



Silver Peak

SD-WAN Deployment Guide

Updated 8/17/2017

Copyright and Trademarks

Silver Peak SD-WAN Deployment Guide

Date: January 2017

Copyright © 2017 Silver Peak Systems, Inc. All rights reserved. Information in this document is subject to change at any time. Use of this documentation is restricted as specified in the End User License Agreement. No part of this documentation can be reproduced, except as noted in the End User License Agreement, in whole or in part, without the written consent of Silver Peak Systems, Inc.

Trademark Notification

The following are trademarks of Silver Peak Systems, Inc.: Silver Peak Systems™, the Silver Peak logo, Network Memory™, Silver Peak NX-Series™, Silver Peak VX-Series™, Silver Peak VRX-Series™, Silver Peak Silver Peak Unity EdgeConnect™, and Silver Peak Orchestrator™. All trademark rights reserved. All other brand or product names are trademarks or registered trademarks of their respective companies or organizations.

Warranties and Disclaimers

THIS DOCUMENTATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. SILVER PEAK SYSTEMS, INC. ASSUMES NO RESPONSIBILITY FOR ERRORS OR OMISSIONS IN THIS DOCUMENTATION OR OTHER DOCUMENTS WHICH ARE REFERENCED BY OR LINKED TO THIS DOCUMENTATION. REFERENCES TO CORPORATIONS, THEIR SERVICES AND PRODUCTS, ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED. IN NO EVENT SHALL SILVER PEAK SYSTEMS, INC. BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER, INCLUDING, WITHOUT LIMITATION, THOSE RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS DOCUMENTATION. THIS DOCUMENTATION MAY INCLUDE TECHNICAL OR OTHER INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE DOCUMENTATION. SILVER PEAK SYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENTATION AT ANY TIME.

Silver Peak Systems, Inc.
2860 De La Cruz Boulevard, Suite 100
Santa Clara, CA 95050

1.877.210.7325 (toll-free in USA)
+1.408.935.1850

<http://www.silver-peak.com/support>

Support

For product and technical support, contact Silver Peak Systems at either of the following:

1.877.210.7325 (toll-free in USA)

+1.408.935.1850

www.silver-peak.com/support

We're dedicated to continually improving the usability of our products and documentation.

- If you have suggestions or feedback for our documentation, please send an e-mail to techpubs@silver-peak.com.
- If you have comments or feedback about the interface, please send an e-mail to usability@silver-peak.com.

Contents

Copyright and Trademarks	2
Support	3
Prerequisites	7
The Parts of SD-WAN Deployment	9
Supported Technologies	10
Unsupported Technologies	11
Deployment Parameters	12
Deployment Checklist	13
Planning	13
Orchestrator	14
Appliances	14
Deploy Orchestrator	16
Installing Orchestrator	17
Configuring Orchestrator	18
Licensing	20
License & Account Key	21
NTP Setup	23
Using a Proxy Server	24
Licensing Physical Appliances	25
Licensing Virtual Appliances	26
Template Groups	28
Interface Labels	29
Deployment Profiles	32
Configuring Deployment Profiles	33
WAN Hardening	36
Deployment Modes	37
In-Line Deployments	38
In-Line Bridge vs In-Line Router	40
In-Line Overview	41

Out-of-Path Deployments	42
Fail Safe Behavior	42
Policy-Based Routing (PBR)	43
PBR Overview	44
PBR Example	45
PBR Configuration Checklist	47
Web Cache Communication Protocol (WCCP)	49
WCCP Overview	50
WCCP Best Practice	51
WCCP Example	54
WCCP Configuration Checklist	56
Virtual Router Redundancy Protocol (VRRP)	59
VRRP Peering to a WAN Router	59
VRRP with PBR Configuration Checklist	63
VRRP Redundant Appliances with PBR	64
VRRP Peering to WAN Configuration Checklist	68
Host-Based Redirection	70
Server Mode Deployments	72
Deploying in NAT Environments	74
Example Deployments	76
Router Mode Considerations	77
In-Line Router Mode (Router + Firewall)	78
In-Line Router Mode (Router + Direct Internet)	79
In-Line Router Mode (Single Direct Internet)	80
In-Line Router Mode (Dual Direct Internet)	81
In-Line Router Mode (Dual MPLS)	82
Bridge Mode (Router + Direct Internet)	83
Bridge Mode (Router + Firewall)	84
Bridge Mode (Dual MPLS)	85
Router Mode MPLS + Internet	86
Router Mode HA (MPLS + MPLS)	87
Dual Home Router Mode (MPLS + Internet)	88
Dual Home Router Mode HA (MPLS + Internet)	90
Business Intent Overlays (BIO)	91
Overlays vs Underlays	93
Building an Overlay	94
Tunnels in an Overlay	96
Tunnel Reporting & Visibility	97

Stateful Firewall with Internet Breakout	98
Deployments	99
Silver Peak Stateful Firewall Use Case	99
Deploying Stateful Firewall	100
Add Source NAT for Internet Breakout Tunnels	100
ACL - Breakout Sanctioned SaaS Apps	101
Route Policies	102
Unity EdgeConnect SD-WAN Data Sheet-v1.3	104
Unity EdgeConnect XL Specifications-v1.3	106
Unity EdgeConnect L Specifications-v1.3	108
Unity EdgeConnect M Specifications-v1.3	110
Unity EdgeConnect S Specifications-v1.3	112
Unity EdgeConnect XS Specifications-v1.3	114
Add Appliances	116
Appliance Provisioning (ZTP)	117
Enable Subnet Sharing	117
Quality of Service (QoS)	119
Dynamic Rate Control (DRC)	121
Troubleshooting & Testing	124
Verify Appliance Connectivity	125
Tunnels Not Showing on Tunnels Page	127
Verify Traffic	128
Videos	130

Prerequisites

This SD-WAN Deployment Guide describes setting up a new generic environment, rather than migrating from an older release. You might experience different results in your environment.

Silver Peak SD-WAN offers automated setup—simply plug in your devices and Orchestrator automatically builds the network. For extra performance, you can add Boost.

The primary components of a Silver Peak Unity SD-WAN are:

- **Unity Orchestrator** - Runs as a virtual machine on VMware ESX/ESXi, Microsoft Hyper-V, Citrix XenServer, and the open source KVM hypervisor.
- **EdgeConnect appliances**
 - The EdgeConnect family scales to support everything from small branches to large data centers.
 - Available as hardware or software appliances (can be mixed and matched per requirements).
 - Hardware appliances are available as 1U rack-mountable appliances.
 - See the [Hardware Reference Guide](#) for more details on the hardware.

Prerequisites:

- Install Orchestrator. We recommend that all new EdgeConnect and Orchestrator deployments be running VXOA version 8.0 or later. You will receive a download link from Silver Peak by email.
- Complete the [Deployment Checklist](#) and gather your Deployment Parameters.
- Register Orchestrator to the Silver Peak Cloud Portal. Orchestrator (IP address) must be able to reach the Silver Peak Cloud Portal via the Internet. Allocate an appropriate IP address for the Orchestrator appliance and allow it access through any security components in the environment to the **cloudportal.silver-peak.com** domain. The Orchestrator requires port 443 access.

- If using VX or NX, obtain a pool of licenses for your appliances. You will receive these by email. If using EdgeConnect, you will use the same account key as Orchestrator.
-

The Parts of SD-WAN Deployment

Your Silver Peak SD-WAN is made of three elements:

- **The Cloud Portal** - manages all licenses without user action
- **Orchestrator** - must connect to Cloud Portal; management and configuration software for Silver Peak EdgeConnect devices and NX/VX appliances
- **EdgeConnect** - must connect to Cloud Portal; creates network connections and moves data securely.

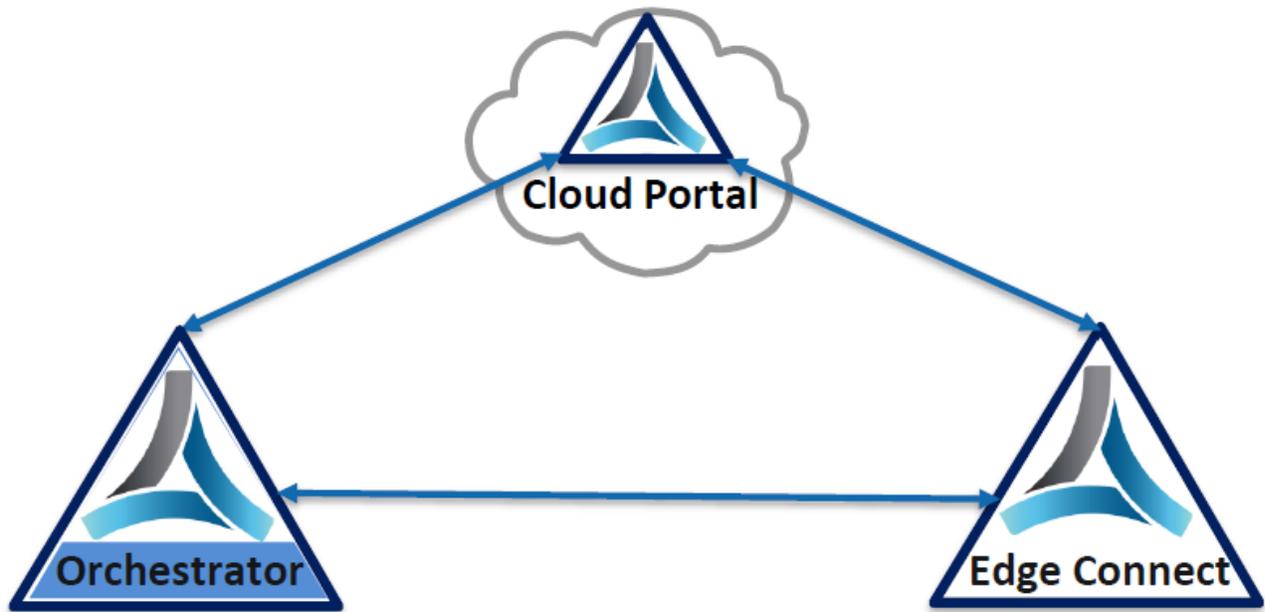


Figure 1. The Silver Peak Network

Your network will incorporate all three elements, so be sure to take this into consideration when designing your network.

Use the [Deployment Checklist](#) to help you.

Supported Technologies

Silver Peak currently supports the following:

- Any interface combination with In-Line Router Mode.
- Standard bridge mode using only WAN0/LAN0, WAN1/LAN1, TWAN0/TLAN0, TWAN1/TLAN1.
- High Availability (HA) when using technologies such as WCCP, VRRP, and PBR to forward packets (not with cross-connect enabled).

Unsupported Technologies

These technologies are **NOT** currently supported:

- Mixing 1Gbps and 10Gbps interfaces when using bridge or router mode.
- 10Gbps Bonding.

Deployment Parameters

Print this page and fill in your information. You will need to refer to this page during deployment.

SD-WAN Deployment Parameters

Site Name	Example	Your Site
Silver Peak Hostname	hostname	
Deployment Mode	4-port, Bridge	
Management IP (mgmt0)	MGMT0_IP	
mgmt0 Default GW	MGMT0_DEFAULT_GW	
Appliance IP (wan0)	WAN0_IP	
wan0 Next Hop	WAN0_IP_DEFAULT_GW	
lan0 Next Hop	LAN0_IP_DEFAULT_GW	
Appliance IP (wan1)	WAN1_IP	
wan1 Next Hop	WAN1_IP_DEFAULT_GW	
lan1 Next Hop	LAN1_IP_DEFAULT_GW	

Deployment Checklist

You can print out this page and use it as a reference.

Planning

- Network topology. Have detailed diagrams of the relevant networks (such as MPLS, Internet, or LTE). The installation will proceed more efficiently if you have WAN link, router/switch, and firewall configurations clearly documented ahead of time.
- Identify if Network Address Translation (NAT) or Port Address Translation (PAT) is in use on the network. This affects addressing schemes and traffic flow across the network.
- Consider traffic flow across the WAN. For example, should VOIP traffic always traverse MPLS, or should file sharing be load balanced across available paths?
- Consider how EdgeConnect appliances will be deployed. Which sites will have in-line appliances? Which appliances are out-of-path?
- Identify any hub sites and their role. (such as Replication hub, email hub, or VOIP hub).
- Write down specific attributes, such as inbound BW, outbound BW, all ports used, or Boost required.
 - See [In-Line Deployments](#) for related requirements.
 - See [Out-of-Path Deployments](#) for related requirements.
- Obtain the appropriate licensing. See [Licensing](#) on how to use your license. Consult your Sales Engineer for choosing the best options for your network.

Orchestrator

- Install Orchestrator, log in, enter the license information, and confirm it registers with the Cloud Portal. See [Deploy Orchestrator](#).
- Immediately change your user name and password, and keep them in a safe place. Failure to do so could subject you to hacking.
- Troubleshoot any alarms before proceeding.
- (Optional) Create Template Groups. See [Template Groups](#)
- Create interface labels. Labels enable you to easily identify each interface. Orchestrator treats interfaces with the same label the same way. See [Interface Labels](#).
- Create an Access List. Access lists are used to match traffic destined for an Overlay. See [Access Control Lists \(ACL\)](#).
- Configure Deployment Profiles using the Template Groups and interface Labels you have created. As Orchestrator discovers your appliances, you can manage how each one is deployed. See [Deployment Profiles](#).
- Create an Overlay (BIO) to manage how particular traffic should be handled in an overlay network, and which traffic will be affected by that overlay. See [Business Intent Overlays \(BIO\)](#).
- Create Route Policies, if needed, within Configuration Templates. If a Route Policy is intended to route traffic into a BIO, the BIO must exist first.

Appliances

- Install appliances and check to see each one can reach the Cloud Portal.

- Apply Deployment Profiles, Business Intent Overlays, and Configuration Templates.
- Define exceptions and perform any other individual configurations.

Deploy Orchestrator

You must deploy Orchestrator when setting up your network. Orchestrator manages all devices on your network and can push network-wide policies on a manual or scheduled basis.

Orchestrator manages your physical, virtual and cloud-based EdgeConnect appliances seamlessly from a single console.

Orchestrator is only offered as a virtual appliance and, therefore, requires a suitable host to run on. You must identify an appropriate host machine with adequate resources to host Orchestrator. Typical deployment locations for Orchestrator would be in a Network Operations Center (NOC) or Data Center, though any location with efficient access to the WAN devices could be suitable. For more information on Orchestrator requirements, refer to [Orchestrator Host System Requirements](#).

Orchestrator automatically detects network devices, but you must manually approve each discovered device.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
 - Within Orchestrator, click the **Question Mark**  on the page for details about each field.
-

Installing Orchestrator

To install Orchestrator for the first time:

1. From the email you received from SD-WAN, click the link to install Orchestrator. The Silver Peak website appears.
2. Log in and click **Download Software**. Save the software to your local machine.
3. After downloading, run the file to install the software.

Silver Peak supports all major hypervisors, including VMware, Microsoft Hyper-V, Citrix XenServer, and KVM.

Configuring Orchestrator

The first time you log into Orchestrator, the **Getting Started Wizard** automatically appears.

- After initial configuration, you can change your initial settings using the Orchestrator interface.
- To change the Admin password, use the **Getting Started Wizard**. To restart the wizard, go to the **Orchestrator Administration** tab, Under **General** choose **Getting Started Wizard**.

To run the Getting Started Wizard

When you have acquired your licenses, you are ready to configure Orchestrator.

1. Within a browser, open Orchestrator. Use the IP address that you set up for this purpose. The **Sign in** page appears.
2. Enter the username and password, then click **Login**.

The default username and password is *admin*.

The **Getting Started Wizard** appears.

3. The first page of the wizard shows the Host name, DHCP, and Password.

- Choose **Static** to enter the IP address manually (recommended best practice).
 - **CHANGE THE PASSWORD to a non-guessable password and put it somewhere secure. Leaving the default password could cause your network to be vulnerable to hackers.**
4. Enter the License number, Account Name, and Account Key given to you by Silver Peak. See [License & Account Key](#).
 5. Click **Next**.

Orchestrator sends the License, Account Name, and Account Key to the Cloud Portal and sends back a Registered message when successful.

To verify, go to **Orchestrator Administration > Silver Peak Cloud Portal** to view the registration status. It should read Registered = Yes.

6. Continue through the wizard until you are **Done**.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

Licensing

Unity Edge Connect licensing is offered with various levels of performance to match your network and avoid costly overhead. All licenses are managed by the Cloud Portal.

Silver Peak licenses are consumed as a subscription, sold in 1, 3, 5 and 7 year increments:

- **Unity EdgeConnect** - use any physical or virtual device; Orchestrator is a required virtual appliance included in your subscription at no cost.
 - **BASE** - Floating license for Zero Touch Provisioning (see [Appliance Provisioning \(ZTP\)](#)), Dynamic Path Control (DPC), and Path Conditioning including FEC and POC.
 - **BOOST** - License for WAN Optimization (per 100Mbps); includes Silver Peak Network Acceleration and Network Memory.
 - **PLUS** - License for over 200Mbps throughput.
 - **SaaS** - Fabric-wide license for SaaS and IaaS optimization.
-

License & Account Key

When first deploying Orchestrator, the **Getting Started Wizard** appears. The license information is entered on page 2 of the wizard.

Getting Started Wizard

1 Hostname, DHCP, Password | **2 License and Registration** | 3 Date/Time | 4 Email | 5 Add Appliances | 6 Backup

EdgeConnect Registration *(Also required for CPX and SaaS)*

Account Name

Account Key

Contact

Registered No

Orchestrator License for NX/VX Appliances

License

This license allows you to manage up to 10 appliances. Visit the Silver Peak Support Portal for upgrade options.

< Previous | Next > | Apply

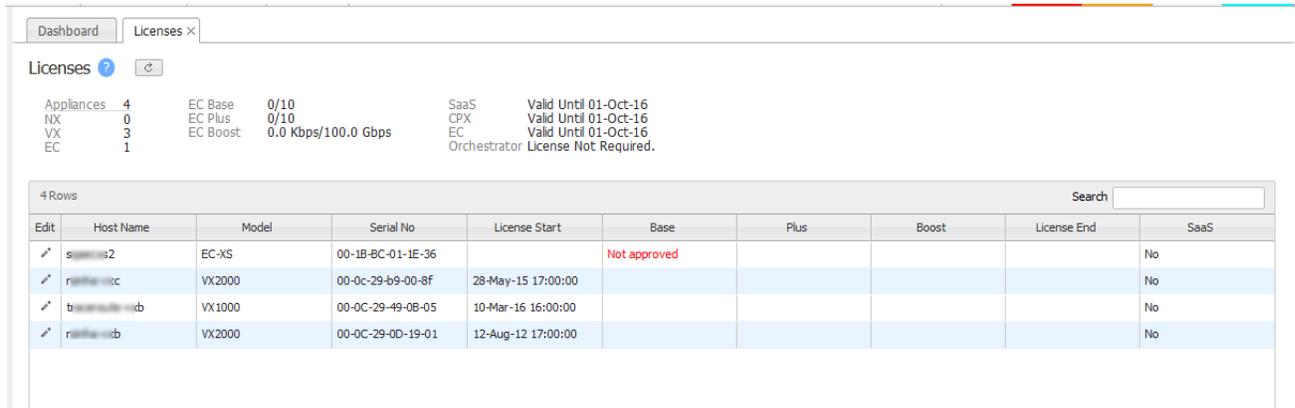
Figure 2. License page of the Getting Started Wizard

If the wizard does not start up automatically, you can access it again within Orchestrator. Go to **Orchestrator Administration > Getting Started Wizard**.

- Enter your license and account key into the appropriate fields, then continue with the wizard. To obtain these, log into the Support Portal.
- The License (usually 60 characters) allows you to use Orchestrator.
- The Account Name and Account Key (usually 32 characters) are required to use EdgeConnect. You can also enter this information within Orchestrator. Go to **Orchestrator Administration > Silver Peak Cloud Portal**.
- When the EdgeConnect reaches the Internet and is approved, it appears in the Orchestrator tree view within the chosen folder.

To view existing Licenses

To view existing licenses, within Orchestrator, go to **Configuration > Licenses**.



The screenshot shows the 'Licenses' page in the Orchestrator interface. At the top, there are tabs for 'Dashboard' and 'Licenses'. Below the tabs, there is a summary section with the following data:

Appliances	4	EC Base	0/10	SaaS	Valid Until 01-Oct-16
NX	0	EC Plus	0/10	CPX	Valid Until 01-Oct-16
VX	3	EC Boost	0.0 Kbps/100.0 Gbps	EC	Valid Until 01-Oct-16
EC	1			Orchestrator	License Not Required.

