



Silver Peak

Optimized Network Guide

2018

Copyright and Trademarks

Silver Peak Optimized Network Guide

Date: February 2018

Copyright © 2018 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 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
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, send an e-mail to techpubs@silver-peak.com.
- If you have comments or feedback about the interface, send an e-mail to usability@silver-peak.com.

Related Documentation

- **Release Notes** provide information on new software features, system bugs, and software compatibility.
- All user documentation is available at <http://www.silver-peak.com>.

Getting Started

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

1. Silver Peak appliances must be visible to traffic that requires optimization. To be optimized, all traffic must flow through the appliances. Do this in one of these ways:
 - **In-line Router** mode (in-path) - Silver Peak Appliances are deployed in-line between the LAN infrastructure and the WAN router, becoming the next hop towards the WAN edge. Appliance LAN/WAN interfaces can be deployed physically in-line, or via the use of VLANs on a dot1q trunk. This mode gives you the most flexibility in designing your network. See [In-line Router Mode](#).
 - **In-line Router** mode (out-of-path) - Silver Peak Appliances are deployed with the LAN side interface connected to an out-of-band subnet. A method (such as BGP, VRRP, PBR or WCCP) is used to redirect traffic to the appliance LAN interface. Appliance LAN/WAN interfaces can be physically segmented, or via the use of VLANs on a dot1q trunk. See [Out-of-Path Router Mode](#).
 - **Bridge** mode (in-line) - This mode is the easiest to setup out of the box. Silver Peak appliances are deployed as a “bump in the wire” in between the LAN infrastructure and the WAN router. See [Bridge Mode](#).
 - **Server** mode (replication only) - The management path and the data path both use the same interface and the same IP address. Static routes are used on replication hosts to redirect traffic to the appliance. Although easy to use, this mode is not recommended best practice. See [Server Mode](#).
2. Silver Peak WAN optimization is a “symmetric” solution. To optimize traffic, Silver Peak appliances are required on both ends of the WAN.
3. Silver Peak TCP Acceleration requires that the appliances be visible to both the transmit and receive direction of a TCP flow. If not, the flow is considered “asymmetric” and TCP Acceleration can’t occur, although Network Integrity and Network Memory continue to work.

Prerequisites:

- Install Orchestrator. You will receive a download link from Silver Peak by email.
- Use the [Deployment Checklist](#) and gather your Deployment Parameters.

- You can obtain a pool of licenses for your appliances from the Support Portal at <https://www.silver-peak.com/support/customer-login>. You will receive instructions by email.

Licensing

Log into the Silver Peak Customer Support Portal (<https://www.silver-peak.com/support/customer-login>) to view and manage your licenses.

License & Account Key

When you first deploy 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

License page of the Getting Started Wizard

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

- The License (usually 60 characters) enables you to use Orchestrator and Appliance Manager.

To view existing Licenses

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

The screenshot shows the 'Licenses' page in Orchestrator. 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 has a red 'Not approved' status in the Base column.

Edit	Host Name	Model	Serial No	License Start	Base	Plus	Boost	License End	SaaS
	server2	EC-XS	00-1B-BC-01-1E-36		Not approved				No
	server-oc	VX2000	00-0c-29-b9-00-8f	28-May-15 17:00:00					No
	server-usb	VX1000	00-0C-29-49-0B-05	10-Mar-16 16:00:00					No
	server-usb	VX2000	00-0C-29-0D-19-01	12-Aug-12 17:00:00					No

The Licenses page in Orchestrator

Licensing Physical Appliances

Silver Peak physical appliances have the licenses already installed.

You must manually add the Orchestrator IP address to each appliance, from **Administration > Orchestrator**, and then approve using Orchestrator.

Deployment Parameters

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

Example Network Deployment Parameters

Site Name	Example	Your Site
Hostname	hostname	
Deployment Mode	In-line Router mode	
lan0 - Voice	LAN0_IP	
lan0 VLAN IP	LAN0_VLAN_IP	
lan1 - Data	LAN1_IP	
lan1 VLAN IP	LAN1_VLAN_IP	
wan0 IP (MPLS)	WAN0_IP	
wan0 Next Hop	WAN0_IP_DEFAULT_GW	
lan0 Next Hop	LAN0_IP_DEFAULT_GW	
wan1 IP (Internet)	WAN1_IP	
wan1 Next Hop	WAN1_IP_DEFAULT_GW	
lan1 Next Hop	LAN1_IP_DEFAULT_GW	
wan2 IP (LTE)	WAN2_IP	
wan2 Next Hop	WAN2_IP_DEFAULT_GW	
lan1 Next Hop	LAN1_IP_DEFAULT_GW	

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.
- Border Gateway Protocol (BGP) and High Availability (HA) when using technologies such as WCCP, VRRP, and PBR to forward packets.

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?
- 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 Silver Peak SE for more information.

Orchestrator

- If using Orchestrator, install it and enter the license information.
- Immediately change your user name and password, and keep them in a safe place. Failure to do so could subject you to hacking. **Highly recommended:** enable 2-step authentication.
- 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 interface Labels you have created. 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 \(BIO\)](#).

Appliances

- ❑ Install appliances and get approval from Orchestrator.
- ❑ In the configuration wizard, apply [Deployment Profiles](#), [Business Intent Overlays \(BIO\)](#), subnet sharing, and Configuration Templates.
- ❑ Define exceptions and perform any other individual configurations.

Using Physical and Virtual Appliances

1. Use Orchestrator to configure the management interface, **mgmt0** (required for virtual machines, optional for physical appliances).
2. Configure **mgmt0** with a static IP address.

DHCP will work, but as best practice, you should configure a static IP address. Otherwise, you might lose communication with the machine after an outage, upgrade, or reboot.

Configuring the mgmt0 Interface

The physical (NX) Quick Start Guide explains how to access and configure the **mgmt0** interface. Here is a quick, generic review.

NOTE The **mgmt0** next-hop is to an L3 (not L2) switch.

To configure the mgmt0 interface on a physical (NX) appliance

- Refer to the [NX Series Appliances Quick Start Guide](#).

NTP Setup

Manually configure the date and time of an appliance, or use a Network Time Protocol (NTP) server for automatic updates.

To configure the date and time

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

Date / Time Setting ?

Time Zone

Manual
Configured when the template group is applied

NTP Time Synchronization

Server IP	Version
<i>No Data Available</i>	

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**.

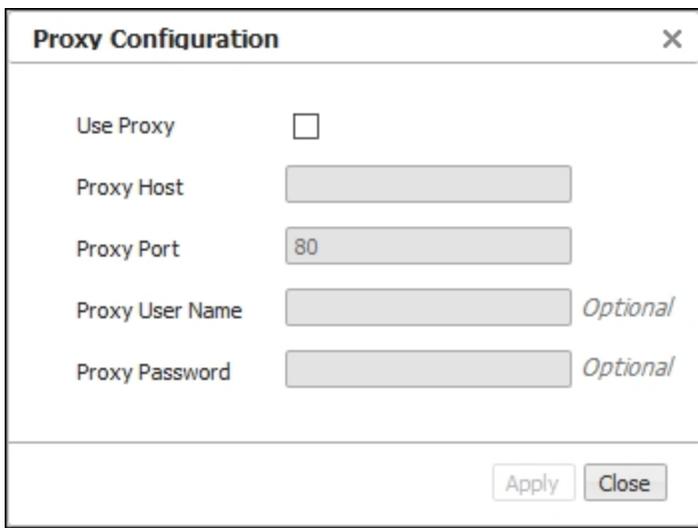
Using a Proxy Server

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.



Use Proxy	<input type="checkbox"/>
Proxy Host	<input type="text"/>
Proxy Port	<input type="text" value="80"/>
Proxy User Name	<input type="text"/> <i>Optional</i>
Proxy Password	<input type="text"/> <i>Optional</i>

Apply Close

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

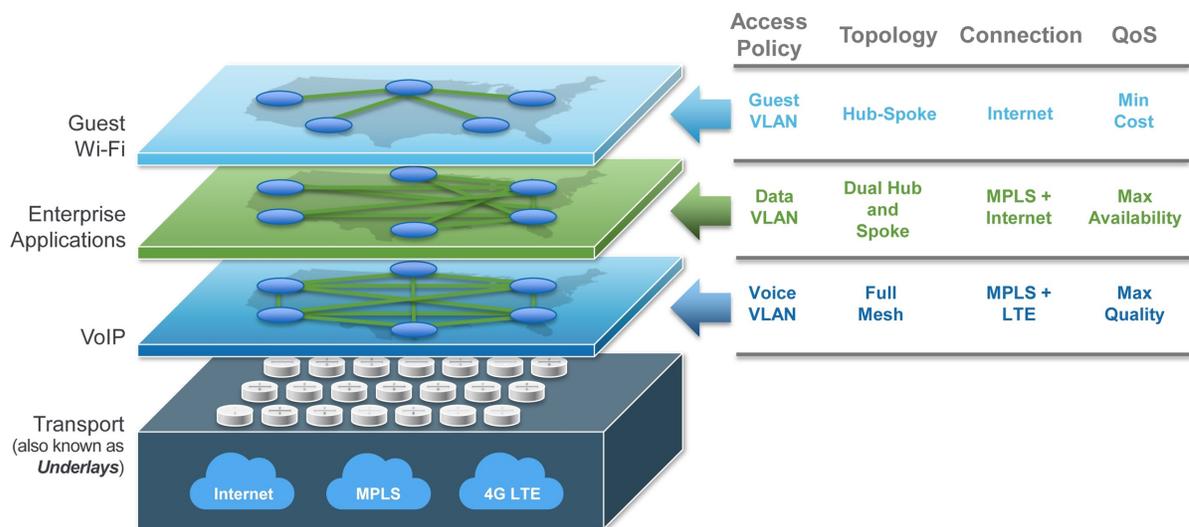
Configuration

Business Intent Overlays (BIO)	16
Overlays vs Underlays	17
Building an Overlay	18
Tunnels in an Overlay	21
Tunnel Reporting & Visibility	22
Template Groups	23
Dynamic Rate Control (DRC)	24
Traffic Redirection	27
Examples of Traffic Redirection	29
When Defaulting to TCP-based or IP-based Auto-Optimization	30
When Specifying a Tunnel	31
High Availability (HA)	32
Asymmetry Mitigation	32
HA Using PBR	33
HA Using WCCP	34
Route Policies	34
Shaper	34
Quality of Service (QoS)	36

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, create virtual network overlays to apply business intent to network segments. Apply profiles to provision a device.



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.

The Business Intent Overlays page automates most network tasks, such as routing, link bonding, traffic shaping, and WAN-OP policies.

For more details on the options on this page, see the Orchestrator Guide.

To Create an Overlay

1. Go to **Configuration > Business Intent Overlays** to open the Business Intent Overlays page.

The screenshot shows the 'Business Intent Overlays' configuration page. The interface includes several sections for configuring network overlays:

- Match Traffic:** Set to 'AnyTraffic' with an 'ACL' dropdown.
- Topology:** Options for 'Mesh' and 'Hub & Spoke' are available. A 'Select Hubs +Add' button is present. The 'Peer Unavailable Action' is set to 'Pass Through Unshaped'.
- WAN Links & Bonding Policy:**
 - Primary:** Internet (checked), MPLS (checked), LTE (unchecked).
 - Backup:** All unchecked.
 - Cross Connect:** All unchecked.
 - Use Backup Ports on:** 'Blackout' and 'Brownout' buttons.
 - Brownout Thresholds:** Loss (no limit, %), Latency (no limit, ms), Jitter (no limit, ms).
- High Availability, High Quality, High Throughput, High Efficiency:** Four columns of policy options. 'High Quality' is currently selected, showing options like 'Failover <1sec', 'Use Best Quality Path', 'Path Conditioning', and 'BW Efficiency >80%'.
- Internet Traffic:**
 - Preferred Policy Order:** 'Break Out Locally' and 'Backhaul Via Overlay' buttons.
 - Break Out Locally Using These Interfaces:** A diagram showing a 'Drop' point.
 - Primary/Backup:** Internet (checked), MPLS (unchecked), LTE (unchecked).
- Traffic Management:**
 - Traffic Class:** '1 (default)' dropdown.
 - LAN DSCP:** 'trust-lan' dropdown.
 - WAN DSCP:** 'trust-lan' dropdown.
 - Boost License:** 'Boost this Traffic' checkbox (unchecked).

At the bottom left, there are 'Save All', 'Save As', and 'Cancel' buttons.

2. Next to the Overlays box, click **+Add**.

The Create Overlay form appears.

3. Enter a descriptive name for the Overlay, such as *Default*, *Voice*, or *AnyTraffic*.
4. Click **Add**.
5. Set the **Match Traffic** overlay to the ACL you created, such as **Any Traffic**.

Match Traffic: AnyTraffic | ACL

6. Choose the **Topology** and **Peer Unavailable Action** you want.

Topology: Mesh | Hub & Spoke | Select Hubs +Add: SPITDC02, SPENGDC01 | Peer Unavailable Action: Pass Through Unshaped

7. *Optional.* Set the **Link Brownout Thresholds**.

Use Backup Ports on: Blackout | Brownout | Brownout Thresholds: Loss (no limit %), Latency (no limit ms), Jitter (no limit ms)

These thresholds determine when traffic should be statefully moved from primary to backup links, based on packet loss, latency and jitter.

A value of ZERO (0) means it is not in use.

8. Choose the **Link Bonding Policy** you want.

WAN Links & Bonding Policy	<table border="1"> <tr> <td></td> <td><i>Primary</i></td> <td><i>Backup</i></td> <td><i>Cross Connect</i></td> </tr> <tr> <td>Internet</td> <td><input checked="" type="checkbox"/></td> <td><input type="checkbox"/></td> <td><input type="checkbox"/></td> </tr> <tr> <td>ATT_NOD</td> <td><input checked="" type="checkbox"/></td> <td><input type="checkbox"/></td> <td><input type="checkbox"/></td> </tr> </table>		<i>Primary</i>	<i>Backup</i>	<i>Cross Connect</i>	Internet	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	ATT_NOD	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Use Backup Ports on <input type="button" value="Blackout"/> <input type="button" value="Brownout"/>	Brownout Thresholds Loss <input type="text" value="no limit"/> % Latency <input type="text" value="no limit"/> ms Jitter <input type="text" value="no limit"/> ms
		<i>Primary</i>	<i>Backup</i>	<i>Cross Connect</i>											
Internet	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>												
ATT_NOD	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>												
<table border="1"> <tr> <td> High Availability <ul style="list-style-type: none"> Failover 0sec Use Best Quality Path Path Conditioning BW Efficiency 50% </td> <td style="border: 2px solid blue;"> High Quality <ul style="list-style-type: none"> Failover <1sec Use Best Quality Path Path Conditioning BW Efficiency >80% </td> <td> High Throughput <ul style="list-style-type: none"> Failover <1sec Load Balance Paths Path Conditioning BW Efficiency >80% </td> <td> High Efficiency <ul style="list-style-type: none"> Failover <1sec Load Balance Paths BW Efficiency 100% </td> </tr> </table>	High Availability <ul style="list-style-type: none"> Failover 0sec Use Best Quality Path Path Conditioning BW Efficiency 50% 	High Quality <ul style="list-style-type: none"> Failover <1sec Use Best Quality Path Path Conditioning BW Efficiency >80% 	High Throughput <ul style="list-style-type: none"> Failover <1sec Load Balance Paths Path Conditioning BW Efficiency >80% 	High Efficiency <ul style="list-style-type: none"> Failover <1sec Load Balance Paths BW Efficiency 100% 											
High Availability <ul style="list-style-type: none"> Failover 0sec Use Best Quality Path Path Conditioning BW Efficiency 50% 	High Quality <ul style="list-style-type: none"> Failover <1sec Use Best Quality Path Path Conditioning BW Efficiency >80% 	High Throughput <ul style="list-style-type: none"> Failover <1sec Load Balance Paths Path Conditioning BW Efficiency >80% 	High Efficiency <ul style="list-style-type: none"> Failover <1sec Load Balance Paths BW Efficiency 100% 												

9. Click **Save All**.

NOTE The Internet Traffic section does not apply to WAN Optimization deployments.

