



SD-WAN Deployment Guide

SteelConnect 2.10

April 2018



© 2018 Riverbed Technology, Inc. All rights reserved.

Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein cannot be used without the prior written consent of Riverbed or their respective owners.

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

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

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

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

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

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

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



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

Part Number
712-00315-01

Contents

Welcome	7
About this guide	7
Audience	7
Related documentation	7
Riverbed SD-WAN terminology	8
Documentation and release notes	8
Contacting Riverbed	9
 1 - Overview of SteelConnect SD-WAN	 11
SteelConnect Manager	12
SteelConnect gateways	12
SteelConnect switches	13
SteelConnect access points	13
 2 - How a Packet Traverses a SteelConnect Network.....	 15
SteelConnect gateway services	16
Routing services overview	16
Traffic forwarding mechanism	17
Routing example: Branch B to Branch A.....	18
Routing example: Branch C to Branch F	19
Routing example: Branch C to Branch D.....	20
Packet operation flow.....	21
Example: traffic sent from site to site.....	22
Example: traffic to the internet.....	23
Example: inbound traffic	24
 3 - SteelConnect and Network Security.....	 25
SteelConnect as a perimeter firewall.....	25
SteelConnect security features	26
Application-defined security with granular filtering rules	26
WAN hardening and network zoning	27
VPN capabilities.....	28

Inbound NAT capabilities.....	30
4 - LAN Topologies	33
Spanning tree on LAN ports.....	33
Layer 2: Access port	34
Layer 2: Trunk port.....	37
Layer 2: High availability.....	38
Layer 3: Switch behind a SteelConnect gateway with static routing.....	40
Layer 3: Switch behind a SteelConnect gateway with LAN-side OSPF	44
Layer 3: High availability.....	47
Integration with firewalls.....	49
Integration with Riverbed SteelHead	51
5 - WAN Topologies	53
Important configuration information for all WAN topologies	53
Single internet connection	54
Multiple internet connections	58
Secondary 3G/4G uplink.....	63
Integration with private networks.....	65
Integration with MPLS CE router	69
MPLS CE router replacement: ASBR-like deployment.....	71
Hybrid WAN: internet and MPLS	76
Integration with WAN-facing firewall	78
Highly secured WAN	79
WAN high availability.....	82
High-availability WAN topology with two routers	83
6 - Data Center Topologies.....	89
Using HA LAN and WAN topologies in the data center.....	89
Traffic attraction with SDI-5030	90
Cabling the appliances	92
Registering the appliances to the organization.....	93
Enabling the data ports	94
Deploying the appliances in a site	94
Creating a data center cluster of 5030 gateways	95
Using BGP for traffic attraction on SCM.....	96
Configuring BGP on the WAN aggregation routers.....	99
Creating data center uplinks for the cluster	99

7 - AWS Cloud Topologies	101
Deployment into the AWS cloud.....	101
Deploying a HA SteelConnect gateway in AWS	102
SteelConnect gateway with SteelHead in AWS	108
Transit VPC Routing using SteelConnect gateways	112
Deploying an AWS transit VPC topology	113
Deployment over AWS DirectConnect	120
8 - Azure Cloud Topologies.....	125
Deployment into the Azure cloud	125
SteelConnect gateway in Azure	126
Deploying SteelConnect gateway in Azure with SteelHead	131
Deploying over Azure ExpressRoute.....	133
Route-tables and user-defined routes (UDRs)	136
Manual routing	138
9 - End-to-End Topology.....	139
Full topology.....	140
Connecting sites with the same WAN	141
SD-WAN sites communicating between each other	141
Brownfield scenario: SD-WAN site communicating with legacy internet site	141
Brownfield scenario: SD-WAN site communicating with legacy MPLS site.....	147
Connecting sites with different WANs	147
SD-WAN sites: multihub	148
SD-WAN sites: full-mesh mode	151
Brownfield scenario: SD-WAN site communicating with legacy site	155
SD-WAN sites: path redundancy.....	158
Brownfield scenario: path redundancy	160

Welcome

About this guide

Welcome to the *SD-WAN Deployment Guide*. This guide provides advanced configuration and conceptual information about the deployment of SteelConnect SD-WAN products into your network.

Audience

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

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

- SteelConnect Manager (SCM)
- SteelConnect gateways
- SteelConnect switches
- SteelConnect access points
- SteelConnect connectivity to the Infrastructure as a Service (IaaS) cloud with Amazon Web Services (AWS)
- SteelConnect connectivity to the IaaS cloud with Microsoft Azure
- SteelHead SD

Related documentation

This guide supplements the documentation and technical notes available on the SteelConnect page of the Riverbed Support site. Two essential additional references are:

- *SteelConnect Manager User Guide*
- *SteelHead SD User Guide*

The Riverbed splash forum provides a community to share SteelConnect information:

- *Riverbed Splash* at <https://splash.riverbed.com/community/product-lines/steelconnect>

Riverbed SD-WAN terminology

Knowing these common terms will help you understand the deployment details in this guide.

- **SteelConnect** - Provides cloud-based management system software for SD-WAN gateways, WiFi access points, and Ethernet switches.
- **SteelConnect Manager (SCM)** - The centralized and multitenant management portal that provides an intuitive and simplified workflow for designing, deploying, and managing distributed and hybrid networks.
- **Virtual gateway** - provides the functionality of a SteelConnect appliance and can run in AWS or any hypervisor like VMware, Hyper-V, KVM, or Xen.
- **AutoVPN** - Connects multiple sites with a secure, full-mesh VPN.
- **Classic VPN** - Creates a manual VPN tunnel using the standard IPsec IKEv1 or IKEv2 protocol. You can connect to a third-party VPN using classic VPN. A remote gateway is not necessary. Classic VPN configurations can classify traffic based on TCP/UDP port number, providing a more granular approach to traffic steering.
- **RouteVPN** - IPsec VPN tunnels built over the internet WAN.
- **SwitchVPN** - SteelConnect's Layer 2 VPN, based on IPsec. It automatically makes a zone available in a remote site when you use the same zone for multiple sites. No manual configuration is necessary.
- **Organization** - A company representing an end customer. It contains the customer details, sites, devices, zones associated with the devices, the uplinks, and so on.
- **Site** - Physical location of one or more office buildings, a hosting location, or a cloud location that makes up the organization. A site houses a SteelConnect gateway and uses a permanent DNS alias.
- **Zone** - Layer 3 network segments or VLANs within sites that are VLAN-tagged traffic. A zone always has a VLAN tag assigned to it.

Documentation and release notes

The most current versions 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 see the most current version of the release notes, click the question mark (?) in the upper-right corner of SCM, and choose What's new in version x.x.x?. We recommend taking a look at the release notes before you use a new version of SCM.

Contacting Riverbed

This section describes how to contact departments within Riverbed.

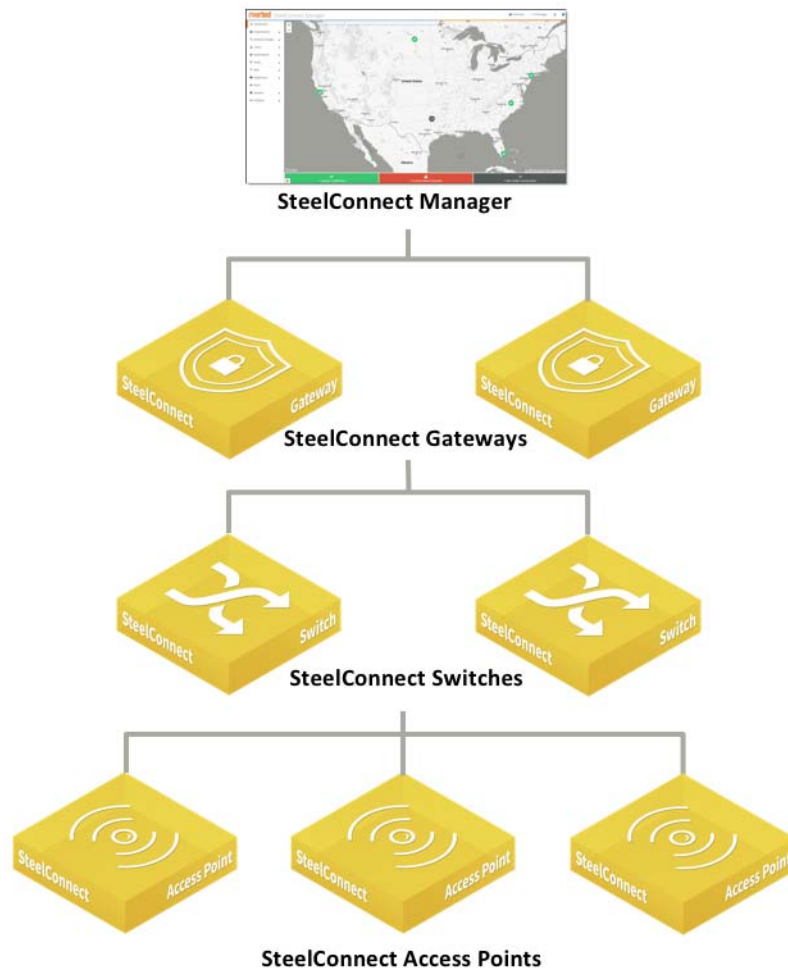
- **Technical support** - Problems installing, using, or replacing Riverbed products? Contact Riverbed Support or your channel partner who provides support. To contact Riverbed Support, open a trouble ticket by calling 1-888-RVBD-TAC (1-888-782-3822) in the United States and Canada or +1 415-247-7381 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.html>.
- **Documentation** - Have suggestions about Riverbed's online documentation or printed materials? Send comments to techpubs@riverbed.com.

Overview of SteelConnect SD-WAN

SD-WAN lets enterprises simplify their network configuration and management. With SD-WAN, enterprises can intuitively manage networks based on parameters relevant to their businesses such as applications, users, locations, performance, and security.

The Riverbed SD-WAN solution provides an intelligent and intuitive approach to designing, deploying, and managing distributed networks for the modern hybrid enterprise. The solution consists of appliances and a centralized management console that the administrator uses to view network health, deploy appliances, and make changes to policies. [Figure 1-1](#) shows the high-level architecture.

Figure 1-1. SteelConnect high-level architecture



SteelConnect components comprise the following:

- [“SteelConnect Manager” on page 12](#)
- [“SteelConnect gateways” on page 12](#)
- [“SteelConnect switches” on page 13](#)
- [“SteelConnect access points” on page 13](#)

SteelConnect Manager

SteelConnect Manager (SCM) is a web-based, management portal that lets you design the network before deploying any hardware. You can use SCM to push configurations to the devices for deployment of infrastructure without the need for an engineer to be on site. After deployment, SCM provides network visibility for manageability and troubleshooting.

The SCM server currently resides in the global Amazon Web Services (AWS) cloud public infrastructure and orchestrates a series of services hosted by Riverbed. Each service has dependencies that function as a part of the collective SteelConnect infrastructure. These services include:

- Management console
- Global certificate authorities (CAs)
- Network Time Protocol (NTP)
- Dynamic Domain Name System (DNS)
- IP address reflectors, which are simple mechanisms for all gateways to find their public IP address per uplink and report the address to SCM.
- Structured Query Language (SQL) relational databases that keep track of which SCMs are associated with which organizations, sites, and devices.

SteelConnect appliances (gateways, switches, and access points) connect to core services, and the services associated with them, to find their assigned server. After an appliance is paired with SCM, it connects only to its corresponding SCM. Each SCM communicates through various services for updates regarding the appliance registration and management changes. All communication between the appliances and SCM, as well as all interoperating services inside of SCM, are authenticated through x509 certificate validation. These Riverbed-owned certificates are exchanged and validated for authenticity.

SCM manages all appliances, including all firmware upgrades.

SteelConnect gateways

SteelConnect gateways are physical and virtual secure WAN gateways that provide unified connectivity (point-to-point and full-mesh) and enforcement of global policy across on-premises and cloud network environments, zero-touch provisioning, and secure automated VPN management.

Gateways are categorized into hardware and software appliances. The gateways automatically map into connected network segments, called *zones*, to:

- provide basic network services.

- handle one or more uplinks, either by concurrent use or as backup.
- enable policy enforcement.
- enforce security.
- enable extended reporting for connected zones.
- connect multiple sites with a secure, full-mesh VPN or a hub-and-spoke VPN deployment.

The SteelConnect gateway comes in various form factors to accommodate a variety of network architectures:

- **Branch gateways** - SDI-130, SD-130W, SDI-330, SDI-1030
- **Data center gateway** - SDI-5030
- **SteelHead SD** - 570-SD, 770-SD, 3070-SD models deliver the benefits of SteelHead WAN optimization and SteelConnect SD-WAN while providing the flexibility of a single box solution.
- **SteelConnect gateway virtual machine (VM)** - Available for various virtualization platforms in these image types:
 - VMware
 - VirtualBox
 - KVM
 - Hyper-V
 - XenCenter
 - Amazon EC2
- **Virtual gateways on IaaS cloud environments** - Instances of the SteelConnect gateway of various sizes can be deployed on AWS or Azure. When you deploy a gateway in your cloud or multiple clouds, the RouteVPN feature lets you connect your data center to the cloud, or even multiple clouds with each other.

SteelConnect switches

- Enable plug-and-play multizone Layer 2 connectivity.
- Provide Power over Ethernet (PoE) to PoE-enabled appliances, including third-party devices.

SteelConnect access points

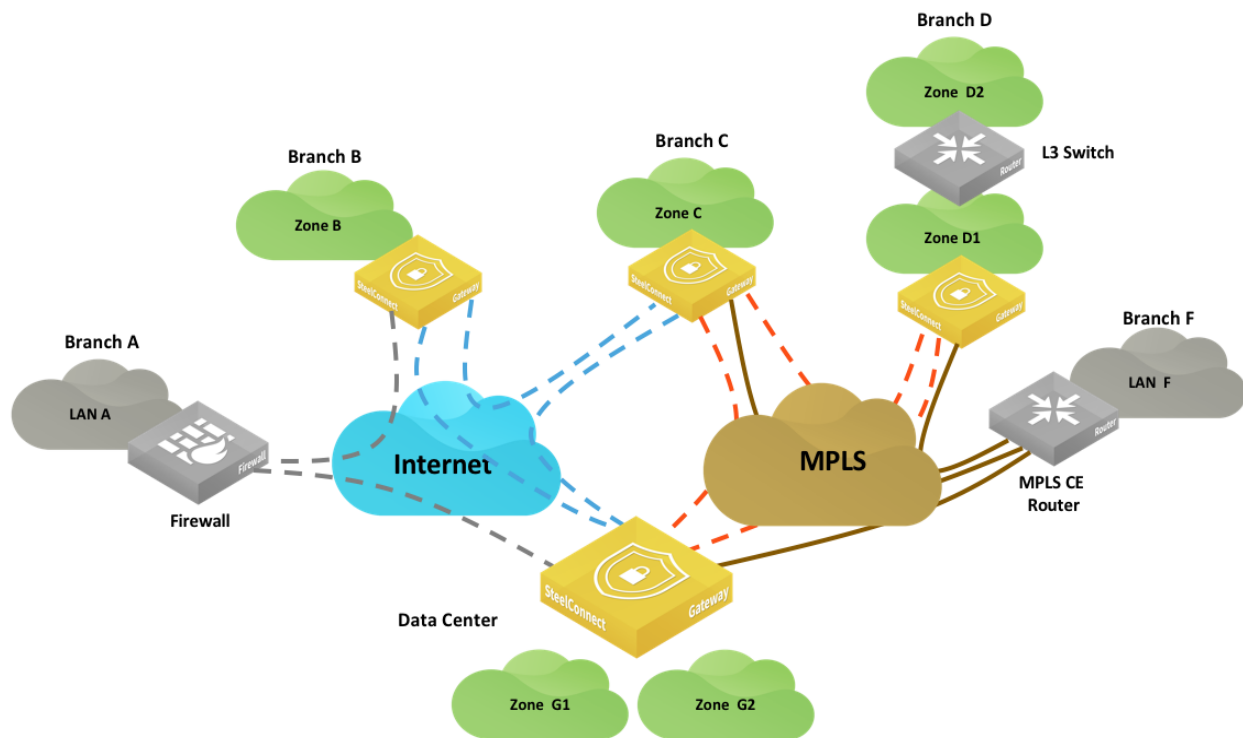
- Provide network access to WiFi clients.
- Prioritize applications and enforce policies at the edge of the network.

How a Packet Traverses a SteelConnect Network

This topic describes how the SteelConnect gateway processes packets as they traverse the network.

The SteelConnect gateway is configured and managed from SteelConnect Manager (SCM). The gateway provides several network services, such as routing services and firewall services, as packets cross the appliance.

Figure 2-1. Sample SteelConnect network



The topology in [Figure 2-1](#) uses these terms and definitions:

- A zone refers to a subnet that is local to the SteelConnect gateway or gateways at the site. The SteelConnect gateway can provide services to the clients such as DNS or DHCP; however, the default gateway or another device might perform these functions.
- In Branch D, D1 is a local zone that is used to reach another zone (Zone D2) that is not directly connected to the gateway at Layer 3. This route is referred to as a *third-party route* in SCM.

The network paths shown in [Figure 2-1](#) use these conventions and definitions:

- Gray dotted lines illustrate a generic IPsec tunnel that has been created between a SteelConnect gateway appliance and a third-party firewall. This configuration is known as a *classic VPN* in SCM.
- Blue dotted lines depict a RouteVPN tunnel. These VPN tunnels are automatically formed over the internet WAN between SteelConnect appliances.
- Orange dotted lines represent the VPN tunnels built over a private WAN (for example, MPLS) when that WAN is configured to use encryption. You can enable encryption in SCM when you create a WAN. This configuration is called *AutoVPN* in SCM.

Orange and blue lines are collectively referred to as the *overlay network*. This term is an abstraction of the internet and WAN in which the gateways communicate with each other. The communication for the overlay network takes place on an *underlay network*. The underlay is the series of network devices owned by a provider or customer making up a network infrastructure.

- Brown solid lines illustrate the connection to the MPLS network and are known as the *underlay network*.

SteelConnect gateway services

The SteelConnect gateway uses these services:

- Routing services. See [“Routing services overview” on page 16](#) for details.
- Firewall services. See [Chapter 3, “SteelConnect and Network Security”](#) for details.

Routing services overview

SteelConnect gateways provide these routing services:

- **Layer 3 routing** - The gateways offer Layer 3 capabilities that are described in [“Traffic forwarding mechanism” on page 17](#).
- **Path selection** - SCM uses traffic path rules to provide path selection. The path is selected by the gateway where the traffic originates and reflected by the gateway on the remote end. For example, traffic from Zone C to Zone E can be configured to prefer the internet path over the private WAN. These rules have automation built in to select the better path based on quality metrics. Latency, loss, and jitter are also used in the calculation, if configured.
- **Quality of Service** - The gateway ensures efficient use of bandwidth and leverages the common applications kept enhanced (CAKE) scheduling mechanism. CAKE uses an advanced fair queue mechanism based on hierarchical priority that distributes bandwidth while considering packet delays. CAKE is a connection-based system that tracks latency on each connection or traffic flow rather than bandwidth per class. CAKE is purposely built for internet-based uplinks so it is ideal for dynamic WAN throughput.

You can configure QoS in the outbound and inbound directions on the uplinks only. Branch gateways support both marking and enforcement.

Note: The SDI-5030 does not support QoS enforcement because it is out-of-path and your QoS enforcement endpoint should be on your existing MPLS routers.

Traffic forwarding mechanism

Important: SteelConnect gateways provide the fundamental subset of router functionality of an edge router but do not have one-to-one parity with a traditional router.

SteelConnect gateways classify different network types by maintaining separate routing tables using these criteria:

Local - The LAN subnets managed by the SteelConnect gateway. In [Figure 2-1 on page 15](#), Zones B, C, D1, E, and F are subnets considered as local for their respective appliance.

Overlay WAN - Subnets that are created as zones on SCM and advertised as reachable on the overlay network. The appliance in site B is aware of remote zones C, D1, E, and F using the overlay WAN.

Underlay WAN - Any WAN forwarding decisions that are made outside of (or “under”) the SteelConnect overlay. For underlay connectivity on the WAN, the site must participate in underlay learning with a routing protocol. (See the *SteelConnect Manager User Guide* for details.) In the example in [Figure 2-1 on page 15](#), Branch C would discover LAN F via a routing protocols such as BGP or OSPF. LAN F’s subnet appears in the underlay routing table of the gateway.

Classic VPN - Creates a manual VPN tunnel using the standard IPsec IKEv1 or IKEv2 protocol. A remote gateway is not necessary. For example, the tunnel could be terminated by a firewall or router or another networking device. LAN A’s subnet would appear in the classic VPN table of Branch B’s gateway.

Statically defined routes - Typically, networks behind Layer 3 routers on the LAN side of the gateway. In SteelConnect terminology, this is called a *third-party routing configuration*. D2’s subnet appears as a third-party/static route in SteelConnect.

Forwarding of traffic is done in two stages:

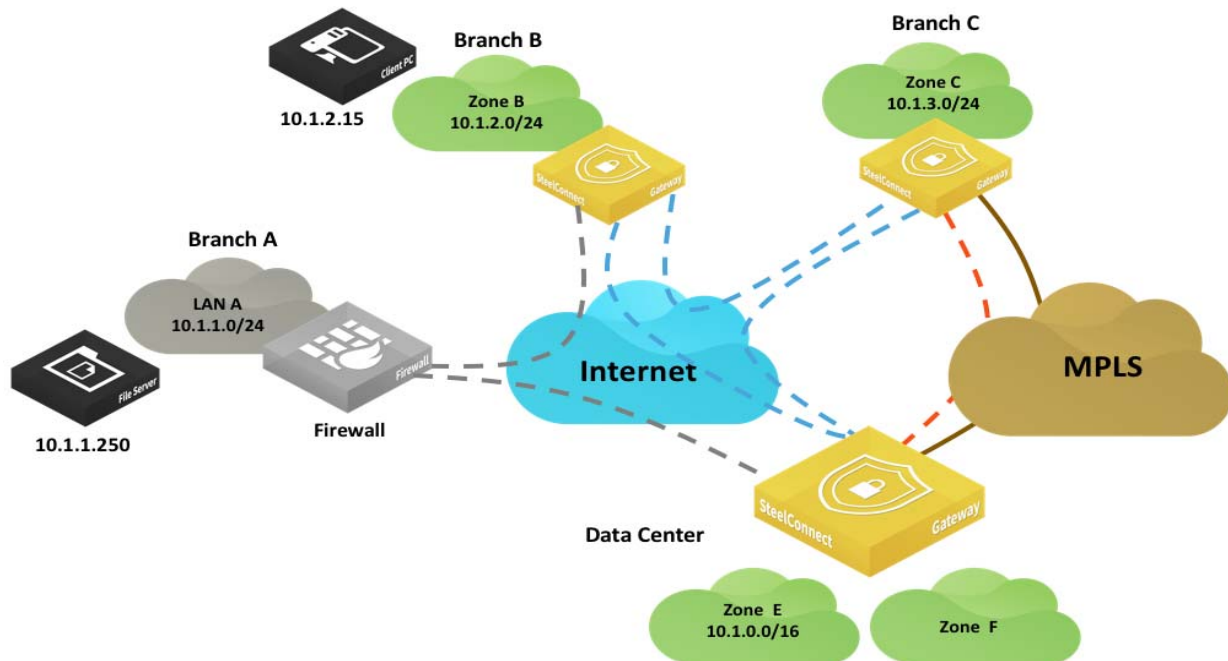
1. A route lookup is performed across all route tables. For a given destination address, the route lookup returns all available paths with the most specific prefix match.
2. The selection of the path is based on policies. You can configure a default path preference at the site or organization level. You can configure more specific traffic path rules to take precedence over less specific rules.

Important: Overlay paths are always preferred over underlay path when routes with equal prefix matches.

Routing example: Branch B to Branch A

In this example, a host (10.1.2.15) in Branch B wants to connect to a server (10.1.1.250) in Branch A.

Figure 2-2. Branch B to Branch A routing example



At the organization level, traffic uses this default path preference:

1. RouteVPN
2. MPLS

The SteelConnect gateway in Branch B has these entries in its routing table:

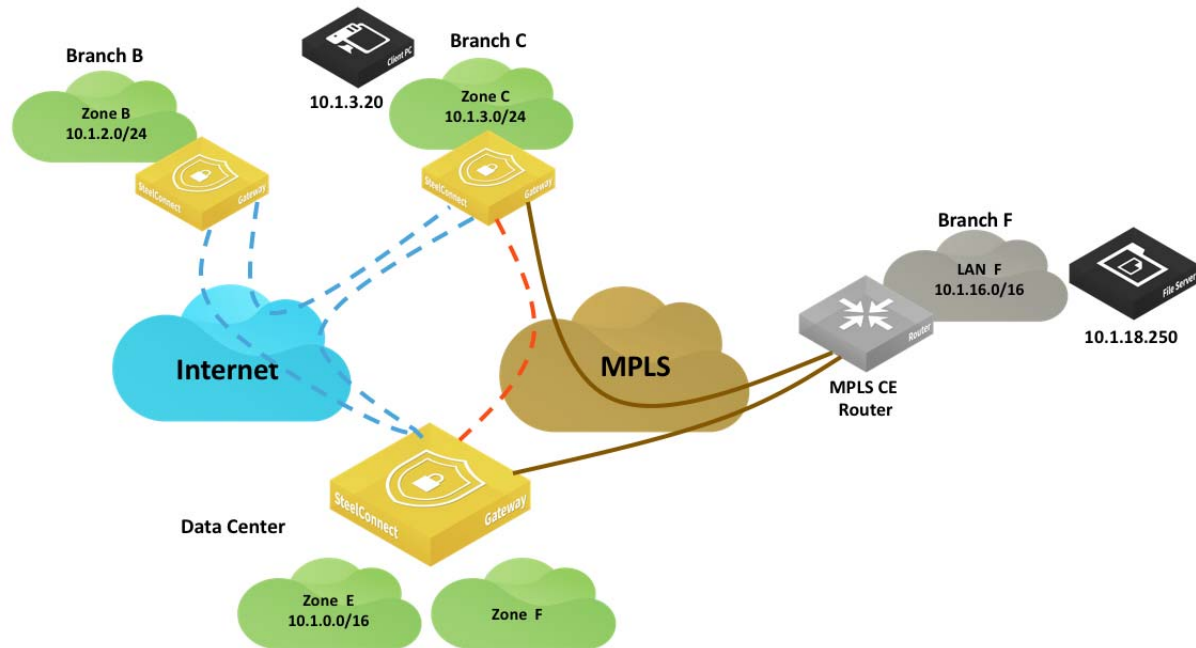
- Local: 10.1.2.0/24 directly connected
- Overlay WAN:
 - 10.1.3.0/24 learned from the Branch C gateway
 - 10.1.0.0/16 learned from the data center gateway
- Underlay WAN: Not used because the two sites do not have a common network. If the overlay fails and because we are using classless addressing, there is not a way to route traffic over the internet.
- Classic VPN: 10.1.1.0/24 learned from the IPsec tunnel to the firewall in Branch A
- Third-party zones/statically defined routes: Not used

The SteelConnect gateway in Branch B checks the routing table and chooses the most specific prefix match, which is the classic VPN connection.

Routing example: Branch C to Branch F

In this example, a host (10.1.3.20) in Branch C wants to connect to a server (10.1.18.250) in Branch F.

Figure 2-3. Branch C to Branch F routing example



At the organization level, traffic uses this default path preference:

1. RouteVPN
2. MPLS

The SteelConnect gateway in Branch C has these entries in its routing table:

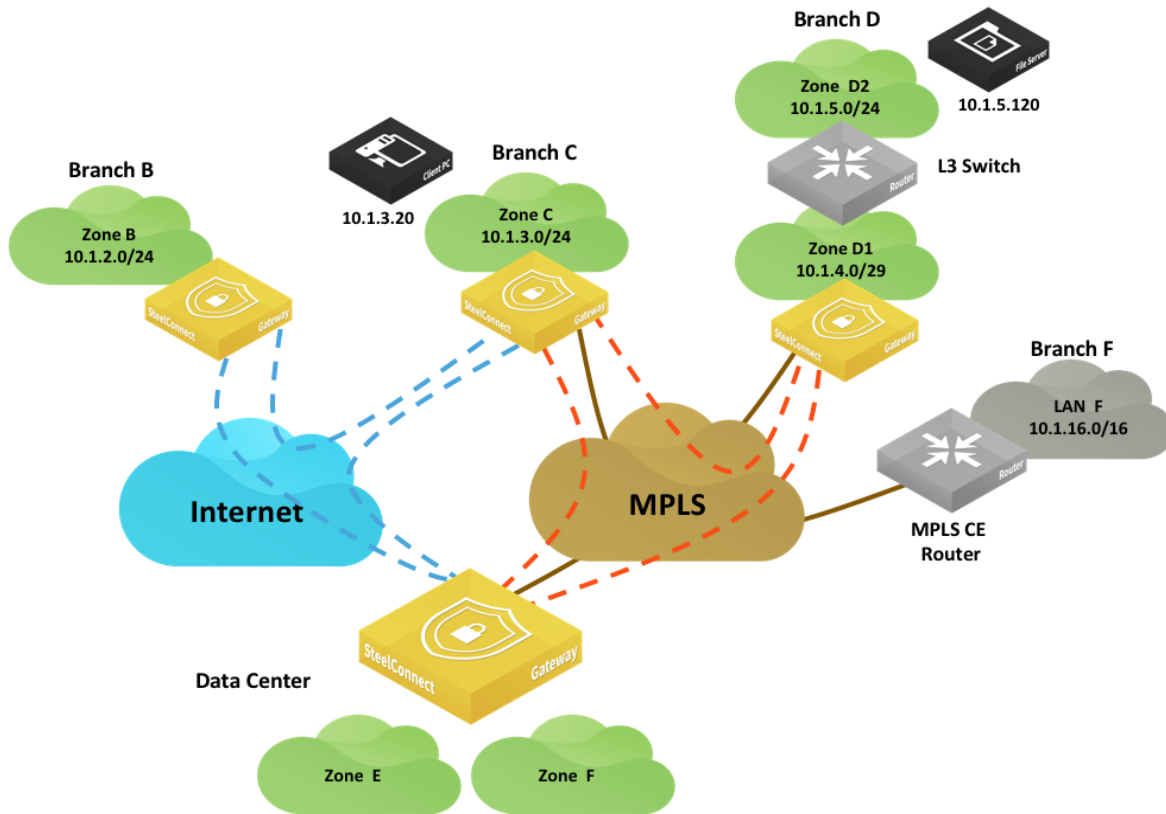
- Local: 10.1.3.0/24 directly connected
- Overlay WAN:
 - 10.1.2.0/24 learned from the Branch B gateway
 - 10.1.0.0/16 learned from the data center gateway
- Underlay WAN: 10.1.16.0/16 learned from the Branch F MPLS CE router
- Classic VPN: Not used
- Third-party zones/statically defined routes: Not used

The SteelConnect gateway in Branch B checks the routing table and chooses the most specific prefix match, which is the underlay network to MPLS.

Routing example: Branch C to Branch D

In this example, a host (10.1.3.20) in Branch C wants to connect to a server (10.1.5.120) in Branch D.

Figure 2-4. Branch C to Branch D routing example



At the organization level, SteelConnect has been changed from the default and uses this path preference:

1. MPLS
2. RouteVPN

The SteelConnect gateway in Branch C has these entries in its routing table:

- Local: 10.1.3.0/24 directly connected
- Overlay WAN:
 - 10.1.2.0/24 learned from the Branch B gateway
 - 10.1.0.0/16 learned from the data center gateway
 - 10.1.4.0/29 and 10.1.5.0/24 learned from the Branch D gateway
- Underlay WAN:
 - 10.1.16.0/16 learned from MPLS CE router in Branch F
 - 10.1.4.0/29 and 10.1.5.0/24

- Classic VPN: Not used
- Statically defined routes: Not used

The route lookup finds two available paths for 10.1.5.120. The SteelConnect gateway in Branch D is advertising its local and third-party subnets on the overlay but it is also advertising on the underlay using BGP or OSPF. See [Chapter 5, “WAN Topologies”](#) for details.

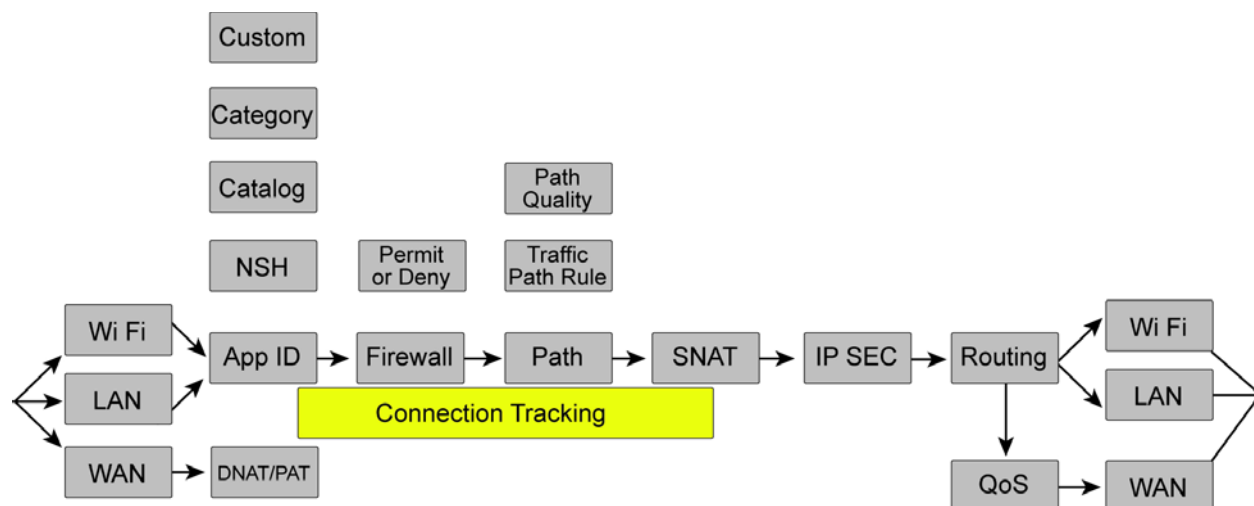
Since Overlay paths are always preferred over the underlay path for equal prefix, the traffic will be routed on the AutoVPN tunnel between Branch C and Branch D.

If a host in Branch B wants to communicate with a server in Branch D while they are in different WANs, it is possible to establish the communication between the two.

Packet operation flow

When a packet is processed by a SteelConnect gateway, the logical order of the operation flow is from left to right as shown in [Figure 2-5](#). Connection tracking, highlighted in yellow, maintains the state for packets of a flow.

Figure 2-5. Packet operation flow



Packets arrive to the gateway on a Wi Fi, LAN, or Ethernet interface. These interfaces differ slightly in terms of how traffic is handled. LAN interfaces on some gateway hardware have a switch inside and traffic between interfaces in the same zone on a LAN interface will be switched at Layer 2, bypassing the features described.

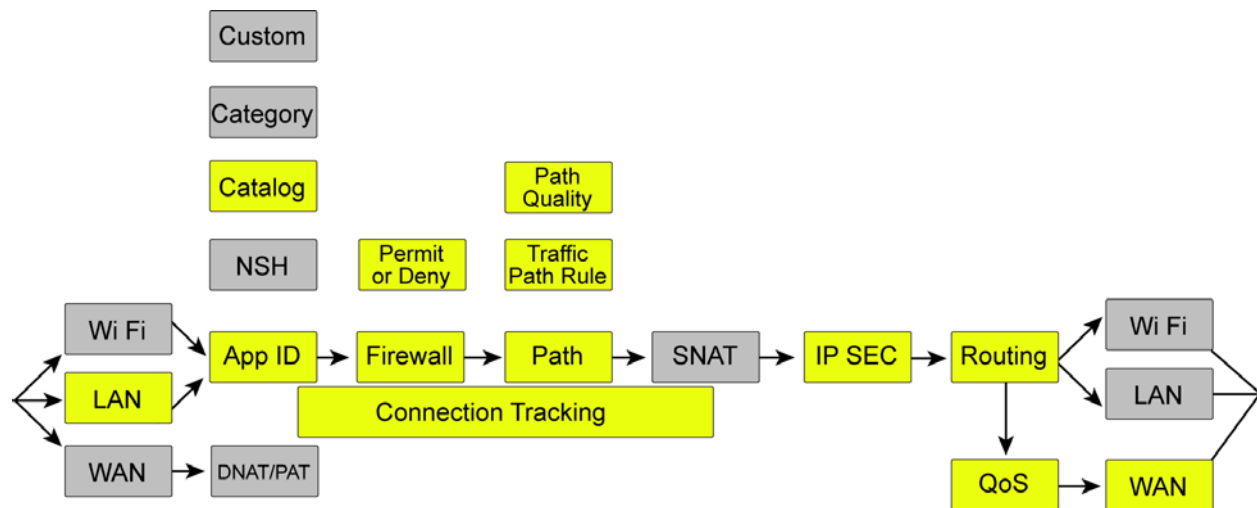
After traffic arrives, SteelConnect classifies it based on the application and tracks this classification for all packets in a flow. The firewall is invoked and will act on packet headers or, if needed, wait until the application is identified. Following the firewall, SteelConnect might select a path for the overlay based on a traffic path rule and quality or an underlay path if there is an LPM match. After it selects the path, it performs any source NAT and the packet might be encrypted. Following the encryption process, the packet is handed over to an interface and the per-interface routing table is consulted. In the case of packets going out a WAN interface (uplink), the packet will have QoS applied, if appropriate, or will retain its downstream marking.

