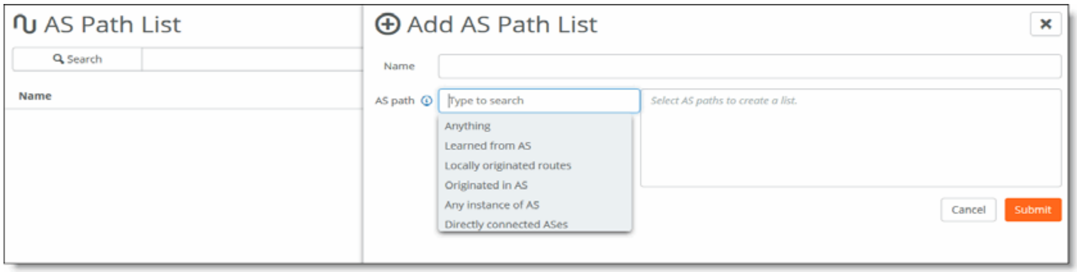

3. Enter a descriptive name for the AS path list.

4. Click the search selector for a list of AS list options. Enter one or more AS numbers from 1 to 4294967295. Separate multiple numbers with a space.

- **Anything** - Specifies the BGP expression ".*", which matches anything. The ".*" matches any single character ("."), and then finds zero or more instances of that single character ("*").
- **Learned from AS** - Enter one or more AS numbers from 1 to 4294967295. Separate multiple numbers with a space.
- **Locally originated routes** - Specifies the BGP expression: "^$", which matches locally originated routes. "^$" means that the string is null. Within the scope of BGP, the only time that the AS path is null is when you are looking at a route within your own AS that you or one of your iBGP peers has originated.
- **Originated in AS** - Enter one or more AS numbers from 1 to 4294967295. Separate multiple numbers with a space.
- **Any instance of AS** - Enter one or more AS numbers from 1 to 4294967295. Separate multiple numbers with a space.
- **Directly connected to AS** - Specifies the BGP expression "^[0-9]+$", which matches all routes originated in any directly connected single AS. These are the routes directly originated by the peers of your AS.

5. Click **Submit.**

6. To edit an AS path list, select the list in the AS Path page, edit the expressions, and click **Submit.**

Figure 6-8. Editing an AS path list

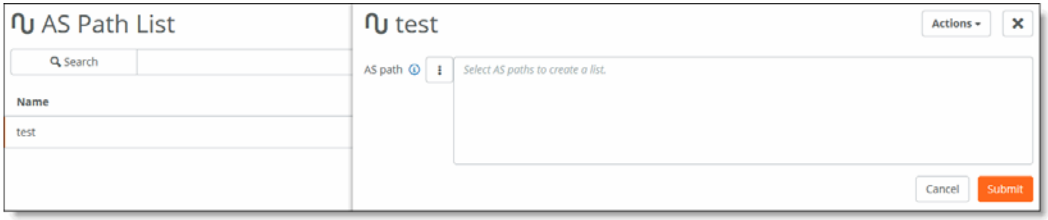

Tip: Click **Actions** to delete a list.

Configuring use case route maps

After you configure the prefix list, AS list, and community lists, these lists can be applied to satisfy the needs of particular use cases.

To create use case route maps

1. Choose Routing > Route Maps.

2. Click **New Route Map**.

Figure 6-9. Creating route maps

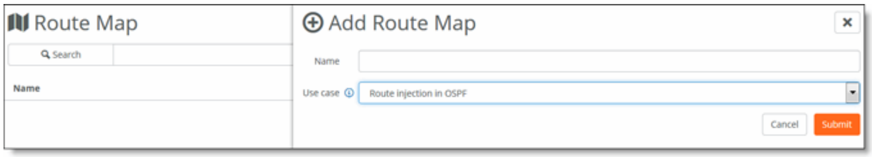

Use case: Route injection in OSPF

3. Specify the name of the route map.

4. Select a use case from the drop-down list:

- **Route injection in OSPF** - Allows the creation of match clauses that can be applied during BGP, static, and connected route injection in OSPF.
- **Default route origination in OSPF** - Allows the creation of match and set clauses that can be used during the default route origination in OSPF.
- **Static and connected route injection in BGP** - Allows the configuration of match and set clauses that can applied while redistributing static and connected routes in BGP.
- **OSPF route injection in BGP** - Allows the creation of match and set clauses that can be applied while redistributing OSPF routes in BGP.
- **Policies at a BGP neighbor level** - Allows the configuration of match and set clauses that can be applied while establishing a BGP neighbor.
- **Default route origination in BGP for a neighbor** - Allows the configuration of match and set clauses that can be applied while advertising a default route to a BGP neighbor.

5. Click **Submit.** The route map is displayed on the Route Map page.

6. In the Route map page. Match Criteria and Set Criteria tabs are displayed depending on the match and set requirements for each use case.

Figure 6-10. Match Criteria and Set Criteria tabs

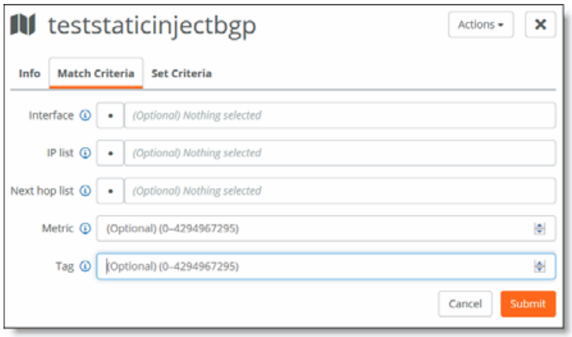
teststaticinjectbgp
Actions ▾
×
Info
Match Criteria
Set Criteria
Interface
(Optional) Nothing selected
IP list
(Optional) Nothing selected
Next hop list
(Optional) Nothing selected
Metric
(Optional) (0–4294967295)
Tag
(Optional) (0–4294967295)
Cancel
Submit

7. Fill out the fields for the Match Criteria and Set Criteria using this table. The criteria differ according to the use case you have chosen.

Use case	Match criteria	Set criteria
Route injection in OSPF	- **Interface**- Optionally, select the interface. When the interface matches the next-hop interface of the route, the route qualifies for redistribution by the router. - **IP list** - Optionally, select an IP list. When a routes prefix address matches a prefix in the list, then that route is qualified for distribution. - **Next hop list** - Optionally, select the next-hop prefix. When the next-hop address matches the selected address, the route qualifies for distribution by the router. - **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. - **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	No set criteria required.
Default route origination in OSPF	- **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. - **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	- **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.

Use case	Match criteria	Set criteria
Static and connected route injection in BGP	- **Interface** - Optionally, click the search selector and select the interface. When the interface matches the next-hop interface of the route, the route qualifies for redistribution by the router. - **IP list** - Optionally, select the IPv4 prefix list. - **Next hop list** - Optionally, select the next-hop prefix. When the next-hop address matches the selected address, the route qualifies for distribution by the router. - **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. - **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	- **AS path** - Click **On** to set the AS path for the route. Specify the AS string as space separated list from 1 to 4294967295. - **Tag** - If **On**, then the value is prepended with the AS path of the BGP route. - **IP next hop** - If **On**, then updates the IP next-hop address of the routes. Enter the IP address to be used as the next hop. - **Self address** - If **On**, under Self address, click **On** to use the self address as the next-hop address. - **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. - **Community** - In addition to the keywords below, you can also configure numbers in the range from 1 to 65535 and numbers in AA:NN format where the range for AA and NN is 1 to 65535. – internet – local-AS – no-export – no-advertise - **Additive** - The specified community string is added to the route's community string.

Use case	Match criteria	Set criteria
OSPF route injection in BGP	▪ **Interface** - Optionally, select the interface. When the interface matches the next-hop interface of the route, the route qualifies for redistribution by the router. ▪ **IP list** - Optionally, select the IPv4 prefix list. ▪ **Next hop list** - Optionally, select the next-hop prefix. When the next-hop address matches the selected address, the route qualifies for distribution by the router. ▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Metric type** - When the type matches the value specified, then that route is qualified to be distributed: – **Type 1** - This type includes the external cost to the destination as well as the cost (metric) to reach the AS boundary router. – **Type 2** - This type uses only the external cost to the destination and ignores the cost (metric) to reach the AS boundary router. ▪ **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	▪ **AS path** - Click **On** to set the AS path for the route. Specify the AS string as space separated list from 1 to 4294967295. ▪ **Tag** - If **On**, then the value is prepended with the AS path of the BGP route. ▪ **IP next hop** - If **On**, then updates the IP next-hop address of the routes. Enter the IP address to be used as the next hop. ▪ **Self address** - If **On**, under Self address, click **On** to use the self address as the next-hop address. ▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Community** - In addition to the keywords below, you can also configure numbers in the range from 1 to 65535 and numbers in AA:NN format where the range for AA and NN is 1 to 65535. – internet – local-AS – no-export – no-advertise ▪ **Additive** - The specified community string is added to the route’s community string.

Use case	Match criteria	Set criteria
Policies at the BGP neighbor level	▪ **Community** - Optionally, select the community list. A BGP route is permitted if it belongs to the specified community string. ▪ **Next hop list** - Optionally, select the next-hop prefix. When the next-hop address matches the selected address, the route qualifies for distribution by the router. ▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	▪ **AS path** - Click **On** to set the AS path for the route. Specify the AS string as space separated list from 1 to 4294967295. ▪ **Tag** - If **On**, then the value is prepended with the AS path of the BGP route. ▪ **IP next hop** - If **On**, then updates the IP next-hop address of the routes. Enter the IP address to be used as the next hop. ▪ **Self address** - If **On**, under Self address, click **On** to use the self address as the next-hop address. ▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Local-preference** - Optionally, enter the value from 0 to 4294967295 to set the value to the received routes. If the iBGP speaker receives multiple routes to the same destination, then the route with the highest value is preferred. ▪ **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router. ▪ **Community** - In addition to the keywords below, you can also configure numbers in the range from 1 to 65535 and numbers in AA:NN format where the range for AA and NN is 1 to 65535. – internet – local-AS – no-export – no-advertise ▪ **Additive** - The specified community string is added to the route's community string.

Use case	Match criteria	Set criteria
Default route origination in BGP for a neighbor	▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router.	▪ **AS path** - Click **On** to set the AS path for the route. Specify the AS string as space separated list from 1 to 4294967295. ▪ **Tag** - If **On**, then the value is prepended with the AS path of the BGP route. ▪ **IP next hop** - If **On**, then updates the IP next-hop address of the routes. Enter the IP address to be used as the next hop. ▪ **Self address** - If **On**, under Self address, click **On** to use the self address as the next-hop address. ▪ **Metric** - Optionally, enter a value from 0 to 4294967295. When a metric value in a route matches this value, the route qualifies for distribution by the router. ▪ **Tag** - Optionally, enter a value from 0 to 4294967295. When a tag in a route matches this value, the route qualifies for distribution by the router. ▪ **Community** - Specify the community list to be set for this route: – internet – local-AS – no-export – no-advertise ▪ **Additive** - The specified community string is added to the route's community string.

8. Click **Submit**.

7

Defining VLAN Trunk Ports on SteelHead SD

This topic describes how to configure VLAN trunk ports for multiple zones on SteelHead SD. It includes these sections:

These procedures describe how to configure multizone VLAN trunk ports on SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch. For details on VLANs, see the *SteelConnect Manager User Guide*.

Introducing multizone VLAN trunk mode on LAN ports

Multiple VLANs are very common in Layer 2 (L2) network environments on the LAN side. With this feature, you can configure multiple VLANs on the same LAN port (that is, trunk port functionality). VLANs are used for segmenting networks at L2 and provide basic security for network traffic by limiting broadcast domains and network flooding.

SteelHead SD 2.0 supports trunk mode VLANs for zones. You can define a trunk port on a zone and configure it with different VLANs. For example, you can have a trunk port on the LAN side with two zones configured with different VLANs.

Figure 7-1. Multiple zones with different VLANs

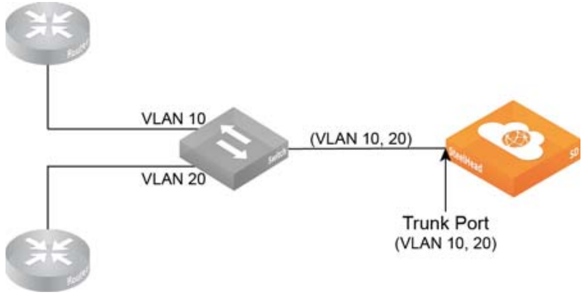

Defining trunk mode on ports

To define trunk mode ports

1. Choose Network Design > Zones to create a zone for the SteelHead SD. For details, see the *SteelConnect Manager User Guide*.