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 box under **Add** to use the overlay you want.

Apply Overlays ? Edit Overlays

View Status

Overlay	Add	Remove
RealTime	<input type="checkbox"/>	<input type="checkbox"/>
Interactive	<input type="checkbox"/>	<input type="checkbox"/>
Critical_Apps	<input type="checkbox"/>	<input type="checkbox"/>
SaaS	<input type="checkbox"/>	<input type="checkbox"/>
Default	<input type="checkbox"/>	<input type="checkbox"/>
DMZFabric	<input type="checkbox"/>	<input type="checkbox"/>

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.

Type	Local Appliance	Remote Appliance	Name	Status	Live View	Historical Charts
17/60 Rows, 1 Selected						
Search <input type="text"/>						
Overlay: Default (12 tunnels)						
Overlay: SaaS (12 tunnels)						
Overlay: Interactive (12 tunnels)						
Overlay: RealTime (12 tunnels)						
overlay	SPENGDC02	SPHQ01	to_SPHQ01_RealTime	up - active		
overlay	SPENGDC02	SPHQ02	to_SPHQ02_RealTime	up - active		
underlay	SPENGDC02	SPHQ02	to_SPHQ02_Internet-Internet	up - active		
underlay	SPENGDC02	SPHQ02	to_SPHQ02_ATT_NOD-ATT_NOD	up - active		
underlay	SPENGDC02	SPHQ01	to_SPHQ01_ATT_NOD-ATT_NOD	up - active		
underlay	SPENGDC02	SPHQ01	to_SPHQ01_Internet-Internet	up - active		
overlay	SPHQ01	SPENGDC02	to_SPENGDC02_RealTime	up - active		
underlay	SPHQ01	SPENGDC02	to_SPENGDC02_Internet-Internet	up - active		

2. From the **Charts** column, click the chart icon . **Live View** shows real-time data, and **Historical Charts** shows data from a specified period of time.

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

- **Live View** - Click to view **Bandwidth, Loss, Jitter, and Latency**.
- **Historical Charts** - Adjust **Range, Granularity, Inbound, Outbound,** and other views of the tunnel by choosing an option at the top of the window.

Template Groups

Templates are configuration values that can be applied to your appliances. Template Groups are a collection of templates that can be applied simultaneously. 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.

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. Consult your Silver Peak SE for more information.

To enable 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. Select 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.

Traffic Redirection

To optimize traffic, the appliance must intercept both the inbound and outbound packets for each flow. When you place an appliance out-of-path, you must redirect traffic from the client to the appliance.

There are three methods for redirecting outbound packets from the client to the appliance (known as **LAN-side redirection**, or **outbound redirection**):

- **BGP** (Border Gateway Protocol) — configured on both the router and the Silver Peak appliance. Also use BGP for appliance redundancy and load balancing.
- **PBR** (Policy-Based Routing) — configured on the router. No other special configuration is required on the appliance. This is also known as Filter-Based Forwarding (**FBF**).

Deploying two Silver Peaks at a site, for redundancy, requires VRRP (Virtual Router Redundancy Protocol) and/or IP SLA configured on the router.

- **WCCP** (Web Cache Communication Protocol) — configured on both the router and the Silver Peak appliance. Also use WCCP for redundancy and load balancing.
- **Host routing** — the server/end station has a default or subnet-based static route that points to the Silver Peak appliance as the 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).

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

Always enable subnet-sharing, to ensure that only inbound redirection from the WAN router (also known as WAN-side redirection) is required.

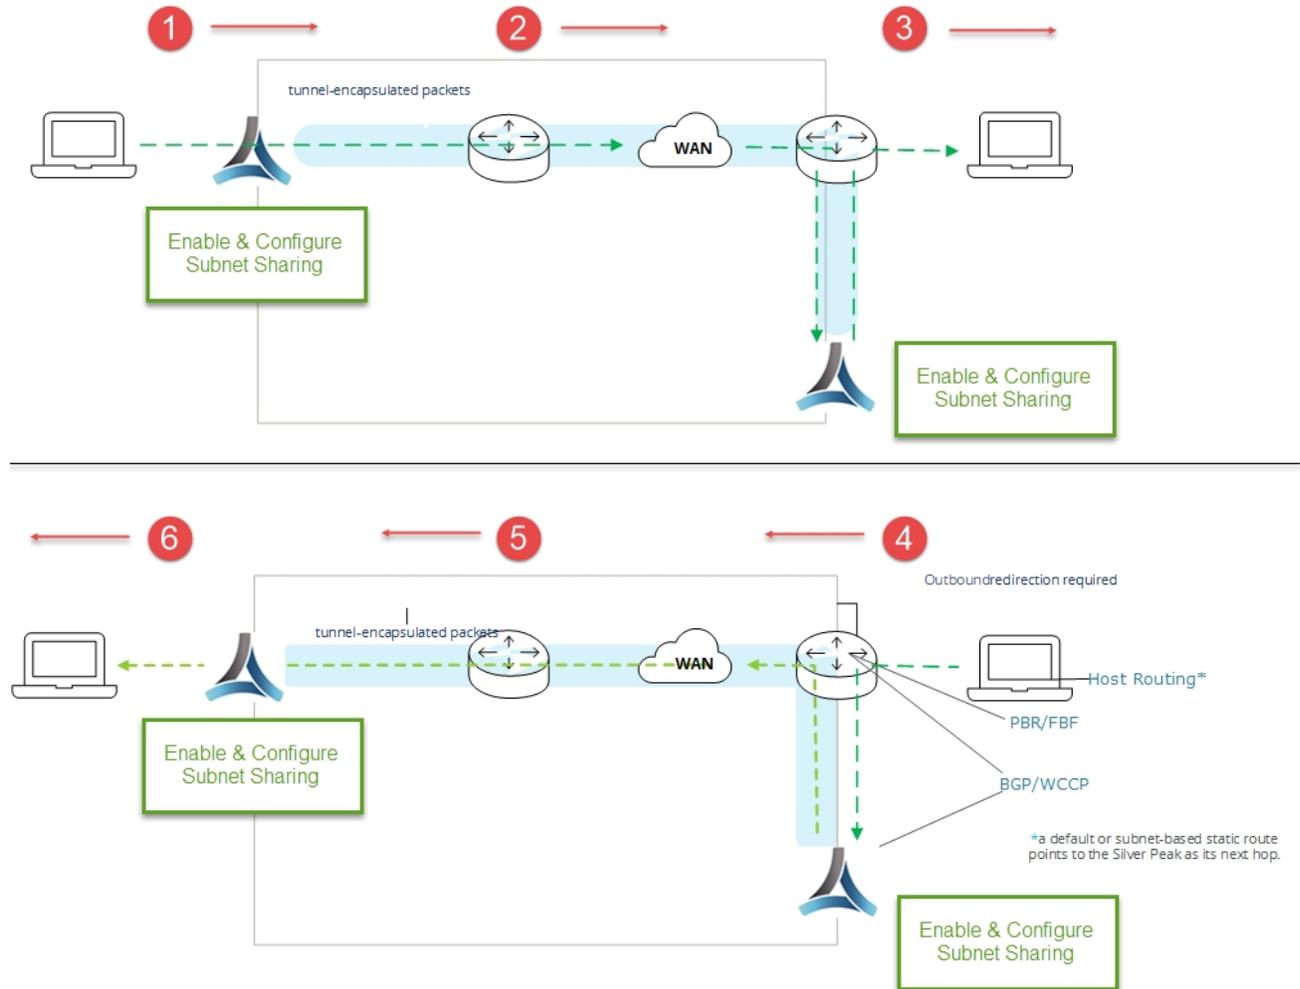
- **Subnet sharing** relies on advertising local subnets between Silver Peak appliances.
- Subnets not directly attached to the Silver Peak might need to be manually added so the local appliance can advertise them to its peers. If those subnets are not reachable via the default LAN-side next-hop router, then you might also need to add a static route to the local Silver Peak, specifying which next-hop router to use to reach a

given subnet. Optionally, BGP to a LAN side router can be used to learn subnets that are not directly connected to the Silver Peak.

- For TCP flows to be optimized, both directions must travel through the same client and server appliances. If the TCP flows are asymmetric—as could occur in a high-availability deployment—you need to configure clusters for flow redirection among local appliances.

Examples of Traffic Redirection

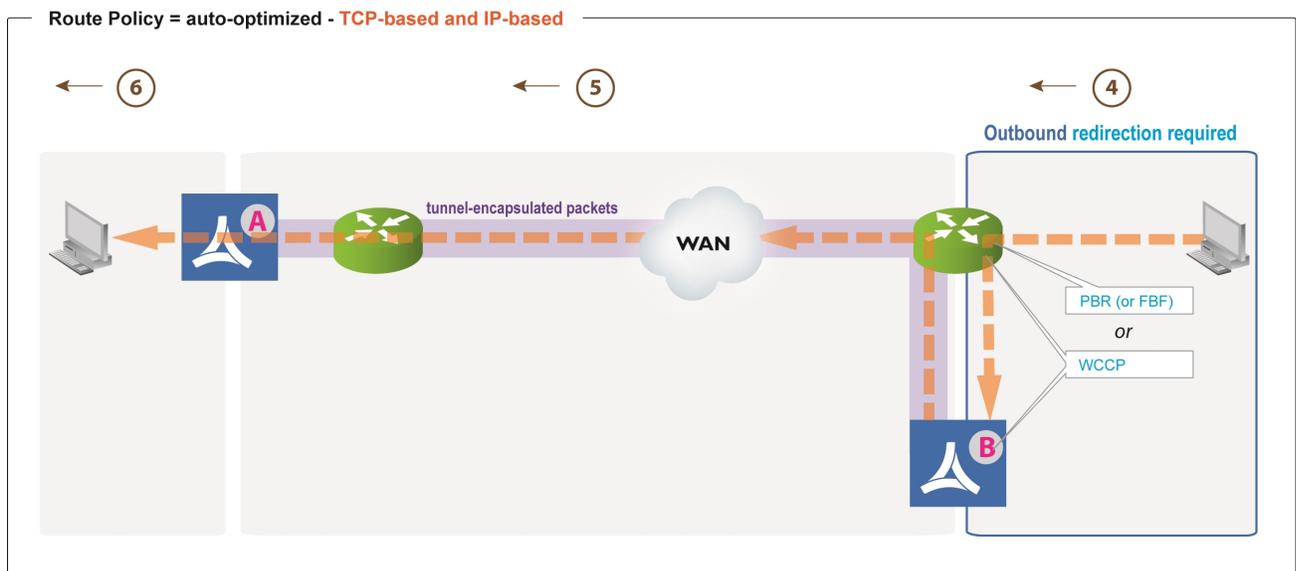
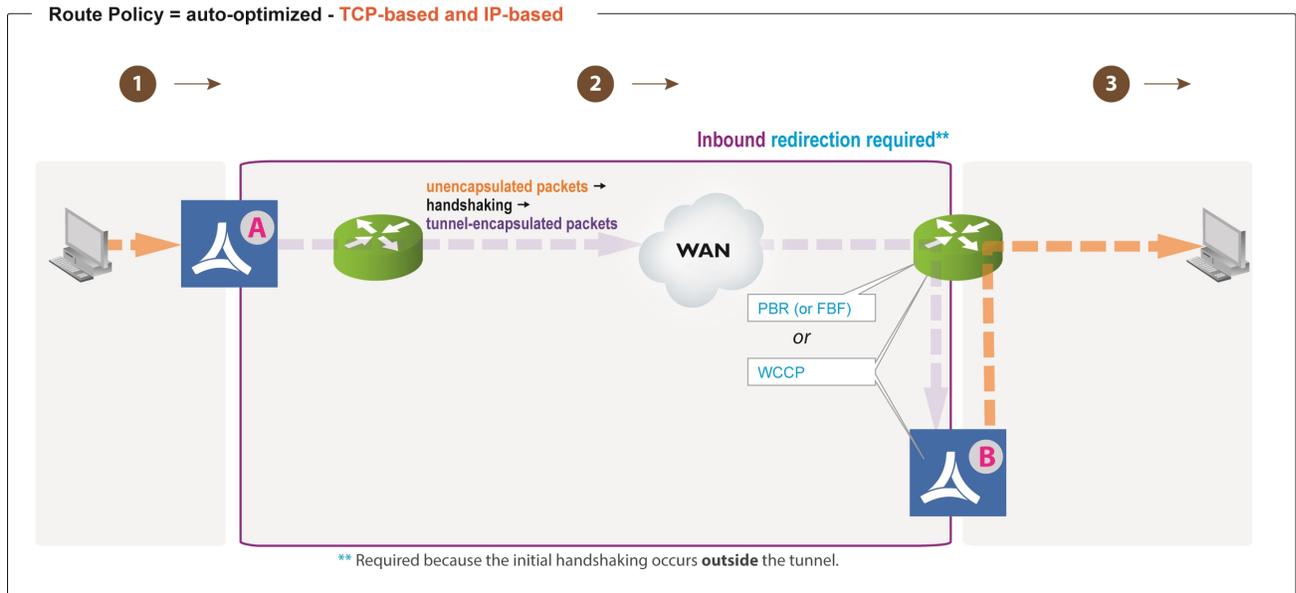
- Enable **subnet sharing** on both the local and remote appliances.
- For outbound redirection to the out-of-path appliance (B), choose from BGP, PBR (or FBF), WCCP, or host routing.
- Host routing only requires configuration on the client — not on the router or appliance.