The following examples show additional details of how a packet is processed by the SteelConnect gateway as it traverses an SD-WAN network. Keep in mind that SteelConnect can skip some steps in the left-to-right flow. In these examples, the functions that are in use are highlighted in yellow and the functions that are skipped are highlighted in gray.

Example: traffic sent from site to site

In this topology, traffic is sent from one site to another over the WAN. **Figure 2-6** shows how packets are processed at the originating SteelConnect appliance (source) before the traffic is sent over the WAN to the remote gateway.

Figure 2-6. Traffic sent from site to site



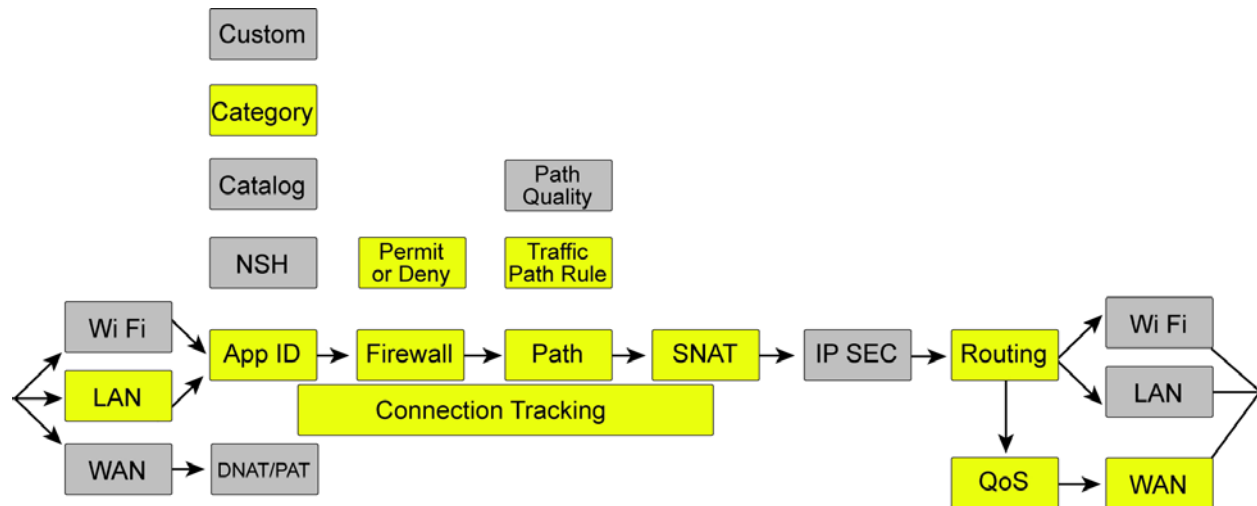
- Traffic arrives on an interface, is not malformed, and passes header validation checks.
- Application ID is checked and the application catalog is used.
- Connection tracking maintains the state for packets of a flow.
- The firewall determines that traffic from the local zone to the remote zone for this application is permitted.
- Path selection can use the organizational default when a zone is a member of more than one WAN but, in this case, a traffic path rule is used to steer the application. In SCM terms, an organization is a company representing an end customer. The traffic path rule has path quality override configured and the path decision is overridden due to quality (latency, loss, and/or jitter).
- Traffic is encrypted using Encapsulating Security Payload (ESP) protocol and encapsulated in a UDP header with a destination for the remote gateway uplink IP address connecting to the WAN that has the better quality metrics.
- Routing determines the next hop out the uplink interface.
- QoS performs scheduling using CAKE to minimize delay.
- Traffic is sent out the WAN interface.

Note: This flow also includes traffic to zones not directly connected to the gateway at Layer 3, referred to as a *third-party route*.

Example: traffic to the internet

In this topology, traffic is destined outbound to the internet.

Figure 2-7. Traffic outbound to the internet

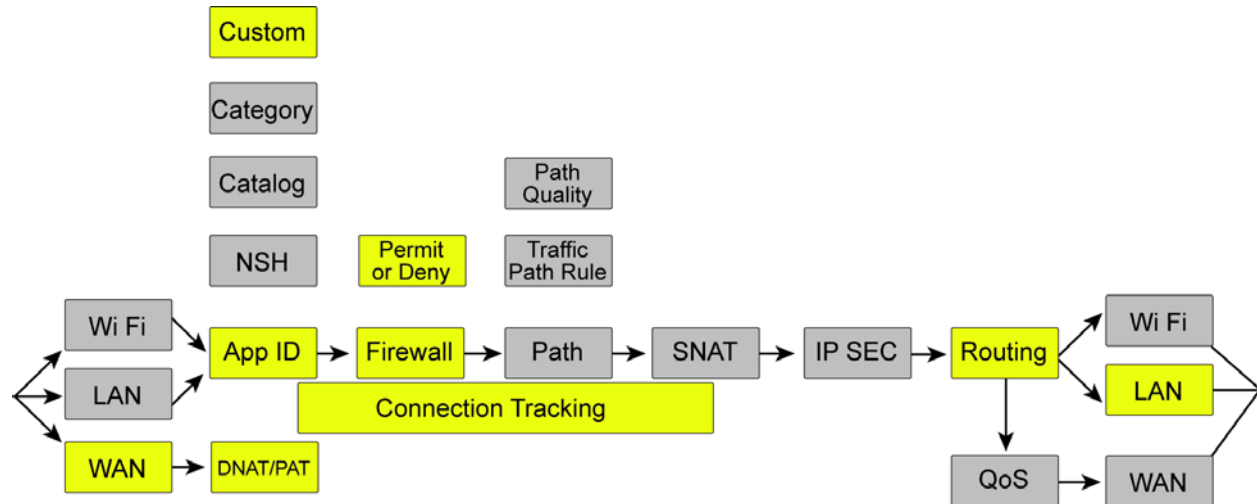


- Traffic arrives on a LAN interface, is not malformed, and passes header validation checks.
- Application ID is learned from the catalog and belongs to a category that is used for determination in the rules.
- Connection tracking maintains the state for packets of a flow.
- The firewall determines that traffic from the local zone to the internet for this application is permitted.
- Path selection uses a traffic path rule to send this application directly to the internet overriding the organizational default for internet egress.
- Source IP should have Network Address Translation.
- Routing determines that the destination is to the internet.
- QoS performs scheduling using CAKE to minimize delay.
- Traffic is sent out the WAN interface (uplink).

Example: inbound traffic

This topology shows how a packet is processed by the SteelConnect gateway when it arrives from the WAN. This topology is for all traffic destined to a server in the local network on the site arriving from the internet. A common use case is an inbound connection to the web server hosted at a site.

Figure 2-8. Traffic arriving inbound from the WAN



- Traffic arrives on the WAN interface for a destination IP address of the uplink and the MAC address of the uplink. The packet passes header validation checks.
- Destination NAT is performed for traffic to reach the internal device.
- The application ID is a custom application created for any traffic to the destination that is a defined device. In SCM, any inbound traffic has to be defined to a device using a custom application unless the WAN is configured as trusted.
- Connection tracking maintains the state for packets of a flow.
- The firewall determines that traffic from the internet to a device in a local zone is permitted and that the destination IP should have NAT to the private IP address.
- Routing determines the destination to the local device.

SteelConnect and Network Security

Providing secure internet access, while preventing malicious content from entering an organization, is critical to maintaining employee productivity and defending the network against threats.

The SteelConnect gateway provides these security capabilities:

- **Perimeter firewall** - Easy-to-configure inbound/outbound firewall rules build a secure wall to control access into and out of the different segments of the internet edge. This functionality (along with a suite of other services, such as network address translation (NAT) and policy-based network zoning) provides effective controls to mitigate the next step of a network intrusion and limits further movement across the network or propagation of a threat. SteelConnect logging capabilities support troubleshooting and policy-compliance auditing.
- **VPN concentrator** - SteelConnect automatically forms secure IPsec VPN tunnels with AES-256 encryption between your sites. IPsec VPN is configurable in full-mesh or hub-and-spoke modes from the centralized SteelConnect Manager.

This topic describes how SteelConnect can act as a robust perimeter firewall for your organization and addresses some common security use cases in the enterprise, including how to enforce firewall policies at branches for access control to zones and the internet.

SteelConnect as a perimeter firewall

Defending the network edge is an essential component for detecting and protecting the network from unwanted traffic, potentially malicious content, and intrusion attempts, and a perimeter firewall is the primary defense in a private network. The perimeter firewall identifies and logs these threats for the network administrator.

The perimeter firewall also blocks incoming network traffic from accessing internal networks and hosts and blocks outgoing traffic from accessing undesirable external networks and hosts. For example, organizations might block access to gaming sites or other social media sites. As such, a perimeter firewall can be considered as having an internal and external interface.

With SteelConnect's networking and security integration, organizations can use internet connections while mitigating security risks. Also, with a centralized cloud-based console, business-intent policies are easy to configure and enforce for your entire organization.

SteelConnect security features

Most enterprises backhaul data from their branch locations to their data centers for thorough inspection of the data going to and from the internet. Alternatively, steering software as a service (SaaS) traffic directly to the internet enables enterprises to offload their WAN network at a lower cost and reduced latency than backhauling traffic to the remote data center. Both backhaul and direct-to-net architectures require a network security stack to protect against malware, phishing, and other security threats. They also provide the ability to apply consistent security policies for all users, no matter where they connect.

Riverbed works with companies that develop specialized security functions to provide a holistic defense. Riverbed has partnered with Palo Alto Networks, Zscaler, and others to provide additional advanced security functions to organizations ensuring a stringent security posture at the branch offices as well as the data centers.

The following table summarizes SteelConnect firewall capabilities in addition to the advanced features provided by partnering with other security vendors.

	Feature	SteelConnect only	SteelConnect + security partner
Perimeter firewall	Application-based stateful firewall	✓	✓
	Granular allow/deny filtering rules	✓	✓
	Robust/hardened gateway	✓	✓
	VPN capabilities	✓	✓
	NAT capabilities	✓	✓
	Logging	✓	✓
Advanced security functions	Intrusion detection/prevention		✓
	Content scrubbing/malware detection		✓
	Acceptable use policy enforcement		✓
	Other InfoSec compliance requirements		✓

Application-defined security with granular filtering rules

Application definitions are a way to attach a business relevancy to all traffic that goes through your network. The Riverbed Deep Packet Inspection (DPI) engine can identify and classify business-critical and nonessential network traffic beginning with the first packet of the traffic flow.

With SteelConnect, you can build a secure firewall policy that regulates who has access and to what resources. Security policies can apply to the entire network, such as a single security policy to turn zone access on and off. You can also make the policy more granular to accommodate specific security needs. For example, you can create firewalled zones that require specific user permission to use specific applications. You can also enable access control and data protection for registered and unregistered mobile phones, tablets, and laptops.

The outbound and internal rules specify a source, a target, and an action. The source can be either a special catch-all selection (such as all registered users) or a custom selection of user groups, device groups, individual users, individual devices, or policy tags. We recommend that you base the outbound and internal rules on user groups and device groups, and then make exceptions using policy tags.

Policy controls are built on two types of rules:

- **Allow/deny firewall rules** - Define the policy for internal users and devices accessing internal or external applications.
- **Inbound access control rules** - Define the policy for external (internet) access to internal applications. Inbound rules offer optional support for NAT, port translations, and an external host whitelist.

You can easily create, review, enable, and disable these policy controls for the whole organization using SCM as opposed to creating tedious access-control lists and mapping them to interfaces on individual routers at each site.

Figure 3-1. Policy controls

Position	Name	Enable	Users / Source	Action	Applications / Target
#01	Outbound_rule_1	On OFF	Hansang Bae	✗	Facebook (1 st packet) Google-Video YouTube (1 st packet) Google-Talk-Video
#02	Outbound_rule_2	On OFF	All users and devices (excluding guests)	✓	Any (Unrestricted access)
#03	Outbound_rule_5	On OFF	CEO Anna Doe Karen Johnson Vivek Ganti	✓	LA → LAN [1000] LA → VoIP [1007] London → LAN [1006] Manaus → LAN [1003]
#04	Outbound_rule_7	On OFF	Guests	✗	BitDefender Bittorrent
#05	Outbound_rule_8	On OFF	Chad Smith Vivek Ganti Win Server Win Server Seoul	✗	Business Time Waste Entertainment / Games Insecure Content / Illegal Activity Shopping Nudity / Adult Content Pornography Social Networking HRApp NATApp Bittorrent Apple iTunes (1 st packet)

For more details, see “Enabling Security Using Rules” in the *SteelConnect Manager User Guide*.

WAN hardening and network zoning

Network hardening is a process to address network security vulnerabilities by implementing software patches, introducing new security systems, and adopting better configuration and operation policies.

You can deploy SteelConnect gateways at the internet edge to harden the enterprise network at branches. No unauthorized outside traffic is allowed to enter the branch unless the administrator permits it by configuring granular inbound and outbound firewall rules.

By default, internet uplinks are hardened and all traffic is denied.

You can create inbound rules to allow only specific traffic on specific TCP/UDP ports to or from specific servers both inside and outside the local branch network. Secure IPsec VPN tunnels allow for data to be encrypted and only secure data can transfer between branches over the internet.

SteelConnect zones help segment the network, and they help maintain control and organization of traffic. The use of zones within NAT rules gives an organization more flexibility and simplicity with its firewall management. In a world where Internet of Things (IoT) devices such as HVAC systems are used as a platform of attack to steal critical customer data, many organizations are taking steps toward network segmentation, such as building multiple subnetworks (overlays) within a shared network (underlay).

SteelConnect makes WAN segmentation radically simpler. You simply need to define a policy describing the underlying network. The policy is then distributed across the nodes in the SD-WAN network, which creates the multipoint (IPsec) tunnels linking these nodes defined in the policy. SteelConnect's built-in software-defined intelligence also allows for automatic firmware updates to gateways that keep security patches and hotfixes current.

Note: As of version 2.10, SteelConnect does not support Layer 3 segmentation and isolation through virtual routing and forwarding (VRF).

VPN capabilities

AutoVPN

AutoVPN is a SteelConnect feature that automatically connects multiple sites with IPsec VPNs, without tedious manual configuration providing a fast way to create a secure and resilient VPN backbone between all your sites.

AutoVPN can be split into two general operating modes:

- **RouteVPN** - A Layer 3 IPsec VPN between internal networks across physical sites.
- **SwitchVPN** - A bridged Layer 2 IPsec VPN that extends zones across multiple sites.

For more details about AutoVPN, see “Secure overlay tunnels” in the *SteelConnect Manager User Guide*.

IPsec VPN with third-party vendor

Some deployments require a SteelConnect gateway connection to an IPsec VPN built by a third-party vendor. You can connect to a third-party VPN using classic VPN. Classic VPN creates a manual VPN tunnel using the standard IPsec IKEv1 or IKEv2 protocol.

Classic VPN configurations can classify traffic based on TCP/UDP port numbers, providing a more granular approach to traffic steering. Classic VPN is an easy and flexible method to use when:

- connecting to a third-party IPsec VPN gateway, such as a firewall or a Unified Threat Management (UTM) appliance.

- migrating from an existing VPN solution to SteelConnect RouteVPN. You can even use the IP subnets of the remote networks and rules.
- integrating sites with overlapping IPv4 addresses, using one-to-one NAT configuration. When connecting networks through VPN that use the same IP addresses on both sides, it's impossible to create a simple IPsec tunnel, as routing through the tunnel doesn't work. Classic VPN uses an integrated NAT layer, in which you can map an overlapping network one-to-one into a virtual network. This means that you can communicate with the remote location using the virtual NAT network and prior to entering the tunnel, the system transparently replaces IPv4 addresses with the matching address from the remote side, allowing both networks to remain unchanged.

For more details about building IPsec tunnels with a third-party vendor, see “Connecting to a third-party VPN” in the *SteelConnect Manager User Guide*.

Important: By default, the outbound rules for a firewall let users and guests access the internet only. Communication between sites is disallowed by default through an explicit “deny” rule.

Figure 3-2 shows the default firewall rules with guest access to the internet.

Figure 3-2. Default firewall rules with guest access to the internet

Position	Name	Enable	Users / Source	Action	Applications / Target
#01	Default_outbound_rule1	On Off	Guests	✓	Internet Access
#02	Default_outbound_rule2	On Off	All users and devices (excluding guests)	✓	Internet Access

To set up access between sites, you can update the second default rule and allow users (excluding guests) to access all sites.

Figure 3-3. Allowing guest access to sites

Edit rule

Position: >> No position change <<

Name: Default_outbound_rule2

Site scope: Apply rule in all sites

Users / Source: All (excluding guests)

Action: ☒ Allow ☐ Deny

Applications / Target: Any

Cancel Submit

You can also define more granular rules depending on security constraints.

Figure 3-4. Defining access details

The screenshot shows a configuration window for defining access details. The fields are as follows:

- Position:** A dropdown menu showing ">> No position change <<".
- Name:** A text input field containing "Default_outbound_rule2".
- Site scope:** A dropdown menu showing "Apply rule in all sites".
- Users / Source:** A dropdown menu showing "Selected zones". Below it is a list of selected zones:

BranchB	LAN [1001]	[trash icon]
BranchD	LAN [1004]	[trash icon]
DatacenterG	LAN [1002]	[trash icon]
DatacenterH	LAN [1003]	[trash icon]
HQ	LAN [1000]	[trash icon]
- Action:** Two buttons: "Allow" (green) and "Deny" (grey).
- Applications / Target:** A dropdown menu showing "Selected zones". Below it is a list of selected zones:

BranchB	LAN [1001]	[trash icon]
BranchD	LAN [1004]	[trash icon]
DatacenterG	LAN [1002]	[trash icon]
DatacenterH	LAN [1003]	[trash icon]
HQ	LAN [1000]	[trash icon]

At the bottom right, there are "Cancel" and "Submit" buttons.

Inbound NAT capabilities

Network address translation (NAT) is the mapping network address to alternate and secure addressing. NAT helps improve security and decrease the number of IP addresses an organization needs.

SteelConnect gateway offers two inbound NAT modes:

- **Dynamic NAT (DNAT)** - Translates the source address for all outbound connections from the private address of the source host to a public address (typically the gateway public IP). This mode replaces the target address. It requires the target system to use the SteelConnect gateway as its default gateway.
- **Full NAT** - Replaces the source address of the incoming connection with the internal gateway address. This mode makes sure the return traffic is sent back to the SteelConnect gateway, even if the system uses a different default gateway.

You can create inbound rules to allow only specific connections to reach your internal NAT servers. While doing so, you are given the option of configuring No NAT, DNAT, or Full NAT as shown in the Mode field.

Figure 3-5. NAT modes

The screenshot shows a 'Create inbound rule' dialog box with the following fields and options:

- Name:** Inbound_rule
- Application:** Please select an internal application
- Uplinks:** Please select one or more uplinks
- Custom WAN IP:** Custom WAN IP (optional, only for single static uplink)
- Mode:** DNAT
- NAT Port mappings:** Please select an application above
- Reflection:** On (selected), Off
- Use external host whitelist:** On (selected), Off

Buttons: Cancel, Submit

Visibility and logging

Visibility is crucial to security. You can't protect what you don't know about.

With SteelConnect Manager, you have a centrally managed system of products designed from the ground up to work together and you have total visibility into your network. You can use the Visibility page to read logs, view a history of the DHCP server IP address assignments, see if and where traffic was blocked, and generate user reports. You can export logs to syslog servers.

SteelCentral Insights for SteelConnect (Insights) also provides visibility into your SD-WAN network. Insights collects data from SteelConnect Manager securely using SteelFlow (a Riverbed-specific form of NetFlow) and REST APIs. Insights uses this data to provide visibility into usage and availability about your overall network, specific sites, servers, applications, and users. With this information, you can make informed policy and deployment decisions, monitor and troubleshoot performance issues, and plan for changes.

For more information see the *SteelConnect Manager User Guide*.

LAN Topologies

This topic explains how to configure the LAN side of the SteelConnect gateway. It includes these possible deployment scenarios:

- [“Spanning tree on LAN ports” on page 33](#)
- [“Layer 2: Access port” on page 34](#)
- [“Layer 2: Trunk port” on page 37](#)
- [“Layer 2: High availability” on page 38](#)
- [“Layer 3: Switch behind a SteelConnect gateway with static routing” on page 40](#)
- [“Layer 3: Switch behind a SteelConnect gateway with LAN-side OSPF” on page 44](#)
- [“Layer 3: High availability” on page 47](#)
- [“Integration with firewalls” on page 49](#)
- [“Integration with Riverbed SteelHead” on page 51](#)

This list of possible scenarios highlights some of the most common deployments, but is not an exhaustive list.

To learn more about how to create and configure sites and VLANs (zones) for your topology, see the *SteelConnect Manager User Guide*.

Spanning tree on LAN ports

Spanning Tree Protocol (STP) is a Layer 2 protocol that passes data back and forth to find out how the switches and gateways are organized on the network and uses this information to create a logical spanning tree. Each site configures a different spanning tree.

STP is active by default on the SteelConnect gateway LAN ports and it prevents network malfunction by blocking ports that cause loops in redundant network paths.

Important: SDI-1030 and VGW on ESX do not support Spanning Tree.

STP is disabled on branch gateways configured for high availability. Due to the lack of STP, we recommend attaching only one switch to one gateway and avoiding mesh topologies with HA deployments.

The original version of Spanning Tree IEEE 802.1D takes between 30 to 50 seconds to converge in a network. Rapid Spanning Tree Protocol (RSTP) defined in 802.1w is an enhanced version with convergence time of 6 seconds (3x2 hello interval or even lower).

SteelConnect gateways and switches implement Multiple Spanning Tree (MST) protocol defined in IEEE 802.1s. MST uses the best features of PVST+ and RSTP, and MST can run a single instance of Spanning Tree Protocol for a group of VLANs. So in a network of 1000 VLANs you can group the first 500 VLANs into MST group 1 and the last 500 VLANs in MST group 2 and reduce 1000 VLAN instances in RPVST+ to just two logical instances.

All spanning tree operations include these four primary steps:

1. Determining a root bridge.
2. Selecting a root port.
3. Selecting designated ports.
4. Blocking ports with loops.

The election of a root bridge is the most crucial step and dictates the movement of Layer 2 traffic. A device with the lowest bridge priority becomes a root bridge in the network. If there is a tie with bridge priority, then the device with lowest MAC address becomes the root bridge. Unlike other vendors, we do not employ default bridge priority of 32768 because SteelConnect access points have the lowest MAC addresses of all Riverbed products.

All SteelConnect devices start with MAC address 6C:98:EB.

The following are the default priorities:

- SDI-S48 - 12288
- SDI-S24 - 16384
- SDI-S12 - 20480
- SDI-330 - 24576
- SDI-130 - 28672
- SDI-AP5/5r - 36864
- SDI-AP3 - 40960

Although access points currently have higher priority, do not change the value or decrease the default priority, which would result in access points becoming a root bridge in your network.

For example, if you have a combination of S48 and S24 switches and a gateway, then the S48 switch will assume the role of root bridge.

When SteelConnect appliances are deployed in a multivendor environment where a root bridge is already elected, we recommend using Root Guard, BPDU Guard, or Loop Guard features or reducing the existing root bridge priority to less than 12288 to avoid or reduce downtime.

Layer 2: Access port

This deployment features switches with access port connections configured in a single zone.

This topology can have a SteelConnect gateway acting as a switch, a single switch, or a pair of switches.

Figure 4-1. Layer 2 deployment with access port connection





Figure 4-2. Layer 2 deployments with dual switches and access port connections



In these topologies, Spanning Tree is active by default on the SteelConnect gateway LAN ports and prevents network malfunction by blocking ports that cause loops.

In this topology, you configure the switch uplink port as an access port.

The following SteelConnect gateways support this topology.


SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓		

You can deploy the SteelConnect virtual gateway (SDI-VGW) on a universal customer premises equipment (uCPE) or a Riverbed SteelFusion appliance as long as physical ports are attached to the ports of the virtual appliance.

Layer 2: Trunk port

This deployment features switches with trunk port connections for multizone environments.

Figure 4-3. Layer 2 deployments with trunk port



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

Important: This topology does not support native VLAN. All frames must be tagged, and untagged frames are ignored/dropped.

To configure the switch and SCM to support this topology

1. Physically connect the Layer 2 switch to a LAN port on the SteelConnect gateway.
2. Configure the connected port as a trunk/ 802.1q tagged port.

In SCM, choose Appliances > Ports, select a LAN port and select Multizone (VLAN tagged) from the drop-down Port mode menu.

Figure 4-4. LAN port mode

3. Configure the core switch with a trunk port to connect to the SteelConnect gateway.

Here is a sample configuration:

```
interface GigabitEthernet1/2
description ToSteelConnectGateway
switchport trunk encapsulation dot1q
switchport trunk allowed vlan all
switchport mode trunk
vlan dot1q tag native
```

Layer 2: High availability

High availability on the LAN provides network redundancy and reliability and maintains uninterrupted service in the event of a power, hardware, software, or WAN uplink failure. Gateways in an HA pair must be the same model.

Figure 4-5. LAN high-availability topology




Figure 4-5 shows the two supported configurations. The interconnections between the SteelConnect HA pair and the Layer 2 switch can be configured as access mode (singlezone) or trunk/802.1q (multizone).

The Spanning Tree Protocol (STP) is disabled on branch gateways configured for high availability. Without STP, we recommend you attach only one switch to a gateway and avoid mesh topologies.

The HA pair uses the Dedicated Control Port to communicate with each other, and that port serves as the VRRP link between the two gateways.

Note: To avoid traffic loops, do not use the dedicated control port to transmit data.

For model SDI-1030 gateways, one port must be configured as the Dedicated Control Port by assigning the Control Port role.

Figure 4-6. Setting the Control Port

The screenshot shows the configuration interface for an SDI-1030 gateway. The 'HA' tab is selected under the 'High availability settings' section. The configuration includes:

- High availability partner appliance:** Zurich -> SDI-1030 [Shadow]
- HA Control Link:** Dedicated Control Port
- Control Port:** Zurich -> SDI-1030 [Shadow] -> 9
- Tracked Uplinks:** Please select one or more uplinks
- Tracked Ports:** Please select one or more ports

Buttons for 'Cancel' and 'Submit' are at the bottom right.

You can use tracked ports to trigger failover if the network goes down on the LAN side of the gateway (for example, if a switch fails) to avoid a black hole situation.

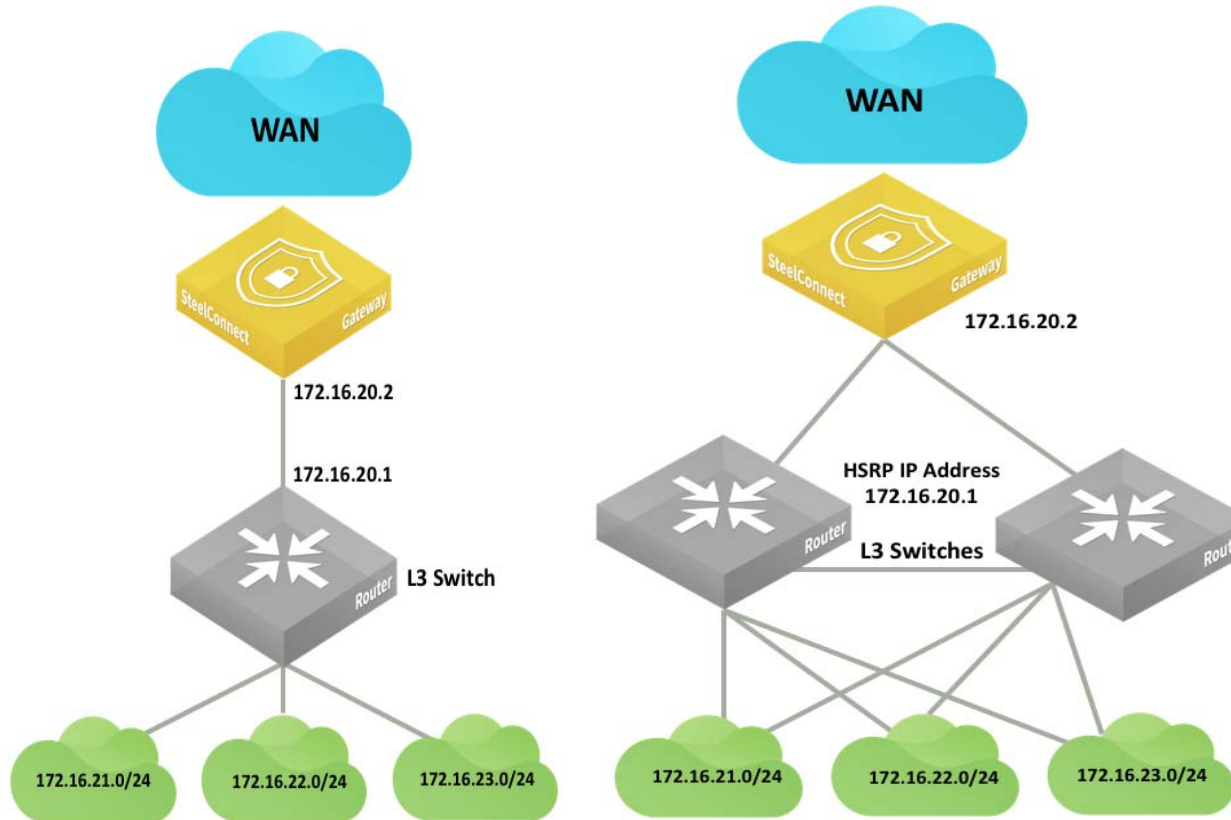
The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

Layer 3: Switch behind a SteelConnect gateway with static routing

A typical use case for Enterprise networks uses a core Layer 3 switch (or a pair of switches) handling inter-VLAN routing for the local network.

Figure 4-7. LAN topology with Layer 3 switch behind SteelConnect gateway with static routing



You can configure static routes:

- on the core switches to route traffic to the WAN through the SteelConnect gateway.
- on the SteelConnect gateway to identify the local subnets and be able to advertise them on the WAN.

Note: The IP addresses and subnets are examples only. You can configure addresses appropriate for your network.

The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

To configure the switch and SCM to support this topology

1. Create a zone for the LAN interconnection or update the default LAN zone assigned to the site.

Figure 4-8. Creating a zone

The 'Create Zone' window shows the following configuration:

- Site:** Zurich (Zurich)
- Name:** LAN_Interco
- Guest Zone:** On (selected), Off (disabled)
- IPv4 Network:** 172.16.20.0/29
- IPv4 Gateway:** 172.16.20.2
- Default Gateway configuration:** Automatic (disabled), Manual (selected)
- SteelConnect Gateway Appliance:** :: None ::
- VLAN Tag:** Optional (can be auto-assigned).

Buttons: Cancel, Submit

2. Physically connect the Layer 2 switch to a LAN port on the SteelConnect gateway.
3. For the newly configured zone, configure the connected port under Port mode to be Singlezone.

Figure 4-9. Setting the port mode

The 'Mode' window shows the following configuration:

- Port mode:** Singlezone
- Zone:** Zurich > LAN_Interco
- Patch Label:** Patch label (limited to 16 character)

Buttons: Cancel, Submit

4. Configure the switch port.

Here is a sample configuration:

```
Switch(config)# interface giga 0/1
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.16.20.1 255.255.255.248
Switch(config-if)# no shutdown
Switch(config)# ip route 0.0.0.0 0.0.0.0 172.16.20.2
```

5. Test the connectivity between the two devices.

6. On SCM, for each local subnet, create a zone with a manual default gateway configuration.

Figure 4-10. Creating a zone with a manual default gateway configuration

Create Zone

Site ⓘ Zurich (Zurich)

Name ⓘ Users

Guest Zone ⓘ **On** **Off**

IPv4 Network ⓘ 172.16.21.0/24

IPv4 Gateway ⓘ 172.16.21.1

Default Gateway configuration Automatic **Manual**

SteelConnect Gateway Appliance ⓘ :: None ::

VLAN Tag ⓘ Optional (can be auto-assigned).

Cancel Submit

7. For the newly created zones, edit the configuration under the Gateways tab and delete the Gateway assignment.

Figure 4-11. Deleting the gateway assignment

Users Actions X

IP **Gateways** DHCP VLAN WAN/AutoVPN Settings ADDL Networks

Discovered Networks

Automatic SteelConnect default gateway

When turning this option on, a SteelConnect gateway appliance deployed in the site will be automatically configured as the default gateway for this zone. It will then **use the default gateway IP addresses specified on the 'IP' tab**. If you want to control all gateway assignments for this zone manually, or you want to use a third-party default gateway for this zone, please turn this option off.

Default Gateway configuration Automatic **Manual**

Gateway assignments

This table shows all SteelConnect gateways that are members of this zone. You can create several memberships, also in remote sites. Every member gateway will be able to route into the zone's network. Default gateway entries that have been added automatically cannot be edited or deleted - if you want to control all gateway parameters, turn off automatic default gateway assignment and create a default gateway manually.

+ Add assignment

Type / Appliance	IPs	Flags
No gateway memberships present		

8. Select the Settings tab and create a third-party route for the gateway to reach the VLANs through the core switch.

Figure 4-12. 3rd-party gateway route

New 3rd-party gateway route


From Zone/Net ? • Zurich > LAN_Interco [1004]

IPv4 gateway ? 172.16.20.1

Layer 3: Switch behind a SteelConnect gateway with LAN-side OSPF

An alternative to the previous configuration is to leverage dynamic routing protocols such as OSPF to discover subnets behind a Layer 3 switch.

Figure 4-13. LAN topology with Layer 3 switch behind SteelConnect gateway with dynamic routing



In this deployment, you must configure OSPF on the Layer 3 switch as well as on the SteelConnect Manager so that:

- the switch advertises the LAN subnets to the SteelConnect gateway.
- the SteelConnect gateway can learn the LAN subnets and eventually redistribute them on the WAN side as well as making SteelConnect Manager aware of those subnets in that site.
- by default, the gateway advertises a default route to the Layer 3 switch.

The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

To configure SCM to support this topology

1. Create a zone for the LAN interconnection, or update the default LAN zone assigned to the site.

Figure 4-14. Creating a zone

Create Zone

Site: Zurich (Zurich)

Name: LAN_Interco

Guest Zone: ☒ On ☐ Off

IPv4 Network: 172.16.20.0/29

IPv4 Gateway: 172.16.20.2

Default Gateway configuration: ☐ Automatic ☒ Manual

SteelConnect Gateway Appliance: :: None ::

VLAN Tag: Optional (can be auto-assigned).

Cancel Submit

2. Physically connect the Layer 2 switch to a LAN port on the SteelConnect gateway.
3. For the newly configured zone, set the connected port under Port mode to be Singlezone.

Figure 4-15. Setting the port mode

Mode

Port mode: Singlezone

Zone: Zurich > LAN_Interco

Cancel Submit

Patch Label

Patch Label: Patch label (limited to 16 character)

4. Configure the switch port accordingly.

Here is a sample configuration:

```
Switch(config)# interface giga 0/1
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.16.20.1 255.255.255.248
Switch(config-if)# no shutdown

router ospf 1
network 172.16.21.0 0.0.0.255 area 0
```

```
network 172.16.22.0 0.0.0.255 area 0
network 172.16.23.0 0.0.0.255 area 0
```

5. Test connectivity between the two devices.
6. On SCM, choose Routing > OSPF and create a new OSPF network for the site.

Figure 4-16. Creating an OSPF network

Create OSPF Net...

Site Zurich (Zurich)

Name Zurich-ospf

Default Area Name Zurich-area

Default Area ID 1

Inherit Org Defaults On Off

Password Enter password

Hello Interval 10

Dead Interval 40

Priority 1

Cost 10

Cancel Submit

7. Attach the interface to learn the routes.

In this example, the LAN_Interco Zone.

Figure 4-17. Attach the interface to learn the routes

The screenshot shows a configuration window titled "Attach OSPF Interface (LAN_Interco)". It contains the following fields and controls:

- OSPF Area:** A dropdown menu showing "Zurich-area [1]".
- Inherit Area Values:** A toggle switch set to "On".
- Password:** A text input field with the placeholder "Enter password" and a visibility icon.
- Hello Interval:** A numeric input field with the value "10".
- Dead Interval:** A numeric input field with the value "40".
- Priority:** A numeric input field with the value "1".
- Cost:** A numeric input field with the value "10".
- Buttons:** "Cancel" and "Submit" buttons at the bottom right.