Below the summary is a table with 4 rows and 10 columns. The columns are: Edit, Host Name, Model, Serial No, License Start, Base, Plus, Boost, License End, and SaaS. The first row shows a license for 's...' with model 'EC-XS' and serial '00-1B-BC-01-1E-36', which is 'Not approved'. The other three rows show licenses for 'r...' with models 'VX2000' and 'VX1000'.

Edit	Host Name	Model	Serial No	License Start	Base	Plus	Boost	License End	SaaS
	s...	EC-XS	00-1B-BC-01-1E-36		Not approved				No
	r...	VX2000	00-0c-29-b9-00-8f	28-May-15 17:00:00					No
	t...	VX1000	00-0c-29-49-06-05	10-Mar-16 16:00:00					No
	r...	VX2000	00-0c-29-0d-19-01	12-Aug-12 17:00:00					No

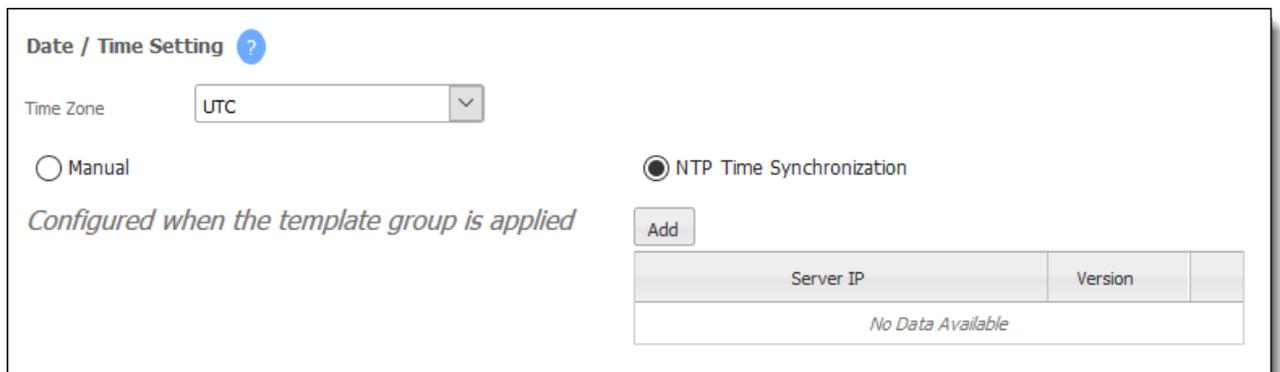
Figure 3. The Licenses page in Orchestrator

NTP Setup

You can manually configure the date and time of an appliance, or use a Network Time Protocol (NTP) server to automatically update date and time.

To configure the date and time

1. From the Templates page, choose **Date/Time**.



Date / Time Setting ?

Time Zone: UTC

Manual NTP Time Synchronization

Configured when the template group is applied

Add

Server IP	Version
No Data Available	

2. Choose a time zone from the drop down list.
3. Select **NTP Time Synchronization**, then enter the server IP address.

Manual matches the appliance time to the client system time of the template.

NTP enables the Appliance Manager to choose servers in the listed order, from top down.

4. Click **Add**.

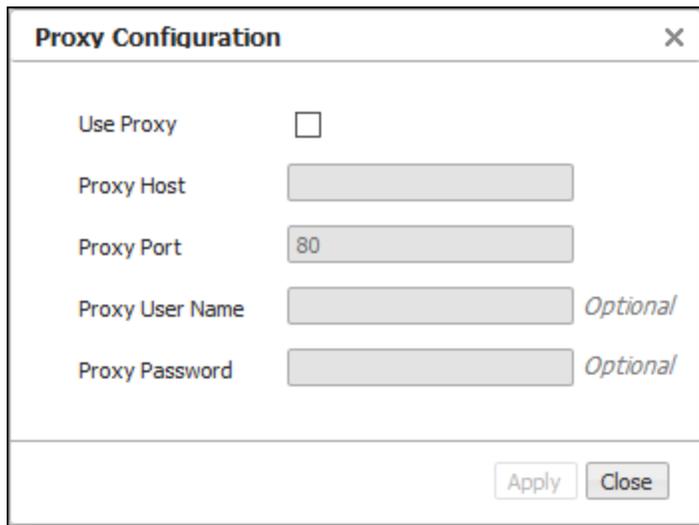
Using a Proxy Server

You can configure a proxy (for example, to overcome firewall issues) to reach the Silver Peak Cloud Portal.

To set up a proxy server

1. Within Orchestrator, go to **Orchestrator Administration > Proxy Configuration**.

The Proxy Configuration form appears.



The image shows a 'Proxy Configuration' dialog box. It has a title bar with the text 'Proxy Configuration' and a close button (X). The main area contains the following fields:

- Use Proxy**: A checkbox that is currently unchecked.
- Proxy Host**: A text input field that is empty.
- Proxy Port**: A text input field containing the value '80'.
- Proxy User Name**: A text input field that is empty, with the word 'Optional' to its right.
- Proxy Password**: A text input field that is empty, with the word 'Optional' to its right.

At the bottom right of the dialog box, there are two buttons: 'Apply' and 'Close'.

2. Check **Use Proxy** and enter the Proxy Host IP address.
Enter the proxy **User Name** and **Password**, if needed.
 3. Click **Apply**.
-

Licensing Physical Appliances

Silver Peak physical appliances have the licenses already installed. Simply add them to the network and let Orchestrator discover them.



NOTE: If the appliance can obtain an IP address via DHCP, use DNS to resolve the domain name of the Cloud Portal, and reach the Cloud Portal via port 443. No additional setup is required.

Licensing Virtual Appliances

Virtual appliances run under a hypervisor installed on a customer server. They don't come with a serial number but require your account information to be entered manually when first booted. After adding your account information, the appliances can then register with the Cloud Portal and obtain a license.

EdgeConnect virtual appliances (EC-V) require a license for speeds above 200Mbps. Your hardware must support the performance. For more information, see the [EdgeConnect virtual appliance requirements](#).

To manage the license of a virtual appliance

1. From the tree view, right-click the appliance, then choose **Appliance Manager** from the context menu.

A new window for the appliances opens in your browser.

2. From the Administration tab, select **License & Registration**. The License and Administration page appears.
3. Enter or modify the license information, if needed. You can also enter an **Account Name** and **Account Key** for EdgeConnect and SaaS Optimization appliances.
4. Click **Apply**.

The updated license information shows on the page.

License & Registration ? [Monitoring](#)

License Key

HvzB-
To retrieve or upgrade a license, go to the [Silver Peak Support Portal](#)

Current License

Start Date	Wed Aug 24th 2016
End Date	None
Portal Registration	Not registered
Model	EC-V
Serial Number	000C2997A7C8

Registration (Only required for cloud-based features and products, including EdgeConnect and SaaS Optimization)

Account Name	<input type="text"/>
Account Key	<input type="text"/>
Appliance Tag	<input type="text" value="Optional"/>
Contact	Available after registration



NOTE: If the appliance can obtain an IP address via DHCP, use DNS to resolve the domain name of the Cloud Portal, and reach the Cloud Portal via port 443. No additional setup is required.

Template Groups

Templates are configuration values that can be applied to one or more appliances. Template Groups are a collection of templates that can be applied simultaneously to designated appliances. Using Template Groups ensures consistency and reduces potential configuration errors throughout your network.



CAUTION:

- Best practice is to edit configuration settings only within a Template or Template Group.
- Some templates will REPLACE all settings on the appliance with the template settings unless the MERGE option is selected. MERGE keeps previously set values while replacing modified default values. Except for special circumstances, best practice is to use REPLACE.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

Interface Labels

To make it easier for you to identify your connections, you can create descriptive interface labels for each link type in your environment. Labels are used for configuration tasks.

There are two types of labels:

- **LAN** - for application traffic, such as Voice or WiFi
- **WAN** - for service traffic, type of network or connection, such as MPLS, Internet, or LTE.

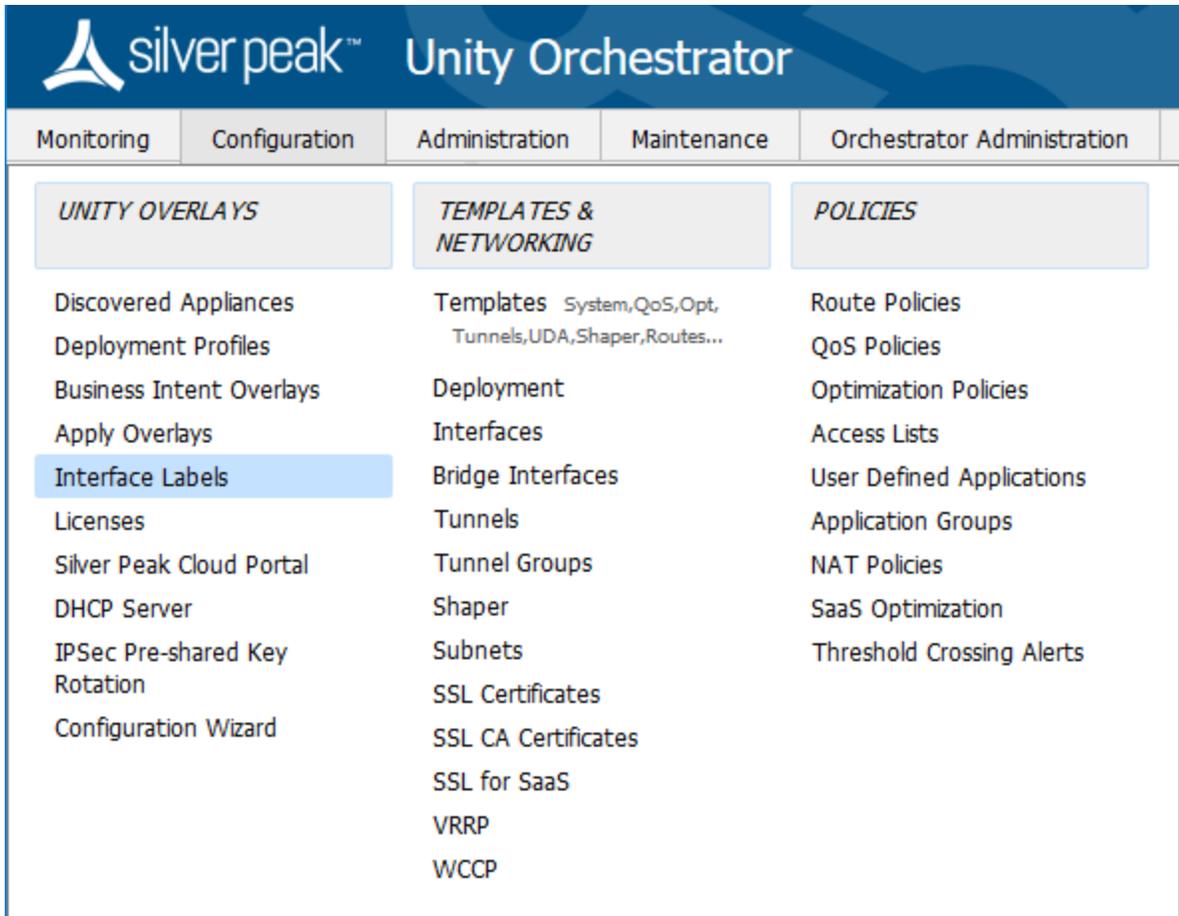
For example, you would choose labels for a traffic access policy in a Business Intent Overlay (BIO), which in turn is applied to an appliance with those **LAN** labels. All traffic matching those interfaces are automatically processed by that BIO. If you use an ACL for a traffic access policy, then the **LAN** label is ignored for that BIO.

The **WAN** labels are used by Orchestrator and BIOs to determine which interfaces on different appliances should be connected by tunnels built by Orchestrator.

Orchestrator automatically pushes the labels to the appliances it manages so you don't have to.

To Create a Label

1. From the Configuration tab, choose **Interface Labels**.



The Interface Labels form appears.

2. Choose either WAN or LAN from the **Type** drop down list, then enter a descriptive name in the **Label** field..

The screenshot shows a window titled "Interface Labels" with a close button (X) in the top right corner. Below the title bar, there are two input fields: "Type" with a dropdown menu showing "wan" and "Label" with an empty text box. To the right of the "Label" field is an "Add" button. Below these fields is a table with two columns: "Type" and "Label". The table has five rows of data, each with a close button (X) in the third column. The rows are: lan Data, lan Voice, wan LTE, wan Internet, and wan MPLS. The "lan" and "wan" columns have a small upward-pointing triangle next to them. At the bottom right of the window are "Save" and "Close" buttons.

Type ▲	Label	
lan	Data	×
lan	Voice	×
wan	LTE	×
wan	Internet	×
wan	MPLS	×

3. Click **Add**.
4. Repeat for each link type, then click **Save**. The form closes and the labels are applied.

Deployment Profiles

Use Deployment Profiles to standardize your deployments, preventing errors and bad configuration values. Instead of configuring each appliance separately, you can create various Deployment Profiles and provision a device by applying the profile you want.

You can use Deployment Profiles to simplify provisioning, whether or not you choose to create and use Business Intent Overlays (BIO).

Configuring Deployment Profiles

Deployment Profiles are used during the deployment wizard to help streamline the installation process, gathering all required information.

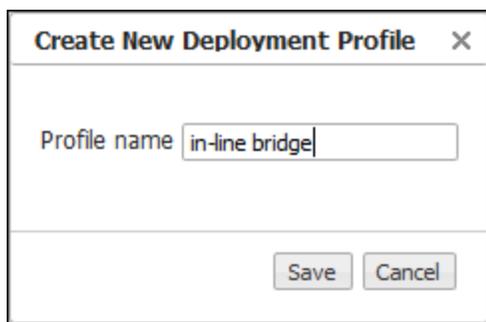
To configure a Deployment Profile

1. From the Configuration tab, select **Deployment Profiles**. The Deployment Profiles page appears.

2. At the top of the page, click **+Add**.

The **Create New Deployment Profile** form opens.

3. Enter a descriptive name for your profile, such as "in-line bridge", then click **Save**.



The form closes.

4. From the mode selector, choose **Bridge** (or whatever mode you want to use).

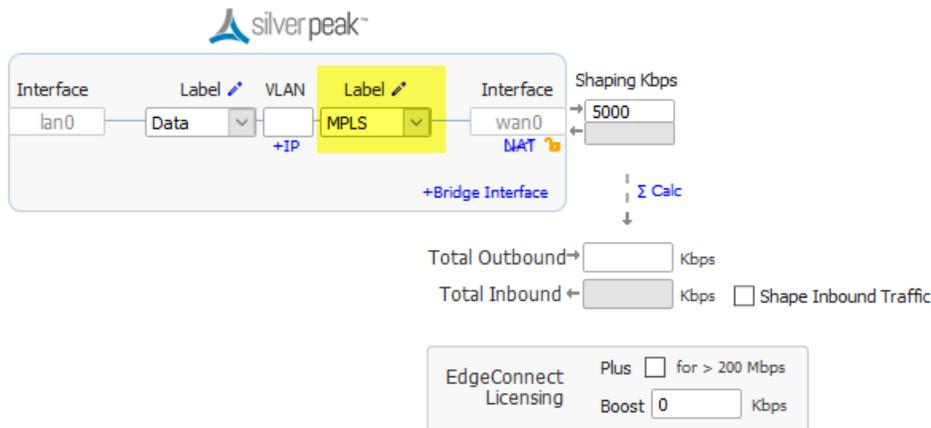


- **Router:** Single or dual WAN interfaces share LAN and WAN data traffic.
- **Bridge:** Uses a virtual interface (**bvi**) created by binding the WAN and LAN interfaces
- **Server:** Both management and data traffic use the **mgmt0** interface.

The mode form appears.

5. Map the labels you created in [Interface Labels](#) to your interface.

For example, in Bridge mode you would map the MPLS label to the WAN interface.



- **LAN:** identifies the traffic type, such as data, VoIP, or replication.
- **WAN:** identifies the service, such as MPLS or Internet.
- **NAT:** If the appliance is behind a NAT-ed interface, select NAT (without the strike-through). When using NAT, use In-Line Router Mode to ensure that addressing works properly. That means you configure paired single or dual WAN and LAN interfaces on the appliance.
- **WAN interface hardening** : In Router mode and in Bridge mode, you can provide security on any WAN-side interface by hardening the interface. This means:
 - For traffic inbound from the WAN, the appliance accepts only IPSec tunnel packets.
 - For traffic outbound to the WAN, the appliance only allows IPSec tunnel packets and management traffic.
 - Click the lock icon to toggle between hardening and unhardening an interface.

See the topic on [WAN Hardening](#).

6. Enter a value for **Shaping Kbps**.

If you are using asymmetric bandwidths, check **Shape Inbound Traffic**.

The Shaper shapes traffic by allocating bandwidth as a percentage of the system bandwidth. This page shows the actual inbound or outbound Shaping in kbps.

7. To add an internet interface, click **Bridge Interface**.

A new interface line appears.

8. Map this to the Internet interface.

The screenshot shows the Silver Peak configuration interface for WAN interfaces. It features a table with columns for Interface, Label, VLAN, another Label, and another Interface. Below the table are summary fields for Total Outbound and Total Inbound bandwidth, and an EdgeConnect Licensing section.

Interface	Label	VLAN	Label	Interface	Shaping Kbps
lan0	Data		MPLS	wan0	5,000
lan1	Voice		Internet	wan1	50,000

Total Outbound → 55,000 Kbps
 Total Inbound ← Kbps Shape Inbound Traffic

EdgeConnect Licensing Plus for > 200 Mbps
 Boost 0 Kbps

9. Enter a value for **Shaping Kbps**.
10. Click **Σ Calc** to automatically sum the shaping Kbps.
11. Click **Save**.

Your profile is created.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark** on the page for details about each field.

WAN Hardening

In Router or Bridge mode, you can provide security on any WAN-side interface by hardening the interface. This means:

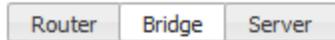
- For traffic inbound from the WAN, the appliance accepts only IPsec tunnel packets.
- For traffic outbound to the WAN, the appliance only allows IPsec tunnel packets and management traffic.

To configure WAN hardening

- From **Configuration > Deployment**, click the lock icon to toggle between hardening and unhardening an interface.

The screenshot shows the Silver Peak configuration interface for Deployment Profiles. At the top, there are tabs for 'Topology', 'Templates', and 'Deployment Profiles'. Below these, the 'Deployment Profiles' section is active, showing a dropdown for 'Profile Name' set to 'CA-Profile' and buttons for '+Add', 'Rename', and 'Delete'. There are also tabs for 'Router', 'Bridge', and 'Server'. The main configuration area is divided into 'LAN Interfaces +Add' and 'WAN Interfaces +Add'. The 'WAN Interfaces' table has columns for 'IP/Mask', 'Label', 'VLAN', 'Interface', and 'Shaping Kbps'. The 'wan0' interface is selected, and a yellow callout box with a lock icon and the text 'Interface hardened' is overlaid on the 'Interface' column. Below the table, there are fields for 'Total Outbound' and 'Total Inbound' in Kbps, and a checkbox for 'Shape Inbound Traffic'. At the bottom, there is a section for 'EdgeConnect Licensing' with a 'Plus' checkbox and a 'Boost' field set to '0'.

Deployment Modes



The versatility of Orchestrator enables you to configure your deployment in various modes:

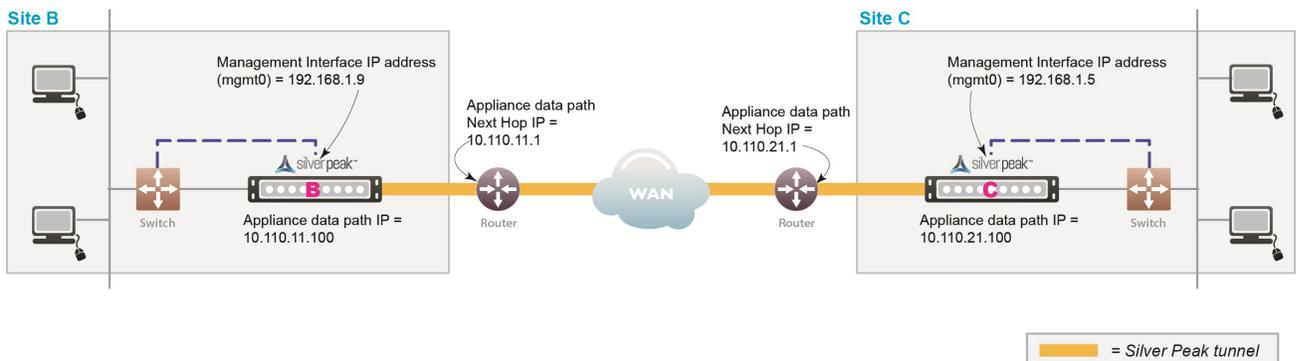
- **Router** - Can be in-line or out-of-path and uses multiple ports. Best for complex networks.
 - Do you want your traffic to be in-path or out-of-path? This mode supports both deployments. In-path deployment offers much simpler configuration.
 - Does your router support VRRP, WCCP, or PBR? If so, you may want to consider out-of-path Router mode deployment. You can set up more complex configurations, which offer load balancing between appliances and high availability.
 - If your LAN destination is behind a router or L3 switch, you need to add a LAN-side route (a LAN next-hop).
- **Bridge** - Can only be in-line and uses 2 or 4 ports. Easiest setup for simple networks.
 - Do you have a physical appliance or a virtual appliance?
 - A virtual appliance has no fail-to-wire, so you need a redundant network path to maintain connectivity if the appliance fails.
 - If your LAN destination is behind a router or L3 switch, you need to add a LAN-side route (a LAN next-hop).
 - If the appliance is on a VLAN trunk, then you need to configure VLANs on the Silver Peak so that the appliance can tag traffic with the appropriate VLAN tag.
- **Server** - Can only be out-of-path and uses 1 port (**mgmt0**). Best for backup configurations.