2. Click the VLAN tab.

Figure 7-2. Creating a VLAN trunk

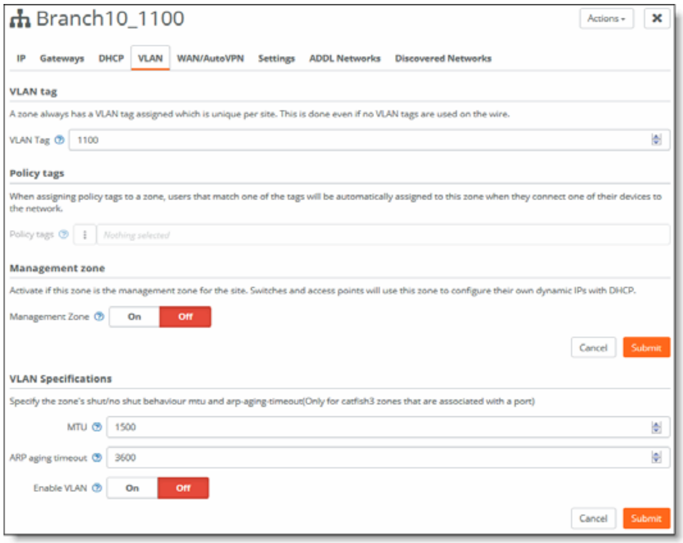

VLAN tag

A zone always has a VLAN tag assigned which is unique per site. This is done even if no VLAN tags are used on the wire.

VLAN Tag: 1100

Policy tags

When assigning policy tags to a zone, users that match one of the tags will be automatically assigned to this zone when they connect one of their devices to the network.

Policy tags: Nothing selected

Management zone

Activate if this zone is the management zone for the site. Switches and access points will use this zone to configure their own dynamic IPs with DHCP.

Management Zone: On [✓ Off]

VLAN Specifications

Specify the zone's shut/no shut behaviour mtu and arp-aging-timeout(Only for catfish3 zones that are associated with a port)

MTU: 1500

ARP aging timeout: 3600

Enable VLAN: On [✓ Off]

3. Specify a VLAN tag, if necessary. Every zone has a VLAN tag assigned. If you leave this field empty, the system picks a free VLAN ID from the pool.

4. Choose Appliances > Ports to configure the trunk port.

5. Select the site and appliance from the drop-down list.

6. Click the port for which you want to create the VLAN trunk. For example, LAN0_0.

Figure 7-3. Creating a LAN trunk port

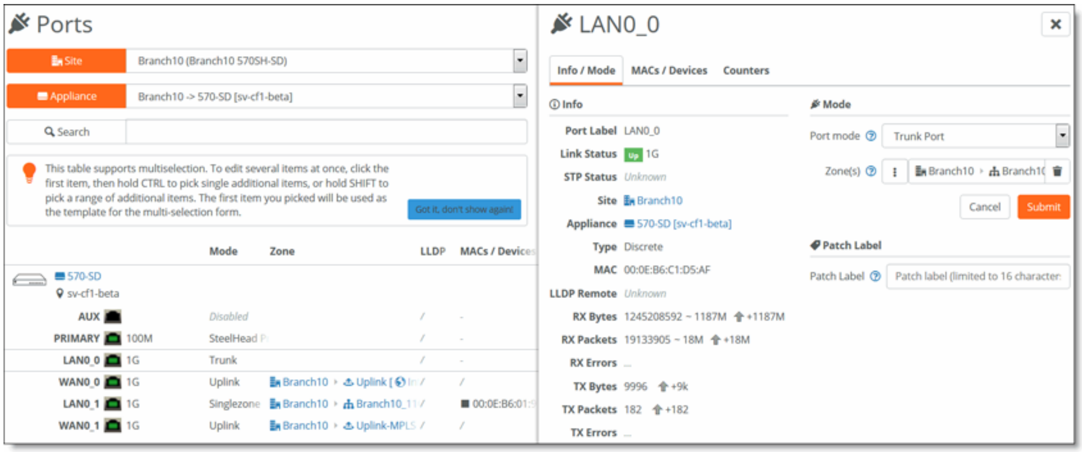

Site: Branch10 (Branch10 570SH-SD)

Appliance: Branch10 -> 570-SD [sv-cf1-beta]

This table supports multiselection. To edit several items at once, click the first item, then hold CTRL to pick single additional items, or hold SHIFT to pick a range of additional items. The first item you picked will be used as the template for the multi-selection form.

570-SD
sv-cf1-beta

		Mode	Zone	LLDP	MACs / Devices
AUX		Disabled		/	-
PRIMARY	100M	SteelHead P		/	-
LAN0_0	1G	Trunk		/	-
WAN0_0	1G	Uplink	Branch10 › Uplink [In	/	/
LAN0_1	1G	Singlezone	Branch10 › Branch10_11	/	00:0E:B6:01:
WAN0_1	1G	Uplink	Branch10 › Uplink-MPLS	/	/

Info

Port Label: LAN0_0
Link Status: Up 1G
STP Status: Unknown
Site: Branch10
Appliance: 570-SD [sv-cf1-beta]
Type: Discrete
MAC: 00:0E:B6:C1:D5:AF
LLDP Remote: Unknown
RX Bytes: 1245208592 ~ 1187M +1187M
RX Packets: 19133905 ~ 18M +18M
RX Errors: ...
TX Bytes: 9996 +9k
TX Packets: 182 +182
TX Errors: ...

Mode

Port mode: Trunk Port
Zone(s): Branch10 › Branch1(

Patch Label

Patch Label: Patch label (limited to 16 character:

7. Under Port mode, select Trunk Port from the drop-down list.

Important: For 2030 appliances, if the port has already been set to either Singlezone or Trunk Port, you must first disable the port before making a change to the Port mode. For example, if the port is already set to Singlezone, you must first disable the port, then set the port to Trunk Port.

8. Click **Submit.**

9. Navigate back to Appliances > Zones to enable VLAN on the configured port.

10. Under Management Zones, click **On** and **Submit** to activate multizone (VLAN trunk) connectivity for this zone.

11. Under VLAN Specifications: Enable VLAN, click **On** and **Submit** to enable VLAN on the trunk port you have configured.

12. Optionally, you can define:

- **Management Zone** - Click **On** if this zone is the management zone for the site. Switches and access points will use this zone to configure their own dynamic IP addresses with DHCP.
- **MTU** - The MTU is the largest physical packet size, measured in bytes, that a network can send. The default value is 1500.
- **ARP aging timeout** - Sets how long, in seconds an ARP entry stays in the cache before the cache refreshes. The default value is 1500.

Figure 7-4. VLAN enabled on the trunk port for the zone

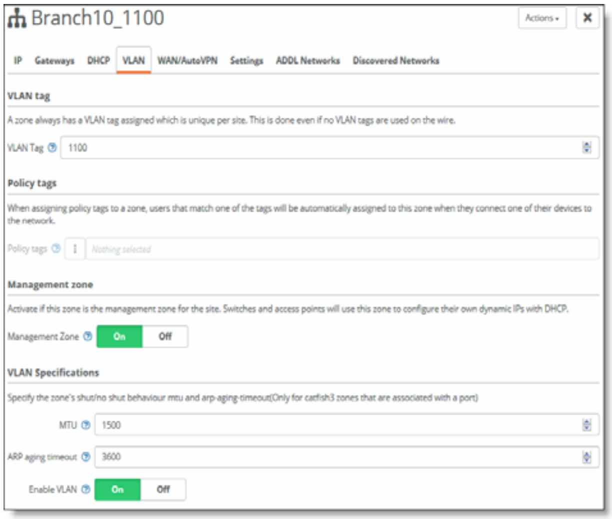

Branch10_1100

VLAN tag

A zone always has a VLAN tag assigned which is unique per site. This is done even if no VLAN tags are used on the wire.

VLAN Tag: 1100

Policy tags

When assigning policy tags to a zone, users that match one of the tags will be automatically assigned to this zone when they connect one of their devices to the network.

Policy tags: Nothing selected

Management zone

Activate if this zone is the management zone for the site. Switches and access points will use this zone to configure their own dynamic IPs with DHCP.

Management Zone: [✓ On] Off

VLAN Specifications

Specify the zone's shut/no shut behaviour mtu and arp-aging-timeout(Only for catfish3 zones that are associated with a port)

MTU: 1500

ARP aging timeout: 3600

Enable VLAN: [✓ On] Off

13. Click **Submit.**

14. Repeat Step 1 through Step 13 to create additional trunk ports with different VLANs.

8

Configuring High Availability on SteelHead SD

This topic describes how to configure high availability (HA) on SteelHead SD 2.0. It includes these sections:

- "Overview" on page 75
- "Prerequisites" on page 79
- "Configuring a SteelHead SD HA pair" on page 80
- "Monitoring a high-availability pair" on page 83
- "Troubleshooting" on page 85

Note: Previous versions of SteelHead SD supported an active-passive HA scheme. Since SteelHead SD 2.0 supports active-active HA, you can't upgrade your SteelHead SD 1.0 HA seamlessly to SteelHead SD 2.0 HA. You must first manually unpair your master and backup appliances in SCM, upgrade to SteelConnect 2.11, and reconfigure HA in SCM.

Overview

SteelHead SD 2.0 provides active-active HA for 570-SD, 770-SD, 3070-SD appliances.

Note: SteelConnect 2.11 provides active-active HA for SteelConnect SDI-2030 appliances located at the data center.

With active-active HA support, when a fault is detected, traffic is immediately routed to the peer appliance so that both appliances function in tandem. Traffic can be sent over any uplink regardless of the role assigned to the SteelHead SD appliance (that is, master or backup appliance). Active-active HA simplifies the configuration of uplinks for the HA pair of appliances.

Figure 8-1 shows an example of a symmetric deployment where the SteelHead SD HA pair are both connected to WAN 1 and WAN 2 via four uplinks.

Figure 8-1. Active-active HA deployment at the branch

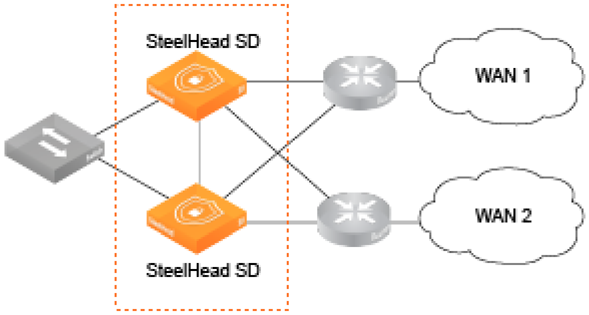

SteelHead SD also supports asymmetric HA deployments.

Figure 8-2. Asymmetric HA deployment

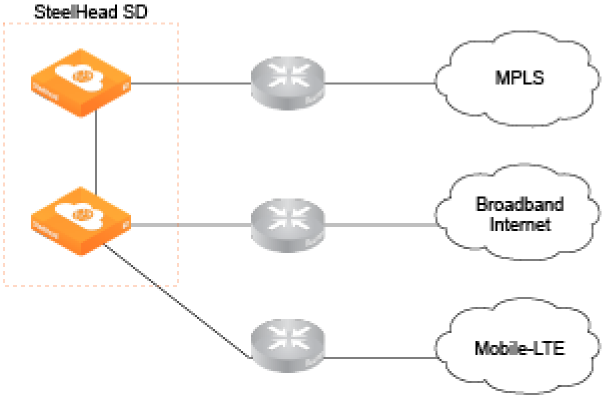

SteelHead SD version 2.0 includes these features:

- Support for OSPF and BGP where SteelHead SD can peer with a router.
- Support for symmetric and asymmetric connectivity.
- Support for Layer 2 (L2) and Layer 3 (L3) LAN topologies.
- Dedicated HA link for the SteelHead SD HA pair so that the peer appliances operate as a single logical unit.
- Autoconfiguration of the HA partner for bootstrapping when SCM connectivity with a peer is not accessible.
- Integration with SCM Health Check for advanced visibility and troubleshooting.
- Zscaler support for HA deployments.

Symmetric and asymmetric uplink connectivity

SteelHead SD version 2.0 provides symmetric and asymmetric uplink connectivity:

- **Symmetric** - In symmetric mode, each peer appliance is connected to all uplinks so that they essentially act as a single appliance. For example, you can have the 2 WAN uplinks connected to the peer appliances with four uplinks. Each uplink operates as a separate tunnel with separate IP addresses assigned to each uplink. If there is an uplink failure, the tunnel on that uplink goes down and the traffic is moved to the backup appliance. The 3070-SD supports up to 6 uplinks, where you can have 1 internet and 2 MPLS WAN uplinks for a total of 6 uplinks.
- **Asymmetric** - In asymmetric mode, different WANs are connected to the peer appliances. If there is an appliance failure or a LAN-side fail over, the master appliance becomes to peer appliance.