Route Policy: determined by Business Intent Overlay - subnet sharing

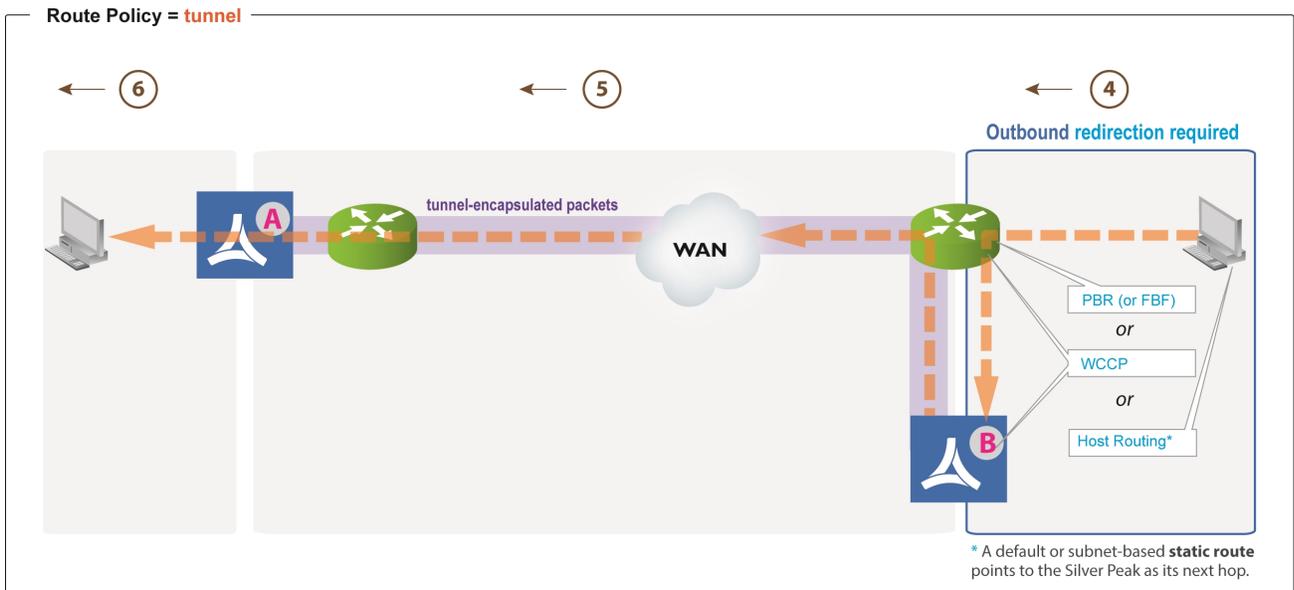
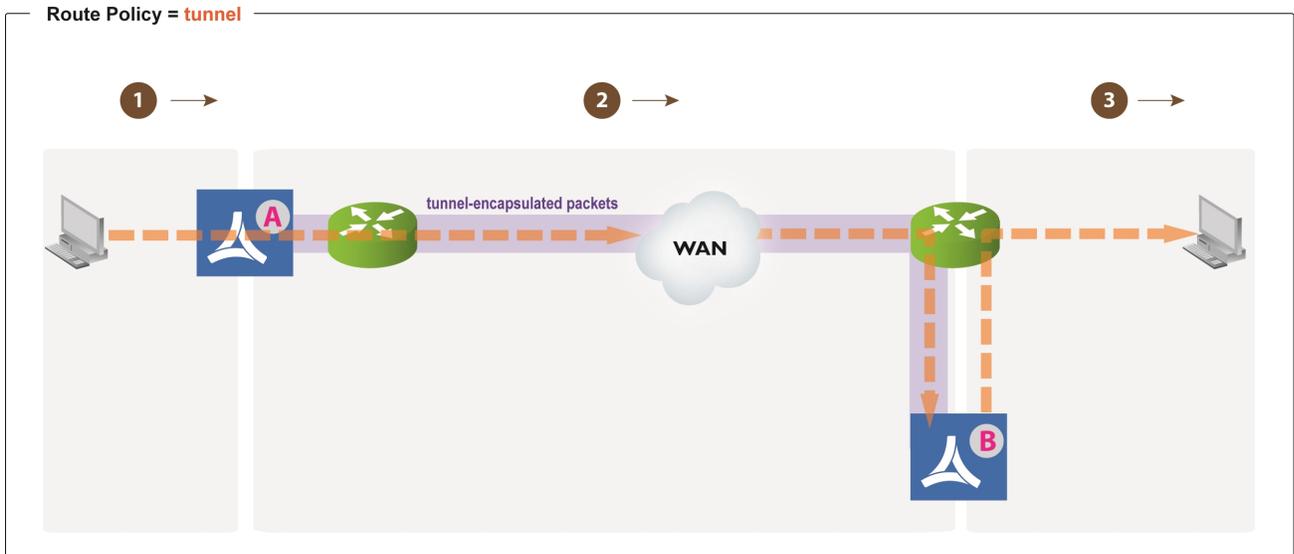
When Defaulting to TCP-based or IP-based Auto-Optimization

- Initial handshaking between appliances happens outside the tunnel, requiring inbound redirection for packet routing.
- For inbound and outbound redirection to the out-of-path appliance (B), choose from PBR (or FBF) or WCCP.



When Specifying a Tunnel

- For outbound redirection to the out-of-path appliance (B), choose from PBR (or FBF), WCCP, or host routing.
- With host routing, the outbound redirection is configured on the client, as opposed to on the router and/or appliance.
- Host routing only requires configuration on the client — not on the router or appliance.



High Availability (HA)

High Availability (HA) means using a pair of Silver Peak appliances to provide continuous acceleration service in case of failure. Each HA pair terminates a single WAN link, while providing automation through BIOs and link resiliency through tunnel bonding.

In High Availability (HA) configurations, the redundant Silver Peak appliances are deployed in router mode, and either BGP, WCCP or PBR redirects flows from the routers to the appliances.

The redundant appliances can be configured **Active/Active** or **Active/Backup**. This is determined by how the BGP, WCCP or PBR redirection is configured on the routers and the appliances.

For our examples, assume that HA is configured in the same location. Refer to the non-redundant appliances as “client-side”.

Asymmetry Mitigation

Flow redirection can prevent TCP asymmetry in high availability environments. For the appliances, this requires configuring HA (or redundant) appliances as peers, and enabling flow redirection. Both tasks are on the **Configuration - Flow Redirection** page.

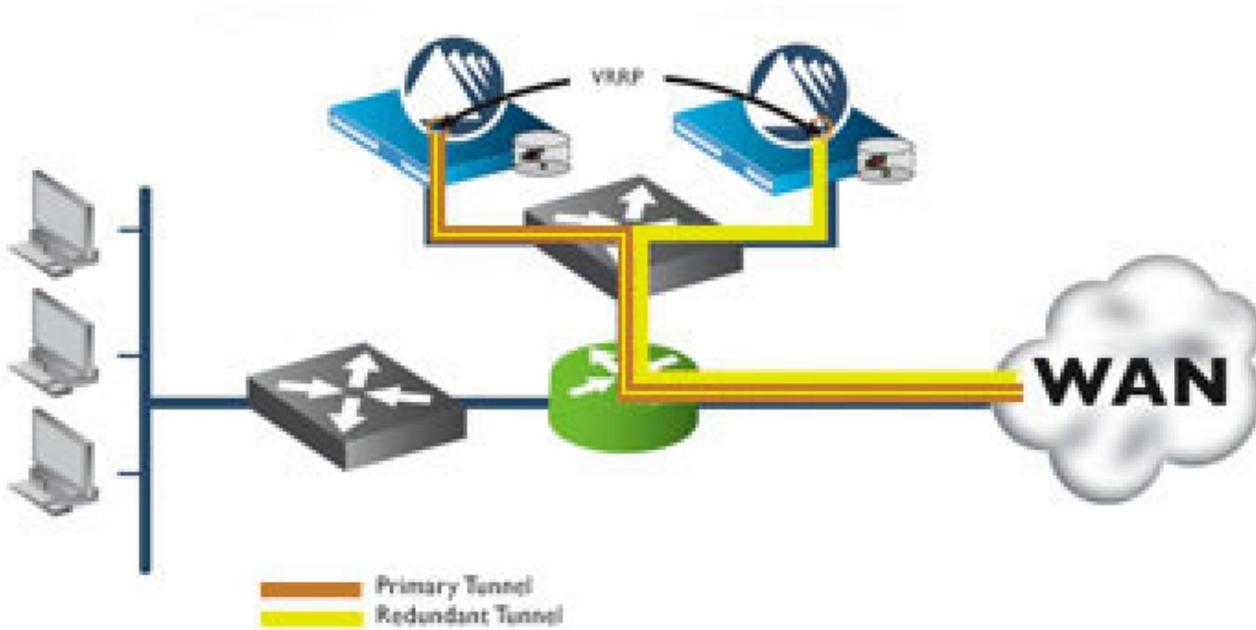
To set up a HA Cluster

1. A single HA link connects the appliances as a cluster.
 - Create a VLAN interface for each WAN interface to support creating tunnels from the other device in the cluster.
 - Assign private IPs to these interfaces
 - The VLAN interfaces must be configured with the same bandwidth as the WAN link it pushes traffic to.
2. Set up the source NAT on each device to harden the tunnel traffic from the private IP of the other device.

3. Build tunnels on the WAN link of each device.
 - Build the tunnel using the private IP and NAT with the public IP of the other device.
4. Create bonded tunnels on each device.
5. Create overlays on each device.
6. Add a route policy to route traffic using the overlays.
7. Might need to adjust the reclassify timer switchover time.

HA Using PBR

Policy-Based-Routing redirection can be used to redirect traffic to an HA pair of appliances.



Silver Peak HA Cluster using Policy-Based Routing Redirection

HA Using WCCP

Web Cache Coordination Protocol (WCCP) can be used to redirect traffic to a redundant pair of Silver Peak NX Series appliances.



Silver Peak HA Cluster using Web Cache Coordination Protocol

Route Policies

BIOs automate the creation of Route Policies, and generally determine to which destination a packet is routed. Route Policy settings are used for exceptions to the BIO configuration.

Shaper

The Shaper provides a simplified way to globally configure QoS on the appliances.

- It shapes traffic by allocating bandwidth as a percentage of the system bandwidth.
- The Shaper 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.

- An inbound Shaper can also be enabled, and is configured independently from outbound Shaping.
- To manage Shaper settings for an appliance system-level WAN Shaper, use the Shaper Template.

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

Quality of Service (QoS)

BIOs automate the creation of QoS policies, and determine which traffic class to place a packet. Quality of Service (QoS) settings are used for exceptions to the BIO configuration.

Set QoS options from a template. The QoS Policy determines how flows are queued and marked. The QoS Policy 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 manually 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.
Or, use the current or default map.

3. Click **Add Rule**.

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

Priority

- 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. Select 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.

Deployment

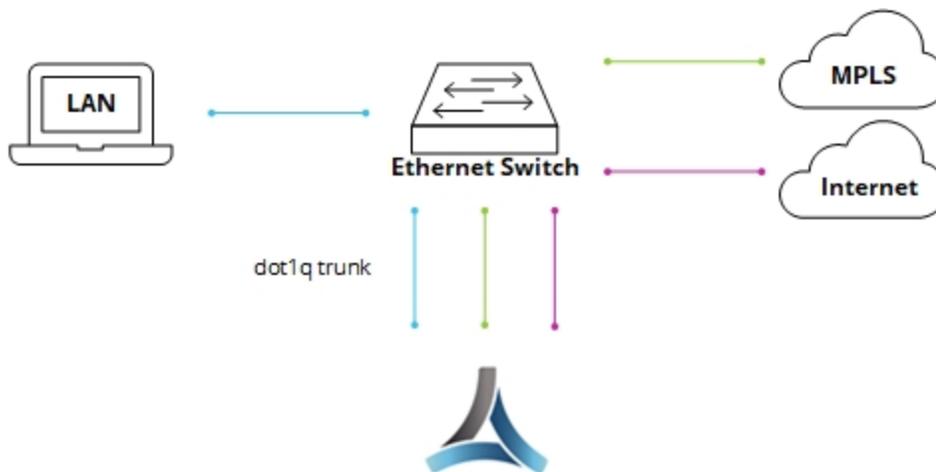
Router Bridge Server

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

- In-line Router** mode - - Uses multiple ports (Example 1) or multiple VLANs (Example 2). Routes between different subnets on LAN/WAN interfaces. Best for complex networks..



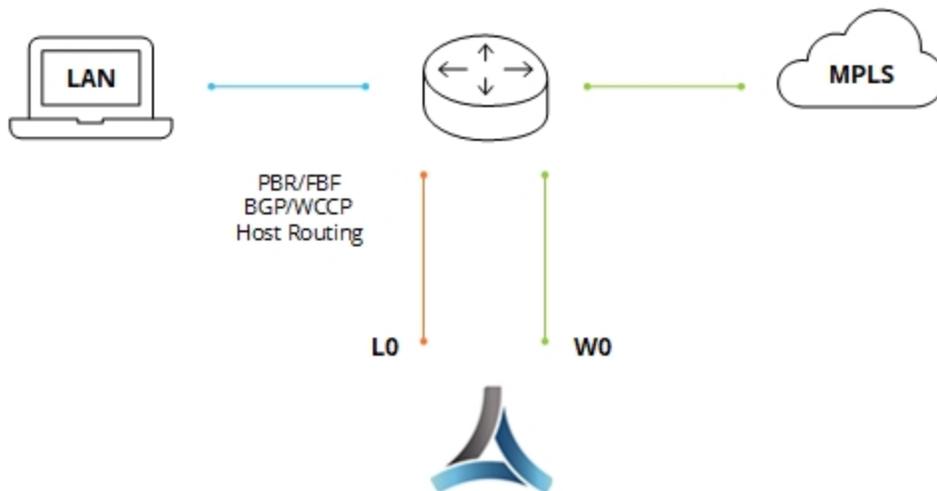
In-line Router Mode Example 1



In-line Router Mode Example 2

- Provides 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).
- In-Line Router** (out-of-path) mode - The LAN interface is configured with an out-of-band subnet. Requires traffic to be redirected to the Silver Peak using BGP, VRRP,

WCCP, or PBR.



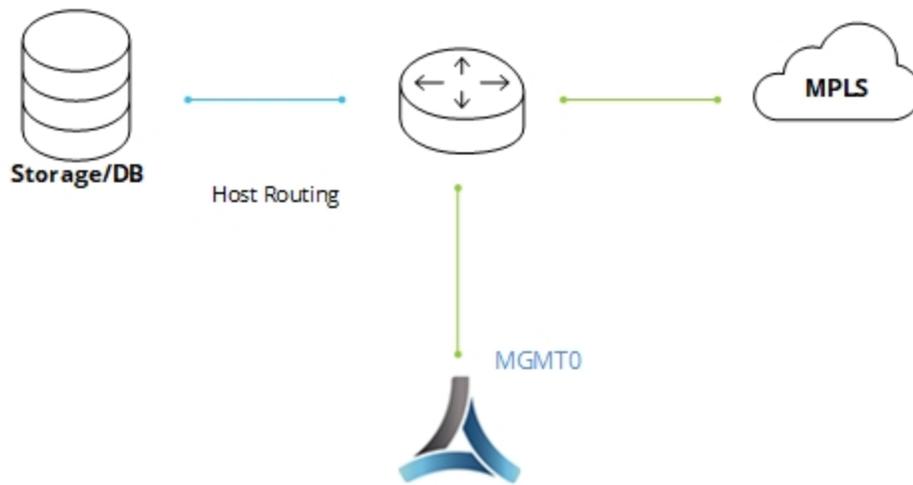
In-Line Router Mode (out-of-path) Example

- **Bridge** mode - Can only be in-line and uses 2 or 4 ports. Bridges two halves of the same subnet on the LAN/WAN interfaces, intercepting traffic according to BIO policy. Good for simple networks.



Simple Bridge Mode Example

- 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.



Server Mode Example

Deployment Profiles

Use Deployment Profiles to standardize your deployments, preventing preventing human error and saving time on repetitive tasks. Instead of configuring each appliance separately, create various Deployment Profiles and provision a device by applying the profile you want.

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.

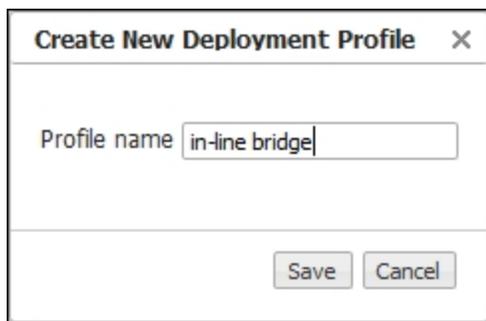
To configure a Deployment Profile

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:** Traffic is routed between separate LAN/WAN interfaces.
- **Bridge:** Uses a virtual interface (**bvi**) created by binding the WAN and LAN interfaces across the same network segment.
- **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.