In-Line Deployments

In-Line deployments can be either Bridge mode or Router mode.

In an in-line deployment, the Silver Peak appliance is inserted in-line between the WAN router and the Ethernet switch on the LAN side of the network. In this mode, the appliance intercepts all packets destined for the WAN. Based on the Route Policy's MATCH criteria, or using Subnet Sharing-enabled auto-optimization, the appliance optimizes all flows that are directed to a tunnel. If the BIO overlay Down action is pass-through, all other traffic passes through the appliance without optimization. Else, the default action is to drop.

In Bridge mode, a failed appliance acts as a crossover cable. Best practice is to use a crossover cable between the appliance and the WAN-side router, and a standard Ethernet cable between the appliance and the LAN-side switch. Verify the physical layer connectivity between the L2 switch and router with the appliance turned off. If you don't receive a link on the router or switch, you need to correct the cabling.



Before deploying, gather information about your network, as shown in the following example:

Sample In-line Deployment Parameters

Hostname	B	C
Mode	In-line (Bridge)	In-line (Bridge)
Admin Password: Old	admin	admin
Admin Password: New / Confirm		

Hostname	B	C
mgmt1 IP Address / Mask	---	---
Time Zone		
NTP Server IP Address		
License (for virtual appliance only)		
mgmt0 IP Address / Mask	192.168.1.9/24	192.168.1.5/24
mgmt0 Next-hop IP Address	192.168.1.1	192.168.1.1
LAN Next-hop IP Address (optional)	---	---
Appliance data path IP Address / Mask	10.110.11.100/24	10.110.21.100/24
Appliance data path Next-hop IP	10.110.11.1/24	10.110.21.1/24

In-Line Bridge vs In-Line Router

In-line mode basically behaves the same in either Bridge or Router Mode with the following differences.



CAUTION: Installing or replacing an appliance in in-line mode requires taking down the network link.

- **In-Line Bridge Mode** is simpler and easier to set up.
 - Requires **mgmt0**.
 - No Ethernet LAN switch or WAN router configuration changes are required. All data flows through the appliance, no redirection is needed.
 - Connects (bridges) two halves of a single subnet.
 - Data Path addresses are not assigned to specific WAN or LAN interfaces. Multicast traffic is bridged as pass-through.
 - Pass-through traffic is automatically forwarded between LAN/WAN pairs.
- **In-Line Router Mode** works for larger, complex networks and is more flexible.
 - Needs at least two interfaces. Up to 6 are supported.
 - Connects to a different subnet on each interface.
 - Data Path addresses are assigned to each LAN or WAN interface. Multicast Traffic is dropped.
 - Pass-through traffic is forwarded between all interfaces if the destination subnet is known.
 - When the destination subnet is not known:
 - **LAN → WAN** goes to the first WAN interface next hop.
 - **WAN → LAN** goes to the first LAN interface next hop (if a next hop has been designated).

In-Line Overview

Following is how to deploy in-line Bridge mode using Subnet Sharing. In this scenario, the Silver Peak Appliance sits between the WAN router and the Ethernet switch.

Appliance Placement	Appliance placed in-line between Ethernet LAN switch and WAN router <ul style="list-style-type: none">• Appliance lan0 interface connects to Ethernet LAN switch• Appliance wan0 interface connects to WAN router
Fail-Safe Behavior	Fail-to-Wire (copper) & Fail-to-Glass (fiber): The appliance behaves as a crossover cable between the Ethernet LAN switch and the WAN router in any failure scenario (hardware, software, power). <ul style="list-style-type: none">• IMPORTANT: Ensure that the Ethernet LAN switch and the WAN router have compatible Ethernet interface physical configuration settings (speed and duplex settings can be found on the Configuration > Interfaces page). This is to ensure that traffic flows correctly if the Silver Peak appliance “Fails-to-wire”.
IP Addresses	This deployment model requires two IP addresses (on the same or separate subnets) <ul style="list-style-type: none">• Silver Peak Appliance data path IP address (to originate and terminate tunnel)• Silver Peak Management IP Address (for appliance configuration and management)

Out-of-Path Deployments

Whenever you place an appliance out-of-path, you must redirect traffic from the client to the appliance. Out-of-path deployments can be in either Server mode or Router mode.

Before deploying, gather the information about your network.

Fail Safe Behavior

Fail-safe behavior should always be tested before production deployment by ensuring that traffic continues to flow in each of the following cases:

- With the appliance in **bypass** state.
- With the appliance **powered off**.
- With the tunnels administratively **down**.

Policy-Based Routing (PBR)

PBR is configured on the router. No other special configuration is required on the appliance. This is also known as FBFB (Filter-Based Forwarding).

To deploy two Silver Peaks at the site, for redundancy or load balancing, you should use VRRP (Virtual Router Redundancy Protocol).

Site A with Peered Silver Peak Appliances OUT-OF-PATH

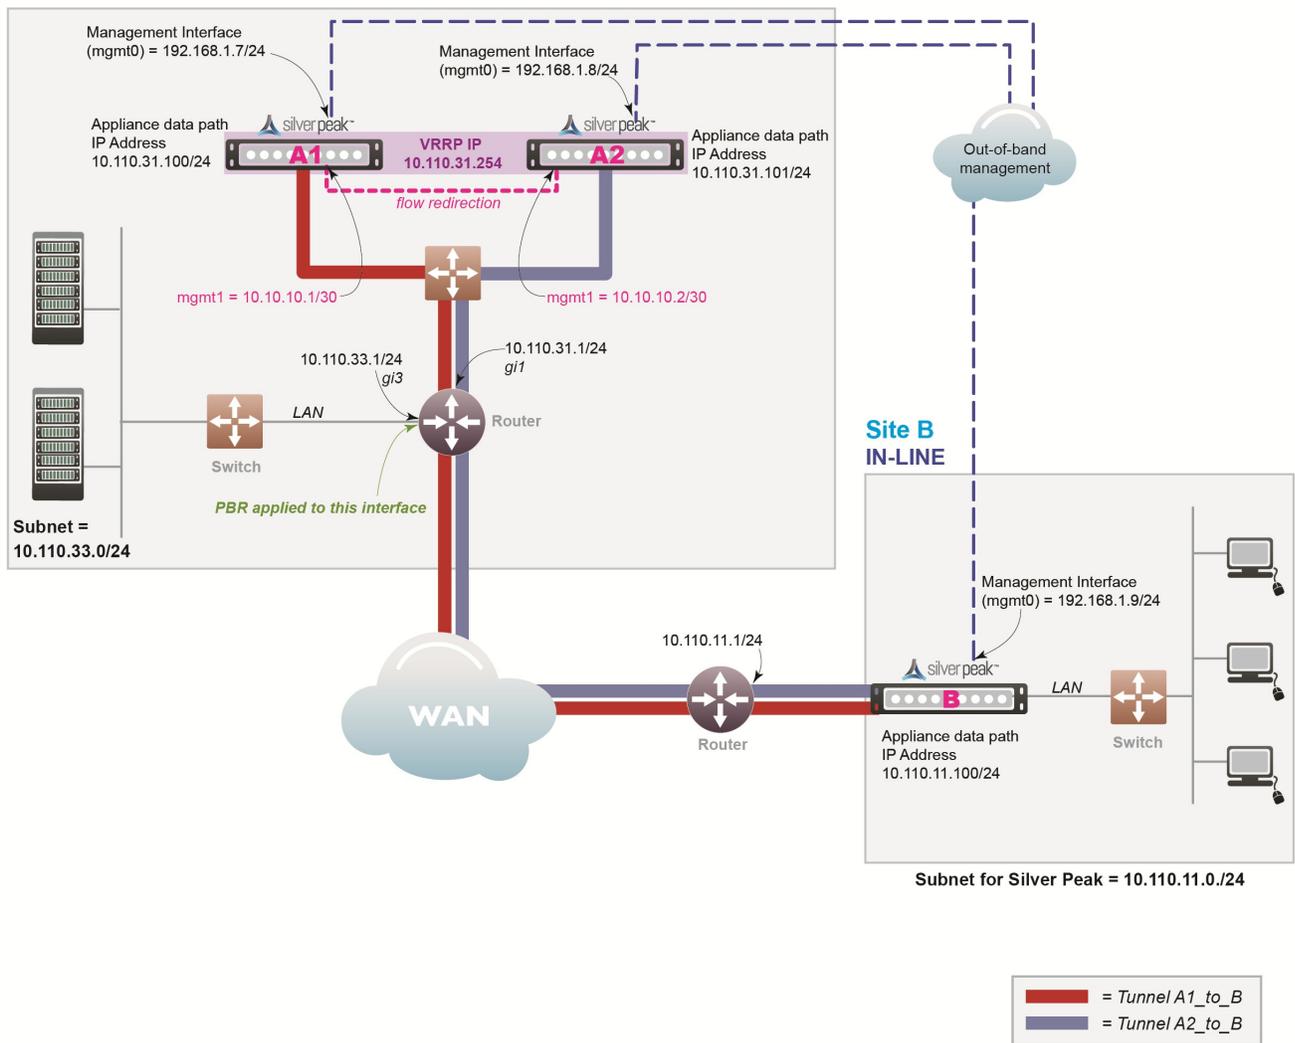


Figure 4. PBR Network

In this example, the Silver Peak appliance optimizes traffic to/from **10.110.33.0/24** and **10.110.11.0/24**.

PBR Overview

PBR Considerations

Appliance Placement	Both appliances are attached to the same available subnet via an Ethernet LAN switch: <ul style="list-style-type: none">• Each appliance wan0 interface connects to the Ethernet switch that is connected to the available WAN interface• Do not connect lan0 interface of either appliance
Failure Method	Fails Open: <ul style="list-style-type: none">• The failed appliance behaves as unconnected port in all failure cases (hardware, software, power).• The redundant Silver Peak appliance assumes the Silver Peak Appliance Virtual IP Address.• Remote appliances switch to the redundant appliance.
IP Addresses	This deployment model requires seven IP addresses: <ul style="list-style-type: none">• Each appliance needs a Silver Peak Appliance IP data path address (to originate and terminate tunnels).• The two appliances share one Silver Peak Appliance Virtual IP Address for VRRP.• Each appliance needs a Silver Peak Management IP Address (for appliance configuration and management).• If using flow redirection, need two more addresses. <p>Configure PBR on WAN router</p> <ul style="list-style-type: none">• Direct traffic from LAN (subnet/interface) destined for WAN to Silver Peak Appliances' Virtual IP Address• Do NOT enable this PBR on the interface to which the Silver Peak appliances connect

Fail-safe behavior should always be tested before production deployment by ensuring that traffic continues to flow in each of the following cases:

- With the appliance in **bypass** state.
- With the appliance **powered off**.

- With the tunnels administratively **down**.

PBR Example

Givens:

- You're not using DHCP.
- For all interfaces, speed and duplex are left at the default, which is auto-negotiation.
- Although not required, best practice is to use different subnets for **mgmt0** and the Appliance IP.

Hostname	A1	A2	B
Mode	Router / Out-of-Path	Router / Out-of-Path	Bridge / In-line
Admin Password: Old	admin	admin	admin
Admin Password: New / Confirm			
Time Zone			
NTP Server IP Address			
License (for virtual appliance only)			
mgmt1 IP Address / Mask	10.10.10.1/30	10.10.10.2/30	---
mgmt0 IP Address / Mask	192.168.1.7/24	192.168.1.8/24	192.168.1.9/24
	In this example, all mgmt0 IP addresses are in the same subnet. In your actual network, it's likely that mgmt0 IP addresses are in different subnets.		
mgmt0 Next-hop IP Address	192.168.1.1	192.168.1.1	192.168.1.1
Appliance data path IP Address / Mask	10.110.31.100/24	10.110.31.101/24	10.110.11.100/24
Appliance data path Next-hop IP	10.110.31.1/24	10.110.31.1/24	10.110.11.1/24
LAN Next-hop IP Address (optional)	not applicable	not applicable	---

Hostname	A1	A2	B
	LAN next-hop IP is only required when there are subnets for which the Silver Peak appliance does not have a configured IP address.		
VRRP Group ID	1	1	---
VRRP Virtual IP Address (VIP)	10.110.31.254	10.110.31.254	not applicable
VRRP Priority	130	128	not applicable

PBR Configuration Checklist

You can print out this page and use it as a reference.

- Gather all the IP addresses needed for setup. See [PBR Example](#).

- Install the appliance into the network
 - Physical appliance:** Connect both appliances to the same available subnet via an Ethernet LAN switch. Verify connectivity, connect power, and verify LEDs.

 - Virtual appliance:** Configure the hypervisor, with the required interfaces.

- Configure the Site A appliances
 - In a browser, access and use the **Initial Config Wizard** to configure each appliance. See [Add Appliances](#).

 - Reboot the appliances after finishing the configuration.

- Configure VRRP for the Site A peers
 - Configure one appliance to be the Master, and the other to be the Backup.

- Configure flow redirection for the Site A peers
 - When you create a cluster, the peers keep track of which appliance owns each flow. If the path between client and server isn't the same in both directions, the flow is redirected to the appliance that first saw it and "owns" it.

- Configure Site B appliances
 - In a browser, access and use the **Initial Config Wizard** to configure the appliance. See [Add Appliances](#).
 - Reboot the appliances after finishing the configuration.
- Verify appliance connectivity. See [Verify Appliance Connectivity](#).
 - Do NOT proceed until you verify connectivity.
- Enable subnet sharing. See [Enable Subnet Sharing](#).
 - This prepares each appliance to share local subnets.
- Apply Deployment Profiles, Business Intent Overlays, and Configuration Templates. See [Template Groups](#).
- Configure the router
 - Access the router command line interface, and configure the router for policy-based routing.
- Test the connectivity from both ends. See [Verify Traffic](#).
 - Verify the tunnel is up and that flows are being optimized.

Web Cache Communication Protocol (WCCP)

WCCP is configured on both the router and the Silver Peak appliance. You can also use WCCP for redundancy and load balancing.

In the following scenario, the Silver Peak appliances are not connected in the direct path of the network traffic. As a result, a network traffic redirection technique is used to forward traffic to the appliance.

WCCP supports the redirection of any TCP or UDP connections to appliances participating in WCCP Service Groups. The appliance intercepts only those packets that have been redirected to it. The appliance accelerates traffic flows that the Route Policy directs to a tunnel; all other traffic passes through the appliance unmodified. (Traffic might be dropped depending on overlay action for the default overlay.)

In the unlikely event that the appliance fails, WCCP on the WAN router removes the appliance from the WCCP Service Group and resumes forwarding traffic normally, according to its routing tables.

At Site A, both the router and the participating appliance require a separate WCCP service group for each protocol used in the tunnel. So, if a tunnel uses both TCP and UDP, you must create a separate WCCP Service Group for each protocol (TCP and UDP) used in the A-to-B tunnel.

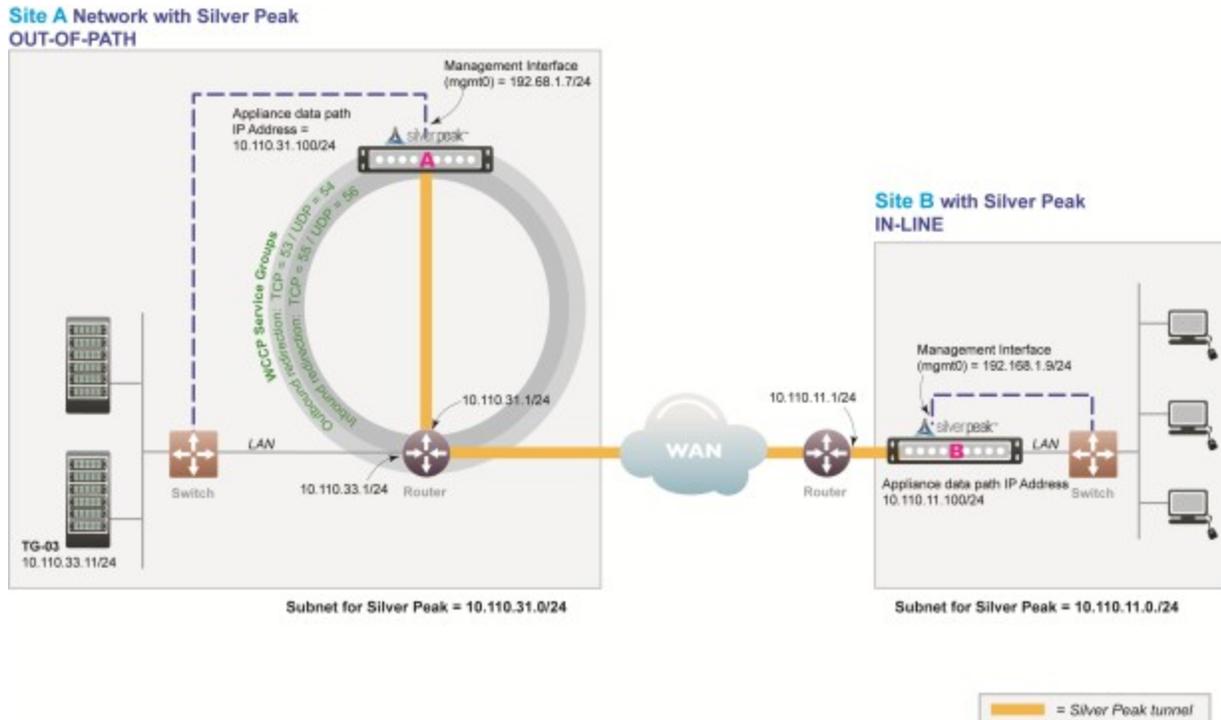


Figure 5. WCCP Network

In this example, the Silver Peak appliances optimize traffic to/from **10.110.33.0/24** and **10.110.11.0/24**.



NOTE: You don't need a spare router port for this configuration. The Silver Peak appliance can be connected to an existing or newly configured sub-interface on the router via a VLAN trunk such that a spare port on the LAN switch can be used for the physical connection.

WCCP Overview

WCCP Considerations

- | | |
|----------------------------|---|
| Appliance Placement | <p>Appliance attached in network, reachable by WAN router</p> <ul style="list-style-type: none"> • Appliance wan0 interface connects to network • Do not connect lan0 interface |
|----------------------------|---|

Fail-Safe Behavior	<p>WCCP recognizes failed appliance</p> <ul style="list-style-type: none">• Appliance removed from WCCP v2 Service Group• WAN router resumes forwarding traffic normally according to its routing tables
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets)</p> <ul style="list-style-type: none">• Silver Peak Appliance data path IP address (to originate and terminate tunnels)• Silver Peak Management IP Address (for appliance configuration and management) <p>Configure WCCP on the Silver Peak appliance and the WAN router. Service Group IDs on the router and appliance must match.</p> <ul style="list-style-type: none">• Configure two WCCP v2 Service Groups on the Silver Peak appliance (one for TCP and one for UDP)• Configure two WCCP v2 Service Groups on the WAN router (one for TCP and one for UDP)

Fail-safe behavior should always be tested before production deployment by ensuring that traffic continues to flow in each of the following cases:

- With the appliance in **bypass** state.
- With the appliance **powered off**.
- With the tunnels administratively **down**.

WCCP Best Practice

Tips for Deployment

- Inbound WCCP redirection is preferred over outbound [also known as ingress/egress] redirection because inbound redirection is less CPU-intensive on the router. Inbound redirection is done in hardware where as outbound is done in software.
 - For Catalyst 6000/76xx deployments, use only inbound redirection to avoid using "redirection exclude in", which is not understood by the switch hardware and must be processed in software.
 - For Catalyst 6000/76xx deployments, use L2 redirection for near line-rate redirection. Silver Peak appliances automatically negotiate assignment and forwarding methods with all routers and L3 switches from Cisco to the best possible combination that the router or L3 switch supports.