The SteelConnect gateway connects to its OSPF peer, in this case the Layer 3 switch on the LAN side, and starts discovering the LAN subnets.

Layer 3: High availability

Configure the HA pair and the pair's connection to the Layer 3 switch or router in the same way you would for a Layer 2 topology. See ["Layer 2: High availability" on page 38](#).

Figure 4-18 and Figure 4-19 illustrate two supported Layer 3 high-availability topologies.

Figure 4-18. High-availability pair of gateways connected to a single Layer 3 switch or router





Figure 4-19. High-availability pair of gateways connected to two Layer 3 switches or routers



The following SteelConnect gateways support these topologies.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

Integration with firewalls

LAN-side integration with firewalls is very similar to LAN topologies with Layer 3 switches. You can configure static routes or dynamic routing to reach subnets behind the firewalls.

Figure 4-20. LAN topologies with firewalls


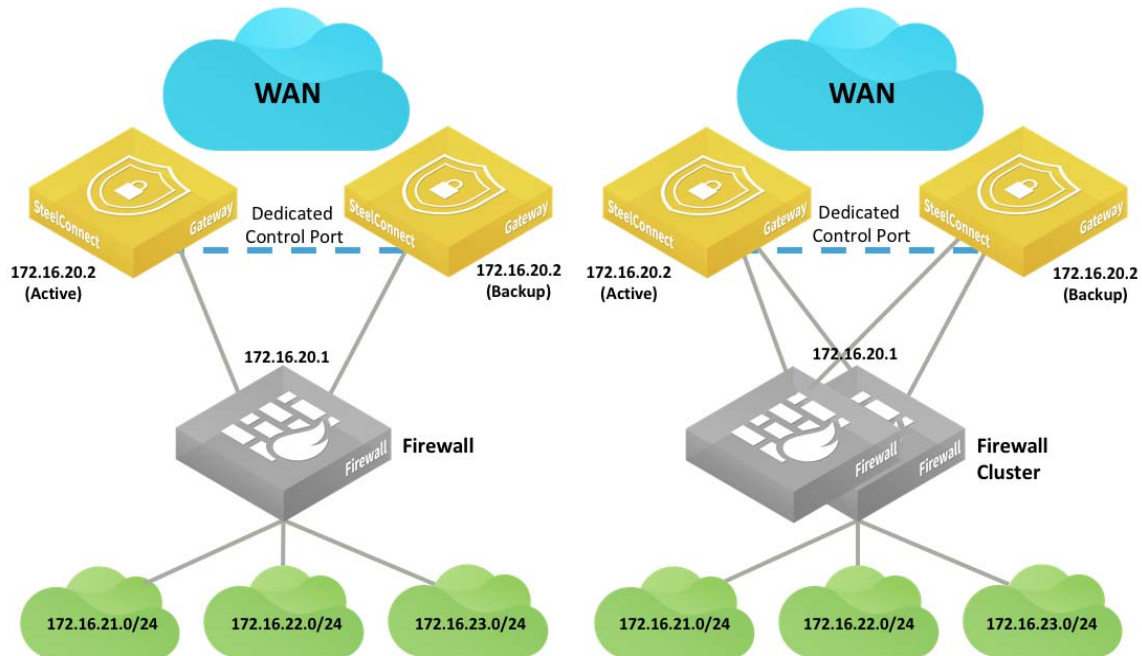


Figure 4-21. HA LAN topologies with firewalls



If the firewall applies NAT-to-LAN traffic, the SteelConnect gateway only sees the firewall IP address, so there is no need to configure static routes or OSPF on the LAN side.

Important: Double NAT can cause issues with some application traffic. You can configure gateway uplinks to skip outbound NAT to resolve double NAT issues.

Firewall clusters typically operate in active/passive mode. Each appliance has the same configuration, and if one goes down the other activates and resumes operation with the same IP addresses. The gateways do not detect a traffic loop.

Integration with Riverbed SteelHead

You can deploy SteelConnect gateways with SteelHeads for WAN optimization.

Figure 4-22. SteelHead topologies

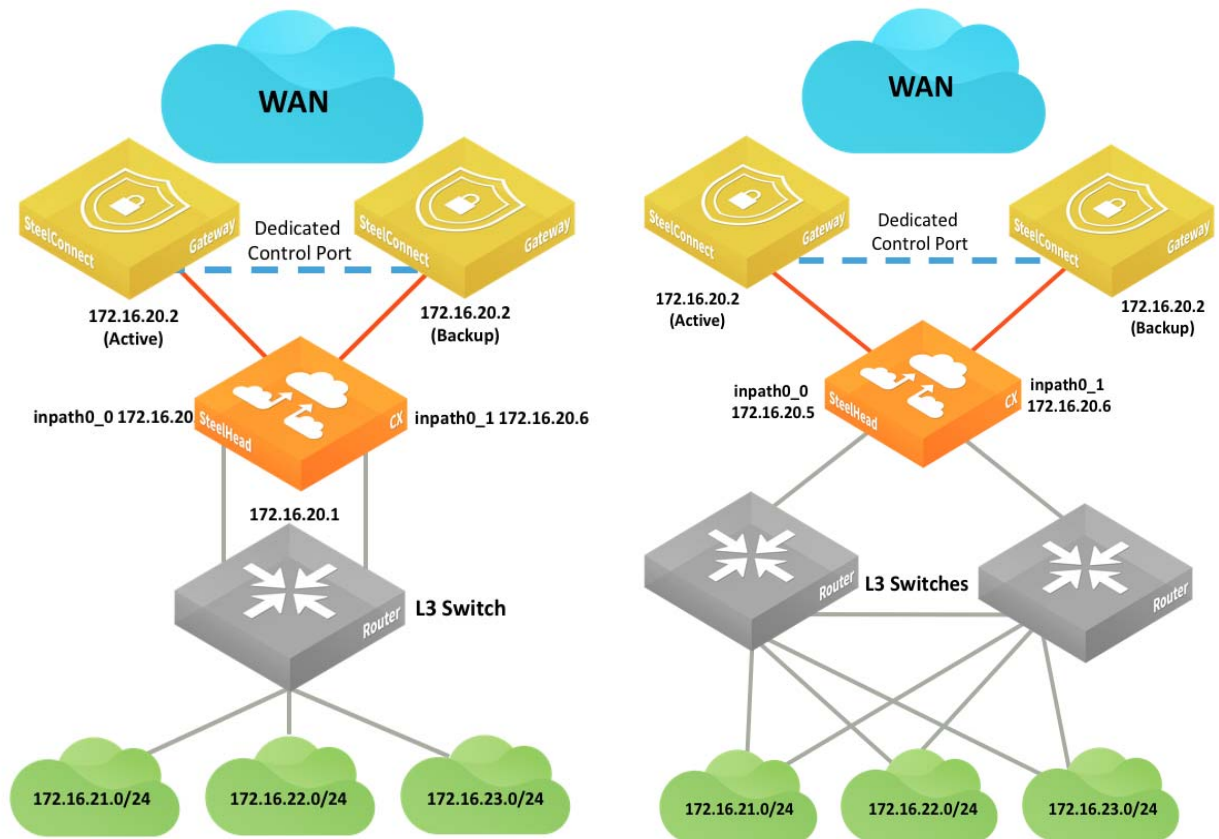



Figure 4-23. Another SteelHead topology



By default, SteelConnect appliances have the SteelHead compatibility feature on.

To enable SteelConnect compatibility on a SteelHead CX, enter this Riverbed command-line interface (CLI) command: `steelhead steel-connect compatibility enable`

See the *SteelConnect Manager User Guide* for more information about compatibility.

See the *SteelHead Deployment Guide* and *SteelHead User Guide* for more details about the SteelHead appliance configuration and deployment options.

WAN Topologies

This topic details common configurations of the WAN side of the gateway. It includes the following scenarios:

- “Important configuration information for all WAN topologies” on page 53
- “Single internet connection” on page 54
- “Multiple internet connections” on page 58
- “Secondary 3G/4G uplink” on page 63
- “Integration with private networks” on page 65
- “Integration with WAN-facing firewall” on page 78
- “Highly secured WAN” on page 79
- “WAN high availability” on page 82

This list of possible scenarios highlights some of the most common deployments, but is not an exhaustive list.

We recommend reading “SteelConnect and Network Security” on page 25 before reading this section.

Important configuration information for all WAN topologies

In all topologies, make sure that the LAN subnets (zones) are given the appropriate permissions. Each topology needs:

- Access resources on the internet as defined with the Outbound rules.

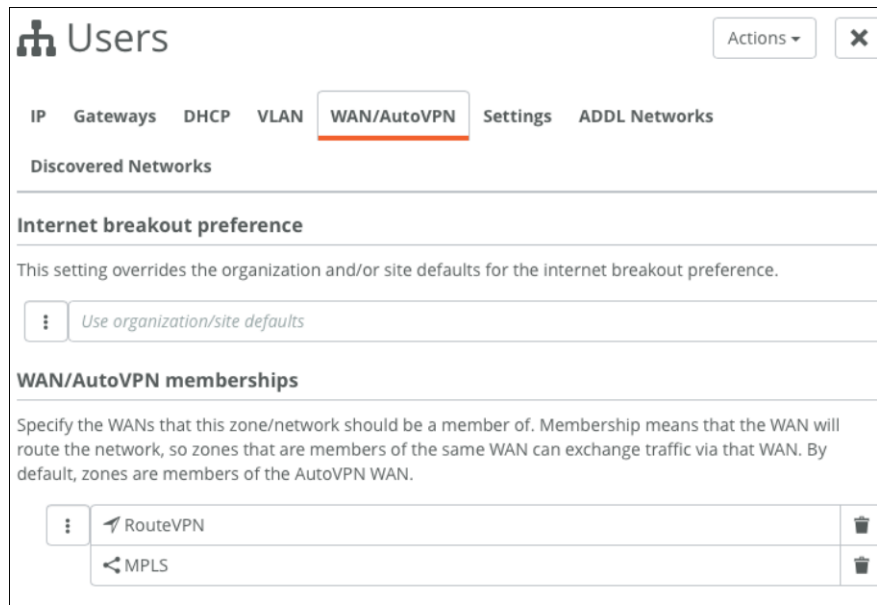
Figure 5-1. Defining access with outbound rules

Outbound / Internal rules						
Search						
Position	Name	Enable	Users / Source	Action	Applications / Target	
#01	Outbound_rule_1	On Off	Hansang Bae	✗	Facebook (1 st packet) Google-Video YouTube (1 st packet) Google-Talk-Video	
#02	Outbound_rule_2	On Off	All users and devices (excluding guests)	✓	Any (Unrestricted access)	
#03	Outbound_rule_5	On Off	CEO Anna Doe Karen Johnson Vivek Ganti	✓	LA → LAN [1000] LA → VoIP [1007] London → LAN [1006] Manaus → LAN [1003]	

- Be advertised and send traffic on different WANs.

Check the WAN/AutoVPN membership in the Zone configuration.


Figure 5-2. Specify WAN/AutoVPN memberships



Single internet connection

This internet-based WAN topology is a typical scenario for small to medium businesses that have a single broadband internet connection from their service provider.

Figure 5-3. SteelConnect gateway deployed physically on the LAN side of the customer premises equipment



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

To configure an internet-based WAN topology

1. Connect the Ethernet LAN port of the internet modem to the WAN port of the SteelConnect gateway.
2. If the modem uses a DHCP address on its LAN interface, set the uplink for the gateway in the Type field on SCM to DHCP client.

Figure 5-4. Setting DHCP for the gateway uplink

The screenshot displays the SteelConnect configuration interface. On the left, the 'Uplinks' section shows a list of sites with their respective uplink configurations. The 'SF' site is highlighted, showing it is configured as a 'DHCP client' uplink. On the right, the 'Uplink' settings for this site are shown. The 'Name' is 'Uplink', the 'WAN' is 'Internet', and the 'Type' is 'DHCP client'. The 'VLAN Tag' is set to 'VLAN tag (1-4049). Usually not required.' and the 'MTU' is 'MTU, default is 1500'. The 'Use as backup only' option is set to 'Off'. The 'Submit' button is highlighted in orange.

Site	Uplink	Type	Status
Minneapolis	Uplink [Internet]	Internet	On
NYC	Uplink [Internet]	Internet	On
Phoenix	Uplink [Internet]	Internet	On
Seattle	Uplink [Internet]	Internet	On
SF	Uplink [Internet]	DHCP client	On

Uplink Settings:

- Name: Uplink
- WAN: Internet
- Type: DHCP client
- VLAN Tag: VLAN tag (1-4049). Usually not required.
- MTU: MTU, default is 1500
- Use as backup only: Off

Buttons: Cancel, Submit

3. If the modem does not use DHCP, select Static IP and enter the IP address, address mask, and default gateway.

Figure 5-5. Setting static network details

The screenshot shows the 'Uplinks' configuration page. On the left, there is a list of uplinks for different sites: Minneapolis, NYC, Phoenix, Seattle, and SF. Each uplink is configured with a Static IP and is currently set to 'On'. The 'SF' uplink is highlighted. On the right, the detailed configuration for the selected 'Uplink' is shown. The 'Settings' tab is active, displaying fields for Name (Uplink), WAN (Internet), Type (Static IP), IPv4 Address (192.168.1.15), IPv4 Gateway (192.168.1.1), IPv6 Address (Static IPv6 address with optional netmask), IPv6 Gateway (Static IPv6 gateway for the uplink), VLAN Tag (VLAN tag (1-4049). Usually not required.), and MTU (MTU, default is 1500). There is also a 'Use as backup only' toggle set to 'Off'. At the bottom right, there are 'Cancel' and 'Submit' buttons.

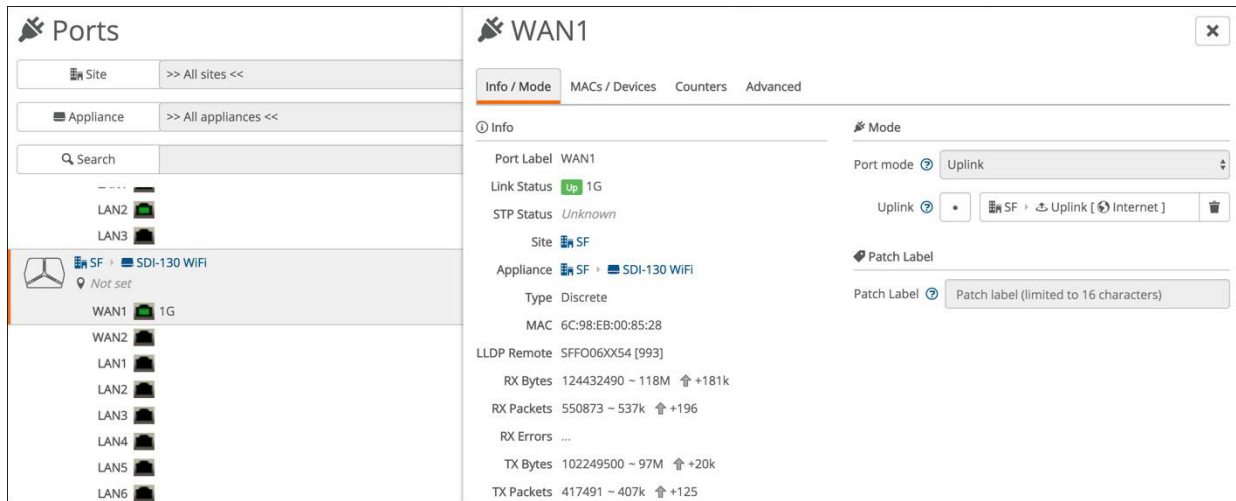
To configure a static IP address on the uplink connection for the first time, you must configure the gateway using the USB method. By default, the gateway tries to get a DHCP IP address on its uplink interface and then communicates with the core services over the internet connection to determine its configuration. To set up a static IP address on an uplink, you can set up the configuration on SCM, even if the hardware is currently not present at the related site. You'll need the serial number of the new gateway to create an offline provisioning configuration file.

To provision an appliance with a static IP address using the USB method

- Select the uplink for the site and enter the IPv4 address as shown in Step 2.
- Register the appliance under Appliances > Add Appliances > Register Hardware Appliance. Enter the serial number of the SteelConnect gateway.
- Select the new hardware appliance, click **Actions**, and select Download config. The system downloads a configuration file named with the gateway serial number. Copy the file to a FAT32 formatted USB stick. The system does not support other file system types like Linux ext2,3,4, NTFS, and so on.
- Deploy the gateway on the site and power on the appliance. Wait at least 30 seconds until the new appliance powers up correctly before plugging in the USB stick. Plug in the USB stick to restore the configuration. Because the gateway does not mount the stick during boot up, it won't import the configuration automatically. The gateway connects to SCM with the previously set up configuration.

4. Ensure that the Port mode for the WAN port of the appliance at that site is set to Uplink.

Figure 5-6. Setting the port mode




In this topology, the gateway is NATed behind the cable modem. For SteelConnect to form an AutoVPN tunnel between two sites, at least one of the gateways must be reachable through a public IP address. When the SteelConnect gateways on both sites are behind a NAT gateway, you might need to configure port forwarding for UDP port 4500 on the modem/CPE. For more information about how to configure port forwarding for some commonly deployed modems/SOHO routers go to <https://portforward.com/router.htm>.

Multiple internet connections

This internet-based WAN topology features multiple internet connections from different service providers.

Figure 5-7. Internet-based WAN topology with multiple connections



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

You can configure SDI-VGW and SDI-1030 ports with up to seven uplinks. SDI-130 and SDI-330 have two uplink ports.

With this topology, you can implement the following configurations:

- Flow-distribution load-balancing traffic over the uplinks.
You create uplinks that belong to the same internet WAN.

Figure 5-8. Multiple internet uplinks



AutoVPN is created on all uplinks.

There is no intelligence about the network conditions in terms of link utilization, packet drops, or latency. The system only monitors link availability.

- Prioritize one link over others.

You can set an AutoVPN priority for uplinks connected to the same WAN that determines the order in which they are used. For example, when you create two uplinks and set the first uplink to high priority and set the other to normal, SCM will command the SteelConnect gateway to steer traffic over the tunnel(s) formed on the uplink with the highest priority.

Figure 5-9. Setting the AutoVPN priority

The screenshot shows the SteelConnect Manager (SCM) interface for configuring ISP1. The 'AutoVPN' tab is selected. Under the 'AutoVPN priority' section, a dropdown menu is open, showing options: 'Don't use this uplink for AutoVPN', 'Low', 'Normal' (selected), and 'High'. Below this, the 'AutoVPN IPv4 target address' section is visible, with a dropdown menu set to 'External Internet IPv4 (recommended for internet uplinks)'. The interface includes 'Cancel' and 'Submit' buttons.

This priority applies to site-to-site traffic using the overlay (AutoVPN) only. Overlay tunnels are always formed. Internet traffic is load balanced over the different uplinks as long as its destined for the internet (by source address) or the same overlay VPN with equal priority (source/destination based).

- Set up a master/backup.

Set up an uplink as a backup to be activated only when the other uplinks of a same WAN are unavailable.

Figure 5-10. Setting an uplink as a backup

ISP2

Info **Settings** AutoVPN DynDNS QoS L2 Settings Networks

Settings

Name ⓘ ISP2

WAN ⓘ Internet

Type ⓘ DHCP client

VLAN Tag ⓘ VLAN tag (1-4049). Usually not required.

MTU ⓘ MTU, default is 1500

Use as backup only ⓘ **On** Off

Skip outbound NAT ⓘ **On** Off

Cancel Submit

- Define application-based traffic rules.
You can treat internet uplinks separately with traffic rules.

To define application-based traffic rules

1. Create a new WAN type on SCM with the appropriate settings.

Figure 5-11. Creating a new WAN

Create WAN

Name ⓘ Internet_ISP2

Long Name ⓘ ISP2

Internet Breakout ⓘ **On** Off

Internet NAT ⓘ **On** Off

Breakout sites ⓘ No sites, WAN provides own breakout

Encryption ⓘ **On** Off

Trusted ⓘ **On** Off

Ping check IP ⓘ Ping check IP

Transfer Networks ⓘ Transfer networks, separated with space (optional)

Cancel Submit

2. Configure the second uplink to be the same type as the new WAN created in Step 1, in that example Internet_ISP2.

Figure 5-12. Configure the second uplink

The screenshot shows the 'ISP2' configuration page. The 'Settings' tab is active. The 'Name' field is 'ISP2'. The 'WAN' dropdown is set to 'Internet_ISP2'. The 'Type' dropdown is set to 'DHCP client'. The 'VLAN Tag' field has a placeholder text 'VLAN tag (1-4049). Usually not required.' and a spinner icon. The 'MTU' field has a placeholder text 'MTU, default is 1500' and a spinner icon. There are two toggle switches: 'Use as backup only' and 'Skip outbound NAT', both with 'On' and 'Off' buttons. The 'Off' buttons are highlighted in red. At the bottom right are 'Cancel' and 'Submit' buttons.

3. Update network preferences at the organization level and/or site level.

Note: This change impacts the traffic distribution for internet.

Figure 5-13. Network preferences

The screenshot shows the 'RetailNow' 'Networking Defaults' page. The 'Networking Defaults' tab is active. Below the tabs are 'Appliance Threshold' and 'Appliance Login' sections. The 'Internet breakout preference' section contains a table with two rows. The first row has a selection icon, the number '1', an up/down arrow, and the text 'Internet'. The second row has a selection icon, the number '2', an up/down arrow, and the text 'Internet_ISP2'. The '2' option is selected.

⋮	1	▲ ▼	Internet
	2	▲ ▼	Internet_ISP2

4. Choose Rules>Traffic Path rules and create traffic path rules to reflect your business logic.

You can set path preference and you can enable path quality options.

Figure 5-14. Path quality and path preferences

Create Traffic Rule

Position: >> Top <<

Name: Traffic_rule

Site scope: London

Users / Source: All (excluding guests)

Applications / Target: Selected applications or groups
SAP

Path preference:

Order	Path	Action
1	Internet_ISP2	Remove
2	RouteVPN	Remove

Path Quality profile: Latency sensitive metrics

Rule fall through: ☒ On ☐ Off

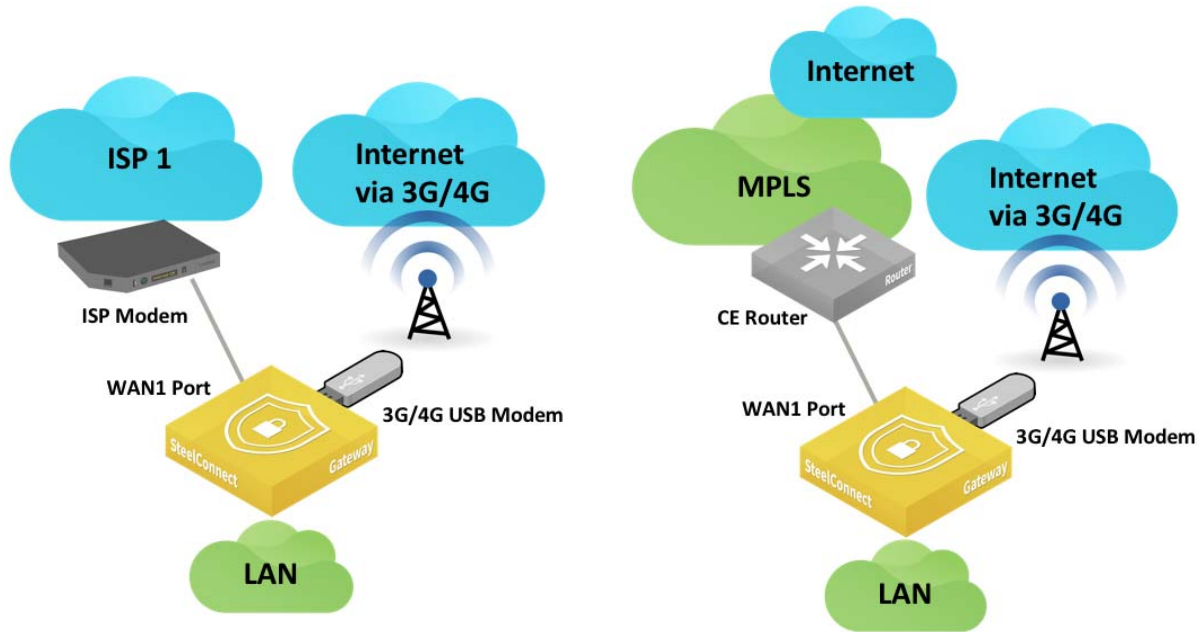
Traffic rules supersede internet and WAN preferences. For more details about traffic rules, see the *SteelConnect Manager User Guide*.

Important: RouteVPN is the name for the overlay on an internet WAN uplink (and only for that type of WAN). Since this WAN enabled encryption, an overlay tunnel is formed on top of uplinks of that type. Overlay tunnels are called *AutoVPN* on SteelConnect.

Secondary 3G/4G uplink

An alternative to a dual-homed internet topology is to rely on 3G/4G mobile connectivity as a second active or backup uplink.

Figure 5-15. WAN topology with 3G/4G uplink through USB port



You can configure SteelConnect appliances that have a USB port to support this topology.

The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
	✓	✓	✓	

To set up a topology with 3G/4G as a second or backup link

1. Configure a new uplink of type USB cellular.

Choose Network Design > Uplinks and click **New Uplink** to create an additional uplink for the site for the 4G connection.

Figure 5-16. New uplink for cellular traffic

The screenshot shows the 'Settings' tab for a new uplink configuration. The tabs at the top are Info, Settings (selected), AutoVPN, DynDNS, QoS, L2 Settings, and Networks. The 'Settings' section includes the following fields and options:

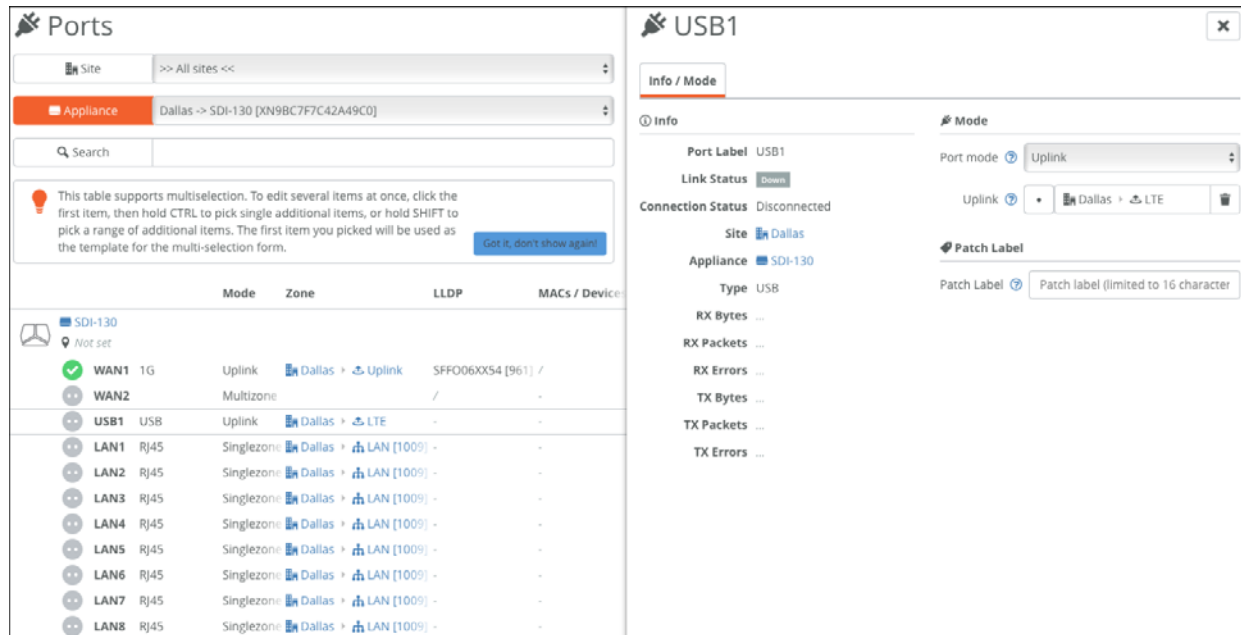
- Name:** A text input field containing 'LTE'.
- WAN:** A dropdown menu set to 'Internet'.
- Type:** A dropdown menu set to 'USB cellular'.
- Warning:** A lightbulb icon followed by the text: 'USB cellular uplinks can only be used with SteelConnect gateways that have a cellular modem plugged into the USB port. Firmly plug the modem into the USB port in order to use this uplink.' A blue button labeled 'Got it, don't show again!' is to the right.
- SIM APN:** A text input field containing 'net'.
- SIM PIN:** A text input field containing four dots, with an eye icon to the right.
- Time for PPP-reconnect:** A dropdown menu set to 'None'.
- Use as backup only:** A toggle switch with 'On' (green) and 'Off' (grey) options.
- Skip outbound NAT:** A toggle switch with 'On' (grey) and 'Off' (red) options.

At the bottom right, there are 'Cancel' and 'Submit' buttons.

The default internet WAN is used for both uplinks—terrestrial and 3G/4G. Because the 3G/4G connection is expensive, it's used as a backup only. This informs SCM not to auto load balance between the two available uplinks for the internet WAN. When enabled, the **Use as backup only** option enables this uplink only when the other uplink (default terrestrial uplink) is not working.

- The uplink is assigned to the USB port of the appliance.

Figure 5-17. Cellular uplink assigned to USB port



- If you don't want to use the 4G uplink as backup only, but want to give it a lower priority, choose Network Design > Uplinks > AutoVPN and set the priority to normal.

Note: Carrier charges might apply for cellular networks.

Integration with private networks

You can deploy SteelConnect gateways in private networks, such as MPLS. However, the gateways still require an internet breakout within the private network so the appliance can register and communicate with SteelConnect Manager.

Note: When choosing the on-premises version of SteelConnect Manager, the appliance still needs to register with the core service.

To create a WAN in this topology

1. In SCM, choose Network Design > WANs to create an additional WAN. Click **New WAN** and configure these WAN settings:
 - Turn on Internet Breakout to indicate that this WAN can be used to transport traffic destined to the internet.
 - It is not required to NAT the traffic. In some cases, such as when using proxies, it is not even recommended.
 - **Breakout sites** - Select the headquarters (HQ) site to break all traffic by default over this MPLS WAN to the internet through the headquarters site. Keep in mind that the WAN selection can be overridden for specific sites, VLANs, or applications by traffic path rules.

- Turn on **Encryption** to create an additional overlay network of VPN tunnels between internal zone-to-zone encrypted WAN traffic over noninternet uplinks. You can deploy an MPLS overlay that uses encryption and turn encryption off for another MPLS overlay.
- Turn on **Trusted** to permit all unencrypted traffic originating from a WAN to communicate into the WAN and LAN zones of the gateway. For example, this setting allows SteelConnect sites and legacy router sites within the WAN to communicate with each other. Instead of creating multiple inbound rules to permit trust, simply enable the Trusted setting. When enabled, all WAN transfer networks and eBGP learned networks are allowed to communicate into the gateway LAN zones.

Figure 5-18. Internet breakout configuration

The screenshot shows a configuration window for an MPLS overlay. The 'Name' and 'Long Name' fields are both set to 'MPLS'. The 'Internet Breakout' toggle is turned 'On' (green). The 'Internet NAT' toggle is turned 'Off' (red). The 'Breakout sites' section shows a list with '1' and 'HQ'. The 'Encryption' and 'Trusted' toggles are both turned 'On' (green). The 'Member zones/networks' section lists two entries: 'NYC > LAN [1002]' and 'HQ > LAN [1000]'. The 'Ping check IP' field contains the text 'Ping check IP'. The 'Transfer Networks' field contains the text 'Transfer networks, separated with space (optional)'. At the bottom right, there are 'Cancel' and 'Submit' buttons.

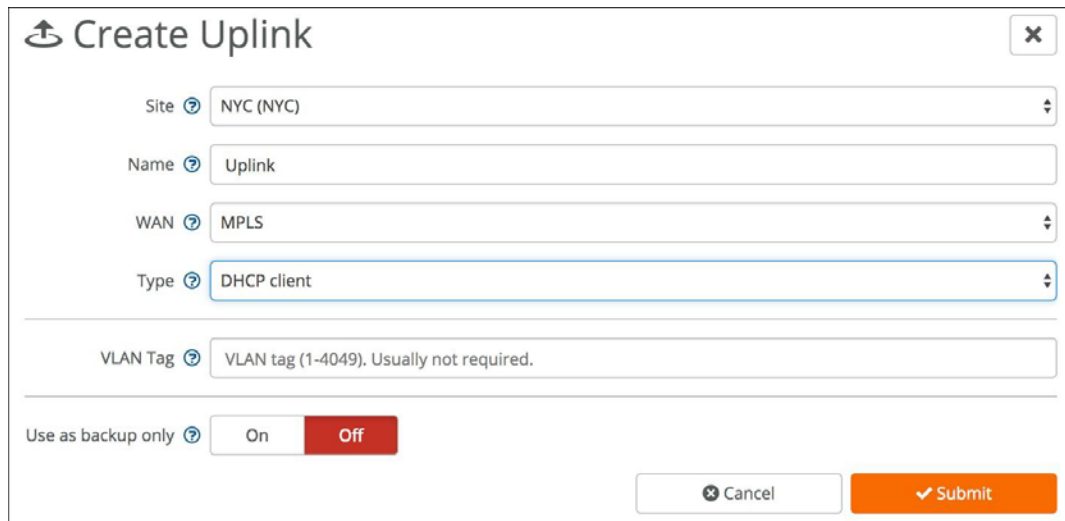
Important: In the WAN configuration, an important option for private networks is Ping Check IP. This configures an IPv4 address that checks the uplink status through ICMP ping. If left blank, it tries to reach 8.8.8.8 (a public IP address that belongs to Google). In the case of a large deployment, all gateways try to ping Google and if all traffic goes through a Central Internet Gateway, Google may throttle down or even blacklist the public IP address to prevent a potential DDOS attack. This would result in bad behavior of the MPLS uplink. Therefore, we recommended specifying an IP address of a stable equipment within the MPLS network.

Important: Proxies in the internet breakout could potentially impact the behavior of the SteelConnect gateways. While a TCP proxy would not, SSL proxy would break the strong security model we have implemented to ensure strong authentication between the SteelConnect Manager and the gateways.

2. Add an uplink for your sites.

Keep in mind that SCM creates an uplink by default for your internet WAN but not for any additional WANs. You need to create these manually.

Figure 5-19. Creating an uplink for your internet WAN



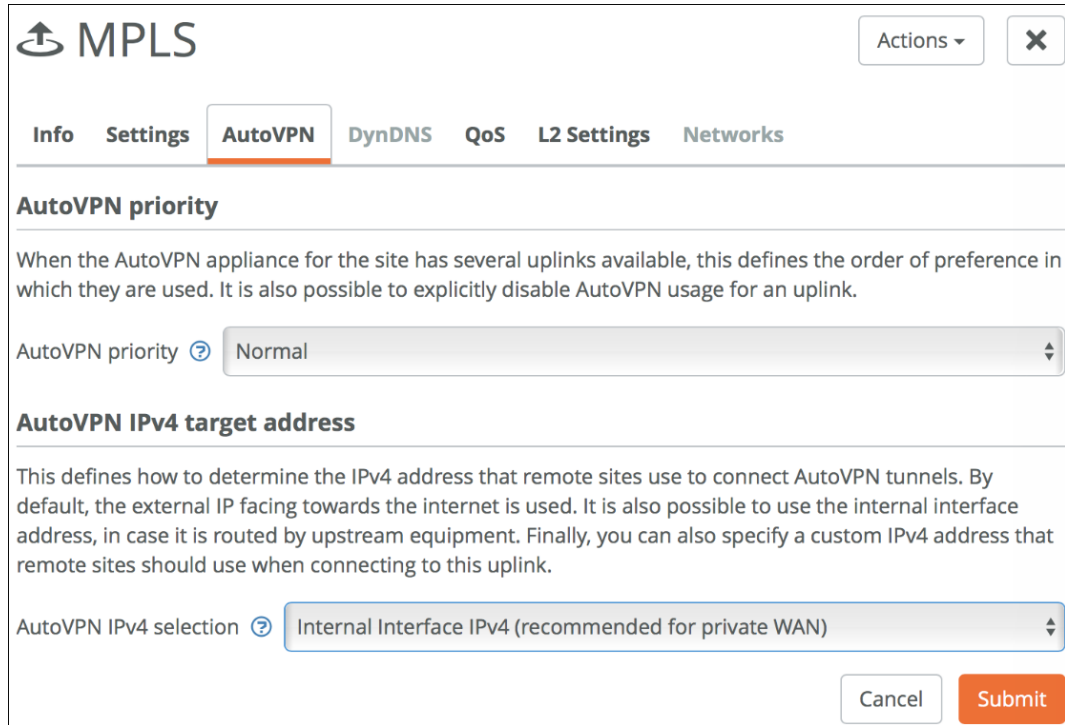
The screenshot shows a 'Create Uplink' dialog box with the following fields and options:

- Site:** NYC (NYC)
- Name:** Uplink
- WAN:** MPLS
- Type:** DHCP client
- VLAN Tag:** VLAN tag (1-4049). Usually not required.
- Use as backup only:** On (selected), Off
- Buttons:** Cancel, Submit

3. Set up AutoVPN over the MPLS network.

Configure the uplink of the gateway to propagate the internal IP address that the remote sites use as the endpoint target for the AutoVPN tunnel. By default, SCM uses the external IP address facing the internet. To configure an uplink to use an internal interface, select the uplink, select the AutoVPN tab, and choose Internal Interface IPv4 from the drop-down list.