The screenshot shows the Silver Peak configuration interface for a shaper. At the top, the Silver Peak logo is visible. Below it, there are two main sections. The first section is a table-like interface for mapping interfaces and labels. It has columns for 'Interface', 'Label', 'VLAN', 'Label', and 'Interface'. The first row shows 'lan0' mapped to 'Data' label, with a '+IP' button below it. The second row shows 'wan0' mapped to 'MPLS' label, with a '+Bridge Interface' button below it. To the right of this table is a 'Shaping Kbps' section with a value of '5000' and a 'Σ Calc' button. Below this is a 'Total Outbound' and 'Total Inbound' section, both with input fields and 'Kbps' labels. There is a checkbox for 'Shape Inbound Traffic'. At the bottom, there is an 'EdgeConnect Licensing' section with a 'Plus' checkbox for '> 200 Mbps' and a 'Boost' input field set to '0' with a 'Kbps' label.

- **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, security is provided 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 (TCP 443 to Orchestrator).
 - Select **harden** from the dropdown list to harden the 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.

- Map this to the Internet interface.

The screenshot shows the Silver Peak configuration interface for traffic shaping. It features a table with columns for Interface, Label, VLAN, another Label, and Shaping Kbps. Two rows are visible: one for 'lan0' with 'Data' label and 'wan0' interface, and another for 'lan1' with 'Voice' label and 'wan1' interface. The 'wan1' interface is highlighted in yellow. Below the table, there are fields for 'Total Outbound' (55,000 Kbps) and 'Total Inbound' (empty), with a 'Σ Calc' button. At the bottom, there are 'EdgeConnect Licensing' options for 'Plus' (checkbox for > 200 Mbps) and 'Boost' (0 Kbps).

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

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

Your profile is created.

For more information:

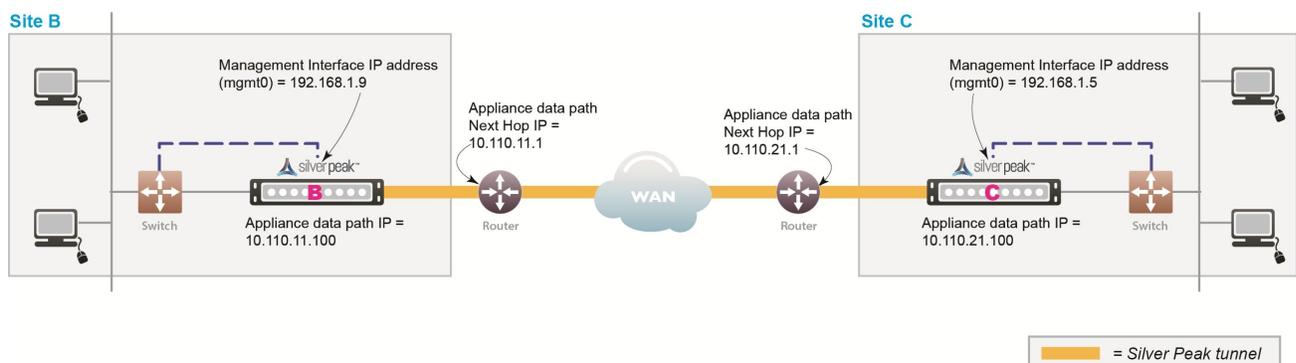
- See the *Unity Orchestrator Operator's Guide* for a more thorough discussion of this topic.
- Within Orchestrator, click the **Question Mark**  on the page for details about each field.

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 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		
mgmt1 IP Address / Mask	---	---

Hostname	B	C
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	<p>Appliance placed in-line between Ethernet LAN switch and WAN router</p> <ul style="list-style-type: none"> ■ Appliance lan0 interface connects to Ethernet LAN switch ■ Appliance wan0 interface connects to WAN router
Fail-Safe Behavior	<p>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).</p> <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	<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 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 (for replication only) 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:

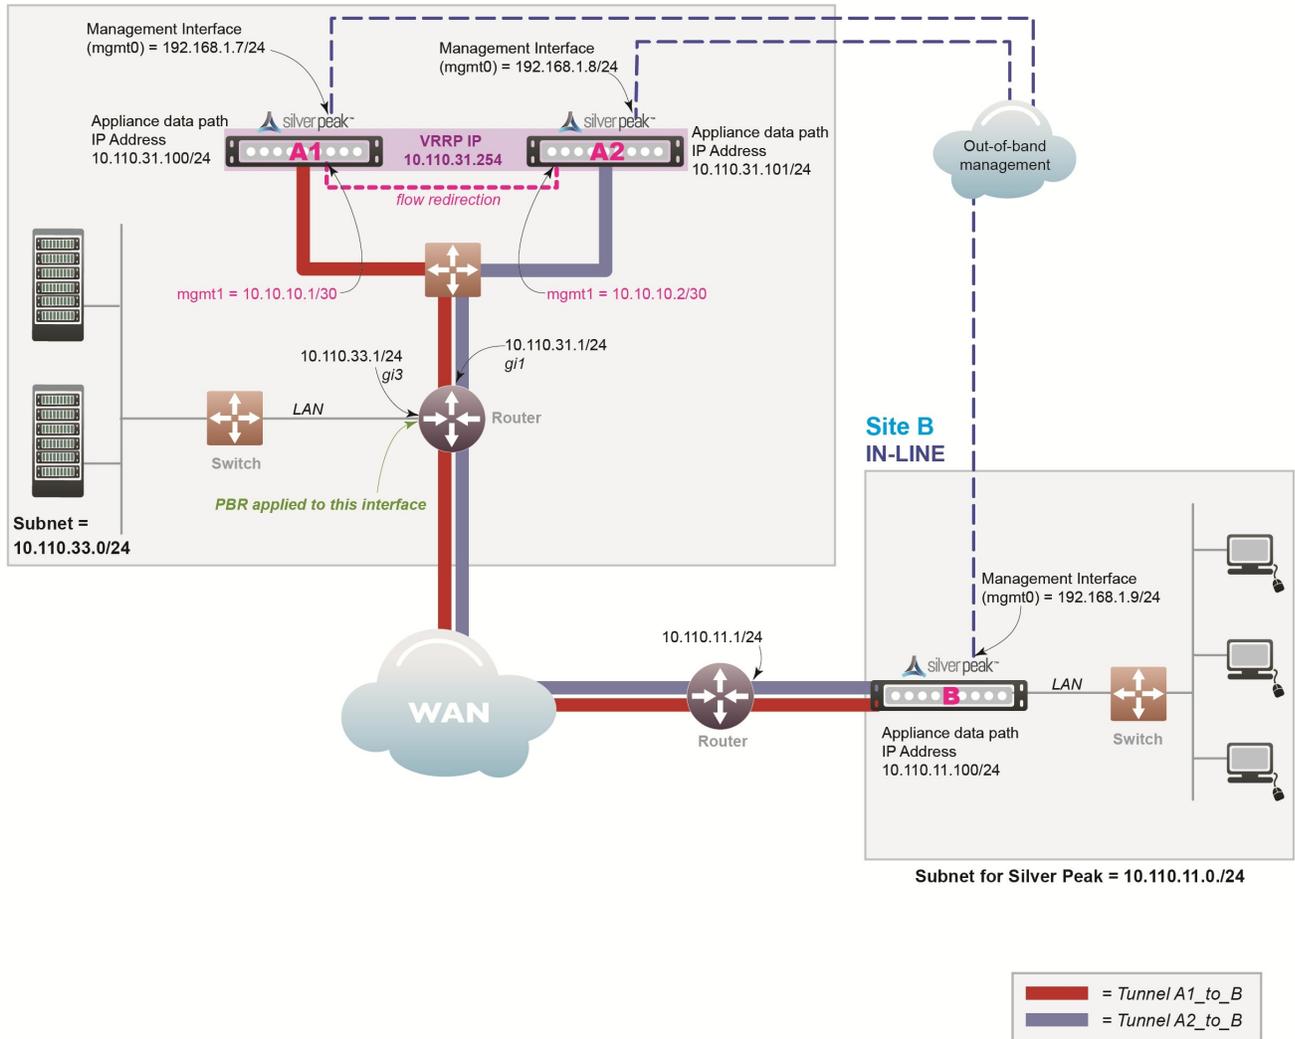
- With the appliance in **bypass** state.
- With the appliance **powered off**.
- With the tunnels administratively **down**.

Policy-Based Routing (PBR)

PBR is configured on the router. No other special configuration is required on the appliance. This is also known as FBF (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



PBR Network

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

PBR Overview

PBR Considerations

Appliance Placement	<p>Both appliances are attached to the same available subnet via an Ethernet LAN switch:</p> <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	<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 seven 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). ■ 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 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	---
	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	---

Hostname	A1	A2	B
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. Also, configure 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.

**Site A Network with Silver Peak
OUT-OF-PATH**

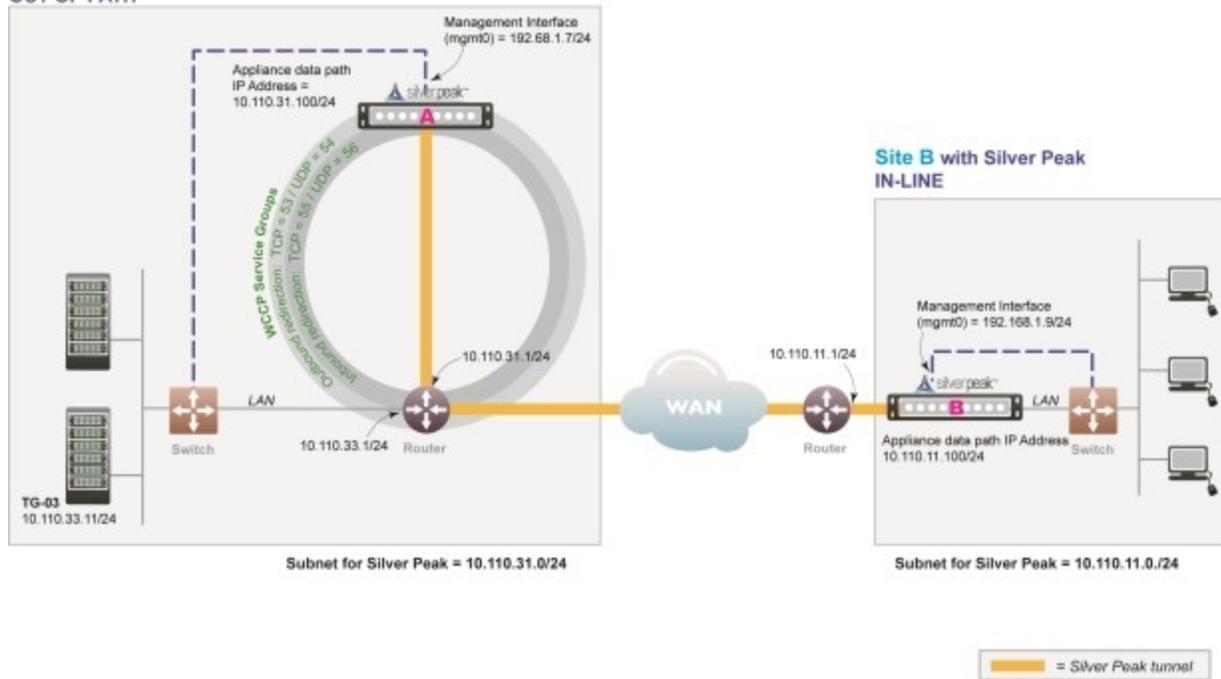


Figure 1. WCCP Network

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.

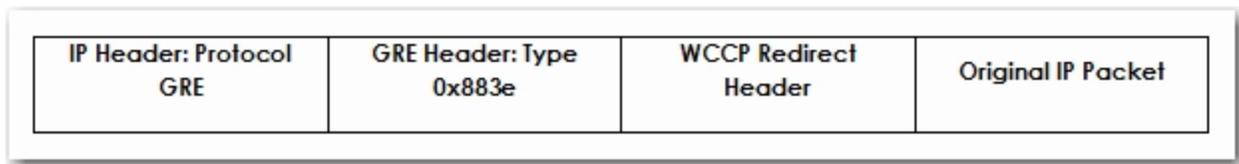
GRE and L2 Redirection

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.

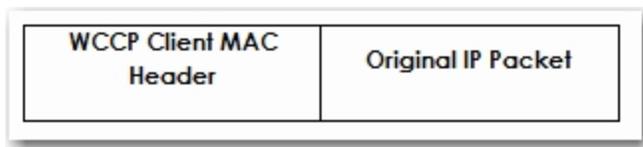
L2 Redirection

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)
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 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)	---

Hostname	A	B
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 might or might 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 to use this feature.

VRRP routing is often used (with subnet sharing) when

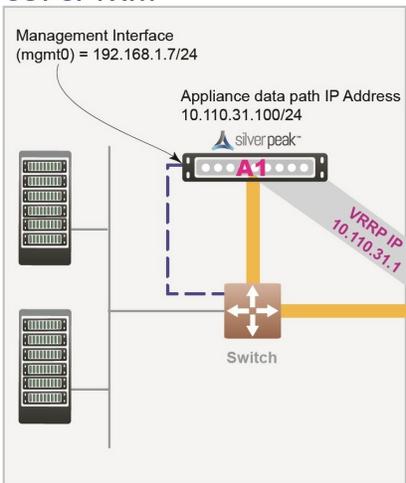
- a Silver Peak appliance uses an existing router when no spare router port is available.
 - OR -
- using redundant Silver Peak appliances with PBR.