- WCCPv2 interception forwards all packets from the router or L3 switch to the appliance. Special care should be taken when traffic redirected to the appliance has to be returned back to the router or L3 switch. For many routers the return traffic is delivered via L2 so there is no CPU impact. However, Catalyst 6000/76xx switches returns via GRE so the CPU can be negatively impacted unless Force L2 return is enabled on the appliance.
 - **Force L2** Return should only be enabled when the interface/VLAN that the appliance is connected to is not also an interface with the redirection applied to.
- The appliance should always be connected to an interface/VLAN that does not have redirection enabled – preferably a separate interface/VLAN would be provided for the appliance.
- The appliance and Catalyst switch negotiate which redirect and return method to use when the service group is formed. There can be many access VLANs on the aggregation switches. Redirection is configured on all VLANs that need optimization. Layer 2 switching ports, including trunk ports, are not eligible for redirection.
- If **Auto Optimization** is used for matching traffic to be optimized via the appliance, WCCP redirection must also be applied on the uplinks of the router or L3 switch to the core/WAN.
- If WCCP redirection is needed on both the WAN and the LAN, the preferred configuration on the appliance is to set the WCCP group configured on the WAN to **wan-ingress** and the group configured on the LAN to **lan-ingress**.
 - The configuration of wan-ingress and lan-ingress ensures that load balancing is symmetrical in both directions of a flow.
 - **wan-ingress** uses the destination address for distribution in the router/L3 switch table
 - **lan-ingress** uses the source address for distribution.
- If Route Policies are used for matching traffic to be optimized via the appliance, WCCP redirection is not required on the core uplinks, only the access/LAN links. If Active/Active redistribution is enabled with route policies, then flow redirection is required to handle asymmetrical flows caused by load balancing. Flow redirection can handle millions of flows and ensures that the owner of a given flow always receives the TCP flow for processing.

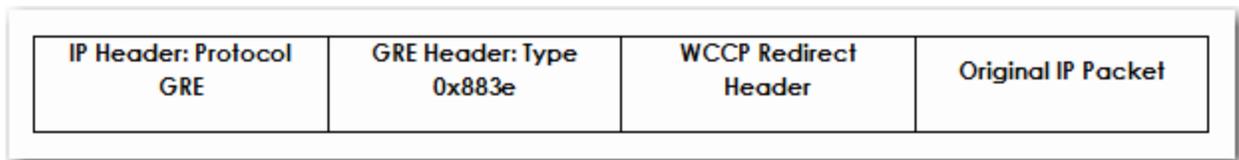
GRE and L2 Redirection

Packet redirection is the process of forwarding packets from the router or L3 switch to the appliance. The router or L3 switch intercepts the packet and forwards it to the appliance for optimization. The two methods of redirecting packets are Generic Routing Encapsulation (GRE) and L2 redirection. GRE is processed at Layer 3 while L2 is processed at Layer 2.

- Silver Peak appliances support both GRE and L2 Redirection.
- Silver Peak appliances support both Mask and Hash assignments.
 - Additional mask and hash assignment adjustment can help fine-tune the distribution of traffic to the appliances. The advanced configuration for fine-tuning can be found in the **Advanced Settings** feature of the WCCP configuration on the appliance.
 - Mask assignments are set up on the appliance. The first appliance that joins the WCCP service group determines the redirection method and masking value – this appliance is referred to as the “designated” appliance. Subsequent appliances that join the group must have the same redirection and mask value setup; otherwise, they are not active participants in the WCCP group.
 - Appliances support both Hash and Mask capabilities for optimal throughput. The preferred WCCP configuration on the appliance is to leave both assignment and forwarding method to “either” which will allow the preferred negotiation to happen between the appliance and the router or L3 switch when WCCP is first enabled.

GRE

GRE is a protocol that carries other protocols as its payload:



In this case, the payload is a packet from the router to the appliance. GRE works on routing and switching platforms. It allows the WCCP clients to be separate from the router via multiple hops. Because GRE is processed in software, router CPU utilization increases with

GRE redirection. Hardware-assisted GRE redirection is available on the Catalyst 6500 with Sup720.

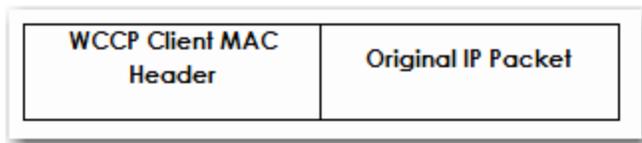
L2 Redirection

L2 redirection requires the appliance to be in the same subnet as the router or switch (L2 adjacency).

The switch rewrites the destination L2 MAC header with the appliance MAC address. The packet is forwarded without additional lookup.

L2 redirection is done in hardware and is available on the Catalyst 6500/7600 platforms. CPU utilization is not impacted because L2 redirection is hardware-assisted; only the first packet is switched by the Multilayer Switch Feature Card (MSFC) with hashing.

After the MSFC populates the NetFlow table, subsequent packets are switched in hardware. L2 redirection is preferred over GRE because of lower CPU utilization.



There are two methods to load balance appliances with L2 redirection: hashing and masking.

WCCP Example

Givens:

- You're not using DHCP.
- Speed and duplex for all interfaces are left at the default, auto-negotiation.
- Although not required, best practice is to use different subnets for **mgmt0** and the Appliance IP.
- Silver Peak Appliance peered with an L3 router using WCCP

Hostname	A	B
Mode	Out-of-path (Router)	In-line (Bridge)

Hostname	A	B
Admin Password: Old	admin	admin
Admin Password: New / Confirm		
Time Zone		
NTP Server IP Address		
License (for virtual appliance only)		
mgmt1 IP Address / Mask	10.10.10.1/24	---
mgmt0 IP Address / Mask	192.168.1.7/24	192.168.1.9/24
	In this example, all mgmt0 IP addresses are in the same subnet. In your actual network, it's likely that mgmt0 IP addresses are in different subnets.	
mgmt0 Next-hop IP Address	192.168.1.1	192.168.1.1
Appliance data path IP Address / Mask	10.110.31.100/24	10.110.11.100/24
Appliance data path Next-hop IP	10.110.31.1/24	10.110.11.1/24
LAN Next-hop IP Address (optional)	not applicable	---
	LAN next-hop IP is only required when there are subnets for which the Silver Peak appliance does not have a configured IP address.	
WCCP Service Groups for outbound redirection	53. (TCP) 54. (UDP)	---
WCCP Service Groups for inbound redirection	55. (TCP) 56. (UDP)	---
WCCP Weight (default)	100	not applicable

WCCP Configuration Checklist

You can print out this page and use it as a reference.

- Gather all IP addresses needed for setup. See [WCCP Example](#). Install the appliance into the network.
 - Physical appliance:** Connect the Site A appliance to the Site A router, and insert the Site B appliance between its WAN edge router and the Ethernet switch. Verify connectivity, connect power, and verify LEDs.
 - Virtual appliance:** Configure the hypervisor, with the required interfaces.
- Configure the Site A router for WCCP.
 - Configure an ACL that redirects all traffic from the Site A subnet to the Site B subnet
 - Configure two WCCP Service Groups — one for UDP, one for TCP
 - Associate the ACL with the Service Group
 - Enable WCCP on the appropriate router interface
- Configure the Site A appliance for out-of-path deployment.
 - Access the Initial Config Wizard to assign Appliance IP and Management IP addresses for the Site A appliance. See [Add Appliances](#).
 - Reboot the appliance.

- Configure the WCCP Service Groups on the Site A appliance.
 - Create a pair of Service Groups (TCP, UDP) for outbound redirection.
 - Inbound redirection isn't needed when using subnet sharing..
- Configure the Site B appliance for in-line deployment.
- IMPORTANT:** The WAN Next Hop IP Address must be the IP address of the WAN edge router. This may or may not be the same as the Management Interface Next Hop IP Address for hosts on the LAN side of your network. If in doubt, check with your network administrator.
 - Run the **Initial Config Wizard** to set up the Site B appliance in Bridge mode.
 - Reboot the appliance.
- Verify the appliance is connected. See [Verify Appliance Connectivity](#).
 - Ensure that the cable connections are secure and that all IP addresses are configured correctly..
 - Do NOT proceed until you have verified connectivity.**
- Enable subnet sharing. See [Enable Subnet Sharing](#).
 - This prepares each appliance to share local subnets.
- Apply Deployment Profiles, Business Intent Overlays, and Configuration Templates. See [Template Groups](#).
- Test the connectivity from both ends. See [Verify Traffic](#).

- Verify that the tunnel is up and that flows are being optimized.

Virtual Router Redundancy Protocol (VRRP)

Virtual Router Redundancy Protocol (VRRP) gives you redundancy (backup) with minimal equipment. The Silver Peak appliances must be deployed in Router mode (which EdgeConnect does automatically by default) to use this feature.

VRRP routing is often used (with subnet sharing) when

- a Silver Peak appliance uses an existing router when no spare router port is available.
- OR -
 - using redundant Silver Peak appliances with PBR.
-

VRRP Peering to a WAN Router

This out-of-path deployment method configures VRRP on a common virtual interface. The possible scenarios are:

- When no spare router port is available, a single appliance uses VRRP to peer with a router (or Layer 3 switch). This is appropriate for an out-of-path deployment where no redundancy is needed.
- A pair of active, redundant appliances use VRRP to share a common, virtual IP address at their site. This deployment assigns one appliance a higher priority than the other, thereby making it the Master appliance, and the other, the Backup.

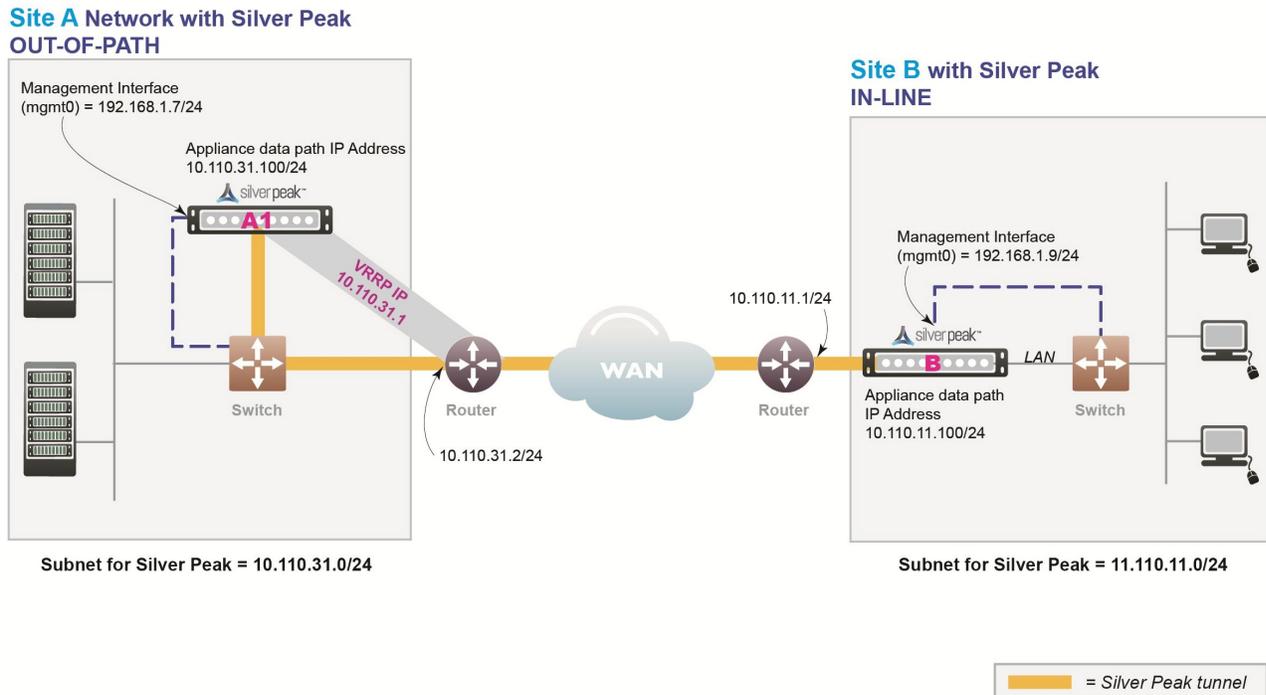


Figure 6. VRRP Peering to a WAN Router Diagram

In this example, the Silver Peak appliance optimizes traffic to/from **10.110.31.0/24** and **10.110.11.0/24**.

VRRP Peering to a WAN Router Considerations

Appliance Placement	Appliance shares LAN segment with existing equipment <ul style="list-style-type: none"> • Appliance wan0 interface connects to Ethernet LAN switch • Do not connect lan0 interface
Failure Method	Fails-Open: <ul style="list-style-type: none"> • The appliance behaves as an unconnected port in all failure cases (hardware, software, power) • WAN router assumes Virtual IP Address and forwards traffic normally

IP Addresses

This deployment model requires three IP addresses:

- Silver Peak Appliance data path IP address (to originate and terminate tunnel)
- Silver Peak Management IP Address (for appliance configuration and management)
- Virtual IP Address (VIP) shared by Silver Peak appliance and the WAN router

The VIP must be the default gateway for the clients and servers on the LAN subnet.

NOTE: Typically, this would be the current default gateway, to avoid client reconfigurations.

- The Silver Peak appliance must share the default gateway VIP with WAN router using VRRP.
 - The Silver Peak appliance must be configured with higher priority and preemption ensure VRRP reverts to the appliance.
-

Fail-safe behavior should always be tested before production deployment by ensuring that traffic continues to flow in each of the following cases:

- With the appliance in **bypass** state.
- With the appliance **powered off**.
- With the tunnels administratively **down**.

VRRP to WAN Example

Givens:

- The IP address of the router must be changed and the VRRP VIP (Virtual IP) address added to the router.
- The VIP address uses the existing router address; this means you don't need to modify the client default gateway.
- The Silver Peak appliance becomes the primary default gateway for all users in that network.
- In the unlikely event that the Silver Peak appliance fails, the router automatically becomes the default gateway.

VRRP - Example settings

Hostname	A1	B
Mode	Out-of-Path (Router)	In-Line (Bridge)
Admin Password: Old	admin	admin
Admin Password: New / Confirm		
Time Zone		
NTP Server IP Address		
License (for virtual appliance only)		
mgmt0 IP Address / Mask	192.168.1.7/24	192.168.1.9/24
	In this example, all mgmt0 IP addresses are in the same subnet. In your actual network, it's likely that mgmt0 IP addresses are in different subnets.	
mgmt0 Next-hop IP Address	192.168.1.1	192.168.1.1
Appliance data path IP Address / Mask	10.110.31.100/24	10.110.11.100/24
Appliance data path Next-hop IP	10.110.31.2/24	10.110.11.1/24
LAN Next-hop IP Address (optional)	not applicable	---
	LAN next-hop IP is only required when there are subnets for which the Silver Peak appliance does not have a configured IP address.	
VRRP Group ID	1	---
VRRP Virtual IP Address (VIP)	10.110.31.1	not applicable
VRRP Priority	130	not applicable

VRRP with PBR Configuration Checklist

You can print out this page and use it as a reference.

- Gather all the IP addresses needed for setup. See [VRRP with PBR Example](#). Install the appliance into the network.
- Start Orchestrator and let it find the appliances.
- Verify appliance connectivity. See [Verify Appliance Connectivity](#).
 - Do NOT proceed until you have verified connectivity.**
- Enable subnet sharing. See [Enable Subnet Sharing](#).
 - This prepares each appliance to share local subnets.
- Apply Deployment Profiles, Business Intent Overlays, and Configuration Templates. See [Template Groups](#).
 - Set up the VRRP template.
- Optional.* Manually add non-local subnets that aren't directly connected to an appliance interface. This option is rarely used.
- Configure the router.
 - Access the router's command line interface, and configure the router for policy-based routing.
- Test the connectivity from both ends. See [Verify Traffic](#).
 - Verify that the tunnel is up and that flows are being optimized.

VRRP Redundant Appliances with PBR

In this example, Site A deploys two redundant appliances out-of-path (Router mode), used as Active and Standby. Site B deploys a single appliance in-line (Bridge mode).

The peered appliances at Site A use the Virtual Router Redundancy Protocol (VRRP) to create and share a common IP address, called the Virtual IP (VIP) address. Configuring for high availability assigns one appliance a higher priority than the other appliance, thereby making it the Master, and the other, the backup.

The appliance at Site B has separate tunnels going to each of the two appliances at Site A:

- If one of the appliances at Site A is down, then Site B only sends traffic to the appliance (tunnel) that is up.
- If both appliances at Site A are up, then Site B sends traffic to the tunnel (appliance) that has higher VRRP priority.

**Site A with Peered Silver Peak Appliances
OUT-OF-PATH**

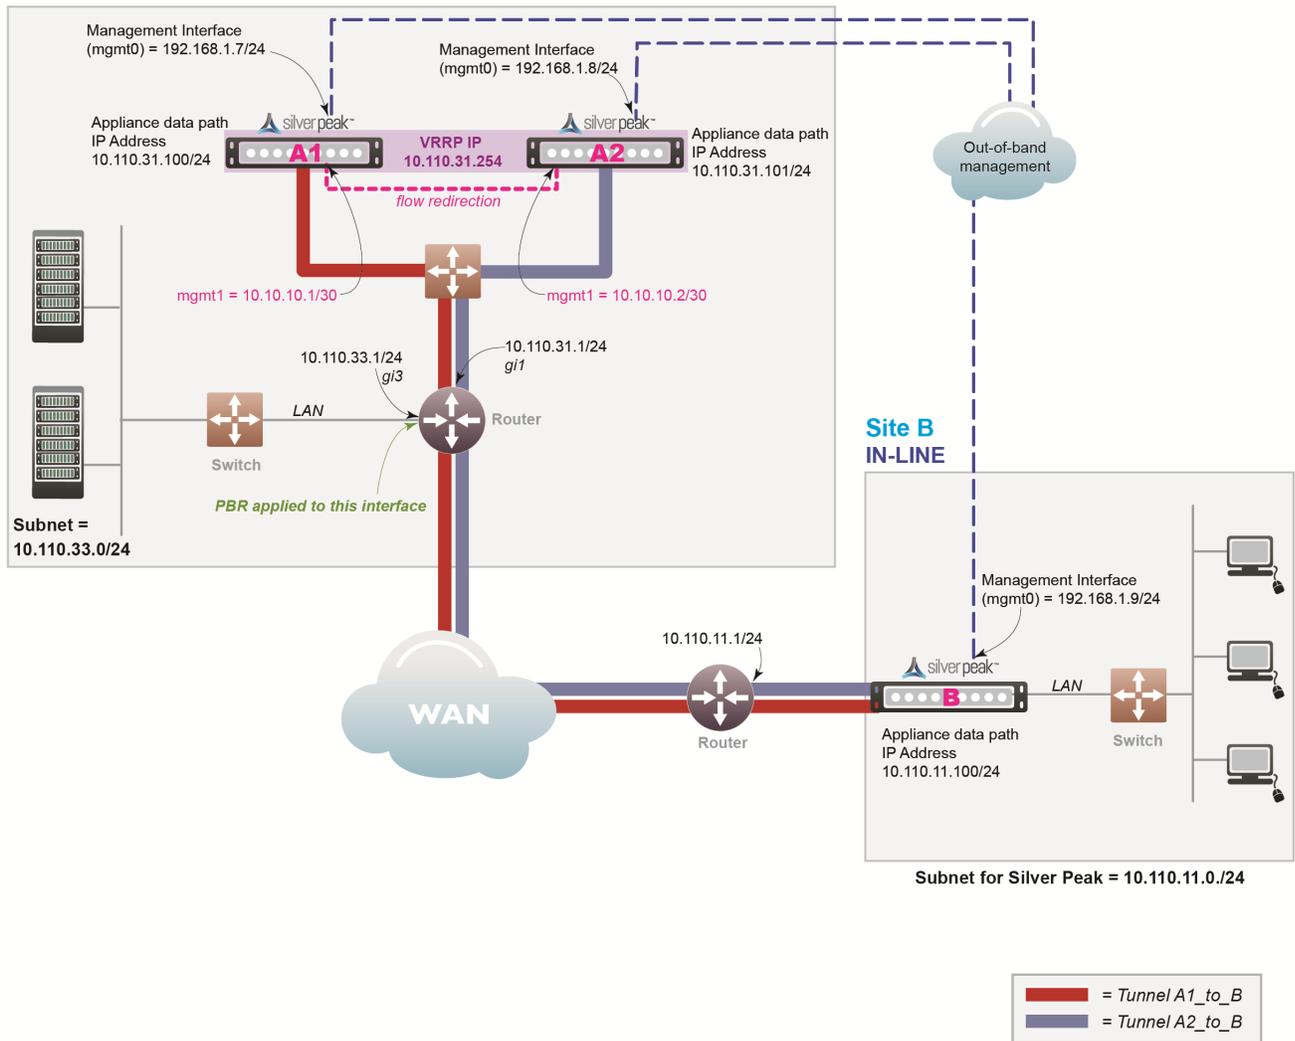


Figure 7. VRRP with PBR Appliances Diagram

In this example, the Silver Peak appliances optimize traffic to/from **10.110.31.0/24** and **10.110.11.0/24**.