Figure 5-20. Uplink AutoVPN settings



The screenshot shows the MPLS configuration interface with the 'AutoVPN' tab selected. The interface includes a top navigation bar with tabs for 'Info', 'Settings', 'AutoVPN', 'DynDNS', 'QoS', 'L2 Settings', and 'Networks'. The 'AutoVPN' tab is active, displaying two sections: 'AutoVPN priority' and 'AutoVPN IPv4 target address'. The 'AutoVPN priority' section has a description and a dropdown menu set to 'Normal'. The 'AutoVPN IPv4 target address' section has a description and a dropdown menu set to 'Internal Interface IPv4 (recommended for private WAN)'. At the bottom right, there are 'Cancel' and 'Submit' buttons.

MPLS Actions X

Info Settings **AutoVPN** DynDNS QoS L2 Settings Networks

AutoVPN priority

When the AutoVPN appliance for the site has several uplinks available, this defines the order of preference in which they are used. It is also possible to explicitly disable AutoVPN usage for an uplink.

AutoVPN priority ? Normal

AutoVPN IPv4 target address


This defines how to determine the IPv4 address that remote sites use to connect AutoVPN tunnels. By default, the external IP facing towards the internet is used. It is also possible to use the internal interface address, in case it is routed by upstream equipment. Finally, you can also specify a custom IPv4 address that remote sites should use when connecting to this uplink.

AutoVPN IPv4 selection ? Internal Interface IPv4 (recommended for private WAN)

Cancel Submit

Integration with MPLS CE router

Figure 5-21. WAN topology with MPLS CE router



To ensure that the zones on the LAN side of the gateway can be reached, we recommend that you configure OSPF on the MPLS (CE) router. You could use static routes, but managing the static configurations becomes very time consuming as network routing choices expand, thus OSPF is recommended instead.


A SteelConnect branch gateway supports OSPF version 2 in a broadcast network for dynamic routing. The gateway uses OSPF zone interfaces connected to LAN segments to learn routes dynamically from other routing devices.

The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

The SteelConnect gateway is not a strict ABR implementation as it can operate only within one OSPF area.


Figure 5-22. WAN topology with MPLS CE router with OSPF area



To configure OSPF routing on a SteelConnect gateway on both the LAN and WAN side, you create an OSPF network, define a single area (or use an existing area), and attach a LAN-side zone and the uplink interface to the area. The basic steps to enable OSPF routing are:

- Select a site and create an OSPF network for that location that includes one area.
- Attach a zone and the uplink to the OSPF area.

Figure 5-23. OSPF network




Important: To configure OSPF (or BGP) on an uplink, you must allocate a static IP address.

For additional OSPF details, refer to [“Layer 3: Switch behind a SteelConnect gateway with LAN-side OSPF” on page 44](#) and the *SteelConnect Manager User Guide*. For more information about how to redistribute OSPF routes into BGP on routers, see the vendor documentation.

MPLS CE router replacement: ASBR-like deployment

In this topology, the SteelConnect gateway connects directly to the MPLS provider router. The gateway can replace a CE router to advertise and receive route announcements from the MPLS underlay.

Figure 5-24. WAN topology with MPLS CE router similar to ASBR



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

Important: SteelConnect gateways only support Ethernet handoffs.

To establish point-to-point connections between neighbors, you configure an eBGP session on a static IP uplink for a SteelConnect gateway and define its eBGP neighbor. Uplinks using DHCP don't support eBGP.

To configure BGP routing

1. Select the branch gateway, choose Appliances, and select the BGP tab.
2. Fill out these required session attributes:
 - **Router ID** - Specify the router IPv4 or IPv6 address to uniquely identify the SteelConnect gateway. This example uses the IPv4 address 172.29.1.2 for the uplink.

- **Local AS** - Specify the autonomous system (AS) number the SteelConnect gateway belongs to. This example uses the default setting of 65000.

Figure 5-25. Local AS setting

The screenshot shows the configuration page for a SteelConnect gateway named [HA] SDI-1030. The 'BGP' tab is active, displaying global BGP settings. The 'Router ID' is set to 172.29.1.2 and the 'Local AS' is set to 65000. There are 'Cancel' and 'Submit' buttons at the bottom right.

3. Create a new BPG neighbor under Network Design > Routing > BGP. Specify the following values:

- **Name** - Specify the peer name. Each uplink can support one eBGP peer.
- **Appliance** - Select the SteelConnect gateway appliance that acts as the CE router.
- **IPv4 Address** - Specify the peer router's IPv4 address. This example uses 172.29.1.1.
- **Remote AS** - Specify the autonomous system number the peer belongs to. This example uses AS 64962 as the MPLS PE router belonging to the peer.
- **Password** - Type a password to enable MD5 authentication. You must use the same password on both BGP neighbors. If you don't require MD5 authentication, you can leave this field blank.
- **Keep Alive Time** - 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 does not expire. The default setting is 60. Set this value to the value set on the PE router.
- **Hold Time** - 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 seconds. Set this to the value set on the PE router. 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.

Figure 5-26. Creating a new BGP neighbor

4. Attach the uplink to BGP.

Figure 5-27. BGP uplink

BGP			
<input type="text" value="Search"/>			
Appliance	Name	Attached Uplink	Available Uplinks
[HA] SDI-1030 HA	ZurichPE	-	Attach MPLS

On the MPLS (PE) router, the following BGP configuration needs to be entered by your MPLS provider:

```
interface FastEthernet0/1
ip address 172.29.1.1 255.255.255.0
duplex auto
speed auto
!
router bgp 64962
bgp log-neighbor-changes
network 10.33.195.0 netmask 255.255.255.0
network 172.29.1.0
neighbor 172.19.1.1 remote-as 65000
no auto-summary
```

Through SCM, you enable SteelConnect gateways to use BGP to advertise all of their associated LAN zones (IP subnets) to an upstream router in the BGP environment of the MPLS provider. You need to create a routing policy to activate the ASBR functionality.

To configure a routing policy

1. Choose Routing > Policy and create a new dynamic routing policy.

Figure 5-28. Creating a dynamic routing policy

The screenshot shows the 'Dynamic Routing Policy' page on the left and a 'Create Dynamic R...' dialog on the right. The dialog has a 'Name' field with 'ZurichRouting' and a 'Site' dropdown menu with 'Zurich (Zurich)' selected. There are 'Cancel' and 'Submit' buttons at the bottom right of the dialog.

2. Redistribute OSPF to BGP to advertise the LAN zones (site subnets learned through OSPF) in the underlay.

Edit the policy, select the BGP Profile tab, and set Redistribute OSPF to BGP to **On**.

Figure 5-29. Editing the BGP profile

The screenshot shows the 'Dynamic Routing Policy' page on the left and an 'Edit Dynamic Rout...' dialog on the right. The dialog has tabs for 'Info', 'BGP Profile', 'BGP Summarization', and 'OSPF Profile'. The 'BGP Profile' tab is selected. Under 'Redistribute Settings', the 'Redistribute OSPF to BGP' toggle is set to 'On'. There are 'Cancel' and 'Submit' buttons at the bottom right of the dialog.

3. Add a default route in OSPF by enabling the feature under the OSPF Profile tab.

Figure 5-30. Adding a default route

The screenshot shows the 'Edit Dynamic Rout...' dialog with the 'OSPF Profile' tab selected. Under 'Redistribute Settings', the 'Default route originate' toggle is set to 'On'. There are 'Cancel' and 'Submit' buttons at the bottom right of the dialog.

Important: There is no BGP route injection into OSPF on SteelConnect gateways as opposed to a full ASBR implementation. The SteelConnect gateway will advertise 0.0.0.0/0 into OSPF.

4. You can summarize/aggregate routes in BGP under the BGP Summarization tab.

The gateway uses eBGP to advertise its LAN-side subnets into the MPLS AS.

To view BGP learned and advertised routes

1. Choose Network Design > Uplinks.

2. Select the uplink with eBGP enabled.
3. Select the Networks tab.

The networks tab is dimmed when eBGP is not enabled on the gateway.

Figure 5-31. Viewing BGP routes of the gateway


The screenshot shows a web interface with a top navigation bar containing tabs: Info, Settings, AutoVPN, DynDNS, QoS, L2 Settings, and Networks. The Networks tab is selected and highlighted with an orange underline. Below the tabs, the section is titled "BGP Routes". A note states: "These are the reported BGP routes of the gateway. Note: This is an information page only. BGP can be configured on the 'BGP' tab." Below the note, there is a "Show" dropdown menu set to "10" and the text "entries". To the right is a "Search:" input field. Below these is a table with the following columns: "Version" (with a sort icon), "Type" (with a sort icon), and "Network" (with a sort icon). The table contains five rows of data. At the bottom left, it says "Showing 1 to 5 of 5 entries". At the bottom right, there are three buttons: "Previous", "1" (the current page), and "Next".

Version	Type	Network
IPv4	Learned	10.33.195.0/24
IPv4	Advertised	172.16.2.0/24
IPv4	Advertised	172.16.4.0/24
IPv4	Advertised	192.168.100.100/32
IPv4	Advertised	192.168.101.101/32

Hybrid WAN: internet and MPLS

As in the previous topologies, you connect a SteelConnect gateway to the routers by creating two WANs and two uplinks configured with different IP addresses. While SCM automatically creates the internet WAN and uplink, you need to add an uplink for the MPLS WAN after it is created, as explained in “Integration with MPLS CE router” on page 69.

Figure 5-32. Internet and MPLS WAN topology



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

SteelConnect can form VPN tunnels over both the default internet connection and the private MPLS WAN circuit. To set up a VPN over the MPLS network, turn encryption on when creating the WAN and configure the uplink to propagate the internal IP address that the remote sites use as the endpoint target for the AutoVPN tunnel.

For sites that have both circuits, there are two available encrypted tunnels (RouteVPN and MPLS) to choose from for inter-site traffic. You can set a default path preference for traffic under Organization > Networking Defaults.

Figure 5-33. WAN usage preference

WAN usage preference				
When internal zones in different sites can exchange traffic via more than one WAN path, this is the preference that is used.				
1	▲	▼	🚦 RouteVPN	
2	▲	▼	🔗 MPLS	

Note that you can override the default path preference for an individual site, zone, user group, user, or application by going to the individual site and configuring it with a different rule. For example, to ensure that all of the VoIP traffic at the NYC site chooses the MPLS link over the RouteVPN, set the MPLS link to the number one path preference as shown in [Figure 5-34](#). This setting overrides the networking default for the organization as shown in [Figure 5-33](#).

Figure 5-34. Creating a path rule for the NYC site

Create path rule

Position: >> Top <<

Site scope: NYC

Users / Source: All (excluding guests)

Applications / Target: Selected applications or groups
Telephony / Voice / VoIP

Path preference:

1	▲ ▼	MPLS	🗑
2	▲ ▼	RouteVPN	🗑

Path Quality profile: Latency sensitive metrics

QoS priority: Urgent


Cancel Submit

This rule now applies to all traffic for the NYC site. You can apply more granular rules using the other drop-down lists. See the traffic path rules topic in the *SteelConnect Manager User Guide* for more information.

Integration with WAN-facing firewall

This topology includes a firewall on the WAN side of the SteelConnect gateway, as shown in [Figure 5-35](#).

Figure 5-35. SteelConnect gateway with firewall



The following SteelConnect gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

To allow the gateway to form an AutoVPN tunnel to the site across the firewall from the other sites, you create a rule to allow inbound connections on UDP port 4500.

In addition, the SteelConnect gateway needs to ping externally and access SteelConnect services on the internet. The following table shows common services and connection ports.


Service	Outbound connection port
DNS	UDP 53
NTP	UDP 123
Uplink IP Reflector	TCP 80
SteelConnect Manager	TCP 443
Configuration/API	TCP 3900
SSH Proxy	TCP 3903
Tunneled SSH	TCP 3904
Reporting	TCP 3905

For a full list of connection ports that SteelConnect uses, see the “SteelConnect Connection Ports” topic in the *SteelConnect Manager User Guide*.

Highly secured WAN

Some customers operate highly secured WAN where only the LAN is trusted. All traffic coming from either the MPLS network and, of course, the internet is not trusted. Traffic can be encrypted by the firewalls on MPLS.

Figure 5-36. Highly secured topology with firewalls



For such a topology, we recommend you leverage virtualization capabilities of the firewalls. For example, Fortinet calls it Virtual Domains (VDM), Palo Alto Networks uses Virtual Systems, and Cisco has Virtual Context.


The concept has hardware firewall appliances operating one or more logical firewalls independently from each other. Each logical firewall performs specific operations.

The following SteelConnect Manager gateways support this topology.

SDI-VGW	SDI-130/130W	SDI-330	SDI-1030	SDI-5030
✓	✓	✓	✓	

A possible implementation of SteelConnect in that context appears in the following deployment diagram.

Figure 5-37. Highly secured topology with logical firewalls



Important: SDI-130 and SDI-330 have only two WAN ports so they require a switch in this example to connect the SteelConnect gateway to a firewall cluster.

SDI-VGW and SDI-1030 can have up to three WAN ports. In that case, the LAN Virtual System considers all WAN-facing interfaces as untrusted. All traffic transported by SteelConnect overlays will be analyzed by this firewall.

The WAN Virtual System protects the site from internet threats and lets a SteelConnect gateway forming IPsec tunnels (RouteVPN or standard IPsec communications) on the internet.

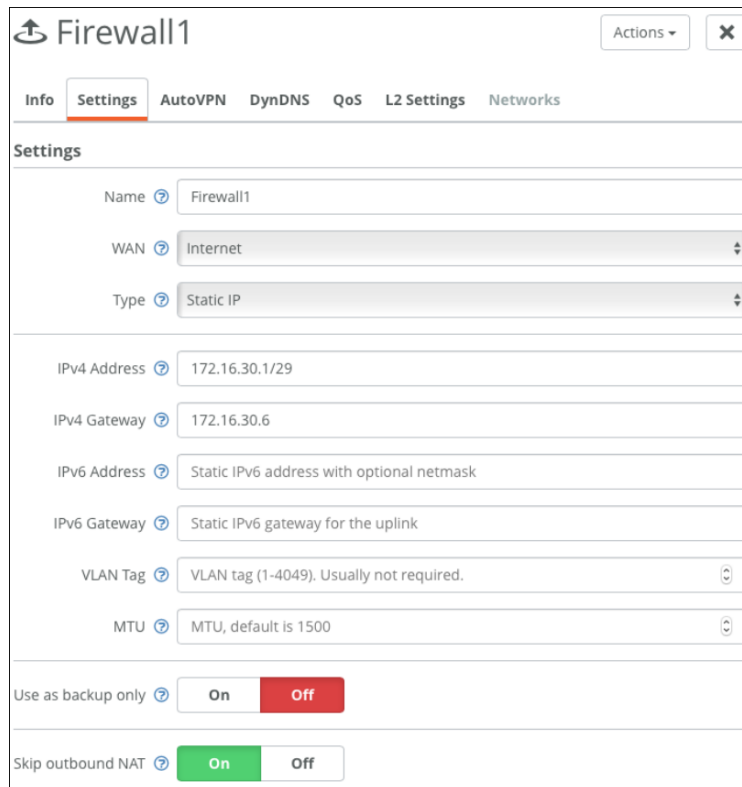
The gateway creates overlay tunnels (AutoVPN) on MPLS between SteelConnect-equipped sites.

To connect SDI-1030 to a firewall cluster (active/backup)

1. Create two internet uplinks that belong to the same internet WAN, have IP addresses in the same subnet, and point to the same gateway IP address (Firewall VIP).

Set the first uplink to active. (Ensure Use as backup only is set to **Off**.)

Figure 5-38. First firewall uplink



Firewall1 Actions X

Info **Settings** AutoVPN DynDNS QoS L2 Settings Networks

Settings

Name Firewall1

WAN Internet

Type Static IP

IPv4 Address 172.16.30.1/29

IPv4 Gateway 172.16.30.6

IPv6 Address Static IPv6 address with optional netmask

IPv6 Gateway Static IPv6 gateway for the uplink

VLAN Tag VLAN tag (1-4049). Usually not required.

MTU MTU, default is 1500

Use as backup only **On** **Off**

Skip outbound NAT **On** **Off**

2. Configure the second uplink as a backup. (Ensure Use as backup only is set to **On**.)

Figure 5-39. Second firewall uplink

The screenshot shows the 'Firewall2' configuration window. The 'Settings' tab is active. The configuration fields are as follows:

- Name: Firewall2
- WAN: Internet
- Type: Static IP
- IPv4 Address: 172.16.30.2/29
- IPv4 Gateway: 172.16.30.6
- IPv6 Address: Static IPv6 address with optional netmask
- IPv6 Gateway: Static IPv6 gateway for the uplink
- VLAN Tag: VLAN tag (1-4049). Usually not required.
- MTU: MTU, default is 1500
- Use as backup only: **On**
- Skip outbound NAT: **On**

Important: Consider the migration strategy: for communication with sites that are not yet equipped with SteelConnect, you can keep creating IPsec tunnels with the Firewalls cluster. Alternatively, you can create IPsec tunnels between a SteelConnect gateway and the remote firewall using the classic VPN feature.

WAN high availability

For complete end-to-end physical and path redundancy, we recommend that you deploy multiple (multihomed) upstream WAN devices along with the HA pair of gateways at the site.

Supported high-availability WAN configurations include MPLS, internet breakout, and VLAN. The SteelConnect HA gateways operate in active/passive mode: the master gateway processes traffic while the backup gateway remains in standby mode, ready to take over if the master gateway fails.

Note: Asymmetric HA is not supported.

Failover is triggered when:

- all uplinks on the master gateway fail.
- the gateway software or hardware fails.


If a specific uplink is critical to site performance, the master gateway can be configured to track that uplink and trigger failover if the connection on that uplink fails. The master gateway can likewise track LAN ports and trigger failover if there is a failure on that port. Those configuration options are located in the SteelConnect Manager GUI on the HA tab for the master gateway.

Note: To clearly distinguish among gateways, enter a unique location in the Location field under the Location tab for the gateway. It is not possible to name a gateway. See the *SteelConnect Manager User Guide* for details.

High-availability WAN topology with two routers


In this topology each gateway is connected directly to each router. Each router must have multiple LAN-facing ports.

Figure 5-40. HA gateways connected directly to WAN edge routers, each router has multiple WAN-side ports




Some routers may have only a single LAN-facing port. In this case use a switch, as shown in [Figure 5-41](#), to connect the two gateways to the WAN. If the WAN is an internet WAN, your ISP can provide a modem with switched ports.

Figure 5-41. HA gateways connected to WAN edge routers, one directly the other through a switch



If all routers each have only a single LAN-facing port, as shown in [Figure 5-42](#), configure the backup gateway with mirrored uplinks. That ensures that the master gateway configuration is mirrored on the backup gateway. The switch is configured with standard access ports.

Figure 5-42. HA gateways connected to a switch connected to multiple routers



To configure mirrored uplinks

1. In SCM, choose Appliances > Ports.
2. Find the backup gateway and select one of its ports. The port you select should be identical to one of the ports you configured on the master gateway. For example, if you configured port WAN1 and port WAN2 on the master gateway, select either port WAN1 or port WAN2 here.

3. In the Info/Mode tab on the port details page, select Mirrored Uplink from the Port Mode drop-down menu.

Figure 5-43. Mirrored uplinks

1

Info / Mode MACs / Devices Counters Advanced

Not online The appliance that this port belongs to is currently **not online**. The reported information listed on these tabs may not be correct any more. It was received **Never**.

Info **Mode**

Port Label 1

Link Status **Down**

STP Status *No link*

Site Zurich

Appliance [HA] SDI-1030 [Gateway 2]

Port mode

Cancel Submit

Patch Label


4. Click **Submit**.
5. Repeat these steps for the other port.
6. Close the port details page. On the Ports page, a blue HA icon appears next to the words Mirrored Uplink under the Mode column for the ports you configured.

Figure 5-44. Port details

		[HA] SDI-1030 Not online	
Gateway 2			
	1 RJ45	Mirrored Uplink	Zurich > WAN1
	2 RJ45	Mirrored Uplink	Zurich > WAN2
	3 RJ45	HA Control Port	
	4 RJ45	Singlezone	Zurich > LAN_Interco [1004]

You can use a SteelConnect switch to connect WANs to your gateways. See [Figure 5-45](#).

Figure 5-45. HA gateways connected to a SteelConnect switch connected to multiple routers



To use a SteelConnect switch to manage the interconnection between gateways and WANs

1. Configure gateways in the same way as if you used a third-party switch. See [Figure 5-42 on page 85](#).
2. Create zones in the site for each WAN interconnection. One zone must be configured as a management zone. When configuring zones, select Manual for Default Gateway configuration.

3. Ensure each zone has no assigned gateways.

Figure 5-46. Assigned gateways

The screenshot shows the 'WAN1-Interconnect' configuration page with the 'Gateways' tab selected. The page has a top navigation bar with tabs: IP, Gateways (active), DHCP, VLAN, WAN/AutoVPN, Settings, ADDL Networks, and Discovered Networks. Below the tabs, there's a section titled 'Automatic SteelConnect default gateway' with a paragraph explaining the feature and a 'Default Gateway configuration' section with 'Automatic' and 'Manual' buttons. The 'Manual' button is selected. Below this is the 'Gateway assignments' section with a paragraph explaining the table. At the bottom right is an 'Add assignment' button. The table has columns 'Type / Appliance', 'IPs', and 'Flags'. The table is currently empty, showing 'No gateway memberships present'.

WAN1-Interconnect Actions X

IP Gateways DHCP VLAN WAN/AutoVPN Settings ADDL Networks Discovered Networks

Automatic SteelConnect default gateway

When turning this option on, a SteelConnect gateway appliance deployed in the site will be automatically configured as the default gateway for this zone. It will then **use the default gateway IP addresses specified on the 'IP' tab**. If you want to control all gateway assignments for this zone manually, or you want to use a third-party default gateway for this zone, please turn this option off.

Default Gateway configuration Automatic Manual

Gateway assignments

This table shows all SteelConnect gateways that are members of this zone. You can create several memberships, also in remote sites. Every member gateway will be able to route into the zone's network. Default gateway entries that have been added automatically cannot be edited or deleted - if you want to control all gateway parameters, turn off automatic default gateway assignment and create a default gateway manually.

Add assignment

Type / Appliance	IPs	Flags
No gateway memberships present		

4. Configure the SteelConnect switch ports with the appropriate zones.

5. Cable the gateways and routers to the SteelConnect switch.

Data Center Topologies

This topic describes advanced topologies that we have qualified for enterprise data centers. It includes these sections:


- [“Using HA LAN and WAN topologies in the data center” on page 89](#)
- [“Traffic attraction with SDI-5030” on page 90](#)
- [“Creating a data center cluster of 5030 gateways” on page 95](#)

Using HA LAN and WAN topologies in the data center

You can leverage the high-availability LAN topologies and WAN topologies presented in the [“LAN Topologies” on page 33](#) and [“WAN Topologies” on page 53](#) topics for designing data center architectures with SteelConnect gateway appliances.

Figure 6-1 is an example of a deployment that can be achieved using the configurations described in the LAN and WAN topology chapters.

Figure 6-1. High availability in the data center



Important: We recommend you use only SDI-VGW and SDI-1030 in the branch for HA data center deployment.

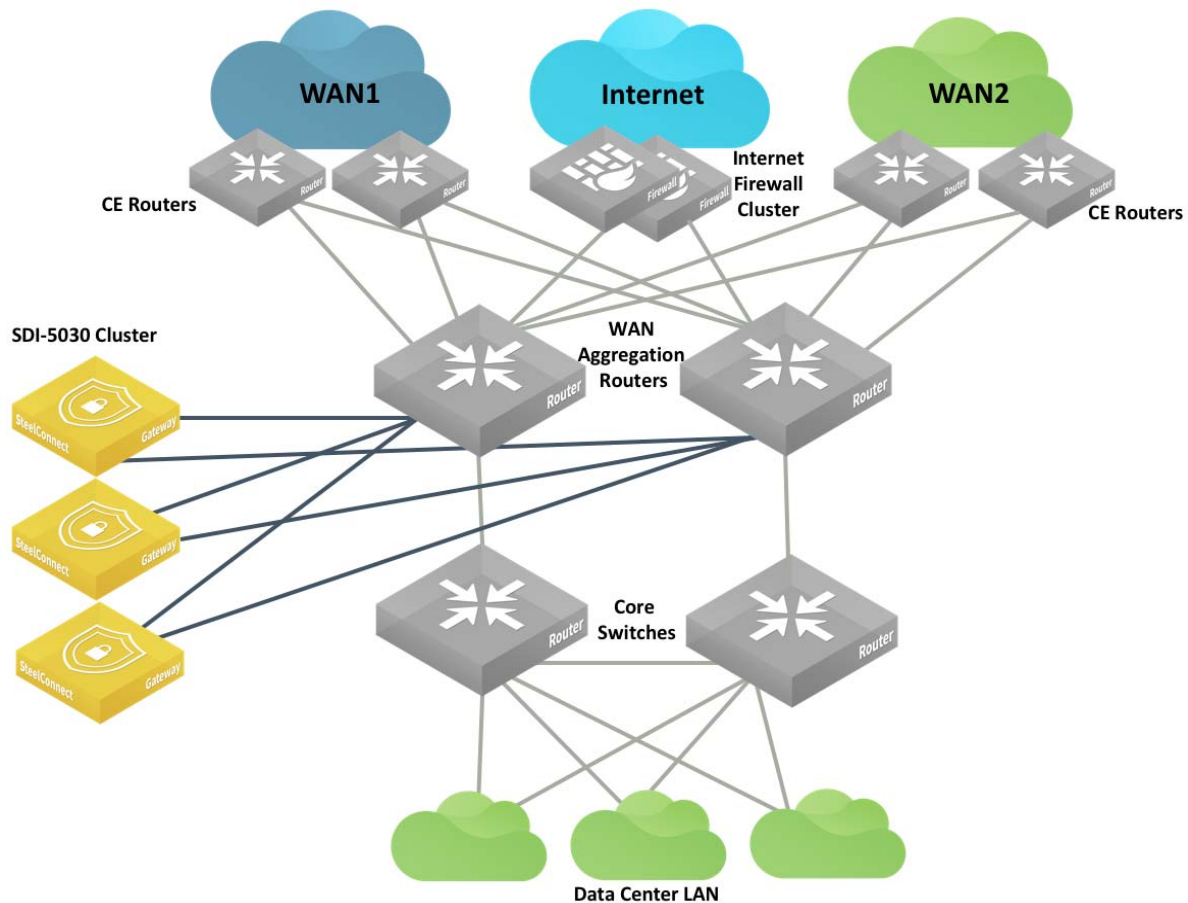
Traffic attraction with SDI-5030

In contrast to branch gateways that handle in-path traffic, the 5030 data center gateways are deployed out of path deep inside the data center network. Because the 5030 gateways are placed physically out of path from the data flow, you can deploy them with no network downtime. The system relies on traffic redirection.

Traffic attraction means that the 5030s exchange routes through eBGP with the aggregation routers. All traffic that reaches these routers destined to the branch is directed to the 5030s, is IPsec encapsulated with our AutoVPN mechanisms, and is sent in the overlay tunnel to the branch.

Note: All ports that have internet connectivity (including the management port) need to be protected by your security infrastructure.

Figure 6-2. Data center and branch end-to-end architecture



At a high level, you need to complete these tasks to configure this topology:

1. “Cabling the appliances” on page 92
2. “Registering the appliances to the organization” on page 93
3. “Enabling the data ports” on page 94
4. “Deploying the appliances in a site” on page 94
5. “Creating a data center cluster of 5030 gateways” on page 95
6. “Using BGP for traffic attraction on SCM” on page 96
7. “Configuring BGP on the WAN aggregation routers” on page 99
8. “Creating data center uplinks for the cluster” on page 99

Cabling the appliances

The first step when setting up a data center deployment is to connect the SDI-5030 appliances to the network. Because the 5030 gateways are placed physically out of path from the data flow, you can deploy them with no network downtime. For a cluster, we recommend identically cabling the gateways for redundancy.

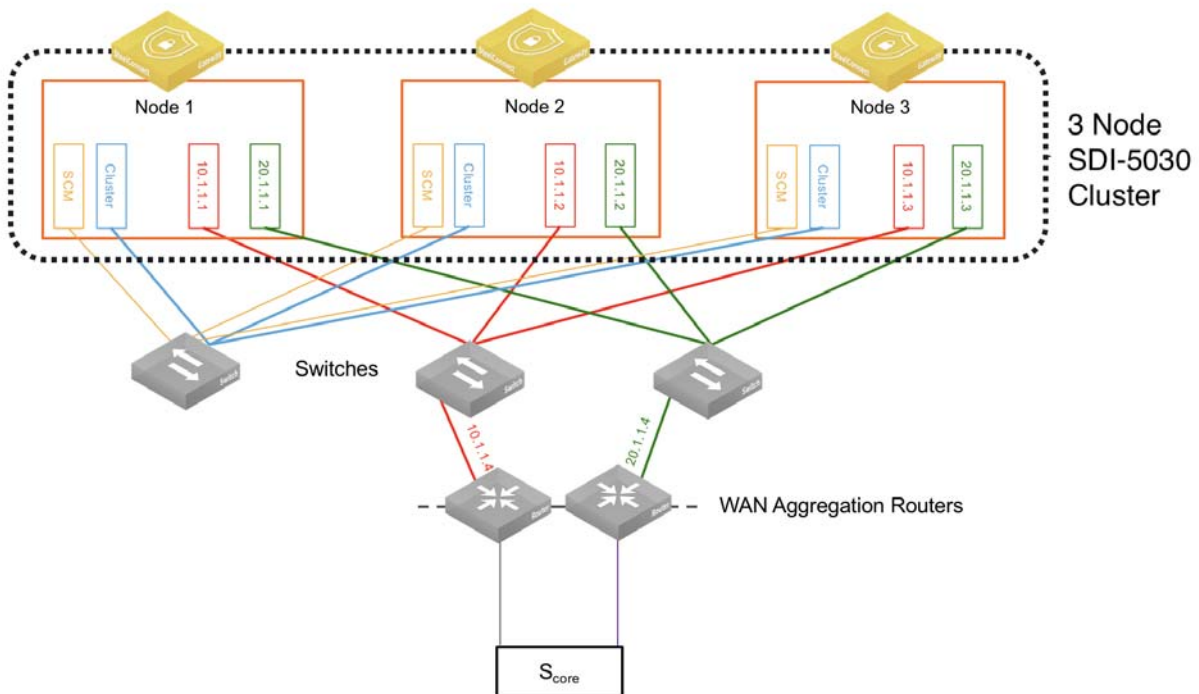
5030 gateways use ports 1 and 2 for system-related tasks. The remaining ports are data ports.

- Port 1 is the management port dedicated to SCM. It is required to use DHCP to allocate dynamically the management IP address. This network must provide connectivity to the internet to allow calls to the SteelConnect Manager. DNS support is mandatory to access Riverbed services. The management port should not be accessible directly from the public internet. This port should be behind the your firewall/security complex and it only requires enabling the documented outbound services.
- Port 2 is a dedicated cluster connectivity port, providing an internal, private physical connection between the 5030 gateways for cluster synchronization and high availability. All 5030 gateways must be Layer 2 adjacent in this network.
- Port 3 is a data port for WAN-facing connectivity. Optionally, port 3 can be used for high availability.
- Ports 4 through 10 are data ports for WAN-facing connectivity.

Important: Whenever possible, we recommended using 10G ports (ports 5, 6, and higher) for the WAN facing connections to guarantee best performance.

For each appliance, one WAN port is dedicated to one WAN. In the following example, there are two WAN connections (shown in this figure in red and green).

Figure 6-3. Three node 5030 cluster



Registering the appliances to the organization

Next, register the appliance.

To register the appliance

1. In SCM, choose Appliances > Overview.
2. Click **Add appliances** and choose Register Hardware Appliance from the drop-down menu.
3. Provide the serial number and click **Submit**.

Figure 6-4. Registering the appliance

The screenshot shows the 'Register Hardware Appliance' dialog box. It contains a text input field for 'Appliance serial #' with a placeholder 'Enter serial #'. Below it is a dropdown menu for 'Deploy into site' with the selected option '>> Do not deploy into a site yet <<'. At the bottom right are 'Cancel' and 'Submit' buttons.

Important: Don't assign the appliance to a site yet; you complete this step later in the process.

When the appliance is registered, it downloads the latest firmware and starts the installation. We highly recommend completing the upgrade process before moving to the next steps.

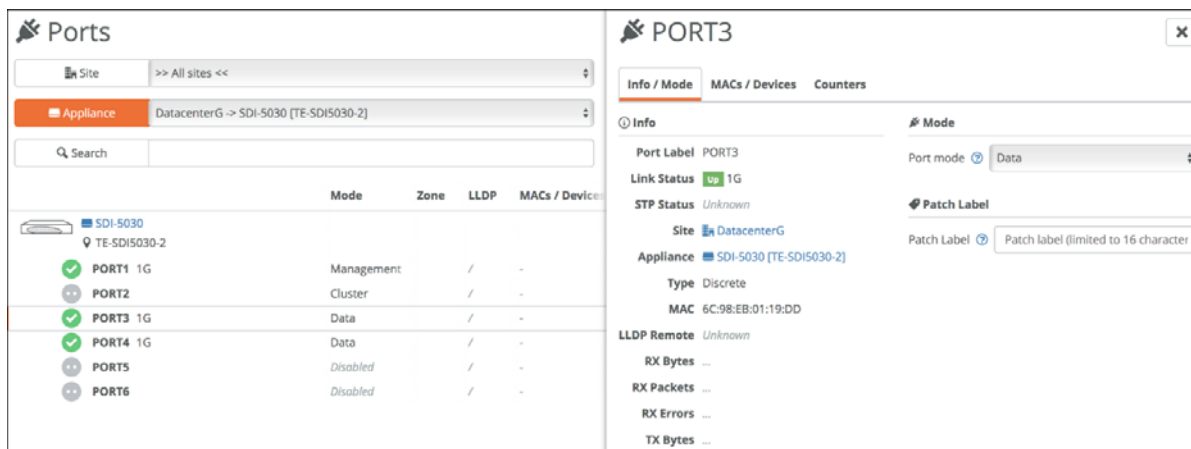
Enabling the data ports

Each 5030 appliance must have at least one data port connected to the WAN.

To enable data ports to use for the 5030 WAN connections

1. Choose Appliances > Ports.
2. Click the 5030 ports to use as a data port.
3. In the Info/Mode tab, set the port mode to Data.

Figure 6-5. Enabling data ports



Deploying the appliances in a site

At this stage of the process, you can assign the appliances to the sites.

To deploy an appliance to a site

1. Choosing Appliances > Overview.
2. Click the appliances to configure and go to the Location tab.

3. Specify the site and a location and click **Submit**.

Figure 6-6. Assigning the site and location

The screenshot shows the configuration page for an SDI-5030 gateway. At the top, there's a header with the device name 'SDI-5030' and an 'Actions' dropdown. Below this is a navigation bar with tabs: 'Live', 'IPs', 'AutoVPN', 'HA', 'Location' (which is selected and underlined), 'BGP', 'Static Routes', and 'Interfaces'. The main content area is titled 'Site & Location'. It contains a text box explaining that appliances need to be assigned to a site before they can operate. Below this, there are two input fields: 'Site' with a dropdown menu showing 'DatacenterG (Datacenter G)' and 'Location' with a text field containing 'TE-SDI5030-2'. At the bottom right, there are 'Cancel' and 'Submit' buttons.

Creating a data center cluster of 5030 gateways

To accommodate data center workloads, 5030 gateways are designed to operate in a cluster. Clusters provide resiliency and reliability in addition to higher bandwidth throughput. A cluster requires a minimum of three 5030 gateways to provide high availability with two gateways as active nodes and one as a spare node.