Figure 8-3. Symmetric and asymmetric HA deployment examples at the branch

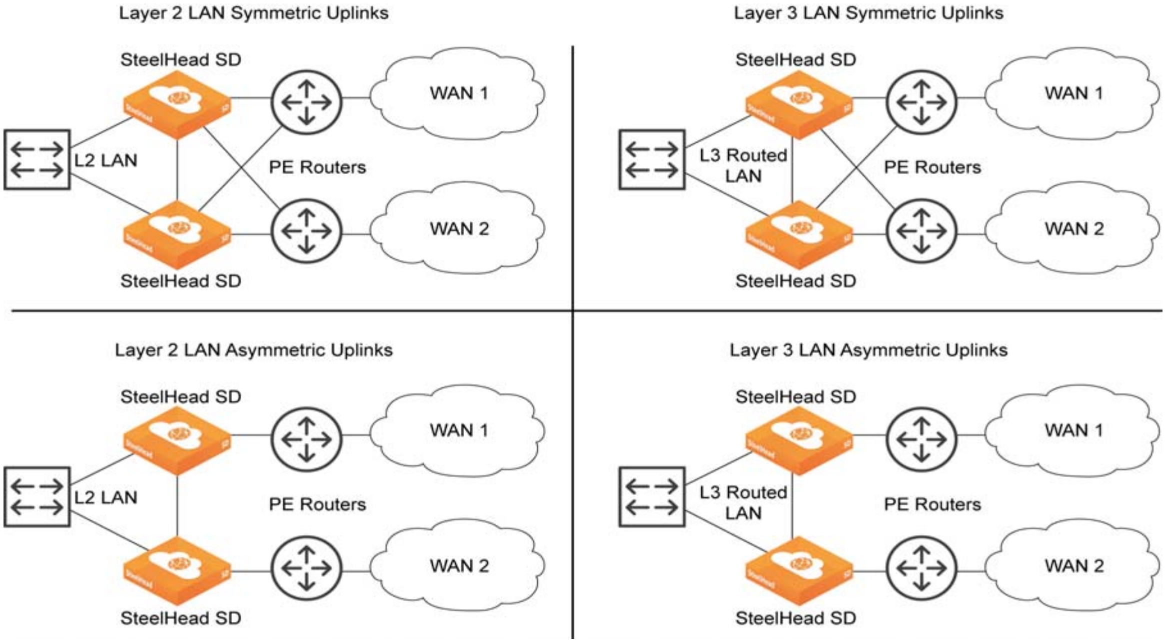

Layer 2 and Layer 3 support at the branch

With SteelHead SD 2.0, you can configure BGP and OSPF on the LAN branch.

You can configure iBGP between SteelHead SD HA peers if you want your overlay network to be advertised between the two appliances so that their routing tables are kept in synchronization. Also, you can have a combination of L2 and L3 zones so that if you have more than one LAN port configured, they can be a mix of L2 and L3. SteelHead SD uses iBGP between the peers to redistribute the overlay and connected routes.

LAN connectivity can be through either Layer 2 (L2) switch domain or Layer 3 (L3). In the case of a L3 LAN, connectivity is established through dynamic routing. SteelHead SD 2.0, supports:

- **L3 LAN** - You can redistribute static, connected, overlay, and WAN routes on both appliances in the HA pair. Your client traffic can go to either appliance in the HA pair. Using route convergence, the master processes the traffic and sends it on the overlay network.

- **L2 LAN** - With L2, you can have a switch on the LAN-side connected to SteelHead SDs that have the same LAN zone with different IP address for each appliance. The system assigns a single virtual IP address (VIP) on the zone that is owned by the master appliance. All traffic goes to the master appliance where it sends it on the overlay network. If there is a failure, the VIP moves to the backup appliance where it becomes the new master.

Multigroup VIP and Virtual Router Redundancy Protocol (VRRP) with a third-party router are not supported at this time.

Failure conditions

SteelHead SD supports appliance, uplink, LAN, and dedicated port failure conditions. The following examples illustrate some typical use cases.

Appliance failure

For failures due to power, hardware, or VM failures, the master role is moved to the peer appliance. The VIP is moved to the new master appliance and L3 advertisements are stopped from the previous master appliance.

Figure 8-4. Appliance failure where VIP is injected into new master appliance

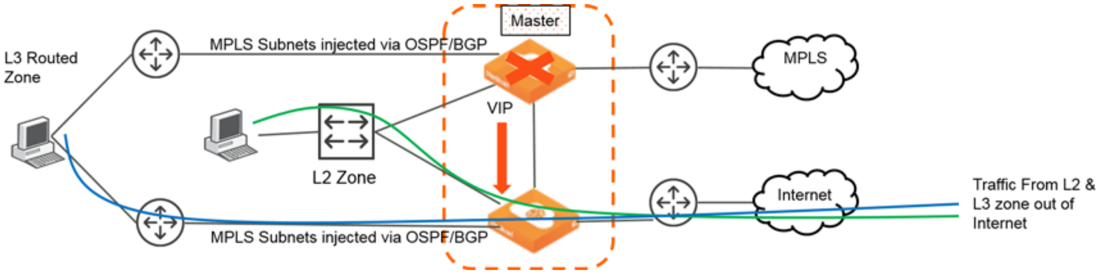
L3 Routed Zone
MPLS Subnets injected via OSPF/BGP
Master
MPLS
VIP
L2 Zone
Internet
Traffic From L2 & L3 zone out of Internet
MPLS Subnets injected via OSPF/BGP

LAN failure

For an L2 LAN failure, the VIP moves to the backup appliance and MPLS connectivity is withdrawn. Traffic is sent through the backup appliance.

For an L3 LAN failure, routing converges to send traffic to backup appliance. Traffic is moved between appliances through the AUX port depending on which uplink the traffic needs to exit the HA pair.

Figure 8-5. L2 link failure where VIP moves to backup but no switchover of the master appliance

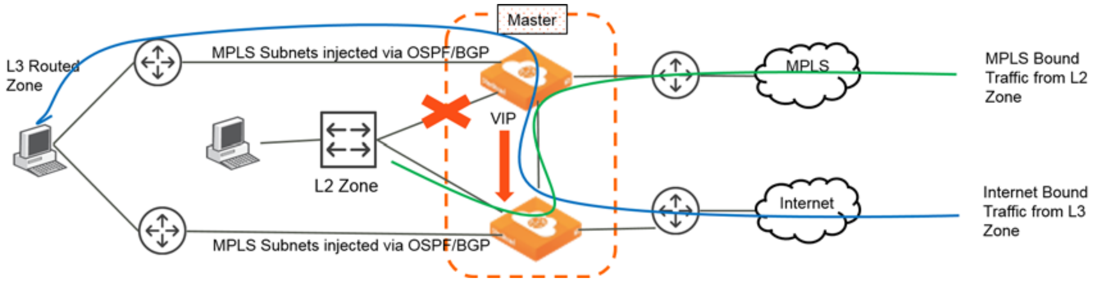
Master
L3 Routed Zone
MPLS Subnets injected via OSPF/BGP
VIP
L2 Zone
MPLS Subnets injected via OSPF/BGP
MPLS
Internet
MPLS Bound Traffic from L2 Zone
Internet Bound Traffic from L3 Zone

Figure 8-6. L3 failure where traffic continues to travel through the master appliance

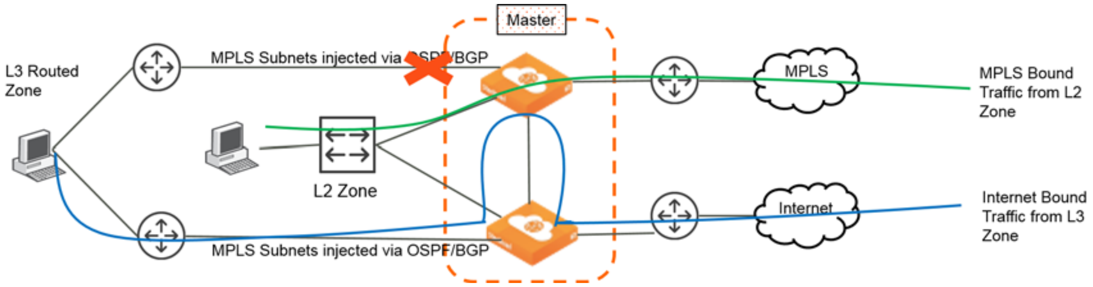
Master
L3 Routed Zone
MPLS Subnets injected via OSPF/BGP
L2 Zone
MPLS Subnets injected via OSPF/BGP
MPLS
Internet
MPLS Bound Traffic from L2 Zone
Internet Bound Traffic from L3 Zone

AUX port failure

If the HA connection fails between the master and backup appliances, each appliance becomes an independent appliance. The VIP continues to be hosted by the master appliance. The MPLS networks are withdrawn on the original master appliance. The backup appliance becomes an independent master appliance and all MPLS traffic travels through it.

Figure 8-7. AUX port failure where each appliance becomes an independent master

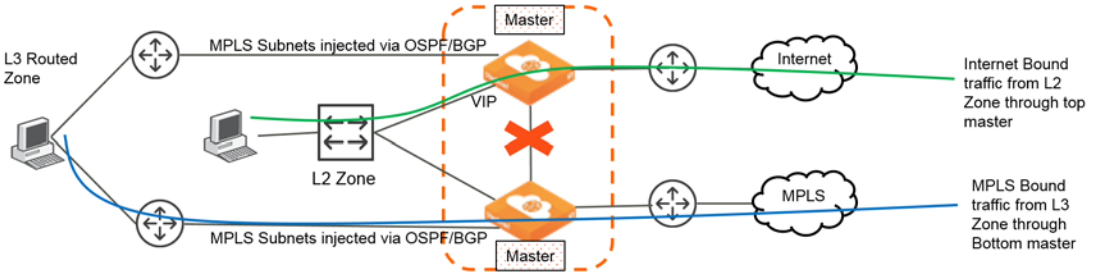
Master
L3 Routed Zone
MPLS Subnets injected via OSPF/BGP
VIP
L2 Zone
MPLS Subnets injected via OSPF/BGP
Master
Internet
MPLS
Internet Bound traffic from L2 Zone through top master
MPLS Bound traffic from L3 Zone through Bottom master

Prerequisites

Before configuring high availability, check these requirements and recommendations. Both appliances must be:

- registered with SCM.
- running the same software version.
- cabled directly on the LAN branch using the AUX port.
- peer appliances must be located in the same zone of the branch network.

Configuring a SteelHead SD HA pair

These steps assume that you have installed, registered, and performed the initial configuration of the SteelHead SD HA pair. You should create your branch site where the HA pair will be located, along with the associated zone and uplinks. For details, see the *SteelConnect Manager User Guide*. This section contains these topics:

- "Configuring the AUX port on the HA pair" on page 80
- "Configuring the LAN zone for the SteelHead SD HA pair" on page 81
- "Configuring the appliances into an HA pair" on page 82

Configuring the AUX port on the HA pair

The first step is to configure the AUX port on the SteelHead SD HA pair. You will select the HA or Cluster mode for the port.

To configure the AUX port on the master and backup SteelHead SD appliances

1. On the first appliance in the pair, choose Appliances > Ports and select the site from the Site drop-down list.

2. Under Appliances, select the appliance. The ports for the appliance are displayed.

3. Select the AUX port to expand the page.

4. Under Mode, select HA from the Port mode drop-down menu.

Figure 8-8. Configuring the AUX port on the HA pair

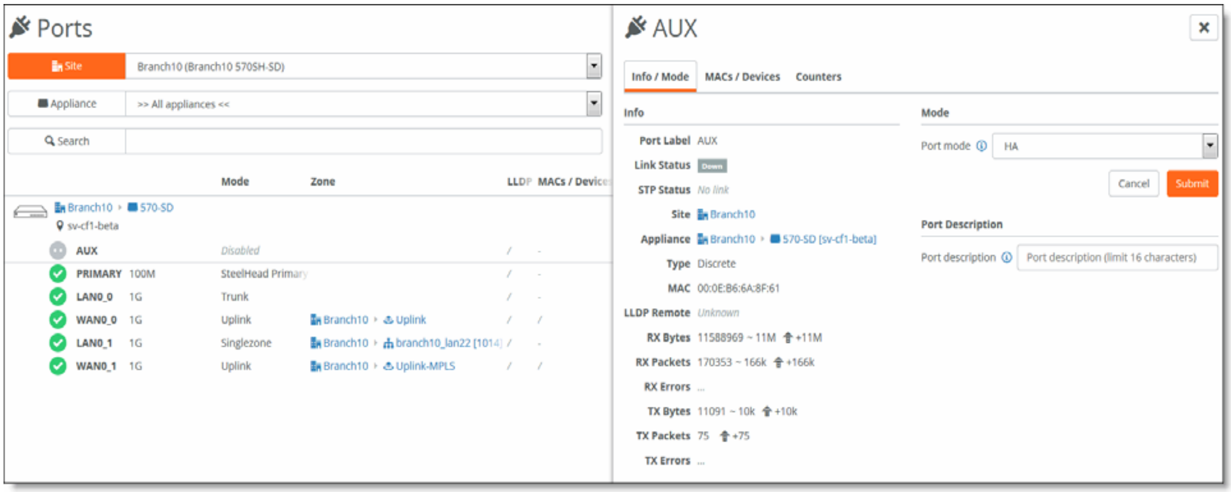

Ports

Site: Branch10 (Branch10 570SH-SD)

Appliance: >> All appliances <<

Search:

Branch10 › 570-SD
sv-cf1-beta

		Mode	Zone	LLDP	MACs / Devices
AUX		Disabled		/	-
PRIMARY	100M	SteelHead Primary		/	-
LAN0_0	1G	Trunk		/	-
WAN0_0	1G	Uplink	Branch10 › Uplink	/	/
LAN0_1	1G	Singlezone	Branch10 › branch10_lan22 [1014]	/	-
WAN0_1	1G	Uplink	Branch10 › Uplink-MPLS	/	/

AUX

Info

Port Label	AUX
Link Status	Down
STP Status	No link
Site	Branch10
Appliance	Branch10 › 570-SD [sv-cf1-beta]
Type	Discrete
MAC	00:0E:B6:6A:8F:61
LLDP Remote	Unknown
RX Bytes	11588969 ~ 11M +11M
RX Packets	170353 ~ 166k +166k
RX Errors	...
TX Bytes	11091 ~ 10k +10k
TX Packets	75 +75
TX Errors	...