Appliance Placement

Both appliances are attached to the same available subnet via an Ethernet LAN switch:

- Each appliance's wan0 interface connects to the Ethernet switch that is connected to the available WAN interface
- Do not connect lan0 interface of either appliance

Failure Method	<p>Fails Open:</p> <ul style="list-style-type: none">• The failed appliance behaves as unconnected port in all failure cases (hardware, software, power).• The redundant Silver Peak appliance assumes the Silver Peak Appliance Virtual IP Address.• Remote appliances switch to the redundant appliance.
IP Addresses	<p>This deployment model requires five IP addresses:</p> <ul style="list-style-type: none">• Each appliance needs a Silver Peak Appliance IP data path address (to originate and terminate tunnels).• The two appliances share one Silver Peak Appliance Virtual IP Address for VRRP.• Each appliance needs a Silver Peak Management IP Address (for appliance configuration and management). <p>Configure PBR on WAN router</p> <ul style="list-style-type: none">• Direct traffic from LAN (subnet/interface) destined for WAN to the Virtual IP Address of the Silver Peak Appliance.• Do NOT enable this PBR on the interface to which the Silver Peak appliances connect

Fail-safe behavior should always be tested before production deployment by ensuring that traffic continues to flow in each of the following cases:

- With the appliance in **bypass** state.
- With the appliance **powered off**.
- With the tunnels administratively **down**.

VRRP with PBR Example

Givens:

- You're not using DHCP.
- For all interfaces, speed and duplex are left at the default, which is auto-negotiation.
- Although it isn't a requirement, it's considered a best practice to use different subnets for mgmt0 and the Appliance IP.

Hostname	A1	A2	B
Mode	Router / Out-of-Path	Router / Out-of-Path	Bridge / In-line
Admin Password: Old	admin	admin	admin
Admin Password: New / Confirm			
Time Zone			
NTP Server IP Address			
License (for virtual appliance only)			
mgmt1 IP Address / Mask	10.10.10.1/30	10.10.10.2/30	---
mgmt0 IP Address / Mask	192.168.1.7/24	192.168.1.8/24	192.168.1.9/24
	In this example, all mgmt0 IP addresses are in the same subnet. In your actual network, it's likely that mgmt0 IP addresses are in different subnets.		
mgmt0 Next-hop IP Address	192.168.1.1	192.168.1.1	192.168.1.1
Appliance data path IP Address / Mask	10.110.31.100/24	10.110.31.101/24	10.110.11.100/24
Appliance data path Next-hop IP	10.110.31.1/24	10.110.31.1/24	10.110.11.1/24
	Only required when there are subnets for which the Silver Peak appliance does not have a configured IP address.		
LAN Next-hop IP Address (optional)	not applicable	not applicable	---
VRRP Group ID	1	1	---
VRRP Virtual IP Address (VIP)	10.110.31.254	10.110.31.254	not applicable
VRRP Priority	130	128	not applicable

VRRP Peering to WAN Configuration Checklist

You can print out this page and use it as a reference.

- Gather all the IP addresses needed for setup. See [VRRP to WAN Example](#). Install the appliance into the network.
 - Physical appliance:** Connect the Site A appliance to the Site A router, and insert the Site B appliance between its WAN edge router and the Ethernet switch. Verify connectivity, connect power, and verify LEDs.
 - Virtual appliance:** Configure the hypervisor, with the required interfaces.
- Start Orchestrator and let it find the appliances.
- Verify appliance connectivity. See [Verify Appliance Connectivity](#).
 - Do NOT proceed until you have verified connectivity.**
- Enable subnet sharing. See [Enable Subnet Sharing](#).
 - This prepares each appliance to share local subnets.
- Apply Deployment Profiles, Business Intent Overlays, and Configuration Templates. See [Template Groups](#).
 - Set up the VRRP Template.
- Configure the router.
 - Access the router's command line interface, and configure the router for policy-based routing.

- Test the connectivity from both ends. See [Verify Traffic](#).
- Verify that the tunnel is up and that flows are being optimized.

Host-Based Redirection

Host routing (also called Host Based Forwarding or Storage Based Forwarding) is when the server/end station has a default or subnet-based static route that points to the Silver Peak appliance as its next hop.

Host routing is the preferred method when a virtual appliance is using a single interface, **mgmt0**, for datapath traffic (also known as Server Mode).

The following example shows the end devices as storage arrays, although they could also be PCs or servers.

- The end devices point to the Silver Peak as next hop via static route or as default gateway.
- The Silver Peaks are performing rate limiting, as opposed to the storage arrays.

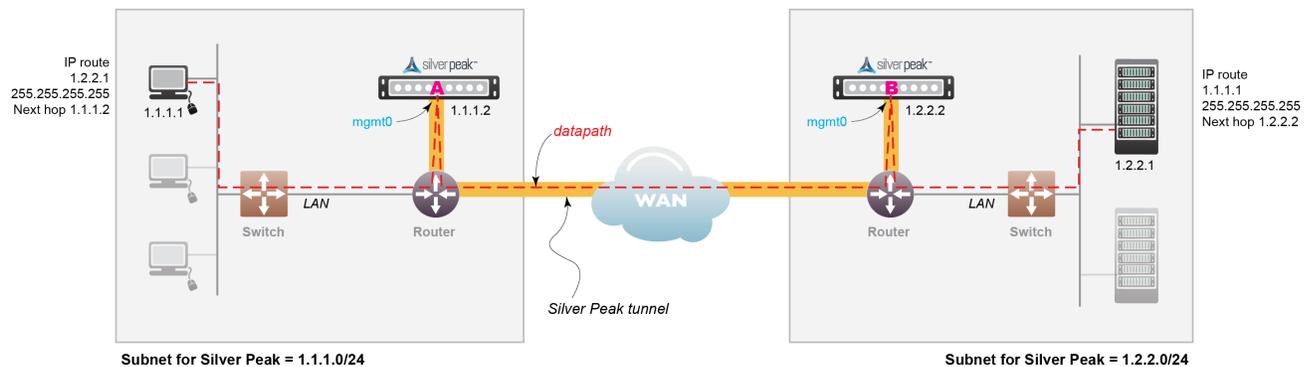


Figure 8. Host-based Redirection Network

In this example:

- The device on the left (IP 1.1.1.1) wants to back up its data to the device on the right (IP 1.2.2.1)
- Two Silver Peaks (1.1.1.2 and 1.2.2.2) are installed with a tunnel (green dashed line) between them.

- Each end device is on the same subnet as its corresponding Silver Peak.
 - Each end device points to its Silver Peak as the next hop using a static route.
 - No changes are needed to switches or routers.
- Rate limit is controlled by the Silver Peaks and not by the routers.
 - Allows you to get maximum performance without exceeding the allocated bandwidth.

To ensure end-to-end connectivity in case of appliance failure, consider using VRRP between the appliance and a router, or the appliance and another redundant Silver Peak.

Server Mode Deployments

Server mode is the default for virtual appliances and:

- Only uses the **mgmt0** interface for management and datapath traffic.
- Uses one IP address.
- Requires traffic redirection.
- Can only be deployed out-of-path.

The screenshot shows the 'Deployment Profiles' configuration interface. At the top, there is a 'Profile Name' dropdown set to 'Main' and a 'View Deployment' link. Below this are three mode selection buttons: 'Router', 'Bridge', and 'Server', with 'Server' being the active mode. The Silver Peak logo is centered above the configuration area. A 'Label' dropdown is set to 'MPLS'. The 'Interface' is 'mgmt0' with 'NAT' indicated below it. To the right, 'Shaping Kbps' is set to '10,000'. A dashed arrow labeled 'Σ Calc' points from the shaping field to the 'Total Outbound' field, which is also set to '10,000' Kbps. The 'Total Inbound' field is empty, and there is a checkbox for 'Shape Inbound Traffic' which is unchecked. At the bottom of the configuration area, there is an 'EdgeConnect Licensing' section with a 'Plus' checkbox (unchecked) for '> 200 Mbps' and a 'Boost' field set to '0' Kbps. At the very bottom of the page are 'Save', 'Save As', and 'Cancel' buttons.

Figure 9. Server mode page

Unlike other Out-of-Path deployments, this mode requires the hosts to configure a route or gateway to the Silver Peak appliance IP address to redirect traffic to the Silver Peak.

Traffic redirected to the Silver Peak is optimized and placed in a tunnel. This includes:

- Traffic arriving in a tunnel from another Silver Peak
- Any non-tunnelized traffic that needs to be directed to the local appliance, such as optimized traffic that has not yet been tunnelized.

Deploying in NAT Environments

If the appliance is behind a Network Address Translation (NAT) interface, select NAT (without the strike-through).

The screenshot displays the 'Deployment Profiles' configuration window for a Silver Peak SD-WAN appliance. The profile is named 'MediumBranch'. Under the 'Router' tab, the configuration is split into LAN and WAN interfaces. The LAN interface 'lan0' is configured with 'None' as the label and 'No DHCP' as the DHCP setting. The WAN interface 'wan0' is configured with 'MPLS' as the label and 'NAT' as the connection type. Below the interface configurations, the 'Total Outbound' and 'Total Inbound' traffic limits are both set to 150,000 Kbps. The 'EdgeConnect Licensing' section shows 'Plus' selected for devices with more than 200 Mbps and a 'Boost' of 0 Kbps. At the bottom, there are 'Save', 'Save As', and 'Cancel' buttons.

Figure 10. NAT Deployment

For deployments in the cloud, best practice is to NAT all traffic — either inbound (WAN-to-LAN) or outbound (LAN-to-WAN), depending on the direction of initiating request. This avoids black-holing that can result from cloud-specific IP address requirements.

- Enabling **NAT** applies NAT policies to all traffic—pass-through as well as optimized traffic—which ensures that black-holing doesn't occur. **NAT** on outbound only applies to pass-through traffic.
- If Fallback is enabled, the appliance moves to the next IP (if available) when ports are exhausted on the current NAT IP.

In general, when applying NAT policies, configure separate WAN and LAN interfaces to ensure that NAT works properly. Do this by deploying the appliance in Router mode in-path with two (or four) interfaces.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

Example Deployments

Here are some examples of possible deployment scenarios.

Router Mode Considerations

- In-Line Router Mode is recommended whenever possible vs. traditional router mode deployment.
- Limitations when using plain router mode over In-Line Router Mode:
 - Unable to act as a firewall (can't distinguish trusted and untrusted zones).
 - Inaccurate pass-through stats (because inbound and outbound are combined into outbound).
 - Inaccurate shaping/QoS (because inbound and outbound are munged into outbound).
 - Unable to use pass-through tunnels in future, or enable service chaining.
- Some limitations to ILRM in 8.0.3 don't allow control of pass-through traffic out of any interface except for **wan0** and in some cases might cause a network loop.



NOTE: Router mode is suggested for out-of-path deployments such as PBR and WCCP until these limitations are fixed or when the topology is not susceptible to creating a loop.

In-Line Router Mode (Router + Firewall)

- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- The NAT option must be enabled on the **wan1** interface.
- Local internet access is disabled if you enable WAN hardening. Other interfaces are available for more WAN links.

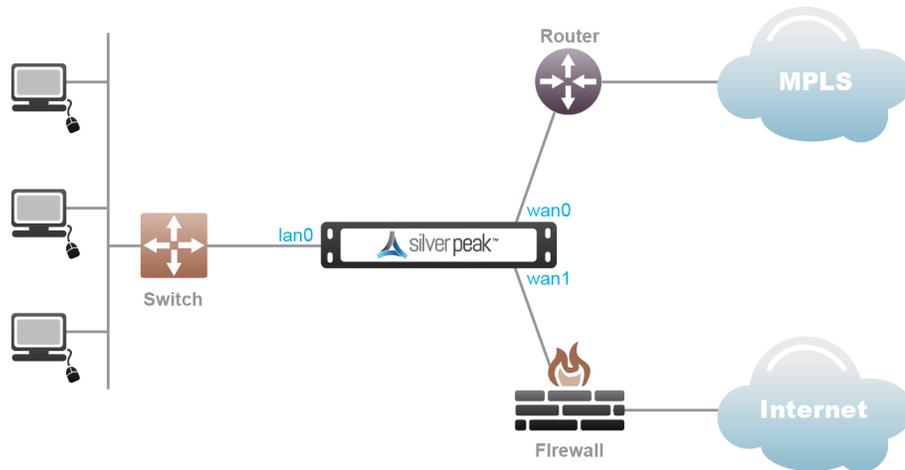


Figure 11.

Figure 12. In-Line Router Mode (Router + Firewall)

In-Line Router Mode (Router + Direct Internet)

- A default route to the internet must be advertised in subnet sharing from the data center.
- The **wan1** interface must be hardened.
- Local internet access is not available in this configuration due to WAN hardening.
- Other interfaces available for more WAN links. All pass-through traffic goes out **wan0**.
- Multicast traffic will be dropped in this configuration and routing will not work between the routers and switch.

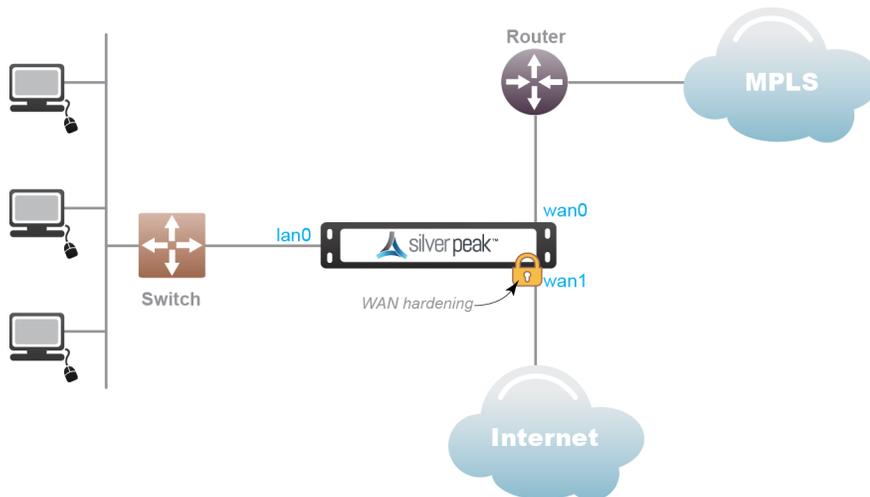


Figure 13.

Figure 14. In-Line Router Mode (Router + Direct Internet)

In-Line Router Mode (Single Direct Internet)

- A default route to the internet must be advertised in subnet sharing from the data center. The **wan0** interface must be hardened.
- Local internet access is not available in this configuration due to WAN hardening.
- Other interfaces are available for more WAN links.



Figure 15.

Figure 16. In-Line Router Mode (Single Direct Internet)

In-Line Router Mode (Dual Direct Internet)

- A default route to the internet must be advertised in subnet sharing from the data center. The **wan0** and **wan1** interfaces must be hardened.
- Local internet access is not available in this configuration due to WAN hardening.
- Other interfaces are available for more WAN links.

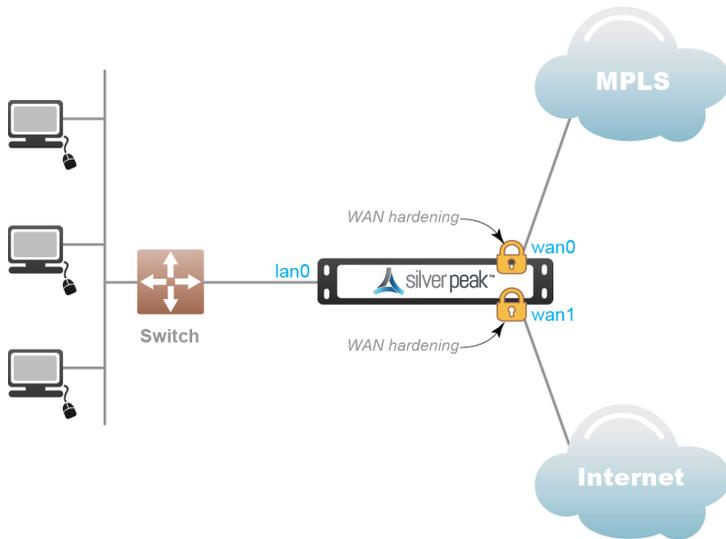


Figure 17. In-Line Router Mode (Dual Direct Internet)

In-Line Router Mode (Dual MPLS)

- A default route to the internet must be advertised in subnet sharing from the data center. Other interfaces are available for more WAN links.
- Multicast traffic will be dropped in this configuration and routing will not work between the routers and switch.
- All pass-through traffic goes out **wan0**.

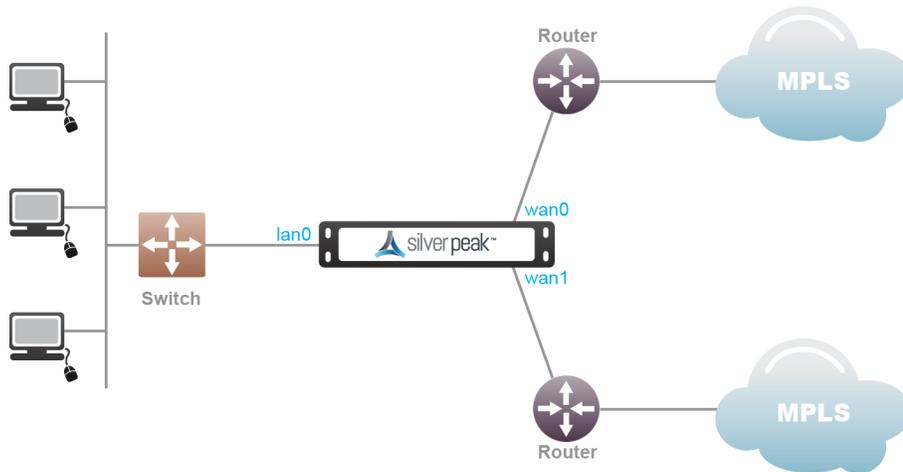


Figure 18.

Figure 19. In-Line Router Mode (Dual MPLS)

Bridge Mode (Router + Direct Internet)

- **lan1** must be left disconnected.
- Propagate link down must be disabled. **wan1** must be hardened.
- Local internet access is disabled when you enable WAN hardening.
- Other interfaces are not available for more WAN links.

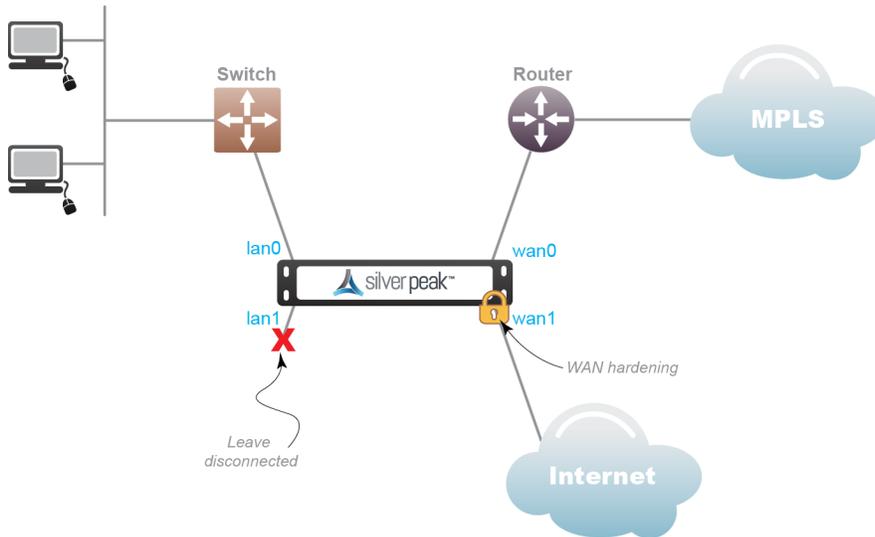


Figure 20. Bridge Mode (Router + Direct Internet)

Bridge Mode (Router + Firewall)

- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- The NAT option must be enabled on the **wan1** interface.
- Local internet access is disabled if you enable WAN hardening .
- Maximum two links allowed in this configuration.