VRRP Peering to a WAN Router

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

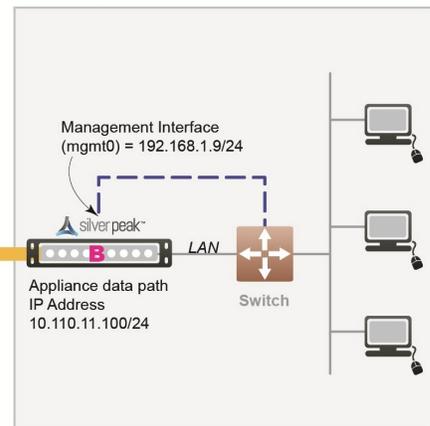
- When no spare router port is available, a single appliance uses VRRP to peer with a router (or Layer 3 switch). This is appropriate for an out-of-path deployment where no redundancy is needed.
- A pair of active, redundant appliances use VRRP to share a common, virtual IP address at their site. This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

Site A Network with Silver Peak OUT-OF-PATH



Subnet for Silver Peak = 10.110.31.0/24

Site B with Silver Peak IN-LINE



Subnet for Silver Peak = 11.110.11.0/24

 = Silver Peak tunnel

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	<p>Appliance shares LAN segment with existing equipment</p> <ul style="list-style-type: none"> ■ Appliance wan0 interface connects to Ethernet LAN switch ■ Do not connect lan0 interface
Failure Method	<p>Fails-Open:</p> <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	<p>This deployment model requires three IP addresses:</p> <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) ■ Virtual IP Address (VIP) shared by Silver Peak appliance and the WAN router <p>The VIP must be the default gateway for the clients and servers on the LAN subnet.</p> <p>NOTE: Typically, this would be the current default gateway, to avoid client reconfigurations.</p> <ul style="list-style-type: none"> ■ 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 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	---

Hostname	A1	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	---
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 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

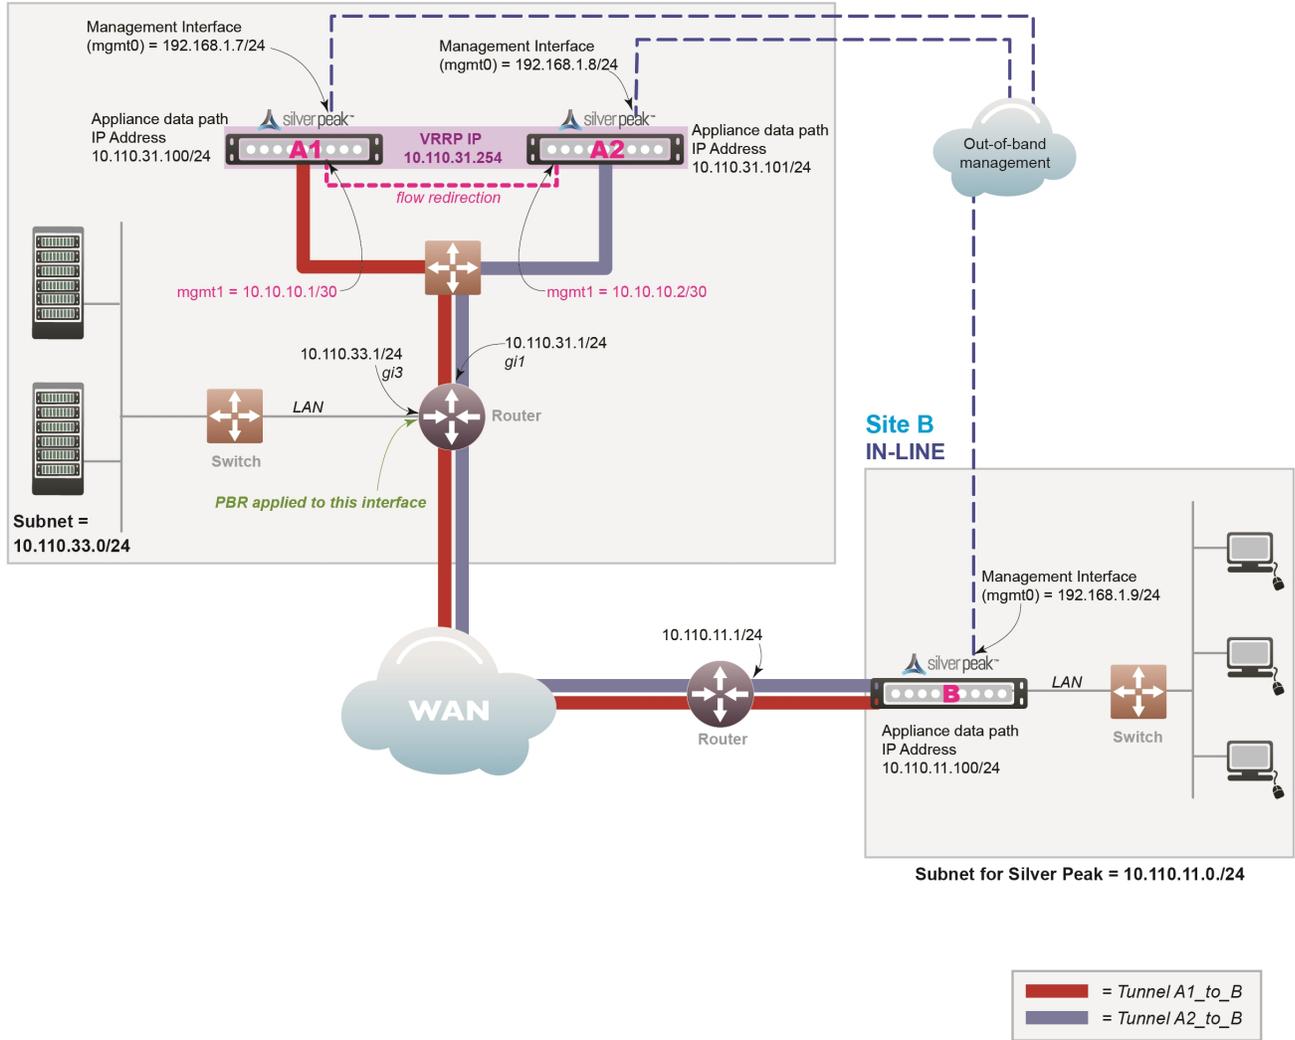
In this example, Site A deploys two redundant appliances out-of-path (Router mode), used as Active and Standby. Site B deploys a single appliance in-line (Bridge mode).

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



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 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 is 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 is 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	---

Hostname	A1	A2	B
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 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.

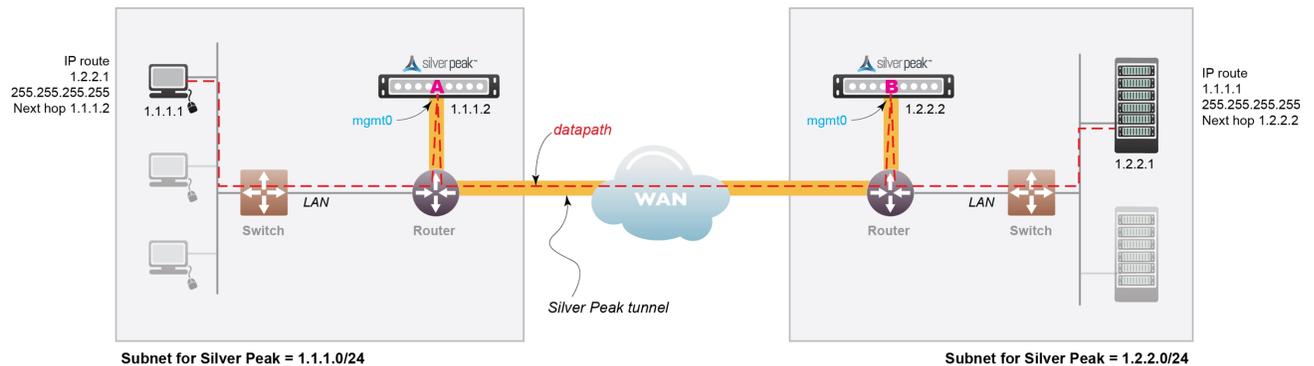
Host-Based Redirection

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.



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 page for a 'Server' mode profile. The profile name is 'Main'. Below the profile name are buttons for '+Add', 'Rename', and 'Delete'. There are three tabs: 'Router', 'Bridge', and 'Server', with 'Server' selected. The Silver Peak logo is displayed above the configuration area. The configuration includes a 'Label' dropdown set to 'MPLS' and an 'Interface' dropdown set to 'mgmt0'. To the right of the interface is a 'Shaping Kbps' field with a value of '10,000'. Below this is a dashed line with a 'Σ Calc' button. Further down are 'Total Outbound' and 'Total Inbound' fields, both with '10,000' Kbps. There is a checkbox for 'Shape Inbound Traffic' which is unchecked. At the bottom, 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 are 'Save', 'Save As', and 'Cancel' buttons.

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 shows the 'Deployment Profiles' configuration page for a 'MediumBranch' profile. It is set to 'Router' mode. The LAN Interfaces section shows 'lan0' with 'None' label and 'No DHCP' option. The WAN Interfaces section shows 'wan0' with 'MPLS' label and 'NAT' selected. Shaping Kbps for 'wan0' is set to 100,000 for both directions. Total Outbound and Inbound are both 150,000 Kbps. EdgeConnect Licensing is set to 'Plus' for > 200 Mbps and 'Boost' is 0 Kbps.

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.

See [WAN Hardening](#).

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.

About Router Mode

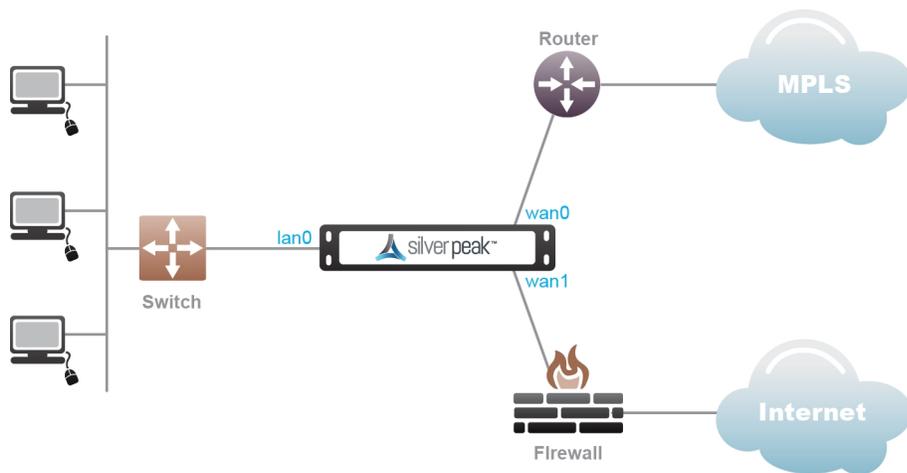
- In-Line Router Mode is recommended whenever possible vs. traditional router mode deployment.
- Limitations when using plain router mode over In-Line Router Mode:
 - Unable to act as a firewall (can't distinguish trusted and untrusted zones).
 - Inaccurate pass-through stats (because inbound and outbound 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)

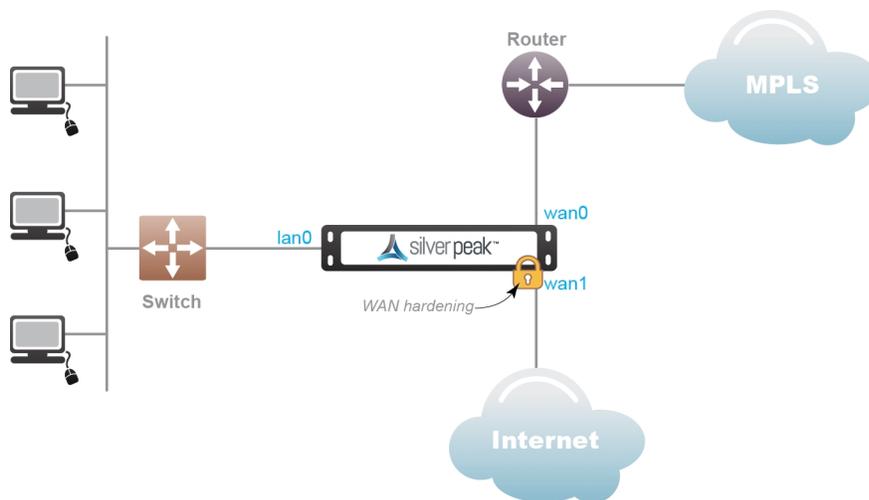
- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- The NAT option must be enabled on the **wan1** interface.
- Local internet access is disabled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



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.



In-Line Router Mode (Router + Direct Internet)

In-Line Router Mode (Single Direct Internet)

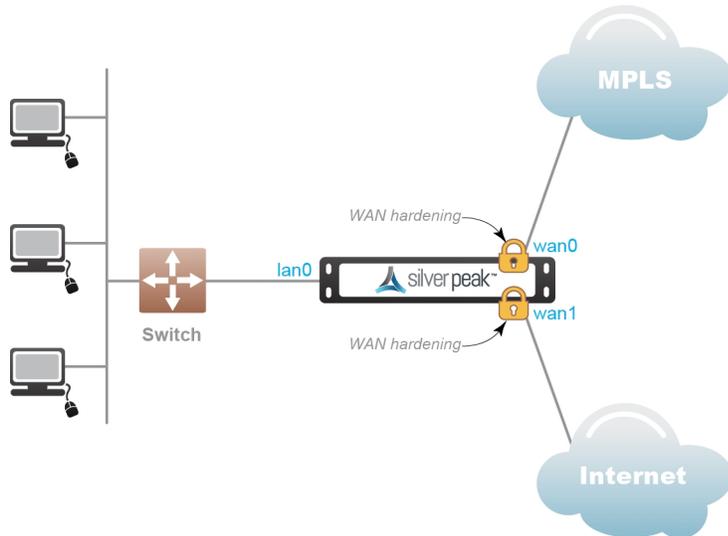
- A default route to the internet must be advertised in subnet sharing from the data center. The **wan0** interface must be hardened.
- Local internet access is not available in this configuration due to WAN hardening.
- Other interfaces are available for more WAN links.



In-Line Router Mode (Single Direct Internet)

In-Line Router Mode (Dual Direct Internet)

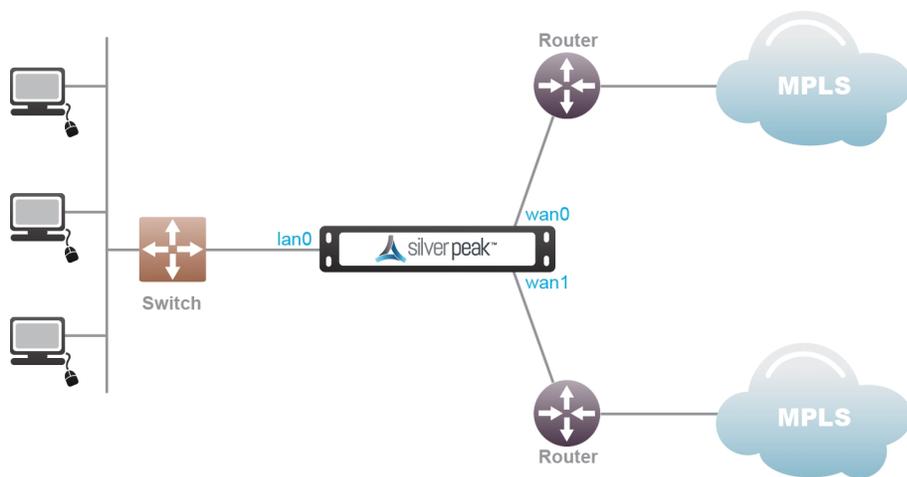
- A default route to the internet must be advertised in subnet sharing from the data center. The **wan0** and **wan1** interfaces must be hardened.
- Local internet access is not available in this configuration due to WAN hardening.
- Other interfaces are available for more WAN links.



In-Line Router Mode (Dual Direct Internet)

In-Line Router Mode (Dual MPLS)

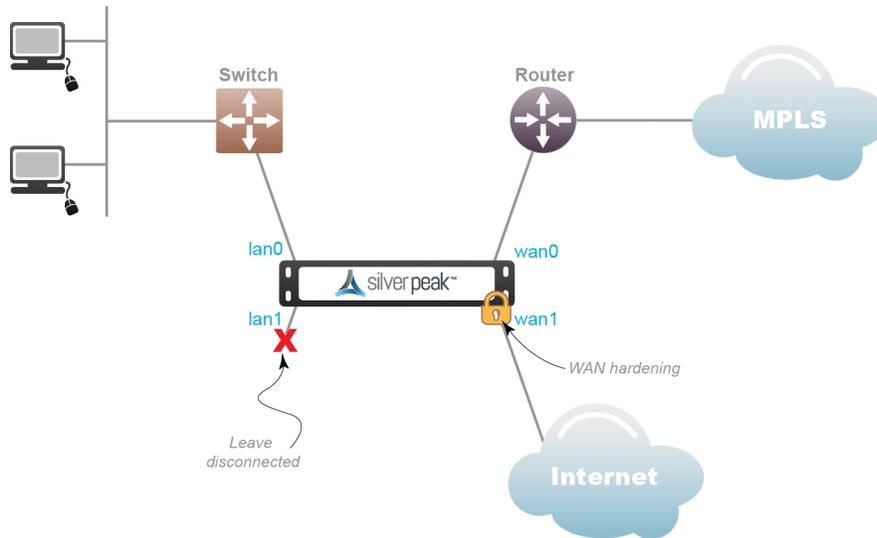
- A default route to the internet must be advertised in subnet sharing from the data center. Other interfaces are available for more WAN links.
- Multicast traffic will be dropped in this configuration and routing will not work between the routers and switch.
- All pass-through traffic goes out **wan0**.