Note: You can have a single 5030 in your data center deployment but you miss the advantages of a cluster and high availability.

To create a new data center cluster of 5030 gateways

1. Choose Network Design > Clusters. Click **New Cluster**.
2. Specify a cluster name. Select the site to deploy the cluster from the drop-down list. In this example, deploy it at DataCenterG.

Tip: In a cluster workflow, it can be difficult to differentiate between data center gateways in a cluster when they are referenced on various SCM pages. We strongly recommend that you always specify a detailed location for the

gateway using the Location field under the Location tab in the appliance page. Setting the location associates a gateway with its location wherever an appliance is referenced, making it easy to identify.

Figure 6-7. Creating a cluster

Create cluster

Before configuring a data center gateway cluster, add at least one 5030 gateway to a site. 5030 gateways must be connected to SCM using a dedicated connection. Register the gateways with SCM and verify that they are running the same firmware version. Also, make sure that the gateways are cabled on the WAN side.

Cluster name ⓘ DataCenterG

Site ⓘ DataCenterG (DataCenter G) ▼

Cluster members ⓘ Nothing selected

Failover nodes ⓘ 0 ▼

Interceptor ⓘ **On** Off

Cluster Proxy IP ⓘ Proxy IP with netmask

Cancel Submit

3. Select the cluster members from the drop-down list.
4. Specify the number of failover nodes. By default, the system creates one failover node in a cluster of three 5030 gateways. For a single 5030 cluster, leave the number of failover nodes at 0.
5. For Interceptor, select **Off** to enable traffic attraction.
Traffic attraction means that the 5030s will exchange routes via eBGP with the aggregation routers so that all traffic that reaches these routers destined to the branch is directed to the 5030s, is IPsec encapsulated, and is sent in the overlay tunnel to the branch.
6. Click **Submit**.

Using BGP for traffic attraction on SCM

The SDI-5030 is not an in-path device and needs BGP to peer with a WAN aggregation stack within a data center. This is required not only for route attraction to occur but also to ensure proper cluster function by prepending BGP advertisements. BGP is critical to ensure the internal and cluster load-balancing.

Use BGP traffic attraction to attract traffic bound for the branches to the data center gateways and to perform traffic tunneling and redirection.

One or more SDI-5030 gateways that are deployed as a data center gateway cluster peer with a WAN distribution router using BGP.

The cluster of SDI-5030 forms an eBGP peering relationship with the WAN distribution routers.

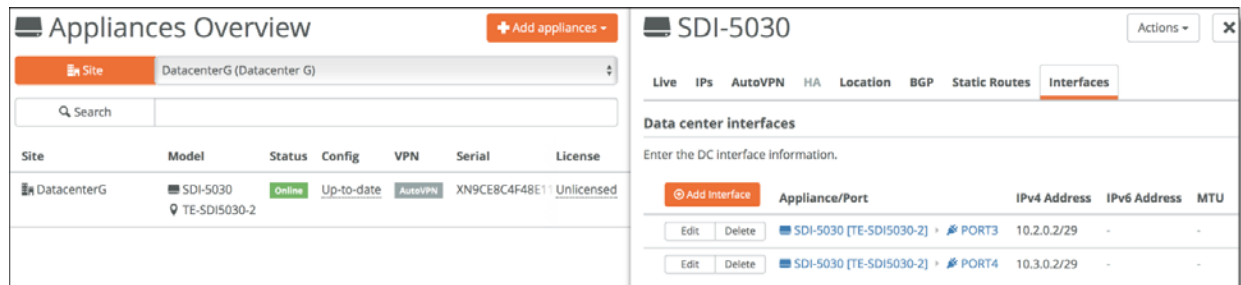
The cluster then advertises Local Tunnel Endpoint (LTEP) subnets as well as the remote subnet prefixes (of the SD-WAN equipped sites) for traffic attraction into the SD-WAN (overlay) network.

In the opposite direction, the data center's aggregation routers communicate to the 5030 gateways all the prefixes they need to route to data center subnets, including branch Tunnel Endpoint (TEP) addresses and the branch subnets for all non-SD-WAN branches.

To configure the individual 5030 gateways with local BGP information

1. Choose Appliances > Overview.
2. Select an appliance.
3. Select the Interfaces tab and configure the IP address for each data port.

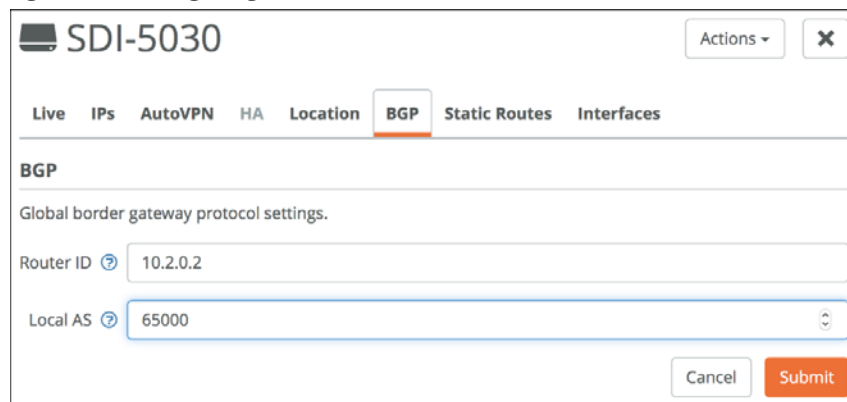
Figure 6-8. Interfaces tab



4. Select the BGP tab.
5. Add the unique Router ID for this appliance and the Local AS.

Important: For each SDI-5030 appliance, you need to allocate a unique BGP AS number. Router ID is the IP address of the Data Interface 3.

Figure 6-9. Configuring local BGP information



6. Repeat these steps for all 5030 gateways in the cluster.

Next, configure the BGP information for the cluster. The cluster needs the router information to communicate with the 5030 gateway.

To configure BGP information for a data center cluster

1. Choose Network Design > Cluster.
2. Select the cluster.
3. Select the BGP tab.
4. Click **Add BGP Neighbor**.
5. Specify a name for the neighbor.
6. Specify the IPv4 address of the neighbor.
7. Specify the Remote AS number the neighbor belongs to. In this example, 5353.
8. In the Password field, type a password to enable MD5 authentication. You must use the same password on both BGP neighbors. This field is optional if you don't require MD5 authentication.
9. Enter the Keep Alive Time and Hold Time in your network. The defaults are 60 and 180, respectively. Use these values if they are acceptable for your environment.
10. Click **Submit**.

Figure 6-10. Configuring BGP information for a data center cluster

The screenshot shows a 'New BGP Neighbor' configuration window. The fields are as follows:

- Name: CoreRouter2
- Appliance: SDI-5030 [TE-SDI5030-2]
- IPv4 Address: 10.3.0.1
- Remote AS: 5353
- Password: Password
- Keep Alive Time: 60
- Hold Time: 180

Buttons: Cancel, Submit

11. Repeat Steps 3 through 10 to create additional neighbors. Because there are three 5030 gateways and two WAN connections for each, there are a total of six BGP neighbors to add.

Important: The SDI-5030 does not yet support LAN-side route retraction and/or dynamic LAN-side zone learning. At this time, the SDI-5030 does not withdraw a route to a remote zone. Once a network is configured in SCM, it is advertised by the SDI-5030 via BGP regardless of the tunnel state to the remote site. The SDI-5030 does not learn routes in the data center.

Configuring BGP on the WAN aggregation routers

The data center aggregation routers must communicate with the 5030 gateways about the data center subnets reachability. Also, these routers need to establish a BGP neighbor relationship with the 5030 gateways so that routes are exchanged and packets destined to the branches are first sent to the 5030 gateway for encapsulation.

Note: With 5030 gateways, you need central VPN keying (C-VPN-K).

In this topology, there are two WAN aggregation routers used for redundancy. Both routers must be configured for the BGP session using the command-line interface as follows:

```
router bgp 5353
bgp log-neighbor-changes
network 192.168.25.0
network 192.168.26.0
network 192.168.27.0
network 172.16.3.0
neighbor 192.168.25.1 remote-as 65000
neighbor 192.168.25.2 remote-as 65002
neighbor 192.168.25.3 remote-as 65004
no auto-summary
```

See your WAN router aggregation vendor documentation for more detailed configuration information.

Creating data center uplinks for the cluster

Next, configure the data center uplinks. On a 5030 gateway, a data port physically connects the cluster to a WAN. Uplinks on a 5030 gateway are logical. A cluster must have at least a single uplink or multiple uplinks to the same WAN and can connect to multiple WANs. This topology has three WANs, so you need to configure two WANs, so you need to configure two data center uplinks: one for the internet WAN and one for the MPLS WAN.

To create data center uplinks for the cluster

1. Choose Network Design > Clusters.
2. Select a cluster to associate with the uplink. Each uplink is cluster specific and its connection type differs between clusters.
3. Select the Datacenter Uplinks tab. Click **Add Datacenter Uplink**.
4. Type the uplink name.
5. Select a WAN.
6. Under Tunnel Endpoints, enter the IP addresses of the tunnel endpoints.

For each appliance, two TEP IP addresses are required per WAN.

A UDP port is required only for the internet WAN and must be allocated. Any internet firewall needs to forward these ports to the TEP IP address.

These IP addresses can be any addresses as long as they are unique in the network. They can be considered as loopback IP addresses.

For three appliances in the cluster, add two tunnel endpoints for each gateway, resulting in six endpoint IP addresses for internet.

7. Specify the public IPv4 address of the uplink with an optional netmask. This address is required for internet WANs and optional for other WANs. In this example topology, the address is 100.100.1.2.

8. Click **Submit**.

Figure 6-11. Data center uplink for the internet WAN

The screenshot shows the 'New Data Center Uplink' configuration window. The 'Name' field is 'DCG Internet Uplink'. The 'WAN' dropdown is set to 'Internet'. The 'Tunnel Endpoints' list contains six entries: 4500->10.0.11.10, 4501->10.0.11.11, 4502->10.0.11.12, 4503->10.0.11.13, 4504->10.0.11.14, and 4505->10.0.11.15. The 'Public IPv4 Address' field is '100.100.1.2'. The 'Cancel' and 'Submit' buttons are at the bottom right.

9. Repeat these steps for the MPLS WAN.

Figure 6-12. Data center uplink for the MPLS WAN

The screenshot shows the 'New Data Center Uplink' configuration window. The 'Name' field is 'DCG MPLS uplink'. The 'WAN' dropdown is set to 'MPLS'. The 'Tunnel Endpoints' list contains six entries: 10.0.12.10, 10.0.12.11, 10.0.12.12, 10.0.12.13, 10.0.12.14, and 10.0.12.15. The 'Public IPv4 Address' field is 'Public IPv4 address of this uplink.'. The 'Cancel' and 'Submit' buttons are at the bottom right.

Note: NAT ports for the tunnel endpoints are only required for the internet WAN. Because the MPLS WAN is a private circuit, no public IPv4 address is needed.

AWS Cloud Topologies

SteelConnect Manager (SCM) provides the ability to orchestrate the deployment of virtual SteelConnect gateways and virtual SteelHeads directly into Amazon Web Services (AWS) Virtual Private Clouds (VPCs) effectively linking branch offices, data centers, and headquarter (HQ) locations and VPCs, both in the same or different regions. After integrating SCM with your AWS account (through Identity and Access Management [IAM] cross-account), the console automatically discovers and displays your subnets, in all VPCs and regions.

With SteelConnect's knowledge of your infrastructure, you can deploy virtual SteelConnect gateways and SteelHeads in all (or individual) VPCs and establish an automated VPN overlay through the internet while also benefiting from optimizing your applications over the WAN. This software-defined WAN automation interconnects VPCs using full-mesh VPN routing—with no manual configuration.

Note: AWS, Azure, and SteelConnect use slightly different terms to refer to similar network concepts. Virtual Private Clouds (AWS) and VNets (Azure) are called *sites* in SteelConnect, and subnets are referred to as *zones* in SteelConnect.

This topic provides a guide and examples to connect AWS VPCs with SteelConnect and deploy the SteelHead for SteelConnect WAN optimization solution in AWS. Traffic can be optimized from your branch offices and data centers to your VPCs and/or directly between VPCs.

Deployment into the AWS cloud

The configuration steps to deploy SteelConnect gateways in the AWS cloud are outlined in the *SteelConnect Manager User Guide*. The steps to do so are:

1. Subscribe to Riverbed AWS products.
2. Configure your AWS accounts with SteelConnect Manager.
3. Import AWS networks.
4. Deploy SteelConnect gateways to your AWS network.
5. Optionally, configure high availability through redundancy.

AWS also offers advanced features, including:

- AWS Direct Connect to configure a private network connection to your cloud deployments
- SteelConnect AWS transit VPC to communicate between multiple Virtual Private Clouds (VPCs)

Once you have read through the steps in the *SteelConnect Manager User Guide* on how to configure your deployment into AWS, you can deploy a SteelConnect gateway in your AWS environment.

Deploying a HA SteelConnect gateway in AWS

You can deploy SteelConnect gateways in AWS in high availability.

An example of this topology has multiple on-premises sites and the organization wants to extend its network to the AWS cloud and connect to its subnets in the Mumbai Region.

Figure 7-1. Network before AWS in Mumbai

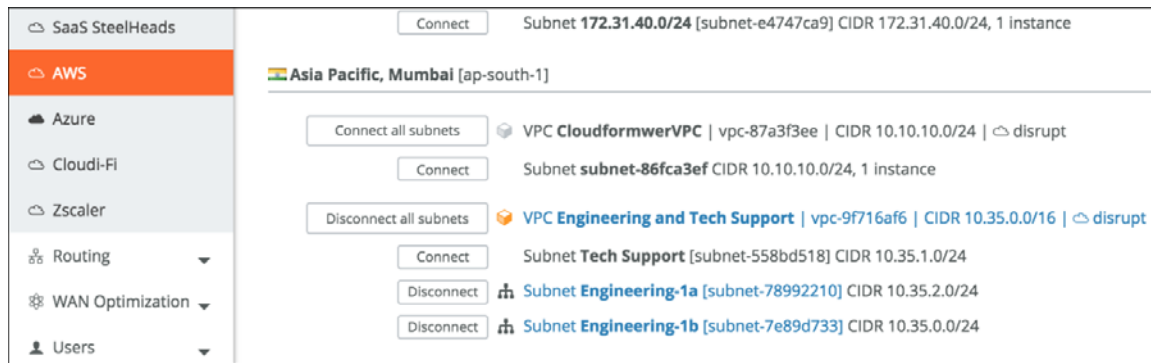


To extend the network to Mumbai with AWS in high availability

1. Connect the subnets.

In SCM, choose Network Design > AWS, select the Import VPCs tab, and click **Connect**.

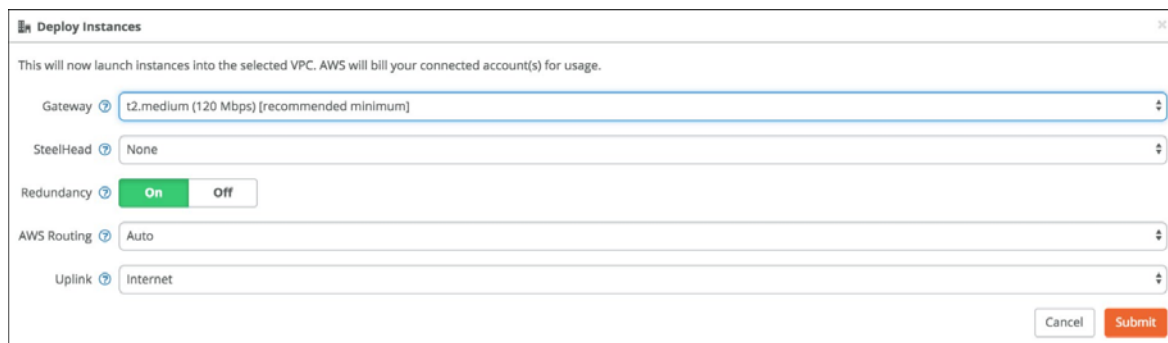
Figure 7-2. Connect VPC



When you connect a subnet in a VPS and AWS region, SCM creates a Site and Zone object for your Organization.

- From the AWS Deploy Instances tab, deploy the gateway and set Redundancy to **On**.

Figure 7-3. Deploy AWS gateway with redundancy

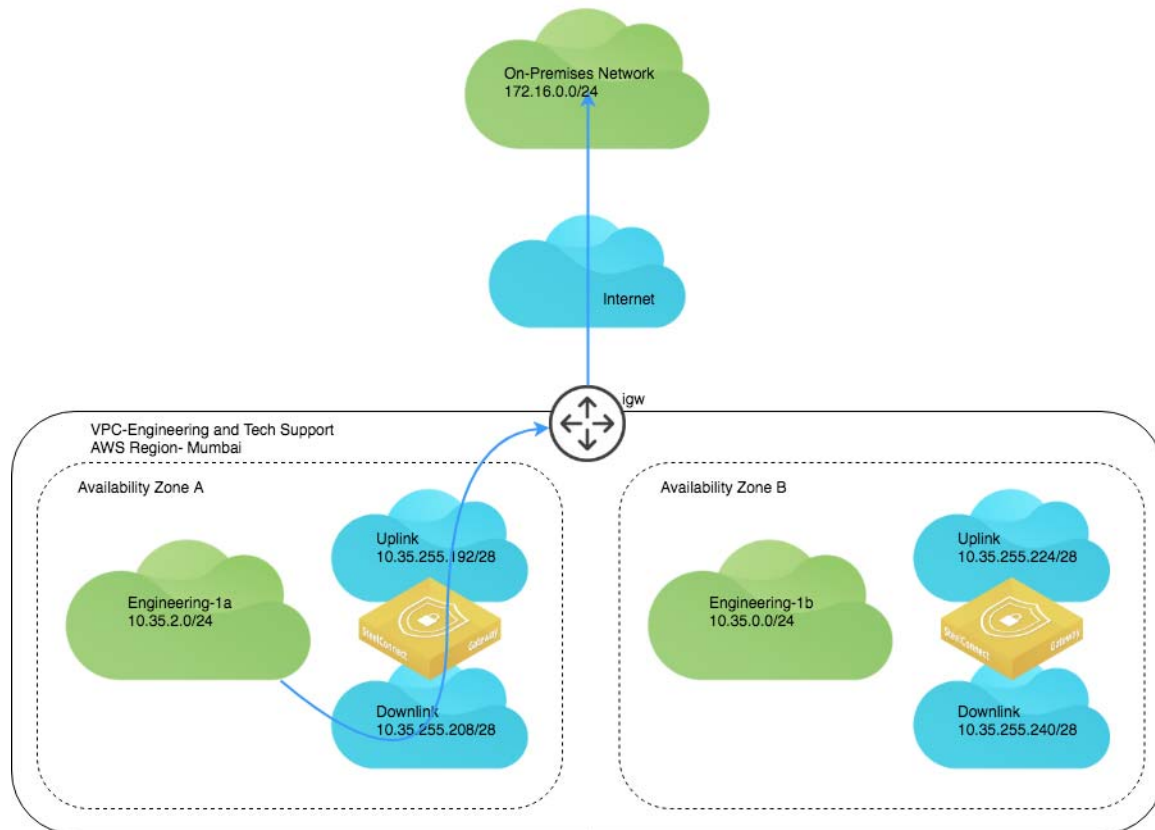


When you click **Submit** to deploy your instance, SCM logs in to your AWS environment, scans for two free /28 subnets: one for the uplink and the other for the downlink subnet of the gateway to deploy. It does the same for the subnet in the other Availability Zone. The gateway boots up, patches itself, gets an Elastic IP address, and connects to the internet gateway (IGW) of the VPC.

Note: You need to have a network configured that has available subnets.

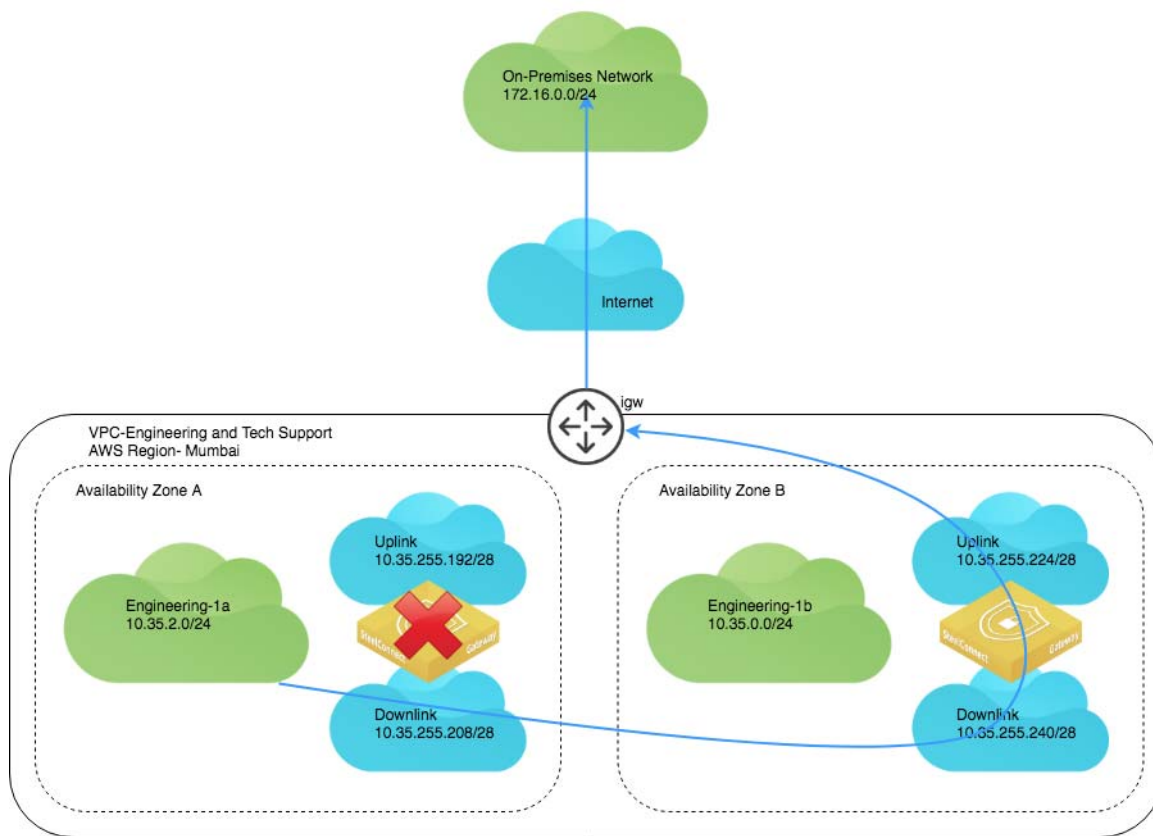
Figure 7-4 details the AWS deployment. Note that only one on-premises network appears for simplicity purposes. In this deployed network, IPsec VPN tunnels are automatically formed from the AWS subnet to all other on-premises networks in a full-mesh.

Figure 7-4. AWS deployment architecture



In case of the failure of the gateway in Availability Zone A, your network can be routed through the gateway deployed in Availability Zone B.

Figure 7-5. Failover for AWS HA



In your VPC, two gateways appear.

Figure 7-6. HA gateways

INSTANCES		i-047b78d1e8e0136...	t2.small	ap-south-1a	running	2/2 chec
Instances		SteelConnect_ap-south-1b_Virtual Gateway	i-0d540d51b6fb633fa	t2.medium	ap-south-1b	2/2 chec
Launch Templates		SteelConnect_ap-south-1a_Virtual Gateway	i-0fd6c74a7dbd81218	t2.medium	ap-south-1a	2/2 chec

The two appliances also appear in HA mode deployed under the Appliances tab on SCM.

Figure 7-7. HA appliances

Appliances Overview							+ Add appliances
Site	>> All sites <<						
Search							
Site	Model	Status	Config	VPN	Serial	License	
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6	[HA] SDI-AWS ap-south-1a	Online	Up-to-date	-	XN1E8DE65BE26142	Not needed	
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6	[HA] SDI-AWS ap-south-1b	Online	Up-to-date	AutoVPN	XND81418F04E1679	Not needed	

In the AWS VPC Dashboard, entries appear for the SteelConnect appliances' uplinks and downlinks.

Figure 7-8. AWS dashboard with SteelConnect uplinks and downlinks

The screenshot shows the AWS VPC Dashboard. On the left, the 'Subnets' link is selected in the navigation menu. The main area displays a table of subnets with columns: Name, Subnet ID, State, and VPC. The subnets listed are:

Name	Subnet ID	State	VPC
SteelConnect_ap-south-1b_Uplink	subnet-863dd6ca	available	vpc-9f716af6 Engineering and T...
SteelConnect_ap-south-1b_Downlink	subnet-373cd77b	available	vpc-9f716af6 Engineering and T...
SteelConnect_ap-south-1a_Uplink	subnet-1a992272	available	vpc-9f716af6 Engineering and T...
SteelConnect_ap-south-1a_Downlink	subnet-4499222c	available	vpc-9f716af6 Engineering and T...
Engineering-1b	subnet-7e89d733	available	vpc-9f716af6 Engineering and T...
Engineering-1a	subnet-78992210	available	vpc-9f716af6 Engineering and T...

With Auto selected for routing when deploying the gateways, all the routes to other networks (on-premises, or any other existing cloud network sites) are added by the script automatically when the gateways start deploying for the subnets on AWS.

Figure 7-9. AWS routes

The screenshot shows the AWS VPC Dashboard with the 'Routes' tab selected for the route table 'rtb-5696503e'. The table lists routes with columns: Destination, Target, Status, and Propagated. The routes are as follows:

Destination	Target	Status	Propagated
10.35.0.0/16	local	Active	No
10.0.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.2.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.3.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.4.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.6.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.7.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.8.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.9.1.16/28	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
10.17.255.240/29	eni-1812c033	Black Hole	No
172.16.0.0/24	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
172.16.1.0/24	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
172.16.2.0/24	eni-51663c7a / i-0d540d51b6fb633fa	Active	No
172.30.255.208/28	eni-1812c033	Black Hole	No

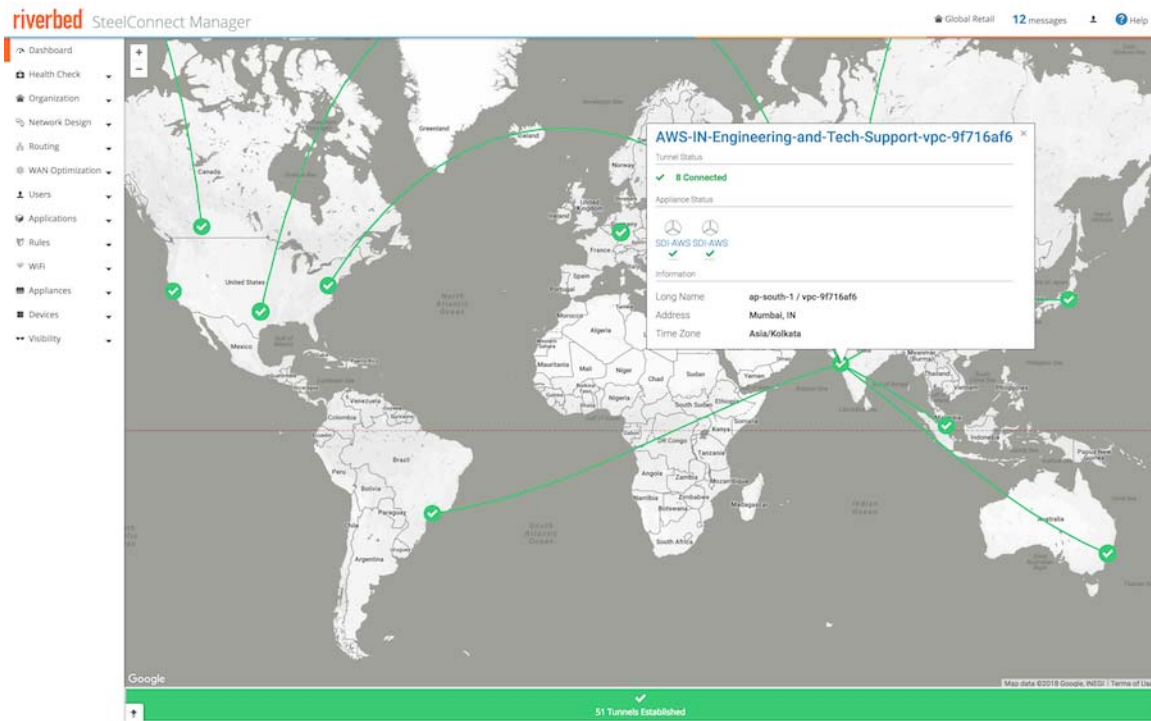
The real-time event log in SCM shows the deployment process.

Figure 7-10. SCM event log with deployment details

2018-03-07 10:19:06	36m ago			AutoVPN tunnel between site Tokyo and site Dallas over Internet on uplink uplink-Uplink-e719ed1a30410c80 came online
2018-03-07 10:18:00	38m ago			AutoVPN tunnel between site Tokyo and site Dallas over Internet on uplink uplink-Uplink-e719ed1a30410c80 went offline
2018-03-07 10:15:29	40m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	External IPv4 address change on Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink (new IP 35.154.159.230)
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Dallas and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-6626b2
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Washington and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-e
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Tokyo and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-e719ed
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Sydney and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-ed6e4
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Singapore and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-a0
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Sao Paulo and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-291
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Heisinki and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-0e82
2018-03-07 10:14:31	41m ago			AutoVPN tunnel between site Calgary and site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 over Internet on uplink uplink-Uplink-cbbd8
2018-03-07 10:14:09	41m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink came online
2018-03-07 10:14:09	41m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink is now in use
2018-03-07 10:14:09	41m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	External IPv4 address change on Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink (new IP)
2018-03-07 10:14:09	41m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance firmware version is now 2.10.0.50-booboo.
2018-03-07 10:14:06	41m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance uploaded new inventory information
2018-03-07 10:13:24	42m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	External IPv4 address change on Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink (new IP 13.127.156.137)
2018-03-07 10:13:01	42m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Appliance firmware version is now 2.10.0.50-booboo.
2018-03-07 10:12:58	43m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Appliance uploaded new inventory information
2018-03-07 10:12:57	43m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Port [HA] SDI-AWS [ap-south-1b] HA Master → LAN2 link state went UP
2018-03-07 10:12:57	43m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink came online
2018-03-07 10:12:57	43m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink is now in use
2018-03-07 10:12:57	43m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	External IPv4 address change on Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink (new IP)
2018-03-07 10:12:50	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance rebooting
2018-03-07 10:12:50	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Sending FIRMWARE_SWITCH (2.10.0.50-booboo).
2018-03-07 10:12:49	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink came online
2018-03-07 10:12:49	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink is now in use
2018-03-07 10:12:49	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	External IPv4 address change on Uplink AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 → Uplink (new IP)
2018-03-07 10:12:46	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Sending FIRMWARE_SWITCH (2.10.0.50-booboo).
2018-03-07 10:12:42	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance has downloaded firmware version 2.10.0.50-booboo.
2018-03-07 10:12:39	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Port [HA] SDI-AWS [ap-south-1a] HA Backup → LAN2 link state went UP
2018-03-07 10:12:21	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance firmware version is now 2.8.2.16-booboo.
2018-03-07 10:12:19	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance came online
2018-03-07 10:12:18	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Port [HA] SDI-AWS [ap-south-1a] HA Backup → LAN2 edited: mac changed from 00:00:00:00:00:00 to 02:4A:CC:16:CC:4A.
2018-03-07 10:12:18	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Port [HA] SDI-AWS [ap-south-1a] HA Backup → LAN1 edited: mac changed from 00:00:00:00:00:00 to 02:61:F0:0D:4D:7E.
2018-03-07 10:12:17	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance uploaded new inventory information
2018-03-07 10:12:16	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Port [HA] SDI-AWS [ap-south-1a] HA Backup → LAN2 link state went DOWN
2018-03-07 10:12:16	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Appliance is downloading new firmware
2018-03-07 10:12:16	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Port [HA] SDI-AWS [ap-south-1a] HA Backup → LAN1 link state went UP
2018-03-07 10:12:16	43m ago	[HA] SDI-AWS [ap-south-1a]	HA Backup	Newer firmware available, sending FIRMWARE_DOWNLOAD (2.10.0.50-booboo)
2018-03-07 10:11:07	44m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Appliance rebooting
2018-03-07 10:11:06	44m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Sending FIRMWARE_SWITCH (2.10.0.50-booboo).
2018-03-07 10:11:03	44m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Appliance has downloaded firmware version 2.10.0.50-booboo.
2018-03-07 10:10:45	45m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Port [HA] SDI-AWS [ap-south-1b] HA Master → LAN2 link state went DOWN
2018-03-07 10:10:45	45m ago	[HA] SDI-AWS [ap-south-1b]	HA Master	Port [HA] SDI-AWS [ap-south-1b] HA Master → LAN1 link state went UP
2018-03-07 10:10:38	45m ago			AWS account disrupt : Adding route 10.3.1.16/28 via eni-51663c7a in site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6
2018-03-07 10:10:37	45m ago			AWS account disrupt : Adding route 10.7.1.16/28 via eni-51663c7a in site AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6

After you deploy the gateways in HA mode, the dashboard displays the new site.

Figure 7-11. SCM dashboard with HA configured



SteelConnect gateway with SteelHead in AWS

In this topology, we deploy SteelHead for end-to-end WAN optimization in the network. SteelHead WAN optimization enhances your SteelConnect Infrastructure as a Service (IaaS) deployment by automatically optimizing all TCP traffic to increase speed and reduce bandwidth over the secure VPN links. In addition to interconnecting your cloud infrastructure to your on-premises sites (or bringing the cloud closer by providing a direct path to your IaaS applications from each branch) you can now use SteelHead to boost performance and reduce traffic. Everything is controlled through a cloud console for a true hybrid-WAN solution using the very latest in software-defined automation.

The process is similar to the steps described in the previous topology, but when you deploy your gateway in AWS, select a specification for the SteelHead. The specifications include throughput limits that indicate the amount of optimized traffic the instance supports.

Figure 7-12. Deploying SteelHead with a gateway

Deploy Instances

This will now launch instances into the selected VPC. AWS will bill your connected account(s) for usage.

Gateway: t2.medium (120 Mbps) [recommended minimum]

SteelHead: t2.medium (30 Mbps throughput)

Redundancy: ☒ On ☐ Off

AWS Routing: Auto

Uplink: Internet

Cancel Submit

Once deployed, the appliance appears in the Appliances Overview page.

Figure 7-13. Appliances Overview page

Appliances Overview ➕ Add appliances

Site: >> All sites <<

Search:

Site	Model	Status	Config	VPN	Serial	License
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6	SDI-AWS ap-south-1b	Online	Up-to-date	AutoVPN	XND81418F04E1679	Not needed

You can see the status of the currently deployed appliance by clicking Manage on the AWS tab on SCM.

Figure 7-14. Managing the SteelHead in AWS

VPC Details

VPC ID: vpc-9f716af6

AWS Routing: Auto

Deployment Type: Internet

Primary Stack

	Serial #	Instance ID	Appliance	Type	Bandwidth (Mbps)
Gateway	XND81418F04E1679	i-0d540d51b6fb633fa	SDI-AWS	t2.medium	120
SteelHead	XNDC9DE1D97505E7	i-061ef919f5f70e86f	--	t2.medium	30

Close

While not required, if you need to make specific configuration changes to the SteelHead for SteelConnect, you can log in to it from the browser of a jump host (Windows/Linux) in your VPC by navigating to its private IP address in the downlink subnet. The initial login credentials for the SteelHead are username admin and your specific instance ID is your password.

Figure 7-15. Configuration changes for the SteelHead for SteelConnect

The screenshot shows the AWS Management Console interface for the EC2 service. The left sidebar contains navigation links for EC2 Dashboard, Events, Tags, Reports, Limits, INSTANCES, IMAGES, ELASTIC BLOCK STORE, and NETWORK & SECURITY. The main content area displays a table of instances. Two instances are listed: 'SteelConnect_ap-south-1b_Virtual Gateway' and 'SteelConnect_ap-south-1b_SteelHead'. The 'SteelHead' instance is selected, and its details are shown below the table. The instance is running in the 'ap-south-1b' availability zone. The details section includes tabs for Description, Status Checks, Monitoring, Tags, and Usage Instructions. The Description tab is active, showing the instance ID, state, type, elastic IPs, availability zone, security groups, and scheduled events. The instance is associated with the 'SteelConnect' security group and has no scheduled events. The instance is running on the 'ebs_scsh_axel_custom_9.6.0a-scsh1_2_dev-919e7efa-e46d-479b-8ab8-3bb1a78f0e73-ami-b8f6d2c3.4 (ami-55b2c93a)' AMI.