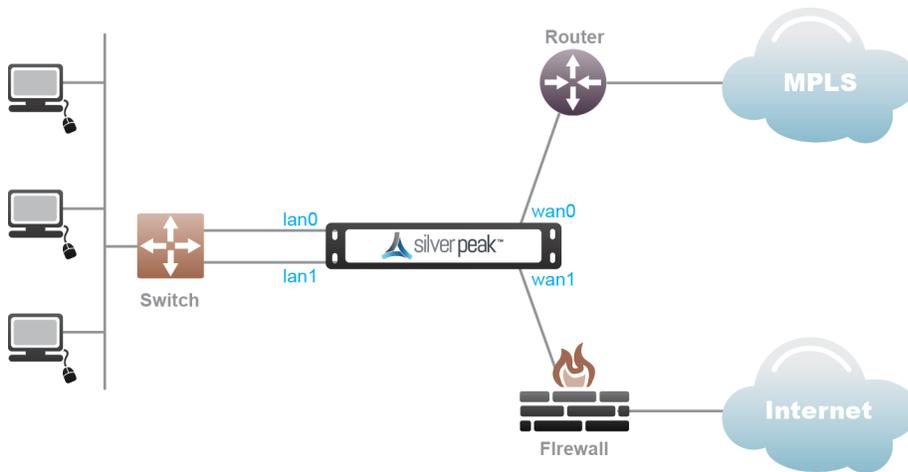


Figure 21. Bridge Mode (Router + Firewall)

Bridge Mode (Dual MPLS)

- A default route must be advertised in subnet sharing from the data center for Internet.
- Only two WAN links are supported - **wan0/wan1**.

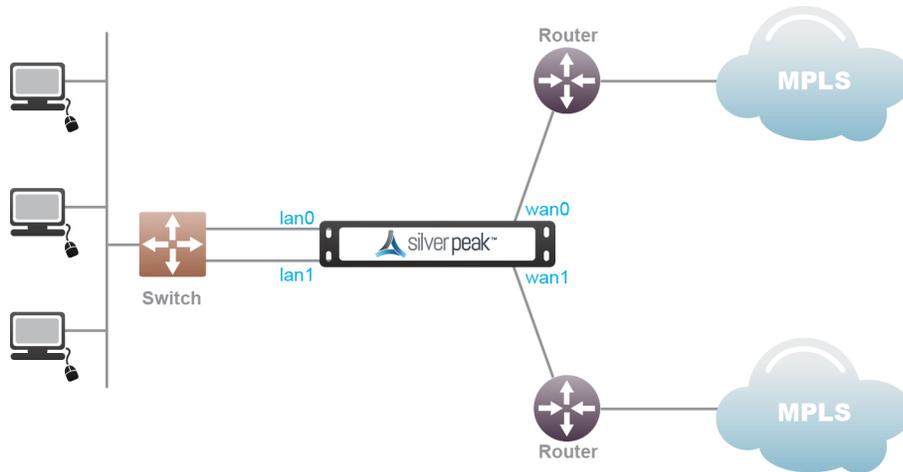


Figure 22. Bridge Mode (Dual MPLS)

Router Mode MPLS + Internet

- A default route must be advertised in subnet sharing from the data center for Internet. Only two WAN links are supported - **wan0/wan1**.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- The NAT option must be enabled on the **wan1** interface.
- Local internet access is disabled if you enable WAN hardening .

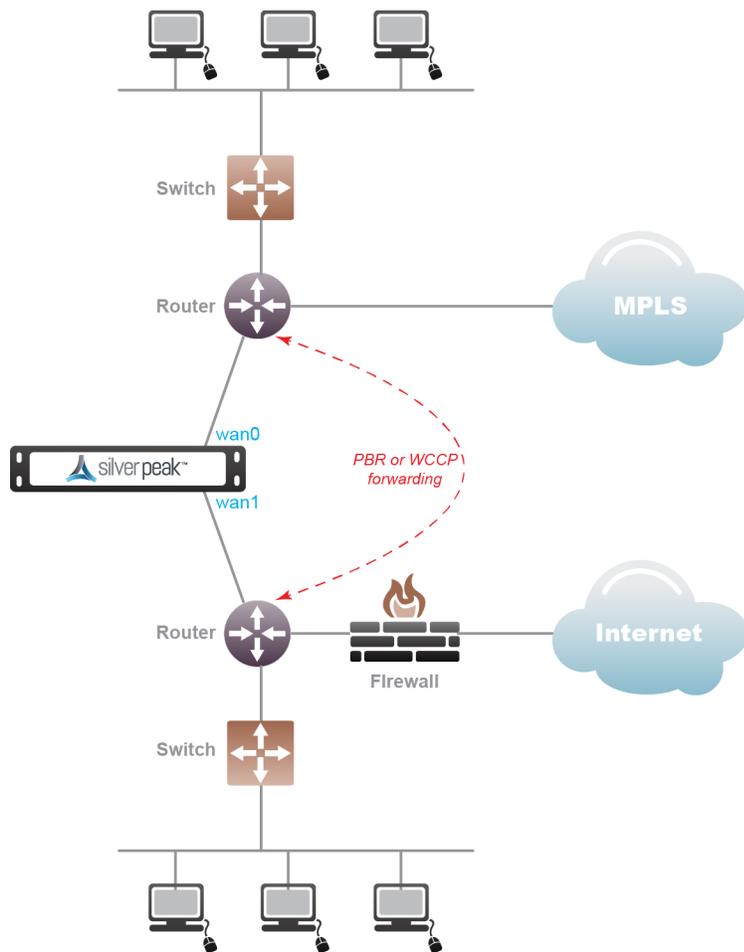


Figure 23. Router Mode MPLS + Internet

Router Mode HA (MPLS + MPLS)

- A default route must be advertised in subnet sharing from the data center for Internet. Only two WAN links are supported - **wan0/wan1**.
- Active/Standby is suggested.

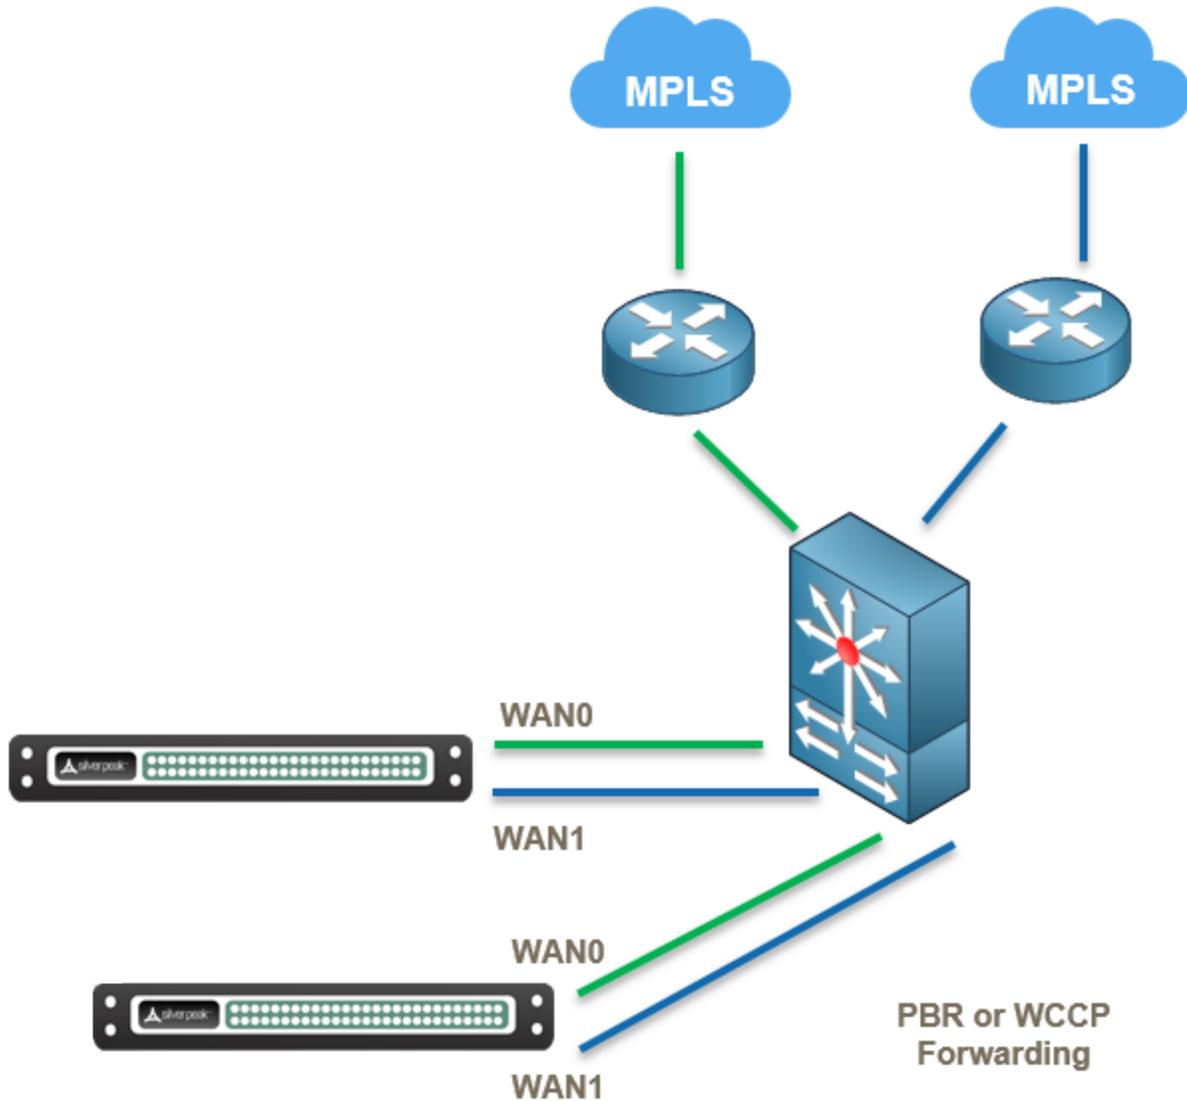


Figure 24. Router Mode HA (MPLS + MPLS)

Dual Home Router Mode (MPLS + Internet)

- A default route must be advertised in subnet sharing from the data center for Internet.
- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- Only two WAN links are supported - **wan0/wan1**.
- Active/Standby is suggested.

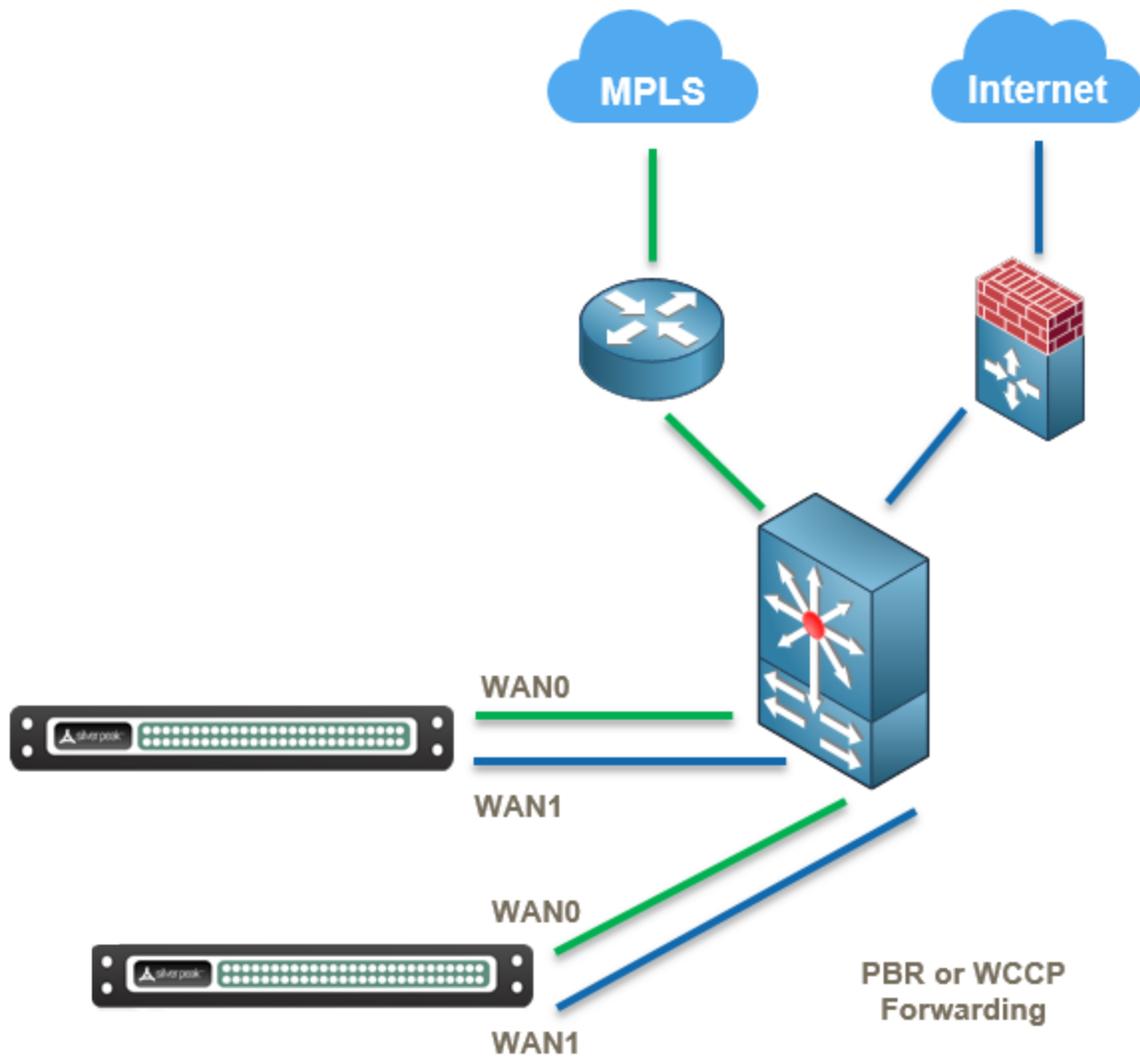


Figure 25. Dual Home Router Mode (MPLS + Internet)

Dual Home Router Mode HA (MPLS + Internet)

- A default route must be advertised in subnet sharing from the data center for Internet.
- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- Only two WAN links are supported - **wan0/wan1**.

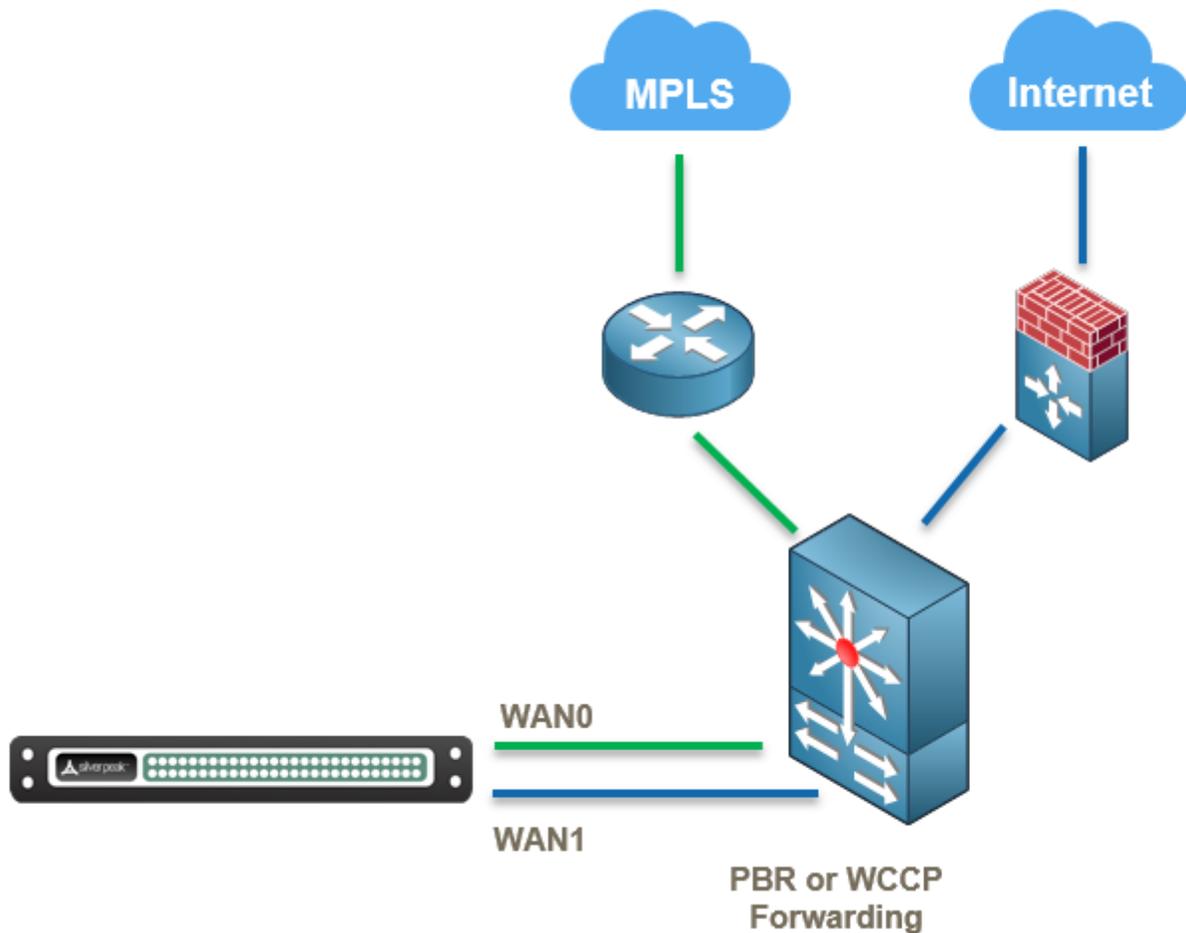


Figure 26. Dual Home Router Mode HA (MPLS + Internet)

Business Intent Overlays (BIO)

A Business Intent Overlay (BIO) specifies how traffic with particular characteristics are handled within the network. Multiple BIOs can be created for different types of traffic. Which traffic matches a particular BIO is determined either by the label on the interface through which it enters the appliance, or by matching traffic to an access list. The BIOs control things like the WAN ports and network types for transmitting traffic, and what to do if the preferred links go down or fail to meet specified performance thresholds. Orchestrator uses BIOs to dynamically build and maintain overlay networks, for example, which sites to build tunnels between and how the network should update the routing of traffic when conditions change.

Within Orchestrator, you can create virtual network overlays to apply business intent to network segments. Provisioning a device is managed by applying profiles.

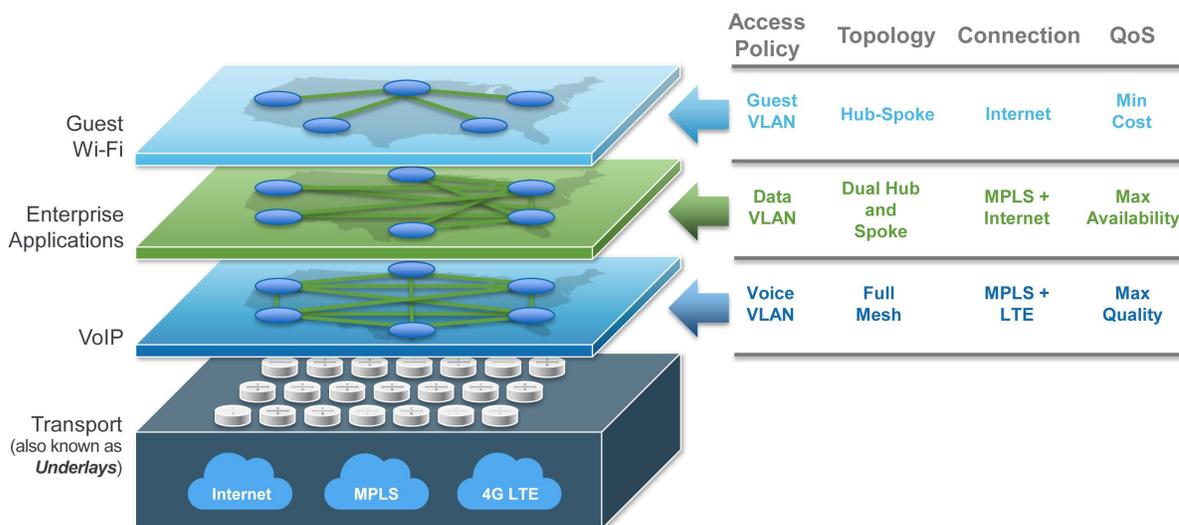


Figure 27. Network with multiple Overlays

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.

- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

Overlays vs Underlays

Overlays are logical tunnels created for different traffic types and policies (such as VoIP).

Underlays are actual IPsec tunnels and physical paths taken (such as MPLS).