Mode

Port mode: HA

Port Description

Port description: Port description (limit 16 characters)

5. Click **Submit**.

6. Repeat Step 1 through Step 5 for the peer appliance in the HA pair.

Configuring the LAN zone for the SteelHead SD HA pair

The next step is to configure the LAN zone for the SteelHead SDHA pair. If it is a Layer 2 or Layer 3 zone, you configure the correct gateway.

To configure the LAN zone

1. Choose Network Design > Zones.

2. Select the Zone for the appliance to expand the page.

3. Under IPv4 network and gateway, specify the Layer 2 or Layer 3 gateway IP address.

Figure 8-9. Configuring the LAN zone gateway

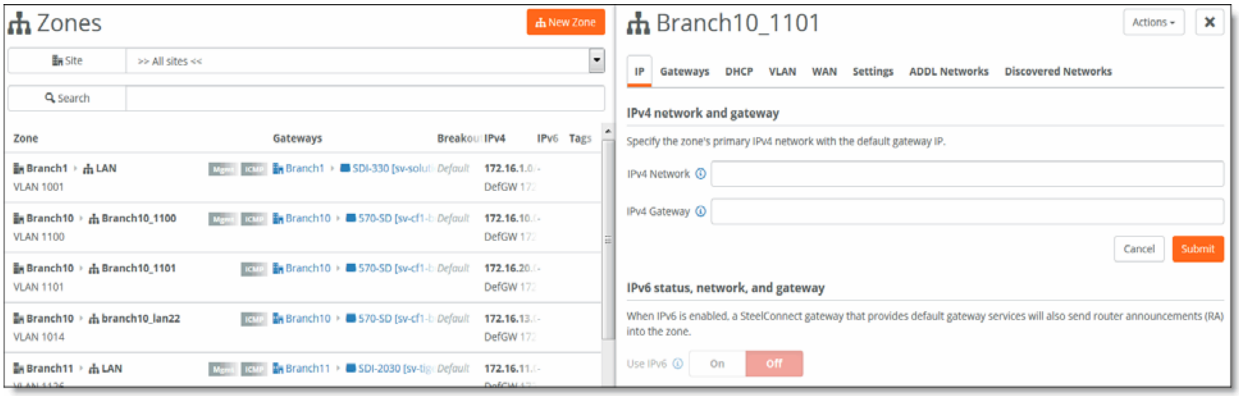

Zone		Gateways	Breakou	IPv4	IPv6	Tags
Branch1 › LAN VLAN 1001	Mgmt ICMP	Branch1 › SDI-330 [sv-soluti	Default	172.16.1.0/ DefGW 172	-	
Branch10 › Branch10_1100 VLAN 1100	Mgmt ICMP	Branch10 › 570-SD [sv-cf1-b	Default	172.16.10.(DefGW 172	-	
Branch10 › Branch10_1101 VLAN 1101	ICMP	Branch10 › 570-SD [sv-cf1-b	Default	172.16.20.(DefGW 172	-	
Branch10 › branch10_lan22 VLAN 1014	ICMP	Branch10 › 570-SD [sv-cf1-b	Default	172.16.13.(DefGW 172	-	
Branch11 › LAN	Mgmt ICMP	Branch11 › SDI-2030 [sv-tig	Default	172.16.11.(	-	

IPv4 network and gateway

Specify the zone's primary IPv4 network with the default gateway IP.

IPv4 Network

IPv4 Gateway

IPv6 status, network, and gateway

When IPv6 is enabled, a SteelConnect gateway that provides default gateway services will also send router announcements (RA) into the zone.

Use IPv6: Off

4. Click **Submit.**

5. Depending on your topology, repeat Step 1 through Step 4 for each zone in the HA branch.

Configuring the LAN zone for the SteelHead SD HA pair

After you configure the LAN zones, you must assign the LAN ports to the zones:

- If the LAN-side network is L2, the zone must to be attached to the LAN port on both appliances.
- If the LAN-side network is L3, the correct zone must be attached to the LAN ports for each of the appliances in the HA pair.

To assign the LAN port to the zone

1. To assign the appliance port to the zone, choose Appliances > Ports.

2. Select the site from the Site list.

3. Select the LAN port to expand the pane.

Figure 8-10. Configuring the LAN port

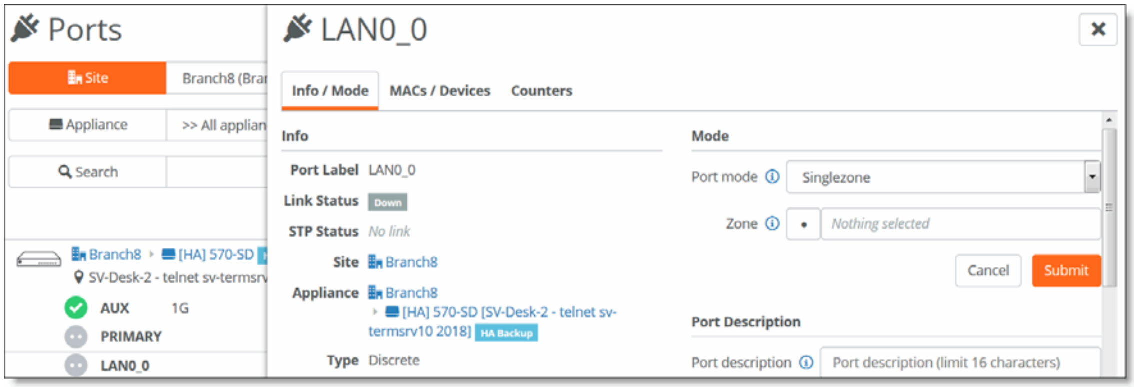

Info

Port Label LAN0_0
Link Status Down
STP Status No link
Site Branch8
Appliance Branch8 › [HA] 570-SD [SV-Desk-2 - telnet sv-termsrv10 2018] HA Backup
Type Discrete

Mode

Port mode Singlezone
Zone Nothing selected

Port Description

Port description Port description (limit 16 characters)

4. Under Mode, select Singlezone or Multizone. If you select Singlezone, select the zone from the drop-down list.

5. Click **Submit**.

6. Depending on your topology, repeat Step 1 through Step 5 for each appliance port that needs to be assigned to a zone.

Configuring the appliances into an HA pair

To configure the appliances into an HA pair

1. Choose Appliances and select the appliance.

2. Select the HA tab.

3. Under High availability settings, select the appliance that is in the branch.

Figure 8-11. Selecting the partner appliance in the branch

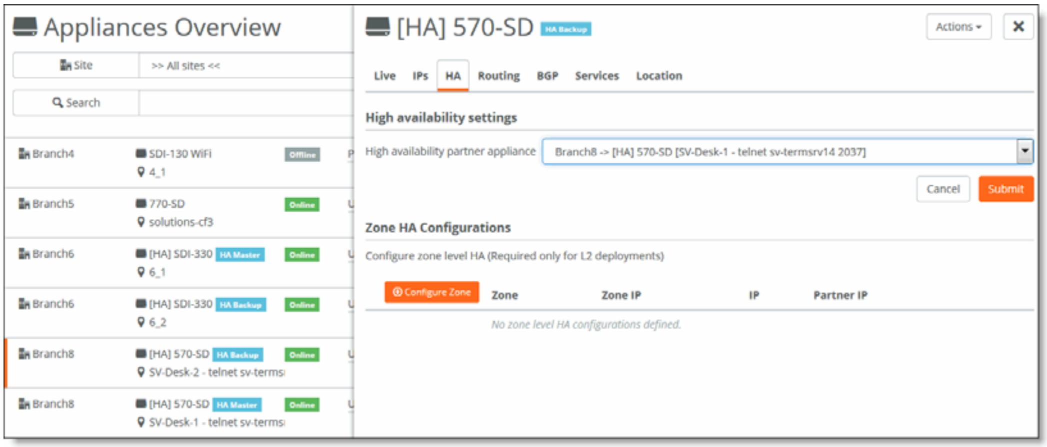

Branch4	SDI-130 WiFi 4_1	Offline
Branch5	770-SD solutions-cf3	Online
Branch6	[HA] SDI-330 HA Master 6_1	Online
Branch6	[HA] SDI-330 HA Backup 6_2	Online
Branch8	[HA] 570-SD HA Backup SV-Desk-2 - telnet sv-terms	Online
Branch8	[HA] 570-SD HA Master SV-Desk-1 - telnet sv-terms	Online

High availability settings

High availability partner appliance Branch8 -> [HA] 570-SD [SV-Desk-1 - telnet sv-termsrv14 2037]

Zone HA Configurations

Configure zone level HA (Required only for L2 deployments)

Zone	Zone IP	IP	Partner IP

No zone level HA configurations defined.

4. Click **Submit**.

Once the two appliances are paired, you can see them negotiate their roles in the Appliances Overview page. The master and backup roles are assigned and appear for the paired appliances.

5. If you have a Layer 2 zone in your network, click **Configure Zone** to configure the LAN interface IP addresses.

Figure 8-12. Configuring the LAN interfaces for Layer 2 zones

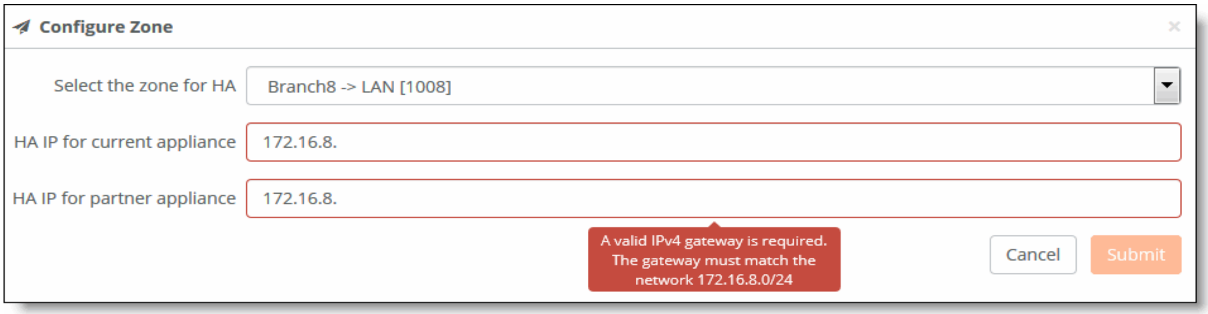

Configure Zone

Select the zone for HA: Branch8 -> LAN [1008]

HA IP for current appliance: 172.16.8.

HA IP for partner appliance: 172.16.8.

A valid IPv4 gateway is required. The gateway must match the network 172.16.8.0/24

6. Select the zone for the HA pair.

7. Enter the HA IP address for the current appliance.

8. Enter the HA IP address for the partner appliance.

9. Click **Submit**.

Monitoring a high-availability pair

SCM displays all appliances belonging to a high-availability pair with a blue HA icon in all views. After the appliance reports its HA state to SCM, the icon indicates whether it is the master or the backup.

When an HA appliance pair lose connectivity, Appliances and Health Check display both the master and backup appliance as HA Master. For SteelHead SD appliances, SCM will not display Offline for an appliance unless the appliance actually goes offline.

Note: Uplink tracking and LAN port tracking is not available on SteelHead SD.

SCM manages both appliances in a pair as one. For example, if you view the ports for an HA pair, they appear together.

Figure 8-13. HA pair ports

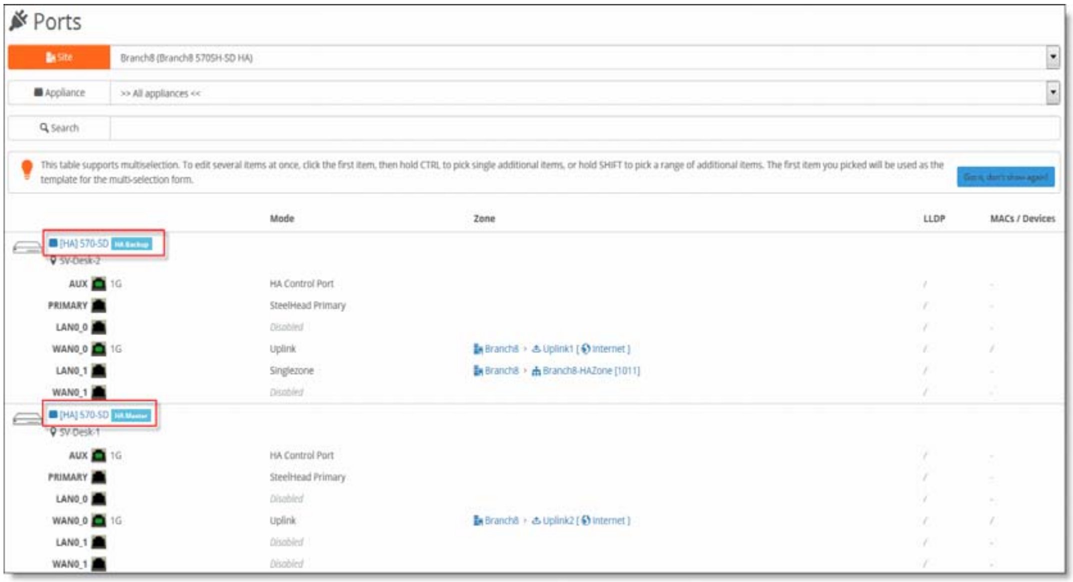

Ports

Site: Branch8 (Branch8 570SH-SD HA)

Appliance: >> All appliances <<

Search

This table supports multiselection. To edit several items at once, click the first item, then hold CTRL to pick single additional items, or hold SHIFT to pick a range of additional items. The first item you picked will be used as the template for the multi-selection form.