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status
SteelConnect_ap-south-1b_Virtual Gateway	i-0d540d51b6fb633fa	t2.medium	ap-south-1b	running	2/2 checks ...	None
SteelConnect_ap-south-1b_SteelHead	i-061ef919f5f70e86f	t2.medium	ap-south-1b	running	2/2 checks ...	None

Instance: **i-061ef919f5f70e86f (SteelConnect_ap-south-1b_SteelHead)** Private IP: 10.35.255.246

Property	Value
Instance ID	i-061ef919f5f70e86f
Instance state	running
Instance type	t2.medium
Elastic IPs	-
Availability zone	ap-south-1b
Security groups	SteelConnect
Scheduled events	No scheduled events
AMI ID	ebs_scsh_axel_custom_9.6.0a-scsh1_2_dev-919e7efa-e46d-479b-8ab8-3bb1a78f0e73-ami-b8f6d2c3.4 (ami-55b2c93a)
Public DNS (IPv4)	-
IPv4 Public IP	-
IPv6 IPs	-
Private DNS	ip-10-35-255-246.ap-south-1.compute.internal
Private IPs	10.35.255.246
Secondary private IPs	-
VPC ID	vpc-9f716af6
Subnet ID	subnet-373cd77b

On AWS, you can review the uplink and downlink subnets for the SteelConnect gateway.

Figure 7-16. Uplink and downlink subnets

The screenshot shows the AWS Management Console interface for the VPC service. The left sidebar contains navigation links for VPC Dashboard, Virtual Private Cloud, Your VPCs, Subnets, Route Tables, Internet Gateways, Egress Only Internet Gateways, DHCP Options Sets, Elastic IPs, Endpoints, Endpoint Services, NAT Gateways, Peering Connections, and Security. The main content area displays a table of subnets. Two subnets are listed: 'SteelConnect_ap-south-1b_Uplink' and 'SteelConnect_ap-south-1b_Downlink'. The 'Uplink' subnet is selected, and its details are shown below the table. The subnet is available in the 'ap-south-1b' availability zone. The details section includes tabs for Summary, Route Table, Network ACL, Flow Logs, and Tags. The Summary tab is active, showing the subnet ID, availability zone, IPv4 CIDR, IPv6 CIDR, state, VPC, and available IPs. The subnet is associated with the 'rtb-3a7e0852 | SteelConnect_ap-south-1b_Internet Routes' route table and the 'acl-d23255ba | SteelConnect ACL' network ACL. The subnet is not the default subnet and does not auto-assign public IP or IPv6 address.

Name	Subnet ID	State	VPC	IPv4 CIDR	Available IPv4
SteelConnect_ap-south-1b_Uplink	subnet-863dd6ca	available	vpc-9f716af6 Engineering and T...	10.35.255.224/28	10
SteelConnect_ap-south-1b_Downlink	subnet-373cd77b	available	vpc-9f716af6 Engineering and T...	10.35.255.240/28	9
Engineering-1b	subnet-7e89d733	available	vpc-9f716af6 Engineering and T...	10.35.0.0/24	251
Engineering-1a	subnet-78992210	available	vpc-9f716af6 Engineering and T...	10.35.2.0/24	251

subnet-863dd6ca | SteelConnect_ap-south-1b_Uplink

Property	Value
Subnet ID	subnet-863dd6ca SteelConnect_ap-south-1b_Uplink
Availability Zone	ap-south-1b
IPv4 CIDR	10.35.255.224/28
IPv6 CIDR	-
State	available
VPC	vpc-9f716af6 Engineering and Tech Support
Available IPs	10
Route table	rtb-3a7e0852 SteelConnect_ap-south-1b_Internet Routes
Network ACL	acl-d23255ba SteelConnect ACL
Default subnet	no
Auto-assign Public IP	no
Auto-assign IPv6 address	no

The routes are added to all other networks in the route tables on AWS.

Figure 7-17. Route tables in AWS

VPC Dashboard

Filter by VPC:

Virtual Private Cloud

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

DHCP Options Sets

Elastic IPs

Endpoints

Endpoint Services

NAT Gateways

Peering Connections

Security

Network ACLs

Security Groups

VPN Connections

Customer Gateways

Virtual Private Gateways

VPN Connections

Create Route Table Delete Route Table Set As Main Table

Search Route Tables and their X

Name	Route Table ID	Explicitly Associated With	Main
	rtb-1a3a1173	0 Subnets	Yes
	rtb-5696503e	0 Subnets	Yes

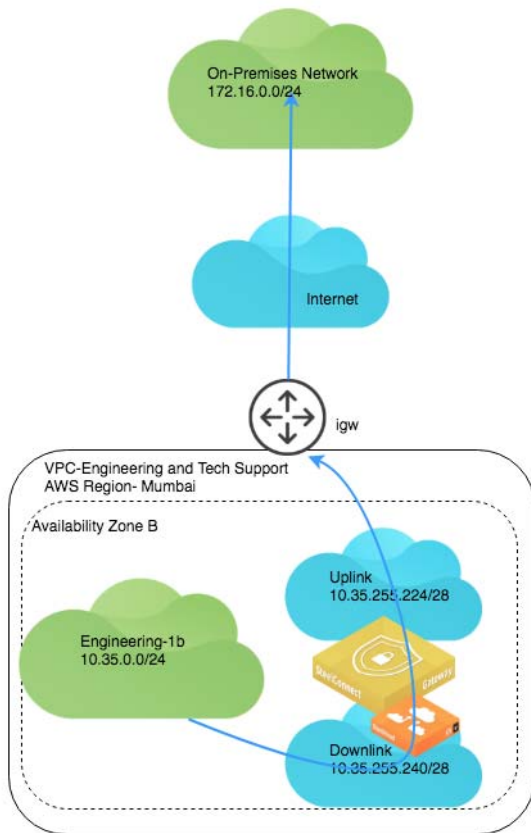
Edit

View: All rules

Destination	Target	Status	Propagated
10.35.0.0/16	local	Active	No
10.0.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.2.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.3.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.4.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.6.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.7.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.8.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.9.1.16/28	eni-6893c943 / i-061ef919f5f70e86f	Active	No
10.17.255.240/29	eni-1812c033	Black Hole	No
172.16.0.0/24	eni-6893c943 / i-061ef919f5f70e86f	Active	No
172.16.1.0/24	eni-6893c943 / i-061ef919f5f70e86f	Active	No
172.16.2.0/24	eni-6893c943 / i-061ef919f5f70e86f	Active	No
172.30.255.208/28	eni-1812c033	Black Hole	No

The architecture for the deployment is show in **Figure 7-18**.

Figure 7-18. AWS architecture with SteelHead



Transit VPC Routing using SteelConnect gateways

The AWS transit VPC feature in SteelConnect enables communication among multiple Virtual Private Clouds (VPCs) within the same AWS region by using a hub-and-spoke topology. One VPC in an AWS region is designated as the *hub* and all other VPCs in the region are *spokes*. All VPCs are connected to each other through the hub by using SteelConnect gateways and AutoVPN tunnels over VPC peering connections. The hub gateway is also automatically configured with an uplink to an AWS internet gateway, which enables AutoVPN tunnels to branch offices over the internet.

The benefits of this feature include:

- **Enhanced security** - The hub can be configured as a bastioned VPC, protecting its spokes behind a firewall or other security measures.
- **Overcoming AWS VPC peering limitations** - AWS VPC peering allows communication between two directly connected VPCs. SteelConnect AWS transit VPC enables network traffic to travel among multiple VPCs without a direct connection between them. Transitive peering and edge-to-edge routing are possible with SteelConnect AWS transit VPC.
- **Simpler operations** - It is much easier to configure and manage networks from a single hub rather than have a full-mesh infrastructure when applicable.

- **Visibility** - Since all traffic flows through the hub site, it is much easier to analyze this consolidated traffic and assess traffic patterns in your networks in a comprehensive manner.

The following features are interoperable and supported with AWS transit VPC topologies:

- High availability (redundancy) for hub gateways and spoke gateways.
- WAN optimization using SteelHead in both hub sites and spoke sites.

Your topology must meet the following requirements to use this feature:

- The hub VPC and all spoke VPCs must be in the same AWS region.
- The minimum instance size for the hub gateway is t2.small.
- In SteelConnect, the hub site must not be configured as an AutoVPN leaf of any other site.
- In SteelConnect, spoke sites must not be configured as a master site.
- In SteelConnect, the hub site must be deployed prior to deploying spoke sites.

Performing parallel deploy, undeploy, or manage operations on a transit VPC hub site or spoke site is not recommended.

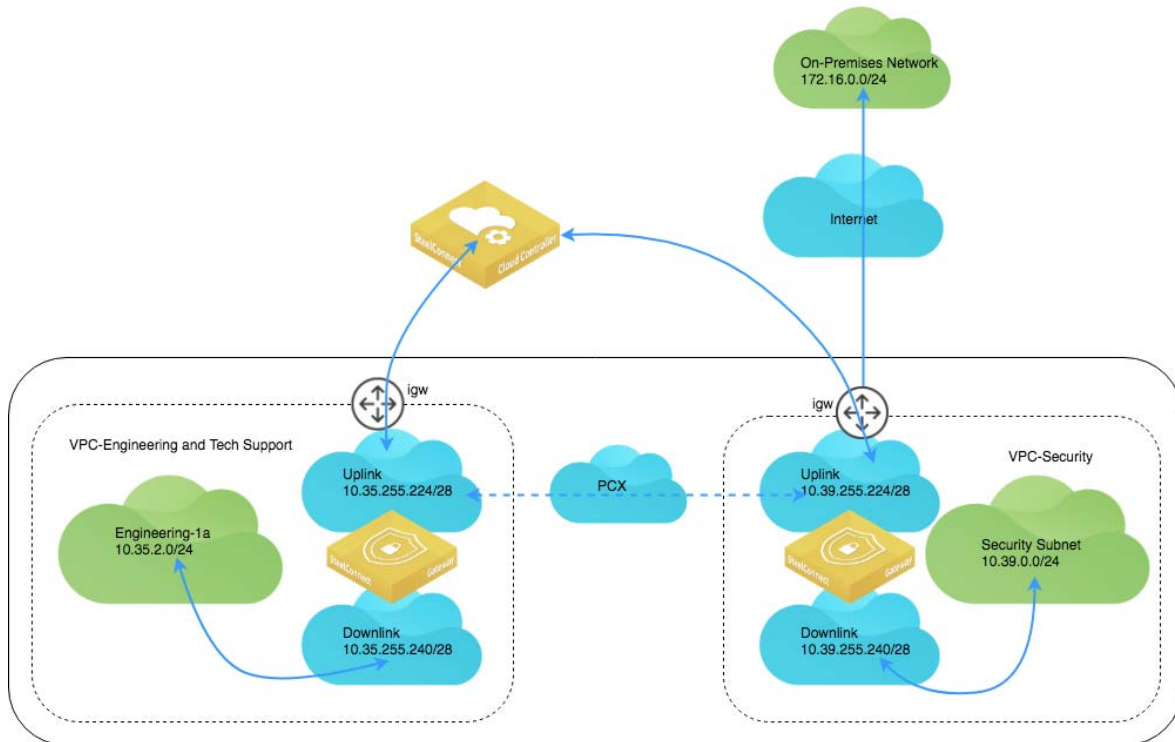
Deploying an AWS transit VPC topology

Before you deploy an AWS transit VPC topology, ensure that:

- you have identified which of your VPCs will serve as the hub site.
- you have identified which of your VPCs will serve as spoke sites.
- all VPCs are imported into SteelConnect Manager.
- the hub site must be deployed prior to deploying spoke sites.

As a sample topology, consider two VPCs: one for applications for Engineering and the other a Security VPC. **Figure 7-19** demonstrates a gateway deployed in the Security VPC as a hub site and the Engineering VPC as a spoke site.

Figure 7-19. Sample AWS transit VPC site

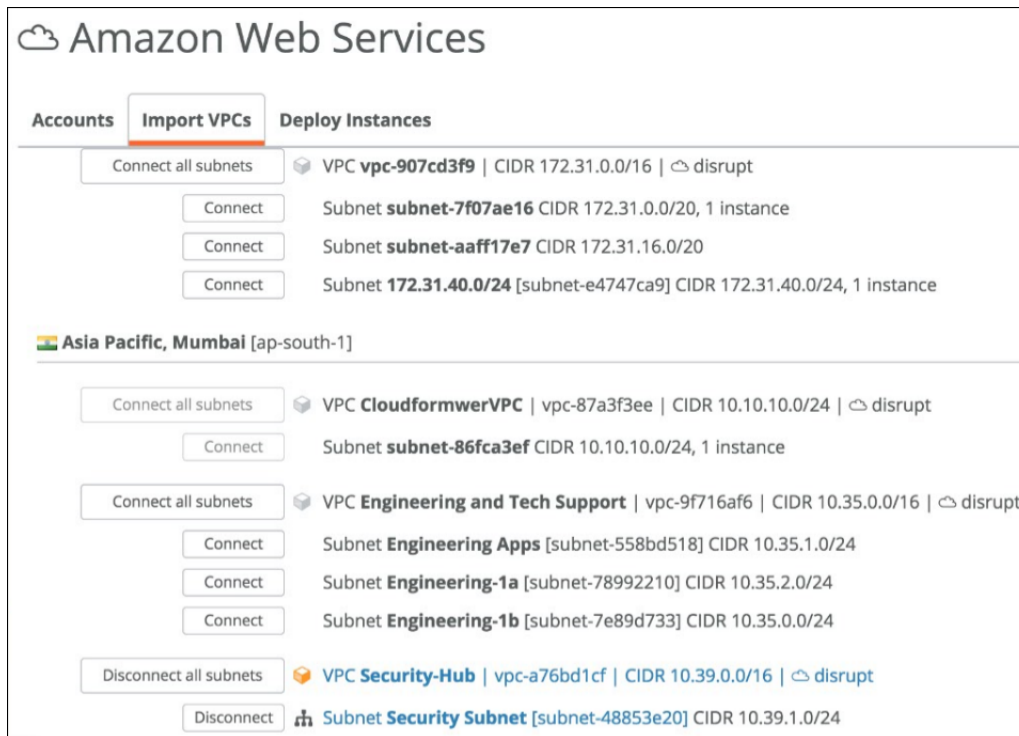


To deploy an AWS transit VPC topology with this sample topology

1. Connect the Security Hub VPC.

In SCM, choose Network Design > AWS. Select the Import VPCs tab and click **Connect** for the Security subnet.

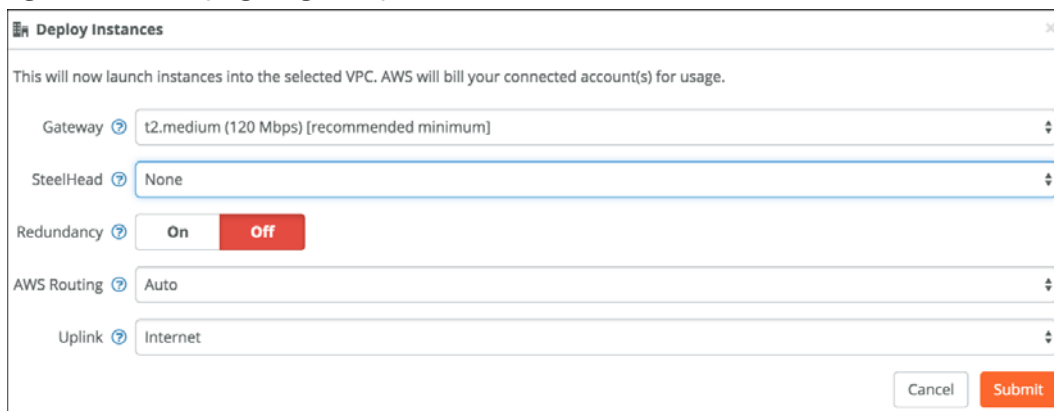
Figure 7-20. Connect to the VPC



2. Deploy the gateway instance for the hub site.

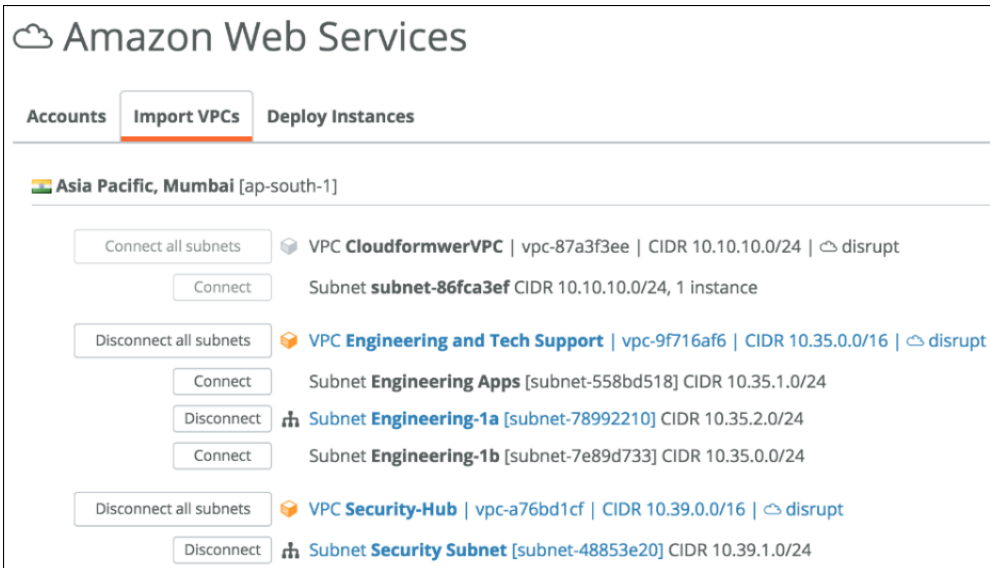
From the Deploy Instances tab, find the site to use as a hub and click **Deploy**.

Figure 7-21. Deploying the gateway for the hub site



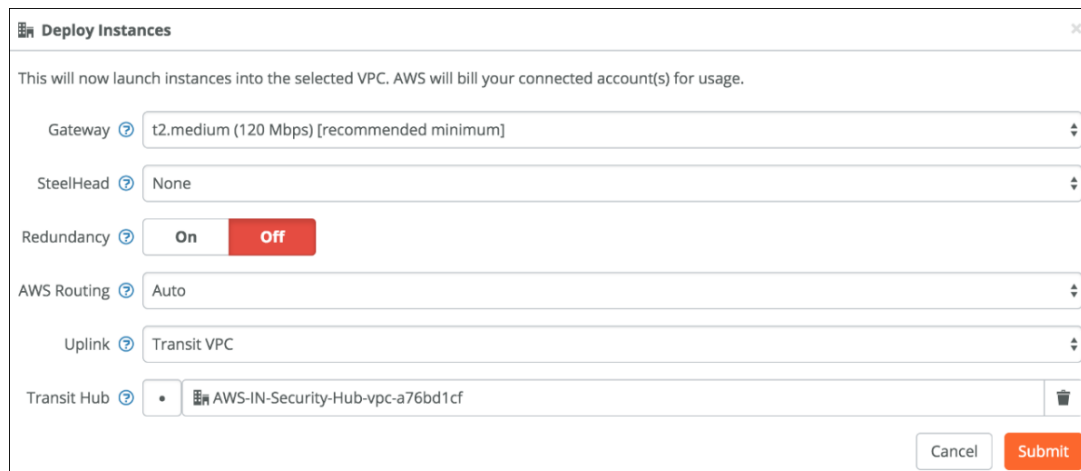
3. Connect the spoke subnet.

Figure 7-22. Connecting the spoke subnet



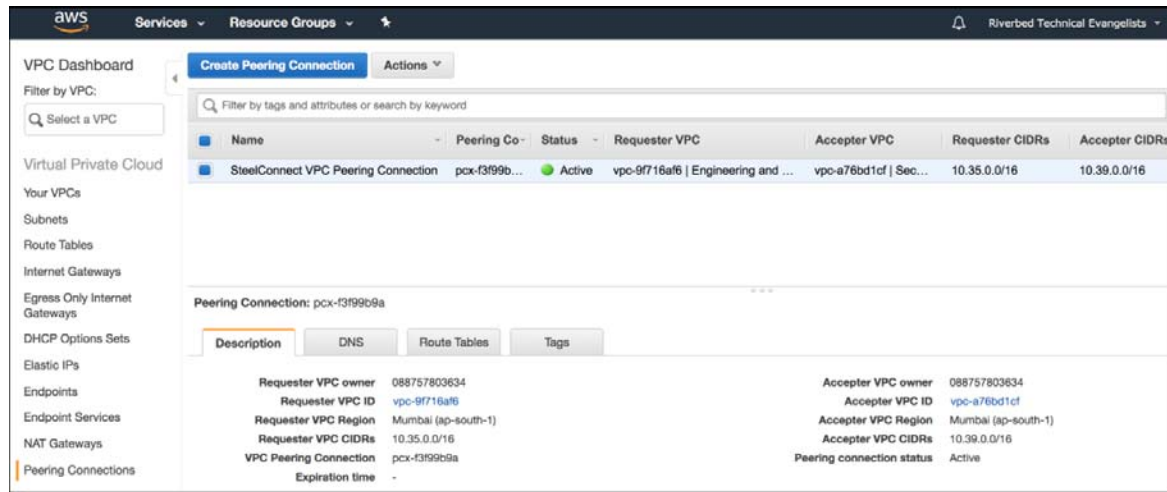
4. Deploy the gateway for the spoke subnet. From the Uplink drop-down, select Transit VPC and in the Transit Hub field set the VPC of your hub subnet. (The hub subnet should appear as an available option in the Transit Hub field.)

Figure 7-23. Deploying the spoke subnet



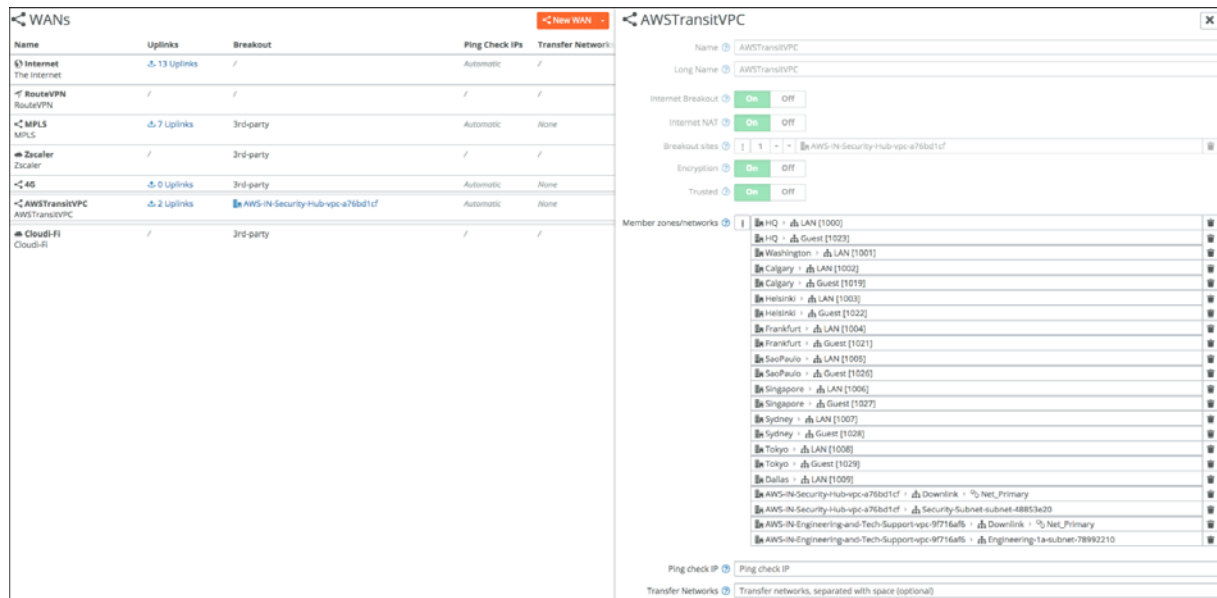
On AWS, the Peering Connection appears under the VPC Dashboard (SCM created this).

Figure 7-24. AWS peering connection



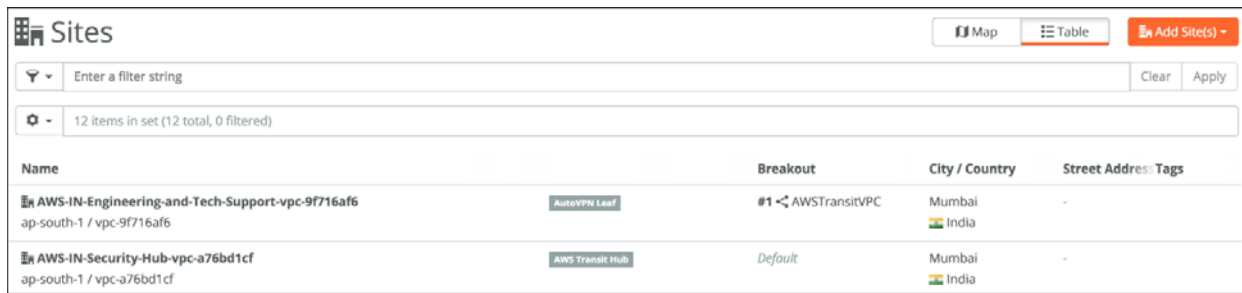
On SCM, a new WAN object for Transit VPC appears.

Figure 7-25. WAN object for Transit VPC



Under Sites, the two new AWS sites for the region you deployed in hub (AWS Transit Hub) and spoke (AutoVPN Leaf) mode appear.

Figure 7-26. AWS hub-and-spoke sites



Name	Breakout	City / Country	Street Address: Tags
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 ap-south-1 / vpc-9f716af6		#1 AWSTransitVPC Mumbai	-
AWS-IN-Security-Hub-vpc-a76bd1cf ap-south-1 / vpc-a76bd1cf		Default Mumbai	-

Under Appliances on SCM, you can see the new appliances deployed on SCM.

Figure 7-27. AWS appliances in SCM



Site	Model	Status	Config	VPN	Serial	License
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 ap-south-1b	SDI-AWS 		Up-to-date		XN8DCF07D71A7F0D	Not needed
AWS-IN-Security-Hub-vpc-a76bd1cf ap-south-1a	SDI-AWS 		Up-to-date		XNADE468ACF8FAF9	Not needed

From the AWS console, you can review the route tables for the Security subnet and Engineering subnets populated with all the routes for the Organization's networks (both on-premises and cloud).

Figure 7-28. Route table for Security subnet

The screenshot shows the AWS Management Console interface for the 'Security Subnet'. The left sidebar lists various AWS services, with 'Subnets' selected. The main content area shows the 'Route Table' for 'subnet-48853e20 | Security Subnet'. The route table is titled 'Route Table: rtb-5b473133' and contains the following routes:

Destination	Target
10.39.0.0/16	local
10.0.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.2.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.3.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.4.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.6.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.7.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.8.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.9.1.16/28	eni-2d44b273 / i-0bc01b1a87a0770a1
10.35.2.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
10.35.255.240/28	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.0.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.1.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.2.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.3.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.4.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.5.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.6.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1
172.16.7.0/24	eni-2d44b273 / i-0bc01b1a87a0770a1

Figure 7-29. Route table for Engineering subnet

The screenshot shows the AWS Management Console interface for the route table of subnet-78992210 (Engineering-1a). The left sidebar lists various AWS services, with 'Route Tables' selected. The main panel shows the 'Route Table' tab for 'rtb-5696503e'. Below the tabs, a table lists the routes. All routes have a 'local' target, indicating they are for local delivery within the VPC.

Destination	Target
10.35.0.0/16	local
10.0.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.2.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.3.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.4.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.6.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.7.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.8.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.9.1.16/28	eni-fda9f3d6 / i-04e78bf4868483c10
10.17.255.240/29	eni-1812c033
10.39.1.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
10.39.255.240/28	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.0.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.1.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.2.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.3.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.4.0/24	eni-fda9f3d6 / i-04e78bf4868483c10
172.16.5.0/24	eni-fda9f3d6 / i-04e78bf4868483c10

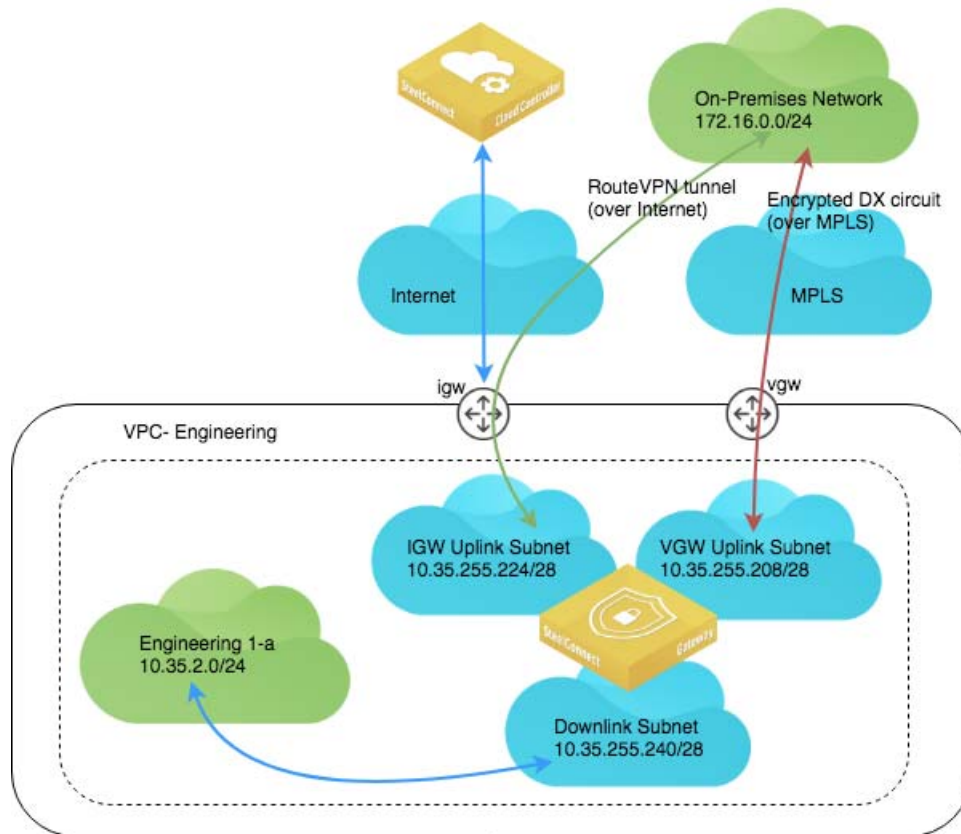
Note: There is no default route for the subnets. You need to add a default route if you want your subnets to reach the internet. To do so, click **Edit** for the Route Table and add an entry for the Destination 0.0.0.0/0 Target IGW (Amazon’s internet gateway) if you want to send the traffic direct to internet, or pcx (peering connection) if over the peering connection (recommended for traffic from the Engineering to Security subnet), or to the Uplink interface of the local gateway.

Deployment over AWS DirectConnect

For customers that already have an existing AWS Direct Connect (DX) circuit and want to extend their on-premises networks into their AWS cloud network over this private DX connection, SteelConnect can provide a secure overlay over the AWS backbone. This allows for SD-WAN features to be implemented at sites that are connected via Direct Connect.

In this deployment, at the on-premises branch office, the SteelConnect gateway is configured with two uplinks: one for the internet connection and the other for the private high-bandwidth DX connection. On AWS, you need to create a virtual gateway (VGW) in the appropriate region and then associate it with the DX circuit along with the VPC that is to be deployed from SCM.

Figure 7-30. AWS Direct Connect architecture



When Direct Connect is selected as an Uplink, the system selects a free /26 subnet and breaks it into 3 /28's, (leaving one of the /28 subnets unused out of the four). One /28 subnet is allocated for the IGW uplink, and another for the VGW uplink used with Direct Connect. The downlink subnet consumes the last /28 subnet.

It is possible to have several VPCs connected to a central hub using the transit VPC OR normal IGW Leaf options, and then use the hub as the dual-linked gateway with Internet + Direct Connect uplinks.

If you use manual routing, no route table entries will be created in by SCM.

When automatic routing is selected, specific routes will be written to the AWS route tables of connected subnets in order to attract traffic. Static routes and propagated routes from the VGW Direct Connect, and will take priority, thus existing routes do not need to be deleted.

Note: In SCM 2.10, the WAN that is selected to be used with Direct Connect must have encryption (overlay) enabled. Underlay WANs are not supported.

To deploy AWS Direct Connect

Note: The *SteelConnect Manager User Guide* has detailed steps for Direct Connect configuration.

1. Connect to the Direct Connect VPC.
2. Deploy the gateway and specify the Uplink to be Internet + Direct Connect.

Figure 7-31. Setting the gateway uplink type for Direct Connect

The screenshot shows the 'Deploy Instances' dialog box. It contains the following fields and options:

- Gateway:** t2.medium (120 Mbps) [recommended minimum]
- SteelHead:** None
- Redundancy:** On (selected), Off
- AWS Routing:** Auto
- Uplink:** A dropdown menu is open, showing three options: Internet (checked), Internet + Direct Connect (highlighted in blue), and Transit VPC.

Buttons for 'Cancel' and 'Submit' are visible at the bottom right.

3. In the Direct Connect WAN field, specify an overlay network.

The Direct Connect WAN field appears once you specify Internet + Direct Connect as the Uplink.

SCM only lets you connect to encrypted WANs and only encrypted WANs appear in this field.

In our sample topology, the DX circuit is over an encrypted MPLS connection.

Figure 7-32. Direct Connect WAN

The screenshot shows the 'Deploy Instances' dialog box with the Uplink set to 'Internet + Direct Connect'. The 'Direct Connect WAN' field is now visible and populated with 'MPLS'. The fields and options are:

- Gateway:** t2.medium (120 Mbps) [recommended minimum]
- SteelHead:** None
- Redundancy:** On (selected), Off
- AWS Routing:** Auto
- Uplink:** Internet + Direct Connect
- Direct Connect WAN:** A field containing 'MPLS' with a trash icon to its right.

'Cancel' and 'Submit' buttons are at the bottom right.

When you deploy, SCM logs in to your AWS console and creates an overlay tunnel over your existing DX circuit from your VPC.

SCM creates two uplinks for the new AWS site: one over the internet and one over the DX circuit.

Figure 7-33. Direct Connect uplinks

Name	WAN	Operation	Gateway
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 > DirectConnect DHCP client	MPLS	On Off	SDI-AWS Not set
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 > Uplink DHCP client	Internet	On Off	SDI-AWS Not set

Once deployed, the AWS gateway appears under Appliances. When you click it and view the IPs tab, you will see two IP addresses for the two Uplink connections over the IGW and VGW.

Figure 7-34. Direct Connect gateway IP addresses

Site	Model	Status	Config	VPN
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6	SDI-AWS ap-south-1b	Online	Up-to-date	AutoVPN
Calgary	SDI-VGW Montreal	Online	Up-to-date	AutoVPN

Uplink	IP
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 > Uplink	10.35.255.235
AWS-IN-Engineering-and-Tech-Support-vpc-9f716af6 > DirectConnect	10.35.255.219

In AWS, you will see an extra uplink subnet created for the DX connection, which routes over Amazon's VGW.

Figure 7-35. Uplink subnet in AWS

Name	Subnet ID	State	VPC	IPv4 CIDR	Available IPv4
SteelConnect_ap-south-1b_Uplink	subnet-6545af29	available	vpc-9f716af6 Engineering and T...	10.35.255.224/28	10
SteelConnect_ap-south-1b_Downlink	subnet-c144ae8d	available	vpc-9f716af6 Engineering and T...	10.35.255.240/28	10
SteelConnect_ap-south-1b_Direct_Conne...	subnet-795ab035	available	vpc-9f716af6 Engineering and T...	10.35.255.208/28	10

Azure Cloud Topologies

SteelConnect offers strong integration with Microsoft Azure. You can quickly and easily deploy virtual SteelConnect gateways in Azure to create a secure full-mesh virtual private network (VPN) fabric between your sites and/or VNets anywhere in the world. Connect your Azure account to SCM, and SteelConnect will find all your subnets, in all networks, in all regions. Importing an entire network, or individual subnets, into SCM is as easy as clicking a button. You can then deploy virtual SteelConnect gateways in Azure—with optional SteelHead WAN optimization—onto your SCM-managed subnets to build an interconnected, full-mesh VPN.

Note: Redundancy is currently only available on AWS.

Note: AWS, Azure, and SteelConnect use slightly different terms to refer to similar network concepts. Virtual Private Clouds (AWS) and VNets (Azure) are called *sites* in SteelConnect, and subnets are referred to as *zones* in SteelConnect.

Deployment into the Azure cloud

The configuration steps to deploy SteelConnect gateways in the Azure cloud are outlined in the *SteelConnect Manager User Guide*. The steps to do so are:

1. Subscribe to Riverbed Azure products.
2. Configure your Azure accounts with SteelConnect Manager.
3. Import Azure networks.
4. Deploy SteelConnect gateways to your Azure network.