Building an Overlay

BIOs use the Labels specified in Deployment Profiles to define how traffic is routed and optimized between sites. These overlays define the paths that traffic will take.

To Create an Overlay

1. From the Configuration tab, click **Business Intent Overlays**. The Business Intent Overlays page appears.
2. Next to the Overlays box, click **+Add**.
The Create Overlay form appears.
3. Enter a descriptive name for the Overlay, such as Default or Voice.
4. Click **Add**.
5. Set the **Traffic Access Policy** to the ACL you created, such as Default_ACL.
6. *Optional.* Set the **Link Brownout Thresholds**.

These thresholds determine when traffic should be limited or even dropped, based on Loss of data, Latency of service, and Jitter.

A value of ZERO means it is not in use.

7. Set **Route Matched Traffic to these WAN Ports** for Primary to both MPLS and Internet.

8. Choose the **Link Bonding Policy** and **Overlay Down Action** you want.
 9. Optional. Check **Boost this Traffic** if this was included in your license.
 10. Click **Save**.
-

To Apply the Overlays

1. From the **Configuration** tab, select **Apply Overlays**.
The Apply Overlays page appears.
 2. From the tree view, select the Appliances to which you want to apply the overlays.
 3. Check the **Add** box next to the overlay you want to use.

 4. Click **Apply**.
-

Tunnels in an Overlay

Overlay tunnels consist of bonded underlay tunnels. Tunnels are created automatically, and you don't need to manually configure them. BIOs use the Labels created in Deployment Profiles to define how traffic is routed and optimized between sites.

Tunnel Reporting & Visibility

The Topology page allows you to quickly see the status of Overlay and Underlay tunnels.

1. Within Orchestrator on the **Topology** page, click a tunnel on the map.

The Tunnels status window appears, showing which tunnels are up or down, active or inactive.

2. From the **Charts** column, click the icon.

A Tunnels Charts appears showing a graphical view of the tunnel traffic.

3. You can adjust the Range, Granularity, Inbound, Outbound, and other views of the tunnel by selecting an option at the top of the window.

Stateful Firewall with Internet Breakout

The Silver Peak SD-WAN supports Stateful Firewall.

This means:

- Less application latency
- Replaces branch firewalls
- Does NOT replace the hub firewall
- Controls access and policy

Stateful Firewall means:

- No need for another branch firewall
- Is used for Internet Breakout, such as sanctioned SaaS and guest WiFi

8.1.4	In this version, you must manually create the Route Map for sanctioned SaaS and guest Wifi.
--------------	---

8.1.5	In this version and beyond, no manual configuration is required.
--------------	--

- Supports NAT/PAT in the branch edge
- Your existing overlays are intact
- When integrated with EdgeConnect, no unauthorized outside traffic can enter the branch. Branch-initiated sessions are allowed, enabling secure Internet Breakout.

Internet Breakout means:

- Is achieved through 'pass-through tunnels'
- Has the ability to choose any egress interface and next-hop
- Automatic load-balance between Internet Breakout tunnels (except with no_encap)
- You can monitor Internet Breakout tunnel statistics

Deployments

■ In-Line Router Mode

- In-line as a next-hop gateway between L2 switch and WAN router or broadband
- internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall

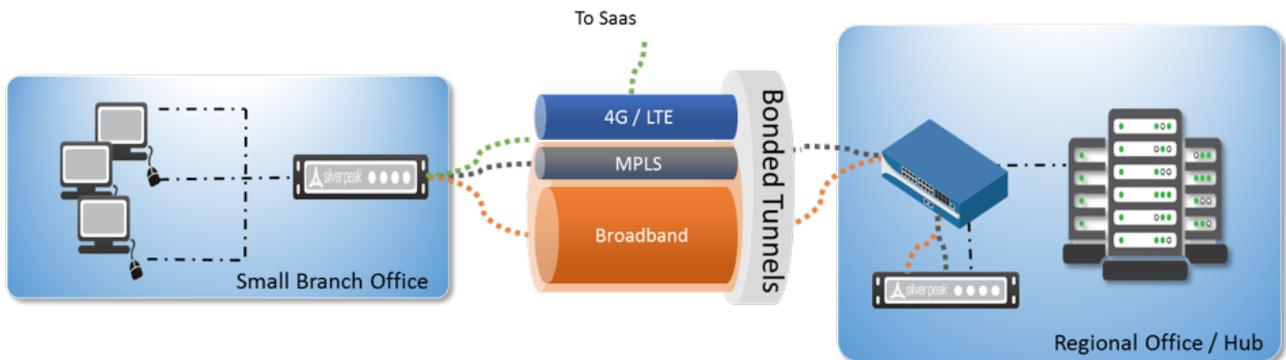
■ In-line Bridge Mode

- In-line between L2 switch and WAN router with fail-to-wire/fail-to-glass, with
- WAN hardening or stateful firewall

■ Out-of-Path (Router) Mode

- Attached to WAN router out-of-path with PBR re-direction, WCCP, and VRRP

Silver Peak Stateful Firewall Use Case



Deploying Stateful Firewall

The screenshot displays the Silver Peak deployment configuration for a Router. It shows LAN and WAN interface settings. The WAN section is annotated with red numbers 1 through 4:

- 1: Firewall dropdown set to **Allow All**.
- 2: Firewall dropdown set to **Harden**.
- 3: Firewall dropdown set to **Stateful**.
- 4: NAT checkbox checked.

Additional configuration details include:

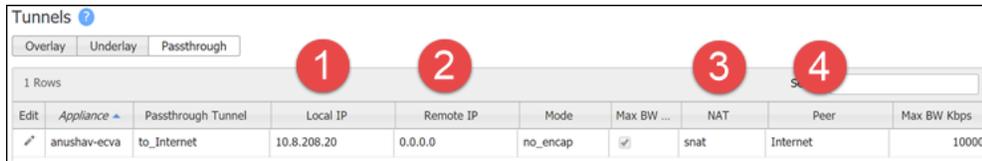
- LAN Interfaces: lan0 (VLAN: +IP, Label: lan, IP/Mask: 10.8.184.20/24, No DHCP).
- WAN Interfaces: wan0 (IP/Mask: 10.8.185.20/24, Label: MPLS, Firewall: Allow All, Interface: wan0, NAT: checked, Bandwidth: 4,000 Kbps, Next Hop: 10.8.185.1).
- WAN Interfaces: wan1 (IP/Mask: 10.8.94.20/24, Label: Internet, Firewall: Harden, Interface: wan1, NAT: checked, Bandwidth: 4,000 Kbps, Next Hop: 10.8.94.1).
- WAN Interfaces: wan2 (IP/Mask: 10.8.208.20/24, Label: Internet2, Firewall: Stateful, Interface: wan2, NAT: checked, Bandwidth: 4,000 Kbps, Next Hop: 10.8.208.254).
- Total Outbound: 10,000 Kbps (≤ 10,000,000 Kbps).
- Total Inbound: (empty) Kbps (Shape Inbound Traffic).
- EdgeConnect Licensing: Plus for > 200 Mbps, Boost 5,000,000 Kbps.

1. MPLS bound traffic can be left as **Allow All**.
2. The Overlays on the Internet link are secured with WAN hardening. No internet breakout is possible.
3. Direct internet breakout traffic is secured with **Stateful Firewall**.
4. The deployment NAT flag indicates to Orchestrator that EdgeConnect is behind a NAT and helps it build tunnels using the public IP. Usually leave this on for internet-bound interface.

8.1.4 In this version, you must manually create the Route Map for sanctioned SaaS and guest Wifi.

8.1.5 In this version and beyond, no manual configuration is required.

Add Source NAT for Internet Breakout Tunnels



Tunnels									
Overlay Underlay Passthrough									
1 Rows									
Edit	Appliance	Passthrough Tunnel	Local IP	Remote IP	Mode	Max BW ...	NAT	Peer	Max BW Kbps
	anushav-ecva	to_Internet	10.8.208.20	0.0.0.0	no_encap	<input checked="" type="checkbox"/>	snat	Internet	10000

1. Internet-bound interface.
2. Default value.
3. Source NAT on for all LAN-side traffic. No other NAT policy configs are needed.
4. Two tunnels to the same peer are used to load-balance flows.

ACL - Breakout Sanctioned SaaS Apps

The following shows settings that define SaaS or other applications that can breakout locally.

Access Lists ?

Sanctioned_Saas ▼ Add ACL Delete ACL Rename ACL

Add Rule Merge Replace

6 Rows

Priority ▲	Match Criteria	Set Actions
		Permit
1030	Domain *conkursolutions.com	permit
1050	IP Intelligence Box	permit
1060	IP Intelligence Dropbox	permit
1070	IP Intelligence Salesforce	permit
1080	IP Intelligence SkypeForBusiness	permit
1090	IP Intelligence Office365	permit

Route Policies

Sample route policies for sanctioned SaaS or guest WiFi setup.

Route Policies ?

7 Rows Search

Edit	Appliance...	Map	Prior...	Match Criteria	Set Actions			Comment
					Destination	Path	Fallback	
	anushav...	map1 (ac...	10	ACL Sanctioned_Saas 1	to_Internet 2		pass-through	
	anushav...	map1 (ac...	20000	ACL Voice_Video	Voice_Video		pass-through	Voice_Video overlay
	anushav...	map1 (ac...	20001	ACL Datacenter_Apps	Datacenter_Apps		pass-through	Datacenter_Apps overlay
	anushav...	map1 (ac...	20002	ACL Scavenger	Scavenger		drop	Scavenger overlay
	anushav...	map1 (ac...	20003	ACL AllWeb	Best_Effort		drop	Best_Effort overlay
	anushav...	map1 (ac...	20004	ACL CatchAll	CatchAll		pass-through	CatchAll overlay
	anushav...	map1 (ac...	65535	Match Everything	auto optimized	default	pass-through	

1. Sample route map to breakout SaaS apps locally.
To breakout guest WiFi, use source subnet or VLAN ID for match criteria.
2. Choose **Internet** breakout tunnel.
3. Business Intent Overlays are intact.



Unity EdgeConnect SD-WAN Family

Unity EdgeConnect Hardware Platforms

	EdgeConnect XS	EdgeConnect S	EdgeConnect M	EdgeConnect L	EdgeConnect XL
Part Identifier	EX-XS	EG-S	EG-M	EG-L	EG-XL
Typical Deployment	Small Branch	Large Branch	Head Office Small Hub	Data Center Large Hub	Data Center Large Hub
Typical WAN Bandwidth	2 - 200 Mbps	10 - 1000 Mbps	50 - 2000 Mbps	1 - 5 Gbps	2 - 10 Gbps
Simultaneous Connections	256,000	256,000	2,000,000	2,000,000	2,000,000
Recommend Boost up to	50 Mbps	200 Mbps	500 Mbps	1 Gbps	5 Gbps
Redundancy / FRUs	No	No	Power and SSD	Power and SSD	Power and SSD
Datapath Interfaces	4 x RJ45 10 / 100 / 1000	6 x RJ45 2 x 1/10G Fiber Option	4 x RJ45 2 x 1/10G Fiber	4 x RJ45 2 x 1/10G Fiber	4 x 1/10G Fiber

Unity EdgeConnect Technical Support

Term	Support is included as part of the EdgeConnect Base subscription license
Web-based Support Portal	Unlimited access 24 / 7 / 365 includes software downloads, technical documentation, and online knowledge base
Software Updates	Major and minor features releases; maintenance releases
Technical Support	24 / 7 / 365 Phone / E-mail / Web
Response Time	2 Hours
Extended Warranty	EdgeConnect hardware purchase options include a 1, 3 or 5-year warranty. Advanced replacement hardware ships the same business day via Priority Overnight Shipment if submitted and verified by 12:00PM local time of the supporting depot.

Flexible Deployment Models

- EdgeConnect Virtual (EC-V) – Download and install EdgeConnect from anywhere in the world. The software runs on all common hypervisors, including VMware vSphere, Microsoft Hyper-V, Citrix XenServer, and KVM.
- EdgeConnect Physical (EC) – For enterprises that are not virtualized in the branch, choose one-of-five EdgeConnect hardware appliance models for plug-and-play deployment.

Subscription Licensing

Subscription licensing for Unity EdgeConnect, which includes Unity Orchestrator, begins at \$199 per-site, per-month. An additional Plus license is required for sites requiring more than 200 megabits-per-second (Mbps) throughput. Unity Boost is an optional performance pack that may be ordered on-demand and is \$5 per-Mbps, per-month. An optional SaaS Optimization license provides internet mapping of optimal egress points to SaaS services.



Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050



Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800



Online

Email: info@silver-peak.com
Website: www.silver-peak.com

© Silver Peak Systems, Inc. All rights reserved. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.



Unity EdgeConnect XL

The Unity EdgeConnect XL is a 1U rack-mountable platform that serves to build an SD-WAN fabric using Zero Touch Provisioning. It supports MPLS, 4G/LTE, and Internet-based hybrid WAN data paths and a control plane that is automated and secured by the Unity Orchestrator software to provide policy-based virtual network segmentation and acceleration of on-premise and SaaS cloud applications. Key fea-

tures of the EdgeConnect XL include high availability through power supply and storage redundancy, available/optional high performance PCIe-based flash network memory for application acceleration, and quad 1/10 Gbps Short Reach (SR) fiber optical ports, making it an ideal platform for data center and large hub environments with typical WAN bandwidth from 2 to 10 Gbps.

Capacity	
WAN Bandwidth Capacity (all features + encryption)	2 to 10 Gbps
Simultaneous Connections	2,000,000
Resiliency	
Power Supplies	Dual redundant
Network	1+1 and N+1 with VRRP or WCCP
RAM	ECC memory
MTBF	109,300 hours (12.5 years) (XL) 89,700 hours (10.2 years) (XL-NM)
Security	
Disk Encryption	128-bit AES
Network Encryption	IPsec (256-bit AES)
Connectivity	
LAN/WAN Ethernet	4 x 1/10 Gb fiber Ethernet data ports (2 LAN / 2 WAN) as per IEEE 802.3z, 802.3ae, 802.1Q
Management	2 x 10/100/1000; Console port

Deployment	
In-Line Router Mode	In-line as next-hop gateway between L2 switch and WAN router or broadband internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall
In-line Bridge Mode	In-line between L2 switch and WAN router with fail-to-glass, with WAN hardening or stateful firewall
Out-of-Path (Router) Mode	Attached to WAN router out-of-path with PBR redirection, WCCP, and VRRP
Power	
Requirement	474W / 1617 BTU (XL) 537W / 1832 BTU (XL-NM)
Power Adapter	1+1 Redundant
Dimensions	
Height	1.69 in (42.8 mm) / 1U
Width	17.34 in (440.5 mm)
Depth	28.57 in (725.8 mm)
Weight	36.0 lbs (16.3 kg)
Regulatory	
EMC	FCC Part 15 Class A, EN 55022 Class A, VCCI Class A, EN 61000-3-2/3-3, EN 55024
Safety	UL/cUL/60950, EN 60950
Environmental	
Temperature (Operating)	10° C to 35° C (50° F to 95° F)
Temperature (Storage)	-40° C to 65° C (-40° F to 149° F)
Humidity (Operating)	10% to 80% relative humidity, non-condensing
Humidity (Storage)	8% to 95% relative humidity, non-condensing
Altitude (Operating)	Up to 10,000 ft (3,048 m)
Altitude (Storage)	Up to 40,000 ft (12,192 m)
Management	
CLI	Full-featured CLI available over DB-9 console port or SSH
GUI	Unity Orchestrator provides centralized configuration, monitoring and management of multiple EdgeConnect appliances
SNMP	SNMPv2c, SNMPv3
Secure Access	SSH and HTTPS
Logging	Syslog with configurable levels. Email alerts
Authentication	Local database, RADIUS, TACACS+
Statistics	Graphing and monitoring, real-time and historical

Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050

Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800

Online

Email: info@silver-peak.com
Website: www.silver-peak.com



Unity EdgeConnect L

The Unity EdgeConnect L is a 1U rack-mountable platform that serves to build an SD-WAN fabric using Zero Touch Provisioning. It supports MPLS, 4G/LTE, and Internet-based hybrid WAN data paths and a control plane that is automated and secured by the Unity Orchestrator software to provide policy-based virtual network segmentation and acceleration of on-premise and SaaS cloud applications. Key fea-

tures of the EdgeConnect L include high availability through power supply and storage redundancy, solid state drive-based network memory (all the appliances use SSDs) for application acceleration, and dual 1/10 Gbps Short Reach (SR) fiber optical ports, making it an ideal platform for data center and large hub environments with typical WAN bandwidth from 1 to 5 Gbps.

Capacity	
WAN Bandwidth Capacity (all features + encryption)	1 to 5 Gbps
Simultaneous Connections	2,000,000
Resiliency	
Power Supplies	Dual redundant
Network	1+1 and N+1 with VRRP or WCCP
RAM	ECC memory
MTBF	113,300 hours (12.9 years) (L) 84,600 hours (9.7 years) (L-NM)
Security	
Disk Encryption	128-bit AES
Network Encryption	IPsec (256-bit AES)
Connectivity	
LAN/WAN Ethernet	4 x 10/100/1000 LAN/WAN (2 LAN/ 2 WAN) as per IEEE 802.3i, 802.3u, 802.3ab, 802.1Q 2 x 1/10 Gbps fiber Ethernet data ports (1 LAN /1 WAN) as per IEEE 802.3z, 802.3ae, 802.1Q
Management	2 x 10/100/1000; Console port

Deployment	
In-Line Router Mode	In-line as next-hop gateway between L2 switch and WAN router or broadband internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall
In-line Bridge Mode	In-line between L2 switch and WAN router with fail-to-wire/fail-to-glass, with WAN hardening or stateful firewall
Out-of-Path (Router) Mode	Attached to WAN router out-of-path with PBR redirection, WCCP, and VRRP
Power	
Requirement	401 W / 1368 BTU (L) 440 W / 1501 BTU (L-NM)
Power Adapter	1+1 Redundant
Dimensions	
Height	1.69 in (42.8 mm) / 1U
Width	17.34 in (440.5 mm)
Depth	28.57 in (725.8 mm)
Weight	36.0 lbs (16.3 kg)
Regulatory	
EMC	FCC Part 15 Class A, EN 55022 Class A, VCCI Class A, EN 61000-3-2/3-3, EN 55024
Safety	UL/cUL/60950, EN 60950
Environmental	
Temperature (Operating)	10° C to 35° C (50° F to 95° F)
Temperature (Storage)	-40° C to 65° C (-40° F to 149° F)
Humidity (Operating)	10% to 80% relative humidity, non-condensing
Humidity (Storage)	8% to 95% relative humidity, non-condensing
Altitude (Operating)	Up to 10,000 ft (3,048 m)
Altitude (Storage)	Up to 40,000 ft (12,192 m)
Management	
CLI	Full-featured CLI available over DB-9 console port or SSH
GUI	Unity Orchestrator provides centralized configuration, monitoring and management of multiple EdgeConnect appliances
SNMP	SNMPv2c, SNMPv3
Secure Access	SSH and HTTPS
Logging	Syslog with configurable levels. Email alerts
Authentication	Local database, RADIUS, TACACS+
Statistics	Graphing and monitoring, real-time and historical

Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050

Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800

Online

Email: info@silver-peak.com
Website: www.silver-peak.com



Unity EdgeConnect M

The Unity EdgeConnect M is a 1U rack-mountable platform that serves to build an SD-WAN fabric using Zero Touch Provisioning. It supports MPLS, 4G/LTE, and Internet-based hybrid WAN data paths and a control plane that is automated and secured by the Unity Orchestrator software to provide policy-based virtual network segmentation and acceleration of

on-premise and SaaS cloud applications. Key features of the EdgeConnect M include high-availability through power supply and storage redundancy, and dual 1/10 Gbps Short Reach (SR) fiber optical ports, making it an ideal platform for head office and small hub environments with typical WAN bandwidth from 50 Mbps to 2 Gbps.