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.



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.

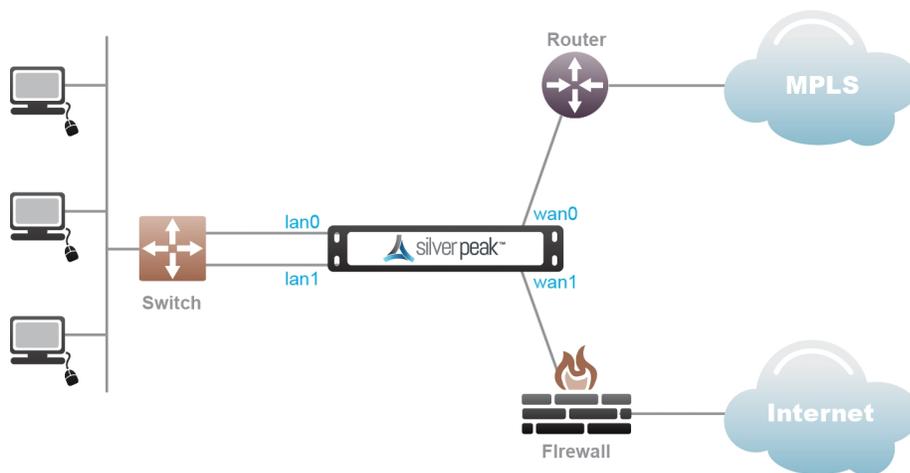
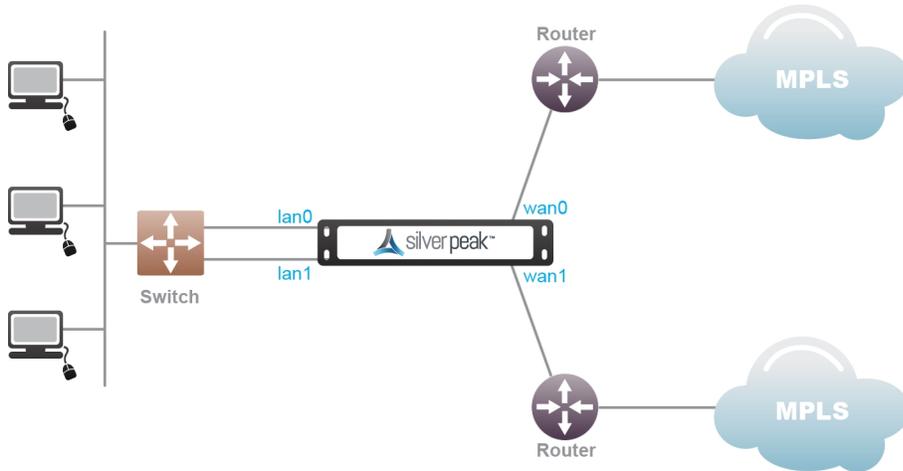


Figure 2. Bridge Mode (Router + Firewall)

Bridge Mode (Dual MPLS)

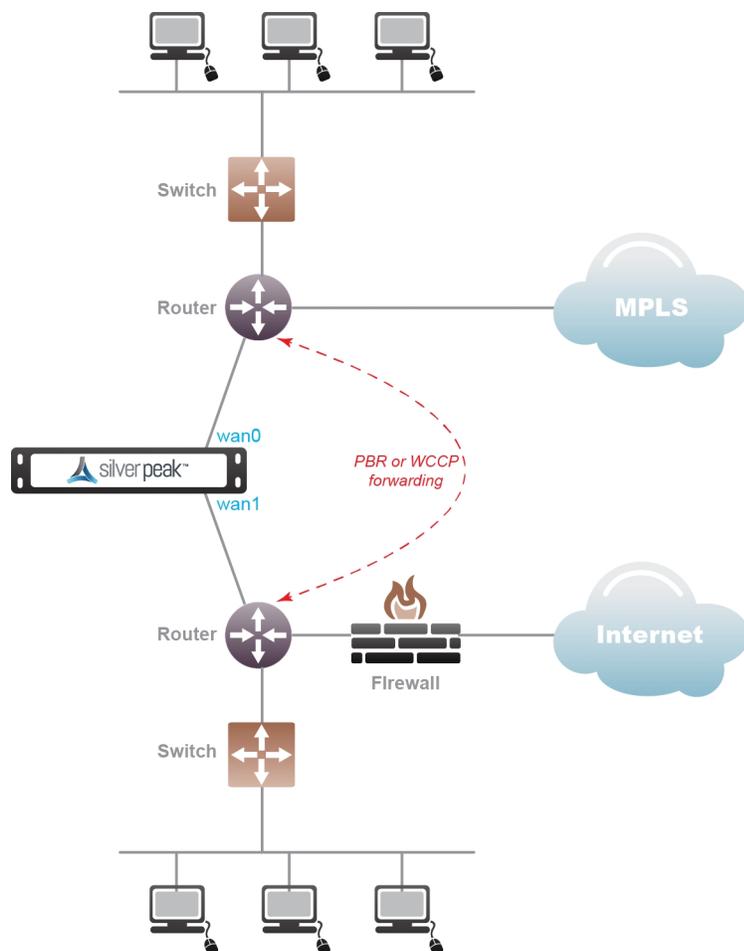
- A default route must be advertised in subnet sharing from the data center for Internet.
- Only two WAN links are supported - **wan0/wan1**.



Bridge Mode (Dual MPLS)

Router Mode MPLS + Internet

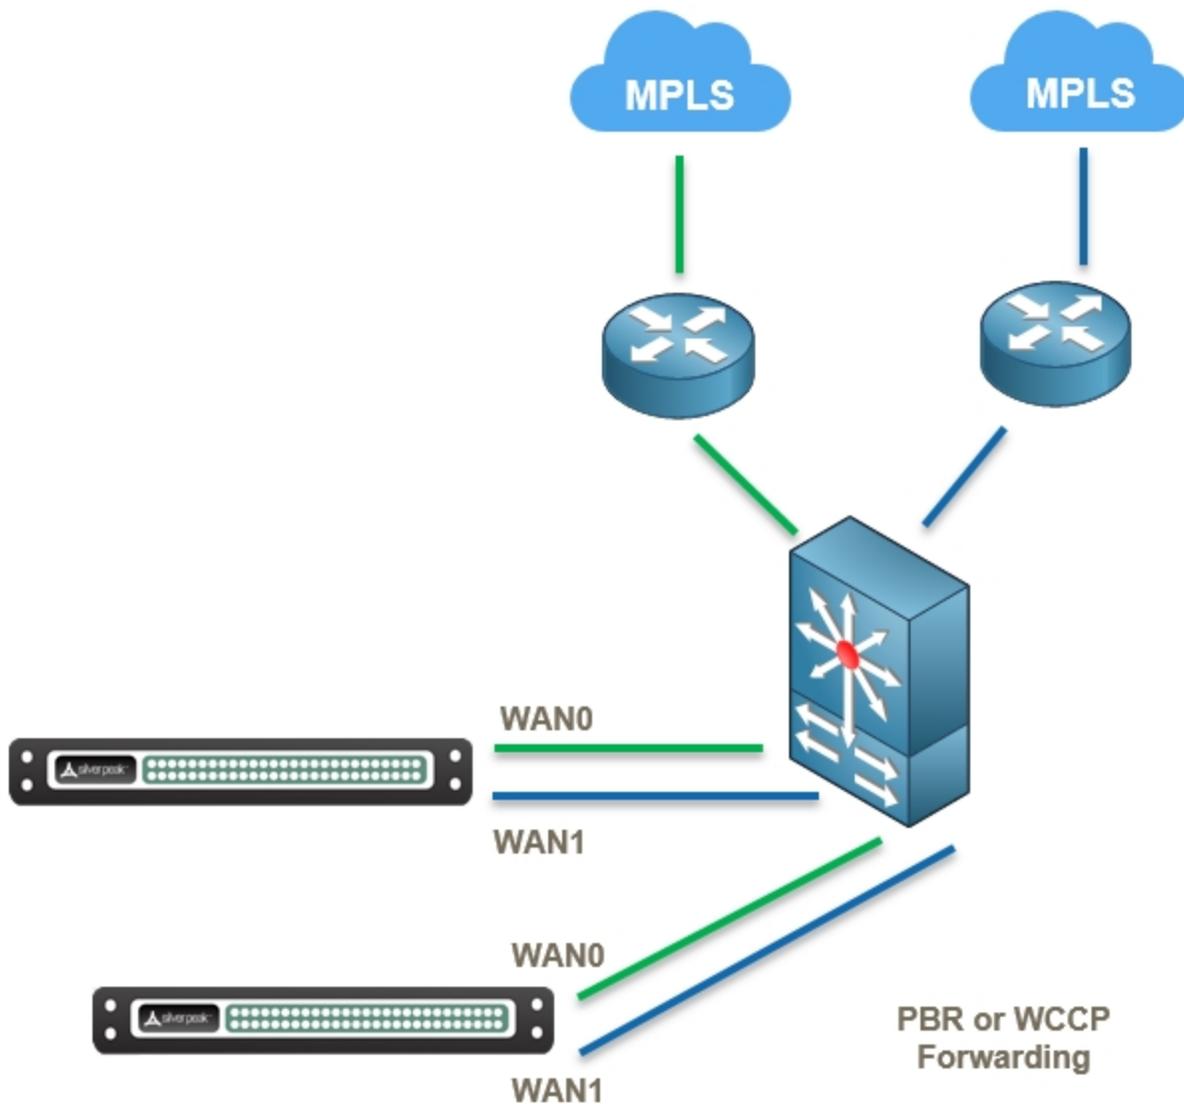
- A default route must be advertised in subnet sharing from the data center for Internet. Only two WAN links are supported - **wan0/wan1**.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- The firewall should have IP protocol 50, UDP 500, and UDP 4500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



Router Mode MPLS + Internet

Router Mode HA (MPLS + MPLS)

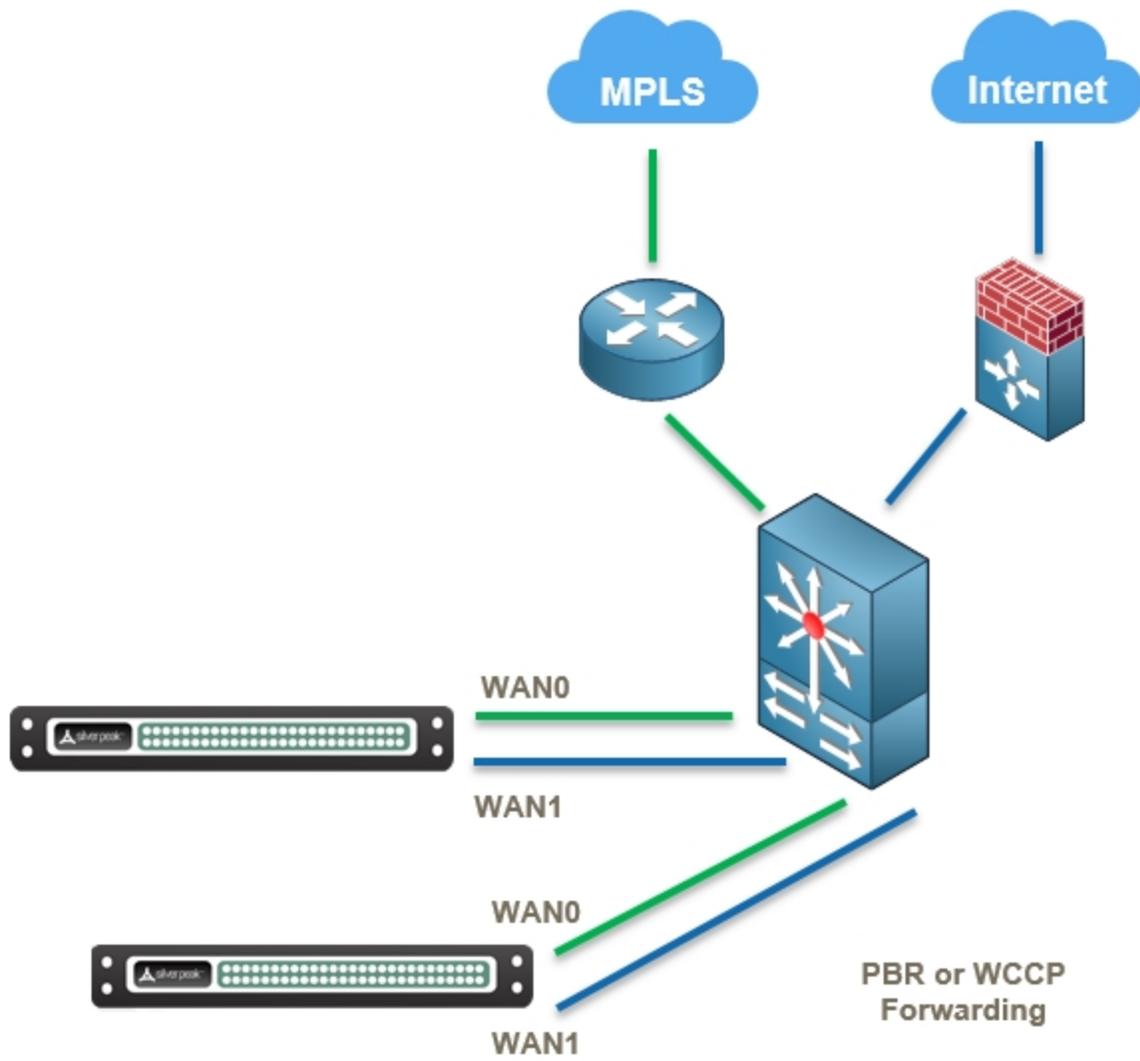
- A default route must be advertised in subnet sharing from the data center for Internet. Only two WAN links are supported - **wan0/wan1**.
- Active/Standby is suggested.