	Mode	Zone	LLDP	MACs / Devices
[HA] 570-SD HA Backup SV-Desk-2				
AUX 1G	HA Control Port		/	-
PRIMARY	SteelHead Primary		/	-
LAN0_0	Disabled		/	-
WAN0_0 1G	Uplink	Branch8 › Uplink1 [Internet]	/	/
LAN0_1	Singlezone	Branch8 › Branch8-HAZone [1011]	/	-
WAN0_1	Disabled		/	-
[HA] 570-SD HA Master SV-Desk-1				
AUX 1G	HA Control Port		/	-
PRIMARY	SteelHead Primary		/	-
LAN0_0	Disabled		/	-
WAN0_0 1G	Uplink	Branch8 › Uplink2 [Internet]	/	/
LAN0_1	Disabled		/	-
WAN0_1	Disabled		/	-

To view appliance health of an HA pair

1. Choose Health Check > Appliance Health.

Figure 8-14. Appliance health in an HA pair

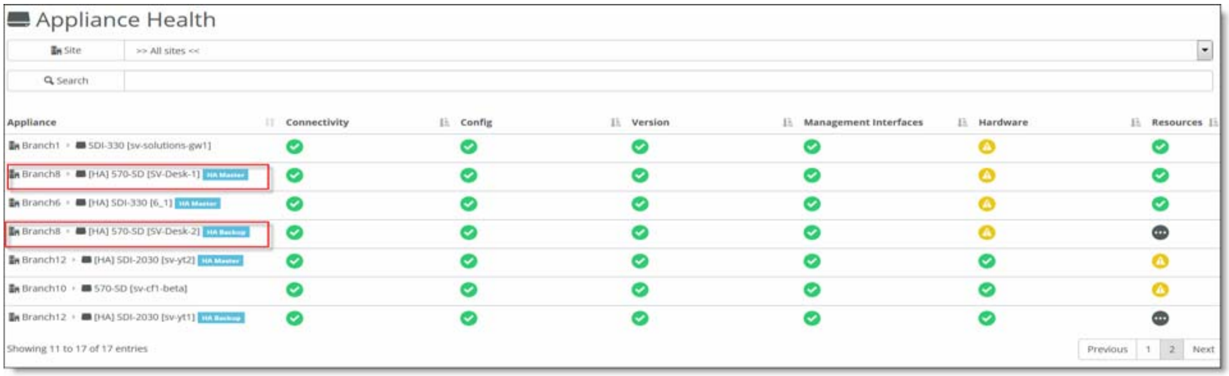

Appliance Health

Site: >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resources
Branch1 › SDI-330 [sv-solutions-gw1]	✓	✓	✓	✓	⚠	✓
Branch8 › [HA] 570-SD [SV-Desk-1] HA Master	✓	✓	✓	✓	⚠	✓
Branch6 › [HA] SDI-330 [6_1] HA Master	✓	✓	✓	✓	⚠	✓
Branch8 › [HA] 570-SD [SV-Desk-2] HA Backup	✓	✓	✓	✓	⚠	⋯
Branch12 › [HA] SDI-2030 [sv-yt2] HA Master	✓	✓	✓	✓	✓	⚠
Branch10 › 570-SD [sv-cf1-beta]	✓	✓	✓	✓	✓	⚠
Branch12 › [HA] SDI-2030 [sv-yt1] HA Backup	✓	✓	✓	✓	✓	⋯

Showing 11 to 17 of 17 entries

2. Select the appliance to expand the page.

Figure 8-15. Viewing HA pair health details

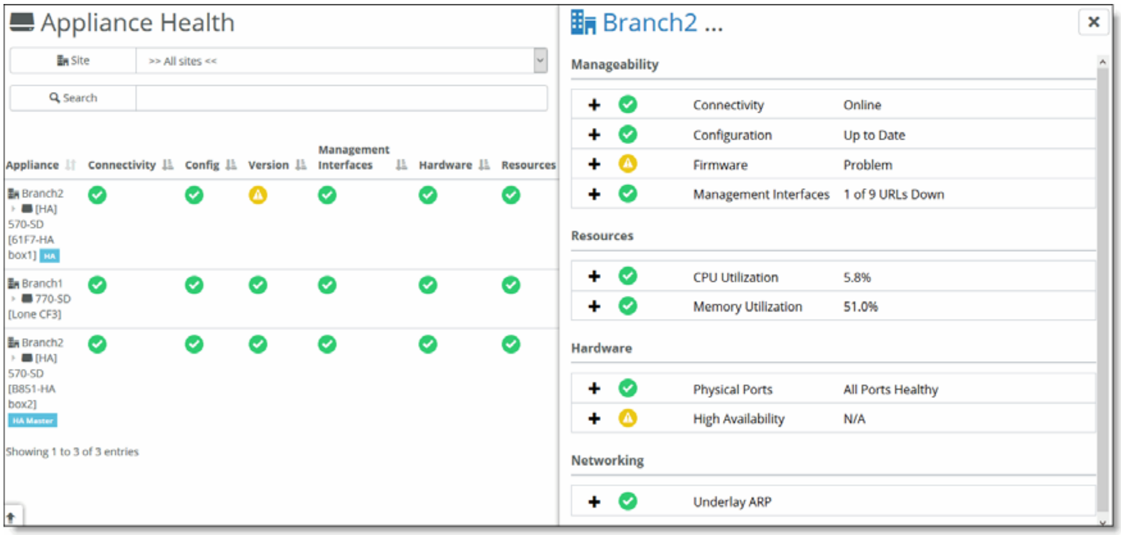

Appliance Health

Site: >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resources
Branch2 [HA] 570-SD [61F7-HA box1] HA	✓	✓	⚠	✓	✓	✓
Branch1 770-SD [Lone CF3]	✓	✓	✓	✓	✓	✓
Branch2 [HA] 570-SD [B851-HA box2] HA Master	✓	✓	✓	✓	✓	✓

Showing 1 to 3 of 3 entries

Branch2 ...

Manageability

+	✓	Connectivity	Online
+	✓	Configuration	Up to Date
+	⚠	Firmware	Problem
+	✓	Management Interfaces	1 of 9 URLs Down

Resources

+	✓	CPU Utilization	5.8%
+	✓	Memory Utilization	51.0%

Hardware

+	✓	Physical Ports	All Ports Healthy
+	⚠	High Availability	N/A

Networking

+	✓	Underlay ARP	

3. Click the plus sign (+) to expand the field. For example, under Hardware, click the plus sign to the left of High Availability to view the HA IP address and status for the selected appliance.

Troubleshooting

- Make sure the roles are displayed correctly on the appliances in the Appliances > Overview page.
- All the tunnels must be up and should be using the uplinks for both the HA appliances.
- If the appliance HA role is *Unknown* or if the appliance pair is listed as Master/Master, make sure the AUX port (that is, the dedicated HA port) is enabled and it is configured as HA mode. If the AUX port is configured and enabled, then collect a system dump from the appliances and contact Riverbed Support at https://support.riverbed.com.
- The HA role is established with a daemon named **keepalived**. Search the logs for “keepalived” to debug HA issues.
- Some useful CLI commands to analyze are:

```
get_keepalived.sh
show_ha_info
```

9

Configuring QoS Shaping on SteelHead SD

This topic describes how to configure QoS shaping on SteelHead SD 2.0. It includes these sections:

- "Introducing QoS shaping for SteelHead SD appliances" on page 87
- "If you set the QoS priority in a traffic rule" on page 88
- "Configuring QoS shaping on SteelHead SD" on page 89

These procedures describe QoS shaping for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

Introducing QoS shaping for SteelHead SD appliances

SteelHead SD 2.0 supports QoS shaping to allocate bandwidth and prioritize traffic. With SteelHead SD QoS there are no classes to configure: you simply set the bandwidth with a fixed value to provide QoS shaping on inbound and outbound traffic. QoS shaping is supported on site uplinks of a given WAN.

You configure QoS as a policy on a WAN uplink to ensure that your uplink doesn't exceed a set bandwidth. For example, if you know your uplink can't handle more than 100 Mbps, you set your limit at something lower than 100 Mbps to ensure that you have a minimal level of service for that uplink. If traffic exceeds the configured bandwidth, it is buffered and shaped. If traffic exceeds the buffer capacity, it is dropped. The QoS shaper throttles and limits traffic to your configured bandwidth.

QoS prioritizes traffic flowing on:

- **Outbound QoS** - Outbound QoS prioritizes and shapes outbound traffic on a WAN uplink. Egress traffic on a WAN uplink is provisioned with an outbound QoS policy.
- **Inbound QoS** - Inbound QoS prioritizes and shapes inbound traffic on a WAN uplink, ingress traffic on a WAN uplink is provisioned with an inbound QoS policy.

Note: QoS marking is currently supported on SteelConnect when you define a traffic rule to match on a specified site or application to apply an outbound DSCP mark. For additional information on QoS marking, see the *SteelConnect Manager User Guide*.

SteelHead SD QoS shaper

SteelHead SD QoS uses a class-based queuing implementation that assigns packets with a DSCP mark to one of the four dedicated service class queues and distributes bandwidth between them using a deficit weighted round robin (DWRR) scheme. Traffic scheduling and shaping uses a basic single rate and bucket algorithm to regulate the data transmission rate and drain the queues. QoS shaping:

- classifies traffic based on the DSCP mark and shapes it according to a fixed bandwidth allocation designated for each traffic class to ensure that their aggregate bandwidth doesn't exceed the configured rate.
- is per packet; it is not flow based.
- applies only on WAN interfaces, not LAN interfaces. It is not site aware.

PQ class ID	QoS class (bandwidth)	Example of traffic type	DSCP values
Urgent (3)	Latency Sensitive - 40%	VoIP	Class Selector CS7(56), CS6(48), EF(46), VA(44), CS5(40), CS4(32)
High (2)	Streaming Media - 30%	Video	AF4x(34,36,38), AF3x(26,28,30), CS3(24), AF2x(18,20,22), TOS4(4), CS2(16), TOS1(1)
Normal (1)	Best Effort - 20%	MAPI	CS0, AF1x, TOS2, and if the DSCP value is not specified
Low (0)	Background Traffic - 10%	YouTube	CS1

Note: The DSCP standards for QoS have been specified and respecified many times. For the latest standards, see DSCP RFC 2474, RFC 3168, RFC 3260, and RFC 5865.

If all queues have equal size packets, and all queues are full, they are not handled equally. Urgent receives 40% of the bandwidth, High receives 30%, Normal receives 20%, and Low receives 10%. This bandwidth distribution occurs when you have different traffic classes shared on the same WAN uplink and under congestion conditions they are competing for the same bandwidth rate. If all your traffic is in the Low traffic class, then it's not competing with the other queue classes, so it receives 100% of the bandwidth.

Traffic scheduling and shaping uses a basic single rate to regulate the data transmission rate and drain the queues. Each round turn traverses only traffic class queues that are in the active bitmap and always in the order from higher to lower priority queues. The token bucket shaper then throttles the rate at which packets are transmitted by determining when a packet selected by the DWRR can be sent.

To ensure that each class-based queue doesn't overflow, when the queue is filled to its maximum capacity, the newly arriving packets are dropped until the queue has enough room to accept incoming traffic.

To view the TOS/DSCP and QoS traffic classes table, see Appendix B, "TOS, DSCP, QoS Traffic Class Table."

If you set the QoS priority in a traffic rule

You can also configure custom DSCP marks in the SCM when you configure the QoS priority in traffic rules (Rules > Traffic Rules > New Traffic Rule). Traffic rules allow you to mark traffic to specific DSCP value. Only DSCP marked traffic is placed in the QoS priority queues. Traffic must have a DSCP marking value for QoS traffic to be classified. For details, see the *SteelConnect Manager User Guide*.

If the QoS priority is configured in a traffic rule, then:

- inbound QoS shaping is done before the QoS mark is applied.
- outbound QoS shaping is done after the QoS mark is applied.

Setting the QoS priority in the traffic rule marks the traffic with the configured DSCP value upon egress, which executes independently of QoS shaping.

For example, if the original DSCP mark on the traffic is NORMAL priority and matches the traffic rule with the QoS priority set to URGENT, then QoS shaping will be influenced as follows:

- Inbound QoS shaping queues and processes the traffic as NORMAL priority, before the traffic rule changes the DSCP mark.
- Outbound QoS shaping queues and processes the traffic as URGENT priority, after the traffic rule changes the DSCP mark.