Azure also offers advanced features, including:

- Azure ExpressRoute to configure a private network connection to your cloud deployments

Note: To integrate your Azure account with SCM for cross-account access, you will also need to follow the steps outlined in this knowledge base article: <https://supportkb.riverbed.com/support/index?page=content&id=s29078>

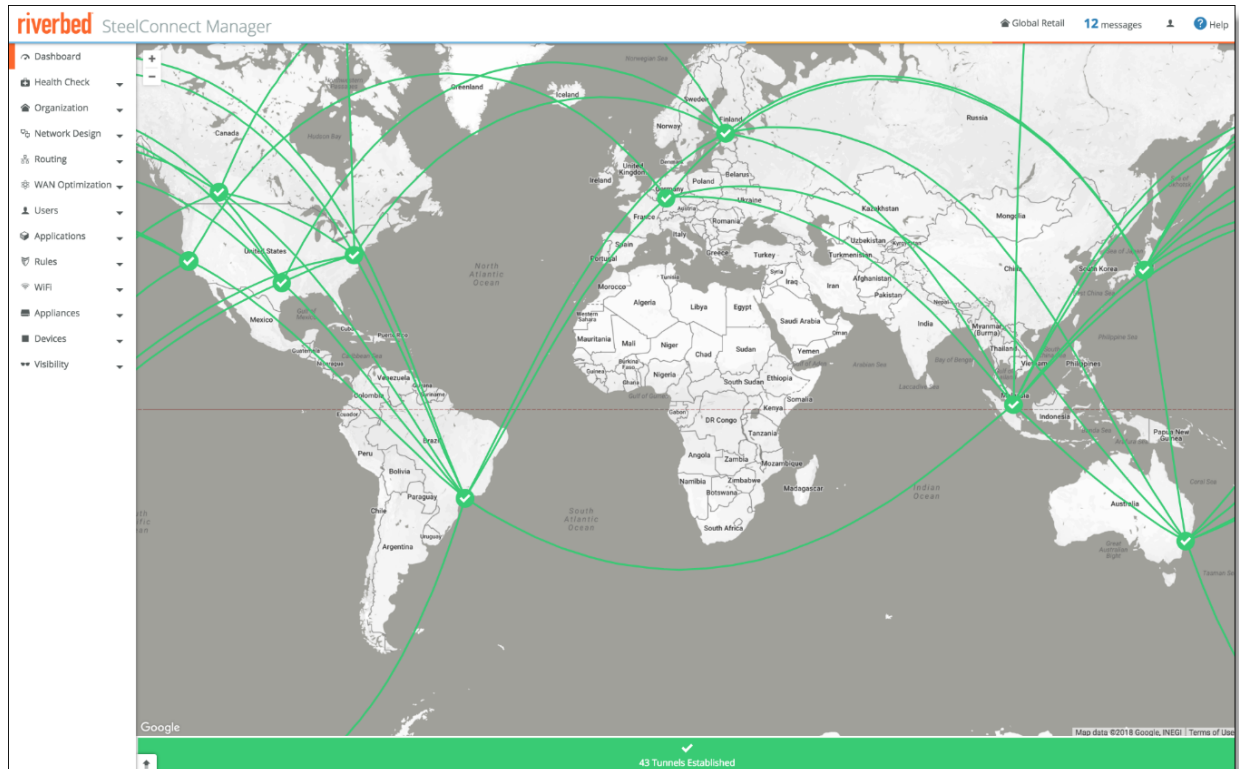
Once you have read through the above steps in the *SteelConnect Manager User Guide* on how to configure your deployment into Azure, read through the example deployment in this chapter as well as the routing notes and SteelHead with Azure notes.

SteelConnect gateway in Azure

This section steps through a sample deployment of SteelConnect gateways in Azure.

This example starts with an on-premises topology for an organization that has a cloud data center network in the Ireland region of Azure.

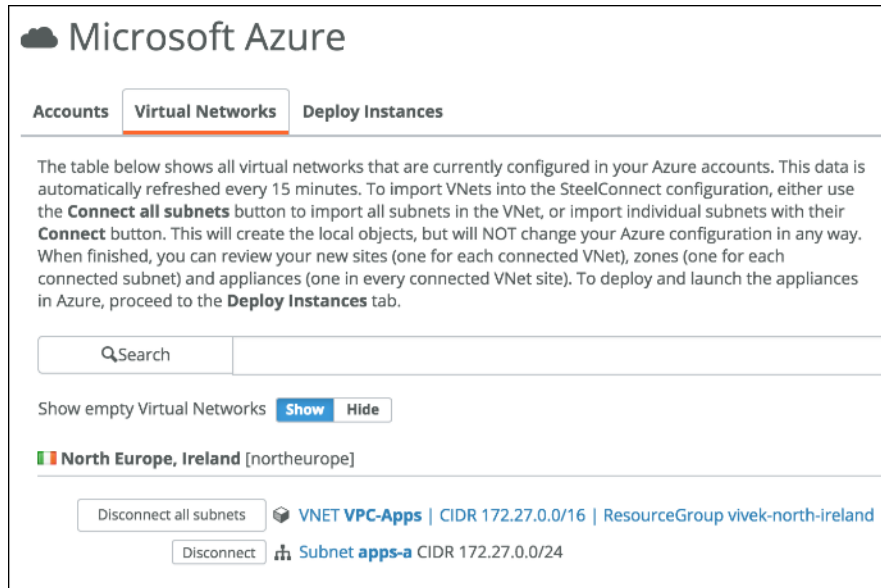
Figure 8-1. Azure network sample dashboard



To connect all sites to the Azure cloud

1. Choose Network Design > Azure and then select the Virtual Networks tab. Click **Connect All** for the North Europe, Ireland region.

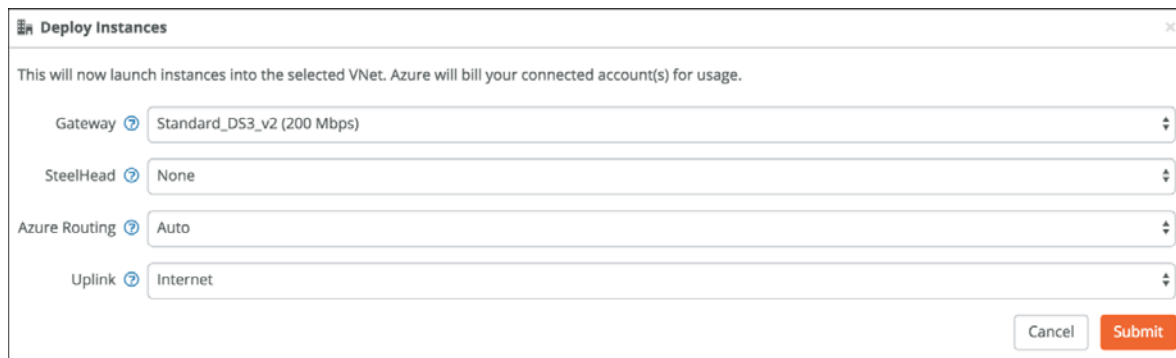
Figure 8-2. Connect all Azure virtual networks



2. Select the Deploy Instances tab and click **Deploy** next to an instance.

Choose the parameters you want. Note, there is no option to deploy a gateway in Azure in HA mode in SCM version 2.10.

Figure 8-3. Azure deployment options




During deployment, log messages detail the status.

Figure 8-4. Deployment log messages

2018-03-12 15:50:56	7s ago	SDI-AZURE [Dublin]	Uplink Azure-IE-VPC-Apps-vivek-north-ireland > Uplink came online
2018-03-12 15:50:56	7s ago	SDI-AZURE [Dublin]	Uplink Azure-IE-VPC-Apps-vivek-north-ireland > Uplink is now in use
2018-03-12 15:50:56	7s ago	SDI-AZURE [Dublin]	External IPv4 address change on Uplink Azure-IE-VPC-Apps-vivek-north-ireland > Uplink (new IP 52.178.162.121)
2018-03-12 15:50:33	30s ago	SDI-AZURE [Dublin]	Appliance has downloaded firmware version 2.10.0.50-paddington.
2018-03-12 15:50:16	47s ago	SDI-AZURE [Dublin]	Appliance is downloading new firmware
2018-03-12 15:50:16	47s ago	SDI-AZURE [Dublin]	Newer firmware available, sending FIRMWARE_DOWNLOAD (2.10.0.50-paddington)
2018-03-12 15:49:17	106s ago	SDI-AZURE [Dublin]	Appliance rebooting
2018-03-12 15:49:16	107s ago	SDI-AZURE [Dublin]	Appliance is downloading new firmware
2018-03-12 15:49:15	108s ago	SDI-AZURE [Dublin]	Newer firmware available, sending FIRMWARE_DOWNLOAD (2.10.0.50-paddington)

Figure 8-5 shows the deployment architecture.

Figure 8-5. Azure deployment architecture



The Azure gateway appears in the Appliances Overview page on SCM.

Figure 8-6. Azure in Appliances Overview

The screenshot shows the 'Appliances Overview' page with a table of appliances. The 'SDI-AZURE' gateway is highlighted. To the right, a detailed view of the 'SDI-AZURE' gateway is shown, including its status (Online), configuration (Up-to-date), and a list of events. The events list shows various AutoVPN tunnels between different sites and the gateway, as well as updates to the gateway's configuration and status.

On the Azure portal, you can confirm the deployment of the gateway along with all its associated resources as shown in Figure 8-7. You can see the routing table and review the routes to all your sites in your network.

Figure 8-7. Azure portal with routing details

The screenshot shows the Microsoft Azure portal interface. The left sidebar contains navigation options like 'Dashboard', 'All resources', 'Resource groups', 'App Services', 'SQL databases', 'SQL data warehouses', 'Azure Cosmos DB', 'Load balancers', 'Storage accounts', 'Virtual machines', 'Virtual networks', 'Azure Active Directory', 'Monitor', 'Advisor', 'Security Center', 'Cost Management + Billing', 'Help + support', and 'Subscriptions'. The main area displays the 'rvbd-scgw-rt-ad3f17' routing table. The table lists various routes with their names, address prefixes, and next hops. The 'rvbd-scgw-rt-ad3f17' route is highlighted. Below the routing table, there is a section for 'Subnets' showing the 'apps-a' subnet with its address range, virtual network, and security group.

With this deployment, there is no default route to the internet. You need to manually add that to your routing table if you want your internal Azure subnet to be able to go out to the internet.

You can review the uplink and downlink interfaces of the gateway under All resources from the Azure portal.

Figure 8-8. Uplink interface in the Azure console

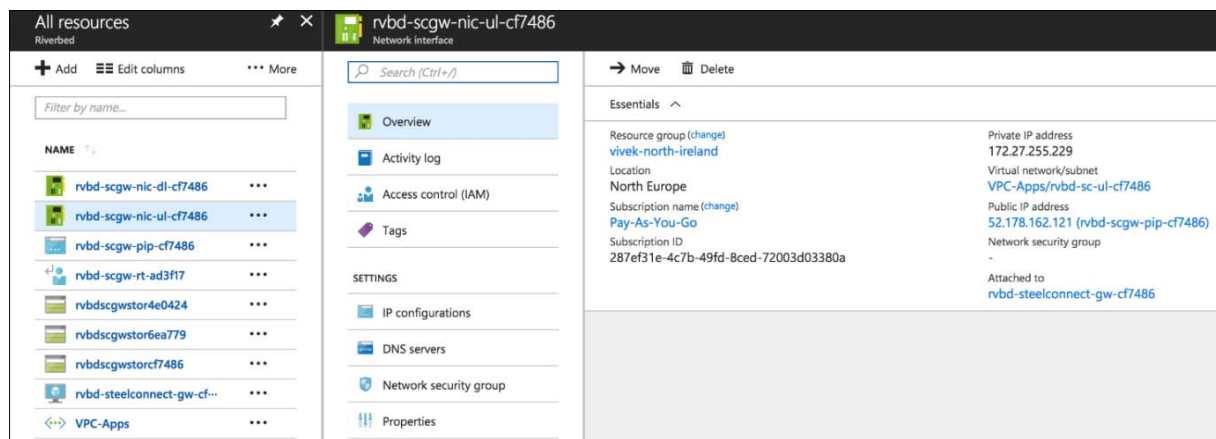
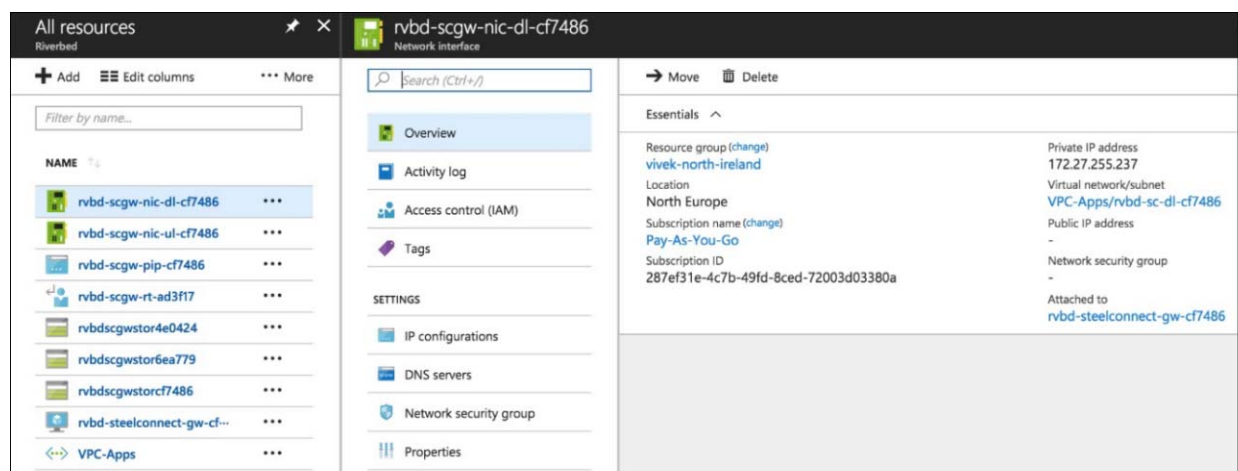
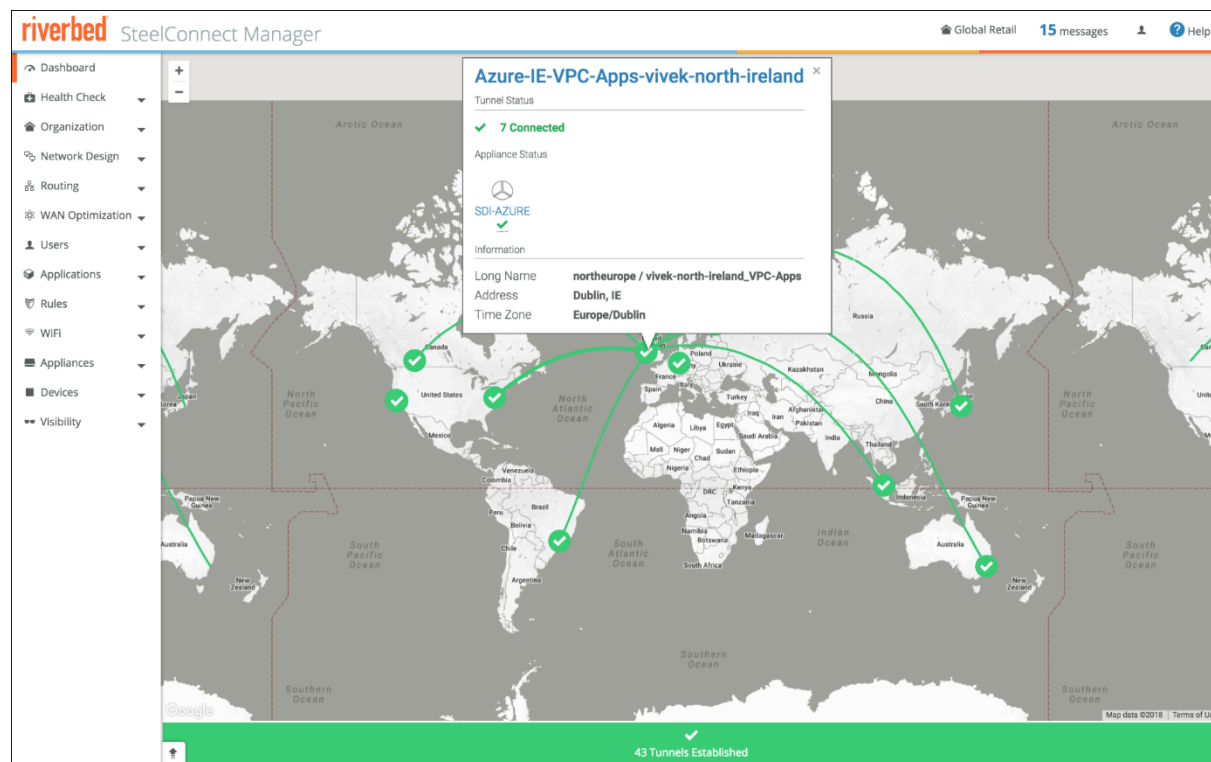


Figure 8-9. Downlink interface in the Azure console



When fully deployed, the gateway appears as online in the Azure region and the tunnels appear from the newly provisioned cloud site to your other sites (cloud or on-premises) depending on your network configuration (leaf mode or full-mesh).

Figure 8-10. Dashboard with new Azure connections



Deploying SteelConnect gateway in Azure with SteelHead

Similar to the deployment of gateway into AWS, you can also deploy SteelHead in Azure. SteelHead WAN optimization enhances your SteelConnect Infrastructure as a Service (IaaS) deployment by automatically optimizing all TCP traffic to increase speed and reduce bandwidth over the secure VPN links. In addition to interconnecting your cloud infrastructure to your on-premises sites (or bringing the cloud closer by providing a direct path to your IaaS applications from each branch) you can now use SteelHead to boost performance and reduce traffic.

To deploy a SteelConnect gateway in Azure with SteelHead optimization, step through the standard deployment process and specify a specification for the SteelHead deployment. The specifications include throughput limits that indicates the amount of optimized traffic the instance supports.

Note: You can add this feature during initial deployment or at any time.

Figure 8-11. Adding SteelHead optimization to a gateway in Azure

Deploy Instances

This will now launch instances into the selected VNet. Azure will bill your connected account(s) for usage.

Gateway ⓘ Standard_DS3_v2 (200 Mbps)

SteelHead ⓘ Standard_DS2_v2 (100 Mbps)

Azure Routing ⓘ Auto

Uplink ⓘ Internet

Cancel Submit

When deployed, the VNET details for appliances appear in SCM.

Figure 8-12. Gateway and SteelHead appliance in Azure dashboard

VNET Details

VNET ID: vivek-west-us_Apps-Engineering

Azure Routing: Auto


Deployment Type: Internet

	Serial #	Instance ID	Appliance	Type	Bandwidth (Mbps)
Gateway	XNE68E0319CF293A	rvbd-steelconnect-gw-d7c5d4	SDI-Azure	Standard_DS3_v2	200
SteelHead	XN90675B617B6536	rvbd-steelconnect-sh-5279de	--	Standard_DS2_v2	100

✓ Close

The deployed appliances also appear in the Appliances pages of SCM. Select the gateway and click the SteelHead tab to review the details of the deployed SteelHead instance.

Figure 8-13. SteelHead tab for a gateway



The deployed topology matches that of AWS. The gateway and SteelHead appliances appear in the Azure portal in the Resources section.

Figure 8-14. Azure portal resources

<input type="checkbox"/>	 rvbd-steelconnect-gw-d7c5d4	Virtual machine	vivek-west-us	West US
<input type="checkbox"/>	 rvbd-steelconnect-sh-5279de	Virtual machine	vivek-west-us	West US

Deploying over Azure ExpressRoute

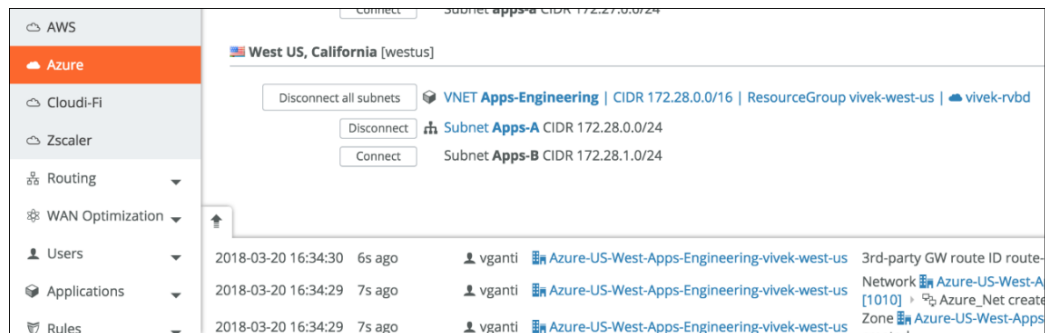
You can use Azure ExpressRoute to create private connections between Azure data centers and infrastructure on your premises or in a colocation environment. The ExpressRoute topology and workflow is the same as with AWS Direct Connect.

Note: The *SteelConnect Manager User Guide* has detailed steps for Azure ExpressRoute configuration.

To deploy Azure ExpressRoute

1. Connect to a subnet in an Azure VNET.

Figure 8-15. Connect a VNET



2. Deploy the gateway and specify the Uplink to be Internet + ExpressRoute.

Figure 8-16. Setting the gateway uplink type for ExpressRoute

3. In the ExpressRoute WAN field, specify an overlay network.

The ExpressRoute WAN field appears once you specify Internet + ExpressRoute as the Uplink.

SCM only lets you connect to encrypted WANs and only encrypted WANs appear in this field.

In our sample topology, the ExpressRoute connection is over an encrypted MPLS connection.

When you deploy, SCM creates two uplinks for the new AWS site: one over the internet and one over the MPLS WAN.

Figure 8-17. ExpressRoute uplinks

	Name	WAN	Operation	Gateway
Split site/Dual-hub				
WANS	Azure-US-West-Apps-Engineering-vivek-west-us > ExpressRoute Static IP	MPLS	On Off	Unassigned
Uplinks	Azure-US-West-Apps-Engineering-vivek-west-us > Uplink DHCP client	Internet	On Off	SDI-AZURE Not set
Zones				

Once deployed, the Azure gateway appears under Appliances. When you click it and view the IPs tab, you will see two IP addresses for the two Uplink connections.

Figure 8-18. ExpressRoute IP addresses

Live

IPs

AutoVPN

Debug

SteelHead

Uplink (external) IPs

These are the reported uplink IP addresses of the gateway. Only actual IPs reported by the gateway are shown! Note: This is an information page only. Uplinks can be configured on the 'Uplinks' overview page and attached to the gateway on the 'Ports' overview page.

Uplink	IP	External IP
Azure-US-West-Apps-Engineering-vivek-west-us > Uplink	172.28.255.229	13.93.203.148
Azure-US-West-Apps-Engineering-vivek-west-us > ExpressRoute	172.28.255.253	13.93.204.180

Management IP

This is the reported management IP address of the panther. Only actual IP reported by the panther is shown! Note: This is an information page only.

Port	IP
------	----

Zone gateway assignment (internal) IPs

These are the current internal IP addresses of the gateway, as set with zone gateway assignments. Note: This is an information page only. Adding, removing or changing gateway assignments and their IPs is done on the 'Zones' overview page.

Zone	Role	IPs	Flags
Azure-US-West-Apps-Engineering-vivek-west-us > Downlink_Azure [1009]	Member Gateway	172.28.255.237 -	DNS Server

On the Azure portal, two uplinks for the gateway and one for the downlink appear.

Figure 8-19. Uplinks and downlink in Azure portal

Home > All resources			
All resources			
Riverbed			
+ Add Edit columns Refresh Assign Tags Delete			
Subscriptions: Pay-As-You-Go – Don't see a subscription? Switch directories			
<input type="text" value="Filter by name..."/> All resource groups All types All locations			
13 items <input type="checkbox"/> Show hidden types			
<input type="checkbox"/> NAME ↑↓	TYPE ↑↓	RESOURCE GROUP ↑↓	LOCATION
<input type="checkbox"/> Apps-Engineering	Virtual network	vivek-west-us	West US
<input type="checkbox"/> rvbd-scgw-nic-dl-f70983	Network interface	vivek-west-us	West US
<input type="checkbox"/> rvbd-scgw-nic-ul2-f70983	Network interface	vivek-west-us	West US
<input type="checkbox"/> rvbd-scgw-nic-ul-f70983	Network interface	vivek-west-us	West US
<input type="checkbox"/> rvbd-scgw-pip-f70983	Public IP address	vivek-west-us	West US

From the Azure portal, you can also review the route table on Azure populated with routes to all of your sites.

Figure 8-20. Route table on Azure portal

NAME	ADDRESS PREFIX	NEXT HOP
rvbd-scgw-route-1eea81	10.3.1.16/28	172.28.255.237
rvbd-scgw-route-1ff8b5	172.16.1.0/24	172.28.255.237
rvbd-scgw-route-2213a2	10.6.1.16/28	172.28.255.237
rvbd-scgw-route-36ab29	10.9.1.16/28	172.28.255.237
rvbd-scgw-route-3bf053	10.7.1.16/28	172.28.255.237
rvbd-scgw-route-695fe7	172.16.0.0/24	172.28.255.237
rvbd-scgw-route-79f32a	10.4.1.16/28	172.28.255.237
rvbd-scgw-route-7ff58f	10.8.1.16/28	172.28.255.237
rvbd-scgw-route-a1e3fa	172.16.4.0/24	172.28.255.237
rvbd-scgw-route-b4e50d	10.2.1.16/28	172.28.255.237
rvbd-scgw-route-e4e949	10.0.1.16/28	172.28.255.237

Route-tables and user-defined routes (UDRs)

When using automatic routing, the connected subnets route table is modified by SCM automatically to control the way traffic is routed to other SteelConnect sites. The SteelHead for SteelConnect automatically uses the SteelConnect Gateway as its upstream exit point, and when present, the routes point to the SteelHead interface that traffic is routed through the SteelHead before being forwarded to the SteelConnect gateway for transit. To allow the network interfaces to receive and forward traffic, SCM enables IP forwarding automatically for the SteelConnect gateway and SteelHead VNets in Azure.

Three separate subnets are created by SteelConnect: the gateway has interfaces in the uplink and downlink subnet while the SteelHead is in a third subnet. The UDR configured on the SteelHead subnet makes sure traffic outbound from the SteelHead is redirected to the gateway. You can see the gateway name assignment in the SCM zones menu (see [Figure 8-21](#)). The downlink subnet is created in Azure by Riverbed to house SteelConnect instances and infrastructure without interfering with your cloud infrastructure in any way. The downlink subnet is displayed in SCM for clarity. There's no need to change anything nor should you deploy your own Azure instances in this subnet.

Figure 8-21. The SteelConnect gateway running in Azure is displayed in the SCM zones

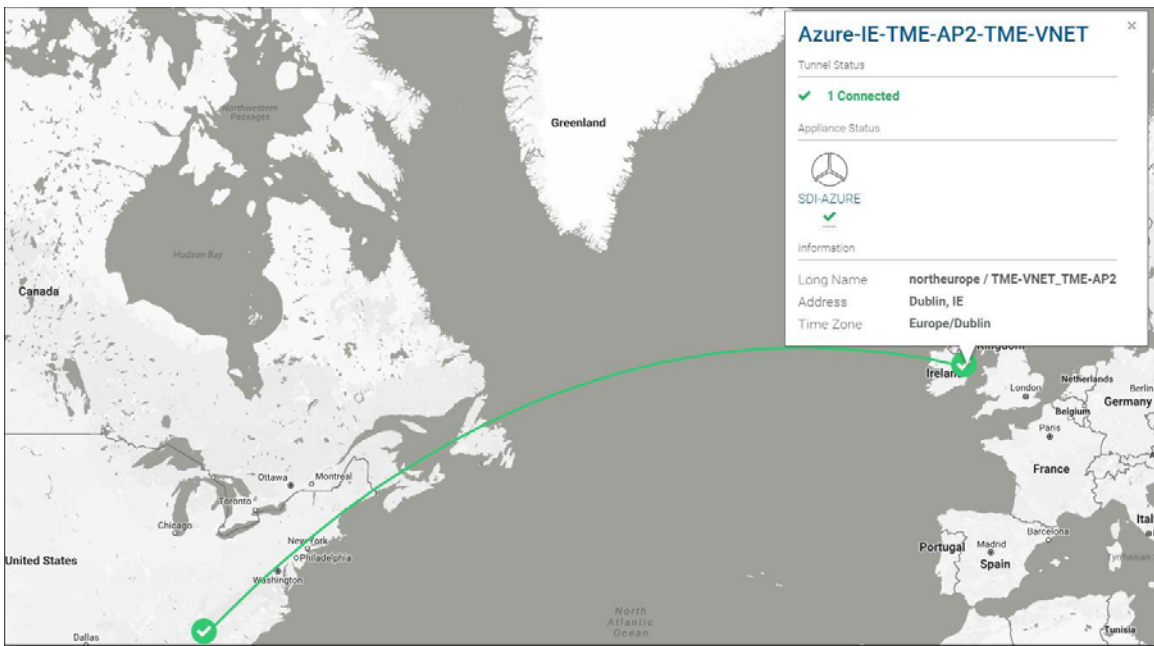
Zone	Gateways	WANs	Breakout IPv4	IPv6
Azure-IE-TME-AP2-TME-VNET > Downlink_Azure VLAN 1001	Azure-IE-TME-AP2-TME-VNET	RouteVF Default	10.17.255.240/-	-
Azure-IE-TME-AP2-TME-VNET > default-TME-AP2-TME-VNET VLAN 1002	No SteelConnect gateway 1 route	RouteVF Default	10.17.1.0/24	-
Azure-IE-TME-AP2-TME-VNET > Backend-TME-AP2-TME-VNET VLAN 1004	No SteelConnect gateway 1 route	RouteVF Default	10.17.2.0/24	-

SCM requires a /27 subnet space to create the subnets needed to house the gateway and the SteelHead. SteelConnect subnets are carved out from the overall address space assigned to individual VNets. SCM needs enough space to create these subnets.

For example, if the address space assigned to a particular VNet is configured as 10.10.0.0/16, and there are 255 (/24) client subnets ranging from 10.10.1.0/24 through 10.10.254.0/24, SCM can spin /27 downlink, uplink, and SteelHead subnets in the nonallocated address space 10.10.255.0/24. However, if all 256 (/24) subnets are consumed by client subnets, SCM doesn't have enough space to spin up the uplink and downlink subnets and displays a "No free subnets available" error.

In **Figure 8-22**, the SCM dashboard shows one other site at Atlanta, Georgia to which the Azure (Dublin) site is connected.

Figure 8-22. Green line indicates VPN tunnel is successfully established



To handle traffic between your Azure client subnets and the internet, you can create a user-defined route in Azure. When different route types are present in a UDR route table, user-defined routes are preferred over the default system routes. Each subnet can only be associated to a single route table. All VMs and cloud services in a subnet use the route table associated to that subnet. For further details on creating user-defined routes in Azure go to:

<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-udr-overview>

In general, when working with gateway subnets, avoid associating a network security group (NSG) to the gateway subnet. Associating an NSG to a gateway subnet might cause your VPN gateway to stop functioning as expected.

Manual routing

When deploying the SteelConnect gateway, the manual routing option can be enabled for a VNET/site by setting the Azure routing to manual. When this option is enabled, the route tables of the connected subnets will not be modified by SCM and you must add the required routes.

Manual routing support is useful for deployments that involve a large number of sites. Because there is a hard limit on the number of routes that can be added to an Azure route table (default limit is 50 routes per route table), SCM will not be able to add all the required routes. In such a case, you can use the manual routing option and manually add the summarized routes to the AWS route tables. This option is also useful if you want to control what gets routed over the SteelConnect AutoVPN with your own scripts and designs.

End-to-End Topology

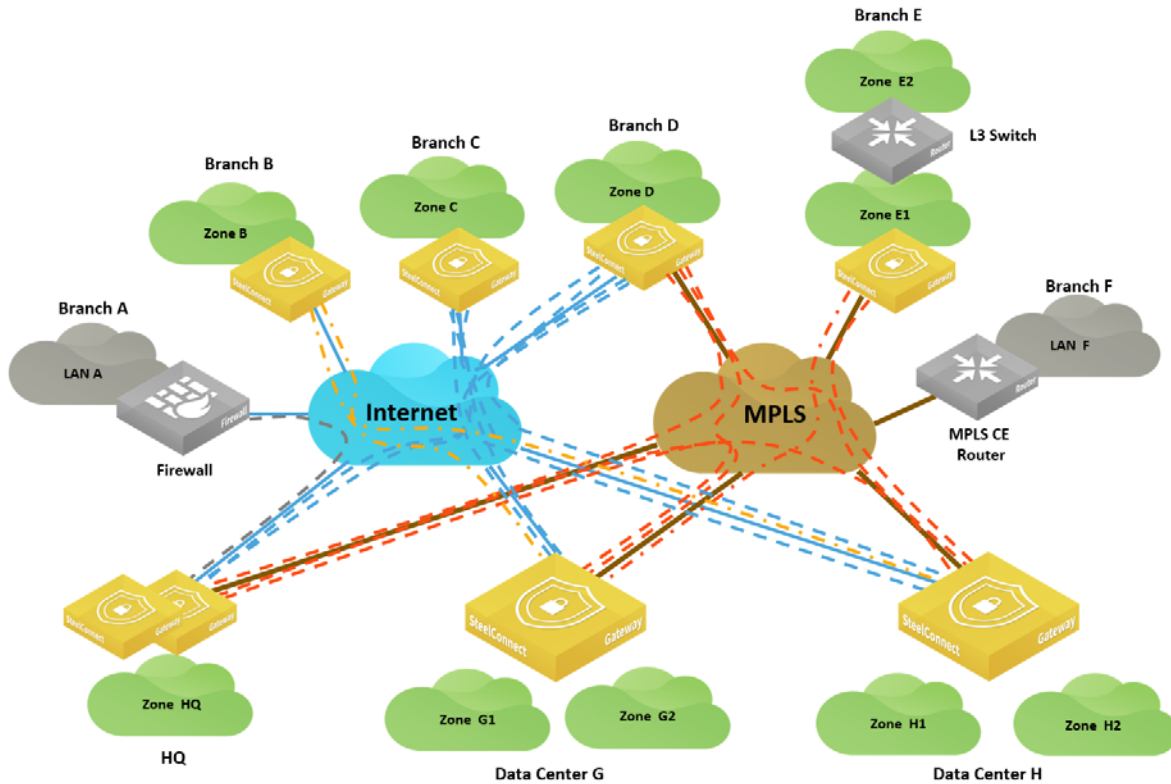
This topic presents a detailed, complete topology. It describes the following scenarios:

- “Full topology” on page 140
- “Connecting sites with the same WAN” on page 141
 - “SD-WAN sites communicating between each other” on page 141
 - “Brownfield scenario: SD-WAN site communicating with legacy internet site” on page 141
 - “Brownfield scenario: SD-WAN site communicating with legacy MPLS site” on page 147
- “Connecting sites with different WANs” on page 147
 - “SD-WAN sites: multihub” on page 148
 - “SD-WAN sites: full-mesh mode” on page 151
 - “Brownfield scenario: SD-WAN site communicating with legacy site” on page 155
 - “SD-WAN sites: path redundancy” on page 158
 - “Brownfield scenario: path redundancy” on page 160

Full topology

In this topic, we step through elements of this network, which is representative of many Riverbed SD-WAN customers. All the scenarios in this chapter are derived from this topology.

Figure 9-1. Full sample network



The network paths shown in Figure 9-1 use these conventions:

- Gray dotted lines represent a classic VPN in SCM.
- Blue dotted lines represent a RouteVPN tunnel.
- Orange dotted lines represent the AutoVPN in full-mesh connectivity.
- Yellow dotted lines represent RouteVPN connections in hub-and-spoke mode.

The topology has these components:

- Branch A is a legacy site connected to HQ (headquarters) through an IPsec tunnel from its firewall.
- Branch B is an SD-WAN site with only internet connectivity. It is configured as a leaf with two hubs, Data Center G and Data Center H.
- Branch C is an SD-WAN site with only internet connectivity. As opposed to Branch B, Branch C participates in a full-mesh overlay on the internet.
- Branch D is an SD-WAN site with MPLS and internet connectivity.
- Branch E is an SD-WAN site connected through MPLS only.
- Branch F is a legacy site with a third-party router connected to MPLS only.

- HQ is the headquarters connected through the internet and MPLS and is equipped with a pair of SDI-1030 appliances in a high-availability configuration.
- Data Center G and Data Center H are connected to MPLS and internet. The data centers are equipped with SDI-5030 appliances.