Router Mode HA (MPLS + MPLS)

Dual Home Router Mode (MPLS + Internet)

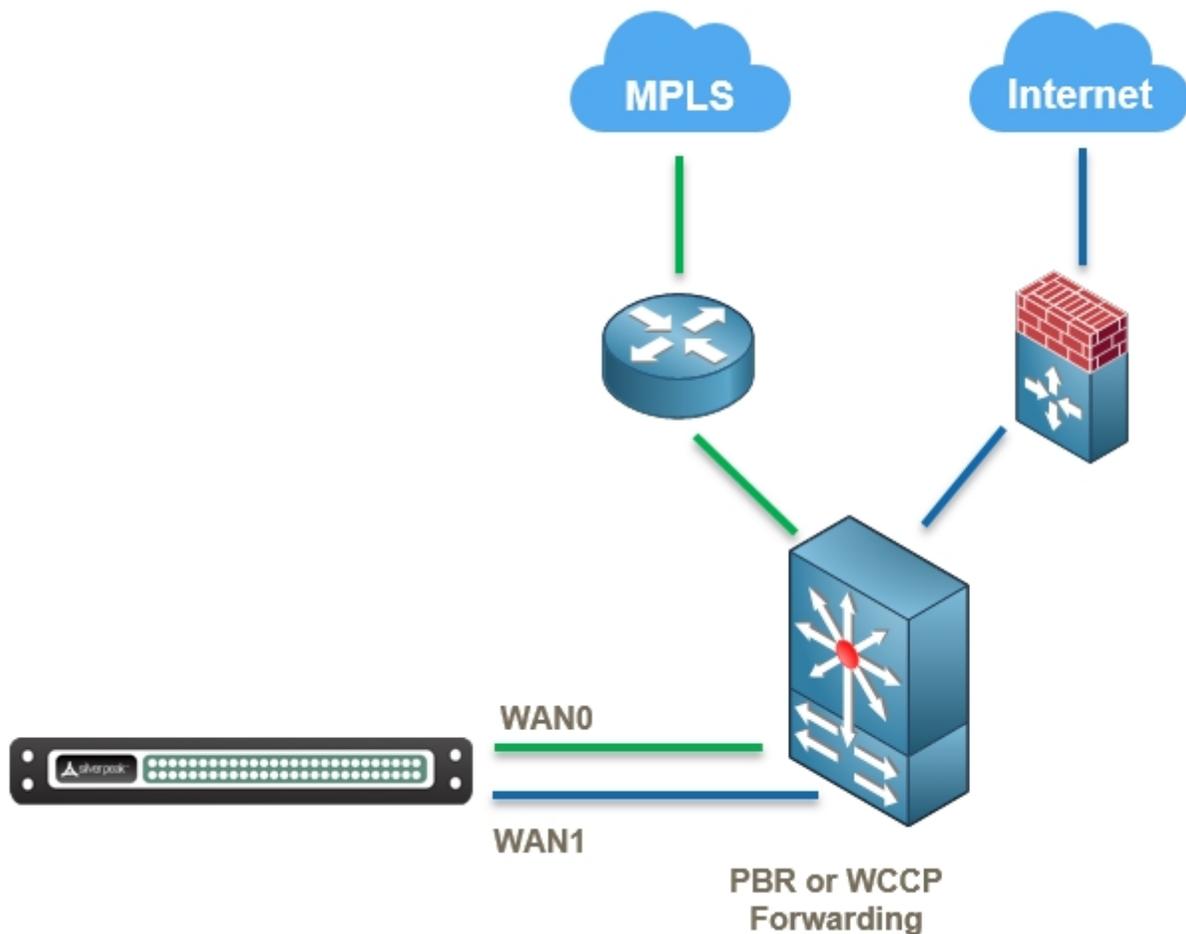
- A default route must be advertised in subnet sharing from the data center for Internet.
- The firewall should have IP protocol 50, UDP 500, and UDP 4500 open to the Silver Peak **wan1** IP address.
- If doing a 1:1 NAT to the Silver Peak **wan1** interface IP, the interface should be hardened.
- Only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



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**.



Dual Home Router Mode HA (MPLS + Internet)

Deploying 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	100
Configuring Orchestrator	101

Installing Orchestrator

To install Orchestrator for the first time:

1. From the email you received from Optimized, 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, initial settings can be changed 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. Using 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.

Add Appliances

After configuring Orchestrator and your overlays, the appliances can be deployed.

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

To Add Appliances

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.

Enable Subnet Sharing 104

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.

Troubleshooting

When things aren't working right, it's either the network, the hardware, or the software.

Network issues are usually caused by hardware, but could also be a result of incorrect software configuration.

Tools for checking the Network are [Ping and Traceroute](#).

Hardware issues that are identified usually need to be repaired or replaced with updated equipment.

Software issues can be addressed by Silver Peak Support, which is 24/7.

From the **Support** tab:

- The [Technical Support](#) page contains links and phone numbers for immediate help.
- Access the [Support Portal Login](#) to open or manage a Silver Peak case. Response time from the Support team depends on the priority you have set.
- [Debug Files](#) enables you to generate a system dump file (which might take a few minutes)

Technical Support

This page lists various ways to contact Silver Peak Technical Support.

- Registered users can open and track their support tickets from the Support Portal (website).
- Other users can request access to the Support Portal by phone or email.

Support

Technical Support

Model	VX-3000 205002006000 Rev 44261
Serial Number	00-0C-29-47-5A-FE
Release	VXOA 7.3.0.0_55922

Website www.silver-peak.com/Support

User Documentation www.silver-peak.com/Support/user_docs.asp

Email support@silver-peak.com

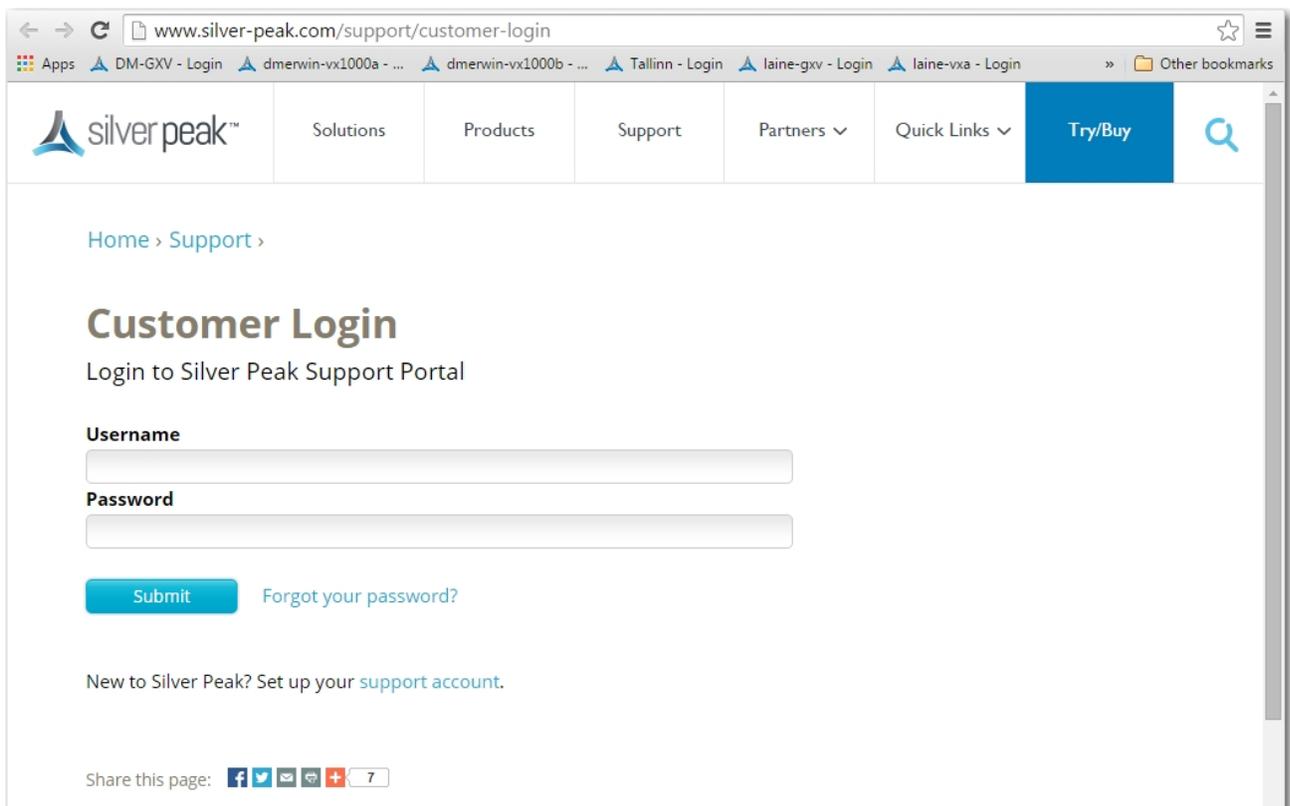
Phone

North America (USA/CAN)	+1 877 210 7325
Global	+1 408 935 1850

Support Portal Login

This link opens a new tab in your browser to the Silver Peak Support Portal.

- Registered users can view and manage support tickets, updates, and other account information from the Support Portal.
- Users can request access to the Support Portal by phone, email, or click **Support account**.



Debug Files

The appliance automatically creates and stores a number of non-configuration data files as a result of normal events, traffic monitoring, system crashes, and testing.

Silver Peak uses these files for evaluation and debugging.

Debug Files Disk Usage
Usage Percent 43% 10,984 MB

Files: All Logs Sys Dump Snapshot TCP Dump Show Tech Generate Sys Dump Generate Show Tech Download Selection to My Computer Refresh 2 mins ago

Show 25 Search

File Type	File Name	Last Modified	File Size	
Logs	messages	Wed, 15 Jul 2015 21:44:26 GMT	6.4MB	
Logs	web_access_log	Wed, 15 Jul 2015 21:44:26 GMT	24.5MB	
Logs	auditlog	Wed, 15 Jul 2015 21:43:11 GMT	12.5MB	
Sys Dump	tunbug-20150715.tar	Wed, 15 Jul 2015 21:00:05 GMT	501.8KB	X
Sys Dump	tunbug.lock	Wed, 15 Jul 2015 21:00:00 GMT	0B	X
Logs	messages.1.gz	Wed, 15 Jul 2015 20:26:01 GMT	2.9MB	
Logs	web_error_log	Wed, 15 Jul 2015 19:59:42 GMT	40.5MB	
Logs	messages.2.gz	Wed, 15 Jul 2015 09:50:02 GMT	3.0MB	
Sys Dump	tunbug-20150714.tar.gz	Wed, 15 Jul 2015 06:00:05 GMT	790.3KB	X
Logs	messages.3.gz	Tue, 14 Jul 2015 23:02:01 GMT	3.1MB	
Logs	alerts	Tue, 14 Jul 2015 22:49:55 GMT	3.2MB	
Logs	messages.4.gz	Tue, 14 Jul 2015 12:04:01 GMT	3.0MB	
Sys Dump	tunbug-20150713.tar.gz	Tue, 14 Jul 2015 06:00:05 GMT	812.6KB	X
Logs	messages.5.gz	Tue, 14 Jul 2015 00:30:02 GMT	3.0MB	
Logs	messages.6.gz	Mon, 13 Jul 2015 13:18:02 GMT	3.0MB	

Showing 1 to 25 of 101 entries First Previous 1 2 3 4 5 Next Last

High Packet Loss

The **Flows** page shows a high-level view of your network.

The screenshot shows the 'Flows' page interface. At the top, there are tabs for 'Dashboard' and 'Flows'. Below this, there are several filter sections: 'Flow Categories' (All, Pass-Through, Asymmetric, Stale), 'Flow Timing' (Active, Active + Ended Last 5min, Started Last 5min, Ended Last 5min, Ended), 'Bytes Transferred' (Total, Last 5m), and 'Flow Characteristics' (Flows to Slow Devices). There are also input fields for Application, Domain, IP Intelligence, IP1, IP2, Port1, Port2, Protocol, VLAN id, and Traffic. A 'Max Flows' field is set to 7000. Below the filters, there are buttons for 'Reset Flows', 'Reclassify Flows', 'Customize', and 'Export'. A summary line shows 'Active 1 | Pass-Through 1 | Asymmetric 1 | Stale 0 | Displayed 1 | Matched 1'. The main table has one row with the following data:

IP1	Port1	IP2	Port2	Detail	Application	Location	IP Intelligence	Bytes
10.10.128.20	36916	52.70.165.138	443		https	United States, Ashb...		104

Go to **Monitoring > Loss > Summary** to see the results of any overlay tunnels between appliances.

Go to **Monitoring > Out-of-Order Packets > Summary** to see which packets are sent out of order.

Log Settings

Configuring local and remote logging requires that you specify the minimum security level of an event to log.



- Set up local logging in the **Log Configuration** section.
- Set up remote logging by using the **Log Facilities Configuration** and **Remote Log Receivers** sections.

Minimum Severity Levels

In decreasing order of severity, the levels are as follows.

Level	Description
EMERGENCY	The system is unusable.
ALERT	Includes all alarms the appliance generates: CRITICAL , MAJOR , MINOR , and WARNING
CRITICAL	A critical event
ERROR	An error. This is a non-urgent failure.
WARNING	A warning condition. Indicates an error will occur if action is not taken.
NOTICE	A normal, but significant, condition. No immediate action required.
INFO RMATIONAL	Informational. Used by Silver Peak for debugging.
DEBUG	Used by Silver Peak for debugging
NONE	If you select NONE , then no events are logged.

- The bolded part of the name is what displays in Silver Peak's logs.
- If you select **NOTICE** (the default), then the log records any event with a severity of **NOTICE, WARNING, ERROR, CRITICAL, ALERT, and EMERGENCY**.
- These are purely related to event logging levels, **not** alarm severities, even though some naming conventions overlap. Events and alarms have different sources. Alarms, once they clear, list as the **ALERT** level in the **Event Log**.

Configuring Remote Logging

- Configure the appliance to forward all events at and above a specified severity, to a remote syslog server.
- A syslog server is independently configured for the minimum severity level that it will accept. Without reconfiguring, it might not accept as low a severity level as you are forwarding to it.
- In the **Log Facilities Configuration** section, assign each message/event type (**System / Audit / Flow**) to a syslog facility level (**local0** to **local7**).
- For each remote syslog server that you add to receive the events, specify the receiver's IP address, along with the messages' minimum severity level and facility level.

Log Viewer

This page displays three types of logs for troubleshooting --- **Audit**, **Event**, and **Alarm**.

Log Viewer ?

Filters

Type Event End Date Now Record Count 100 (Max 10000) Retrieve Logs

Show 100 Search

Date/Time	Hostname	Level	Process	PID	TID	
07/15/2015 14:11:10	Seattle-EC	NOTICE	jsond	2204	139992985995024	Action set successful: /action
07/15/2015 14:11:10	Seattle-EC	NOTICE	mgmtd	2029	139819964372752	Action performed by user 2204-0-0:
07/15/2015 14:11:10	Seattle-EC	NOTICE	mgmtd	2029	139819964372752	Ignoring preferred GW request (not
07/15/2015 14:11:10	Seattle-EC	INFO	mgmtd	2029	139819964372752	ACTION: [0] ifname = mgmt0 (string
07/15/2015 14:11:10	Seattle-EC	INFO	mgmtd	2029	139819964372752	ACTION: /net/routes/actions/overla
07/15/2015 14:11:10	Seattle-EC	INFO	mgmtd	2029	139819964372752	Handling ACTION request
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Forking then execing binary "/sbin/if
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Launched process /sbin/ifconfig with
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Forking then execing binary "/sbin/if
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Launched process /sbin/ifconfig with
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Forking then execing binary "/sbin/if
07/15/2015 14:11:09	Seattle-EC	INFO	mgmtd	2029	139819964372752	Launched process /sbin/ifconfig with

Showing 1 to 100 of 100 entries First Previous 1 Next Last

After changing any **Filters** parameter, click **Retrieve Logs** to update the table.

Ping and Traceroute

Use the **ping** and **traceroute** commands to help diagnose network connectivity problems.

Ping/Traceroute ?