Configuring QoS shaping on SteelHead SD

To configure QoS shaping

1. To configure inbound/outbound QoS on the branch WAN uplink, choose Network Design > Uplinks.

2. Select the QoS tab.

Figure 9-1. Configuring inbound and outbound QoS

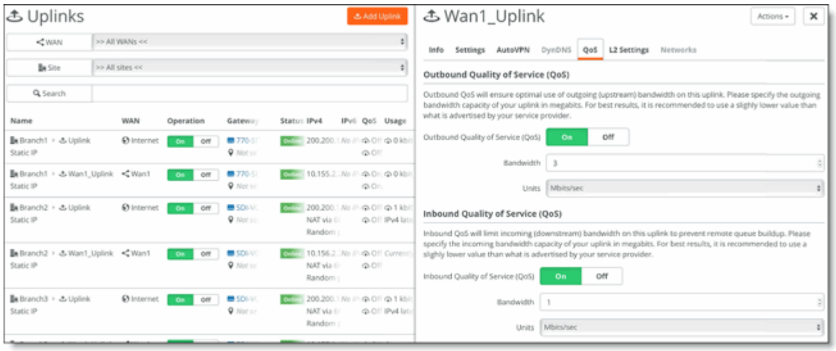

Name	WAN	Operation	Gateway	Status	IPv4	IPv6	QoS	Usage
Branch1 › Uplink Static IP	Internet	On Off	770-S Not s	Online	200.200.1	No IP	Off Off	0 kbit
Branch1 › Wan1_Uplink Static IP	Wan1	On Off	770-S Not s	Online	10.155.2.	No IP	On On	0 kbit
Branch2 › Uplink Static IP	Internet	On Off	SD-V Not s	Online	200.200.1 NAT via 6 Random	No IP	Off Off	1 kbit IPv4 lat
Branch2 › Wan1_Uplink Static IP	Wan1	On Off	SD-V Not s	Online	10.156.2. NAT via 6 Random	No IP	Off Off	Current
Branch3 › Uplink Static IP	Internet	On Off	SD-V Not s	Online	200.200.1 NAT via 6 Random	No IP	Off Off	1 kbit IPv4 lat

Outbound Quality of Service (QoS)

Outbound QoS will ensure optimal use of outgoing (upstream) bandwidth on this uplink. Please specify the outgoing bandwidth capacity of your uplink in megabits. For best results, it is recommended to use a slighly lower value than what is advertised by your service provider.

Outbound Quality of Service (QoS): ● On ○ Off

Bandwidth: 3

Units: Mbits/sec

Inbound Quality of Service (QoS)

Inbound QoS will limit incoming (downstream) bandwidth on this uplink to prevent remote queue buildup. Please specify the incoming bandwidth capacity of your uplink in megabits. For best results, it is recommended to use a slighly lower value than what is advertised by your service provider.

Inbound Quality of Service (QoS): ● On ○ Off

Bandwidth: 1

Units: Mbits/sec

3. Specify the bandwidth for QoS shaping and click **Submit**.

   Let's say your egress throughput traffic rate on the Wan1_Uplink is 5 Mbps and you want to constrain it to not exceed 3 Mbps. Set the Outbound QoS Bandwidth to 3Mbps to ensure that the aggregate throughput upon the egress of the Wan1_Uplink doesn't exceed your configured bandwidth.

   Different traffic classes (Urgent, High, Normal, Low) sent on the same WAN uplink will share the bandwidth at a ratio of (4:3:2:1) respectively. Their combined bandwidth with not exceed the configured rate.

   If you configure an inbound policy as well as an outbound policy, it is rate limited independent of the outbound policy.

4. To check if the policy has been applied, choose Appliances > Overview. The Config column will change from Pending to Shipped.

10

Health Check and Reporting on SteelHead SD

This topic describes the health-check and reporting features on SteelHead SD 2.0. It includes these sections:

These procedures describe health-check and reporting tools for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch. For details on health check and visibility in SCM, see the *SteelConnect Manager User Guide*.

Checking SteelHead SD connectivity to SCM

You can check SteelHead SD connectivity to SCM in the Health Check > Appliance Health page.

To view uplink connectivity for SteelHead SD appliances

1. Choose Health Check > Appliance Health.

2. Select the SteelHead SD appliance to expand the page.

Figure 10-1. Viewing appliance health

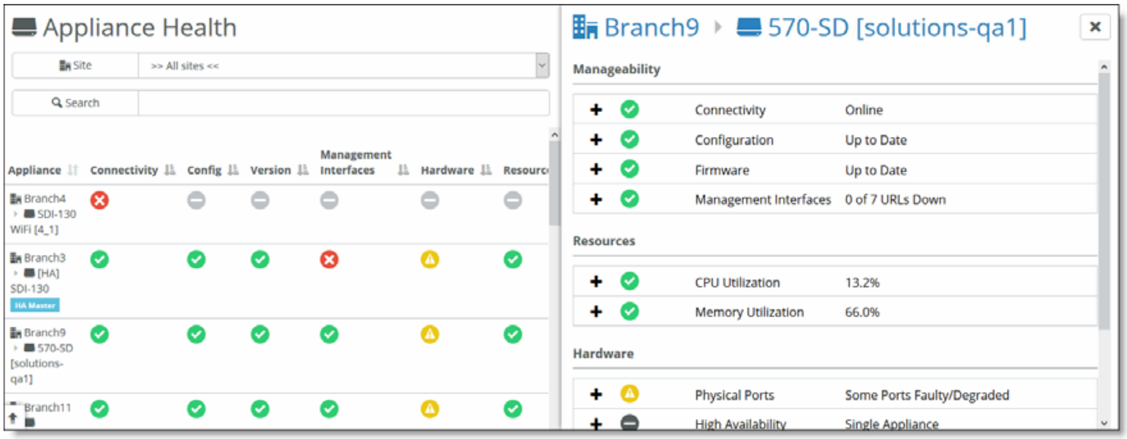

Appliance Health

Site: >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resourc
Branch4 SDI-130 WiFi [4_1]	✗	⊖	⊖	⊖	⊖	⊖
Branch3 [HA] SDI-130 HA Master	✓	✓	✓	✗	⚠	✓
Branch9 570-SD [solutions-qa1]	✓	✓	✓	✓	⚠	✓
Branch11	✓	✓	✓	✓	⚠	✓

Branch9 > 570-SD [solutions-qa1]

Manageability

+	✓	Connectivity	Online
+	✓	Configuration	Up to Date
+	✓	Firmware	Up to Date
+	✓	Management Interfaces	0 of 7 URLs Down

Resources

+	✓	CPU Utilization	13.2%
+	✓	Memory Utilization	66.0%

Hardware

+	⚠	Physical Ports	Some Ports Faulty/Degraded
+	⊖	High Availability	Single Appliance

3. Under Manageability: Connectivity, click the plus sign (+). The current status for the appliance is displayed.

Figure 10-2. Viewing if SteelHead SD is connected to SCM

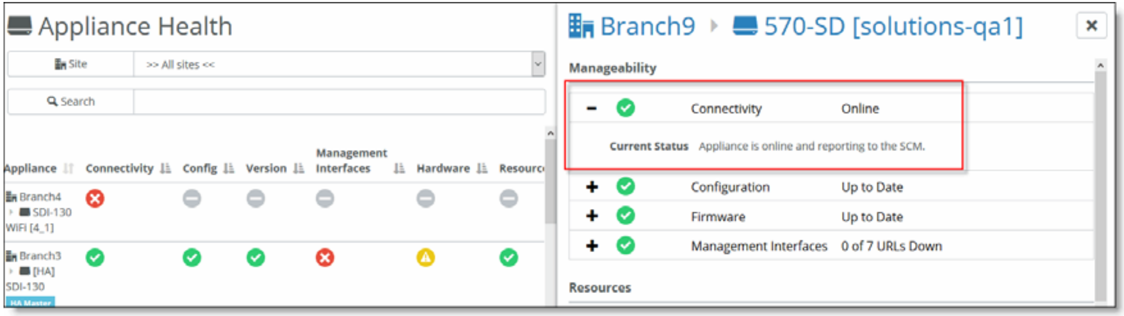

Appliance Health

Site: >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resourc
Branch4 SDI-130 WiFi [4_1]	✗	⊖	⊖	⊖	⊖	⊖
Branch3 [HA] SDI-130 HA Master	✓	✓	✓	✗	⚠	✓

Branch9 > 570-SD [solutions-qa1]

Manageability

–	✓	Connectivity	Online
		Current Status	Appliance is online and reporting to the SCM.
+	✓	Configuration	Up to Date
+	✓	Firmware	Up to Date
+	✓	Management Interfaces	0 of 7 URLs Down

Resources

Viewing the SteelHead SD HA IP address

You can view the SteelHead SD high availability (HA) IP address and current status of the appliance in the Health Check > Appliance Health page.

To view the SteelHead SD HA IP address and status

1. Choose Health Check > Appliance Health.

2. Select the SteelHead SD appliance to expand the page.

3. Under Manageability: Hardware, click the plus sign (+). The current HA IP address and status for the appliance is displayed.

Figure 10-3. Viewing the SteelHead SD HA IP address and status

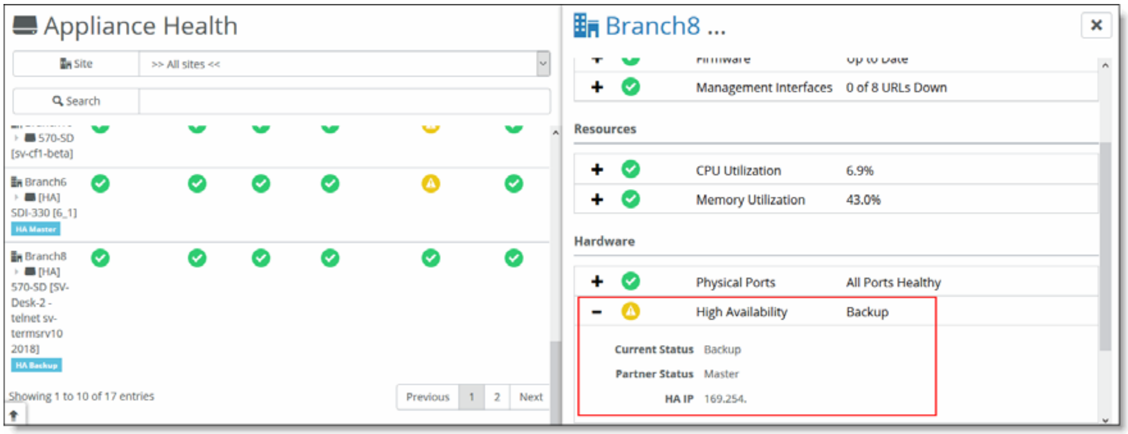

Appliance Health

Site >> All sites <<

Search

› 570-SD [sv-cf1-beta]	✓	✓	✓	✓	⚠	✓
Branch6 › [HA] SDI-330 [6_1] HA Master	✓	✓	✓	✓	⚠	✓
Branch8 › [HA] 570-SD [SV-Desk-2 - telnet sv-termsrv10 2018] HA Backup	✓	✓	✓	✓	✓	✓

Showing 1 to 10 of 17 entries

Previous 1 2 Next

Branch8 ...

+	✓	Firmware	Up to Date
+	✓	Management Interfaces	0 of 8 URLs Down

Resources

+	✓	CPU Utilization	6.9%
+	✓	Memory Utilization	43.0%

Hardware

+	✓	Physical Ports	All Ports Healthy
−	⚠	High Availability	Backup

Current Status Backup

Partner Status Master

HA IP 169.254.

Displaying underlay FIB and ARP tables

SCM displays the underlay Forward Information (FIB) and Address Resolution Protocol (ARP) tables for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

To display the underlay FIB and ARP tables

1. Choose Health Check > Appliance Health.

2. Select the SteelHead SD appliance to expand the page.

3. Under Networking, click Underlay ARP to display the FIB table.

Figure 10-4. Displaying underlay ARP table

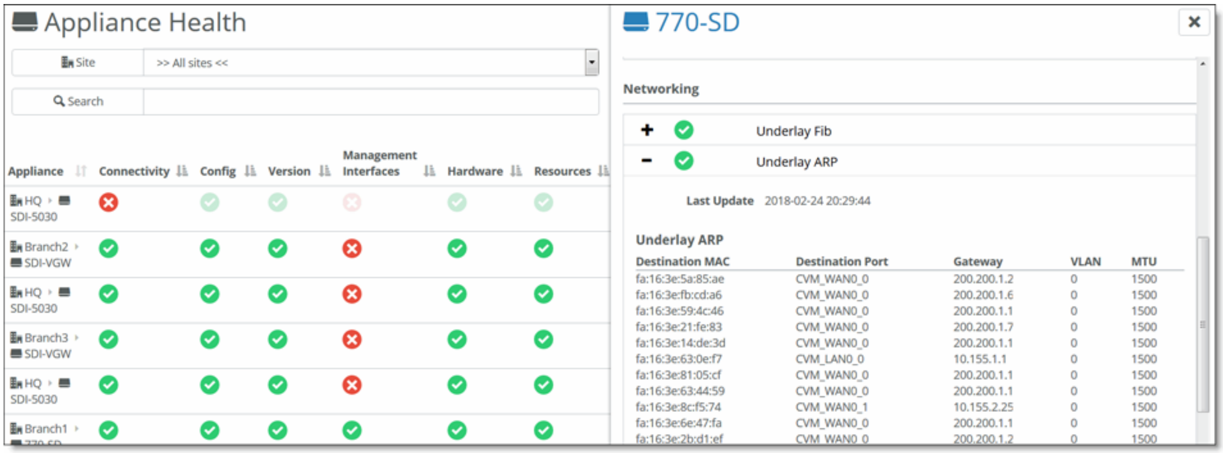

Appliance Health

Site >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resources
HQ › SDI-5030	✗	✓	✓	✗	✓	✓
Branch2 › SDI-VGW	✓	✓	✓	✗	✓	✓
HQ › SDI-5030	✓	✓	✓	✗	✓	✓
Branch3 › SDI-VGW	✓	✓	✓	✗	✓	✓
HQ › SDI-5030	✓	✓	✓	✗	✓	✓
Branch1 › 770-SD	✓	✓	✓	✓	✓	✓