Note: The diagram is not suggesting a possible in-path deployment of SDI-5030 appliances. For clarity and simplification of the topology in the data centers, internet firewalls, aggregation routers, and BGP neighbor adjacencies are not depicted. Furthermore, SDI-5030 can be deployed as a single-unit cluster or n+k-units cluster. To get more details on SDI-5030 implementation, see [Chapter 6, “Data Center Topologies.”](#)

The following sections detail the different use cases and configurations to implement this network and explain the different brownfield scenarios, meaning how an SD-WAN site can communicate with a non-SD-WAN site. The brownfield scenarios are extremely important to consider during an SD-WAN project, especially during the transformation phase.

Connecting sites with the same WAN

The topics in this section show how sites using the same WAN can connect.

SD-WAN sites communicating between each other

By default, SD-WAN sites operate in a full-mesh mode on the internet. Tunnels are formed automatically and they learn the subnets of each other through the overlay network.

On a private network, such as MPLS, it is possible to use the underlay directly; however, we recommend you enable encryption when defining the WAN MPLS to benefit from all SD-WAN capabilities (including quality-based path selection). When encryption is on, SD-WAN appliances form the overlay tunnels automatically. For more details, see [“Integration with private networks” on page 65.](#)

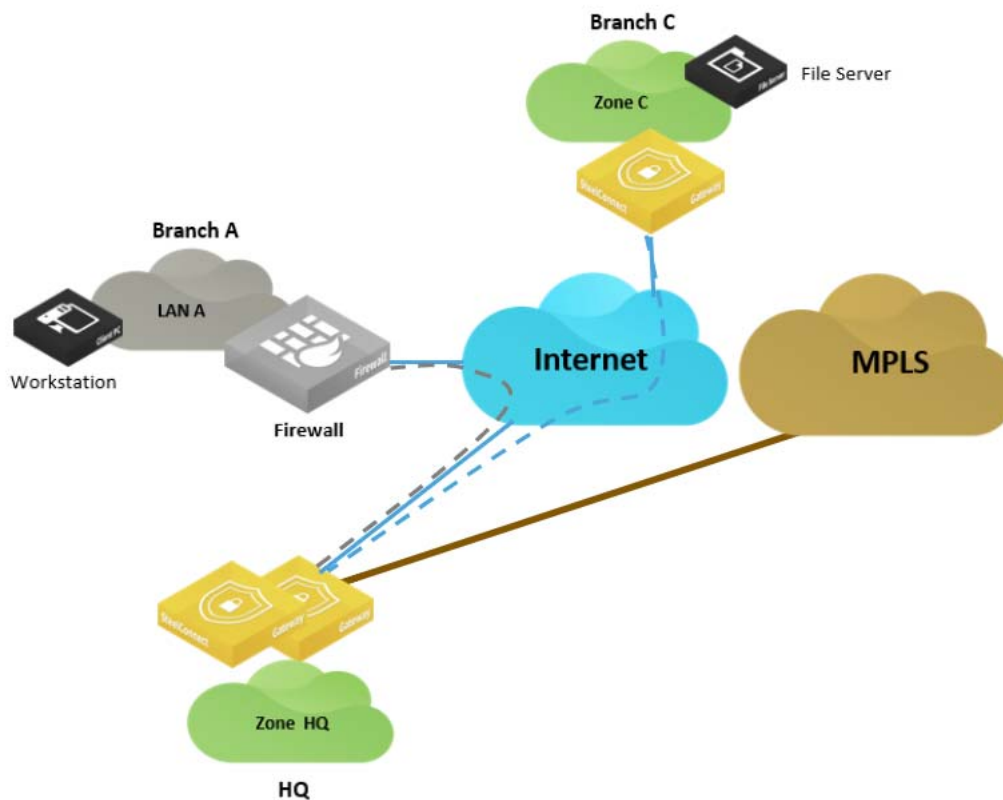
Brownfield scenario: SD-WAN site communicating with legacy internet site

In this scenario, a workstation in Branch A wants to connect to a server hosted in Branch C. Branch A is equipped with a third-party firewall. An IPsec tunnel is configured between the firewall and the SDI-1030 appliances at the headquarters (HQ) office.

You could create another IPsec tunnel between the firewall in Branch A and the gateway in Branch C. However, depending on the number of sites, it might be difficult to maintain and scale this deployment.

Let's consider how to transit traffic through the HQ classic VPN tunnel from our branch sites. This requires minor modifications on the third-party device and minimal configuration in SCM.

Figure 9-2. SD-WAN site with legacy branch



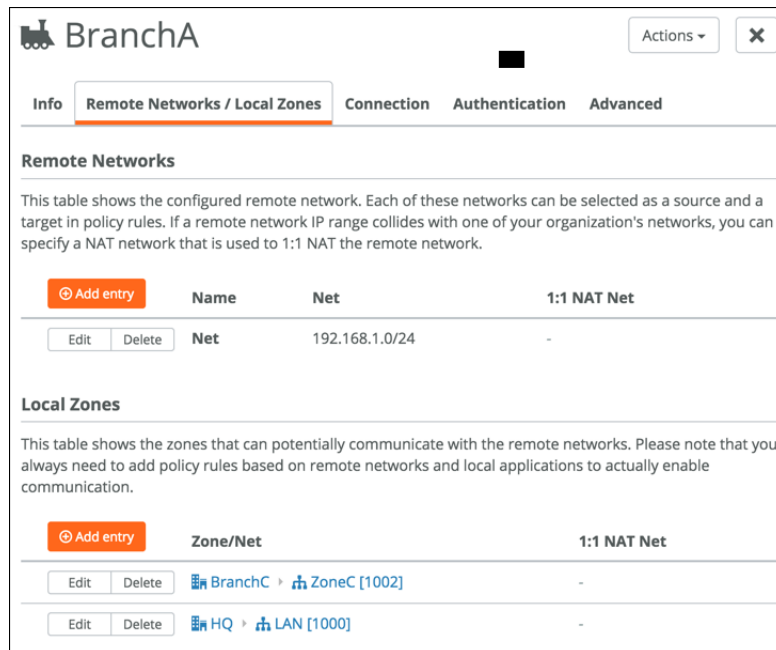
In this situation, the traffic between Zone C and Zone A travels through the RouteVPN. The HQ gateway routes traffic based on longest prefix match over the classic VPN to Zone A, a LAN segment in Branch A.

To configure this brownfield scenario

1. Create a classic VPN between the HQ gateway and the firewall at Branch A.
Choose Network Design > ClassicVPN and click **New ClassicVPN connection**.
2. Define traffic path rules for traffic from the branch to traverse the RouteVPN.
Choose Rules > Traffic Rules and click **New Traffic Rule**.
3. Add an outbound rule allowing traffic from Zone C to the classic VPN.
 - Choose Rules > Outbound / Internal and click **New policy rule**.
 - The remote network definition is the subnets behind the Branch A firewall.
 - The local zones define Branch C, Zone C, and HQ LAN 1000. This creates a proxy ACL defining what traffic needs to be encrypted between HQ and Branch A. This configuration must be a mirror image of the configuration on the firewall.
4. Choose Network Design > ClassicVPN and add a new classic VPN connection.

5. Modify the remote networks to reflect the local and remote networks to be encrypted.

Figure 9-3. Remote and local networks



Here is an example of the firewall configuration (in VyOS):

```
tunnel 0 {
    allow-nat-networks disable
    allow-public-networks disable
    esp-group office-srv-esp
    local {
        prefix 192.168.1.0/24
    }
    remote {
        prefix 172.16.0.0/24
    }
}

tunnel 1 {
    allow-nat-networks disable
    allow-public-networks disable
    esp-group office-srv-esp
    local {
        prefix 192.168.1.0/24
    }
    remote {
        prefix 172.16.4.0/24
    }
}
```

Connection settings define the remote gateway (Branch A firewall) and HQ as the tunnel endpoints.

6. Confirm the connection terminates at HQ.

If you select the branch site instead of the HQ site, the tunnel fails to establish.

Figure 9-4. Connection settings

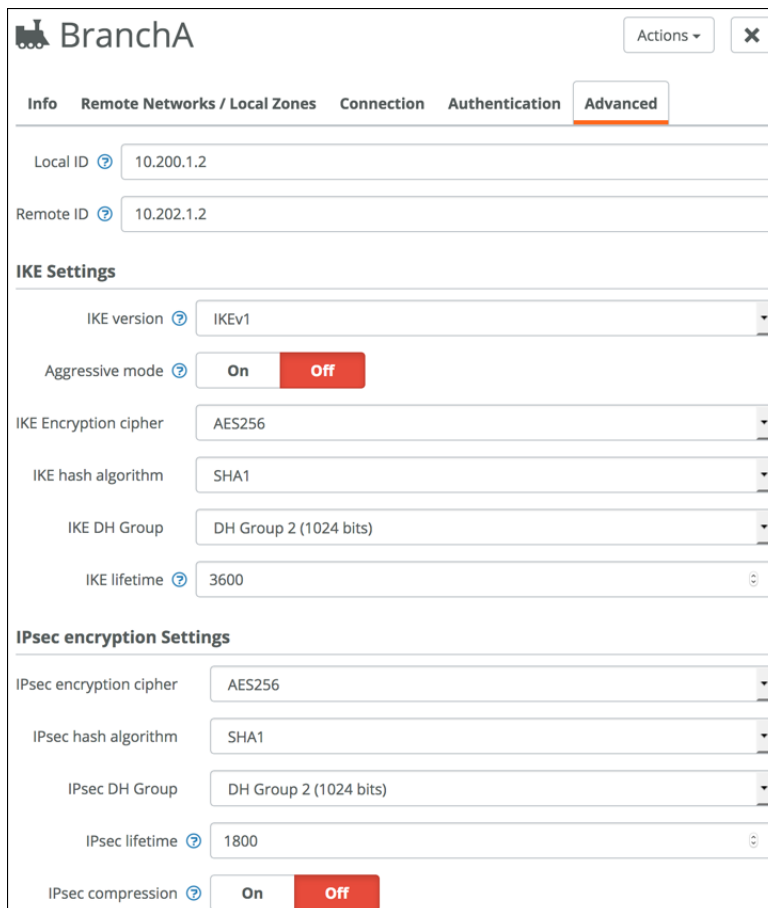


The screenshot shows the 'BranchA' configuration page with the 'Connection' tab selected. The 'Connection settings' section includes the following fields:

- Name:** BranchA
- Remote Gateway:** 10.202.1.2
- Site:** HQ (selected from a dropdown menu)
- Uplink:** Use AutoVPN uplink selection (selected from a dropdown menu)

7. In the Advanced tab, match the VPN parameters between the firewall and the Branch A configuration.

Figure 9-5. Advanced settings



The screenshot shows the 'BranchA' configuration page with the 'Advanced' tab selected. The 'Advanced' section includes the following settings:

- Local ID:** 10.200.1.2
- Remote ID:** 10.202.1.2
- IKE Settings:**
 - IKE version:** IKEv1
 - Aggressive mode:** On (selected from a dropdown menu)
 - IKE Encryption cipher:** AES256
 - IKE hash algorithm:** SHA1
 - IKE DH Group:** DH Group 2 (1024 bits)
 - IKE lifetime:** 3600
- IPsec encryption Settings:**
 - IPsec encryption cipher:** AES256
 - IPsec hash algorithm:** SHA1
 - IPsec DH Group:** DH Group 2 (1024 bits)
 - IPsec lifetime:** 1800
 - IPsec compression:** On (selected from a dropdown menu)

It is important that the IKE settings and the IPsec settings match. Here is an example configuration from the firewall (VyOS):

```

ipsec {
    esp-group rvbd-esp { <----- Matches the IPsec settings in the classic VPN configuration.
        lifetime 1800
        proposal 1 {
            encryption aes256
            hash sha1
        }
    }
    ike-group rvbd-ike { <----- Matches the IKE settings in the classic VPN configuration.
        lifetime 3600
        proposal 1 {
            encryption aes256
            hash sha1
        }
    }
    ipsec-interfaces {
        interface eth0
    }
}

```

Additionally, the local and remote ID is overridden in the classic VPN configuration so that FQDN is not used to establish the tunnel.

8. Choose Rules > Traffic Rules and create a traffic rule to send traffic between the branches over RouteVPN.

Figure 9-6. Creating a traffic rule

Edit Traffic Rule Actions X

Position: >> No position change <<

Name: Traffic_rule

Site scope: Apply rule in all sites

Users / Source: Selected ClassicVPN remote networks

- BranchA > Net [192.168.1.0/24]

Applications / Target: Selected zones

- BranchC > ZoneC [1002]

Path preference: 1 RouteVPN

Path Quality profile: None

Rule fall through: On Off

QoS priority: >> Automatic <<

WAN Optimization: Use organization defaults

9. To get Branch C to forward traffic to Branch A, a route must be learned on the Branch C gateway. To accomplish this, add a zone to HQ advertising a supernet of the subnet at Branch A.
10. Choose Network Design > Zones and add a zone to advertise the prefix of the branch sites.

Figure 9-7. Zone to advertise the prefix of the branch sites

The screenshot shows the 'Test' zone configuration page in the SteelConnect interface. The 'IP' tab is active, displaying the 'Discovered Networks' section. Under 'IPv4 network and gateway', the IPv4 Network is set to 192.168.0.0/16 and the IPv4 Gateway is set to 192.168.0.1. The 'IPv6 status, network, and gateway' section shows the 'Use IPv6' toggle set to 'Off'.

Test Actions X

IP Gateways DHCP VLAN WAN/AutoVPN Settings ADDL Networks

Discovered Networks

IPv4 network and gateway

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

IPv4 Network ? 192.168.0.0/16

IPv4 Gateway ? 192.168.0.1

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

Brownfield scenario: SD-WAN site communicating with legacy MPLS site

In this scenario, a workstation in Branch F requests files on the server hosted in Branch E.

Figure 9-8. Legacy MPLS site



The Riverbed SD-WAN solution can send and receive traffic with the underlay network. As long as the SD-WAN gateway is configured to advertise its local subnets on MPLS and can learn routes from the MPLS network, it can communicate through the underlay. For more information, see [“Integration with MPLS CE router” on page 69](#) and [“MPLS CE router replacement: ASBR-like deployment” on page 71](#).


Connecting sites with different WANs

The topics in this section show how sites using different WANs can connect.

SD-WAN sites: multihub

In this scenario, a workstation in Branch B wants to connect to a server hosted in Branch E. Both sites are in different WANs: Branch B connected to the internet and Branch E connected only to MPLS.

Figure 9-9. Connecting WANs of different types




This deployment requires a transit site that interconnects both WANs. In this example, it could be either the HQ (headquarters) or either of the data centers.

For this example, we'll configure Branch B as a leaf in a hub-and-spoke topology, Data Center G as the primary hub, and Data Center H as the secondary.

The following figure rearranges the topology for better visualization.

Figure 9-10. Visualizing the transit site deployment



To configure this scenario supporting different WANs

1. Create a dual-hub configuration and assign it to the site.

In SCM, choose Network Design > Split site/Dual-hub and click **New Split/Dual-hub Group**. For this example, select Dual-hub group as the type of site group and select DatacenterG and DatacenterH as members of the group.

Figure 9-11. Dual-hub configuration

Name	Group type	Members
No data to display		

Create Split/Dual-hub...

Name: DualHUB

Type of the sitegroup: Dual-hub group

Member preference	Member
1	DatacenterG
2	DatacenterH

Cancel Submit

2. Assign the dual-hub configuration to a site.

Choose Network Design > Sites, select Branch B and select the WAN/AutoVPN tab. Configure Branch B as a leaf, select dual-hub group as the master type, and select the group created in the previous step.

Figure 9-12. Configuring Branch B in AutoVPN leaf mode

The screenshot shows the configuration page for Branch B, specifically the WAN/AutoVPN tab. The page has a header with the site name 'BranchB' and an 'Actions' dropdown. Below the header is a navigation bar with tabs: Info, Zones, xLAN, DNS, WAN/AutoVPN (selected), Agents, RADIUS, Size, Cloudi-Fi, and Zscaler. The main content area is divided into sections: 'Location', 'Internet breakout preference', 'Outbound NAT', and 'AutoVPN Leaf Mode'. The 'Internet breakout preference' section has a dropdown menu set to 'Use organization default'. The 'Outbound NAT' section has a toggle switch for 'Skip outbound NAT' set to 'Off'. The 'AutoVPN Leaf Mode' section has a toggle switch set to 'On', a 'Master type' dropdown set to 'Dual-hub group', and a 'Master site group' dropdown set to 'DualHUB'. At the bottom right are 'Cancel' and 'Submit' buttons.

BranchB Actions

Info Zones xLAN DNS **WAN/AutoVPN** Agents RADIUS Size Cloudi-Fi Zscaler

Location

Internet breakout preference

This setting overrides the organization default for the internet breakout preference. When set, it is valid for all zones in this site. Note: You can also override this setting per zone.

⋮ Use organization default

Outbound NAT

When active, traffic exiting this site will not be NATed. Enable this setting if an upstream gateway is handling NAT or if the network uses public IP addresses.

Skip outbound NAT **On** **Off**

AutoVPN Leaf Mode

When active, this site will use another (non-leaf) master site or group to send and receive AutoVPN traffic. This can be used to build hub-and-spoke topologies, or to enable connectivity between sites that are both behind restrictive firewalls or third-party NAT devices. This setting will be disabled if the site is a hub.

AutoVPN Leaf Mode **On** **Off**

Master type Single site **Dual-hub group**

Master site group DualHUB


Cancel Submit

With this configuration, Branch B can access the resources in Branch E (and vice versa) but also access resources in HQ through Data Center G and alternatively through Data Center H. Traffic is steered on the overlay network.

SD-WAN sites: full-mesh mode

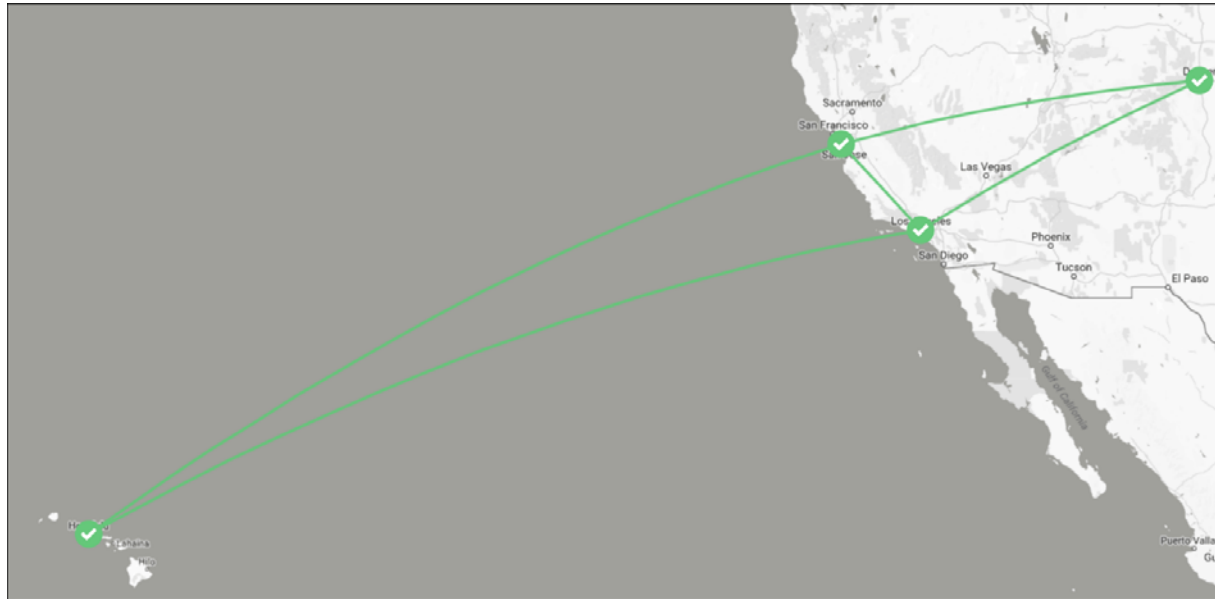
This scenario is similar to the previous one, except Branch C is in full-mesh mode over the internet. By default, the gateway in Branch C has no knowledge of Branch E subnets.

Figure 9-13. Different WANs with full mesh



For this scenario, assume that Branch C is already connected to both data centers over the internet. Branch E is also connected to the data centers over MPLS. (Branch C is in Honolulu, Branch E is in Denver, Data Center G is in Los Angeles, and Data Center H is in Sunnyvale.)

Figure 9-14. Branch C and Branch E connected to data centers



To configure this scenario

1. Configure Data Center G as an alternate hub for Branch C.

Choose Network Design > Sites and select Branch C. Under the WAN/AutoVPN tab, add DatacenterG as an alternate hub.

Figure 9-15. Adding an alternate hub

The screenshot displays the Palo Alto Networks SD-WAN configuration interface. On the left, the 'Sites' table lists five sites: Branch C, Branch E, Datacenter G, Datacenter H, and Headquarters. The 'BranchC' configuration page is open on the right, with the 'WAN/AutoVPN' tab selected. The 'Alternate Hub' field is set to 'DatacenterG'.

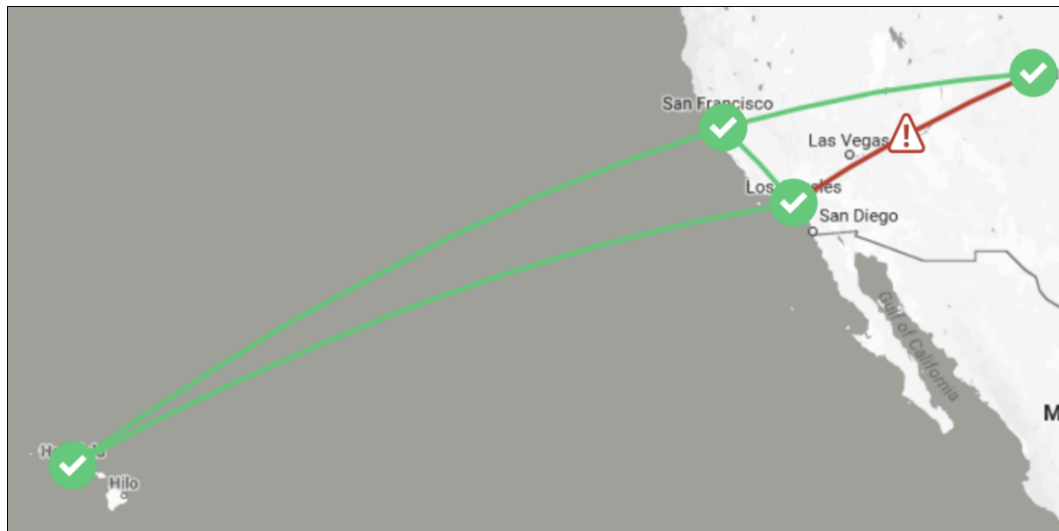
Name	Breakout	City / Country	Street Address	Tags
BranchC Branch C	Default	Honolulu United States	-	
BranchE Branch E	Default	Denver United States	-	
DatacenterG Datacenter G	Default	Los Angeles United States	-	
DatacenterH Datacenter H	Default	Sunnivale United States	-	
HQ Headquarters	Default	Portland United States	-	Agent home

BranchC Configuration (WAN/AutoVPN tab):

- Location:** Internet breakout preference (Use organization default)
- Outbound NAT:** Skip outbound NAT (On/Off)
- AutoVPN Leaf Mode:** AutoVPN Leaf Mode (On/Off)
- Alternate Hub:** DatacenterG
- AutoVPN advanced settings:** AutoVPN Port (4500)

At this stage of the configuration, the workstation in Branch C can access the server in Branch E (and vice versa). If the internet or MPLS connection in Data Center G goes down, the branch sites C and E are no longer able to connect.

Figure 9-16. Connection down between Branch C and Branch E



To provide more resiliency, configure Data Center H as the alternate hub for Data Center G so if MPLS goes down, Data Center G knows that it can send traffic destined to the MPLS network through Data Center H.

Figure 9-17. Configuring Data Center H as the alternate hub

Sites

Map
Table
Add Site(s)

Enter a filter string
Clear
Apply

5 Items in set (5 total, 0 filtered)

Name	Breakout	City / Country	Street Address	Tags
BranchC Branch C	Default	Honolulu United States	-	
BranchE Branch E	Default	Denver United States	-	
DatacenterG Datacenter G	Default	Los Angeles United States	-	
DatacenterH Datacenter H	Default	Sunnivale United States	-	
HQ Headquarters	Agent home	Default Portland United States	-	

DatacenterG

Info
Zones
xLAN
DNS
WAN/AutoVPN
Agents
RADIUS
Size

Location

Internet breakout preference

This setting overrides the organization default for the internet breakout preference, all zones in this site. Note: You can also override this setting per zone.

Use organization default

Outbound NAT

When active, traffic exiting this site will not be NATed. Enable this setting if an upstream NAT or if the network uses public IP addresses.

Skip outbound NAT On Off

AutoVPN Leaf Mode

When active, this site will use another (non-leaf) master site or group to send and receive traffic. This can be used to build hub-and-spoke topologies, or to enable connectivity between sites behind restrictive firewalls or third-party NAT devices. This setting will be disabled if the site is a hub.

AutoVPN Leaf Mode On Off

Alternate Hub • DatacenterH


Important: In version 2.10, you can configure only one alternate hub for a site. If Data Center G goes down, Branch C won't be able to communicate with Branch E anymore. This will be enhanced in a future release.

Brownfield scenario: SD-WAN site communicating with legacy site

In this scenario, a user in Branch C needs to access files on a server hosted in Branch F. Branch F is not equipped with Riverbed SD-WAN technology. It has a traditional CE router providing connectivity to the MPLS network.

In this case, it is not possible to leverage the power of an overlay network exclusively. You have to rely on the underlay and its routing protocols.

Figure 9-18. SD-WAN and legacy site



On SteelConnect version 2.10, SDI-5030 appliances cannot discover routes locally through BGP. They can only announce locally the subnets (zones) that are configured on SteelConnect Manager to attract the traffic destined to them.

Tip: An enhancement is scheduled for the next version (2.11) that will gracefully handle that scenario.

In version 2.10, you can use a workaround to attract the traffic in the data centers so they transit traffic between the two WANs. For that to happen, create a supernet in Data Center G that will be announced in the overlay for Branch C to send the traffic to Data Center G. The supernet will be advertised through BPG on the MPLS so Branch F sends the traffic the Data Center G.

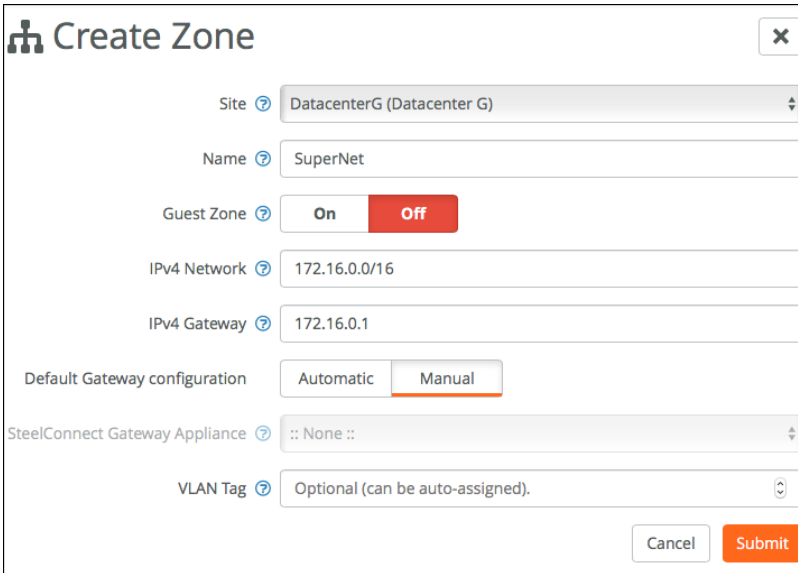
To ensure redundancy and make sure that if Data Center G is down, Branch C and Branch F are able to communicate between each other, configure both data centers as split data centers. Configure the supernet zone in both data centers.

To configure this topology

1. Create a supernet zone in Data Center G.

Choose Network Design > Zones and click **New Zone**.

Figure 9-19. Creating a zone in Data Center G



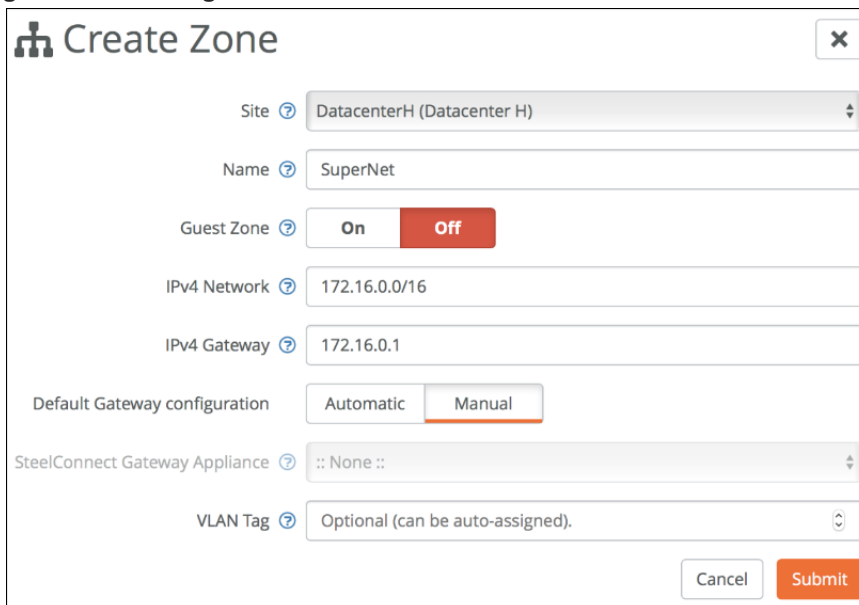
The screenshot shows the 'Create Zone' dialog box with the following configuration:

- Site: DatacenterG (Datacenter G)
- Name: SuperNet
- Guest Zone: Off (The 'Off' button is highlighted in red)
- IPv4 Network: 172.16.0.0/16
- IPv4 Gateway: 172.16.0.1
- Default Gateway configuration: Manual (The 'Manual' button is highlighted in orange)
- SteelConnect Gateway Appliance: :: None ::
- VLAN Tag: Optional (can be auto-assigned).

At the bottom right, there are 'Cancel' and 'Submit' buttons.

2. Create the same zone in Data Center H.

Figure 9-20. Creating a zone in Data Center H



The screenshot shows the 'Create Zone' dialog box with the following configuration:

- Site: DatacenterH (Datacenter H)
- Name: SuperNet
- Guest Zone: Off (The 'Off' button is highlighted in red)
- IPv4 Network: 172.16.0.0/16
- IPv4 Gateway: 172.16.0.1
- Default Gateway configuration: Manual (The 'Manual' button is highlighted in orange)
- SteelConnect Gateway Appliance: :: None ::
- VLAN Tag: Optional (can be auto-assigned).

At the bottom right, there are 'Cancel' and 'Submit' buttons.

3. Choose Network Design > Split site/Dual Hub and create a new Split Site group with both data centers.

Figure 9-21. Creating a split site group for both data centers

Create Split/Dual-hub...

Name ?

Datacenter

Type of the sitegroup ?

Split Site group

Member preference

⋮	1	▲	▼	🏠 DatacenterG	🗑
	2	▲	▼	🏠 DatacenterH	🗑


Cancel

Submit

SD-WAN sites: path redundancy

In this scenario, the MPLS connection in Branch D goes down and you want to guarantee business continuity for a workstation in Branch D that needs to retrieve files from a server hosted in Branch E.

Figure 9-22. Path redundancy for SD-WAN sites



In this case, configure Branch D to use data center G as an alternate hub.

Choose Network Design > Sites, select Branch D and in the WAN/AutoVPN tab specify DatacenterG as the alternate hub.

Figure 9-23. Configuring as an alternate hub

Sites

Clear Apply

7 items in set (7 total, 0 filtered)

Name	Breakout	City / Country	Street Address	Tags
BranchB Branch B	Default	Mexico Mexico	-	
BranchC Branch C	Default	Honolulu United States	-	
BranchD Branch D	Default	San Antonio United States	-	
BranchE Branch E	Default	Denver United States	-	
DatacenterG Datacenter G	Default	Los Angeles United States	-	
DatacenterH Datacenter H	Default	Sunnivale United States	-	
HQ Headquarters	Agent home Default	Portland United States	-	

BranchD

Info Zones xLAN DNS **WAN/AutoVPN** Agents RADIUS Size

Location

Internet breakout preference
This setting overrides the organization default for the internet breakout preference. all zones in this site. Note: You can also override this setting per zone.

Outbound NAT
When active, traffic exiting this site will not be NATed. Enable this setting if an upstream NAT or if the network uses public IP addresses.
Skip outbound NAT

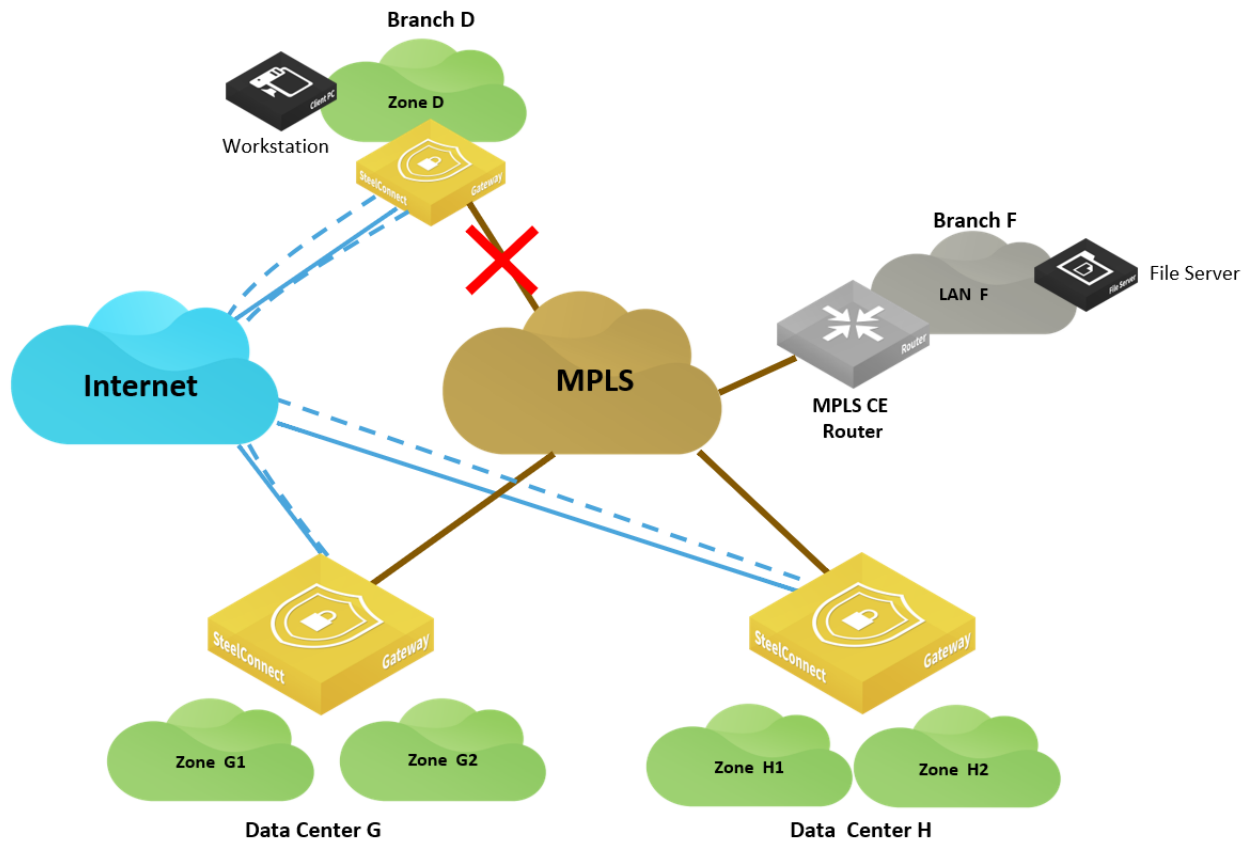
AutoVPN Leaf Mode
When active, this site will use another (non-leaf) master site or group to send and receive traffic. This can be used to build hub-and-spoke topologies, or to enable connectivity between sites behind restrictive firewalls or third-party NAT devices. This setting will be disabled if the site is a hub.
AutoVPN Leaf Mode

Alternate Hub DatacenterG

Brownfield scenario: path redundancy

In this scenario, the MPLS connection in Branch D goes down and you want to guarantee business continuity for a workstation in Branch D that needs to retrieve files from a server hosted in Branch F.

Figure 9-24. Path redundancy for mixed sites



This is the same situation as the “[Brownfield scenario: SD-WAN site communicating with legacy MPLS site](#)” on [page 147](#) and you can apply the same configuration.