Network Connectivity

IP/Hostname

Options

Output

```

PING 10.0.238.136 (10.0.238.136) 56(84) bytes of data.
64 bytes from 10.0.238.136: icmp_seq=1 ttl=64 time=1.89 ms
64 bytes from 10.0.238.136: icmp_seq=2 ttl=64 time=0.138 ms
64 bytes from 10.0.238.136: icmp_seq=3 ttl=64 time=0.205 ms
64 bytes from 10.0.238.136: icmp_seq=4 ttl=64 time=0.392 ms
64 bytes from 10.0.238.136: icmp_seq=5 ttl=64 time=0.100 ms

--- 10.0.238.136 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4003ms
rtt min/avg/max/mdev = 0.100/0.545/1.894/0.682 ms

```

Using ping

- Use the **ping** command to send Internet Control Message Protocol (ICMP) echo requests to a specified host.
- By default, the **ping** command uses the **mgmt0** interface. If you want to ping out of a datapath interface, use the **-I** option with the local appliance IP address. For example:



■ SYNOPSIS:

```
ping [ -LRUbdfnqrvVaAB] [ -c count] [ -i interval] [ -l
preload] [ -p pattern] [ -s packetsize] [ -t ttl] [ -w
deadline] [ -F flowlabel] [ -I interface] [ -M hint] [ -Q
tos] [ -S sndbuf] [ -T timestamp option] [ -W timeout] [
hop ...] destination
```

The following **ping** options are supported:

Option	Description
-A	Adaptive ping. Interpacket interval adapts to round-trip time, so that effectively not more than one (or more, if preload is set) unanswered probes present in the network. Minimal interval is 200 msec if not super-user. On networks with low rtt this mode is essentially equivalent to flood mode.
-b	Allow pinging a broadcast address.
-B	Do not allow ping to change source address of probes. The address is bound to one selected when ping starts.
-c	count: Stop after sending count ECHO_REQUEST packets. With deadline option, ping waits for count ECHO_REPLY packets, until the time-out expires.
-d	Set the SO_DEBUG option on the socket being used. Essentially, this socket option is not used by Linux kernel.
-F	flow label: Allocate and set 20 bit flow label on echo request packets. (Only ping6). If value is zero, kernel allocates random flow label.
-i	interval: Wait interval seconds between sending each packet. The default is to wait for one second between each packet normally, or not to wait in flood mode. Only a super-user can set interval to values less 0.2 seconds.
-I	interface address: Set source address to specified interface address. Argument can be a numeric IP address or name of device. When pinging IPv6 link-local address this option is required.
-l	preload: If preload is specified, ping sends that many packets not waiting for reply. Only the super-user can select preload more than 3.

Option	Description
-L	Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.
-M	MTU discovery <i>hint</i> : Select Path MTU Discovery strategy. The hint can be either do (prohibit fragmentation, even local one), want (do PMTU discovery, fragment locally when packet size is large), or dont (do not set DF flag).
-n	Numeric output only. No attempt is made to lookup symbolic names for host addresses.
-p	pattern : You can specify up to 16 "pad" bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, -p ff will cause the sent packet to be filled with all ones.
-Q	<ul style="list-style-type: none"> ■ tos: Set Quality of Service -related bits in ICMP datagrams. tos can be either decimal or hex number. ■ Traditionally (RFC1349), these have been interpreted as: 0 for reserved (currently being redefined as congestion control), 1-4 for Type of Service and 5-7 for Precedence. ■ Possible settings for Type of Service are: minimal cost: 0x02, reliability: 0x04, throughput: 0x08, low delay: 0x10. ■ Multiple TOS bits should not be set simultaneously. ■ Possible settings for special Precedence range from priority (0x20) to net control (0xe0). You must be root (CAP_NET_ADMIN capability) to use Critical or higher precedence value. Bit 0x01 (reserved) can't be set unless ECN has been enabled in the kernel. ■ In RFC2474, these fields has been redefined as 8-bit Differentiated Services (DS), consisting of: bits 0-1 of separate data (ECN will be used, here), and bits 2-7 of Differentiated Services Codepoint (DSCP).
-q	Quiet output. Nothing is displayed except the summary lines at startup time and when finished.
-R	Record route. Includes the RECORD_ROUTE option in the ECHO_REQUEST packet and displays the route buffer on returned packets. Note that the IP header is only large enough for nine such routes. Many hosts ignore or discard this option.
-r	Bypass the normal routing tables and send directly to a host on an attached interface. If the host is not on a directly attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it provided the option -I is also used.
-s	packetsize : Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.
-S	sndbuf : Set socket sndbuf. If not specified, it is selected to buffer not more than one packet.
-t ttl	Set the IP Time to Live.

Option	Description
-T	timestamp option: Set special IP timestamp options. timestamp option can be either tsonly (only timestamps), tsandaddr (timestamps and addresses) or tsprospec host1 [host2 [host3 [host4]]] (timestamp prespecified hops).
-U	Print full user-to-user latency (the old behavior). Normally ping prints network round trip time, which can be different f.e. due to DNS failures.
-v	Verbose output.
-V	Show version and exit.
-w	deadline: Specify a timeout, in seconds, before ping exits regardless of how many packets have been sent or received. In this case ping does not stop after count packet are sent, it waits either for deadline expire or until count probes are answered or for some error notification from network.

Using traceroute

- Use the **traceroute** command to trace the route that packets take to a destination.
- When latency is high, use this command to troubleshoot.
- SYNOPSIS:

```
traceroute [ -dFInrvx ] [ -f first_ttl ] [ -g gateway ] [
-i iface ] [ -m max_ttl ] [ -p port ] [ -q nqueries ] [ -s
src_addr ] [ -t tos ] [ -w waittime ] [ -z pausesecs ]
host [ packetlen ]
```

The following **traceroute** options are supported:

Option	Description
-d	Enable socket level debugging.
-f	Set the initial time-to-live used in the first outgoing probe packet.
-F	Set the “don’t fragment” bit.
-g	Specify a loose source route gateway (8 maximum).
-i	Specify a network interface to obtain the source IP address for outgoing probe packets. This is normally only useful on a multi-homed host. (See the -s flag for another way to do this.)
-I	Use ICMP ECHO instead of UDP datagrams.

Option	Description
-m	Set the max time-to-live (max number of hops) used in outgoing probe packets. The default is 30 hops (the same default used for TCP connections).
-n	Print hop addresses numerically rather than symbolically and numerically (saves a nameserver address-to-name lookup for each gateway found on the path).
-p	Set the base UDP port number used in probes (default is 33434). Traceroute hopes that nothing is listening on UDP ports base to base + nhops - 1 at the destination host (so an ICMP PORT_UNREACHABLE message will be returned to terminate the route tracing). If something is listening on a port in the default range, this option can be used to pick an unused port range.
-q	nqueries
-r	Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it (for example, after the interface was dropped by routed (8C)).
-s	Use the following IP address (which usually is given as an IP number, not a hostname) as the source address in outgoing probe packets. On multi-homed hosts (those with more than one IP address), this option can be used to force the source address to be something other than the IP address of the interface the probe packet is sent on. If the IP address is not one of this machine's interface addresses, an error is returned and nothing is sent. (See the -i flag for another way to do this.)
-t	Set the type-of-service in probe packets to the following value (default zero). The value must be a decimal integer in the range 0 to 255. This option can be used to see if different types-of-service result in different paths. (If you are not running 4.4bsd, this might not matter since the normal network services like telnet and ftp don't let you control the TOS). Not all values of TOS are legal or meaningful - see the IP spec for definitions. Useful values are probably 16 (low delay) and 8 (high throughput). If TOS value is changed by intermediate routers, (TOS=<value>!) will be printed once: value is the decimal value of the changed TOS byte.
-v	Verbose output. Received ICMP packets other than TIME_EXCEEDED and UNREACHABLEs are listed.
-w	Set the time (in seconds) to wait for a response to a probe (default 5 sec.).
-x	Toggle ip checksums. Normally, this prevents traceroute from calculating ip checksums. In some cases, the operating system can overwrite parts of the outgoing packet but not recalculate the checksum (so in some cases the default is to not calculate checksums and using -x causes them to be calculated). Note that checksums are usually required for the last hop when using ICMP ECHO probes (-I). So they are always calculated when using ICMP.
-z	Set the time (in milliseconds) to pause between probes (default 0). Some systems such as Solaris and routers such as Ciscos rate limit icmp messages. A good value to use with this is 500 (e.g. 1/2 second).

Tunnels Not Showing on Tunnels Page

If your tunnels are not showing, confirm these items:

- 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 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 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, 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 data path address as the ping source address, use the **-I** option (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.

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.

Orchestrator Reports

To access custom reports generated by Orchestrator, go to **Monitoring > Status & Reporting > Schedule & Run Reports**. For a detailed description of reporting features, see the [Unity Orchestrator Operator Guide](#).

To access all reports residing on the Orchestrator server, click **View Reports**. Orchestrator retains reports and zipped .csv files for 30 days.

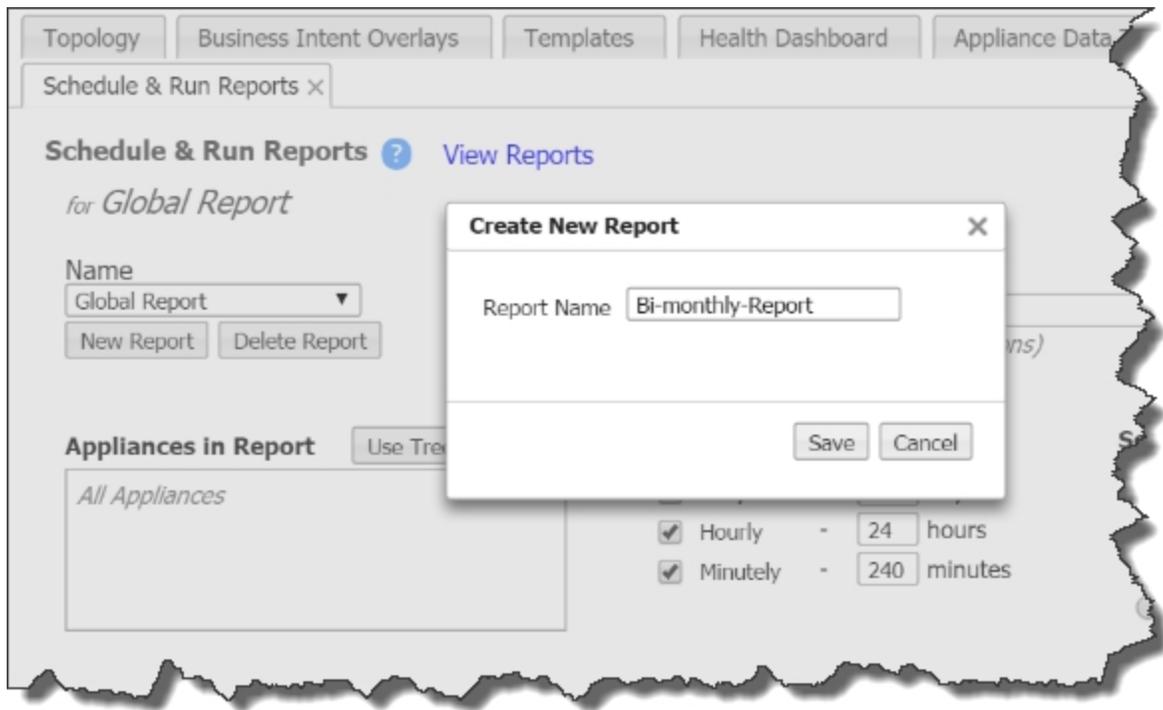
Scheduling a Report

Scheduled or on-demand reports can be generated Daily, Hourly, or by the Minute containing user-selected charts. Check the reports you want to run, then click **Run Now** to run a report immediately, or **Edit** the Run Scheduled Report option. You can have multiple reports, each running at different times.

To schedule a report

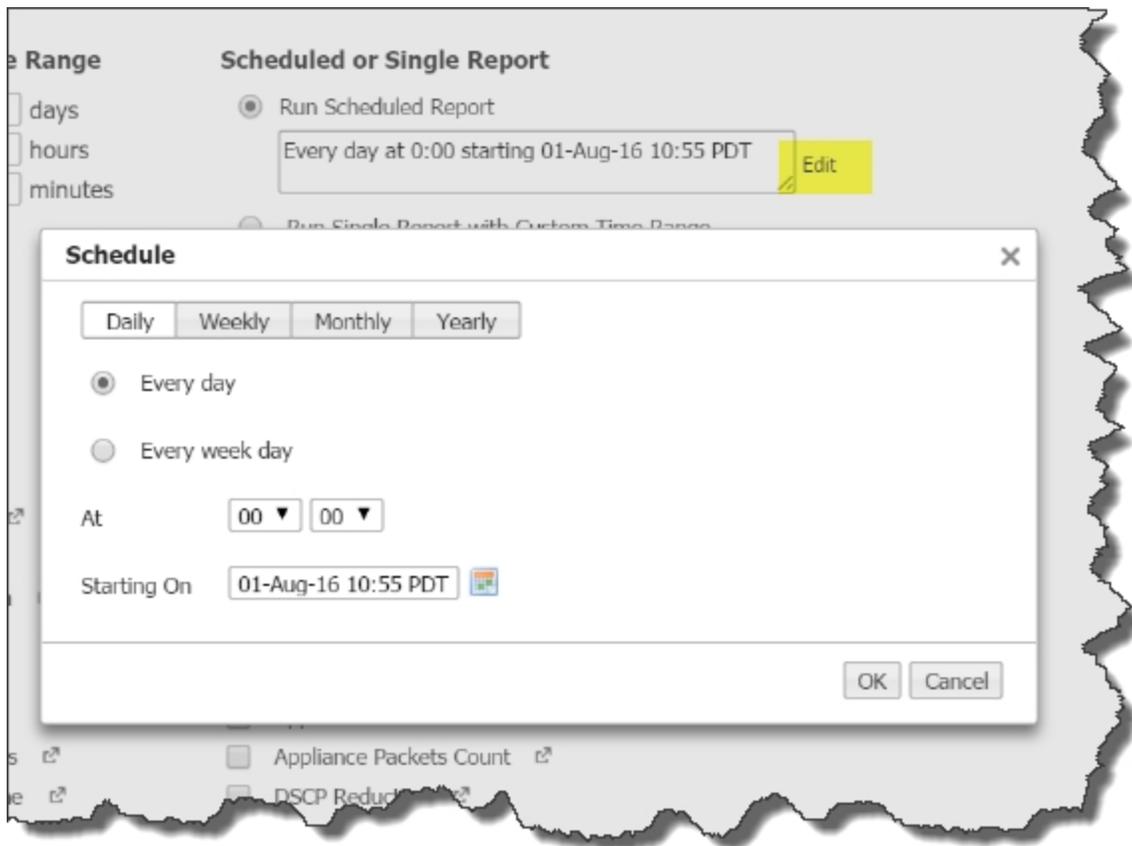
For example, if you want to run a report on the 1st and 15th of the month:

1. In the Name section, click **New Report**. This ensures you don't delete any previously defined reports.
2. Type a name into the field, then click **Save**. You can use letters, numbers, and certain symbols, but not spaces.



3. In the Scheduled or Single Report section, click **Edit**.

The Schedule form appears.



4. Under Monthly, set Day 1 of every 1 month. Then click **OK**.
This runs the report on the 1st of each month.

Schedule [X]

First day of the month
 Last day of the month
 Day of every month(s)
 The of every month(s)

At

Starting On [Calendar Icon]

[OK] [Cancel]

5. Choose your report options, then click **Save**.
6. Repeat steps 1-4, creating a new report and opening the Schedule box.
7. Under Monthly, set Day 15 of every 1 month. Then click **OK**.

This runs the report on the 15th of each month.

8. Choose your report options, then click **Save**.

A Success message appears at the bottom of the page each time you have successfully saved your report settings.