770-SD

Networking

+	✓	Underlay Fib
−	✓	Underlay ARP

Last Update 2018-02-24 20:29:44

Underlay ARP

Destination MAC	Destination Port	Gateway	VLAN	MTU
fa:16:3e:5a:85:ae	CVM_WAN0_0	200.200.1.2	0	1500
fa:16:3e:fb:cd:a6	CVM_WAN0_0	200.200.1.6	0	1500
fa:16:3e:59:4c:46	CVM_WAN0_0	200.200.1.1	0	1500
fa:16:3e:21:fe:83	CVM_WAN0_0	200.200.1.7	0	1500
fa:16:3e:14:de:3d	CVM_WAN0_0	200.200.1.1	0	1500
fa:16:3e:63:0e:f7	CVM_LAN0_0	10.155.1.1	0	1500
fa:16:3e:81:05:cf	CVM_WAN0_0	200.200.1.1	0	1500
fa:16:3e:63:44:59	CVM_WAN0_0	200.200.1.1	0	1500
fa:16:3e:8c:f5:74	CVM_WAN0_1	10.155.2.25	0	1500
fa:16:3e:6e:47:fa	CVM_WAN0_0	200.200.1.1	0	1500
fa:16:3e:2b:d1:ef	CVM_WAN0_0	200.200.1.2	0	1500

Displaying FIB tables for an organization

SCM displays the FIB tables at the organization level for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

To display FIB tables for an organization

1. Choose Health Check > Routing Tables.

2. Select the FIB Tables tab to display the FIB tables for all the appliances in the organization.

Figure 10-5. FIB tables

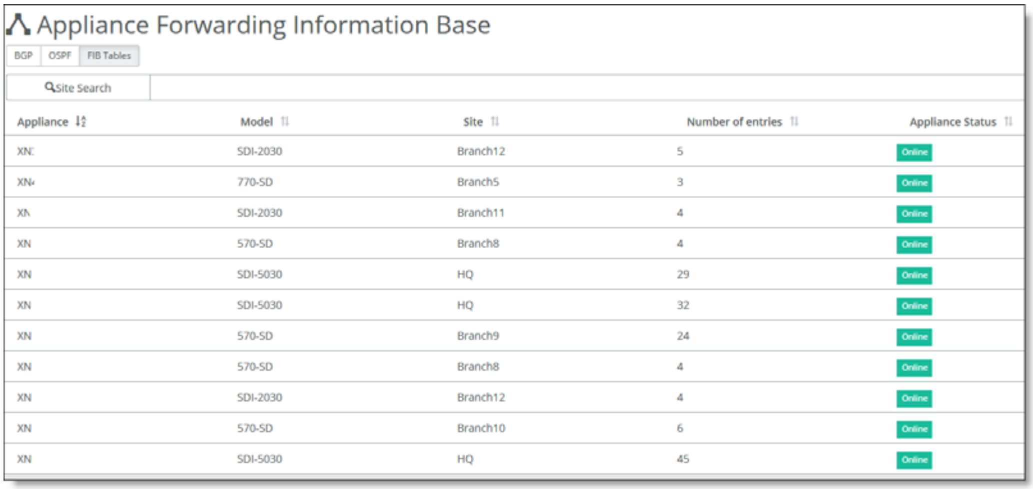

Appliance	Model	Site	Number of entries	Appliance Status
XN	SDI-2030	Branch12	5	Online
XN	770-SD	Branch5	3	Online
XN	SDI-2030	Branch11	4	Online
XN	570-SD	Branch8	4	Online
XN	SDI-5030	HQ	29	Online
XN	SDI-5030	HQ	32	Online
XN	570-SD	Branch9	24	Online
XN	570-SD	Branch8	4	Online
XN	SDI-2030	Branch12	4	Online
XN	570-SD	Branch10	6	Online
XN	SDI-5030	HQ	45	Online

3. To view the details for the FIB tables, click the appliance that has FIB entries.

Figure 10-6. Viewing FIB details by appliance

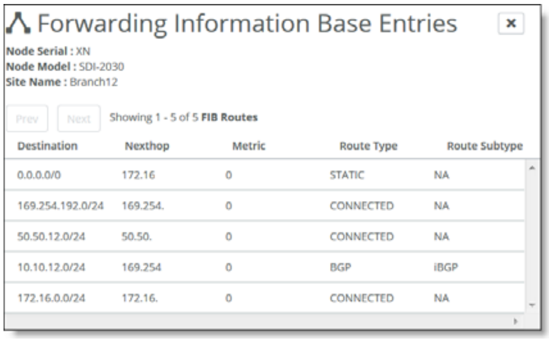

Node Serial : XN
Node Model : SDI-2030
Site Name : Branch12

Showing 1 - 5 of 5 **FIB Routes**

Destination	Nexthop	Metric	Route Type	Route Subtype
0.0.0.0/0	172.16	0	STATIC	NA
169.254.192.0/24	169.254.	0	CONNECTED	NA
50.50.12.0/24	50.50.	0	CONNECTED	NA
10.10.12.0/24	169.254	0	BGP	iBGP
172.16.0.0/24	172.16.	0	CONNECTED	NA

The Route Type specifies whether the route is directly connected. It can be connected either by OSPF or BGP.

Displaying BGP peer tables

SCM displays the BGP peer and routing tables for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

To display a BGP peer table

1. Choose Health Check > Routing Tables.
2. Select the BGP tab.
3. To display the BGP tables for all the appliances in the organization, select the BGP Tables tab. All the BGP learned and advertised routes are displayed.

Figure 10-7. BGP tables and details

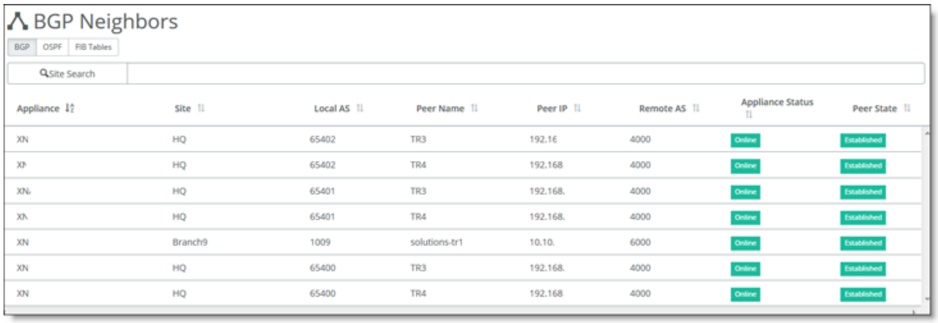

Appliance	Site	Local AS	Peer Name	Peer IP	Remote AS	Appliance Status	Peer State
XN	HQ	65402	TR3	192.16	4000	Online	Established
XN	HQ	65402	TR4	192.168	4000	Online	Established
XN	HQ	65401	TR3	192.168.	4000	Online	Established
XN	HQ	65401	TR4	192.168.	4000	Online	Established
XN	Branch9	1009	solutions-tr1	10.10.	6000	Online	Established
XN	HQ	65400	TR3	192.168.	4000	Online	Established
XN	HQ	65400	TR4	192.168	4000	Online	Established

4. Select the appliance to expand the page.

Figure 10-8. Viewing BGP learned and advertised routes

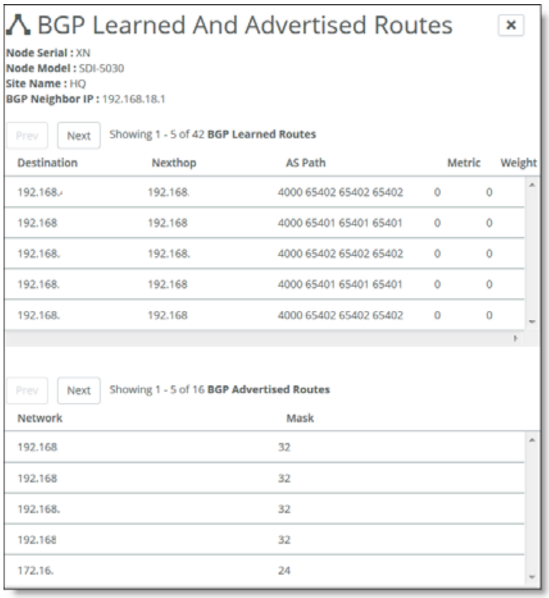

Node Serial : XN
Node Model : SDI-5030
Site Name : HQ
BGP Neighbor IP : 192.168.18.1

Showing 1 - 5 of 42 **BGP Learned Routes**

Destination	Nexthop	AS Path	Metric	Weight
192.168.	192.168.	4000 65402 65402 65402	0	0
192.168	192.168	4000 65401 65401 65401	0	0
192.168.	192.168.	4000 65402 65402 65402	0	0
192.168.	192.168	4000 65401 65401 65401	0	0
192.168.	192.168	4000 65402 65402 65402	0	0

Showing 1 - 5 of 16 **BGP Advertised Routes**

Network	Mask
192.168	32
192.168	32
192.168.	32
192.168	32
172.16.	24

You can navigate using the Prev/Next buttons. The total number of learned BGP routes is displayed in the label beside the buttons.

The lower half of the page displays the BGP Advertised Routes table (if any).

Displaying OSPF nodes and routes

SCM displays the OSPF neighbors and routing tables for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

To display OSPF nodes and routes

1. Choose Health Check > Routing Tables.

2. To display the OSPF tables for all the appliances in the organization, select the OSPF Tables tab. All the OSPF nodes are displayed.

Figure 10-9. Displaying OSPF appliances

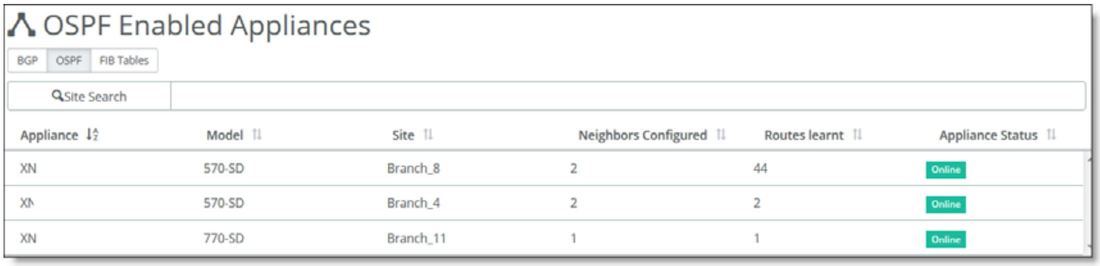

OSPF Enabled Appliances

Appliance	Model	Site	Neighbors Configured	Routes learnt	Appliance Status
XN	570-SD	Branch_8	2	44	Online
XN	570-SD	Branch_4	2	2	Online
XN	770-SD	Branch_11	1	1	Online

3. Select an OSPF appliance to display the OSPF nodes and routes for the appliance.

Figure 10-10. OSPF appliance neighbors and learned routes

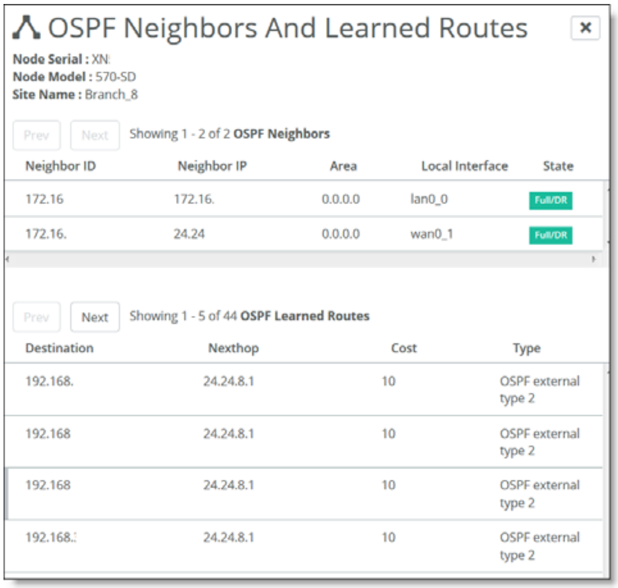

OSPF Neighbors And Learned Routes

Node Serial : XN
Node Model : 570-SD
Site Name : Branch_8

Showing 1 - 2 of 2 **OSPF Neighbors**

Neighbor ID	Neighbor IP	Area	Local Interface	State
172.16	172.16.	0.0.0.0	lan0_0	Full/DR
172.16.	24.24	0.0.0.0	wan0_1	Full/DR

Showing 1 - 5 of 44 **OSPF Learned Routes**

Destination	Nexthop	Cost	Type
192.168.	24.24.8.1	10	OSPF external type 2
192.168	24.24.8.1	10	OSPF external type 2
192.168	24.24.8.1	10	OSPF external type 2
192.168.	24.24.8.1	10	OSPF external type 2

You can navigate using the Prev/Next buttons. The total number of OSPF neighbors is displayed in the label beside the buttons.

The lower half of the page displays the OSPF advertised routes (if any). The total number of learned OSPF routes is displayed in the label beside the buttons.

Displaying NTP server status

SCM displays NTP server status for SteelHead SD 570-SD, 770-SD, 3070-SD appliances and the SteelConnect SDI-2030 gateway located at the branch.