Capacity	
WAN Bandwidth Capacity (all features + encryption)	50 to 2,000 Mbps
Simultaneous Connections	2,000,000
Resiliency	
Power Supplies	Dual redundant
Network	1+1 and N+1 with VRRP or WCCP
RAM	ECC memory
MTBF	162,700 hours (18.6 years)
Security	
Disk Encryption	128-bit AES
Network Encryption	IPsec (256-bit AES)
Connectivity	
LAN/WAN Ethernet	4 x 10/100/1000 LAN/WAN (2 LAN/ 2 WAN) as per IEEE 802.3i, 802.3u, 802.3ab, 802.1Q 2 x 1/10 Gb fiber Ethernet data ports (1 LAN /1 WAN) as per IEEE 802.3z, 802.3ae, 802.1Q
Management	2 x 10/100/1000; Console port

Deployment	
In-Line Router Mode	In-line as next-hop gateway between L2 switch and WAN router or broadband internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall
In-line Bridge Mode	In-line between L2 switch and WAN router with fail-to-wire/fail-to-glass, with WAN hardening or stateful firewall
Out-of-Path (Router) Mode	Attached to WAN router out-of-path with PBR re-direction, WCCP, and VRRP
Power	
Requirement	150W or less / 512 BTU or less
Power Adapter	1+1 Redundant
Dimensions	
Height	1.69 in (42.8 mm) / 1U
Width	17.1 in (434 mm)
Depth	26.1 in (663 mm) or less
Weight	26.0 lbs (11.8 kg) or less
Regulatory	
EMC	FCC Part 15 Class A, EN 55022 Class A, VCCI Class A, EN 61000-3-2/3-3, EN 55024
Safety	UL/cUL/60950, EN 60950
Environmental	
Temperature (Operating)	10° C to 35° C (50° F to 95° F)
Temperature (Storage)	-40° C to 65° C (-40° F to 149° F)
Humidity (Operating)	10% to 80% relative humidity, non-condensing
Humidity (Storage)	8% to 95% relative humidity, non-condensing
Altitude (Operating)	Up to 10,000 ft (3,048 m)
Altitude (Storage)	Up to 40,000 ft (12,192 m)
Management	
CLI	Full-featured CLI available over DB-9 console port or SSH
GUI	Unity Orchestrator provides centralized configuration, monitoring and management of multiple EdgeConnect appliances
SNMP	SNMPv2c, SNMPv3
Secure Access	SSH and HTTPS
Logging	Syslog with configurable levels. Email alerts
Authentication	Local database, RADIUS, TACACS+
Statistics	Graphing and monitoring, real-time and historical

Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050

Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800

Online

Email: info@silver-peak.com
Website: www.silver-peak.com



Unity EdgeConnect S

The Unity EdgeConnect S is a compact form factor Thin Edge appliance that serves to build an SD-WAN fabric using Zero Touch Provisioning. It supports MPLS, 4G/LTE, and Internet-based hybrid WAN data paths and a control plane that is automated and secured by the Unity Orchestrator software to provide policy-based virtual network segmentation and acceleration of on-premise and SaaS cloud ap-

plications. Key features of the EdgeConnect S include quiet operations and flexible mounting options, making it an ideal platform for large branch and remote office environments with typical WAN bandwidth from 10 Mbps to 1 Gbps. The EC-S can optionally be configured with a dual 1/10 Gbps Short Reach (SR) or Long Reach (LR) fiber optic module.

Capacity	
WAN Bandwidth Capacity (all features + encryption)	10 to 1,000 Mbps
Simultaneous Connections	256,000
Resiliency	
Disk Drives	Two
Power Supplies	Single
Network	1+1 and N+1 with VRRP or WCCP
RAM	ECC memory
MTBF	50,827 hours (5.8 years)
Security	
Disk Encryption	128-bit AES
Network Encryption	IPsec (256-bit AES)
Connectivity	
LAN/WAN Ethernet	6 x 10/100/1000 LAN WAN (3 LAN/3 WAN) as per IEEE 802.3i, 802.3u, 802.3ab, 802.1Q Optical module for 2 x 1/10Gbps Ethernet fiber data ports (1 LAN/1 WAN) as per IEEE 802.3z, 802.3ae, 802.1Q
Management	2 x 10/100/1000; Console port

Deployment	
In-Line Router Mode	In-line as next-hop gateway between L2 switch and WAN router or broadband internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall
In-line Bridge Mode	In-line between L2 switch and WAN router with fail-to-wire (copper only) in case of failure, with WAN hardening or stateful firewall
Out-of-Path (Router) Mode	Attached to WAN router out-of-path with PBR re-direction, WCCP, and VRRP
Power	
Requirement	100-240VAC / 47-63Hz / 100W
Power Supplies	Single
Dimensions	
Height	1.73 in (44 mm)
Width	17 in (431 m)
Depth	12 in (305 mm)
Weight	11 lbs (5 kg)
Regulatory	
EMC	FCC Part 15 Class A, EN 55022 Class A, VCCI Class A, EN 61000-3-2/3-3, EN 55024
Safety	UL/cUL/60950, EN 60950
Environmental	
Temperature (Operating)	10° C to 35° C (50° F to 95° F)
Temperature (Storage)	-40° C to 65° C (-40° F to 149° F)
Humidity (Operating)	8% to 90% relative humidity, non-condensing
Humidity (Storage)	8% to 95% relative humidity, non-condensing
Altitude (Operating)	Up to 10,000 ft (3,048 m)
Altitude (Storage)	Up to 40,000 ft (12,192 m)
Management	
CLI	Full-featured CLI available over DB-9 console port or SSH
GUI	Unity Orchestrator provides centralized configuration, monitoring and management of multiple EdgeConnect appliances
SNMP	SNMPv2c, SNMPv3
Secure Access	SSH and HTTPS
Logging	Syslog with configurable levels. Email alerts
Authentication	Local database, RADIUS, TACACS+
Statistics	Graphing and monitoring, real-time and historical

Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050

Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800

Online

Email: info@silver-peak.com
Website: www.silver-peak.com



Unity EdgeConnect XS

The Unity EdgeConnect XS is a very compact form factor Thin Edge appliance that serves to build an SD-WAN fabric using Zero Touch Provisioning. It supports MPLS, 4G/LTE and Internet-based hybrid WAN data paths and a control plane that is automated and secured by the Unity Orchestrator software to provide policy-based virtual network segmentation and accel-

eration of on-premise and SaaS cloud applications. Key features of the EdgeConnect XS include quiet operations and flexible mounting options, making it an ideal platform for small branch and remote office environments with typical WAN bandwidth from 2 to 200 Mbps.

Capacity	
WAN Bandwidth Capacity (all features + encryption)	2 to 200 Mbps
Simultaneous Connections	256,000
Resiliency	
Disk Drives	Single
Power Supplies	Single
Network	1+1 and N+1 with VRRP or WCCP
RAM	ECC memory
MTBF	62,787 hours (7.16 years)
Security	
Disk Encryption	128-bit AES
Network Encryption	IPsec (256-bit AES)
Connectivity	
LAN/WAN Ethernet	4 x 10/100/1000 LAN/WAN (2 LAN/ 2 WAN) as per IEEE 802.3i, 802.3u, 802.3ab, 802.1Q
Management	2 x 10/100/1000; Console port
Deployment	
In-Line Router Mode	In-line as next-hop gateway between L2 switch and WAN router or broadband internet modem for Zero Touch Provisioning, with WAN hardening or stateful firewall

In-line Bridge Mode	In-line between L2 switch and WAN router with fail-to-wire in case of failure, with WAN hardening or stateful firewall
Out-of-Path (Router) Mode	Attached to WAN router out-of-path with PBR re-direction, WCCP, and VRRP
Power	
Requirement	100-240VAC / 50-60Hz / 23W / 78.5BTU
Power Supplies	Single 23W/78.5
Dimensions	
Height	1.7 in (44 mm)
Width	9.4 in (240 m)
Depth	6.5 in (166 mm)
Weight	3.0 lbs (1.4 kg)
Regulatory	
EMC	FCC Part 15 Class A, EN 55022 Class A, VCCI Class A, EN 61000-3-2/3-3, EN 55024
Safety	UL/cUL/60950, EN 60950
Environmental	
Temperature (Operating)	10° C to 35° C (50° F to 95° F)
Temperature (Storage)	-40° C to 65° C (-40° F to 149° F)
Humidity (Operating)	8% to 90% relative humidity, non-condensing
Humidity (Storage)	8% to 95% relative humidity, non-condensing
Altitude (Operating)	Up to 10,000 ft (3,048 m)
Altitude (Storage)	Up to 40,000 ft (12,192 m)
Management	
CLI	Full-featured CLI available over console port or SSH
GUI	Unity Orchestrator provides centralized configuration, monitoring and management of multiple EdgeConnect appliances
SNMP	SNMPv2c, SNMPv3
Secure Access	SSH and HTTPS
Logging	Syslog with configurable levels. Email alerts
Authentication	Local database, RADIUS, TACACS+
Statistics	Graphing and monitoring, real-time and historical

Company Address

Silver Peak Systems, Inc
2860 De La Cruz Blvd.
Santa Clara, CA 95050

Phone & Fax

Phone: +1 888 598 7325
Local: +1 408 935 1800

Online

Email: info@silver-peak.com
Website: www.silver-peak.com

Add Appliances

After configuring Orchestrator and your overlays, you can deploy your appliances. Orchestrator **MUST** be configured and deployed before you connect any EdgeConnect appliance.

Orchestrator can automatically detect your appliances, or you can deploy them manually. You can find a Quick Start Guide for your specific appliance from our documentation page at [Silver Peak User Documentation](#).

To Add Appliances using EdgeConnect

1. After following the instructions given in the Quick Start Guide, connect the **mgmt0** port to a DHCP capable switch port and power the unit on.

DO NOT connect any LAN or WAN ports until approved, licensed and configured.

2. Log into Orchestrator.

Orchestrator and the appliance both contact the Silver Peak Cloud portal. This might take a couple of minutes.

When successful, the **Appliances Discovered** button appears at the top of the page.

3. Click **Appliances Discovered**.



4. For each appliance you want to manage, click **Approve**.

The **Appliance Setup Wizard** appears.

5. Follow the wizard using the IP addresses you identified in [Deployment Parameters](#).

After the appliance is licensed, approved, and configured, schedule the downtime needed to connect the Silver Peak to the appropriate WAN connection points.

Appliance Provisioning (ZTP)

Unity EdgeConnect enables you to automatically detect physical appliances in your network using Zero Touch Provisioning (ZTP).

For example, when you connect a new appliance to the network:

1. The appliance automatically connects to the Cloud Portal and registers itself.
2. The Cloud Portal tells Orchestrator about the new appliance.
3. Orchestrator authenticates the new appliance to the fabric and adds it. New appliances can be added in seconds.

Enable Subnet Sharing

Using auto subnet sharing is a recommended best practice. If you choose not to use subnet sharing, you must also configure inbound redirection on the WAN router (or L3 switch) to avoid creating asymmetric flows that cannot be accelerated when an appliance is deployed out-of-path.

Subnet information is not shared between appliances until a tunnel comes up between them.

Subnet sharing is enabled through the Orchestrator **Initial Config Wizard** (see [Add Appliances](#)), but no subnet information is actually shared until the tunnels are brought up.

To enable subnet sharing, do the following on each appliance:

1. Within the appliance, go to **Configuration > Subnets**. The Subnets tab appears.
 - Check **Use shared subnet information**.
 - Check **Automatically include local subnets**.
 - Set the **Metric for automatically added subnets**. The default is **50**. A lower metric (such as **40**) has a higher priority.

- Best practice is to set up appliances on the same subnet with different metrics. For example, set the first appliance to 40 and leave the second to the default (50).

2. Click **Apply**.

The subnet table updates to include the local subnet. If it doesn't, try refreshing the page.

3. **Save** the changes.

Quality of Service (QoS)

Generally, BIOS determine to which traffic class a packet is placed. Quality of Service (QoS) settings are used for exceptions to the BIOS configuration.

The Shaper provides a simplified way to globally configure QoS (Quality of Service) on the appliances.

- It shapes traffic by allocating bandwidth as a percentage of the system bandwidth.
- The Shaper's parameters are organized into ten traffic classes. Four traffic classes are preconfigured and named --- realtime, interactive, default, and best effort.
- The system applies these QoS settings globally after compressing (deduplicating) all the tunneled and pass-through-shaped traffic --- shaping it as it exits to the WAN.
- To manage Shaper settings for an appliance's system-level WAN Shaper, use the Shaper Template.

You can set QoS options from a template. The QoS Policy determines how flows are queued and marked. The QoS Policy's SET actions determine two things:

- What traffic class a shaped flow —whether optimized or pass-through—is assigned.
- Whether to trust incoming DSCP markings for LAN QoS and WAN QoS, or to remark them as they leave for the WAN

Use the Shaper to define, prioritize, and name traffic classes. Then use the QoS Policy to assign packets to traffic classes for processing.

To create a QoS policy

1. From the Templates page, select **QoS Policies**.

The QoS Policy page appears.

2. To create a new map, click **Add Map** and enter a name.

You can use the current or default map, instead.

3. Click **Add Rule**.

A new row appears with pre-filled criteria. Modify as needed.

Priority

- You can create rules with any priority between 1 and 65534.
- Reserve priorities from 1000 to 9999 inclusive for Orchestrator. If using Orchestrator templates to add route map entries, Orchestrator will delete these entries before applying its policies.
- The lowest priority number (such as 1) has first priority. Best practice is to use priority 65534 as your default.
- Best practice is to add priority numbers in increments of 10, leaving room for you to easily insert new rules.

Source:Dest Port

- An IP address can specify a subnet—for example: 10.10.10.0/24.
 - To allow any IP address, use 0.0.0.0/0.
 - Ports are available only for the protocols TCP, UDP, and TCP/UDP.
 - To allow any port, use 0.
4. Click to check the **QoS Policies** check box, then click **Save** at the bottom of the list.
 5. To push your template values to your appliances, check **Shaper** from the template list, then click **Apply Templates** at the bottom of the list.

Dynamic Rate Control (DRC)

When multiple appliances are simultaneously transmitting into a hub, each at maximum tunnel bandwidth, the Hub could be overrun. This could cause congestion and traffic slow-down. Use the Shaper with Dynamic Rate Control (DRC) to prevent this issue.

To enable this DRC

1. From the Templates page, choose **Shaper**. The Shaper page appears.
See [Quality of Service \(QoS\)](#) for information on the Shaper.
2. Click **Inbound**.

Shaper ?

Inbound Outbound Shaper Add Shaper Delete Shaper

Click Add Shaper button to add inbound shapers

Dynamic Rate Control

Enable Dynamic Rate Control

Inbound Bandwidth Limit

Pass-through Shaped Traffic

Max Bandwidth

3. Check **Enable Dynamic Rate Control**. That allows the Hub to regulate the tunnel traffic by lowering each remote appliance Tunnel Max Bandwidth. The smallest possible value is that appliance Tunnel Min(imum) Bandwidth.
Enter a number for the **Inbound Bandwidth Limit**. This caps how much the appliance can receive. Zero (0) means no limit.
4. Check the **Shaper** check box and click **Save** at the bottom of the list.

5. Push your template values to your appliances. Click **Apply Templates** at the bottom of the list.

For more information:

- See the [Unity Orchestrator Operator's Guide](#) for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

Troubleshooting & Testing

Orchestrator automatically sets up your network and connects to your devices. Sometimes, however, you might run into an issue.

Here are some ways to test and troubleshoot your setup.

Verify Appliance Connectivity	125
Tunnels Not Showing on Tunnels Page	127
Verify Traffic	128

Verify Appliance Connectivity

Check the connectivity of each appliance to verify that the cables are working and the IP address is configured correctly.

To verify appliances are connected

1. From the Tree View, right-click the appliance from the list, then choose **Appliance Manager** from the context menu.

Make sure you do NOT have **Block Popups** enabled in your browser.

The Appliance Manager opens in a new window or tab showing graphs for Bandwidth, Top Applications used, Latency, and Loss.

2. From the **Maintenance** tab, choose **Ping/Traceroute**. The Ping/Traceroute page appears.
3. Select **Ping**.
4. Enter the IP address of a remote appliance in the **IP/Hostname** field.
5. In the **Option** field, you can enter the local appliance IP address. By default, Silver Peak uses the **mgmt0** IP address as the source address for a ping. To specify the local device's data path address as the ping's source address, use the **-I** option (that is, uppercase I as in India, not lowercase L).

Ping/Traceroute ?

Network Connectivity

1 Ping Traceroute IP/Hostname 10.0.183.22 **2** remote appliance

Options

3 -I 10.0.183.22 local appliance (optional)

4 Start Clear

Output

```
PING 10.0.183.22 (10.0.183.22) from 10.0.183.21 : 56(84) bytes of data.
64 bytes from 10.0.183.22: icmp_seq=1 ttl=64 time=0.664 ms
64 bytes from 10.0.183.22: icmp_seq=2 ttl=64 time=0.100 ms
64 bytes from 10.0.183.22: icmp_seq=3 ttl=64 time=0.100 ms
64 bytes from 10.0.183.22: icmp_seq=4 ttl=64 time=0.074 ms
64 bytes from 10.0.183.22: icmp_seq=5 ttl=64 time=0.095 ms

--- 10.0.183.22 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms
rtt min/avg/max/mdev = 0.074/0.206/0.664/0.229 ms
```

5 results

6. Click **Start**.

Run the test as long as you want, then click **Stop**.

The results appear at the bottom of the **Output** field.



TIP: Before putting a bridge mode appliance into production, best practice is to test the connectivity with the appliance in bypass mode to make sure the network still functions if the Silver Peak device fails to wire.

Tunnels Not Showing on Tunnels Page

Are you using Business Intent Overlays and not seeing the tunnels you expect on the Tunnels page?

- Have you created and applied the Overlay to all the appliances on which you're expecting tunnels to be built?

Verify this in the **Apply Overlays** tab.

- Are the appliances on which you're expecting the Overlays to be built using Release 8.0 or later?

View the active software releases on **Maintenance > System Information**.

- Do you have at least one WAN Label selected as a Primary port in the Overlay Policy?

Verify this in the **Business Intent Overlay** tab, in the **Route Matched Traffic to these WAN Ports** section.

- Are the same WAN labels selected in the Overlay assigned to the WAN interfaces on the appliances?

Verify that at least one of the Primary Labels selected in the Business Intent Overlay is identical to a Label assigned on the appliance's Deployment page. Tunnels are built between matching Labels on all appliances participating in the overlay.

- Do any two or more appliances have the same Site Name?

We only assign the same Site Name if we don't want those appliances to connect directly. To view the list of Site Names, go to the **Configuration > Tunnels** tab and click **Sites** at the top.

Verify Traffic

Subnet sharing enables Silver Peak devices that are connected by tunnels to automatically share subnet information and direct all IP traffic to the appropriate destinations.

To verify traffic

1. Verify that each appliance is learning subnets from the other appliance.
 - At each appliance, go to **Configuration > Subnets**, then verify that local subnets are being advertised to peers.
 - Verify that the subnet table lists subnets learned from the remote appliance.

The local appliance uses this learned subnet information. When auto optimization is enabled (this is the default Route Policy), LAN-to-WAN flows are examined for the destination address. If the destination address matches a subnet learned by the local appliance, the flow is routed into the tunnel that terminates at the Silver Peak advertising the subnet.

Subnets ?

Use shared subnet information
 Automatically include local subnets
 Metric for automatically added subnets:

Show Search

Subnet/Mask ▲	Metric	Is Local	Advertise to Peers	Type	Learned from Peer
10.110.11.0/24	50	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer	10.110.11.100
10.110.21.0/24	50	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Auto (added by system)	

Showing 1 to 2 of 2 entries

2. Verify that traffic is being optimized.
 - Bring up a connection between two devices on the end subnets—in this case, hosts on the 10.110.21.0 and 10.110.11.0 subnets. This could be as simple as pinging between them.

```
Administrator: Command Prompt
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping 10.110.21.11

Pinging 10.110.21.11 with 32 bytes of data:
Reply from 10.110.21.11: bytes=32 time=55ms TTL=128
Reply from 10.110.21.11: bytes=32 time=1ms TTL=128
Reply from 10.110.21.11: bytes=32 time<1ms TTL=128
Reply from 10.110.21.11: bytes=32 time=1ms TTL=128

Ping statistics for 10.110.21.11:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 55ms, Average = 14ms

C:\Users\Administrator>_
```

For continuous pinging, use **ping 4**.

- While the ping is running, go to **Monitoring > Current Flows**.

You should see the flow between the two end devices. To refresh the screen, click **Apply**.

When flows stop, they quickly age out of the table. So when the pinging stops, the flow soon disappears.

3. Verify connectivity for pass-through traffic.

As a best practice, always verify connectivity for all devices in the network. For example, if you've configured a route policy to cause certain traffic from certain devices to be handled as pass-through or pass-through unshaped, you should also verify connectivity for these devices.

4. Test network connectivity by using your applications, such as a CIFS mount or an FTP transfer.

Videos

- The *SD-WAN Lightboard Series* includes:
 - - How to Migrate from Legacy Routers to SD-WAN
 - How to Build a Thin Branch with SD-WAN
 - Secure Applications No Matter Where They Reside with SD-WAN
 - How First-packet iQ Application Classification Enables SD-WAN Internet Breakout
 - Assure Predictable Application Performance Over Any Transport with SD-WAN
- *Silver Peak Internet Breakout Automation*
- *SD-WAN in Action*