To display NTP server status

1. Choose Health Check > Appliance Health.

2. Select the SteelHead SD to expand the pane.

Figure 10-11. Viewing appliance health

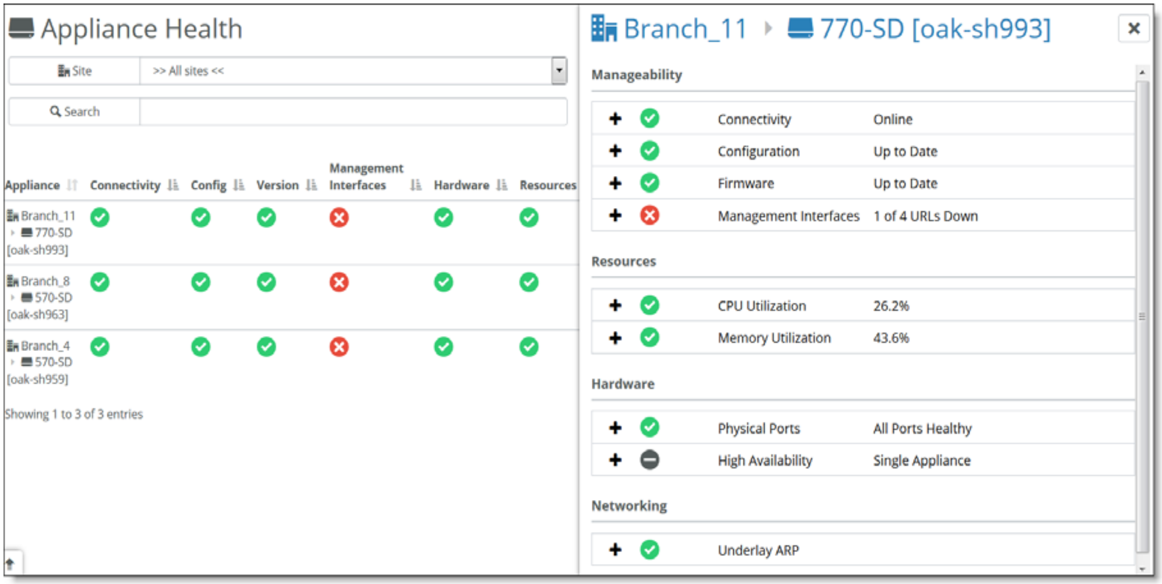

Appliance Health

Site: >> All sites <<

Search

Appliance	Connectivity	Config	Version	Management Interfaces	Hardware	Resources
Branch_11 › 770-SD [oak-sh993]	✓	✓	✓	✗	✓	✓
Branch_8 › 570-SD [oak-sh963]	✓	✓	✓	✗	✓	✓
Branch_4 › 570-SD [oak-sh959]	✓	✓	✓	✗	✓	✓

Showing 1 to 3 of 3 entries

Branch_11 › 770-SD [oak-sh993]

Manageability

+	✓	Connectivity	Online
+	✓	Configuration	Up to Date
+	✓	Firmware	Up to Date
+	✗	Management Interfaces	1 of 4 URLs Down

Resources

+	✓	CPU Utilization	26.2%
+	✓	Memory Utilization	43.6%

Hardware

+	✓	Physical Ports	All Ports Healthy
+	⊖	High Availability	Single Appliance

Networking

+	✓	Underlay ARP	

3. Under Manageability, select Management Interfaces.

Figure 10-12. Displaying NTP server status

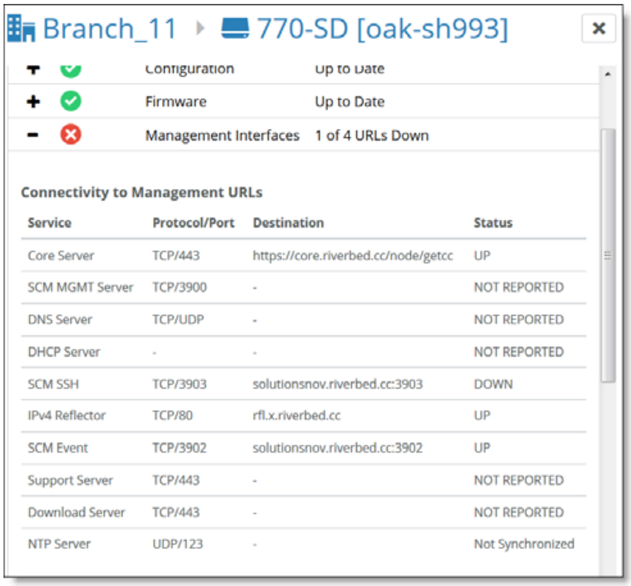

Branch_11 › 770-SD [oak-sh993]

+	✓	Configuration	Up to Date
+	✓	Firmware	Up to Date
−	✗	Management Interfaces	1 of 4 URLs Down

Connectivity to Management URLs

Service	Protocol/Port	Destination	Status
Core Server	TCP/443	https://core.riverbed.cc/node/getcc	UP
SCM MGMT Server	TCP/3900	-	NOT REPORTED
DNS Server	TCP/UDP	-	NOT REPORTED
DHCP Server	-	-	NOT REPORTED
SCM SSH	TCP/3903	solutionsnov.riverbed.cc:3903	DOWN
IPv4 Reflector	TCP/80	rfl.x.riverbed.cc	UP
SCM Event	TCP/3902	solutionsnov.riverbed.cc:3902	UP
Support Server	TCP/443	-	NOT REPORTED
Download Server	TCP/443	-	NOT REPORTED
NTP Server	UDP/123	-	Not Synchronized

Enabling SNMP reporting and logging

SNMP reporting is supported on SteelHead SD SD-570, SD-770, SD-3070 and 2030-SDI appliances located at the branch.

When direct SNMP reporting is enabled, your network management system (NMS) initiates the SNMP poll to all individual appliances in a realm. The appliances send SNMP data directly to the NMS. You can override this setting to limit the SNMP data to all gateways within an organization.

To enable SNMP operations (such as SNMP GET and SNMP WALK or asynchronous traps from a SteelHead SD appliance), you need to provide the NMS interface IP address through which the appliance is reachable. The IP address of the appliance is required:

- For a 2030-SDI located at the data center, use the management IP address. To view the management IP, choose Appliances > IPs tab.
- For the SD-570, SD-770, SD-3070, and 2030-SDI at the branch, use the zone IP address for the appliance. To view the zone IP address, choose Appliances. Select the appliance and click the IPs tab. Scroll down to Under Zone gateway assignment (internal) IPs, to view the zone IP address.

The port number to query the SteelHead SDs is the standard SNMP port 161. You need to specify this port if the NMS doesn't use it by default.

For details on SNMP supported versions and how to configure SNMP, see the *SteelConnect Manager User Guide*.

Exporting syslog messages to an external syslog server

You can export syslogs to an external server using SCM for SteelHead SD SD-570, SD-770, SD-3070, and 2030-SDI appliances located at the branch. For syslog reporting, the only requirement is that the appliance reach the remote syslog server.

For details on sending syslog data to a remote server, see the *SteelConnect Manager User Guide*.

A

Port Mapping for SteelHead SD

This appendix summarizes the port mapping for SteelHead SD appliances. It includes these sections:

- "SteelHead SD 570-SD and 770-SD appliances" on page 99
- "SteelHead SD 3070-SD appliance" on page 101

SteelHead SD 570-SD and 770-SD appliances

Physical ports

The SteelHead SD 570-SD and 770-SD appliances have these ports:

- AUX, PRI, LAN0_0, WAN0_0, LAN0_1, WAN0_1

CVM ports

The CVM has these ports:

- knet2, knet3, knet4, knet5, knet6, knet7

Physical port to flows port mapping

Physical port	AUX	Primary	LAN0_0	WAN0_0	LAN0_1	WAN0_1
Flows port	8	9	10	11	12	13

Service chain virtual machines

Virtual machine (VM)	Pod name	Function
Service virtual machine (SVM)	catfish_secure_node0	Overlay tunnels, QoS, NAT, etc.
Routing virtual machine (RVM)	routing_pod0	Routing protocols, DNS service
Virtual SteelHead (vSH)	vsh_node0	WAN optimization

SteelHead SD dynamically allocates vSwitch ports based on service chain configuration and the WAN optimization toggle.

vSwitch mapped VM ports

The vSwitch port mapping state can be fetched at runtime using this command on the CVM:

```
XNXXXXD8XXXA9FF9-CVM:>orchestrator-agent --get_port_interface_mapping
```

Node name	Interface name	Port
cvm	knet2	AUX
cvm	knet3	PRI
cvm	knet4	LAN0_0
cvm	knet5	WAN0_0
cvm	knet6	LAN0_1
cvm	knet7	WAN0_1
catfish_secure_node0	knet22	WAN0_1
catfish_secure_node0	knet23	WAN0_0
catfish_secure_node0	knet24.1101	LAN0_0
catfish_secure_node0	knet24.1100	LAN0_0
catfish_secure_node0	knet25	LAN0_1
catfish_secure_node0	knet26	— (binds to vSHLAN0_0)
catfish_secure_node0	knet27	— (binds to vSH WAN0_0)
routing_pod0	knet18	LAN0_1
routing_pod0	knet19.1101	LAN0_0
routing_pod0	knet19.1100	LAN0_0
routing_pod0	knet20	WAN0_1
routing_pod0	knet21	WAN0_0
vsh_node0	knet14	PRI
vsh_node0	knet15	AUX
vsh_node0	knet16	LAN0_0
vsh_node0	knet17	WAN0_0

Bridged VM ports for internal communication

Source	Port name	IP address	Protocol	Remote end	Purpose
CVM	port1	169.254.0.2	Static	Hypervisor mgmt_br bridge	Connects to hypervisor
	port2	169.254.169.254	Static	Hypervisor linklocal_br bridge	Connects to service chain VMs
SVM	port1	—*	Static*	Hypervisor linklocal_br bridge	Connects to CVM
RVM	port1	—*	Static*	Hypervisor linklocal_br bridge	Connects to CVM
vSH	hpn	—*	DHCP	Hypervisor linklocal_br bridge	Connects to CVM

* Allocated at runtime.

SteelHead SD 3070-SD appliance

Physical ports

The SteelHead SD 3070-SD appliance has these physical ports:

- AUX, PRI, LAN3_0, LAN3_0, WAN3_1, WAN3_1

These ports are present only if you have installed an add-on NIC:

- LAN2_0, WAN2_0, LAN2_1, WAN2_1

CVM ports

The CVM has these ports:

- knet2, knet3, knet4, knet5, knet6, knet7

These ports are present only if you have installed an add-on NIC:

- knet8, knet9, knet10, knet11

Physical port to flows port mapping

Physical port	AUX	Primary	LAN3_0	WAN3_0	LAN3_1	WAN3_1
Flows port	8	9	10	11	12	13

Note: The 3070-SD appliance supports add-on NICs. The presence of an add-on NIC can change the total NIC count on the appliance and can also result in different flows port mapping accordingly. Each add-on NIC can carry either two or four NICs. For details on add-on NICs, see "NIC support" on page 18.

SVM ports

There are four more virtual NICs in SVM for each physical add-on NIC.

RVM ports

There are four more virtual NICs in RVM for each physical add-on NIC.

vSH ports

The vSH has these ports:

- hpn, PRI, AUX, LAN0_0, WAN0_0, inpath0_0

vSH has only one LAN-WAN pair and will not change with the addition of any physical add-on NIC.

B

TOS, DSCP, QoS Traffic Class Table

This appendix contains the TOS, DSCP, and QoS traffic Classes table. For details on configuring QoS shaping for SteelHead SD appliances, see "Configuring QoS Shaping on SteelHead SD" on page 87.

TOS, DSCP, and QoS Traffic Classes Table

TOS Value	DSCP Value	Traffic Class ID	Traffic Class Priority
0	0	1	Normal
4	1	2	High
8	2	1	Normal
12	3	1	Normal
16	4	2	High
20	5	1	Normal
24	6	1	Normal
28	7	1	Normal
32	8	0	Low
36	9	1	Normal
40	10	1	Normal
44	11	1	Normal
48	12	1	Normal
52	13	1	Normal
56	14	1	Normal
60	15	1	Normal
64	16	2	High
68	17	1	Normal
72	18	2	High
76	19	1	Normal
80	20	2	High
84	21	1	Normal

TOS Value	DSCP Value	Traffic Class ID	Traffic Class Priority
88	22	2	High
92	23	1	Normal
96	24	2	High
100	25	1	Normal
104	26	2	High
108	27	1	Normal
112	28	2	High
116	29	1	Normal
120	30	2	High
124	31	1	Normal
128	32	3	Urgent
132	33	1	Normal
136	34	2	High
140	35	1	Normal
144	36	2	High
148	37	1	Normal
152	38	2	High
156	39	1	Normal
160	40	3	Urgent
164	41	1	Normal
168	42	1	Normal
172	43	1	Normal
176	44	3	Urgent
180	45	1	Normal
184	46	3	Urgent
188	47	1	Normal
192	48	3	Urgent
196	49	1	Normal
200	50	1	Normal
204	51	1	Normal
208	52	1	Normal
212	53	1	Normal
216	54	1	Normal
220	55	1	Normal

TOS Value	DSCP Value	Traffic Class ID	Traffic Class Priority
224	56	3	Urgent
228	57	1	Normal
232	58	1	Normal
236	59	1	Normal
240	60	1	Normal
244	61	1	Normal
248	62	1	Normal
252	63	1	Normal